

For: KNILES

Printed on: Thu, May 23, 1991 10:36:27

From book: Frontcover in 7X9SLD_5/12

Document: final_cover

Last saved on: Mon, May 20, 1991 08:52:11

Document: Title_final

Last saved on: Wed, May 22, 1991 13:05:44

Document: copyright_hw

Last saved on: Mon, May 20, 1991 08:59:15

Document: TOC

Last saved on: Wed, May 22, 1991 13:01:16

Document: Pref

Last saved on: Mon, May 20, 1991 09:03:14

Document: CHAP1

Last saved on: Wed, May 22, 1991 16:37:51

Document: CHAP2

Last saved on: Wed, May 22, 1991 11:47:57

Document: CHAP3

Last saved on: Wed, May 22, 1991 12:14:49

Document: CHAP4

Last saved on: Wed, May 22, 1991 12:13:43

Document: CHAP5

Last saved on: Wed, May 22, 1991 12:12:49

Document: CHAP6

Last saved on: Thu, May 16, 1991 13:29:42

(...)

Fiber Distributed Data Interface



System Level Description

Order Number: EK-FDSLD-SD-002

Fiber Distributed Data Interface

System Level Description

May 1991

This manual introduces and describes Digital's implementation of Fiber Distributed Data Interface.



Order Number: EK-FDSLD-SD-002

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

Copyright © 1991 by Digital Equipment Corporation
All Rights Reserved.
Printed in U.S.A.

The following are trademarks of Digital Equipment Corporation:

DEBET	DECnet	TURBOchannel
DEC	DECstation	ULTRIX
DEC FDDIcontroller	DECsystem	UNIBUS
DECbridge	DECUS	VAX
DECconcentrator	DECwindows	VAXcluster
DECconnect	LAN Traffic Monitor	VAXstation
DECelms	PDP	VMS
digital	ThinWire	

This manual was produced by Telecommunications and Networks Publications.

Contents

Preface

1 Overview

1.1	FDDI Defined	1-2
1.2	Background to FDDI	1-4
1.2.1	ASC X3T9.5 Task Group	1-4
1.2.2	ANSI Specifications	1-7
1.2.3	FDDI and the OSI Model	1-8
1.3	FDDI Devices	1-9
1.3.1	Single Attachment Stations (SAS)	1-10
1.3.2	Dual Attachment Stations (DAS)	1-10
1.3.3	Concentrators	1-11
1.4	Application Environments	1-12
1.4.1	Backbone	1-13
1.4.2	Workgroup	1-14
1.5	FDDI Topologies	1-14
1.5.1	Standalone Concentrator	1-16
1.5.2	Tree of Concentrators	1-17
1.5.3	Dual Ring	1-18
1.5.4	Dual Ring of Trees	1-20
1.5.5	Dual Homing	1-22
1.6	Dual Ring of Trees, Topology of Choice	1-23
1.6.1	Reliability	1-27
1.6.2	Scalability	1-28

1.6.3	Maintainability	1-28
1.7	Digital Product Line	1-28
1.7.1	Concentrators	1-29
1.7.2	FDDI Bridges	1-29
1.7.3	Adapters	1-29
1.7.4	LAN Management Software	1-30
1.8	FDDI Specifications	1-30
1.9	General Planning and Configuration	1-31

2 Concepts

2.1	Overview	2-1
2.2	FDDI Structure	2-2
2.3	Physical Layer Medium Dependent (PMD)	2-3
2.3.1	Fiber Cable Requirements	2-4
2.3.2	Optical Bypass Relay	2-6
2.3.3	Connector Keying	2-7
2.4	Physical Layer Protocol (PHY)	2-7
2.4.1	Symbols	2-8
2.4.2	Line States	2-11
2.4.3	4B/5B Data Encoding	2-12
2.4.4	Modified NRZI Encoding	2-13
2.4.5	Clock Recovery and the Elasticity Buffer	2-13
2.4.6	Smoothing Function	2-14
2.4.7	Repeat Filter	2-14
2.4.8	Scrubbing	2-14
2.5	Media Access Control (MAC)	2-15
2.5.1	Ring Scheduling	2-16
2.5.2	Data Link Addressing Scheme	2-16
2.5.3	Universally Administered Addressing	2-17
2.5.4	Canonical Addresses and the Order of Transmission	2-18
2.6	Protocol Data Units	2-19
2.6.1	Tokens	2-20
2.6.2	Frames	2-21
2.6.3	Frame and Token Fields	2-21

2.7	Station Management (SMT)	2-23
2.7.1	Connection Management (CMT)	2-23
2.7.2	Connection Rules	2-25
2.7.3	Ring Management (RMT)	2-27
2.7.4	SMT Frame Services	2-28
2.7.5	SMT Frames	2-29
2.8	Logical Link Control (LLC)	2-29
2.9	Ring Operation	2-30
2.9.1	Asynchronous Ring Transmission	2-30
2.9.2	Timers	2-30
2.9.3	Target Token Rotation Time, Claiming, and Initialization ...	2-31
2.9.4	Steady State Operation	2-33
2.10	Ring Maintenance	2-36
2.10.1	Beaconing	2-37
2.10.2	Trace	2-38
2.11	Digital Enhancements	2-40
2.11.1	Frame Content Independent Stripping Algorithm	2-40
2.11.2	Ring Purger Algorithm	2-41

3 Concentrators

3.1	Description	3-1
3.2	Concentrator Functions	3-5
3.2.1	MAC and MAC-less Concentrators	3-5
3.2.2	Station Bypass	3-6
3.2.3	Network Management	3-6
3.3	Concentrator Operation	3-7
3.4	Concentrator Planning Considerations	3-7
3.5	Configuration Guidelines	3-7
3.5.1	Concentrator Connection Rules	3-9
3.5.2	Concentrators and Structured Wiring	3-10

4 Bridges

4.1	Description	4-2
-----	-------------------	-----

4.2	Bridge Functions	4-4
4.2.1	Translation Not Encapsulation	4-4
4.2.2	Frame Forwarding and Filtering	4-5
4.2.3	Source Address Tracking	4-5
4.2.4	Spanning Tree	4-6
4.2.5	Network Management	4-6
4.2.6	Internet Protocol Fragmentation	4-7
4.3	Bridge Operation	4-7
4.3.1	FDDI Port	4-8
4.4	Performance Considerations	4-9
4.5	Bridge Planning Considerations	4-10
4.6	Bridge Configuration Guidelines	4-10
4.6.1	General Connection Rules	4-10

5 Adapters

5.1	Description	5-1
5.2	Adapter Functionality	5-1
5.2.1	Host Interface	5-2
5.2.2	Adapter Manager	5-3
5.2.3	Packet Memory	5-3
5.2.4	FDDI Port	5-3
5.3	Operation	5-3
5.3.1	Initialization	5-4
5.3.2	Accessing the Ring	5-4
5.3.3	Network Management Functions	5-4
5.4	Adapter Planning Considerations	5-5
5.5	Adapter Configuration Guidelines	5-5

6 Network Management

6.1	Network Management Benefits	6-2
6.2	Network Management Functions	6-2
6.2.1	Network Configuration Management	6-3
6.2.2	Fault Management	6-3

6.2.3	Performance Management	6-3
6.2.4	Accounting Management	6-3
6.2.5	Security Management	6-4
6.3	Digital's Network Management Capabilities	6-4
6.3.1	Device Naming	6-4
6.3.2	Device Registry	6-4
6.3.3	Event Messages and Event Log	6-5
6.3.4	Security	6-5
6.3.5	Configuration Parameters	6-5
6.3.6	Initialization	6-6
6.3.7	Bridge Port Enabling and Configuring	6-6
6.3.8	PHY Port Configuration and Control	6-6
6.3.9	Address Filtering	6-6
6.3.10	Protocol Filtering	6-7
6.3.11	Network Monitoring	6-7
6.3.12	FDDI Physical Ring Mapping	6-7
6.3.13	SMT Neighborhood Information Frame (NIF) Display	6-7
6.3.14	SMT Station Information Frame (SIF) Display	6-7
6.3.15	Troubleshooting	6-8

A OSI Reference Model

A.1	Physical Layer	A-3
A.1.1	Data Link Layer	A-3
A.1.2	Network Layer	A-3
A.1.3	Transport Layer	A-3
A.1.4	Session Layer	A-4
A.1.5	Presentation Layer	A-4
A.1.6	Application Layer	A-4

B Additional Configurations

B.1	Standalone Configurations Using Concentrators	B-1
B.2	Tree Configurations Using Concentrators and Bridges	B-3
B.3	Dual Ring Configurations	B-6

B.4	FDDI Dual Homing	B-8
B.5	Single-Mode Fiber Configurations	B-10

C Related Documents

Glossary

Index

Figures

1-1	Dual Ring of Trees Topology with Extended LANs	1-3
1-2	FDDI Model	1-6
1-3	FDDI and Other LAN Standards	1-7
1-4	Relationship Between the OSI Model and FDDI	1-9
1-5	Single Attachment Station Model	1-10
1-6	Dual Attachment Station Model	1-11
1-7	FDDI Concentrator Model	1-12
1-8	Backbone Configuration	1-13
1-9	Standalone Workgroup Configuration	1-14
1-10	FDDI Physical Connection	1-15
1-11	Standalone Concentrator Topology	1-16
1-12	Tree of Concentrators Topology	1-17
1-13	Dual Ring—Conceptual View	1-18
1-14	Dual Ring—Isolating a Link Failure	1-19
1-15	Dual Ring—Isolating a Station Failure	1-20
1-16	Dual Ring of Trees Topology	1-21
1-17	Dual Homing Topology Using Concentrators	1-22
1-18	Dual Homing Topology Using Concentrators and a DAS Bridge	1-23
1-19	Single Ring Failure of a non-FDDI Topology	1-24
1-20	Dual Ring Failure of an FDDI Topology	1-25
1-21	Concentrators in a Dual Ring of Trees Configuration	1-27

2-1	FDDI Network Connection Model	2-3
2-2	Optical Bypass	2-6
2-3	MIC Receptacle Keying	2-7
2-4	Symbol Flow Through the PHY	2-9
2-5	NRZI Encoding Process	2-13
2-6	48-Bit Address Format	2-16
2-7	Universal Address Format	2-17
2-8	Order of Transmission for Addresses	2-18
2-9	Token Format	2-20
2-10	FDDI Frame Format	2-21
2-11	Station Management Architecture	2-24
2-12	PHY- Port Types	2-25
2-13	FDDI Connection Rules Matrix	2-26
2-14	SMT Frame Format	2-29
2-15	Logical Link Control and FDDI	2-29
2-16	Claiming Process	2-32
2-17	Capturing the Token	2-34
2-18	FDDI Typical Ring Recovery	2-37
2-19	Detection of Defective Station	2-39
3-1	Concentrator in a Standalone Topology	3-2
3-2	Concentrators in a Dual Ring of Trees Topology	3-3
3-3	Concentrator Architecture	3-4
3-4	Dual Ring of Concentrators in a Backbone Wiring System	3-8
3-5	Concentrators Cascaded in Equipment Rooms	3-11
3-6	Dual Ring of Trees in a Campus Environment	3-12
4-1	Extended LAN Examples	4-2
4-2	FDDI Bridge Functional Diagram	4-3
4-3	SAS Bridge-to-Concentrator Connection	4-8
4-4	Bridge Distance Requirements for Multimode Connections ...	4-11
5-1	Adapter Model	5-2
5-2	Adapter Distance Requirements for Multimode Fiber	5-6
A-1	OSI Seven-Layer Model	A-2
B-1	Standalone Workgroup Installation	B-2
B-2	Dedicated Network in a Dual Ring Installation	B-3

B-3	FDDI Tree of Concentrators and Bridges, Campus Configuration	B-4
B-4	Tree of Concentrators Installed in a Building	B-5
B-5	Dual Ring of Concentrators Installed in a Building	B-6
B-6	Dual Ring of Concentrators and Bridges Installed in a Campus	B-7
B-7	Concentrators Installed for Dual Homing	B-8
B-8	DAS Bridge Installed for Dual Homing	B-9
B-9	Connecting to the Dual Ring Through a Single-Mode Link ...	B-10
B-10	Dual Ring with Mixed Link Lengths	B-11

Tables

1-1	ANSI and ISO Standards for FDDI	1-5
1-2	FDDI General Specifications	1-30
1-3	Recommended Fiber Counts	1-32
2-1	Multimode Fiber Cable Specifications	2-4
2-2	Alternate Multimode Fiber Cable Specifications for FDDI	2-4
2-3	Single-Mode Fiber Cable Specifications	2-5
2-4	Symbol Encoding	2-10
2-5	FDDI Frames	2-19

Preface

The *Fiber Distributed Data Interface System Level Description* describes the American National Standards Institute's (ANSI) Fiber Distributed Data Interface (FDDI), 100-Mb/s, timed-token ring network standard. This manual also describes Digital Equipment Corporation's implementation of the ANSI standard. It discusses the various FDDI components designed and built by Digital and how these devices connect to form an FDDI network.

This manual discusses the topologies recommended by Digital and the benefits to users who implement them. In addition, this manual explains Digital's approach to network management and describes the facilities provided by network management software.

Audience

This manual is intended for users who are familiar with basic network concepts and who want an in-depth understanding of FDDI. By explaining FDDI and Digital's FDDI product set, this manual provides users with an understanding of the complexity of FDDI and the power of Digital's products. After reading this document, users will better understand the benefits that FDDI provides in the workplace.

Document Structure

This manual consists of six chapters and three appendixes as follows:

Chapter 1	Defines ANSI FDDI and provides background information on how it evolved. This chapter also describes the FDDI topologies recommended by Digital and lists Digital's FDDI product set. An overview of FDDI configuration guidelines is included at the end of this chapter.
-----------	--

Chapter 2	Discusses the basic design features of FDDI. This chapter lists and defines the four standards that, taken as a whole, make up ANSI FDDI. Topics covered include frame and token structure, clocking, the physical connection of FDDI stations, and basic ring operation.
Chapter 3	Contains a general description of concentrators and how they are used in an FDDI environment. It emphasizes the importance of the concentrator in Digital's FDDI topology. Configuration guidelines for concentrators are included at the end of this chapter.
Chapter 4	Describes bridges and their role in the FDDI environment. Bridge configuration guidelines are included at the end of this chapter.
Chapter 5	Provides a general description of adapters and their use in an FDDI environment. Adapter configuration guidelines are included at the end of this chapter.
Chapter 6	Describes network management as it applies to FDDI and the extended LAN environment. This chapter provides an overview of network management capabilities. It also discusses concepts of Digital's Enterprise Management Architecture (EMA).
Appendix A	Describes the International Organization for Standardization (ISO) seven-layer Open System Interconnect (OSI) network architecture.
Appendix B	Contains examples of FDDI network configurations built on Digital's implementation of the commercial building wiring scheme.
Appendix C	Contains a list of related documents that can be ordered from Digital Equipment Corporation.

Overview

In an effort to standardize high-speed networking protocols, the American National Standards Institute (ANSI) has defined a 100-Megabit-per-second (Mb/s) local area network (LAN). Fiber optic cable is the first transmission medium defined in the ANSI specification. Called Fiber Distributed Data Interface (FDDI), this ANSI specification consists of a set of standards that define the components of FDDI: Physical Layer Medium Dependent (PMD), Physical Layer Protocol (PHY), Media Access Control (MAC), and Station Management (SMT). This manual describes FDDI and addresses questions, including:

- Why did ANSI develop FDDI?
- What are the goals for FDDI?
- What problem is ANSI solving with FDDI?

This manual also discusses the benefits of adopting Digital's implementation of FDDI, based on the dual ring of trees network topology. These benefits include the following:

- Reliability
- Scalability
- Maintainability

Digital's implementation of FDDI complies with the ANSI standards as they are written. This chapter provides a brief history of the FDDI specification. It also in-

cludes a description of the physical environments where FDDI is used and an introduction to the basic components of Digital's implementation of FDDI.

Throughout this manual, the words station, device, and node are used extensively in descriptions of equipment connected to the FDDI ring. Keep the following definitions for these words in mind:

Node	Any attachment device to the FDDI network. It can be a single attachment station (SAS), a dual attachment station (DAS), or a concentrator.
Station	As used in the FDDI environment, any node that possesses a MAC entity.
Device	Generic hardware that connects to the FDDI network. Can be a station, a node, or the equipment that connects to the network through the FDDI components.

1.1 FDDI Defined


FDDI consists of a set of standards that, when taken together, define a 100-Mb/s, timed-token passing, local area network that uses fiber optic medium to transmit data between attached devices. Note that fiber optic medium is the first medium defined by the standards. Work is underway to add other media. The standards define several types of networking devices. These include concentrators, dual attachment stations, and single attachment stations. These devices allow users to construct various network configurations based on the FDDI standard.

Digital Equipment Corporation recommends a basic network topology, called a dual ring of trees, that provides a backup physical link for increased reliability. In this topology, an FDDI ring is constructed as trees of stations connected to the FDDI dual ring. Figure 1–1 illustrates this dual ring of trees topology. This figure shows how the various FDDI components connect to form an integrated, fully functional network. In this topology, dual attachment stations and single attachment stations connect to the FDDI ring through concentrators forming the branches of the trees.

Figure 1–1 emphasizes the importance of the concentrator as the key component in a dual ring of trees topology. In Figure 1–1, the dual ring connects concentrators. Stations connect to the dual ring through a tree of concentrators. This topology offers a campus-wide network free from disruption by everyday events.

Network management software provides the tools to maintain and manage the network. Other sections of this manual provide details on the benefits of this topology and how SAS and DAS bridges connect FDDI LANs to extended LANs.

Figure 1-1: Dual Ring of Trees Topology with Extended LANs



LKG-5102-91I

Digital, by making its implementation fully compliant with the ANSI standards, achieves interoperability with other vendors whose products are also truly ANSI compliant. By maintaining compliance with the standard, Digital also makes it easier for users to design networks using off-the-shelf components, rather than exotic, proprietary hardware and software.

1.2 Background to FDDI

FDDI is an ANSI standard for a high-speed, general-purpose network for connecting large computers, workstations, desktop computers, and other equipment. ANSI defines and oversees the standards creation process and also charts other organizations to write standards for domestic use. This section explains the various ANSI specifications that make up FDDI. Section 1.2.3 maps the ANSI standards to the ISO/OSI model.

1.2.1 ASC X3T9.5 Task Group

ASC X3T9.5 is the ANSI Task Group responsible for the design of the overall ANSI FDDI specification. The Accredited Standards Committee (ASC), X3, develops standards for Information Processing Systems. Technical Committee X3T9 is responsible for writing standards for I/O interfaces. The X3T9.5 Task Group is chartered to write the FDDI standard.


The ANSI standards, and related ISO versions, that compose the FDDI specification are listed in Table 1–1 and described in the following sections. These sections are overviews of the individual standards. For further information refer to the individual ANSI or ISO document for each standard.

Table 1–1: ANSI and ISO Standards for FDDI

ANSI/ISO Standard	FDDI Standard
X3.166–1990/ ISO 9314-3:1990	Physical Layer Medium Dependent —This standard corresponds to the lower portion of the OSI Physical layer. PMD defines the transmit/receive power levels, optical transmitter and receiver interface requirements, error rates, and cable and connector specifications.
X3.148–1988/ ISO 9314-1:1989	Physical Layer Protocol —This medium-independent standard corresponds to the upper portion of the Physical layer. PHY defines symbols, line states, encoding/decoding techniques, clocking requirements, and data framing requirements.
X3.139–1987/ ISO 9314-2:1989	Media Access Control —This standard corresponds to the lower portion of the OSI Data Link layer. MAC defines data link addressing, frame formatting and checking, medium access, error detecting, and token handling. Its primary function is to deliver formatted data to nodes attached to the FDDI LAN.
X3.T9.5/84–49 (under development)	Station Management —This draft standard defines the system management services for the FDDI protocols. SMT includes facilities for connection management, node configuration, recovery from error conditions, and the encoding of SMT frames.


Figure 1–2 shows the relationship among the various FDDI standards and how they combine to form the complete FDDI standard. Figure 1–3 shows the relationship among the FDDI standards and other common industry standards. Note that the 802.2 Logical Link Control and IEEE 802.1 standards are common across the various lower layer standards.

Figure 1–2: FDDI Model



LKG-4042-9C

Figure 1–3: FDDI and Other LAN Standards



1.2.2 ANSI Specifications

The ANSI standards that make up the FDDI specification include X3.166–1990, X3.148–1988, X3.139–1987, and X3.T9.5/84–49. Together, these standards define a fully functional high-speed fiber optic network. Chapter 2 describes the individual standards in greater detail. Following are summaries of the FDDI standards.

X3.166–1990 defines the Physical Layer Medium Dependent (PMD) sublayer. This sublayer corresponds to the lower portion of the Physical layer of the OSI seven-layer model. This standard provides the specifications for the optical transmitter and receiver used by FDDI, the power levels, allowable bit error rates, fiber optic transmission wavelength, fiber optic cable, and connectors.

X3.148–1988 defines the Physical Layer Protocol (PHY). This medium-independent layer corresponds to the upper portion of the OSI Physical layer. This standard defines the encoding scheme, clock synchronization, and serial-parallel conversion.

X3.139–1987 defines the Media Access Control (MAC) sublayer. The MAC corresponds to the lowest sublayer of the OSI Data Link layer (DLL). This standard defines medium access, addressing, and error detection protocols. Its primary function is the delivery of frames to attached nodes on the FDDI ring. It also provides access to the ring for higher level protocols, specifically, the Logical Link Control (LLC) protocol.

X3.T9.5/84–49 defines the system management applications for the FDDI protocol layers. Called Station Management (SMT), it provides the specifications for the station-level control necessary to ensure proper node operation in an FDDI ring. These include topology, station configuration, ring management, and connection management.

1.2.3 FDDI and the OSI Model

The Open Systems Interconnection (OSI) model is a reference for defining the architecture necessary to provide communications between open systems. Developed by the International Organization for Standardization (ISO), it provides a framework for the development of international standards for computer communications. The OSI model is based on the layering technique where the various communications functions are partitioned into seven discrete units. This section describes the relationship between the OSI layered model and the FDDI standard. For further information concerning the OSI model, refer to Appendix A.


The FDDI standards apply to the Physical and Data Link layers of the OSI model. Figure 1–4 shows the relationship between the OSI model and the FDDI standards. In the OSI model, the Physical layer defines data transmission over the physical link. Issues such as bit signaling, mechanical and electrical specifications, cabling, and connectors are covered here. In FDDI this is further divided into the PHY and PMD sublayers, which separate the media and transmission details into two distinct parts.

The OSI Data Link layer defines parameters for medium access, frame construction, data link addressing, error detection, and connection to higher layers. FDDI divides this layer into two sublayers. The upper sublayer, Logical Link Control, is provided by the IEEE 802.2 LLC standard. This is the same LLC used in IEEE 802.3, .4, and .5. It provides a means for exchanging data between LLC users across an FDDI LAN.

The lower sublayer is defined in the MAC and SMT standards in FDDI. This sublayer defines the data link packets specifically required by the FDDI token passing

scheme. These include addressing and token handling. In Figure 1–4, Service Access Points (SAPs) are logical connections between layers in the OSI model. They are shown here for informational purposes only.

Figure 1–4: Relationship Between the OSI Model and FDDI



LKG-4044-901

1.3 FDDI Devices

The FDDI standard defines two types of devices: stations and concentrators. Stations can be further divided into single attachment stations (SAS) and dual attachment stations (DAS). The concentrator provides connections for the attachment of multiple nodes to the FDDI ring. The concentrator can be a single attachment concentrator (SAC), dual attachment concentrator (DAC), or a standalone concentrator. All of these device types act as connection points to the FDDI network. Digital implements the SAS, DAS, and DAC in the dual ring of trees topology. The following sections


describe the attachment devices defined by FDDI. Chapter 2 provides a detailed discussion of the components defined by the FDDI standards.

1.3.1 Single Attachment Stations (SAS)

The single attachment station (SAS) permits access to the FDDI ring through the concentrator. The SAS provides a single connection point to the FDDI ring. It does not have the facilities for connecting to the dual ring; therefore, it cannot wrap the ring in case of a station or fiber failure. The concentrator provides fault control. In the dual ring of trees topology, the SAS proves to be a reliable, cost-effective method for connecting to the FDDI network.

The components of the SAS are a single instance of the MAC entity, the PHY, the PMD, and the SMT. Figure 1–5 shows a model of the SAS.

Figure 1–5: Single Attachment Station Model




LKG-4045-90I

1.3.2 Dual Attachment Stations (DAS)

The dual attachment station (DAS) permits a device to connect to the primary and secondary rings of the FDDI dual ring, allowing a device to access the dual ring architecture of FDDI. The DAS consists of two PHYs, two PMDs, one or (optionally) two MACs, one SMT, and an optical bypass relay (optional). Figure 1–6 shows a model of the DAS.

In the dual ring of trees topology employed by Digital, DAS nodes are supported, but their use is not required. Concentrators, SAS and DAS bridges, and SAS nodes compose a typical dual ring of trees topology as implemented by Digital.

Figure 1–6: Dual Attachment Station Model



LKG-4046-901

1.3.3 Concentrators


The concentrator is also specified as part of the ANSI standards. The concentrator allows the attachment of multiple SAS, DAS, or concentrators to the primary (or optionally, secondary) ring of the FDDI network. The concentrator can also be used as the root of a tree topology. As shown in Figure 1–7, the concentrator consists of SMT circuitry, an optional MAC entity, and several PHY/PMD entities. A complete physical connection consists of a PHY/PMD in a concentrator connected by a medium to the PHY/PMD in another node.

An FDDI concentrator retimes data, detects faults, performs fault isolation, and is manageable by network management software. Unlike concentrators found in other networks (IEEE 802.5, for example), the FDDI concentrator is an active device that can actually control the physical topology of a network.

The concentrator can reconfigure the network by inserting or removing connected nodes. The configuration switch shown in Figure 1–7 provides this capability to the concentrator.

In Figure 1–7, the individual PHY/PMD entities have separate paths to the SMT, as indicated by the dashed lines through subsequent blocks in the figure. Chapter 3 provides a closer look at the concentrator as it is used by Digital in the dual ring of trees topology.

Figure 1–7: FDDI Concentrator Model



LKG-4047-901


1.4 Application Environments

Digital's implementation of FDDI supports various LAN environments including backbones and workgroups. These environments provide great flexibility in the design and implementation of an FDDI network. The following sections and supporting illustrations further describe these network environments.

1.4.1 Backbone

Figure 1–8 shows an FDDI backbone configuration. This configuration ties together multiple LANs in a campus or multibuilding environment into one integrated network. The backbone provides a high-bandwidth facility for the connection of FDDI devices. Because attached devices can be subjected to frequent power down by the user, connection to the FDDI ring is optimized by the use of concentrators. Backbones provide high reliability; the use of concentrators maintains ring integrity even when one or all devices attached to it are shut down or unplugged. The backbone supports radial wiring schemes such as the one defined in EIA/TIA 568 (TR–41.8.1).

Figure 1–8: Backbone Configuration



HDF = Horizontal distribution frame
IDF = Intermediate distribution frame
MDF = Main distribution frame

SAS = Single attachment station
DAS = Dual attachment station


LKG-4048-901

1.4.2 Workgroup

The workgroup configuration connects workstations, personal computers, and mini-computers through one or more concentrators. This network is characterized by a relatively small number of attached devices spread over a limited geographical area and with FDDI's ability to transfer large amounts of data quickly and reliably.

The typical workgroup configuration normally has distances between stations less than the 2 kilometers (1.2 miles) specified in the FDDI standards. The dual ring of trees architecture employed in the FDDI ring allows network to sustain the loss of one or more attached nodes and still function. A workgroup can be part of a backbone configuration as shown in Figure 1-8. Figure 1-9 illustrates a standalone workgroup configuration using a concentrator.

Figure 1-9: Standalone Workgroup Configuration




LKG-4050-901

1.5 FDDI Topologies

An FDDI ring network consists of a number of serially attached stations that are connected by a transmission medium to form a closed loop. An active station transmits information serially as a stream of symbols to the next active station on the ring. As each active station receives these symbols, it regenerates and repeats them to the next active device on the ring (its downstream neighbor).

The physical connection (see Figure 1–10) is basic to the FDDI ring design: Through this connection attached devices pass information over the ring. An FDDI physical connection is formed between the Physical layers (PHY/PMD) of two stations connected by the medium. FDDI requires that all connections in the ring be point-to-point, full-duplex, and bidirectional. Each attachment to the network has a transmit and receive component. For example, a valid fiber optic connection contains two fibers providing a point-to-point, full-duplex path.

Figure 1–10: FDDI Physical Connection



As part of the FDDI standard, ANSI permits a number of topologies. Of these topologies, four are of particular importance to the users of Digital's FDDI products. These topologies are as follows:


- Standalone concentrator with attached nodes
- Tree of concentrators
- Dual counter-rotating ring
- Dual ring of trees

These topologies are discussed in detail in the following sections.

1.5.1 Standalone Concentrator

The standalone concentrator topology consists of a single concentrator and its attached stations (Figure 1–11). These stations can be either SAS or DAS devices. Because of their lower cost, SAS devices are the preferred station for this topology. This topology does not require access to the FDDI dual ring. It offers all the advantages associated with the concentrator, including reliability and ease of configuring, to small, independent workgroups. This topology allows the use of star wiring. This means existing structured fiber optic wiring can be used, affording significant cost savings in prewired sites. A standalone concentrator topology grows into a tree of concentrators as the network is expanded.

Figure 1–11: Standalone Concentrator Topology




LKG-4056-90I

1.5.2 Tree of Concentrators

The tree of concentrators, shown in Figure 1–12, is used when wiring large groups of user devices together. Concentrators are wired in a star topology with one concentrator serving as the root of the tree. Additional concentrators can connect to the second tier of concentrators as needed to support new users. This topology can be used to wire a single building or to service many stations on one floor of a building. This topology lends itself to structured wiring systems.

Figure 1–12: Tree of Concentrators Topology




LKG-4236-90I

1.5.3 Dual Ring

One of the basic concepts in the FDDI standard is the dual counter-rotating ring. The dual ring consists of a primary ring and a secondary ring. The direction of data flow in the secondary ring is opposite (counter to) that of the primary ring. The main functions of this secondary path are to assist in the ring initialization and reconfiguration process and to provide backup to the first ring. Inherent in this design is the ability to continue ring transmission if a device on the ring fails. The ring is restored by wrapping to the secondary ring to maintain the transmission path. This redundancy in the ring design provides a degree of fault tolerance not found in other network standards. Figure 1–13 illustrates the basic concepts of the dual ring.

Figure 1–13: Dual Ring—Conceptual View




All devices on the ring are dual attachment stations or dual attachment concentrators.

LKG-4053-901

Dual attachment stations or dual attachment concentrators on the dual ring connect to each other by means of the primary and secondary fiber cables. Under normal conditions, the secondary ring is idle with no data being transmitted. If a fiber failure occurs, the stations on either side of the failure reconfigure by wrapping the primary ring to the secondary ring, effectively isolating the fault. This maintains connectivity between stations in the dual ring, allowing normal operation to continue. Figure 1–14 shows the FDDI technique of isolating a link failure.

Figure 1–14: Dual Ring—Isolating a Link Failure




All devices on the ring are dual attachment stations or dual attachment concentrators.

LKG-4054-90I

If an attached station fails, again the stations on either side of the failing station reconfigure, effectively isolating the failing station from the ring. Figure 1–15 shows the FDDI technique of isolating a station failure.

FDDI limits total fiber length to 200 km (124 mi). Since the dual ring topology effectively doubles media length in the event of a ring wrap, the actual length of each ring is limited to 100 km (62 mi).

Figure 1–15: Dual Ring—Isolating a Station Failure




All devices on the ring are dual attachment stations or dual attachment concentrators.

LKG-4055-90I

1.5.4 Dual Ring of Trees

A fourth topology described in FDDI, the dual ring of trees (DRT), connects concentrators and DAS bridges located at a Main Distribution Frame (MDF) together in a dual ring. This provides the reliability of the dual ring where it is needed most, in the campus backbone, while allowing for the use of structured wiring in all other campus locations. Trees of concentrators then connect to the dual ring, providing connections at the Intermediate Distribution Frames (IDF) and Horizontal Distribution Frames (HDF) of the structured wiring system. Stations then connect to concentrators located in the HDFs of the structured wiring system. The tree can branch as required by simply adding concentrators as long as the station number or ring distance limits are not exceeded. Figure 1–16 shows the DRT topology.

Figure 1-16: Dual Ring of Trees Topology




The concentrator plays a key role in fault isolation in this topology. The dual ring at the backbone guarantees a secondary data path to all concentrators and dual attachment stations in the dual ring. Beyond this level, in the trees, individual stations can be removed from the ring as needed with no interruption of service to the rest of the network. The concentrators can bypass inactive or defective stations, as required.

1.5.5 Dual Homing

The dual homing topology uses backup links to enhance fault tolerance in a tree topology. One link is called the primary link, and the other link is called the backup link. Figure 1–17 shows a dual homed topology using concentrators. If concentrator 2 or the primary link to concentrator 3 fails, the backup link to concentrator 3, through concentrator 1, is activated. This ensures that the devices connected to concentrator 3 will have uninterrupted service. Activation of the individual links follows the FDDI connection rules, described in Section 2.7.2.


Figure 1–17: Dual Homing Topology Using Concentrators



LKG-4812-91I

Figure 1–18 shows a dual homed topology using concentrators and a DAS bridge. If concentrator 2 or the primary link to the DAS bridge fails, the backup link to the DAS bridge, through concentrator 1, is activated. This ensures that the devices connected to the DAS bridge will have uninterrupted service. Activation of the individual links follows the FDDI connection rules, described in Section 2.7.2.

Figure 1–18: Dual Homing Topology Using Concentrators and a DAS Bridge




LKG-4934-911

1.6 Dual Ring of Trees, Topology of Choice

Network failure problems that are inherent in a single ring topology can be corrected by implementing a dual ring topology. However, a dual ring topology creates segments if multiple station or cable failures occur. A dual ring of trees topology using concentrators corrects the problems that are inherent in both single and dual ring designs. This section explains Digital's choice of the dual ring of trees as the recommended topology for FDDI.

A single ring topology consists of a series of point-to-point physical links in a closed loop. If a failure on the ring occurs, the entire network fails. Figure 1–19 shows the single-point failure problem in a non-FDDI single-ring topology. In large installations where the probability of ring or station failure is high, this type of disruption is unacceptable.

Figure 1–19: Single Ring Failure of a non-FDDI Topology




LKG-4057-901

To eliminate the single-point-of-failure problem, the ANSI FDDI standards committee designed the dual ring topology. This topology solves the problem of a single failure disrupting the ring by providing a secondary transmission path. The secondary transmission path assists in start-up, initialization, and reconfiguration of the primary path.

Although the counter-rotating ring design eliminates the problems associated with the single ring, a new issue arises when multiple failures occur. Figure 1–20 illustrates what happens to a ring of DAS devices when multiple ring failures or disconnects occur. The result is two or more (depending on number of failures) segmented rings.

Figure 1–20: Dual Ring Failure of an FDDI Topology



LKG-5095-90I

Although each ring is fully functional, the rings cannot access each other because of the multiple failures. These breaks can be caused simply by a user disconnecting a workstation. Again, in large installations where the risk of multiple failures increases with the number of stations or large number of users added to the network, this is not an acceptable situation.

The dual ring solution can also give rise to other concerns such as:


- The dual ring is not well suited to the use of structured wiring schemes.
- Dual attachment stations require twice the connectors and cables at the drop site.
- The site backbone cabling is accessible to each user on the dual ring, which presents a ring integrity risk at every user location.
- Because each station is part of the backbone wiring, the behavior of each user is critical to the operation of the ring.
- Manual intervention and manipulation of the dual ring cabling is required for all additions, moves, and changes, resulting in periods of ring instability.

Digital recognizes the segmented ring issue inherent with DAS installations and addresses it with the use of FDDI concentrators configured in a dual ring of trees. The DRT topology allows the creation of a large, sophisticated network that can sustain the loss of all stations connected to the concentrators without losing ring integrity. In this topology, concentrators electronically bypass ring disruptions. Connection Management allows the removal of defective links with no disruption of service to remaining nodes. Figure 1–21 illustrates this concept. It shows two concentrators attached to the dual ring with two bypassed stations.

Only network components (concentrators or bridges) should be connected to the dual ring. These devices are typically in locked closets or cabinets and only accessible by network-knowledgeable individuals. End user stations, such as computers or workstations, should only be connected as SAS.

Further inspection of the DRT reveals it to be a more robust and dynamic solution to configuration problems than either the single or dual ring approaches to network design. With the dual ring of trees, the best of both topologies is available to network users. The dual ring implemented at the root of the tree provides transmission path redundancy to attached concentrators. The single ring cabling between concentrators and SAS stations or other concentrators on the tree branches provides significant cost savings compared to the purchase of DAS devices and the installation of multiple dual fiber cables at all user sites.

Figure 1–21: Concentrators in a Dual Ring of Trees Configuration



* Multiple station failures in a dual ring of trees result in no loss of access to the ring. Failing stations are simply removed from the ring.

LKG-4237-90I

Key points in the choice of the dual ring of trees topology are reliability, scalability, and maintainability. These factors provide the power not found in the other topologies.

1.6.1 Reliability

A dual ring of trees provides a backup physical link at the dual ring. In addition, concentrators in the dual ring or in the tree can bypass an inactive or defective station and allow the insertion or removal of stations as needed. This is done without the intervention of the network manager and without disrupting the network.

Because the network is unaffected by single- or multiple-point failures, it is always available to attached stations. Barring a power failure or major natural disaster, there is no ordinary incident that would cause the loss of ring access to other stations connected to the network. The dual ring of trees allows any number of stations to be removed from the ring without denying ring availability to remaining attached stations.

1.6.2 Scalability

Increasing or decreasing ring size is a simple matter of adding or removing concentrators. In a dual ring topology, recabling is necessary every time a station is added or removed. The inherent flexibility of the concentrator allows network growth to occur as needed; recabling is not necessary in a dual ring of trees topology. Starting with a single concentrator and attached SAS devices, a user can expand the network to fit the needs of the site as long as the station number and ring distance limit are not exceeded. The dual ring of trees permits additions, moves, and changes without disrupting the network. This flexibility in wiring complies with the wiring scheme specified in the EIA/TIA 568 standard.

1.6.3 Maintainability

Proper maintenance is critical to the operation of any network. Digital provides intelligent concentrators and distributed networking tools that permit the user to monitor and maintain the network down to individual attached devices. Proper installation in dedicated equipment rooms allows the concentrators to be carefully controlled and maintained. The physical topology recommended by Digital meets the specifications set forth in the EIA/TIA 568 standard.

1.7 Digital Product Line

The Digital implementation of ANSI FDDI allows for great flexibility in the design and construction of a high-speed fiber network. Digital designs and builds a number of products for the construction of an FDDI ring. They are fully compliant with the ANSI FDDI standards. The following sections provide an overview of the FDDI devices produced by Digital.

1.7.1 Concentrators

Concentrators are necessary in large networks where DAS configurations are difficult to configure and maintain, and are expensive and unreliable. Digital implements the DECconcentrator 500 as the recommended device for connecting stations to the FDDI ring. Digital supports the use of a DAS configured into a dual ring or as a member of a tree.

The DECconcentrator 500 complies with the ANSI X3T9.5 standards. It provides for attachment of nodes (SAS, DAS, or concentrator) in a standalone environment. The concentrator also supports the hierarchical building and campus physical wiring topologies discussed in previous sections by allowing concentrators to be interconnected in a tree topology. Finally, it supports the dual ring interconnection of concentrators and DAS, referred to as the dual ring in the ANSI standards. Refer to Chapter 3 for details concerning concentrator functionality.

1.7.2 FDDI Bridges

To preserve previous investments in equipment, the FDDI LAN must be able to interface with existing network technologies. Bridges provide a means to do this. In Digital's implementation of FDDI, the DECbridge 500/600 series provides the interface necessary to connect the 100-Mb/s FDDI ring to the 10-Mb/s, 802.3/Ethernet LAN. The DECbridge 500/600 series contains all necessary buffering and translation logic for interfacing to the 802.3/Ethernet network. In addition, the DECbridge 600 series allows multiple 802.3/Ethernet segments to be connected to the FDDI network.

The DECbridge 500/600 series contains SAS and DAS bridges. SAS and DAS bridges interconnect in a tree topology and in a dual ring of trees topology through a concentrator. Additionally, DAS bridges interconnect in a dual ring topology. Refer to Chapter 4 for further details concerning the bridge.

1.7.3 Adapters

Adapters connect different devices to the FDDI ring through the concentrator. Adapters can be SAS or DAS. An adapter connects to a device such as a workstation or minicomputer with an I/O bus, for example, a TURBOchannel bus. The adapter contains all necessary logic to act as an interface between the FDDI ring and device system memory. Refer to Chapter 5 for more details on adapter functionality.

1.7.4 LAN Management Software

Proper network management is critical to the successful implementation of a LAN. To this end, Digital has developed sophisticated LAN management software. This software allows a user at a management station to control and observe bridges, FDDI concentrators, and other nodes in the extended local area network. Network managers use the management software to monitor and control a network; field service technicians use it to troubleshoot the LAN. Refer to Chapter 6 for details on network management and the Digital management philosophy.

1.8 FDDI Specifications

FDDI allows the creation of a flexible, robust high-performance network. Table 1–2 lists the specifications for FDDI as set forth by Digital and the ANSI standards. For detailed information concerning building or site-cabling procedures and requirements, refer to the *DECconnect System Fiber Optic Planning and Configuration* and *Fiber Optic Installation* manuals.

Table 1–2: FDDI General Specifications

Item	Specification
Transmission rate	125 megabaud (100 Mb/s at the data link)
Physical layer entities	1000 (maximum)
Fiber length	200 km (124 mi) (maximum)
Link loss budget	11.0 dB (maximum loss—multimode fiber optic cable) 22 dB (maximum loss—single-mode fiber optic cable)
Interstation link length*	2 km (1.2 mi) (maximum—multimode fiber optic cable) 40 km (24.8 mi) (maximum—single-mode fiber optic cable)
Transmission medium	Fiber optic cable
Network topology	Dual ring of trees
Media access method	Timed-token passing
*FDDI link length is bandwidth limited to 2 km (1.2 mi) for multimode fiber.	

The number of stations that can be installed in an FDDI LAN depends on the number of Physical layer entities that are present. A Physical layer entity consists of one instance of PHY/PMD and is the equivalent of one PHY port. A station can have one Physical layer entity (SAS), two Physical layer entities (DAS), or multiple Physical layer entities (concentrators).

The ANSI FDDI standard specifies a maximum of 1000 physical layer entities and 1000 MAC entities in an FDDI LAN. The standard allows up to 500 stations depending on the topology. Digital supports the configurations specified in the ANSI FDDI standard. However, Digital recommends that an FDDI LAN generally be limited to 100 stations and 20 concentrators due to performance and reliability characteristics of large rings.

1.9 General Planning and Configuration

This section lists some general points to consider when planning an FDDI network. It does not, however, address all network planning requirements. For further details regarding the planning and installation of Digital fiber optic cable, refer to the following two manuals.

- *DECconnect System Fiber Optic Planning and Configuration*
- *DECconnect System Fiber Optic Installation*

In the initial planning stages, it is important to determine proper fiber counts at FDDI device locations. Consider the following points when determining fiber count:

- Multiple services support
- Level of redundancy
- Planned growth and expansion

Table 1–3 lists the minimum recommendations for fiber optic cable between various network components. The drop locations are defined according to the EIA/TIA 568 standard.

Table 1–3: Recommended Fiber Counts

Drop Location	Minimum Recommended Fiber Count	
	Multimode	Single Mode
MDF to IDF	44	12
IDF to HDF	24	6
HDF to office wallbox	2	n/a
MDF = Main Distribution Frame		
IDF = Intermediate Distribution Frame		
HDF = Horizontal Distribution Frame		
n/a = not applicable		

Concepts

This chapter explains the basic concepts of FDDI, giving a general overview of a complex design. Topics include Protocol Data Units (PDUs), Target Token Rotation Time (TTRT), and signaling techniques. It concludes with a discussion of ring operation. This chapter provides basic information about how FDDI operates and why certain functions were included in the FDDI design. It also discusses FDDI in terms of the four standards that compose FDDI:

- Physical Layer Medium Dependent (PMD)
- Physical Layer Protocol (PHY)
- Media Access Control (MAC)
- Station Management (SMT)

While reading about the intricacies of FDDI, keep in mind that Digital has incorporated the power of FDDI in a standard chipset that is used throughout the Digital family of FDDI products.

2.1 Overview

The ANSI FDDI committee reviewed many options during the development process. Key design concepts of FDDI include the following:

- Timed-Token Protocol
- Symbols for data transmission and line state signaling

- 4B/5B and NRZ/NRZI encoding
- Synchronous, asynchronous, and restricted service
- Token rotation time determined by negotiation among all connected stations
- Immediate transmission of data after token acquisition
- Distributed clocking to support a large number of stations on the ring


Because of the requirement for a 100-Mb/s transmission rate (125 megabaud), FDDI presents a new set of challenges to designers. By making use of the techniques noted above, designers have enabled the FDDI network to maintain a high transmission rate over great distances among a large number of stations.

2.2 FDDI Structure

FDDI is built on four sublayers as defined in the ANSI standards. Sections 2.3 through 2.7 discuss these items in detail. Figure 2–1 illustrates the basic FDDI network connection model and shows how these sublayers interact. The four standards that make up FDDI are the PMD, PHY, MAC, and SMT.

These four sublayers plus the IEEE 802.2 Logical Link Control (LLC) standard combine to provide essential networking services to devices attached to the FDDI network. These services include hardware connection, ring initialization, error detection, token handling, and overall station management.

Figure 2–1: FDDI Network Connection Model



2.3 Physical Layer Medium Dependent (PMD)

The PMD standard defines the lower sublayer of the OSI Physical layer. It describes the physical requirements for nodes that attach to the FDDI network and the hardware connection between these nodes. It also provides all services necessary to transmit an NRZI-encoded digital bit stream between stations. The PMD addresses the following:

- Transmitter and receiver requirements, including power budget
- Transmitter and receiver waveform characteristics
- Conversion of electrical data signals to and from optical signals
- Jitter limits and permissible bit error rates.
- Mechanical mating requirements for the Media Interface Connector (MIC), including the keying methods and connector footprint
- Fiber optic cable requirements

The PMD standard specifies a nominal optical wavelength of 1300 nm for data transmission over 62.5/125 μm multimode, graded index, fiber optic cable. This permits the use of lower-cost LEDs in place of the more expensive laser devices required for single-mode fiber.

2.3.1 Fiber Cable Requirements

Fiber optic cable is available in different sizes and types. The FDDI standard specifies the use of 62.5/125 μm graded index, multimode, fiber optic cable. An appendix to the PMD standard lists cable alternatives. Table 2–1 lists Digital’s recommended fiber cable specifications in both the 1300 nm and 850 nm windows. Digital recommends the use of dual-window fiber to support FDDI 1300 nm requirements and existing or future non-FDDI applications requiring a 850 nm window. Table 2–2 lists alternate multimode fiber cable specifications for FDDI in the 1300 nm window.

Table 2–1: Multimode Fiber Cable Specifications

Recommended Cable	Parameters
Nominal core diameter	$62.5 \pm 3.0 \mu\text{m}$
Cladding diameter	$125.0 \pm 2.0 \mu\text{m}$
Numerical aperture	0.275 ± 0.015
Modal bandwidth minimum	500 MHz•km @ 1300 nm, 160 MHz•km @ 850 nm
Maximum attenuation rate	1.5 dB/km @ 1300 nm, 3.5 dB/km @ 850 nm
Power budget	11.0 dB @ 1300 nm (for FDDI), varies by product @ 850 nm

Table 2–2: Alternate Multimode Fiber Cable Specifications for FDDI

Alternate Cable	100/140	50/125
Nominal core diameter	$100 \mu\text{m}^* \pm 4.0 \mu\text{m}$	$50 \mu\text{m}^* \pm 3.0 \mu\text{m}$
Cladding diameter	$140 \mu\text{m}^* \pm 6.0 \mu\text{m}$	$125 \mu\text{m}^* \pm 2.0 \mu\text{m}$
Optical wavelength	1300 nm	1300 nm
Numerical aperture	0.290 ± 0.015	0.200 ± 0.015
Modal bandwidth minimum	500 MHz•km @ 1300 nm	500 MHz•km @ 1300 nm
Power budget for FDDI	9.0–13.0 dB	6.0–7.0 dB
Maximum distance	1.6 km (.96 mi)	2 km (1.2 mi)
*Power budget contingent on core and cladding tolerances.		

The FDDI single-mode PMD standard (currently under development) describes the traits of single-mode fiber for long-distance links. Table 2–3 lists the single-mode fiber optic cable specifications.


Table 2–3: Single-Mode Fiber Cable Specifications

Recommended Cable	Parameters
Mode Field diameter	8.2 μm to 10.5 μm
Cladding diameter	125 $\mu\text{m} \pm 2 \mu\text{m}$
Fiber cladding noncircularity	2% maximum
Core to cladding concentricity error	1 μm maximum
Nominal operating wavelength	1300 nm
Fiber cutoff wavelength	1270 nm maximum
Zero dispersion wavelength	1300 to 1322 nm
Zero dispersion slope	0.095 ps/(nm ² km) maximum
Optical power attenuation	≤ 0.40 dB per km @ 1310 nm
Minimum power budget	22 dB @ 1300 nm
Minimum required loss	12 dB @ 1300 nm
Maximum distance	40 km (24.8 miles)

2.3.2 Optical Bypass Relay

The optical bypass relay (option) shown in Figure 2–2 can be used to maintain connectivity of the FDDI ring in the absence of power or during fault conditions in a node. The bypass relay allows the light to bypass the optical receiver in the faulty node. In this way, the faulty node is bypassed and the operation of the FDDI ring is maintained.

Figure 2–2: Optical Bypass



LKG-5092-91I

Optical bypass relays have a power penalty, however, which may cause the maximum allowable loss between stations to be exceeded. This limits the number of serially connected relays in the ring.

Other considerations when using optical bypass relays include:

- Bypass relays introduce additional loss in the network, and they do not perform repeater functions of amplifying and restoring the bit stream.
- By bypassing a station, the new distance between adjacent stations may exceed the maximum allowable value.
- Bypass relays, as any mechanical devices, may introduce less than reliable service to the network.


2.3.3 Connector Keying

The Media Interface Connector (MIC) is the device used in FDDI to connect a node to multimode fiber. The MIC is designed to align the fiber cable properly with the transmit and receive optics in the node. The MIC consists of a plug and a receptacle. The receptacle is keyed to prevent improper plug attachment. Proper keying prevents misconnecting port types, which avoids misconfiguring the network.

Four key types are defined for the MIC receptacle: MIC A and MIC B keying provide for the attachment of DAS and DAC into the dual ring; MIC M and MIC S keying provide for attachment of a SAS or a SAC (MIC S) to a concentrator (MIC M). Figure 2-3 shows the MIC receptacle keying design.

The position of the keyway on the MIC plug determines the type of receptacle to which it connects. A MIC S plug fits all receptacles. The keyway on the other MIC plugs and receptacles must match for proper mating.

Figure 2-3: MIC Receptacle Keying



Connector keying:

MIC A	Primary in/secondary out-DAS A port
MIC B	Primary out/secondary in-DAS B port
MIC M	Concentrator M port
MIC S	SAS S port

LKG-5072-90I

2.4 Physical Layer Protocol (PHY)

The PHY standard defines the media independent portion of the Physical layer in FDDI. The PHY establishes clock synchronization with the incoming data stream. The PHY standard defines the following components. The sections that follow discuss these components in greater detail.

- Symbols—The smallest primitive signaling element used by the MAC
- Line state—Indicates connection status between neighboring stations as determined by receiving station and based on symbols received from adjacent neighbor
- Encode/decode process—4B/5B and NRZ/NRZI data encoding
- Data recovery—Recovers data bits from the incoming signal
- Elasticity buffer—Compensates for clock frequency differences between neighboring stations
- Smoothing—Compensates for preamble bit loss
- Repeat filter—Prevents propagation of code violations and certain line states from inbound to outbound links
- Scrubbing—Prevents orphan protocol data units (PDUs) from continually circulating on the ring

To transmit data signals over the fiber, the PMD must first convert the electrical signals from the PHY to an optical data stream. In the optical data stream, a binary 1 is represented by a polarity transition high to low or low to high, a binary 0 is represented by the absence of a transition. A method must be provided for the receiving station to synchronize its clock to the transmitting station clock.


To solve this problem, the PHY encodes data in such a way as to guarantee the presence of transitions in the data stream. To achieve efficiency in data transmission, FDDI combines two types of data encoding: 4B/5B and NRZ/NRZI. This dual coding guarantees at least one transition for each transmitted symbol, ensuring recovery of the clock by the receiving device. Section 2.4.1 describes the use of symbols in FDDI; Sections 2.4.3 and 2.4.4 discuss the PHY encoding methods.

2.4.1 Symbols

Although data is transmitted over the fiber in a serial bit stream, the basic unit of information is the symbol. The symbol transmits information between stations on the FDDI network. The PHY specification defines the symbol set. FDDI defines three types of symbols: line state, data, and control. Each symbol consists of four bits. In the PHY, these four bits are passed through a 4B/5B encoder that generates a code group consisting of five bits.

The five-bit code group establishes the signaling rate for FDDI at 125 Megabaud to meet the 100-Mb/s data rate requirement. Figure 2–4 is a simplified diagram depicting the flow of symbols through the 4B/5B and NRZ/NRZI encoders onto the medium.

Figure 2–4: Symbol Flow Through the PHY



1111 is the symbol for F.
11101 is the code group for F.
The order of bit transmission is leftmost first.

LKG–4239–90i

Table 2–4 lists and defines the symbols used in FDDI. As shown, only 16 of the 32 symbols represent input data; the remaining symbols define control and line state conditions.

Table 2–4: Symbol Encoding

Symbol	Code Group	Assignment	Symbol Type
Q	00000	Quiet	Line state symbol
I	11111	Idle	Line state symbol
H	00100	Halt	Line state symbol
J	11000	J	First of sequential Starting Delimiter Pair
K	10001	K	Second of sequential Starting Delimiter Pair
0	11110	Zero	Data symbol
1	01001	One	Data symbol
2	10100	Two	Data symbol
3	10101	Three	Data symbol
4	01010	Four	Data symbol
5	01011	Five	Data symbol
6	01110	Six	Data symbol
7	01111	Seven	Data symbol
8	10010	Eight	Data symbol
9	10011	Nine	Data symbol
A	10110	A	Data symbol
B	10111	B	Data symbol
C	11010	C	Data symbol
D	11011	D	Data symbol
E	11100	E	Data symbol
F	11101	F	Data symbol
T	01101	T	Ending Delimiter
R	00111	R	Control indicator/logical zero (reset)
S	11001	S	Control indicator/logical one (set)
V or H	00001	V or H	Invalid code assignment
V or H	00010	V or H	Invalid code assignment
V	00011	V	Invalid code assignment
V	00101	V	Invalid code assignment

Table 2–4 (Cont): Symbol Encoding

Symbol	Code Group	Assignment	Symbol Type
V	00110	V	Invalid code assignment
V or H	01000	V or H	Invalid code assignment
V	01100	V	Invalid code assignment
V or H	10000	V or H	Invalid code assignment

Note: Order of transmission is leftmost bit first (12345).

2.4.2 Line States

FDDI reserves the symbols Q, I, and H, for PHY-to-PHY communications. Called line state symbols, they can be combined in various ways to determine and control the status of the physical connection between attached stations.

Line states are determined by the receiving PHY. They indicate a physical link condition that exists longer than can be represented by a single symbol or pair of symbols. SMT Connection Management (CMT) and Ring Management (RMT) services use line state symbols to verify and maintain ring integrity. The various line states are as follows:

Quiet Line State (QLS)	Indicates the absence of a physical connection. A station enters this state upon receipt of at least 16 consecutive Quiet (Q) symbols.
Master Line State (MLS)	Part of the physical connection process, a station enters this state upon receipt of 8 consecutive HQ or QH symbol pairs.
Halt Line State (HLS)	Part of the physical connection process, a station enters this state upon receipt of at least 16 consecutive H symbols.
Idle Line State (ILS)	Establishes and maintains clock synchronization on the outbound physical link. A station enters this state upon receipt of at least 4 consecutive Idle (I) symbols.

Active Line State (ALS)	Indicates that the inbound symbol stream is a MAC frame sequence and that the neighboring PHY has enabled the associated physical connection. A station enters this state upon receipt of a Starting Delimiter (JK) symbol pair.
Noise Line State (NLS)	Indicates noise on the inbound physical link and, if NLS persists, that the associated connection is faulty. A station enters NLS upon receipt of 16 noise events without satisfying the criteria for entry into another line state. Potential noise events include receipt of a Q, H, J, K, or V symbol.

2.4.3 4B/5B Data Encoding


To convert the 4-bit MAC symbols into 5-bit code groups, the PHY standard specifies the use of 4B/5B encoding. The fifth bit added to the symbol group is for clocking purposes and to provide control symbols that are unique from data. Figure 2–4 shows the symbol flow through the 4B/5B encoding logic. Properties of 4B/5B encoding include the following:

- Idle symbols have the maximum clock information for synchronization of station clocks.
- Maximum of one leading 0 and two trailing 0's in a code group to prevent more than three contiguous 0's in a data stream during active line state.
- Starting delimiter, JK, is always recognized regardless of symbol boundary.
- Line state recognition is independent of symbol boundaries.

2.4.4 Modified NRZI Encoding

Once a symbol is converted to a 5-bit code group, it is further encoded using a Non-Return-to-Zero/Non-Return-to-Zero Invert on Ones (NRZ/NRZI) coding scheme. NRZ/NRZI encoding allows data to be transmitted at 100% efficiency by means of polarity changes. A change in amplitude or optical power level indicates a one; no change indicates a zero. Figure 2–5 illustrates this concept.

Figure 2–5: NRZI Encoding Process



LKG-4062-90I

2.4.5 Clock Recovery and the Elasticity Buffer

At 100 Mb/s, the use of a centralized clocking scheme would put critical demands on the internal phase locked loop circuitry of each station. To eliminate the need for expensive circuits, FDDI specifies the use of distributed clocking on the network. Each station has an autonomous clock source for transmitting or repeating information on the ring. Each receiver recovers clock and data from its upstream neighbor.

Within the PHY standard, clocking is defined as part of the data receive function. The clocking circuitry within each station ensures that received data is properly decoded. A phase locked loop circuit derives a 125 MHz clock based on the incoming data. The resulting clock is used to sample and restore incoming data to logic levels.

To compensate for any possible clock frequency differences between stations, each station has an elasticity buffer. Data is clocked into the buffer at the clock rate recovered from the incoming bit stream. It is clocked out of the buffer at the station's internal clock rate. The MAC entity that originates a frame inserts at least 16 Idle symbols before the frame. The elasticity buffer in a subsequent station can change Idle symbol count as necessary to maintain clock synchronicity among attached stations.

The PHY standard specifies a frequency tolerance of .005% for a station, which translates to a .01% frequency variation between the incoming clock and the station internal clock. The standard further specifies a minimum station elasticity of ± 4.5 code bits. Multiplying the .01% frequency difference times the maximum frame size of 4500 bytes (45,000 code bits) results in the 4.5 code-bit elasticity buffer size.

2.4.6 Smoothing Function

As a MAC data frame circles the ring, it is possible that repeated elasticity buffer operations could cause sufficient preamble shrinkage to cause a station to lose frames. To prevent this, a smoother is built into each PHY. This function compensates for the possibility of preamble shrinkage by absorbing symbols from longer preambles and distributing them into shorter preambles. The smoother can insert or delete Idle symbols in the preamble as necessary to ensure the proper reception of frames.

2.4.7 Repeat Filter

The repeat filter prevents propagation of code violations and certain line states from the inbound link to the outbound link while permitting the propagation of valid frames. It also provides for the propagation of lost (damaged) frames so that they can be counted by the next MAC entity on the ring. This filter allows a station to pass data from its upstream neighbor to its downstream neighbor without involving the MAC entity.

2.4.8 Scrubbing

When a ring is reconfigured, attached devices can enter or leave the ring. During this process, orphan protocol data units (PDUs) are sometimes created. These orphan PDUs may not be identified as belonging to any particular attached device because they are generated by a device no longer part of the original configuration. For example, if a station suddenly fails, the station abruptly leaves the ring without completing normal exit procedures. The adjacent station cannot identify the PDUs left behind.

To remove these orphans, the adjacent station scrubs the ring when a station is inserted or removed. The adjacent station generates a series of Idle symbols while stripping all frames and tokens that are received. This operation is performed long enough to guarantee that all frames have been removed from the ring and that the ring will reinitialize. This prevents orphan frames from continually circulating on the ring and prevents old frames from being received long after they were originally sent.

2.5 Media Access Control (MAC)

An FDDI token ring consists of stations connected in series by medium forming a closed loop. Data is transmitted serially as a symbol stream from one attached station to its downstream neighbor. Each station in turn regenerates and repeats each symbol, passing it to its downstream neighbor.

MAC is the lower sublayer of the Data Link layer defined by the OSI reference model. The functions of the MAC include the delivery of LLC PDUs, frame construction, transmission, receiving, repeating, and stripping. The MAC is further responsible for claiming and beaconing. Error detection mechanisms are also defined in the MAC standard. Functions defined by the MAC standard include the following:

- Fair and equal access to the ring through use of the timed token
- Communication between attached devices using frames and tokens
- Ring initialization (claiming)
- Ring fault isolation (beaconing)

To gain the right to transmit on the ring, stations must first acquire the token, which is a unique symbol sequence that circulates around the ring following a data transmission. Once a station acquires the token, it removes the token from the ring and begins transmitting data. At the end of transmission, the station issues a new token onto the ring, providing other stations with the opportunity to transmit.

2.5.1 Ring Scheduling


One of the fundamental concepts of the FDDI network is the use of a timed-token ring architecture and a Timed-Token Protocol (TTP). To guarantee synchronous service time, the MAC standard specifies that a Target Token Rotation Time (TTRT) must be observed by all attached stations. TTRT is determined by mutual agreement among all active stations. The TTP guarantees that the token appears at a station by two times the TTRT. The MAC standard defines the timed token as the means for acquiring access to the ring in order to transmit data from one attached device to the next.

The TTP is a major difference from such commonly used LAN standards as IEEE 802.5. In some other LANs, a station captures the token and flips a bit within the token to indicate the station has captured it. In FDDI, a station captures the token, transmits for its allotted time period, and reissues the token onto the ring once it has completed transmission or has exhausted its available transmission time.

2.5.2 Data Link Addressing Scheme

Digital's FDDI devices employ the 48-bit addressing scheme defined by the FDDI standards. In this scheme, 48 bits specify the destination and source addresses of active nodes. Two of the bits also determine whether the address is an individual or group address, and whether address administration is local or by means of the IEEE address administrator. The 48-bit address must be unique for all nodes. No two nodes in the extended LAN can have identical individual addresses. Figure 2–6 shows the format and defines bits used in the address.

Figure 2–6: 48-Bit Address Format




In the destination address field, a 0 in the I/G bit indicates that the 48-bit address is an individual address; a 1 indicates that it is a group address. A 1 in the U/L bit indicates a locally administered address; a 0 in this bit position indicates a universally administered address. In the source address field when source routing is used, the I/G bit is set to 1. Otherwise, the I/G bit is set to 0.

Digital recommends using universally administered addressing to eliminate the risk of duplicate addressing on the extended network and to eliminate the need for the network manager to assign addresses. Because every address is unique, this technique also eliminates the need to resolve addressing conflicts when LANs are merged.

2.5.3 Universally Administered Addressing

Universally administered addresses ensure that any address assigned on the extended LAN is unique. This is done by the use of a hierarchical address administration. In the 48-bit addressing of FDDI, the 24 Organizationally Unique Identifier (OUI) bits are assigned to the owner of the OUI by the IEEE Address Administrator. The remaining 24 bits are assigned by the owner of the OUI. Controlling address assignments in this manner avoids address duplication on the LAN. Figure 2-7 shows the universal address format.

Figure 2-7: Universal Address Format



Although OUI assignments are specified as 24 bits, the owner of the OUI can set the I/G bit to 0 or 1 as needed to define an individual or group address. The U/L bit is set to 0 to indicate a universally administered address.

2.5.4 Canonical Addresses and the Order of Transmission

The IEEE canonical form for a 48-bit address represents the address as a 6 octet hexadecimal string, with hyphens separating the octets, as shown in the canonical address line of Figure 2–8 part A. The first octet contains the first 8 bits of the address, with the I/G bit as the least significant bit, the U/L bit as the next significant bit, and so on. (In Figure 2–8 part A, the first octet of the binary representation line is the leftmost octet.)

The IEEE strictly defines the order of transmission for addresses in the DA and SA fields. Addresses in the DA and SA fields must be transmitted I/G bit first, no matter what transmission order is used by the LAN. Transmission order is left to right, as shown in the transmission order line of Figure 2–8 part A.

Because FDDI transmits symbols MSB first, the correct symbols must be transmitted to guarantee that the DA and SA fields meet the IEEE transmission order requirements. The FDDI symbols line in Figure 2–8 part A shows the FDDI symbols that must be transmitted to produce the proper bit order. This order of transmission also applies to addresses carried in SMT frames.

Figure 2–8: Order of Transmission for Addresses

A Address in SA and DA fields:

Canonical address	DE–BC–68–A5–34–27					
Binary representation	11011110–10111100–01101000–10100101– 00110100– 00100111					
Transmission order	01111011 00111101 00010110 10100101 00101100 11100100					
FDDI symbols	7B	3D	16	A5	2C	E4

B Address in the INFO field:

Canonical address	DE–BC–68–A5–34–27					
802.3/802.4 transmission	01111011 00111101 00010110 10100101 00101100 11100100					
FDDI/802.5 transmission	11011110 10111100 01101000 10100101 00110100 00100111					
FDDI symbols	DE	BC	68	A5	34	27

LKG–4241–90I

The rules for order of transmission of addresses apply only to the DA and SA fields. Addresses carried as information in data fields should be transmitted in accordance with the transmission order rules for the specific LAN. For FDDI and 802.5 LANs, this is MSB first. For 802.3 and 802.4 LANs, it is LSB first.

The example in Figure 2–8 part B shows the transmission order, left to right, for the different LAN types. In the line labeled 802.3/802.4 transmission, the I/G bit of the address is the leftmost bit of the leftmost octet.

In the line labeled FDDI/802.5 transmission, the I/G bit is the rightmost bit of the leftmost octet. The FDDI symbols that represent this transmission order are shown in the FDDI symbols line. Addresses carried as data in the info field for interchange between higher layer protocols (for example, ISO 9542 or TCP/IP ARP) should be transmitted in their canonical (6-octet string) form.

2.6 Protocol Data Units

Protocol Data Units (PDUs) are the units of transfer between peer entities in the layered model. FDDI MAC defines two PDU formats: tokens and frames. Frames are further divided into MAC frames, Station Management (SMT) frames, Logical Link Control (LLC) frames, Void frames, and Implementer frames. Table 2–5 lists frame types and Frame Control values that define them. Sections 2.6.1, 2.6.2, and 2.7.5 describe each type of PDU in greater detail.

Table 2–5: FDDI Frames


Frame Control (FC) Bit Field	Frame Type	Description
CLFF FFFF		
0x00 0000	Void	Void frame
1000 0000	Token	Nonrestricted token
1100 0000	Token	Restricted token
1L00 0001 to 1111	MAC	MAC frame
1L00 0010	MAC	MAC beacon frame
1L00 0011	MAC	MAC claim frame
0L00 0001 to 1111	SMT	SMT frame
0L00 1111	SMT	Next station addressing frame

Table 2–5 (Cont.): FDDI Frames

Frame Control (FC) Bit Field	Frame Type	Description
CL01 r000 to r111	LLC	LLC frame
0L01 rPPP	LLC	Asynchronous priority transmission, priority determined by PPP in bit field
1L01 rrrr	LLC	Synchronous transmission
CL10 r000 to r111	Implementer	Implementer frame
CL11 rrrr	Reserved	Reserved for future use by standard
C = class r = reserved L = length x = don't care P = priority		

2.6.1 Tokens

The token is the means by which the right to transmit is passed from one attached node to another. In FDDI, stations acquire and absorb the token when they want to transmit data. Immediately after data transmission, the station puts a new token onto the ring. Figure 2–9 illustrates the MAC token structure. Section 2.6.3 defines token fields.

Figure 2–9: Token Format


PA = Preamble — 16 symbols
 SD = Starting delimiter — 2 symbols
 FC = Frame control — 2 symbols
 ED = Ending delimiter — 2 symbols

LKG–4063–901

2.6.2 Frames

Although a token transfers the right to transmit, frames are the means for carrying MAC control data or user information among stations. Frames are divided into several fields, each containing information that defines a specific frame parameter. Figure 2–10 shows the FDDI frame format and component fields. Each frame field consists of a number of FDDI symbols. Refer to Section 2.4 for information concerning symbols and data encoding; Section 2.6.3 provides frame field definitions.

Figure 2–10: FDDI Frame Format



LKG–4066–90I

2.6.3 Frame and Token Fields

Frame and token fields are defined as follows (refer to Table 2–5 for bit and field positions within the frame or token):

Preamble (PA)	This field is used for Receive clock synchronization in frames and tokens. The preamble is transmitted with at least 16 Idle symbols. Frames can be received that contain more or fewer Idles, depending on station behavior.
Starting Delimiter (SD)	A unique 2-symbol pattern that indicates the start of the frame or token. In FDDI this pattern is always the symbols J and K. No PDU is considered valid unless it begins with this sequence.

Frame Control (FC)	A 2-symbol field that defines frame or token type. The bits in this field are defined as follows from left to right: C = Class bit, indicating class of service: synchronous or asynchronous. In Digital's implementation, it is set to 1 indicating asynchronous frames. L = Address Length bit, indicating either 48- or 16-bit addressing. In Digital's implementation, it is set to 1 indicating 48-bit addressing. FF = Format bits. Used in conjunction with the other FC bits, they define the frame type being transmitted: Void, Nonrestricted Token, Restricted Token, SMT, MAC, LLC, Implementer, and Reserved. ZZZZ = Control bits used with the CLFF bits to define the frame.
Destination Address (DA)	This 48-bit field identifies the station to receive the frame. The DA can be a single station address, a broadcast address, or a multicast address.
Source Address (SA)	This 48-bit field identifies the originator of the frame.
INFO	This field contains the data to be transmitted over the ring. It can be either MAC control information or user data. Field length is variable but is limited by the maximum FDDI frame size of 4500 bytes.
Frame Check Sequence	This 32-bit field contains a Cyclic Redundancy Check (CRC) value based on the contents of the FC, DA, SA, INFO, and FCS fields. The receive station uses this value to verify the accuracy of the transmitted frame.
Ending Delimiter (ED)	This field indicates the end of a frame. For tokens, the ED is always two T (Terminate) symbols. For frames, the ED is always one T symbol.

Frame Status (FS)	This field consists of the Error Detected (E), Address Recognized (A), and Frame Copied (C) indicators. These indicators can contain an Reset (R) or Set (S) symbol.
-------------------	--

2.7 Station Management (SMT)

Station Management (SMT) provides the necessary services at the station or local level to monitor and control an FDDI node. It allows stations to work cooperatively over the ring, and ensures proper station operation. SMT provides the following services to an FDDI node:

- Connection management (including insertion and removal)
- Configuration management
- Encoding of SMT frames
- Fault isolation and recovery

Although FDDI stations can have multiple instances of the MAC, PHY, and PMD, all FDDI station types have only one SMT entity. SMT contains three major protocols: Connection Management, Ring Management, and SMT Frame Services. The three SMT protocols are discussed in the following sections. Figure 2–11 shows how SMT is incorporated into the FDDI architecture.


2.7.1 Connection Management (CMT)

Connection Management (CMT) is that portion of SMT that performs Physical layer insertion and removal, and controls the configuration of PHYs and MACs within a node. CMT controls the connection of nodes to the FDDI network, the physical connection of PHYs in separate nodes (including concentrators), and other functions such as scrubbing and trace. Figure 2–11 shows CMT in relation to other elements of SMT. The CMT components are as follows:

- Physical Connection Management (PCM)—Provides for management of the physical connection between adjacent PHYs. This includes connection establishment, link confidence testing, and enforcement of connection rules.

- Configuration Management (CFM)—Provides for the configuration of PHY and MAC entities within a node. There is only one CFM entity in a station; there is only one CFM entity for each PHY in a concentrator.
- Entity Coordination Management (ECM)—Provides for control of bypass relays and signaling to control PCM.

Figure 2–11: Station Management Architecture




LKG-4067-90I

2.7.2 Connection Rules

In the FDDI topology, there are four physical port (PHY-port) types: A, B, M, and S. Figure 2–12, a simplified diagram of a concentrator, shows the various PHY-ports used in FDDI. These PHY ports are defined as follows:

- **PHY A**—Connects to the incoming primary ring and the outgoing secondary ring of the FDDI dual ring. This port is part of a dual attachment station (DAS) or a dual attachment concentrator (DAC).
- **PHY B**—Connects to the outgoing primary ring and the incoming secondary ring of the FDDI dual ring. This port is part of a DAS or a DAC. This port is also used to connect the DAS to the concentrator.
- **PHY M**—Connects a concentrator to a SAS, DAS, or another concentrator. This port is implemented only in the concentrator.
- **PHY S**—Connects a SAS to a concentrator PHY-M port. PHY S also connects to a DAS or another SAS, but these configurations are not always recommended.

Figure 2–12: PHY- Port Types



In the ANSI FDDI standards, connection rules are specified to ensure against the construction of illegal topologies. Connections between stations have the status of in use or not in use. When a cable is connected a decision is made to use or not use that cable in the token path. The ring connection rules for Digital devices are listed in Figure 2–13. In the figure, **This Node** is the reference point for all connection decisions.

Figure 2–13: FDDI Connection Rules Matrix

		This Node				
		PHY port	A	B	M	S
Other Node	A	No	Rule 1	Yes	Yes	
	B	Rule 2	No	Yes	Yes	
	M	Rule 3	Yes Note 1	No	Yes	
	S	Yes	Yes	Yes	Yes	

Yes = Accepts this connection
 No = Does not accept this connection

LKG-4810-91I

- Rule 1: PHY B of **This Node** accepts connection to PHY A of **Other Node** if:
- PHY A of **This Node** is not connected, or
 - PHY A of **This Node** is connected to PHY B or PHY S of another node.
- Rule 2: PHY A of **This Node** accepts connection to PHY B of **Other Node** if:
- PHY B of **This Node** is not connected, or
 - PHY B of **This Node** is connected to PHY A or PHY S of another node.
- See also Note 2.

- Rule 3: PHY A of **This Node** accepts the connection to PHY M of **Other Node** if:
- PHY B of **This Node** is not connected, or
 - PHY B of **This Node** is connected to PHY A or PHY S of another node.
- See also Note 2.

Note 1: **Backup tree connection (dual homing)**

The connection of PHY B of **This Node** to PHY M of **Other Node** always takes precedence over the connection of PHY A of **This Node** to PHY M of another node.

When a connection between PHY A of **This Node** and a PHY M or PHY B of another node exists, it is disabled when PHY B of **This Node** connects to PHY M of **Other Node**.

If the connection breaks between PHY B of **This Node** and PHY M of **Other Node**, the connection is reestablished between PHY A of **This Node** and PHY M of another node.

Note 2: **Tree connections are preferred over dual ring connections**

Tree connections are preferred over dual ring connections when only one of two connections may be accepted.

The connection of PHY A or B of **This Node** to PHY M of **Other Node** always takes precedence over the connection of the remaining PHY (A or B) of **This Node** to PHY A or PHY B of another node.

The connection is disabled between the remaining PHY (A or B) of **This Node** and PHY A or PHY B of another node when PHY A or PHY B of **This Node** connects to PHY M of **Other Node**.

2.7.3 Ring Management (RMT)

The Ring Management (RMT) portion of SMT receives status information from the MAC and CMT configuration logic. RMT then reports this status to the SMT and higher level processes. Services provided by RMT include stuck beacon detection and resolution by means of Tracing, MAC availability for transmission, and duplicate address detection and resolution.

In the case of a stuck beacon (a station is locked into continuous beacon frame transmission), a stuck-beacon timer in RMT measures the length of beacon transmission. If beaconing exceeds the timer limit, RMT begins stuck-beacon recovery procedures to restore the station to normal operation, if possible.

This process begins with the transmission of a directed beacon, which informs the ring of the stuck condition. If the ring has not recovered by the end of this transmission, then RMT initiates the Trace function. Trace provides a recovery mechanism, using Physical layer signaling, for stuck beacon conditions on the FDDI ring when a stuck beacon cannot be localized to a single link. Refer to Section 2.10.2 for a detailed discussion of the Trace function.

Duplicate addresses on the FDDI ring can prevent the ring from becoming operational. If RMT logic detects that its MAC has a duplicate address, RMT takes action to allow the ring to resume normal operation.

2.7.4 SMT Frame Services


SMT frame services provide control and observation mechanisms for the FDDI network. These services are implemented by means of different SMT frame classes. Within each frame class, different frame types are also defined. Frame class identifies the function the frame is performing. Type indicates whether the frame is an announcement, request, or response to a request. The Header field of an SMT frame contains class and type definitions for the SMT frame. These classes include, but are not limited to, the following:

- Neighborhood Information Frames (NIF) are used by a station to announce its address and a description of the station. Stations also use these frames to request this information of a neighbor station.
- Station Information Frames (SIF) are used to request, and provide in response to another request, a station's configuration and operating information. In addition, SIF frames contain information about the attachment status of each port. This information can be used to create a physical ring map that show the token path and the topology.
- Request Denied Frames (RDF) are generated in response to an unsupported frame request.

2.7.5 SMT Frames

The SMT frame consists of three fields as shown in Figure 2–14. The MAC header portion identifies the frame as an SMT frame and contains the source and destination addresses. The SMT header identifies the SMT class and type of the frame. The SMT INFO field contains the SMT parameters to be transmitted.


Figure 2–14: SMT Frame Format



2.8 Logical Link Control (LLC)

Although not part of the standard, FDDI requires LLC for proper ring operation. It resides in the second, or Data Link, layer of the OSI model. As in most Data Link standards, LLC defines the transmission of a frame of data between two stations with no intermediate switching nodes. Figure 2–15 shows the relationship between LLC and FDDI. FDDI assumes implementation of the IEEE 802.2 LLC standard.

Figure 2–15: Logical Link Control and FDDI



Data is transmitted between LLCs by means of LLC PDUs, which are frames with Frame Control fields designating the INFO field as containing LLC information. As part of the LLC PDU, a service access point address is included indicating that the receiver is a particular LLC client. LLC frames can cross bridges into the extended LAN; MAC and SMT frames cannot do this.

2.9 Ring Operation

This section describes basic FDDI ring operation. It provides an overview of ring functionality using the concepts described in the preceding sections. It also discusses timing, steady state operation, ring initialization, and the claiming process. Section 2.10 describes the beacon process and the trace function. Because of the complexity of FDDI, it is beyond the scope of this document to describe every detail of ring operation.

2.9.1 Asynchronous Ring Transmission

In Digital's FDDI implementation, all attached stations use asynchronous data transmission to pass data around the ring. In asynchronous data transmission, all attached stations are dynamically allocated a transmission time based on the TTRT (discussed previously). Stations acquire the token and transmit until the Token Holding Time (THT) expires.

Asynchronous transmission, when used in the context of the FDDI environment, pertains to ring traffic and services that are not extremely delay sensitive. Do not confuse this usage with the asynchronous transmission of data on the physical media, which refers to the intercharacter timing of a transmission as used on low speed lines for devices such as ASCII terminals. Digital's FDDI devices receive synchronous and asynchronous transmissions, but they transmit asynchronous only.

2.9.2 Timers

Each station on the FDDI ring uses three timers to regulate operation of the FDDI ring: the Token Rotation Timer (TRT), the Token Holding Timer (THT), and the Valid Transmission Timer (TVX). These timers are administered locally by the individual stations.

The TRT, which times the period between the receipt of tokens, is critical to the successful operation of the FDDI network. It is initialized to different values depending on the status of the ring. This timer controls ring scheduling during normal operation and it is also used for timing error recovery in the node. The TRT is set during normal operation to the value of the Target Token Rotation Time (TTRT) negotiated by stations attached to the FDDI ring. Section 2.9.3 discusses this in more detail.

The THT controls the length of time during which a station can start transmission of asynchronous frames. Attached stations can begin asynchronous transmissions if THT has not expired. THT is initialized with the amount of time remaining until TTRT expires when the token is captured. The station can then initiate asynchronous transmission until THT expires.


The TVX times the period between valid transmissions on the ring and allows the station to recover from transient ring errors. The TVX is used in the detection of excessive ring noise, token loss, and other faults. It is reset on receipt of a valid frame or token.

2.9.3 Target Token Rotation Time, Claiming, and Initialization

This section describes normal station power up and the physical connection process, which is controlled by CMT using PHY signaling. At power up, stations recognize their neighbors on the fiber by transmission of a defined line state sequence (QLS–HLS–ILS). Once stations acknowledge each other, they exchange information on port type and connection rules. They negotiate link confidence test (LCT) duration, run LCT, and exchange results. If LCT passes, they complete the physical connection by transmitting another defined line state sequence (HLS–MLS–ILS).

Once the physical connection is complete, FDDI requires the establishment of a TTRT to ensure that stations receive guaranteed service time. The TTRT default value in the Digital environment is 8 ms; this value is not normally changed. If necessary, TTRT can be set to a different value based on the number of attached stations, the length of the ring, and the time required by each station to transmit data over the ring. TTRT is negotiated during ring initialization by means of a claim token bidding process (described in the following paragraphs). Figure 2–16 shows how the claiming process operates.

Figure 2–16: Claiming Process



1. All active stations start issuing claim frames.
2. Station 2 receives a shorter claim and yields.
3. Station 3 receives a longer claim and continues sending claim frames.
4. Stations 1 and 4 receive a shorter claim and yield.
5. Station 3 receives its own claim frame signaling that it has won the bidding process.
6. Station 3 issues the token.

LKG-4064-90I

The claiming process determines which station will initialize the ring. The process begins when one or more MACs enter the claim state. In this state, the MAC continually transmits claim frames. The claim frame contains the originator's address and the station's bid (called T_Bid) for TTRT.

Receiving stations take incoming claim frames and compare the received TTRT with their bid for TTRT. If the received T_Bid represents a shorter (in time) TTRT than the station's own bid value, then the station repeats the claim frame and stops transmitting its own claim. If the received bid value represents a longer (in time) TTRT, then the station strips the claim frame and continues transmitting frames with its own T_Bid onto the ring. The claim winner is determined when a station receives its own claim frame. In the event that two or more stations make identical bids, then the station with the longest and highest address wins the bidding.

Eventually, the ring is filled with claim frames from the winning station, which initializes the ring by issuing a nonrestricted token. This token passes around the ring without being captured by any station. Each station in turn, upon receipt of the initial token, resets its TTRT to the negotiated value (T_Neg) and sets its Ring Operational bit. Synchronous transmissions are allowed on the second token rotation. On the third token rotation, nodes can begin asynchronous data transmission. At this point, the ring is in steady state operation.

Because FDDI uses nonprioritized, asynchronous transmission with a nonrestricted token, bandwidth is time-sliced among all active stations on the ring. This mode supports fair access to the ring at a frame-level granularity.


2.9.4 Steady State Operation

In steady state operation, the ring has been initialized and all attached stations are fully functional. A station that is ready to transmit waits for a token. The station captures the token from the ring, which allows it to transmit. The station continues initiating transmissions until either it has no more data to transmit or the THT expires.

Capturing the token in an FDDI ring differs from the process used in an 802.5 ring in which a station ready to transmit captures the token and flips a bit within the token. A station ready to transmit data in FDDI must wait until it detects a token. The station then captures the token, stops the repeat process, and begins transmitting frames. After the station captures the token, there is no token on the ring, effectively prohibiting other stations from transmitting data onto the ring.

As each station receives a frame, it checks for its own address and for frame errors. Stations use the Frame Status field, which consists of a number of symbols called Control Indicators, to signal frame condition. If the frame is destined for the station that has received it, the station sets the A Control Indicator in the Frame Status field. It then copies the frame and sets the C indicator, if the copy is successful. Repeating stations simply retransmit the data frames onto the ring to their downstream neighbor while checking for errors. Figure 2–17 shows the process of capturing the token and transmitting a frame on an FDDI ring.

Figure 2-17: Capturing the Token



LKG-4070-901

The receiving station then inserts a symbol in one or more of the frame status indicators depending on the status of the incoming frame.

- An S symbol in the E indicator signals that an error has been found in the frame.
- An S symbol in the A indicator signals that the receiving station recognized its own address.
- An S symbol in the C indicator signals that the receiving station has copied the frame.

Depending upon actions taken by layers above the Data Link layer, the MAC can make use of, or ignore, the Frame Status Control indicators. If used, the combination of the A indicator set and the C indicator reset can be interpreted as notification that the destination station did not copy the frame. Since bridges do not set the A indicator when forwarding frames, and since the addressed station may exist elsewhere in the extended LAN, it is not correct to interpret the A indicator reset as notification that a station does not exist. In most applications, the A and C indicators are ignored.

All stations, except the transmitter, receive and then repeat each frame from the transmitting station. This results in a delay being introduced on the ring by each repeating station. The total delay introduced by the transmitting station and all repeating stations, plus the cable propagation delay, is the ring latency. As soon as a station finishes transmitting, it places a new token on the ring. If the ring latency is greater than the length of a station transmission, it is possible for a token to be issued before a frame is stripped from the ring.

Frames placed onto the ring by the transmitting station, after circulating around the ring, are stripped by the transmitting station. Stripping causes partial frames, called fragments or remnants, to be left on the ring after a station transmits. Logic circuitry in the FDDI stations ensures that frame remnants do not degrade ring operation. Frame fragments are removed from the ring when they encounter a transmitting station.


2.10 Ring Maintenance

The responsibility for monitoring ring integrity is distributed among all attached stations. Each station monitors the ring for conditions that require ring initialization. These conditions include ring inactivity longer than TVX, invalid frames on the ring, or a physical or logical break in the ring. To detect a lack of ring activity, stations use the internal timer TVX to track the length of time between reception of valid frames or tokens. If this time exceeds the value in TVX, an error condition is indicated and stations enter the claim process. If the claim process fails, the station enters the beacon process.

Stations leave the steady state whenever a ring or station malfunction occurs. Figure 2–18 shows the events that occur when there is a break in the fiber of an FDDI ring. In Figure 2–18, the token path on the primary ring is from station A to station B, and so forth. If a break in the primary ring occurs between station A and station B, PCM at PHY A of station B sees the loss of optical signal on the input fiber. Station B then shuts down the connection, which turns off the optical transmitter on the secondary ring (that transmits to PHY B of station A). Station B then wraps.

PHY B of station A detects loss of signal on the receiver, so it shuts down the connection and wraps. Stations A and B then attempt to restart the physical connection, but the PCM protocol detects the break and does not complete the initialization sequence. The ring remains operational because stations A and B are wrapped until the fault is corrected or until another fault partitions the ring.

Figure 2–18: FDDI Typical Ring Recovery



1. A break in the primary ring causes port A of station B to detect signal loss and wrap.
2. Port B of station A also detects signal loss and wraps.
3. Stations A and B attempt to restart the physical connection. SMT PCM detects the break and does not complete initialization.

Note: All devices on the ring are dual attachment stations or dual attachment concentrators.

LKG-4071-901

2.10.1 Beaconsing

If a ring malfunction causes a claim to fail, or if SMT requests it, an FDDI station can start the beacon process. In this state, a station begins continuous transmission of a beacon frame (refer to Table 2–5 for frame structure). The MAC uses this frame to announce that the ring is broken and to locate and isolate the fault, if possible.

If a station receives a beacon from its upstream neighbor, it stops beaconing and repeats the upstream beacon. If no beacon is received, the station transmits a continuous beacon stream. If a station receives its own beacon, it assumes the logical ring is restored, ceases beacon transmission, and begins the claim process.

2.10.2 Trace

The trace function uses PHY signaling to recover from a stuck beacon condition. In the event of a stuck beacon condition, the fault is localized to the beaconing MAC and its nearest upstream neighbor MAC. The trace function causes all stations and concentrators in the suspect section of the ring to leave the ring and run Pathtest in an attempt to localize the fault.


Stations run Pathtest to check the components of the failed path. This test checks the transmit and receive components of the path, including the optics and switching logic. If Pathtest finds a fault, the station remains off the ring.

Figure 2–19 illustrates the trace function. The transmitter in station A is defective and begins issuing a continuous stream of Idle symbols onto the ring. The TVX in station B times out. Station B attempts to recover by claiming—and failing that—by beaconing. Finally, station B determines that it is stuck beaconing and begins a trace.

The trace signal (MLS) from station B propagates from B to A over the secondary ring. Stations A and B leave the ring and begin running Pathtest. Running Pathtest should allow station A to detect its transmitter fault. If it does not, the fault is detected when stations A and B attempt to restart the connection. PCM will not complete the initialization and the connection remains down (as in the example of Figure 2–18).

Although this example deals with stations on the dual ring, trace operates the same way in a tree topology. The trace always occurs between two MACs on the ring, with the trace being propagated upstream over the transmit fiber of the ring from the beaconing MAC back to the defective station.

Figure 2–19: Detection of Defective Station



1. Defective station A transmits continuous idle symbols.
2. Downstream station B fails the claim process and begins issuing beacon frames.
3. Station B detects a stuck beacon and initiates the trace function by sending MLS to its upstream station.
4. Stations A and B leave the ring and begin Pathtest.
5. If Pathtest detects a fault in station A, station B rejoins the ring and A remains down.
6. If Pathtest fails to detect the fault, stations attempt to rejoin the ring. At this time, PCM prevents defective station A from reconnecting to the ring.

LKG-4243-91I

2.11 Digital Enhancements

Digital has developed two algorithms for FDDI: the Frame Content Independent Stripping algorithm and the Ring Purger algorithm.

2.11.1 Frame Content Independent Stripping Algorithm

On FDDI a MAC entity (also called a MAC) is required to remove (strip) its own transmitted frames after exactly one traversal of the frames around the ring. The ANSI FDDI MAC Standard uses a Source Address Match algorithm for frame stripping, where a received frame or fragment is stripped if the Source Address field matches the MAC's address.

Digital developed the Frame Content Independent Stripping (FCIS) algorithm for use by a station to remove or strip frames that it has transmitted onto the ring. The algorithm is also called the Bridge Stripping algorithm.

The algorithm allows frames to be stripped independently of the content of the frame. This is extremely important for implementations, like bridges, that have to transmit frames with source addresses of other stations on different LANs. The algorithm allows implementations like bridges to strip frames without having to do any address comparisons in real time. The FCIS algorithm interoperates with and is compatible with implementations of the ANSI FDDI MAC standard.

The FCIS algorithm is based on stripping the same number of frames that the station has transmitted since the last token capture. It does so by keeping a local count of frames transmitted but not yet stripped. It strips frames until the count is zero.

Each time a frame is transmitted the algorithm increments the count. Each time an error-free frame is received and stripped, the algorithm decrements the count. Also, if the count is greater than zero, the algorithm strips frames even with a Frame Check Sequence error without decrementing the count.

In addition to the count mechanism, the algorithm requires the transmission of a well formed Void frame (Void with FCS), which is transmitted at the end of transmitting the frames and prior to transmitting the token. When the well formed Void frame is received, and if it is error free, the local count is set to zero unconditionally. On receipt of a token, claim, or beacon frame the count is set to zero. When the count is zero, the MAC strips frames based on the Source Address Match algorithm.

The ANSI FDDI MAC standard defines a well formed Void frame as a special frame that causes the Valid Transmission Timer (TVX) in all stations to be reset. The standard further stipulates that Void frames are not to be copied by any MAC. All well formed Void frames used by FCIS contain the transmitter's address as the source address so that they can be properly stripped by the transmitter. The well formed Void frame is 17 octets long, which is the minimum size frame allowed by the standard.

Digital presented the FCIS algorithm to the ANSI FDDI Committee in November 1988 and in June 1989, where it was highly regarded. The presentation showed the committee that the algorithm has all the right properties and is robust. Digital also published and presented papers on this algorithm at SIGCOMM 90 and at the Local Computer Network Conference in September and October 1990.

2.11.2 Ring Purger Algorithm

On all types of token rings, a frame can circulate indefinitely if it is not removed. A frame that is not removed (stripped) after its first traversal of the ring is called a No Owner frame (NOF). While the occurrence of an NOF may be rare, the consequences of allowing the circulation of NOFs are severe. Station resources are consumed by receiving these No Owner frames, possibly several times. For example, a single NOF on FDDI may cause a station to receive duplicate copies of a frame at the rate of up to 20,000 frames per second and may consume almost 80 Mb/s bandwidth on the ring and the receiving station.

The ANSI FDDI standard does not specify a guaranteed method of removing NOFs. It should be noted that other token rings, such as the IEEE 802.5 token ring, have mechanisms for removing NOFs. In the IEEE 802.5 token ring, one of the tasks performed by the monitor station is to remove NOFs.

Digital has invented a simple, yet robust, algorithm (called the Ring Purger algorithm) that corrects problems posed by NOFs by removing them reliably from the ring. The mechanism used by a Ring Purger is extremely simple, requiring the transmission of Void frames, which are not received by the stations on the ring. The ring bandwidth consumed is small and is an insignificant proportion of the total ring bandwidth.

The Ring Purger algorithm consists of two parts: an election algorithm to designate, in a distributed manner, a station as the purger; and the purging algorithm itself. Any one of the Digital stations in the ring can be elected to act as the ring purger. A unique station capable of being the ring purger is not needed. All Digital stations participate in this set of algorithms.

The primary purpose of the Ring Purger election algorithm is to ensure that there is one, and only one, ring purger operating in the ring most of the time.

The Ring Purger algorithm first elects one Digital station as the ring purger. The designated ring purger station then removes (purges) NOFs and fragments from the ring transparently. Although the ring purger is not part of ANSI FDDI standard, it inter-operates with and is compatible with implementations compliant with the ANSI FDDI standard.

The Ring Purger algorithm recovers from ring initializations and from failure of a station acting as the ring purger. Each time the ring is initialized, a Digital station becomes the ring purger if it is the FDDI MAC Claim token winner or if it was the ring purger prior to the ring initialization.

A ring purger announces its presence periodically by transmitting a Purger Hello frame. If there are multiple stations capable of acting as ring purgers, the station with the highest address is elected the ring purger. If there is no ring purger after ring initialization or any time during ring operation, a new ring purger is elected using a distributed election protocol.

The ring purger election protocol frames (for example, Purger Hello frames) use the SMT Extended Service frames. The election algorithm is designed to elect one, and only one, station to be the ring purger. The overhead of the purger election algorithm is negligible because the ring purger only periodically (approximately every 10 seconds) sends a short SMT Extended Service frame to the other participating stations on the ring.

The removal of NOFs from the ring is achieved by the purging algorithm, which operates as follows. The ring purger waits for a token. Once it captures a token, if it has frames to transmit, it transmits the frames. Once the transmission is completed, the ring purger transmits two well formed Void frames followed by the token. These Void frames are extremely small (17 bytes long) and consume little, if any, of the ring's usable bandwidth. The Void frames demarcate the end of transmission of the ring purger's frames. The Void frames are transmitted with the destination and source address equal to the My Long Address (MLA) of the ring purger.

When the ring is operating normally, the ring purger receives only the frames it sent followed by the Void frames. The ring purger strips the ring of all frames and fragments until it receives one of its Void frames. It then strips the two Void frames it transmitted and ceases stripping frames from the ring. If there are NOFs circulating on the ring, the ring purger also strips these NOFs while stripping all the frames it receives before receiving its Void frames. The ring purger does not begin another purging operation until it receives the token. The ring purger ceases to operate when it sees a Beacon or Claim frame.

An active ring purger ensures that no frame on the ring circulates as an NOF for more than one traversal of the ring. These frames will therefore be received at most twice by the destination stations. The presence of the ring purger strictly limits the degradation in performance of the stations on the ring and the bridged-LAN due to NOFs.

The impact of the ring purger on the performance of the ring and the stations on the ring is negligible. Stations compliant to the ANSI FDDI MAC standard experience no overhead because the Void frames have no effect. The ANSI FDDI standard specifies that stations are not to copy Void frames. This prevents station resources from being consumed by the Void frames.

The ring bandwidth used by the ring purger is also negligible, because the two Void frames are transmitted only when a token is received. Also, the transmission of two Void frames takes only 4.48 microseconds. The Ring Purger algorithm is designed so that the bandwidth consumed becomes smaller as the load on the ring increases. When the stations on the ring use more of the available bandwidth, the overhead consumed by the ring purger decreases.

The ring purger transmits two Void frames each time it sees the token. On an idle ring, the ring purger transmits Void frames at a high rate. However, under this idle condition, the Void frames are only using bandwidth on the ring that is being wasted as a result of no stations transmitting on the ring. When the ring is idle, the consumption of ring bandwidth by the Void frames is of no consequence to the stations, which also do not receive these frames. As soon as the load on the ring is increased, the rate of Void frames decreases, until it finally reduces down to two Void frames every TTRT period.

Under heavy load, the ring purger transmits the Void frames less frequently. For example, under a very heavy load on the ring, it can take up to a TTRT (which can range from 4 to 173 milliseconds) for the token to make one traversal around the ring. Thus, the load posed by the ring purger's Void frames backs off when there is a load on the ring. In the worst case, on a heavily loaded ring, the ring purger's effect on the usable bandwidth is less than 0.22% (approximately).

The ring purger algorithm cleans NOFs from the ring transparently, and therefore, helps to improve the robustness and reliability of the FDDI token ring operation.

Concentrators

The concentrator plays a central role in FDDI topologies. It provides a connection point for end stations or other concentrators in a physical tree. It supports building and campus wiring topologies through the creation of hierarchical trees. It also supports the dual ring design specified in the FDDI standards. This chapter provides a general explanation of the concentrator and its role in FDDI applications. Section 3.5 provides configuration guidelines for using the concentrator in an FDDI network.

Digital's concentrator, the DECconcentrator 500, provides the flexibility to accommodate the diverse building wiring schemes found at customer sites. It is the key component in the Digital dual ring of trees topology. It also supports small workgroup topologies when used in a standalone configuration. The concentrator enhances reliability and maintainability in both large and small networks.

3.1 Description


The concentrator is an optoelectronic device that connects stations or other concentrators to the FDDI network. It maps the logical token ring into a hierarchical physical tree. The concentrator is an active device that decodes, retimes, and modifies, if necessary, the data stream as it passes through the concentrator. By cascading concentrators, network builders can create dual ring of trees topologies. In a standalone topology, it is the hub in a workgroup topology.

The main advantage of the concentrator is that it removes control of ring topology from individual stations and places it in the concentrator itself. This feature increases the reliability of the network for all users. Network managers can use the concentrator to alter the network topology by selectively removing or inserting attached nodes.

The FDDI standard defines two concentrator types: a single attachment concentrator (SAC) and a dual attachment concentrator (DAC).


The DAC has all the necessary components for attaching to the dual ring or into a tree (acting as a SAC). The SAC cannot connect directly to the dual ring because it lacks the necessary hardware and logic. Figure 3–1 and Figure 3–2 show two concentrator topologies, the standalone and the dual ring of trees topology.

Figure 3–1: Concentrator in a Standalone Topology



LKG-4050-90I

Figure 3–2: Concentrators in a Dual Ring of Trees Topology




LKG-4244-901

Figure 3–3 illustrates the basic architecture of a concentrator. The optional MAC entity provides MAC services to the concentrator. If the MAC entity is implemented, it is placed just in front of the outbound port. This helps to ensure that SMT station mapping can be performed accurately. The configuration switch provides internal configuration capabilities to the concentrator, allowing it to disconnect any or all attached nodes from the ring. With the configuration switch, the concentrator can route the data path as required according to the number of attached devices. This switch also disconnects nonfunctional stations attached to the concentrator.

As shown in Figure 3–3, the concentrator contains a number of PHY and PMD entities. Each PHY/PMD combination makes up an FDDI PHY port. These PHY ports are full duplex with a transmit and receive function.

Figure 3–3: Concentrator Architecture



LKG-4073-901

Concentrator PHY ports can be type A, B, M, or S. The DECconcentrator 500 has PHY ports type A, B, and M only. Type A and B ports connect the concentrator to the primary and secondary rings of the dual ring. Type B ports can also connect one concentrator to the M port of another concentrator. Type M ports can connect stations or other concentrators to the FDDI network. Refer to Figure 2–13 for device connection rules.

Concentrators with type S ports are part of a SAC. They connect the SAC to the FDDI network through the M port of another concentrator.

3.2 Concentrator Functions

The concentrator performs three reconfiguration functions: station insertion, station bypassing, and scrubbing. Station insertion and station bypass are performed by the configuration switch, based on inputs from SMT. Ports can be switched into and out of the ring either by control of SMT Connection Management or with a management command. Refer to Section 2.4.8 for a description of the scrubbing function.

Depending on the requirements of a network or the desired application, the concentrator can be equipped with a MAC or it can be MAC-less. Section 3.2.1 describes the two available concentrator types and their differences. The remaining sections describe MAC functions in further detail.

3.2.1 MAC and MAC-less Concentrators

The ANSI standards specify that concentrators can be constructed with or without the MAC. The addition of a MAC allows the concentrator to be managed by a remote management station using the same FDDI (in-band remote management).

Concentrators normally operate without need of manual intervention and provide all basic services independent of MAC capability. A concentrator can be managed through a local management interface. A network manager can observe and control the concentrator directly, or remotely through a communications channel. Remote management may be provided through the same FDDI (referred to as in-band remote management) or through a separate channel (referred to as out-of-band remote management).

Management adds the capabilities of remote monitoring and network configuration. In large networks, this functionality is usually required to ensure proper maintenance and control of the network.

Concentrators automatically configure the ring and connect stations to the ring through the services of SMT Configuration Management (CMT) and Ring Management (RMT). MAC-less concentrators are most useful in cases either where out-of-band remote management is needed for reliability or where local management is sufficient, such as in standalone topologies serving a limited geographical area. In these cases, in-band management is an unnecessary expense.

3.2.2 Station Bypass

Station bypass allows the concentrator to logically and physically disconnect the attached station from the ring. The concentrator does this according to a specific set of conditions as set forth by SMT Physical Connection Management (PCM). The occurrence of any one of these conditions results in the station being removed from the ring. The concentrator will disconnect the station if one of the following events occurs:

- The attached station is powered off or the media is unplugged.
- The concentrator PCM detects a defective connection.
- The concentrator PCM is explicitly requested to reconfigure by a network management entity, removing an attached station from the network.
- The concentrator PCM detects a high error rate on the connection.

PCM, a service of the SMT entity, is part of the Connection Management (CMT) portion of SMT. It initializes the connection of neighboring PHYs. PCM also provides support for connection maintenance and detection of marginal connection conditions.

3.2.3 Network Management

Independent of network management, the concentrator has the capability to remove and insert any connection. For example, the concentrator can disconnect itself, a SAS, or a DAS based on the results of the Link Error Monitor (LEM), which is a function of CMT. The LEM monitors the connection's bit error rate (BER). If a station's LEM determines that the BER is too high, the concentrator removes the connection from the ring.

From a network perspective, network management software allows the manager to configure a concentrator to suit the needs of the network. Under control of the network manager, the concentrator can remove or insert stations as necessary. The network manager can also access the management information maintained by the concentrator. The resultant information can be used for error checking.

3.3 Concentrator Operation

Frames pass through the concentrator by way of the PHY A, B, and M ports. If the concentrator is attached directly to the dual ring, frames enter by way of the PHY A port (if the primary ring is active), flow through the PHY M ports to attached nodes, and exit by way of the PHY B port. If the concentrator is part of a tree and not connected to the backbone, frames from the higher level concentrator enter through the PHY B port, flow through the PHY M ports to attached nodes, and back through the PHY B port to the higher level concentrator.

If the concentrator is dual homed and the backup link is active, frames from the higher level concentrator enter through the PHY A port, flow through the PHY M ports to attached nodes, and exit by way of the PHY A port.

3.4 Concentrator Planning Considerations

When planning for the installation of concentrators, such as the DECconcentrator 500, in an FDDI network, consult the following Digital manuals:

- *DECconnect System Fiber Optic Planning and Configuration*
- *DECconnect System Fiber Optic Installation*


3.5 Configuration Guidelines

This section provides an overview of rules and requirements for the configuring of a concentrator into an FDDI network. In most cases, Digital products autoconfigure. Physically installing a DECconcentrator is described in the installation manuals for the product.

The concentrator can be an integral part of the following network topologies:

- Dual ring of trees in a building or campus wiring system.
- Dual ring of concentrators in a backbone wiring systems (Figure 3–4).
- Standalone workgroup configurations such as LANs connecting high-performance workstations and servers. These workgroups can also be connected to a backbone system.

Figure 3–4: Dual Ring of Concentrators in a Backbone Wiring System



LKG-3812-89I

3.5.1 Concentrator Connection Rules

Use the following rules when connecting a concentrator into an FDDI network.

- Maximum distance between concentrators must not exceed 2 km (1.2 mi) for multimode connection, 1 km (0.62 mi) for low-power multimode connections, and 40 km (24 mi) for single-mode connections. This includes cables connected between the patch panel and the device.
- Total length of the network must not exceed 200 km (124.2 mi). This distance must include the length of the secondary ring, if a dual ring is implemented.
- Distance between concentrators and attached stations must not exceed 2 km (1.2 mi) for multimode connections, 1 km (0.62 mi) for low-power multimode connections, and 40 km (24 mi) for single-mode connections.
- Maximum number of stations in an FDDI LAN is 500 (1000 Physical layer entities). A station can have one Physical layer entity (SAS), two Physical layer entities (DAS), or multiple Physical layer entities (concentrators). Digital supports the configurations specified in the ANSI FDDI standard. However, Digital recommends that an FDDI LAN generally be limited to 100 stations and 20 concentrators due to performance and reliability characteristics of large rings (see also Section 1.8).
- Maximum number of concentrators in an FDDI network (dual ring of trees and dual ring) is limited by the maximum number of Physical layer entities that can connect to the ring (see also Section 1.8).
- Maximum optical attenuation allowable between a concentrator port and its neighbor is 11 dB for multimode connections, 7 dB for low-power multimode connections, and 22 dB for single-mode connections. Minimum allowable optical attenuation for single-mode connections is 12 dB.

Refer to Section 2.7.2 for FDDI device connection rules.

3.5.2 Concentrators and Structured Wiring

Commercial building wiring standards, such as EIA/TIA 568, address all communication needs including voice, data, and video for single or multiple building applications. EIA/TIA 568 defines a wiring system that supports a multivendor, multi-product environment. At the center of the wiring system is a hierarchy of connection equipment that supports the network.

The connection hierarchy specified in Digital's implementation of EIA/TIA 568 is the Main Distribution Frame (MDF), Intermediate Distribution Frame (IDF), Horizontal Distribution Frame (HDF), and the Wallbox. The concentrator can be located anywhere in the structured wiring scheme.


Digital's implementation of FDDI fully supports the EIA/TIA 568 Building Wiring Standard as it applies to fiber optic cabling. The FDDI wiring subsystem structures consist of the following:

- Interbuilding backbone (multibuilding system) that goes from a Main Distribution Frame (MDF) in one building to the Intermediate Distribution Frame (IDF) in one or more buildings.
- Intrabuilding backbone (single building system) that goes from the IDF to one or more HDFs.
- Horizontal wiring system that goes from an HDF to an office Wallbox.

In the Interbuilding Campus Backbone, concentrators are connected in a dual ring at the Main Distribution Frame in one building. From the MDF, fiber cable runs to concentrators at Intermediate Distribution Frames in the other buildings. Within each building, concentrators at the IDFs connect to concentrators at one or more HDFs. From the HDF, fiber cable can be run directly to the end station by means of the Wallbox.

Figure 3–5 is a conceptual view of this wiring scheme using the recommendations set forth in EIA/TIA 568 for structured wiring installations. It is possible to install a custom ring that exceeds the recommendations of this standard as long as FDDI distance and PMD-specific loss limits are not exceeded.

Figure 3–5: Concentrators Cascaded in Equipment Rooms




*Recommendations set forth in EIA/TIA 568 for structured wiring schemes.

LKG-3796-90I

Figure 3–6 shows the concentrator in dual ring of tree topology. Figure 3–6 is a conceptual view of a campus backbone wiring scheme. This figure also shows the PHY port connections for each concentrator.

Figure 3–6: Dual Ring of Trees in a Campus Environment




DACs1 and 2 are connected through PHY ports A and B to the dual ring.
 Concentrators attach to the ring through PHY ports M and B.
 DASs connect to the ring through PHY port B. DAS PHY port A is not
 used in this topology.
 Arrows indicate direction of token flow.

LKG-5097-901

Bridges

Bridges are the connecting devices between local area networks in an extended LAN architecture. Bridges allow the creation of extended LANs by acting as data link relays between similar or dissimilar LAN types (see Figure 4–1). In the FDDI environment, the bridge acts as a link between FDDI LANs or an FDDI LAN and other LAN types. This chapter explains bridge operation and the DECbridge 500/600 series FDDI-to-802.3/Ethernet SAS and DAS bridges. Also included in this chapter is a functional bridge description along with planning and configuration considerations for FDDI bridge. For a more detailed discussion of bridges, refer to the *Bridge and Extended LAN Reference* manual.

Figure 4–1: Extended LAN Examples




LKG-4075-90I

4.1 Description

Bridges are protocol-independent, specialized LAN stations that provide the connection point between two or more local area networks. They are store-and-forward devices operating at the Data Link layer. Bridges receive and store frames from attached LANs and selectively filter the frames or forward them to their destination address. This reduces total traffic on the extended LAN by filtering unnecessary traffic from the network.

In the FDDI environment, SAS bridges connect to the dual ring through a PHY M port connection of a concentrator. DAS bridges connect directly to the dual ring or to concentrators. One or more ports to the other network(s) complete the data path through the bridge. Figure 4–2 is a functional diagram of a typical FDDI bridge. A SAS bridge has a PHY S port. A DAS bridge has a PHY A and a PHY B port. Both SAS and DAS bridges have all the additional elements of a SAS or DAS including the PHY, PMD, MAC, and SMT.

Figure 4–2: FDDI Bridge Functional Diagram



LKG-4076-90I

4.2 Bridge Functions

Basic bridge functions include translation, frame forwarding and filtering, source address tracking, participation in the spanning tree algorithm, network management, and Internet Protocol (IP) fragmentation. FDDI bridge structure includes an FDDI SAS or DAS port(s), one or more ports to other network(s), address table and table lookup logic, packet memories, translation processor, and the forwarding/filtering control processor.

4.2.1 Translation Not Encapsulation

A major reason for choosing FDDI is that it is designed to be an industry standard for high-speed networks. To take advantage of this, users should avoid FDDI components that lock them into a specific vendor implementation. Otherwise, the user becomes tied to a closed, proprietary system with equipment choices limited to a single source.

Encapsulation is a technique used by some vendors to send data over the extended network. In encapsulation, packets received by the bridge from the non-FDDI network become the data in an FDDI frame. Using a proprietary protocol, the bridge surrounds, or encapsulates, the packet with an FDDI header and trailer and retransmits it onto the FDDI ring. Only a bridge, or other receiving device, supplied by the same vendor or one who subscribes to the same proprietary protocol can make use of this encapsulated data. Although this method is fairly simple to implement, it results in the user being forced into a closed, proprietary, single-vendor network architecture.

To avoid this pitfall, and in keeping with the open system philosophy, Digital does not use encapsulation. Digital employs the translation process in its bridge architecture. In the translation process, the bridge modifies the fields of a forwarded packet, making it compliant with the packet format of the network to which it is being forwarded. Although difficult to implement because it must be done in real time, translation allows the use of other vendor components in the network as long as they adhere to the various network standards on both sides of the bridge. See Section 4.3 for further details on translation.

4.2.2 Frame Forwarding and Filtering

The basic purpose of a bridge is to forward frames between its attached LANs, so that the stations on each LAN can communicate with all other stations as if they were on one large network. Bridges do not forward all frames that pass across a LAN. Bridges forward only those frames destined for a station on another port. If a frame is destined for a station known to be on the same port as the sending station, the bridge discards the frame.

Under network management control, bridges can be configured to filter certain frames from crossing the bridge. For example, individual addresses can be set so that the bridge will always filter, or always forward, frames destined to these addresses. The bridge does this regardless of whether or not it recognizes the addresses as belonging on the extended network.

Additionally, the bridge can be instructed to filter frames based on their source address and/or protocol type. These filters can be set on groups of ports, or they can be set globally.

4.2.3 Source Address Tracking

Digital and IEEE 802.1 require that bridges be self-configuring. For a bridge to determine whether to forward or discard a given frame, it must maintain a list of all known station addresses in the extended network, and the port on which the station is located. Each bridge maintains a source address table that it automatically updates. This table, called the Forwarding Database, is built from the source addresses in received frames.

When a bridge receives a frame, it records the source address of the frame into its Forwarding Database (address table), along with an indication of the port from which the frame was received. Because every frame contains the source address of the transmitting node, the bridge is constantly updating its forwarding database, learning the location of every station that sends a frame over the network. In an extended network, every bridge learns where stations are located, relative to its own position in the network, as frames pass through the bridge.

4.2.4 Spanning Tree

The logical topology of an extended LAN must be loop-free to prevent the circulation of packets indefinitely on the extended network. To eliminate this problem, the extended network forms a logical configuration, called a spanning tree, which eliminates logical loops in the extended network. This spanning tree is based on the spanning tree algorithm defined by Digital and implemented with minor changes by IEEE 802.1d. Digital bridges default to the IEEE 802.1d application and use the LAN Bridge 100 spanning tree implementation if they detect bridges that do not support IEEE 802.1d.

The spanning tree algorithm has two primary functions: the creation of one and only one path between any two bridges, and the proper connection of all LANs into a single extended LAN. The algorithm works by having all bridges in a given LAN exchange Hello messages (bridge PDUs). The bridges use the information in these messages to elect which bridges can forward frames and which cannot.

Bridges that form redundant paths in an extended LAN go into backup mode. In backup mode, bridges continue to monitor network traffic but do not forward frames. If the active bridge fails, the backup bridge notices the failure and begins to forward frames.

4.2.5 Network Management

Network management can monitor conditions on the extended network and adjust the logical topology to optimize traffic patterns and to fix things if they go wrong. To these ends, bridges maintain counters to keep track of the number of frames forwarded and discarded, and other such events.

The bridge forwarding database can be remotely observed and modified. This allows network management to determine the location of user stations, or to manage various regions of the extended network. By adjusting parameters in the spanning tree computational process, a network manager can optimize the algorithm performance for specific topologies.

The spanning tree algorithm database also contains useful information about the extended network topology. This information is made available to the network manager through the use of higher-level protocol messages.

Network management is also used to set special filters; for example, protocol filters. Network management entries for these filters are stored in a nonvolatile memory so that they can be restored easily on power up.

4.2.6 Internet Protocol Fragmentation

FDDI permits frames up to 4500 bytes in length. The 802.3/Ethernet LAN cannot handle frame sizes this large. To ensure that frame size restrictions are not violated when frames cross the bridge, the bridge performs a process called fragmentation on large FDDI IP frames. Fragmentation occurs when a frame is larger than the frame size allowed by the outbound port.

A frame is fragmented only if its LLC header indicates that it is an IP frame. The bridge receives the frame and decides how many packets must be generated. It then constructs the new packet(s) with a protocol type indicating IP frame and a new checksum in the IP header. The bridge also manipulates the header bits to indicate that the frame is a fragmented frame. The outbound port transmits the resulting packets with new data link CRCs for each packet. At the destination station, the frames are reassembled in the proper sequence.

4.3 Bridge Operation

The FDDI bridge provides a transparent link between networks. Users on the FDDI ring can freely exchange data with users on other networks. Data is transmitted over the FDDI ring in frames up to 4500 bytes in length. As FDDI data is received in the bridge, it is placed in packet memory buffers arranged in a queue structure. Each buffer is assigned a descriptor containing status information about buffer contents.

The queue manager (QM) determines the ultimate destination of frames received from the FDDI ring. The QM does this by looking at the descriptor, source and destination address, Frame Control field, type and quantity of previously received frames, and the bridge forwarding database. Based on this information, the QM forwards the frames or discards them. Packets destined for another network are placed into the FDDI forward queue for retransmission from the other network port.


At this point, the translation processor takes the packets and descriptors and physically moves them into the packet memory of the other network interface. As they are moved, the packets are converted to the format required by the other network. These packets are then passed to the network interface controller and eventually clocked onto the network.

Data received from the other network is handled in basically the same fashion. The filtering/forwarding processor determines where the received frames are to be sent. Packets are moved into the translation processor and converted to FDDI format. Data to be transmitted to the FDDI ring is placed in the FDDI transmit queue for re-transmission to the FDDI network.

4.3.1 FDDI Port

The DECbridge 500/600 series incorporates a PHY S port or a PHY A and PHY B port in its design, making the bridge a SAS or DAS. FDDI entities in the bridge include a PHY, PMD, MAC, and SMT. Figure 4-1 shows SAS and DAS bridge connections to FDDI. Figure 4-3 shows a SAS bridge connection to a concentrator.

Figure 4-3: SAS Bridge-to-Concentrator Connection



LKG-4247-90I

Data from the FDDI network passes through the FDDI interface where it is converted from optical signals, error checking is performed, and the resultant data is passed to FDDI packet memory.

4.4 Performance Considerations

A major reason for migrating to the FDDI network is its high bandwidth and resultant large data-transmitting capacity. To take advantage of this power, a bridge must be capable of supporting high throughput. It should not be a bottleneck on the ring. To be effective, a bridge must make frame forwarding decisions in real time. Key bridge performance characteristics are forwarding rate, filtering rate, and latency. Digital provides the performance characteristics in all three areas to support the FDDI bandwidth.

The forwarding rate is based on the number of data packets that can be transferred to and from the FDDI network. It should be large enough to allow a high volume of traffic between the FDDI ring and other networks. Forwarding rate is determined by the speed of the bridge in parsing a packet, verifying the destination address by checking its forwarding data base, and determining whether or not to forward the packet.

The filtering rate is the speed at which the bridge processes received packets to determine whether they are to be forwarded across the bridge. A high filtering rate keeps the bridge from clogging with data packets that are not being forwarded. Efficient filtering reduces the probability of overflowing internal storage areas and missing a packet to be forwarded. Efficient filtering reduces the probability of a lost packet.

Latency is the delay incurred when data crosses the bridge. Certain applications and protocols are delay sensitive. High latency results in nodes not being able to access devices on remote networks. High latency can cause the delay time between end nodes in an application to exceed the application's response time requirements, causing poor performance or connection failures. Digital designs its bridges to minimize latency time, thereby enhancing bridge performance.

4.5 Bridge Planning Considerations

When planning for the installation of an FDDI Bridge such as the DECbridge 500/600 series, consult the following Digital manuals:

- *DECconnect System Fiber Optic Planning and Configuration*
- *DECconnect System Fiber Optic Installation*

4.6 Bridge Configuration Guidelines

This section lists general configuration guidelines for the DECbridge 500/600 series. These bridges allow the use of a SAS or DAS connection to the FDDI network, and provide one or three ports that interface to 802.3/Ethernet networks.

The DECbridge 500/600 series connects to the FDDI ring through a concentrator (SAS bridges) or directly to the ring (DAS bridges). Although the FDDI standards allow a PHY S port to connect to any other FDDI port type, Digital recommends that a SAS bridge's PHY S port connect to a PHY M port of the concentrator. Refer to Section 2.7.2 for FDDI device connection rules.

4.6.1 General Connection Rules


Use the following guidelines when connecting one of the DECbridge 500/600 series into an FDDI network:

- SAS bridges should connect to a PHY M port on an FDDI concentrator to work in an FDDI ring. DAS bridges can connect directly to the dual ring or can be dual homed through a PHY M port on two separate concentrators.
- Maximum distance between the bridge and the concentrator or adjacent station must not exceed 2 km (1.2 mi) for multimode connections (refer to Figure 4-4) and 40 km (24 mi) for single-mode connections. This includes cables between the patch panel and the device.
- Maximum number of stations in an FDDI LAN is 500 (1000 Physical layer entities). A station can have one Physical layer entity (SAS), two Physical layer entities (DAS), or multiple Physical layer entities (concentrators).

Digital supports the configurations specified in the ANSI FDDI standard. However, Digital recommends that an FDDI LAN generally be limited to 100 stations and 20 concentrators due to performance and reliability characteristics of large rings (see also Section 1.8).

- Maximum number of bridges (SAS and DAS) in an FDDI network is limited by the maximum number of Physical layer entities that can connect to the ring (see also Section 1.8).
- Maximum number of bridges between two stations on the extended LAN is seven.
- Maximum allowable optical attenuation between the bridge and an adjacent FDDI station is 11 dB for multimode connections and 22 dB for single-mode connections. Minimum allowable optical attenuation for single-mode connections is 12 dB.

Figure 4-4: Bridge Distance Requirements for Multimode Connections



LKG-4077-90I

The following rules apply if a DECbridge 500/600 series unit is being installed using the AUI connector on the 802.3/Ethernet side of the bridge:

- If the bridge connects to an IEEE 802.3 transceiver such as the H4005, the transceiver cable must be an IEEE 802.3 compliant transceiver cable (BNE3H/K/L/M or BNE4C/D).
- IEEE 802.3 transceiver cables and Ethernet transceiver cables cannot be interconnected.
- If the bridge connects to an 802.3/Ethernet network, maximum length of the 802.3/Ethernet transceiver cable cannot exceed 50 m (164 ft). This maximum length is reduced by the internal cabling equivalency of a device (such as a DELNI) that is connected between the bridge and the transceiver, or by the use of an office transceiver cable. The internal cabling equivalency of such a device must be subtracted from the 50-m (164-ft) maximum.

NOTE

Cabling equivalency is a measure of the internal timing delay of a device expressed in meters of transceiver cable, for example:

If a device has a 5-m (16.4-ft) cabling equivalency, then its maximum allowable transceiver cable length is 50 m (164 ft) minus 5 m (16.4 ft) or 45 m (148 ft).

Office transceiver cable (BNE4x-xx), because of its smaller diameter, has a signal loss that is four times that of the (BNE3x-xx) transceiver cable. If office transceiver cable is used, the maximum transceiver cable distance must be divided by 4. The maximum office transceiver cable length allowed is 12.5 m (41 ft).

If a Digital FDDI-to-802.3/Ethernet bridge is being installed using the ThinWire connector, maximum length of the ThinWire cable segment is 185 m (606.95 ft).

For device-specific information related to cabling equivalency, refer to the *DECconnect System Planning and Configuration Guide*.

Adapters

FDDI adapters connect end-user nodes to the FDDI ring. Adapters allow attached devices to act as single attachment stations (SAS) or dual attachment stations (DAS). This chapter provides a functional overview of adapters. It also discusses planning and configuration considerations for these devices.

Several different adapters are available in the Digital product set. This chapter provides a general overview of one type of adapter. Although construction and implementation may vary from one adapter to another, basic functionality and purpose, the connection of a device to the FDDI network, remains the same.

5.1 Description

Digital's family of FDDI adapters are single attachment stations that connect to the FDDI ring by means of the concentrator. Part of the adapter is an interface to a specific workstation or system bus; for example, TURBOchannel, or other standard bus interface (see Figure 5-1). The SAS components of the adapter consist of a PMD, PHY, MAC, and SMT entity. Connection to the FDDI ring is through a PHY S port. In order to properly interface to the ring, the adapter should attach to a PHY M port on a concentrator. Section 5.5 lists basic configuration guidelines for the adapter.


5.2 Adapter Functionality

Adapters connect to the end node's internal bus structure. Digital's FDDI adapters use internal microprocessor-controlled logic to permit data transfer between the end node and the FDDI network without intervention of the host processor.

Digital FDDI adapters can be logically divided into four separate components as follows: (See Figure 5–1.)

- Host interface logic
- Adapter manager
- Packet memory
- FDDI SAS port

Figure 5–1: Adapter Model



5.2.1 Host Interface

The host interface connects the adapter to the host's internal bus. It contains all the necessary logic to pass frames to and from host memory. As part of its structure, buffers in the host interface provide the capability for high-speed data movement to and from the host.

5.2.2 Adapter Manager

The adapter manager oversees the transfer of data between the node and the FDDI ring. Functions performed by the adapter manager include the following:

- Adapter initialization
- Error handling
- Device driver command implementation
- Local management of the FDDI MAC, PHY, and PMD components within the adapter
- Collection and maintenance of node statistics

5.2.3 Packet Memory

Packet memory within the adapter operates in much the same manner as the memory in a bridge. Data from the FDDI ring is placed in queues within this memory space. Virtual addressing is used to access the data as it progresses through the queues. As in the bridge, the data itself never actually moves until its final destination is determined. A descriptor word containing control information for each received packet is actually passed around the entities. This word is changed as required when data ownership changes. Depending on the type of packet or frame received, the data is passed onto the node, forwarded to the adapter manager, or discarded.

Similarly, data from the node moves through packet memory. Once it is determined that the data is meant for transmission onto the ring, it is passed to the FDDI port.

5.2.4 FDDI Port

The adapter incorporates a PHY S port in its design, making the adapter a SAS node. FDDI components in the adapter include a PHY, PMD, MAC, and SMT entity. Data from the node that is destined for the FDDI ring passes to the FDDI port. Here the data is formatted into FDDI frames and transmitted over the ring.

5.3 Operation

Actual adapter operation is a complex process. The following sections provide a summary of the initialization and ring-access processes that an adapter goes through when it is first powered on.

5.3.1 Initialization

The adapter must be initialized when it is powered on or recovering from a fatal error condition. In the initialized state, the adapter has allocated all buffers for internal memory, loaded all registers with appropriate start addresses, and is ready to process commands from the user device. During this process the following occurs:

- Self-test diagnostics are run to ensure adapter integrity. If the tests pass, the remaining events occur. If not, the adapter is not initialized.
- Internal logic devices are initialized.
- The 48-bit adapter LAN address is loaded into data registers.

5.3.2 Accessing the Ring

Once the adapter is initialized, it must access the FDDI ring. This is accomplished as follows:

- After successful initialization, the adapter transitions to the run state and starts FDDI logic testing.
- After successful completion, the FDDI logic requests connection to the ring by generating a series of claim frames.
- After synchronizing with the ring, physical and logical connections are established and the adapter is able to communicate with other attached devices.

5.3.3 Network Management Functions

To permit network management, Digital adapters incorporate a number of functions in their basic design. Proper use of these functions provide the network or system manager with the ability to monitor and control adapters individually, or as a distinct group, from a remote location. These adapter functions include the following:

- Request ID
- System ID
- Request counters
- Remote boot

5.4 Adapter Planning Considerations

When planning for the installation of an FDDI adapter, consult the following Digital manuals:

- *DECconnect System Fiber Optic Planning and Configuration*
- *DECconnect System Fiber Optic Installation*

5.5 Adapter Configuration Guidelines

This section lists general configuration guidelines for Digital's adapters. Digital adapters are single attachment stations that connect to the FDDI ring by means of the concentrator.

Although the FDDI standards allow a PHY S port to connect to any other FDDI port type, Digital recommends that a PHY S port connect to a PHY M port of a concentrator. Refer to Section 2.7.2 for FDDI device connection rules.


Use the following guidelines when connecting an adapter into an FDDI network.

- SAS adapters should connect to a concentrator PHY M port to work in an FDDI ring. In a standalone configuration, the adapter PHY S port can connect to a concentrator PHY A or PHY B port.
- Maximum allowable distance between the concentrator and an adapter is 2 km (1.2 mi) for multimode fiber connections, and 100 m (330 ft) for Thin-Wire and shielded twisted-pair. This includes cables between the patch panel and the device. Figure 5-2 shows the distance requirements for an FDDI adapter with multimode fiber. It also shows the connection to a concentrator.
- Maximum number of stations in an FDDI LAN is 500 (1000 Physical layer entities). A station can have one Physical layer entity (SAS), two Physical layer entities (DAS), or multiple Physical layer entities (concentrators).

Digital supports the configurations specified in the ANSI FDDI standard. However, Digital recommends that an FDDI LAN generally be limited to 100 stations and 20 concentrators due to performance and reliability characteristics of large rings (see also Section 1.8).

- Maximum number of adapters in an FDDI network is limited by the maximum number of physical layer entities that can connect to the ring (see also Section 1.8).
- Maximum number of patch-panel connections between the concentrator and the adapter is two connections.
- Maximum allowable loss between an adapter and the concentrator is 11 dB for multimode connections.

Figure 5-2: Adapter Distance Requirements for Multimode Fiber



LKG-4079-89I

Network Management

Proper management is important to the successful maintenance of a network. Without sophisticated and flexible management software, even the most powerful network can become an unwieldy, hard-to-maintain system.

Network management, as envisioned by Digital, is a means for network and system-level personnel to effectively design, implement, and maintain networks to meet specific organizational goals and technical needs. Digital provides a framework for the management of heterogeneous, multivendor systems linked by a number of network environments. This ability to manage devices over an extended network provides users with tremendous flexibility in terms of network design and cost.

In the Digital scheme, network management is one component of Enterprise Management. This term applies to Digital's ability to manage not just a given network, but a broad spectrum of systems, including voice and data, corporate-wide and personal computing, and distributed computing environments.

Digital believes that the concept of Enterprise Management Architecture is where the future of system management lies. In designing EMA, Digital has followed the model set forth in the OSI system management architecture. This chapter describes part of the Enterprise Management Architecture through a discussion of management over the extended network.

6.1 Network Management Benefits

Digital's network management tools provide the following benefits:

- Flexibility
- Extensibility
- Integration
- Growth Potential

A further benefit is protection of the network investment. Because existing components, as well as any additions, are supported, there is no need to discard existing equipment when expanding network facilities.

6.2 Network Management Functions

Network management functions are those parts of the software system that network and system-level personnel use to monitor, control, and maintain the network. As set forth in the ISO/OSI management framework, these functions are as follows:

- Network configuration management to set or change operating parameters of the network
- Fault management to detect, diagnose, and correct network faults and error conditions
- Performance management to monitor and evaluate network performance
- Accounting management to monitor and evaluate network resources
- Security management to provide a secure computing environment

If management functions such as fault notification and fault logging are implemented independently of the network components being managed, the result is an increase in the scope of the management system. Management functions must be capable of being customized according to user requirements. They must be capable of interaction, so that the output of one function can be used as input to another. This provides for the implementation of more sophisticated management applications to support the network systems throughout their existence, from planning and implementation to operation and change.

The role of the ISO/OSI categories is to serve as the basis for developing usable management applications. For example, to be truly effective, network operation requires the integration of performance, fault, and configuration management information. The following sections describe the various ISO/OSI components in further detail.

6.2.1 Network Configuration Management

Network configuration management provides a means for configuring the network and its components through setting or changing the operating parameters of the local or extended network. This includes turning the parameters on or off, altering their operational status, and collecting and distributing current status information of attached devices.

Ring mapping, a part of FDDI configuration management, enables users to identify all stations on the FDDI network. This feature allows mapping of the FDDI network configuration, providing network managers with valuable ring status information.

6.2.2 Fault Management

Fault management provides a means for detecting, diagnosing, and correcting network faults and error conditions. Fault management capabilities include the receipt of unsolicited error messages from network devices, periodic polling for error messages, and setting error thresholds for individual devices. Through proper setting and monitoring of operational parameters (for example, error alarms and threshold limits), network managers can actually prevent certain fault conditions.

6.2.3 Performance Management

Performance management allows the network manager to monitor and evaluate the performance of local and extended networks by collecting and processing relevant data. The information is gathered from device counters and error reports, then processed and analyzed for network planning and tuning.

6.2.4 Accounting Management

Accounting management allows managers to effectively monitor and evaluate network resources and their use. This function identifies usage costs associated with the network, and provides the necessary tools for identifying and charging users for network time. Accounting management is an application function.

6.2.5 Security Management

As networks become larger and more complex, and as more companies place highly sensitive data on the network, the ability to create a secure environment increases in importance. Security management defines those facilities required to manage network services: authentication, control of access to resources, and the protection of confidential information on both the local and extended network.

6.3 Digital's Network Management Capabilities

Network management is done through the use of a layered software product that runs on a designated host station. This software communicates with corresponding software on network devices, allowing the network manager to remotely configure, manage, and monitor devices on the extended LAN. These devices include the following:

- FDDI and non-FDDI bridges
- FDDI concentrators

The following sections describe some of the capabilities of Digital's network management software.

6.3.1 Device Naming

With the network management software, all manageable bridges and concentrators can be assigned unique ASCII names, as opposed to more cumbersome hardware addresses. These names can then be used in commands to the devices.

6.3.2 Device Registry

The management software creates a network registry containing information about each manageable bridge and concentrator in the extended LAN. This registry contains device type, address, name, and a brief description for each attached device. Registry information can be added or modified as required. Depending on the management software, this information can also be automatically collected as new bridges or concentrators are added to the network.

6.3.3 Event Messages and Event Log

The network management station can display error and event messages. The network manager collects these messages by polling and by error or event logging. Event codes can also be recorded in an event log. The network management software provides a simple mechanism for report generating based on the contents of the event log.

6.3.4 Security

In any open network, security is always a primary concern. Digital's network management software provides the ability to assign passwords to all bridges and concentrators in the extended LAN. Set from the network management station, these passwords prevent the unauthorized control and configuration of network devices. Devices store the passwords in their nonvolatile memory.

6.3.5 Configuration Parameters

Network management software allows the user to remotely set bridge and concentrator configuration parameters from the network management station. Configurable parameters include the following:

- Spanning tree parameters—Controls the extended LAN logical layout and performance by manipulating bridge and line spanning tree parameters.
- Firmware upgrade down-line loading—Enables a bridge or concentrator to accept down-line loading of firmware upgrades, eliminating the need to perform upgrades at the bridge or concentrator site.
- IP fragmentation—Enables an FDDI bridge to break large Internet Protocol (IP) frames received from the FDDI LAN into smaller frames for retransmission across the bridge. For example, fragmentation occurs in Digital FDDI-to-Ethernet bridges because maximum frame size in the FDDI LAN is 4500 bytes, but 802.3/Ethernet can only support a 1500-byte frame size. The DECbridge 500/600 series units come with the default mode set to fragment IP traffic.
- IP fragmentation error checking—Enables an FDDI bridge to perform header checksum verification on large IP frames received from the FDDI LAN to ensure data integrity of frames being sent across the bridge. The DECbridge 500/600 series units automatically perform IP fragmentation error checking.

6.3.6 Initialization

The user can initialize bridges and concentrators from a network management station. Doing this forces the device to reset its internal counters to zero and run its self-test diagnostics. At initialization, bridges perform the spanning tree computation process that determines active and standby bridges. Concentrators perform the initialization process described in Section 2.9.3.

6.3.7 Bridge Port Enabling and Configuring

From the network management station, bridge ports in the extended LAN can be enabled and disabled. As bridges are enabled and disabled, the spanning tree computation process determines which are the active and which are the standby bridges in the network.

In the FDDI environment, network management software can set the maximum token rotation time (T_MAX), the requested token rotation time (T_REQ), and Valid Transmission Time (TVX). Setting T_REQ and T_MAX for all attached devices on the FDDI LAN allows the user to control the value of Target Token Rotation Time (TTRT) on the FDDI LAN.

6.3.8 PHY Port Configuration and Control

The network manager can enable and disable the PHY ports of concentrators. This allows the configuring of networks to fit the needs of the network users. Managers also have the ability to set Link Error Monitoring (LEM) threshold limits for FDDI devices. LEM performs continuous testing of links during ring operation to ensure the bit error rate is being met.

6.3.9 Address Filtering

With Digital's network management software, network managers can add or remove addresses at the forwarding database of a bridge (FDDI and non-FDDI). This capability is used to block certain stations from sending or receiving frames across a bridge, localizing traffic on a specific subnetwork or to stations within a subnetwork.

6.3.10 Protocol Filtering

In addition to address filtering, network management software allows the filtering of protocols at a bridge (FDDI or non-FDDI). This capability instructs the bridge to forward or filter (discard) frames based on the protocol information within the frame. This feature is useful to regionalize traffic or to isolate broadcast storms or other protocol-related problems. Selective protocol filtering allows only selected protocols to cross the bridge. Again, selected protocols are specified using the network management station and supported software.

6.3.11 Network Monitoring

Device status, characteristics, and counter contents are all accessible through the network management software. Bridges, concentrators, and the connection to each port on these devices can be addressed. In the case of bridges, the spanning tree parameter settings and the outcome of the spanning tree computational process are accessible using the software.

6.3.12 FDDI Physical Ring Mapping

The network management software supports creation of a physical ring map of attached devices in the FDDI LAN. This map allows a user to determine the order of token rotation among devices on the ring.

6.3.13 SMT Neighborhood Information Frame (NIF) Display

Besides displaying the NIF contents, network management software can use the NIF frames to show the order in which each station appears in the token path. Other uses of the NIF contents include downstream neighbor information gathering and duplicate address detection.

6.3.14 SMT Station Information Frame (SIF) Display

The network management software can display SIF frames from individual FDDI stations. These frames contain a station's configuration and operating information. Network management software can create a physical ring map using the SIF configuration information. The operating information can be used as an aid in ring fault detection.

6.3.15 Troubleshooting

Digital's network management software provides the network user with a wealth of useful information. Proper use and interpretation of this information results in a powerful troubleshooting tool. By properly analyzing the network performance and error information, a user can diagnose and correct a variety of network problems with minimal effect on overall network performance.

Combining the FDDI dual ring of trees with the power of Digital's network management software results in a dynamic network resistant to failures caused by everyday events. Proper use of the network management tools allows the user to isolate and remove problem network devices with little or no effect on network performance or throughput.

OSI Reference Model


Chapter 1 describes FDDI and its relationship to the OSI reference model. This appendix provides an overview for readers unfamiliar with the OSI reference model.

The International Organization for Standardization (ISO) has established what has become known as the Open Systems Interconnection (OSI) Reference Model. The OSI model implements the design structure known as layering. In this structure, the communications functions are divided into a vertical set of layers, with each layer performing a related subset of the functions required to communicate with another system. Each layer relies on the services of the next lower layer and the ability of the protocols of the lower layer to conceal those functions from it and layers further up the ladder. Layering provides modular design and allows independence of design teams. As shown in Figure A–1, the OSI model consists of seven layers, with an interface between layers .

The layered approach to networking adheres to the following general guidelines:

- Each level should perform a well-defined function.
- The layers should comply with recognized standards.
- Layer interfaces should be designed to minimize data flow between layers.
- Layers should be sized to minimize complexity without making the layer unmanageable.

Figure A-1: OSI Seven-Layer Model



LKG-4080-90

A.1 Physical Layer

At the Physical layer, the physical interface between devices and the rules for bit transmission are established. The Physical layer defines the mechanical, electrical, functional, and procedural elements of the network. Examples of Physical layer standards are RS-232-C, RS-449, RS-422, and RS-423.

A.1.1 Data Link Layer

The Data Link layer provides the raw bit-stream service. It is responsible for providing reliable data transmission from one node to another. It provides the means to activate, maintain, and deactivate the link. This layer also provides error detection and control functions to the higher layers of the OSI model.

If communication is established between two devices not on the same fiber, however, the result is a number of functionally independent data links. In this instance, the higher layer must assume responsibility for overall error detection and correction. Examples include HDLC and LLC.

A.1.2 Network Layer

The Network layer provides for the transparent transfer of data between transport entities. It relieves the Transport layer of the need to be aware of the underlying data transmission and switching technologies used to connect systems. Examples include Internet Protocol and the OSI Router (ISO 9542).

A.1.3 Transport Layer

The Transport layer provides a reliable means for the exchange of data between processes in different systems. This layer ensures that data units (frames) are delivered error free, in sequence, with no losses or duplications. The Transport layer can provide optimization and requested service capabilities. For example, the session layer could specify error rates, delay and priority, and security parameters for a given session. It is the Transport layer that would provide them. TCP is an example of this layer.

A.1.4 Session Layer

The Session layer provides the means for controlling the dialogue between presentation entities. The Session layer permits two presentation entities to establish and use a connection, called a session. In addition, it provides services such as dialog type and recovery. The recovery function provides a mechanism for retrieving data transmitted during a session, in the event of a session failure.

A.1.5 Presentation Layer

The Presentation layer offers applications and terminal handler programs a set of data transformation services. This layer provides such services as data translation, formatting, and syntax selection. Examples of protocols in this layer are text compression and encryption.

A.1.6 Application Layer

The Application layer provides a means for application processes to access the OSI environment. It contains management functions and other mechanisms to support distributed applications. Examples of protocols at this level are virtual file protocol and job transfer and manipulation protocol.

Additional Configurations

The EIA/TIA 568 commercial building wiring standard provides an infrastructure for FDDI networks. This appendix contains examples of FDDI network configurations built on Digital's implementation of the commercial building wiring scheme.


B.1 Standalone Configurations Using Concentrators

Figure B-1 shows a concentrator installed in a standalone workgroup. The workstations can be either SAS or DAS devices. SAS devices are preferred for this configuration. This configuration is useful for compute-intensive applications shared by a limited number of users in a limited geographical area, or where security is paramount.

Figure B-2 shows a dedicated network on a dual ring. The attached workstations, file server, mini-computer, and main frame are SAS devices. This configuration can serve the same purposes as the standalone configuration with the added benefit of extending the geographical area.

These configurations allow the use of star wiring, which means existing structured cabling can be used.


Figure B-1: Standalone Workgroup Installation



SAS = single attachment station

LKG-3797-89I

Figure B-2: Dedicated Network in a Dual Ring Installation




LKG-4456-89

B.2 Tree Configurations Using Concentrators and Bridges

Tree configurations are used when wiring large groups of user devices together. Concentrators and bridges are wired in a star topology with one or more concentrators serving as the root of the tree. Figure B-3 shows a tree of concentrators installed in a campus configuration wired with Digital's implementation of EIA/TIA 568. Figure B-4 shows a tree of concentrators installed in a building with multiple HDFs. These configurations also show the use of single and multiport bridges. Refer to Section 3.5.2 and Figure 3-5 for the definition of distribution frames in structured wiring.


Figure B-3: FDDI Tree of Concentrators and Bridges, Campus Configuration



IDF = Intermediate distribution frame
HDF = Horizontal distribution frame
MDF = Main distribution frame
B = Bridge

LKG-3809-901

Figure B-4: Tree of Concentrators Installed in a Building




HDF = Horizontal distribution frame
IDF = Intermediate distribution frame

LKG-3813-89I

B.3 Dual Ring Configurations


A dual ring of concentrators allows multiple devices to be attached to the dual ring through the concentrators. Figure B-5 shows a dual ring of concentrators installed in a building. The concentrators can be located in the MDF of a building, IDF of a building, or distributed throughout the building in HDFs.

Figure B-5: Dual Ring of Concentrators Installed in a Building



LKG-3812-8

Figure B-6: Dual Ring of Concentrators and Bridges Installed in a Campus



B.4 FDDI Dual Homing

Figure B–7 shows three concentrators installed in a dual homing configuration. In this configuration, if concentrator 2 or the primary link to concentrator 3 fails, the backup link to concentrator 3, through concentrator 1, is activated. This ensures that the devices connected to concentrator 3 will have uninterrupted service. Activation of the individual links follows the FDDI connection rules, described in Section 2.7.2.

Figure B–7: Concentrators Installed for Dual Homing





Figure B-8 shows a DAS bridge installed in a dual homing configuration. In this configuration, if concentrator 2 or the primary link to the bridge fails, the backup link to the bridge, through concentrator 1, is activated. This ensures that the devices on the LANs connected to the bridge will have uninterrupted service. Activation of the individual links follows the FDDI connection rules, described in Section 2.7.2.

Figure B-8: DAS Bridge Installed for Dual Homing



LKG-4935-91I

B.5 Single-Mode Fiber Configurations

Figure B-9 shows a SAS bridge connected to a dual ring through a long distance single-mode fiber cable. This is the recommended method to connect remote sites through FDDI.

Figure B-9: Connecting to the Dual Ring Through a Single-Mode Link





Figure B–10 shows a dual ring which requires some links to be longer than 2 km (1.2 mi). In this configuration, the maximum total length of all links must not exceed 100 km (60 mi). DAS connections count as twice the link length.

Figure B–10: Dual Ring with Mixed Link Lengths



LKG-5000-911

Related Documents

The following documents provide additional information concerning FDDI hardware and software products offered by Digital.

- *DECconcentrator 500 Installation* (Order No. EK-DEFCN-IN)
- *DECconcentrator 500 Problem Solving* (Order No. EK-DEFCN-PS)
- *DECbridge 500/600 Series Installation and Upgrade*
(Order No. EK-DEFEB-IN)
- *DECbridge 500/600 Series Problem Solving* (Order No. EK-DEFEB-PS)
- *DECelms Installation* (Order No. AA-PAK1A-TE)
- *DECelms Use* (Order No. AA-PAK2A-TE)
- *DECelms Reference* (Order No. AA-PBWBA-TE)

For information concerning planning and configuration of fiber networks, consult the following documentation.

- *DECconnect System Fiber Optic Planning and Configuration*
(Order No. EK-DECSY-FP)
- *DECconnect System Fiber Optic Installation* (Order No. EK-DECSY-FI)

The following documents provide information concerning other Digital network offerings.

- *Networks and Communications Product Documentation* (Order No. EK-NACPD-RE)
- *Telecommunications and Networks Buyer's Guide*
- *Bridge and Extended LAN Reference* (Order No. EK-DEBAM-HR)

For additional information concerning Digital products that can connect to an FDDI network and for information on miscellaneous networking topics, refer to the following documents.

- VAX and VAX Workstation document set for hardware platform information
- VAX and DECwindows document set for software platform information
- *Guide to DECnet-VAX Networking* for DECnet node information
- *Introduction to Network Performance* for network performance information
- *Network Troubleshooting* for network problem-solving information

Glossary

attenuation

The difference between transmitted and received power. Expressed in units of decibels (dB).

bridge

A connecting device between local area networks in an extended LAN architecture. A bridge acts as a link between similar and dissimilar LANs permitting the creation of extended LANs.

bypass

The ability of a station to be electronically or optically isolated from the network while maintaining the integrity of the ring.

code bit

The smallest signaling element used by the Physical layer for transmission on the fiber cable.

code group

The specific sequence of five code bits representing a PHY symbol.

concentrator

An FDDI node that provides additional attachment points for stations that are not attached directly to the dual ring, or for other concentrators in a tree structure. A concentrator has two or more Physical Layer entities and may have zero or more Data Link layer entities. The concentrator is the focal point of Digital's dual ring of trees topology.

Connection Management (CMT)

That portion of the Station Management function that controls network insertion, removal, and connection of the PHY and MAC entities within a station.

Dual Attachment Concentrator (DAC)

A concentrator that offers two connections to the FDDI network capable of accommodating the FDDI dual (counter-rotating) ring, and additional ports for connection of other CONs or stations.

Dual Attachment Station (DAS)

An FDDI station that offers two connections to the FDDI network capable of accommodating the FDDI dual (counter-rotating) ring.

entity

A manageable active element within an Open Systems Interconnection (OSI) layer or sublayer.

fiber

A dielectric waveguide that guides light.

Fiber Distributed Data Interface (FDDI)

A set of ANSI/ISO standards that, when taken together, define a 100 Mb/s, timed-to-token protocol, local area network that uses fiber optic cable as the transmission medium. The standards define Physical Layer Medium Dependent, Physical Layer, Media Access Control, and Station Management entities.

fiber optics

A technology whereby signals are transmitted over an optical waveguide medium through the use of light-emitting transmitters and light-detecting receivers.

fiber optic cable

A jacketed fiber or fibers.

frame

A MAC Protocol Data Unit (PDU) transmitted between MAC entities on a ring. A frame consists of a variable number of bytes up to a maximum of 4500, including 2 bytes of preamble.

Horizontal Distribution Frame (HDF)

Located on the floor of a building. Consists of the active, passive, and support components that provide the connection between the building backbone cabling and the horizontal wiring.

Intermediate Distribution Frame (IDF)

Located in an equipment room. Consists of the active, passive, and support components that provide the connection between interbuilding cabling and the intrabuilding cabling for a building.

Main Distribution Frame (MDF)

Located in an equipment room. Consists of the active, passive, and support components that provide the connection of the interbuilding backbone cables between IDFs.

Media Access Control (MAC)

The Data Link layer sublayer responsible for scheduling and routing data transmissions on a shared medium local area network (for example, FDDI).

Media Interface Connector (MIC)

An optical fiber connector pair which links the fiber media to the FDDI node or another cable. The MIC consists of two halves, the plug and the receptacle. The MIC plug terminates an optical fiber cable. The MIC receptacle is associated with the FDDI node.

MIC plug

That portion of the Media Interface Connector which terminates an optical fiber cable.

MIC receptacle

The fixed portion of the Media Interface Connector which is part of an FDDI node.

node

A generic term applying to any FDDI network attachment device (single attachment station, dual attachment station, or concentrator).

Non-Return-to-Zero (NRZ)

A data representation for binary data in which a polarity level high, or low, represents a logical 1 or 0 respectively.

Non-return-to-Zero-Invert on Ones (NRZI)

A data transmission technique in which a polarity transition from low to high, or high to low, represents a logical 1. The absence of a polarity transition represents a 0.

Numerical Aperture (NA)

The sine of the radiation or acceptance half angle of an optical fiber multiplied by the refractive index of the material in contact with the exit or entrance face of a fiber.

optical receiver

An optoelectronic circuit that converts an incoming optical signal to an electrical signal.

optical reference plane

The plane that defines the optical boundary between the MIC plug and the MIC receptacle.

optical transmitter

An optoelectronic circuit that converts an electrical signal to an optical signal.

physical connection

The full-duplex physical layer association between adjacent PHYs in an FDDI ring.

Physical Layer Medium Dependent (PMD)

The Physical layer sublayer that defines the physical requirements for nodes that attach to the FDDI ring. Items defined by the PMD include transmit and receive power levels, connector requirements, and fiber optic cable requirements.

Physical Layer Protocol (PHY)

The Physical layer sublayer that defines the media-independent portion of the Physical layer in FDDI. Items defined by the PHY include symbols, line states, data encode/decode process, and the data-recovery process.

physical link

The path, via PMD and attached cabling, from the transmit logic of one PHY to the receive logic of an adjacent PHY in an FDDI ring. Part of the physical connection.

primitive

A basic element of the services provided by one entity to another.

Protocol Data Unit (PDU)

Information delivered as a unit between peer entities that may contain control information, address information, and data.

receive

The act of a station acquiring a frame, token, or control sequence from the ring.

repeat

The act of a station in receiving a code-bit stream (frame or token) from an upstream station and placing it onto the physical link connected to its downstream neighbor. The repeating station may examine, copy to a buffer, or modify frame status bits in the code-bit stream as appropriate.

repeater

An FDDI node that minimally comprises the functionality of two PMDs and provides only a repeat function. A repeater does not have any MACs or concentrator functionality.

ring

Two or more stations wherein information is passed sequentially between active stations, each station in turn examining or copying the information, finally returning it to the originating station.

services

The services provided by one entity to a higher entity or to Station Management (SMT).

Single Attachment Concentrator (SAC)

A concentrator that offers one attachment to the FDDI network, and extra ports for the attachment of stations or other CONs.

Single Attachment Station (SAS)

An FDDI station that uses one PHY-port of type S to connect to the FDDI network.

station

An addressable node on an FDDI ring capable of transmitting, receiving, and repeating data. A station has exactly one instance of SMT and at least one instance of MAC, PHY, and PMD.

Station Management (SMT)

The entity within a station on the FDDI ring that monitors station activity and exercises overall appropriate control of station activity.

symbol

The smallest signaling element used by the MAC sublayer. The symbol set consists of sixteen data symbols and eight control symbols. Each symbol corresponds to a specific sequence of code bits (code group) to be transmitted by the Physical layer.

transmit

The act of a station that consists of generating a frame, token, or control sequence, and placing it on the ring for receipt by the next (downstream) station.

wallbox

Connects workstation wiring to horizontal wiring. Also called faceplate, information outlet, or telecommunications outlet.

window

As applied to optics, an optical wavelength region of relatively high transmittance, surrounded by regions of low transmittance. For example, the first window is at 850 nm, the second at 1300 nm. Also called Spectral Window or Transmission Window.

Index

Numbers

4B/5B encoding, 2–12

A

Active Line State, 2–12

Adapters, 1–29, 5–1

configuration guidelines, 5–5

description, 5–1

distance requirements, 5–6

FDDI port, 5–3

functionality, 5–1

host interface, 5–2

model, 5–2

operation, 5–3

packet memory, 5–3

planning, 5–5

Address format

48-bit, 2–16

universal, 2–17

Alternate fiber specifications, 2–4

American National Standards Institute

role in FDDI standards, 1–1

X3, 1–4

X3.139 MAC standard, 1–8

X3.148 PHY standard, 1–7

X3.166 PMD standard, 1–7

X3.T9.5 Station management, 1–8

X3T9, 1–4

X3T9.5, 1–4

Application, environments

backbone, 1–12

workgroup, 1–12

Application layer, layer, A–4

Asynchronous transmission, 2–30

B

Backbone, 1–13

Beacon process

description, 2–37

directed beacon, 2–28

function, 2–37

stuck beacon, 2–28

Bridges, 4–1

configuration

example, B–3, B–7, B–9

guidelines, 4–10

definition, 4–1

description, 4–2

distance requirements, 4–11

dual homing, 1–23, B–9

encapsulation, 4–4

- FDDI-to-Ethernet
 - DAS devices, 4-1
 - SAS devices, 4-1
- filtering, 4-5
- forwarding
 - database, 4-5
 - frames, 4-5
- fragmentation, 4-7
- function, 1-29, 4-4
- multiport, 1-29, B-3, B-5, B-9
- network management, 4-6
- operation, 4-7
- performance, 4-9
- planning considerations, 4-10
- queue manager, 4-7
- source address tracking, 4-5
- spanning tree, 4-6
- stripping algorithm, 2-40
- translation
 - process, 4-4
 - processor, 4-8

C

- Canonical addressing
 - description, 2-18
 - fields, 2-18
 - transmission rules, 2-19
- Claim process, 2-32
- Clock recovery
 - description, 2-13
 - frequency tolerance, 2-14
- Code
 - bits, 2-8
 - group, 2-8
- Concentrator
 - architecture, 3-3, 3-4

- configuration guidelines, 3-7
- connection rules, 3-9
- definition, 3-1
- dual
 - homing, 1-22, B-8
 - ring of trees, 3-3
- dual attachment concentrator, 1-9, 3-2
- FDDI device, 1-2
- function, 1-9, 1-11, 3-1, 3-5
- in-band management, 3-5
- network management, 3-6
- operation, 3-7
- out-of-band management, 3-5
- Physical Connection Management, 3-6
- planning, 3-7
- single attachment concentrator, 1-9, 3-2
- standalone topology, 3-2
- station bypass, 3-6
- structured wiring, 3-10
- topologies, 1-29
- types
 - dual attachment concentrator, 3-2
 - single attachment concentrator, 3-2
 - with a MAC, 3-5
 - without a MAC, 3-5
- Configuration
 - campus backbone, B-4
 - dedicated network, B-3
 - dual homing, 1-22, 1-23, 2-27, B-8, B-9
 - logical star, B-4
 - multiport bridge, B-7
 - standalone network, B-1
 - workgroup LAN, B-5
- Configuration Management, 2-24
- Connection Management

- Configuration Management, 2–24
- Entity Coordination Management, 2–24
- function, 2–23
- Physical Connection Management, 2–23
- Connection rules
 - bridge, 4–10
 - concentrator, 3–9
 - PHY port types, 2–25
- Connector keying, 2–7
- Control indicators, 2–33
- Counter-rotating ring, 1–18

D

- DAC. *See* Concentrator
- DAS. *See* Dual attachment station
- Data Link layer, 1–8, A–3
- Data recovery, 2–8
- DECbridge 500/600 series, 1–29, 4–1, 4–10
- DECconcentrator 500, 1–29, 3–1, 3–7
- Destination address, 2–22
- Device, definition, 1–2
- Directed beacon, 2–28
- Documents
 - Digital FDDI products, C–1
 - networks, C–1, C–2
- Dual attachment station, 1–2, 1–9
 - bridges
 - configuration rules, 4–10
 - dual homing, 1–23, B–9
 - concentrator
 - components, 3–2
 - configuration rules, 3–9
 - dual homing, 1–22, B–8
 - FDDI connection, 1–9

- connection rules, 1–23
- model, 1–11
- optical bypass relay, 1–10
- Dual homing
 - configuration
 - guidelines, 2–27
 - with concentrators, 1–22, B–8
 - with DAS bridge, 1–23, B–9
 - description, 1–22
- Dual ring of trees, 1–1, 1–2, 1–20
- topology of choice, 1–23
- Dual-window fiber, 2–4

E

- Elasticity buffer
 - buffer size, 2–14
 - description, 2–13
 - function, 2–8
 - use of idle symbols, 2–14
- Encapsulation, 4–4
- Encode/decode process, 2–8
- Ending delimiter, 2–22
- Enterprise Management, 6–1
- Entity Coordination Management, 2–24

F

- FCIS. *See* Frame Content Independent Stripping
- FDDI. *See* FDDI port types; Fiber Distributed Data Interface
- FDDI port types
 - PHY A, 2–25
 - PHY B, 2–25
 - PHY M, 2–25
 - PHY S, 2–25
- Fiber Distributed Data Interface

- definition, 1–2
- planning and configuration, 1–31
- recommended fiber counts, 1–32
- specifications, 1–1, 1–30, 2–4, 2–5
- Fiber specifications, 2–4, 2–5
- Fragmentation, 4–7, 6–5
- Frame
 - check sequence, 2–22
 - control field
 - address length bit, 2–22
 - class bit, 2–22
 - control bits, 2–22
 - format bits, 2–22
 - description, 2–21
 - ending delimiter, 2–22
 - fields, 2–21
 - format, 2–21
 - fragments
 - creation of, 2–35
 - stripping, 2–40
 - preamble, 2–21
 - status field, 2–23
 - control indicators, 2–33
 - stripping
 - algorithm, 2–40
 - fragments, 2–40
 - process, 2–40
 - transmitting station, 2–35
 - void frames, 2–40
 - types, 2–19
 - void, 2–40
- Frame Content Independent Stripping, 2–40

H

Halt Line State, 2–11

Index–4

Horizontal distribution frame, 1–20, 3–10

I

Idle Line State, 2–11

Idle symbols

- origin of, 2–14

- use in

- clock recovery, 2–14

- scrubbing, 2–15

- trace function, 2–38

IEEE 802.2, 1–8

IEEE 802.3, 1–8

IEEE 802.4, 1–8

IEEE 802.5, 1–8, 1–11

IEEE Address Administrator, 2–17

In-band management, 3–5

INFO field, 2–22

Interbuilding backbone, 3–10

Intermediate distribution frame, 1–20, 3–10

International Organization for Standardization

- FDDI standards, 1–4

- Open Systems Interconnection model, 1–8, A–1

Internet Protocol fragmentation, 4–7, 6–5

Intrabuilding backbone, 3–10

L

LAN management, 1–30

Line states, 2–8, 2–11

- Active Line State, 2–12

- Halt Line State, 2–11

- Idle Line State, 2–11

- Master Line State, 2–11
- Noise Line State, 2–12
- Quiet Line State, 2–11
- Link Confidence Test, 2–31
- Logical Link Control, 1–6, 1–8, 2–29, 2–30

M

- MAC. *See* Media Access Control
- Main distribution frame, 1–20, 3–10
- Maintainability, 1–28
- Master Line State, 2–11, 2–38
- Media Access Control
 - data link addressing, 2–16
 - FDDI standard, 1–1
 - functions, 2–15
 - ring scheduling, 2–16
 - universal addressing, 2–17
- Media Interface Connector
 - description, 2–7
 - keying, 2–7
 - MIC A, 2–7
 - MIC B, 2–7
 - MIC M, 2–7
 - MIC S, 2–7
- Multimode fiber, specifications, 2–4

N

- Neighborhood Information Frame
 - function, 2–28, 6–7
 - uses, 6–7
- Network connection model, 1–15, 2–2
- Network layer, A–3
- Network Management, 6–1
 - accounting management, 6–3

- address filtering, 6–6
- benefits, 6–2
- bridge port
 - configuration, 6–6
 - enabling, 6–6
- capabilities, 6–4
- configuration
 - management, 6–3
 - parameters, 6–5
- configuration management, 6–3
- device
 - naming, 6–4
 - registry, 6–4
- Enterprise Management, 6–1
- event
 - log, 6–5
 - message, 6–5
- fault management, 6–3
- firmware upgrade, 6–5
- initialization, 6–6
- IP fragmentation, 6–5
- management functions, 6–2
- neighborhood information frame display, 6–7
- network monitoring, 6–7
- performance management, 6–3
- PHY port
 - configuration, 6–6
 - control, 6–6
- protocol filtering, 6–7
- ring mapping, 6–3, 6–7
- security
 - management, 6–4
 - password, 6–5
- spanning tree parameters, 6–5
- status information frame display, 6–7
- troubleshooting, 6–8

NIF. *See* Neighborhood Information Frame

No Owner frame
definition, 2–41
removal from ring, 2–41

Node, definition, 1–2

NOF. *See* No Owner frame

Noise Line State, 2–12

NRZ/NRZI encoding, 2–13

O

Open Systems Interconnection
function, 1–8
layers, A–3, A–4
model, A–2

Optical bypass relay, 1–10
description, 2–6
limitation of, 2–6

Organizationally Unique Identifier, 2–17

OSI. *See* Open Systems Interconnection

Out-of-band management, 3–5

P

PHY. *See* Physical Layer Protocol

PHY ports. *See* FDDI port types

Physical Connection Management,
2–23, 2–36, 2–38

Physical layer, 1–8, A–3

Physical layer entities
description, 1–30
number, 1–31, 3–9, 4–10, 5–5

Physical Layer Medium Dependent,
1–1, 2–3

alternate fiber specifications, 2–4
connector keying, 2–7
fiber specifications, 2–4

Media Interface Connector, 2–7

optical bypass relay, 2–6

Physical Layer Protocol, 1–1, 2–7

data recovery, 2–8

elasticity buffer, 2–8

encode/decode process, 2–8

line states, 2–8

repeat filter, 2–8

smoothing, 2–8

symbols, 2–8

PMD. *See* Physical Layer Medium Dependent

Preamble, 2–21

Presentation layer, A–4

Protocol data units

definition, 2–19

formats, 2–19

logical link control units, 2–15, 2–30

removal of orphans, 2–14

role of scrubbing, 2–8

valid format, 2–21

Q

Quiet Line State, 2–11

R

Reliability, 1–27

Repeat filter, 2–8, 2–14

Ring

latency, 2–35

maintenance

beacon process, 2–37

function, 2–36

recovery, 2–37

trace, 2–38

operation, 2–30

- purger
 - algorithm, 2–41
 - election process, 2–42
 - function, 2–41
 - No Owner frame removal, 2–41
 - overhead, 2–42
- scrubbing, 2–14
- Ring Management
 - description, 2–27
 - stuck beacon, 2–28

S

- SAC. *See* Concentrator
- SAS. *See* Single attachment station
- Scalability, 1–28
- Scrubbing, 2–8, 2–14
- Service Access Points, 1–9
- Session layer, A–4
- SIF. *See* Station Information Frame
- Single attachment station, 1–2
 - concentrator, 1–9
 - FDDI devices, 1–9
 - model, 1–10
- Single-mode fiber
 - configurations, B–10, B–11
 - specifications, 2–5
- Smoothing, 2–8, 2–14
- SMT. *See* Station Management
- Source
 - address, 2–22
 - routing, 2–17
- Standalone topology, 1–16
- Starting delimiter, 2–21
- Station
 - definition, 1–2
 - power up, 2–31

- Station bypass, 3–6
- Station Information Frame
 - function, 2–28, 6–7
 - uses, 6–7
- Station Management
 - architecture, 2–24
 - Connection Management, 2–23
 - connection rules, 2–25
 - FDDI standard, 1–1
 - frame
 - fields, 2–29
 - format, 2–29
 - services, 2–28
 - function, 2–23
 - PHY-port types, 2–25
 - protocols, 2–23
 - Ring Management, 2–27
- Steady-state operation
 - capturing the token, 2–33
 - control indicators, 2–33
 - description, 2–33
- Stripping. *See* Frame
- Stuck beacon, 2–28
- Symbol
 - code bits, 2–8
 - code group, 2–8
 - definition, 2–8
 - encoding, 2–10, 2–11
 - function, 2–8

T

- T_Bid, 2–32
- T_Neg, 2–33
- Target Token Rotation Time
 - default value, 2–31
 - function, 2–31
 - MAC requirement, 2–16

- negotiation process, 2–31, 2–32
- operational value, 2–31
- THT. *See* Token Holding Timer
- Timed Token Protocol, 2–16
- Timers, 2–30
 - Target Token Rotation Timer, 2–31
 - Token Holding Timer, 2–30
 - Token Rotation Timer, 2–30
 - Valid Transmission Timer, 2–30, 2–36
- Token
 - acquisition, 2–16, 2–20, 2–34
 - definition, 2–15
 - fields, 2–21
 - format, 2–20
 - function, 2–20
 - issuing, 2–16, 2–20
- Token Holding Timer, 2–30
- Token Rotation Timer, 2–30
- Topologies, 1–14
- Trace, 2–38
- Translation, 4–4

- Transport layer, A–3
- Tree of concentrators, 1–17
- TRT. *See* Token Rotation Timer
- TTRT. *See* Target Token Rotation Time
- TVX. *See* Valid Transmission Timer

U

- Universal address
 - format, 2–17
 - unique identifier, 2–17

V

- Valid Transmission Timer, 2–30, 2–36
- Void frames, 2–40

W

- Wallbox, 3–10
- Workgroup, 1–12, 1–14

HOW TO ORDER ADDITIONAL DOCUMENTATION

DIRECT TELEPHONE ORDERS

In Continental USA
and Puerto Rico
call 800-258-1710

In Canada
call 800-267-6146

In New Hampshire
Alaska or Hawaii
call 603-884-6660

DIRECT MAIL ORDERS (U.S. and Puerto Rico*)

DIGITAL EQUIPMENT CORPORATION
P.O. Box CS2008
Nashua, New Hampshire 03061

DIRECT MAIL ORDERS (Canada)

DIGITAL EQUIPMENT OF CANADA LTD.
940 Belfast Road
Ottawa, Ontario, Canada K1G 4C2
Attn: A&SG Business Manager

INTERNATIONAL

DIGITAL
EQUIPMENT CORPORATION
A&SG Business Manager
c/o Digital's local subsidiary
or approved distributor

Internal orders should be placed through Publishing and Circulation Services (P&CS),
Digital Equipment Corporation, 10 Forbes Road, Northboro, Massachusetts 01532-2597

*Any prepaid order from Puerto Rico must be placed
with the Local Digital Subsidiary:
809-754-7575

READER'S COMMENTS

What do you think of this manual? Your comments and suggestions will help us to improve the quality and usefulness of our publications.

Please rate this manual:

	Poor			Excellent	
Accuracy	1	2	3	4	5
Readability	1	2	3	4	5
Examples	1	2	3	4	5
Organization	1	2	3	4	5
Completeness	1	2	3	4	5

Did you find errors in this manual? If so, please specify the error(s) and page number(s).

General comments:

Suggestions for improvement:

Name

Date

Title

Department

Company


Street

City

State/Country

Zip Code

DO NOT CUT – FOLD HERE AND TAPE



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY LABEL


FIRST CLASS PERMIT NO. 33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

digital™

**Telecommunications and
Networks Publications**

550 King Street
Littleton, MA 01460–1289



DO NOT CUT – FOLD HERE

digital

Printed in U.S.A.