NIST Special Publication 500-212

Computer Systems Technology

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards and Technology

NIST

Planning for the Fiber Distributed Data Interface (FDDI)

William E. Burr

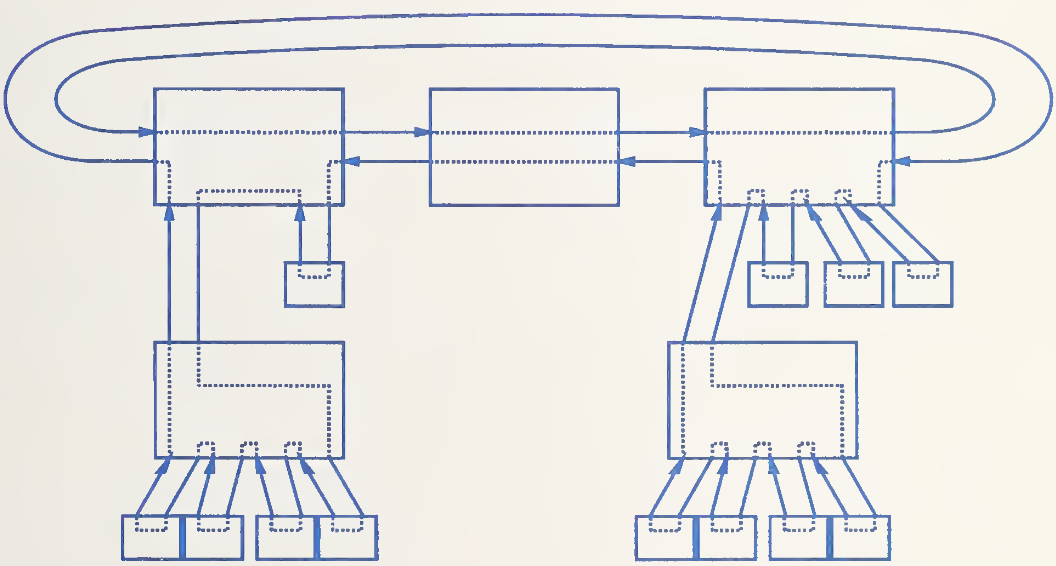

The National Institute of Standards and Technology was established in 1988 by Congress to "assist industry in the development of technology . . . needed to improve product quality, to modernize manufacturing processes, to ensure product reliability . . . and to facilitate rapid commercialization . . . of products based on new scientific discoveries."

NIST, originally founded as the National Bureau of Standards in 1901, works to strengthen U.S. industry's competitiveness; advance science and engineering; and improve public health, safety, and the environment. One of the agency's basic functions is to develop, maintain, and retain custody of the national standards of measurement, and provide the means and methods for comparing standards used in science, engineering, manufacturing, commerce, industry, and education with the standards adopted or recognized by the Federal Government.

As an agency of the U.S. Commerce Department's Technology Administration, NIST conducts basic and applied research in the physical sciences and engineering and performs related services. The Institute does generic and precompetitive work on new and advanced technologies. NIST's research facilities are located at Gaithersburg, MD 20899, and at Boulder, CO 80303. Major technical operating units and their principal activities are listed below. For more information contact the Public Inquiries Desk, 301-975-3058.

Technology Services

- Manufacturing Technology Centers Program
- Standards Services
- Technology Commercialization
- Measurement Services
- Technology Evaluation and Assessment
- Information Services

Electronics and Electrical Engineering Laboratory

- Microelectronics
- Law Enforcement Standards
- Electricity
- Semiconductor Electronics
- Electromagnetic Fields[1]
- Electromagnetic Technology[1]

Chemical Science and Technology Laboratory

- Biotechnology
- Chemical Engineering[1]
- Chemical Kinetics and Thermodynamics
- Inorganic Analytical Research
- Organic Analytical Research
- Process Measurements
- Surface and Microanalysis Science
- Thermophysics[2]

Physics Laboratory

- Electron and Optical Physics
- Atomic Physics
- Molecular Physics
- Radiometric Physics
- Quantum Metrology
- Ionizing Radiation
- Time and Frequency[1]
- Quantum Physics[1]

Manufacturing Engineering Laboratory

- Precision Engineering
- Automated Production Technology
- Robot Systems
- Factory Automation
- Fabrication Technology

Materials Science and Engineering Laboratory

- Intelligent Processing of Materials
- Ceramics
- Materials Reliability[1]
- Polymers
- Metallurgy
- Reactor Radiation

Building and Fire Research Laboratory

- Structures
- Building Materials
- Building Environment
- Fire Science and Engineering
- Fire Measurement and Research

Computer Systems Laboratory

- Information Systems Engineering
- Systems and Software Technology
- Computer Security
- Systems and Network Architecture
- Advanced Systems

Computing and Applied Mathematics Laboratory

- Applied and Computational Mathematics[2]
- Statistical Engineering[2]
- Scientific Computing Environments[2]
- Computer Services[2]
- Computer Systems and Communications[2]
- Information Systems

---

[1] At Boulder, CO 80303.
[2] Some elements at Boulder, CO 80303.

NIST Special Publication 500-212

Planning for the Fiber Distributed Data Interface (FDDI)

William E. Burr

Computer Systems Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899

Sponsored by:
Internal Revenue Service
ISD: I : S : C Room 402
2500 Wilson Boulevard
Arlington, VA 22201

October 1993

U.S. DEPARTMENT OF COMMERCE
Ronald H. Brown, Secretary

TECHNOLOGY ADMINISTRATION
Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
Arati Prabhakar, Director

Reports on Computer Systems Technology

The National Institute of Standards and Technology (NIST) has a unique responsibility for computer systems technology within the Federal government. NIST's Computer Systems Laboratory (CSL) develops standards and guidelines, provides technical assistance, and conducts research for computers and related telecommunications systems to achieve more effective utilization of Federal information technology resources. CSL's responsibilities include development of technical, management, physical, and administrative standards and guidelines for the cost-effective security and privacy of sensitive unclassified information processed in Federal computers. CSL assists agencies in developing security plans and in improving computer security awareness training. This Special Publication 500 series reports CSL research and guidelines to Federal agencies as well as to organizations in industry, government, and academia.

National Institute of Standards and Technology Special Publication 500-212
Natl. Inst. Stand. Technol. Spec. Publ. 500-212, 108 pages (Oct. 1993)
CODEN: NSPUE2

U.S. GOVERNMENT PRINTING OFFICE
WASHINGTON: 1993

For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, DC 20402

Table of Contents

1. Introduction 1

2. Overview of FDDI 2
   2.1 Structure of the FDDI Standards 2
   2.2 FDDI Topology 4
   2.3 Physical Medium Dependent (PMD) 11
   2.4 Physical Layer Protocol (PHY) 15
   2.5 Medium Access Control (MAC) 17
   2.6 Station Management (SMT) 22

3. Overview of FDDI-II 27
   3.1 Structure of the FDDI-II Standards 27
   3.2 Applications and Planning for FDDI-II 31

4. FDDI Network Configuration Considerations 34
   4.1 TTRT and FDDI Performance 34
   4.2 Partitioning FDDI Rings 39

5. FDDI Media 41
   5.1 Optical fiber media 41
   5.2 Twisted Pair Media for FDDI 46
   5.3 Selecting FDDI Media 49

6. FDDI and Standards for Wiring of Office Buildings 50
   6.1 Introduction to Building Wiring Standards 50
   6.2 Summary of FIPS 174 Requirements 50
   6.3 FDDI with structured wiring. 53
   6.4 Wiring Crossovers 54

7. GOSIP, GNMP, Functional Standards and FDDI 55
   7.1 The GOSIP Architecture and FDDI 55
   7.2 FDDI's Role in GOSIP 58
   7.3 FDDI and the GNMP 58
   7.4 Functional Standards. 59

8. Connecting an FDDI LAN to another Network 61
   8.1 Some Definitions 61
   8.2 Bridges and Routers. 63
   8.3 Hubs and Collapsed Backbones. 69

9. FDDI and Newer Emerging Network Technologies 70
   9.1 HIPPI 70
   9.2 Fibre Channel 71
   9.3 SONET 75
   9.4 Frame Relay 77
   9.5 ATM LANs and B-ISDN 78

9.6 DQDB and SMDS 83
9.7 100 Mbps Ethernet 85

Appendix A: Sources for Documents 86

Appendix B: List of Acronyms 88

Appendix C: List of References 93

Index 99

ABSTRACT: The Fiber Distributed Data Interface (FDDI) is an emerging Standard fiber optic Local Area Network (LAN) technology suitable for backbone and high performance workstation applications. FDDI will be widely used by federal agencies in the 1990's. This report describes the FDDI standards and the media that FDDI uses, and provides information about wiring for FDDI LANs and about effectively configuring FDDI LANs. It also describes the relationship of FDDI to the Government Open Systems Interconnection Profile (GOSIP) and discusses connecting FDDI to other networks.

KEYWORDS: CDDI, dual ring network, FDDI fiber optic local network, LAN, local area network, timed token ring.

1. Introduction

The Fiber Distributed Data Interface (FDDI) is a standard for a 100 Mbps fiber optic token ring, now coming into wide use. FDDI's initial application has been primarily as a "backbone" Local Area Network (LAN), connecting together other LANs with lower data rates. Since FDDI supports large physical configurations, it is also seeing use as a Metropolitan Area Network (MAN), connecting LANs together over an entire metropolitan area. These MANs may be private, however some public carriers are also offering FDDI based services. Past costs for FDDI attachment have been too high to allow widespread use of FDDI as a workstation interface, however these costs are falling rapidly as the volume of FDDI sales increases. FDDI is rapidly joining the IEEE 802.3 [ISO88023] and 802.5 [ISO88025] LAN standards as a commodity network interface with attractive price and performance characteristics. While the 802.3 standard seems secure as the most common inexpensive LAN interface, as workstations become more powerful, and imaging applications become commonplace, FDDI will frequently be attached directly to workstations to provide the bandwidth required for image processing.

Section 2 of this report provides an introduction to the FDDI standards and how they work. FDDI has been developed as a layered standard. There are now three published FDDI standards; three others are essentially complete but not yet published. Two remaining alternative physical layer standards are still in development. Products that implement FDDI are now widely available. FDDI-II, an enhanced version of FDDI that includes the FDDI token ring packet service and adds an isochronous service, is discussed in section 3. Three of the four standards that comprise FDDI-II are now essentially complete. FDDI-II shares the same physical media standards as FDDI. The first nascent FDDI-II products are just emerging.

Section 4 describes the practical configuration limits for FDDI rings, and discusses the choice of a value for the *Target Token Rotation Time*, the key user controlled parameter that regulates the operation of the FDDI timed token ring protocol. Section 5 describes the various optical and electrical media that can be used with FDDI and discusses the appropriate use of each. Section 6 then describes the standards for modern structured office building wiring plans and shows how FDDI fits into such a wiring plan.

Section 7 describes how FDDI fits into the Government Open Systems Interconnection Profile (GOSIP) and the Government Network Management Profile (GNMP). FDDI is particularly important as a backbone LAN connecting other LANs. Section 8 therefore describes bridges and routers, the devices that are usually employed to connect LANs together, and describes the specific features added to FDDI to facilitate bridging. Since this is a very rapidly developing market segment, the emphasis is on general principles rather than the specifics of current bridge or router products and protocols.

A variety of newer, standard high speed network technologies are now emerging. These technologies may complement FDDI, compete with FDDI, or provide a growth path beyond FDDI. Section 9 is a brief introduction to several of these technologies, including the *High Performance Parallel Interface (HIPPI)*, the *Fibre Channel*, the *Synchronous Optical Network (SONET)*, *Asynchronous Transfer Mode (ATM)* LANs and the *Distributed Queue Dual Bus (DQDB)* standard.

2. Overview of FDDI

This chapter describes the four FDDI standards, the types of FDDI stations and the FDDI Topology. FDDI is a 100 Mbps *Local Area Network (LAN)* which was designed to use fiber optic media, but may also use copper twisted pair media and be carried over telecommunications transmission trucks, such as SONET. FDDI implements a dual ring of trees topology and uses a timed token medium access protocol. FDDI supports physically large local networks (for example 500 nodes and 60 km of cable). Currently the major use for FDDI is as a “backbone” for other lower rate LANs, however in the future FDDI will commonly be directly attached to high performance workstations and PCs to support data intensive applications, such as image processing.

FDDI uses a token ring network protocol. FDDI stations are logically connected in a closed ring (although they are frequently wired in star or tree fashion). Each FDDI station has an upstream data link, on which it receives data packets (called *frames*), and a downstream data link, on which it transmits frames. Each station continuously monitors all frames received from the upstream link, and, unless the station itself originated the frame, the station repeats those frames on its downstream link. Those received frames originated by the station are not repeated, thus *stripping* them from the ring, after they have gone around the entire ring one time. If a station has frames it wishes to originate, it waits until it receives a special frame called a *token*. The station does not repeat the token until after it has transmitted all the frames it wishes to originate (or until certain time limits have expired). The station then transmits the token on its downstream link.

2.1 Structure of the FDDI Standards

Figure 1 illustrates the seven layer OSI reference model, showing the Physical and Data Link layers each subdivided into two sublayers, as has become customary for LANs. FDDI occupies only the Physical layer and the lower sublayer of the Data Link layer.

FDDI consists of four standards illustrated in Figure 2:

- Physical Layer Medium Dependent (PMD) [PMD] which specifies a 100 Mbps full duplex data channel. Several different variants of PMD have been or are being developed for different transmission media.
- Physical Layer Protocol (PHY) [PHY] which specifies the 4:5 line code and provisions for clocking data around the ring.
- Medium Access Control (MAC) [MAC] which specifies the token passing packet protocol of FDDI.
- Station Management (SMT) [SMT], specifies the protocols for FDDI management, initialization and error recovery.

FDDI supports the *Logical Link Control (LLC)* [ISO88022] sublayer and transports LLC *Protocol Data Units (PDUs).* The LLC PDU transport service provided LLC is interchangeable with the similar service provided by other standard LANs such as the IEEE 802.3 [ISO88023] and IEEE 802.5 [ISO88025] standards. Several different protocol stacks are commonly used above LLC, with FDDI and the other standard LANs, including the OSI stack, TCP/IP and proprietary protocols. In this report we usually assume, particularly for purposes of network management, that layers 3 through 7 are an OSI protocol stack.

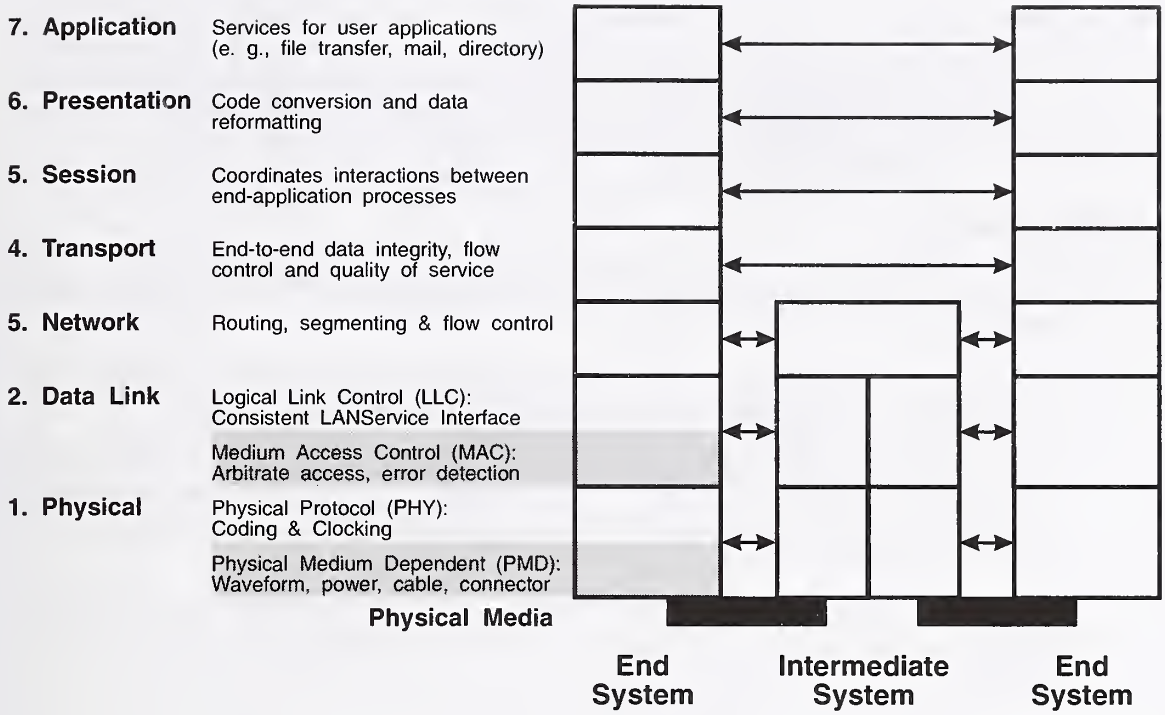

Figure 1 - The OSI Reference Model.

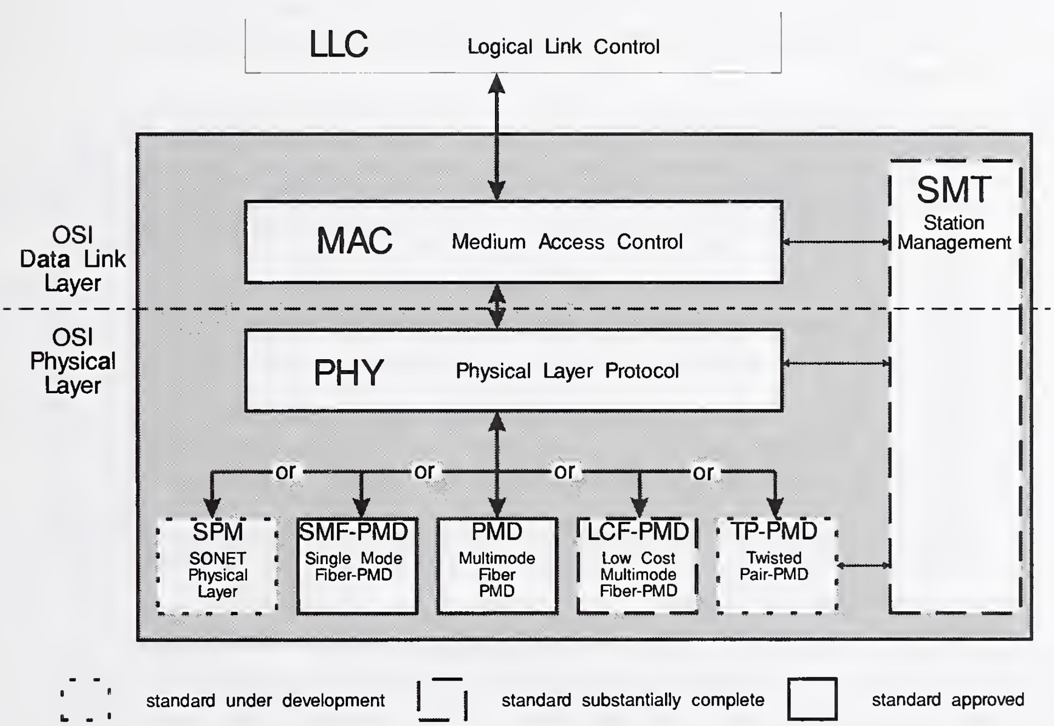

FDDI Standards Architecture

Figure 2 - FDDI Standards Architecture.

Referring again to Figure 1, note that systems are labeled either *end systems* or *intermediate systems*. Intermediate systems, often called *routers*, are used to join separate subnetworks through the routing services of the Network layer. In a single FDDI ring there would be no intermediate system or router and the Network, Data Link and Physical layer entities in the end systems would communicate directly.

2.2 FDDI Topology

The full FDDI topology is called a *dual ring of trees*. FDDI networks are made of *nodes* connected by cables. A node has one or more ports, which consist of a PHY entity and a PMD entity. Four kinds of FDDI nodes are defined:

- Dual Attachment Stations (DAS) : A station with one or two MAC entities which attaches to the dual ring through two ports labeled *A* and *B*
- Dual Attachment Concentrators (DAC) : A node with 0, 1 or 2 MAC entities, which attaches to the dual ring through two ports labeled A and B, and which has one or more *Master (M)* ports through which Single Attachment Stations or Single Attachment Concentrators may be attached to one of the two rings.
- Single Attachment Stations (SAS) : A station with one MAC entity which attaches to one ring through a *Slave (S)* port. A SAS is normally attached to the M port of a concentrator and is the leaf of a tree.
- Single Attachment Concentrators (SAC): A node with 0 or 1 MAC entities which attaches to one ring through an S port and has one or more M ports, through which Single Attachment Stations or Single Attachment Concentrators may be attached to one FDDI ring. A SAS is normally attached to the M port of a concentrator and is the leaf of a tree, however it may also be the root of an entirely tree wired ring with no dual ring.

As we can see above four kinds of ports are defined:

- A Port: A port on a DAS or DAC which should be connected to a B port. The A port has the primary ring in and the secondary ring out. An A-to-A connection causes a crossover of the primary and secondary rings and a twisted dual ring.
- B Port: A port on a DAS or DAC which should be connected to an A port. The B port has the primary ring out and the secondary ring in. A B-to-B connection causes a crossover of the primary and secondary rings and a twisted dual ring.
- M Port: A master port on a concentrator to which should be connected to an S port on another concentrator. M- to-A or M-to-B connections are unusual, but may occasionally be useful. M-to-M connections are forbidden.
- S Port: A port on a single attachment station or concentrator which is normally connected to an M port on a concentrator. S-to-A, S-to-B and S-to-S connections may occasionally be useful.

The simplest form of the FDDI topology is the dual ring, which is composed of Dual Attachment Stations. A simple two node dual ring is illustrated twice in Figure 3. Figure 3 (A) illustrates ports and paths. Each dual attachment station has an A port and a B port. Each station also has two *paths* through which data is repeated. A path may or may not have a MAC; in Figure 3 there are 4 paths through nodes and three have a MAC while one does not. Note that PMD PHY and MAC consist of separate transmitter and receiver functions. A path passes first through the receiver functions and then through the transmitter functions in the order PMD, PHY and MAC (if present) receiver then MAC (if present), PHY and PMD transmitter.

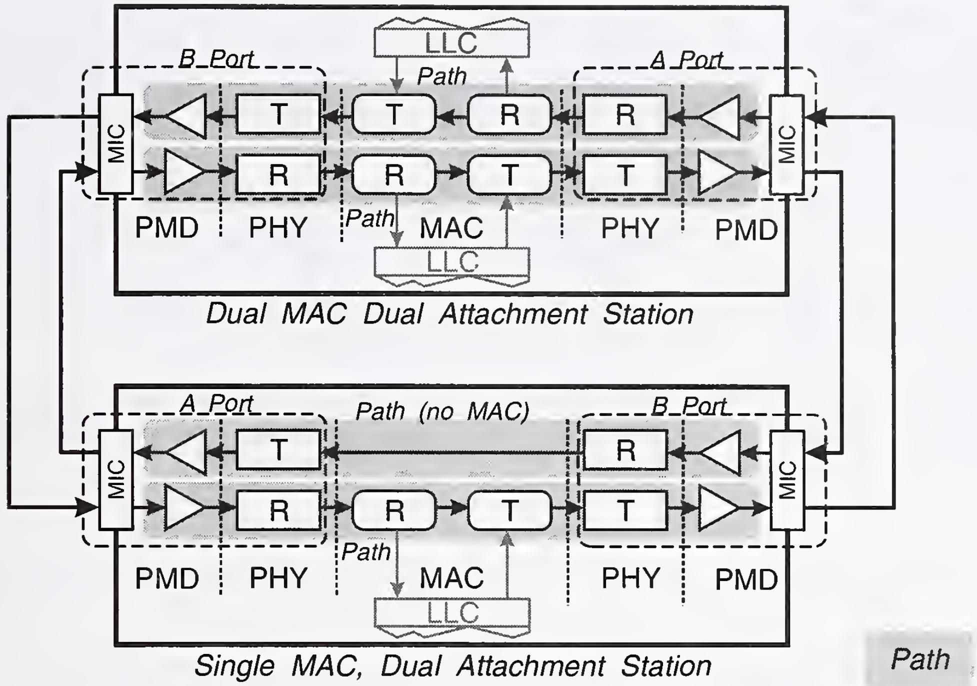

a) Ports and Paths

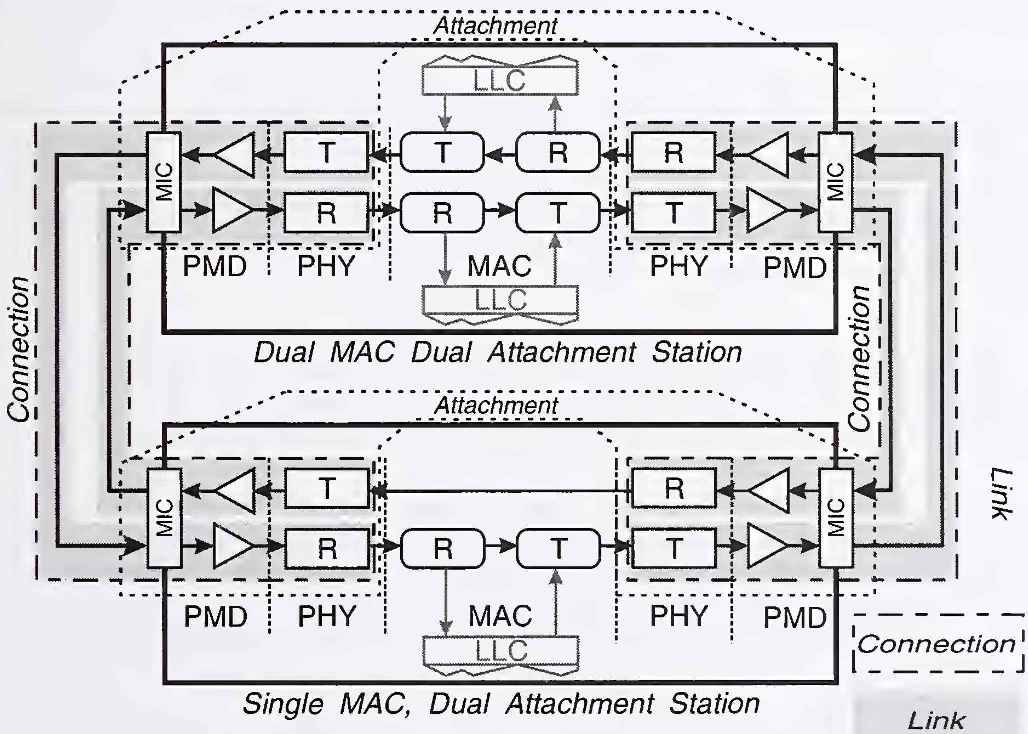

b) Connections, Links and Attachments

Figure 3 - The Parts of an FDDI Ring.

In this figure we show dual attachment stations in the thru state and each path enters through the receive side of one port and exits through the transmit side of the other. In single attachment stations there is only one path and port and the path enters and exits through the same port. Dual attachment stations may also be configured so that a path enters and exits through the same port. Examples of other path configurations can be seen in Figures 4, 5, 6, and 7.

Figure 3 (b) illustrates *attachments*, *connections* and *links*. A port or a pair of ports which is managed as a unit is called an attachment. In Figure 3 both stations are dual attachment stations; a dual attachment always consists of an A port and a B port and may optionally include an optical bypass switch (not shown), which allows the two paths through the station to be optically bypassed when a port or path is broken or the station is turned off.

A PHY and PMD transmitter entity connected by a cable to a PHY and PMD receiver entity is called a *link*. Links are identified by the data transmission medium, that is multimode fiber, single mode fiber, twisted pair and so on.

Two links between two ports form a connection. In Figure 3 (b) there are two connections. Connections are identified by the type of port at each end. Most FDDI networks will consist entirely of A-to-B connections (called peer connections) and M-to-S connections. The duplex connectors on the ends of cables may be keyed to prevent insertion in the wrong type of port. Any kind of connection may use any type of link, and we will discuss the various types of links below in the section on PMD.

Figure 4(a) illustrates a dual ring with 4 nodes. Every connection is an A-to-B peer. All four nodes are in the *thru* state, resulting in a dual counter rotating ring. Note that we have simplified this and the following figures by drawing only the ports and omitting the separate PHY and PMD functions. In this case every path enters through one port and exits through the other.

In Figure 4(a) some stations have one MAC, while other stations have two MACs. Frames are transmitted or received by a MAC entity and one path through a single MAC dual attachment node does not have a MAC. For this reason one ring is designated a *primary* ring while the other ring is designated the *secondary* ring. Single MAC dual attachment nodes usually have their one MAC on the primary ring. This insures that every node with at least one MAC can send frames to every other node.

Other strategies may be followed. Some single MAC stations may be configured with the MAC on one ring while others are configured with the MAC on the other. This may balance the load on two rings rather than one. In this case, however, some mechanism, such as a bridge, may be needed to allow MACs on one ring to communicate with MACs on the other ring.

The dual ring allows a single fault to be bypassed as illustrated in Figure 4(b). Here we see that nodes 2 and 3 are in a *Wrap* state around a connection fault between them. Node 2 is said to be in the *Wrap_A* state since its A port is still in the ring, while node 3 is in the *Wrap_B* state. In the wrap condition paths now enter and exit through the same port. There is now only one ring and it may be nearly twice as long as the original dual ring. Where the total bandwidth on the "unwrapped" dual ring was 200 Mbps, it is only 100 Mbps on the wrapped ring.

With only one fault, there is one network and every node can still communicate with every other node. A second fault, as illustrated in Figure 4 (c), however, can divide the network into two separate, non communicating rings.

The dual ring provides effective protection against disruption of the network due to a single bad connection. It is particularly useful in the backbone of a large building or a campus. However, if two connections in the same dual ring run through the same conduit or cable way, then a single accident can still divide the dual ring. Therefore, if high network availability is a major concern, it is important to ensure that only one FDDI connection in a dual ring uses a single conduit or raceway. This may not be practical in some buildings with a single vertical riser for cables, or where a single tunnel or conduit connects separate buildings, but the dual ring is vulnerable whenever a single accident can cut two connections.

Turning a dual ring node off may have the same effect as a break in a cable: it may cause the ring to wrap. Dual attachment nodes may optionally include an optical bypass switch which allow a powered-off node to remove itself from the ring to prevent a wrap of the dual ring. This, however, complicates link power budget design; several successive wrapped stations may have too much loss in the cables and switches. Backup power supplies for dual attachment nodes also provide some protection from wrapped rings.

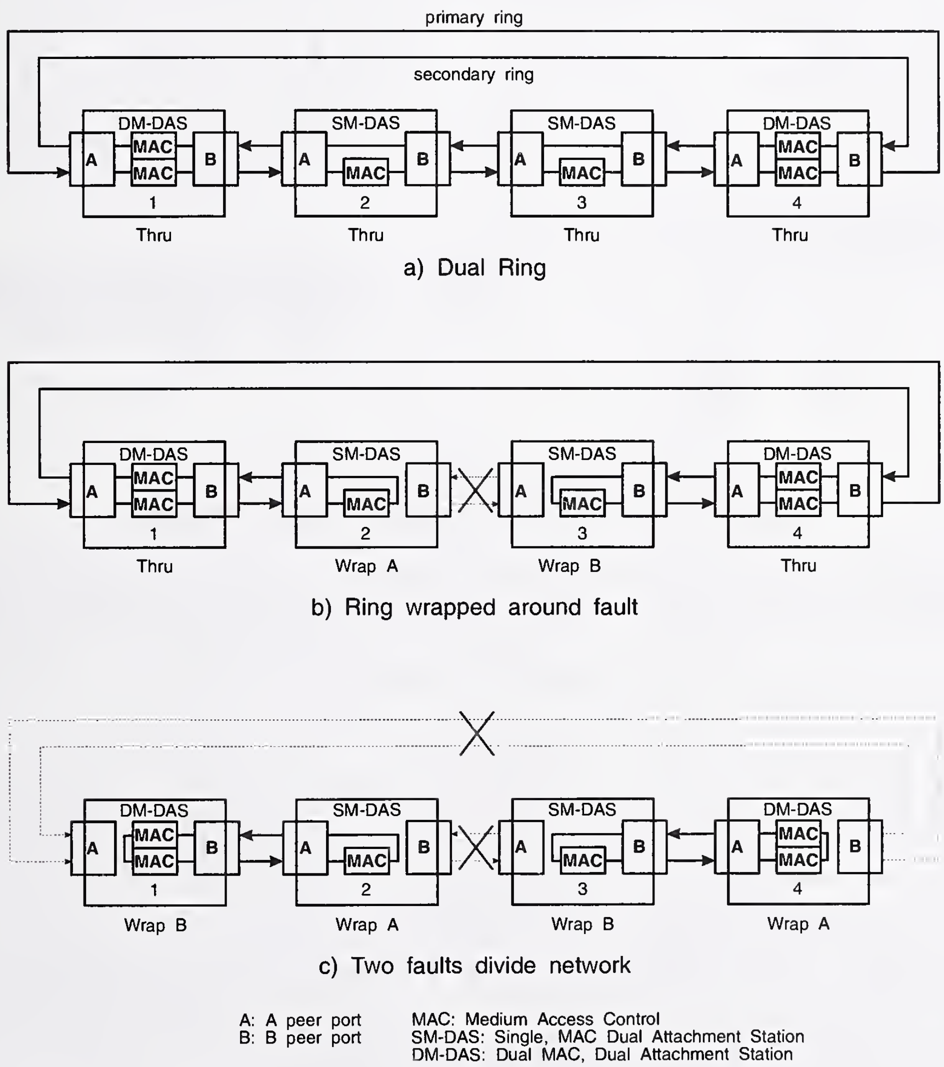

Figure 4 - The FDDI Dual Ring and Fault Recovery.

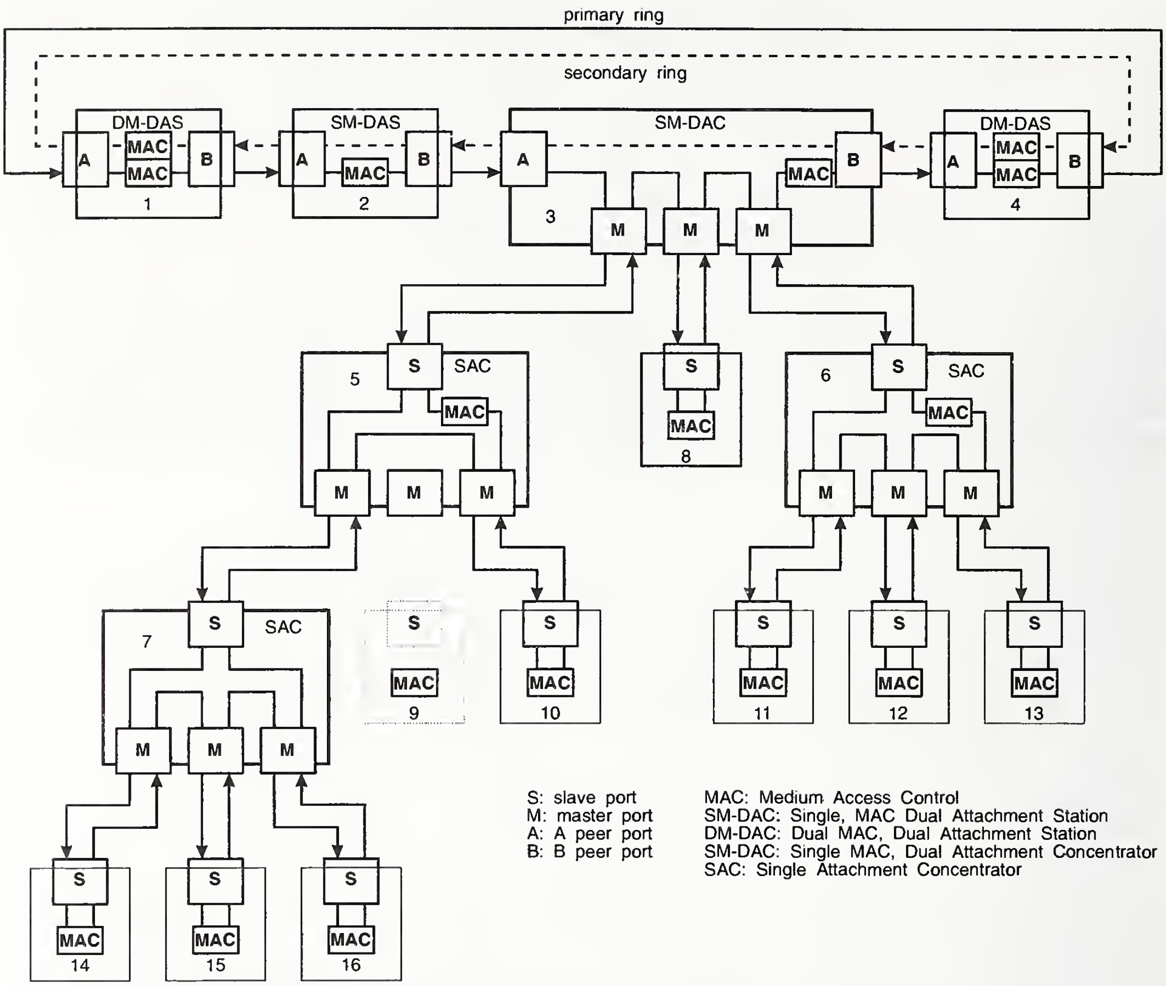

Figure 5 - The FDDI Dual Ring of Trees Topology.

Finally, a user who simply unplugs a cable on the dual ring, will cause a wrap. For this reason it is desirable to limit user access to the dual attachment nodes and their cables, by locating them in wire closets or computer rooms from which untrained personnel are excluded. Dual attachment stations, then should not generally be used to attach workstations to an FDDI network.

FDDI also provides for tree wired rings and allows them to be combined with dual rings to form a topology called a *dual ring of trees* which is illustrated if Figure 5. An FDDI dual ring of trees ordinarily contains only A-to-B and M-to-S connections. In Figure 5, nodes 1, 2, and 4 are DASs in the dual ring. Node 3, a DAC, is the root node of a tree of SACs and SASs. The M port on each concentrator senses when the attached S port is active and inserts or removes it from the ring. In Figure 5, node 9 is an SAS which is not active and has been bypassed by its parent concentrator, node 5. When node 9 is turned on the M port will sense this, and the CMT function in SMT will initialize the connection. The concentrator will then insert the connection into the ring. Note that the paths in single attachment stations exit through the same port in an SAS, while many path configurations are possible in concentrators.

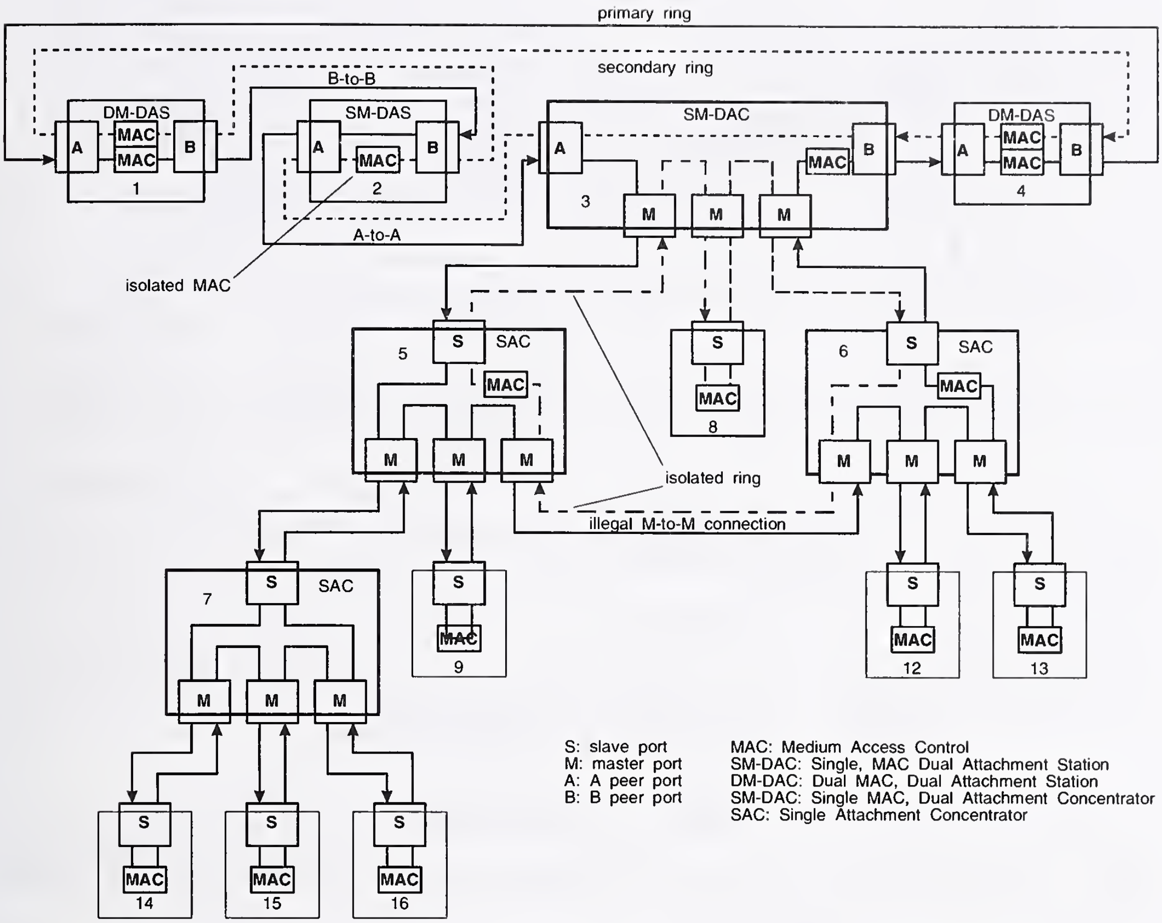

Figure 6 - FDDI Topology Errors.

It is apparent that turning an SAS off, or disconnecting it has relatively little effect on the network, since an SAS is always a leaf on the tree and an inactive S port is simply bypassed by its parent concentrator. Therefore workstations or other user accessible nodes are normally connected to an FDDI network as an SAS.

Turning an SAS on or off, or plugging it in or out of the network does have some consequences to the network, however. Every time a node is inserted in or removed from the network the ring is interrupted and packets may be lost. Moreover the ring must be reinitialized, using a process described in 2.5.4 below. This generally takes only a few milliseconds, but it does prevent communications during this period. In most networks this will be a minor effect, but an SAS which is rapidly cycled on and off could seriously disrupt the network.

We noted above that nearly all FDDI connections will be either A-to-B or M-to-S. In fact, any network formed of only these connections must be a correct dual ring of trees and some other kinds of connection is considered suspect. Figure 6 illustrates some of the possible adverse consequences of other types of connections. In this case node 2, a single MAC DAS, is connected to its neighbors by A-to-A and B-to-B connections. The result is a "twisted" dual ring, where the primary and secondary rings become confused. In this particular case the result is that the MAC in node 2 is twisted onto the secondary ring and is isolated from all the SAS

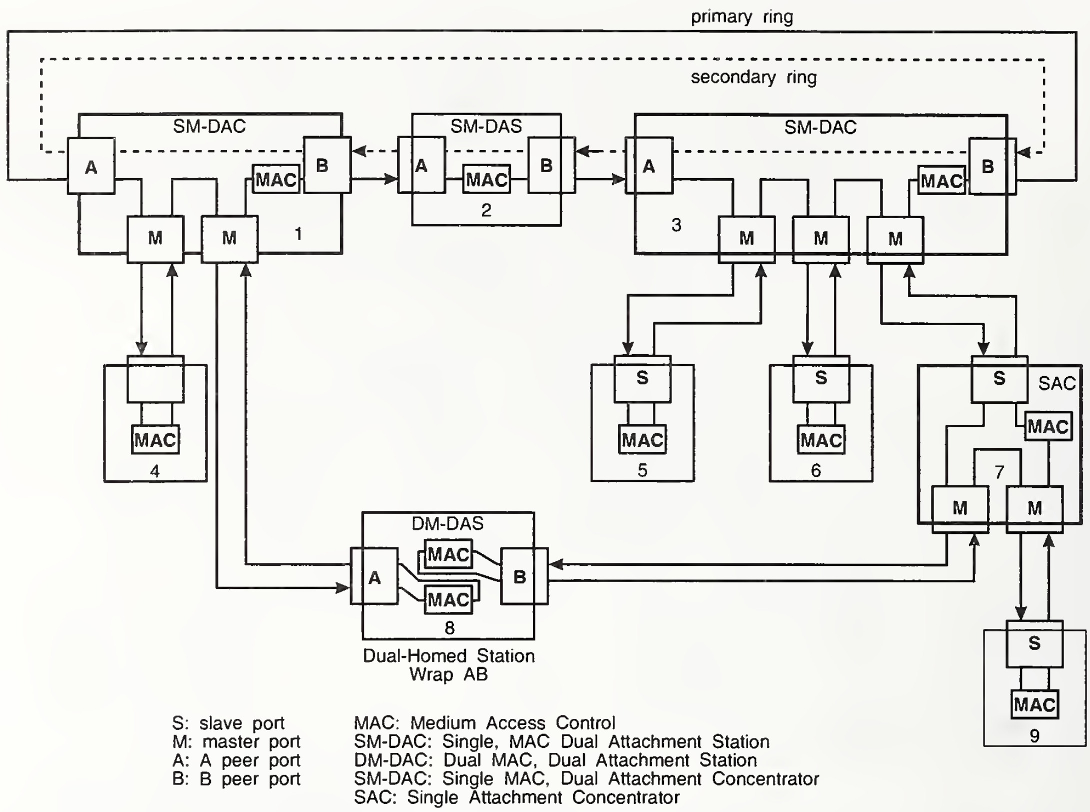

Figure 7 - Dual Homing.

MACs. Note that while for every A-to-A link, there must be a B-to-B link, but they do not have to be adjacent, as they are in the example.

Figure 6 also illustrates a more serious topology violation, caused by an M-to-M link between concentrator nodes 5 and 6. A third, entirely isolated ring is the result and the MACs in nodes 5 and 8 are entirely isolated from the network. There are some occasional uses, however, for all the link types except M-to-M (which is forbidden).

The most significant use for an unusual connection is called *dual homing*, which is illustrated in Figure 7. It uses M-to-A and M-to-B links to provide redundant connections between a DAS and two concentrators. It may be used where it is important that a node not be isolated from the network by a single fault. In a dual homing configuration node 8, a DAS, remains in a *Wrap_AB* state and both paths enter and exit through one port. Node 8 must never go to the Thru state, or the result will be a third ring isolated from the rest of the network.

During connection initialization the *Connection Management (CMT)* function of SMT in each node determines the type of the remote port. An M-to-M connection is forbidden and will not be made active. All other connections will be accepted by CMT, however for every type except A-to-B and M-to-S, M-to-A and M-to-B, SMT will broadcast a *Status Information Frame (SIF)*, specifying an *Undesirble Connection Event*, to notify the network manager of a possible problem.

2.3 Physical Medium Dependent (PMD)

The PMD sublayer provides for the transmission of code bits between FDDI *P*MD entities. Two PMD entities are connected by a cable to form an FDDI link, which is a full duplex channel between the ports. While a connection is identified by the types of ports connected, a link is identified only by the type of medium and any type of link can support any type of connection. This section describes the various PMD interface standards while the fiber or twisted pair media used with PMD are discussed in section 5.

The data transfer rate on and FDDI link is 100 Mbps, however the link signaling rate is generally 125 Mbps, since all FDDI links defined to date use a 4:5 code, described in more detail under PHY below.

Five variants of PMD have been defined for different media and applications. These are described below.

2.3.1 Multimode PMD (PMD)

Multimode fiber PMD [PMD], is the original PMD. This is the version of PMD that is currently available, and references to PMD without further qualification mean multimode fiber PMD. PMD uses 1300 nm Light Emitting Diode transmitters with 62.5/125 μm graded index multimode fiber. The power budget is 9 dB for a *Bit Error Rate (BER)* of 10^{-12}. However, with an 11 dB link loss, the bit error rate must not exceed 2.5×10^{-10}. It is possible to meet the intended BER for a bit traversing the entire ring, which is 10^{-9}, if not more than four connections have a BER between these two values.

Links distances are usually signal dispersion limited rather than power limited and PMD allows for links of up to 2 km. It is likely that PMD will continue to be the variant used in building trunks, but a lower cost variant is being developed for the concentrator to desktop role.

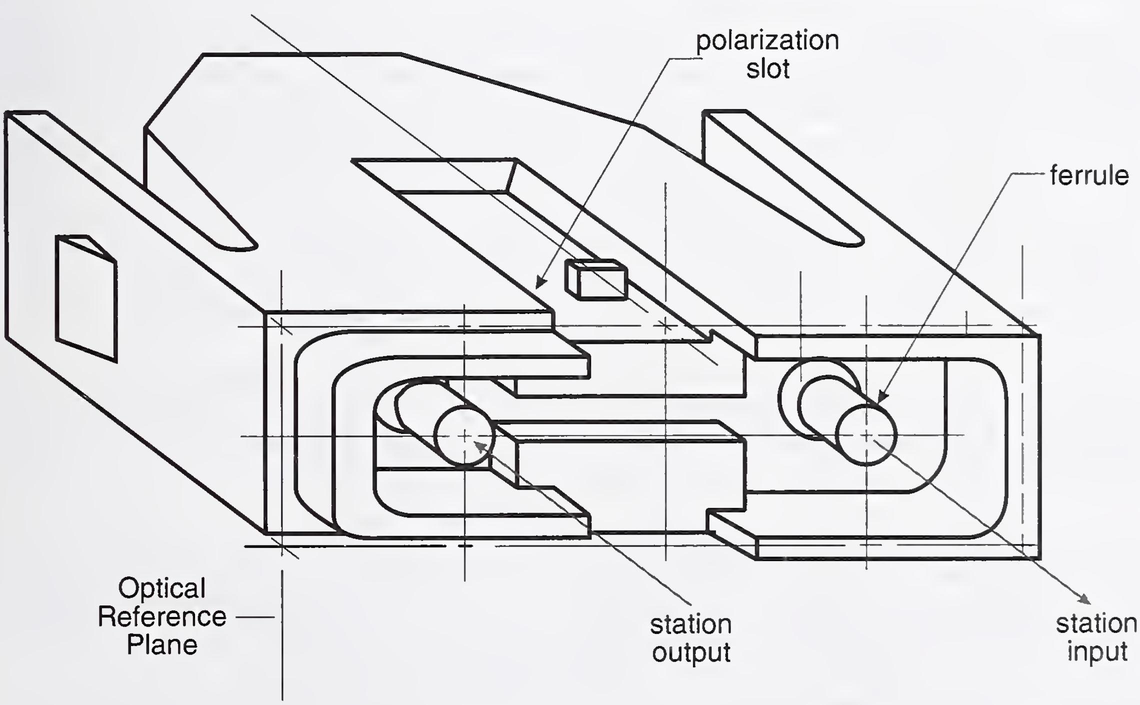

Figure 8 - The FDDI MIC Plug.

The PMD standard defines the receptacle for a *media interface connector (MIC)* connector plug, rather than the plug itself. The MIC plug is depicted in Figure 8. It is a duplex plug with two ferrules containing the fibers. A fixed shroud protects the ferrules, and the plug is polarized with a slot on the top. When the plug is viewed from the open end with the polarizing slot above, light emerges from the right ferrule and enters the left ferrule. Four kinds of ports are defined labeled *A, B, S* and *M* (see sec. 2.2 above) and receptacles are keyed. Plugs may or may not also be keyed. Keying helps to prevent topology errors such as those shown in Figure 6.

2.3.2 Low Cost Fiber PMD (LCF-PMD).

Low Cost Fiber PMD, referred to as *LCF-PMD* [LCF], is intended to provide a lower cost than PMD at the expense of the capability to cover long distances and a reduced capacity to use bypass switches. LCF-PMD will generally be compatible with PMD and uses the same 1300 nm wavelength and 62.5/125 μm multimode fiber. The specifications for the transmitter and receiver parts have been relaxed, however so that the power budget is reduced to 7 dB, and the chromatic characteristics are relaxed so that only 500 m of fiber are allowed. There is no problem connecting LCF_PMD ports to PMD ports, provided the link distance is less than 500 m and the loss does not exceed 7 dB.

The 7 dB power budget is sufficient for 500 m of fiber plus 5 connectors; this is equivalent to passing through two wiring closets and one wall plug, however no power is provided for bypass switches. Although LCF-PMD may find limited application in trunk dual rings, its major intended market is the concentrator to workstation master-slave connections.

A new duplex Low Cost Fiber Media Interface Connector (LCF-MIC) receptacle is also defined for LCF-PMD, based upon the simplex SC connector. The SC connector is a straight pull connector, so it is comparatively easy to join two SC connectors together into a single duplex MIC. Most of the parts remain identical to the simplex SC. Since the simplex SC has come into wide use in building wiring, and is expected to be adopted by IEC building wiring standards, it is believed that the volume of the common parts will result in a less expensive MIC for FDDI. The LCF-MIC is also somewhat smaller the PMD-MIC, since the ferrules are on 1.27 cm (.5 in) centers, rather than 1.79 cm (.7 in) centers. It is quite possible that the LCF_MIC will come to be used in regular PMD stations as well, since it is smaller. Confusingly, Appendix G of the LCF-PMD standard also presents an alternative duplex connector, based upon the ST simplex connector. Although this may be a satisfactory connector, it raises the specter of needlessly confusing users about the correct cables to use with FDDI stations, and it is recommended that the dual SC LCF-MIC be specified when procuring LCF-PMD equipment.

The LCF-MIC is polarized, but not keyed. That is it cannot be inserted backward, but there is no keying to prevent plugs intended for one port type (S, M A or B) from being inserted in another. The LCF-PMD standard, however, recommends a labeling scheme, which is applicable to all FDDI ports and connectors. That scheme is shown in Table 2 below.

2.3.3 Single Mode Fiber PMD (SMF-PMD)

Single Mode Fiber PMD [SMFPMD], is referred to as *SMF-PMD*. SMF-PMD uses single mode optical fiber, and 1300 nm laser sources. Two classes of single mode transmitters and receivers are defined, category 1 and category 2. Category 1 transmitters transmit less power than category 2 transmitters and category 1 receivers are less sensitive than category 2 receivers. The receiver and transmitter categories need not be the same in a port, that is it is possible to have a port with a category 1 transmitter and a category 2 receiver, or a category 2 transmitter and a category 1 receiver. Similarly, a category 1 transmitter can transmit to a category 2 receiver and vice versa.

Category 2 receivers are sensitive and can be overdriven by both category 1 and category 2 transmitters, and category 2 transmitters are also able to overdrive category 1 receivers. Therefore, except for SMF-PMD links that are category 2 on both ends, it is necessary to ensure a minimum loss in the cable to avoid overdriving the receiver. A category 1 transmitter connected to a category 1 receiver provides a link power budget of 11 dB (enough for at least 20 km of fiber), while a category 2 transmitter connected to a category 2 receiver requires a minimum cable loss of 11 dB and provides a link power budget of 33 dB (which is enough for more than 50 km of fiber). For cost reasons SMF-PMD is normally used only where links greater than 2 km are needed. A MIC receptacle similar to the PMD connector is specified.

2.3.4 Twisted Pair PMD (TP-PMD)

Twisted Pair PMD is referred to as *TP-PMD*. Twisted pair media includes shielded twisted pair, unshielded twisted pair and special "category 5 data grade" unshielded twisted pair copper wires. Several proprietary commercial products currently exist for FDDI over either shielded or unshielded twisted pair. The FDDI standards committee is developing a standard solution that works over either shielded twisted pair media or type 5 data grade unshielded twisted pair wires for distances of up to 100 m. While type 5 cable is now commercially available and is suitable for other LANs and telephone applications, few existing twisted pair cable plants have this high quality cable already installed. A three level code, called *MLT-3*, has been selected (rather than the NRZI code used in other PMD variants) to reduce the bandwidth required, while scrambling will spread the spectrum to avoid emission peaks (a difficult challenge for FDDI over unshielded twisted pair is to meet FCC emission limits). At the time this report was written, the general nature or the TP-PMD standard was clear, and there appeared to be little doubt about the feasibility of the standard, but many details were still not resolved.

The FDDI committee is also planning to develop a second variant of TP-PMD that uses sophisticated digital signal processing techniques (this is obviously challenging at 100 Mbps) to permit the use of old, lower quality twisted pair wiring, or perhaps allows longer distances with the new category 5 data grade shielded twisted pair.

Table 1 - Characteristics of PMD Variants

Name	Medium	Maximum Distance	Power Budget	Transmitter
PMD	62.5/125 μm multimode optical fiber, NA: 0.275	2 km	11 dB	1300 nm LED
LCF-PMD		500 m	7 dB	
SMF-PMD, Cat. I	125 μm cladding single mode fiber	20 km	10 dB	1300 nm laser
SMF-PMD, Cat. II		64 km	32 dB	
TP-PMD	150Ω shielded Cu twisted pair or 100Ω data grade category 5 unshielded Cu twisted pair	100 m	dispersion & EMI limited.	electrical
SPM	SONET STS-3	unlimited	N/A	N/A

Table 2 - Labels for FDDI Ports

Code Meaning	Code Letter
Connector Type	
MIC	M
LCF-MIC SC type	2
LCF-MIC ST type	T
RJ-45	4
DB-9	D
PMD Type	
Original multimode fiber PMD	4
LCF-PMD	L
TP-PMD (DTP)	D
TP-PMD (STP)	S
TP-PMD (UTP)	4
SMF-PMD (rcv/xmit type 1)	T
SMF-PMD (rcv/xmit type 2)	2
SMF-PMD (rcv type 1, xmt type 2)	4
SMF-PMD (rcv type 2, xmit type 1)	4
SPM	4
Port Type	
A port	A
B port	B
Master Port	M
Slave Port	S

2.3.5 Sonnet Physical Mapping (SPM)

SONET Physical Mapping, referred to as SPM. The *Synchronous Optical NETwork (SONET)* is a new standard for high speed long distance data communications trunks (see 9.3). It is being widely implemented in the public network. SONET implements a hierarchy of carriers, including STS-1 with a capacity of approximately 50 Mbps, STS-3 at approximately 150 Mbps and STS-12 with a capacity of approximately 600 Mbps. SPM will map FDDI code bits requiring a bandwidth of 125 Mbps over an STS-3 carrier, leaving about 25 Mbps free for other use. When the SPM is used there is no limit to the length of an FDDI link, however token passing delays will limit the performance of very long FDDI networks. The international equivalent to SONET is called the *Synchronous Data Hierarchy (SDH).*

2.3.6 Summary of FDDI PMD Types

The characteristics of the five PMD variants are summarized in Table 1.

The multiplicity of FDDI media, connector types and port types leads to confusion, particularly when users simply wish to order an appropriate cable. As a result, Annex H of the LCF-PMD standard provides a system for labeling FDDI MIC receptacles. Table 2, shows the recommended labels. The port label format is:

FDDI-XYZ, where

X is the connector type,

Y is the PMD type
Z is the Port type

For example:

FDDI-CLS refers to an **SC** type **LCF-MIC** connector on an **LCF-PMD** interface, which is acting as **Slave** port.

FDDI-MPM refers to a PMD MIC connector on an original multimode PMD interface, which is a **Master** port.

2.4 Physical Layer Protocol (PHY)

PHY handles the coding and clocking of data. Figure 9 shows a data path through an FDDI node and illustrates the the various PHY functions. There is always at least one path through an FDDI node and there may be many paths. Received code bits are input from PMD to PHY. In all FDDI PMD variants, except for twisted pair, code bits are transmitted as *a Non-Return to Zero Inverted (NRZI)* signal where a transition represents a 1 bit value and no transition represents a 0 bit value. Twisted pair FDDI uses a three level code, called *MLT-3*, to signal code bits. The PHY clock recovery function first derives the clock from the incoming NRZI bit stream. The FDDI 4:5 code ensures that a signal transition (code bit 1) is available for clock recovery at intervals of not more than 4 code bit times. There is no master clock in an FDDI ring and data is reclocked in each station using a local 67.5 MHz oscillator.

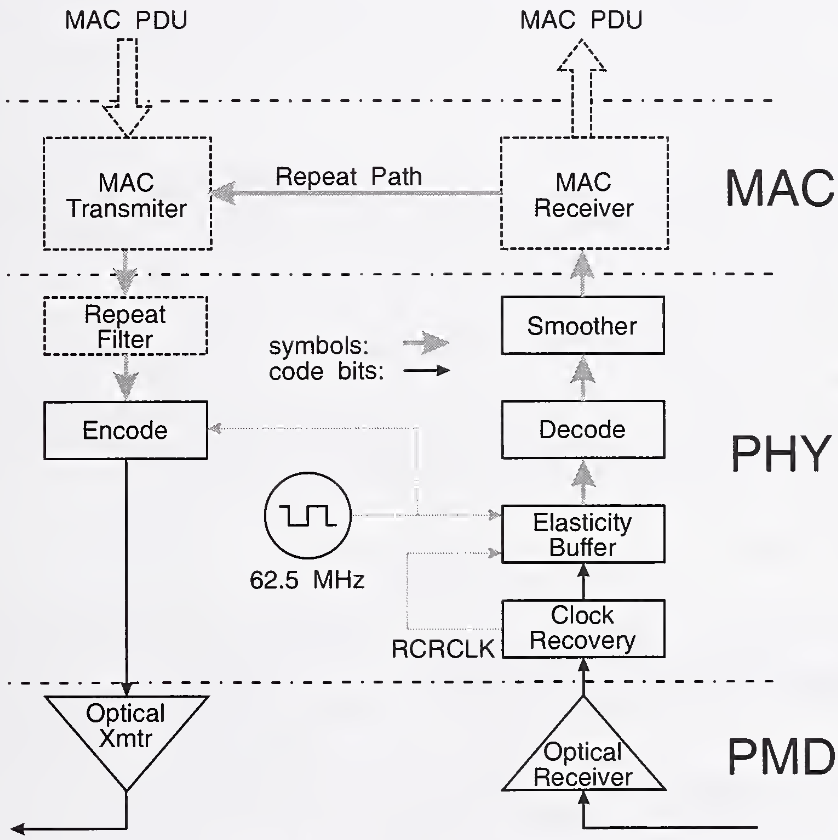

Figure 9 - The FDDI Repeat Path.

Table 3 - FDDI Symbol Set

Binary Value	Notation	Assignment	Binary Value	Notation	Assignment
00000	Q	Quiet	10000	V	Violation
00001	V	Violation	10001	K	Start delimiter 2
00010	V	Violation	10010	8	Data
00011	V	Violation	10011	9	Data
00100	H	Halt	10100	2	Data
00101	V	Violation	10101	3	Data
00110	V	Violation	10110	A	Data
00111	R	Reset	10111	B	Data
01000	V	Violation	11000	J	Start delimiter 1
01001	1	Data	11001	S	Set
01010	4	Data	11010	C	Data
01011	5	Data	11011	D	Data
01100	V	Violation	11100	E	Data
01101	T	Ending Delimiter	11101	F	Data
01110	6	Data	11110	0	Data
01111	7	Data	11111	I	Idle

The elasticity buffer is a FIFO which clocks data in using the clock derived from the input data stream and clocks data out using the local clock. It inserts or removes bits from the Idle sequence between packets, which is a continuous stream of code bit ones, to account for the differences there may be between the local clock rate and the upstream clock rate. It is the size of the required elasticity buffer and the specifications on clock tolerances that limit the maximum frame, described under MAC below.

The Decoder converts the NRZI code bits into 5-bit FDDI symbols. The FDDI 4:5 code is given in table 1. It defines 16 data symbols, 0 through F and the special control symbols labeled H, I, J, K, Q, R, S, and T. Eight illegal or Violation symbols are also defined; these symbols violate the clocking and DC balance constraints of the code. The I symbol is all ones and is used for the idle stream between packets. The J, K, T, R, and S symbols are used to form frames as described in 2.5.1 below. In particular the unique symbol pair JK marks the start of a frame and provides data byte alignment. The H symbol is a special symbol used during connection initialization to mark certain errors during ring operation. Q is "Quiet," meaning that there are no transitions and the connection is not operational. Figure 10 illustrates the relation of FDDI symbols to code bits and the NRZI and MLT-3 transmission codes.

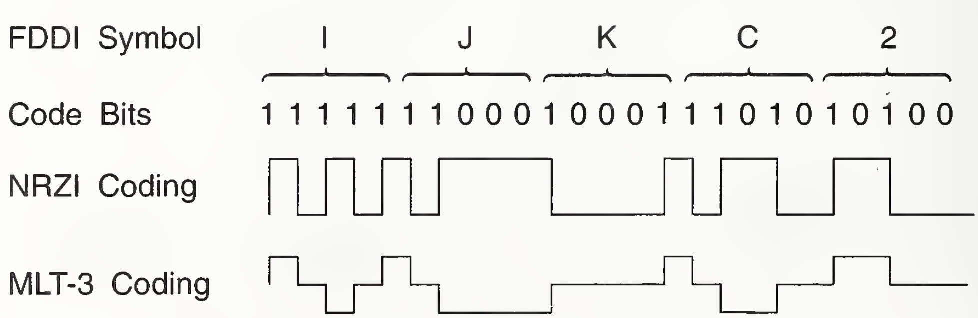

Figure 10 - FDDI Code Example.

The smoother adds bits to the idle stream between frames if it becomes too short (because of the action of the elasticity buffer) and removes bits from the idle streams with excess symbols to make up for symbols it adds to the short idle streams. Without the smother there is a chance that the operation of successive smoothing buffers could destroy frames by removing the idle stream between the frames.

A MAC may or may not be located on an FDDI path. The operation of the MAC, if present, is discussed in the MAC section below.

The repeat filter must be present if a MAC is not included on a path. It is optional if a MAC is present; in effect it duplicates some of the functions of a MAC associated with repeating data. The repeat filter repeats all symbols following the special JK starting delimiter until an I symbol, a Violation symbol or a lone J or K symbol are encountered. After an I symbol the repeat filter transmits I symbols until the next JK pair is received. In the Case of a Violation or individual J and K symbols, the repeat filter transmits four H symbols then transmits I symbols until a new JK pair is received. This prevents the propagation of errors and facilitates the determination of the connection where bit errors causing violations occurred.

The encoder converts the FDDI symbols into bit serial NRZI form and passes this to the PMD transmitter.

2.5 Medium Access Control (MAC)

FDDI MAC is a sublayer of the Data Link Layer (layer 2) of the Open Systems Interconnection Reference Model (OSI) [ISO7498]. MAC provides a packet data transmission service to the layer above it, typically the IEEE 802.2 Logical Link Control (LLC) sublayer.

2.5.1 FDDI Frame

A packet is called a *frame* in FDDI, and in OSI terms a frame is a MAC *Protocol Data Unit (PDU)*. The FDDI frame is illustrated in Figure 11. The frame is composed of FDDI *symbols*, which have code 5 bits and represent the 4-bit hexadecimal values 0 to F, plus certain control symbols. Symbols are discussed in more detail in 2.4 above. An FDDI frame consists of:

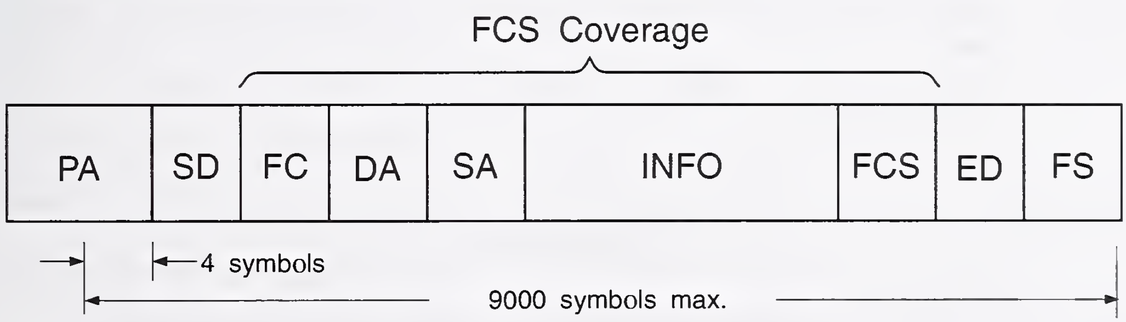

PA: Preamble (16 or more symbols)
SD: Starting Delimiter (2 symbols)
FC: Frame Control (2 symbols)
DA: Destination Address (12 symbols)
SA: Source Address (12 symbols)
INFO: Information (0 or more symbol pairs)
FCS: Frame Check Sequence (8 symbols)
ED: Ending Delimiter (1 symbol)
FS: Frame Status (3 or more symbols)

One FDDI symbol is 4 data bits and 5 code bits

Figure 11 - FDDI Frame Format.

- A *Preamble*, consisting entirely of the special Idle symbol. The preamble maintains clock synchronism between frames. The minimum preamble transmitted is 16 symbols, and a MAC is not required to be able to copy a frame addressed to it with a preamble of less than 12 symbols.

- A *Starting Delimiter (SD)*, consisting of the special unique symbol pair, JK. The starting delimiter establishes frame and byte alignment. A MAC repeats all symbols following an SD until an ED (see below) is encountered, or until an invalid symbol is encountered, unless it is required to strip the frame (see SA below).

- A *Frame Control (FC)* field, that establishes the type of frame, for example a MAC frame, an SMT frame, or one of the special control frames such as the *Token*, the *Beacon Frame*, or the *Claim Frame* discussed below. The FC field, like all fields between the SD and ED (see below) uses the 16 data symbols 0 through F, therefore one symbol is equivalent to 4 bits.

- A *Destination Address (DA)* field, that gives the address of the destination MAC. If it has the available buffer space a MAC copies every frame with its DA. FDDI addresses are normally 48 bits (12 symbols) as shown in Figure 12, although a 16 bit (4 symbol) variant is defined.

- A *Source Address (SA)* field which MACs monitor to determine if they are required to strip frames. A MAC strips any frame with its own SA. The MAC strips frames by replacing every symbol after the SA field with an Idle symbol (previous fields have already been repeated). This leaves a fragment of SD, FC, DA and SA on the ring. MACs receiving these fragments recognize them as fragments and ignore them. Eventually the fragment arrives at a MAC holding the token and is removed completely, since MACs do not repeat frames received while they are holding the token.

- An *Information (I)* field of from 0 to 4472 bytes. A byte is 2 symbols.

- A *Frame Check Sequence (FCS)* field which is 8 symbols. This field contains a 32 bit cyclic checksum computed over the FC, DA, SA, and INFO fields. A MAC generates the FCS for every frame it transmits and checks the FCS of every frame it repeats.

- An *Ending Delimiter (ED)* field consisting of a special T symbol. The ED marks the end of the frame.

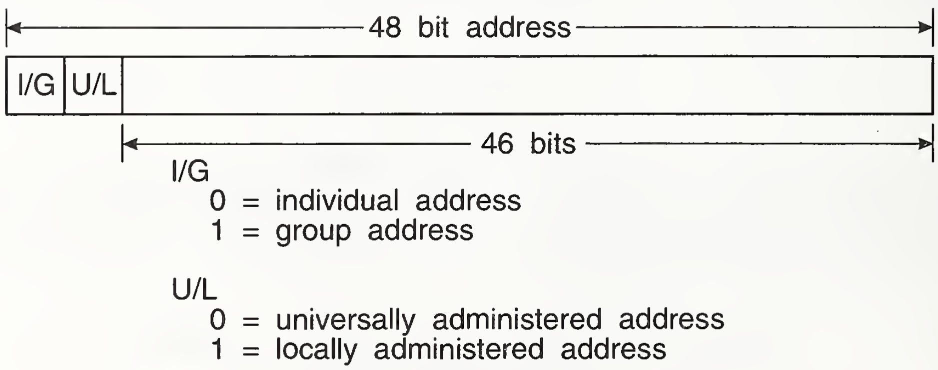

Figure 12 - 48-bit Universal Address Format.

- A *Frame Status (FS)* field consisting of three or more symbols. These symbols are called *trailing indicators* and must be either an S symbol, signifying "set" or an R symbol, signifying "reset." A MAC transmits a frame with at least three trailing indicators, all set to R symbols. The first trailing indicator, called *E*, is changed to an S by any MAC detecting an FCS error. The second trailing indicator, called *A*, is changed to an S by any MAC recognizing the Destination Address as its address. The third trailing indicator, called *C*, is changed to an S by any MAC which copies a frame addressed to that MAC into its buffer memory. Additional trailing indicators are allowed, and must be repeated by MACs, but their meaning is not defined in the standard.

The order of bit and byte transmission in FDDI frames is that the most significant (leftmost in Figure 11) bit or byte is transmitted first. The total length of the frame may not exceed 4500 bytes. The frame information field may be from 0 to 4472 bytes long.

2.5.2 MAC Addresses

Both 16 and 48-bit addresses are defined in the standard, and the type of address used is determined in the Frame control field. The 16-bit address, originally included for compatibility with the IEEE 802.5 token ring standard, has become an anachronism and is not ordinarily used. In the 48-bit addresses the most significant bit determines whether the addresses are Individual or Group addresses. Group addresses may be associated with more than one MAC on a ring, however the standard does not specify how group addresses are assigned. The all ones address is the Broadcast address and frames with this destination address are copied by all MACs.

The second most significant bit of 48-bit addresses determines address administration for individual addresses. A zero means that the address is "universally administered," while a one means that the address is "locally administered." Universally administered addresses are assigned by a registration authority (the IEEE) to vendors, and each vendor assigns a unique address to each MAC built by the vendor. Therefore every FDDI MAC comes with a preassigned universally administered address which is unique. It is not only unique for FDDI MACs but is also unique for IEEE 802.3 CSMA-CD and IEEE 802.5 token ring stations as well. Therefore it is possible to use MAC level bridges to combine FDDI, 802.3 and 802.5 LANs into a single extended LAN with the assurance that every MAC has a universally administered address which is not duplicated elsewhere in the extended LAN.

2.5.3 The Timed Token Protocol

FDDI MAC implements a timed token ring protocol. The token is a special FDDI frame. An FDDI MAC may transmit frames only while it holds the token. A MAC may transmit more than one frame while it holds the token before releasing the token. The length of time a station may hold the token is determined by:

- A *Target Token Rotation Time (TTRT)* which is the same for all stations and which is determined during ring initialization.
- The *Token Holding Timer (THT)* of the MAC which measures the length of time that the MAC may hold the token while transmitting frames.
- The *Token Rotation Timer (TRT)* of the MAC which measures the time interval between successive arrivals of the token at the MAC.
- The *Late_Ct* which is used to record the occurrence of a "late" token.

FDDI defines two traffic modes, *asynchronous* and *synchronous*. Asynchronous traffic is the normal mode for ordinary computer packet traffic.

When a station receives a token, it loads the THT with the value of TRT and then resets the TRT. If Late_Ct is zero, the MAC may then transmit queued asynchronous frames until either the THT exceeds TTRT or the queue is exhausted. If the Late_Ct is not zero, the MAC does not transmit queued frames and immediately passes the token.

Whenever the TRT exceeds TTRT while the Late_Flag is zero, the Late_Flag is then incremented the TRT is restarted. Whenever the TRT exceeds TTRT and the Late_Ct is not zero, then this indicates a "lost token" and the MAC causes the ring to be reinitialized using the *Claim* process. (see 2.5.4 below).

After a MAC transmits frames it releases the token without waiting for its frames to be returned to it (in contrast to some token ring protocols, such as the 802.5 token ring, [ISO88025, BUX81] which hold the token until the packets the station transmitted are returned to it). A MAC strips its own frames by monitoring the *Source Address (SA)* of each received frame and removing any frames with its SA. The stripping leaves a short fragment circulating which is eliminated whenever it arrives at a MAC holding a token, since a MAC does not repeat any frame or fragment which arrives while the MAC holds the token.

An optional priority scheme for asynchronous traffic is defined in FDDI by using priority timers *(T_Pri_n)* in each station, which are set to some value less than TTRT. The priority timers govern the operation of lower priority message transmission. They are treated in a manner similar to the TRT. Less time is available for the transmission of queued lower priority messages than for packets which may use the full TTRT. Up to 8 priority timers are defined in the FDDI standard, but none are required by the standard. A station without priority timers can send only at the highest priority. Since most higher level protocols do not use this facility, and since some FDDI MACs do not implement priorities, the use of FDDI priorities will probably be limited to special applications rather than general purpose networks and we will not consider FDDI asynchronous priorities further in this report.

Scheduling for asynchronous traffic is round-robin and, at the same priority, the operation of FDDI is generally fair (that is every station has the same opportunity to access the ring). In an n station network, each station is guaranteed that it will receive a token usable for asynchronous traffic within a time which cannot exceed n times TTRT, whatever the load [SEVI, JAIN].

Synchronous mode bandwidth is pre-allocated to stations by a manager. It is allocated in units of time and the sum of all the synchronous bandwidth allocations in the entire ring must not exceed TTRT. A station holding a synchronous bandwidth allocation is allowed to send synchronous traffic for a period of time not exceeding its allocation whenever it captures the token, whether the token is late or not. A station holding a synchronous bandwidth allocation is guaranteed to never have to wait longer than twice TTRT between tokens, and to not wait longer than TTRT on average, whatever the asynchronous load on the network. The smaller the TTRT, the smaller the maximum delay for synchronous traffic.

Synchronous mode is intended for applications such as voice, video and telemetry, which may require a constant continuous bandwidth, in contrast to the more "bursty" nature of most computer data traffic. Consequently, most computer FDDI ports will probably never send synchronous mode packets, unless they are originating voice or video signals. Bandwidth allocated to synchronous traffic but not used by the station to which it is allocated becomes available for

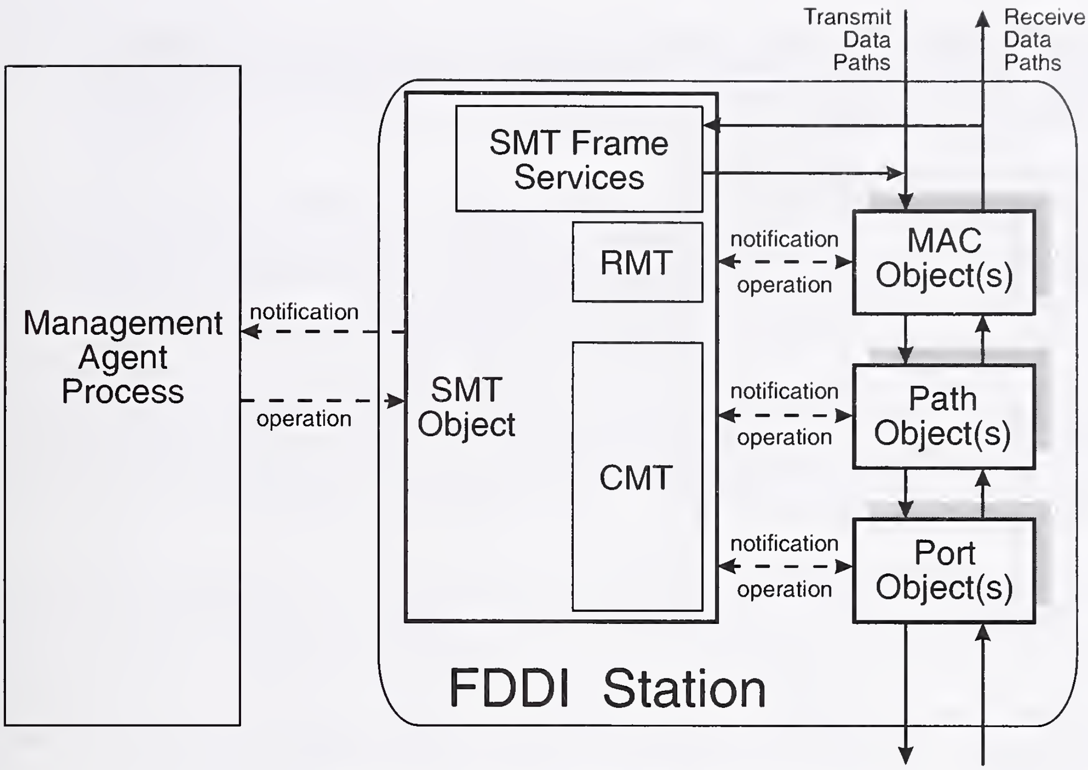

Figure 13 - SMT Management Model.

asynchronous traffic. SMT defines a frame for requesting a synchronous bandwidth allocation from a bandwidth manager, but does not specify how this manager operates (see 2.6.2).

A new FDDI standard, under development, informally called FDDI-II, provides an 8 kHz, synchronized telephone network compatible isochronous service. FDDI-II is discussed in section 3.

2.5.4 Ring Initialization

Initialization of an FDDI ring uses two processes, *Beacon* and *Claim*. Whenever SMT senses that a new connection exists with another station, it directs MAC to Beacon. MAC begins transmitting a special frame defined as a Beacon. Every MAC receiving a Beacon from upstream defers to the upstream Beacon and repeats it. When a MAC receives its own Beacon it then knows that the ring is complete and begins the Claim process.

Each MAC has a parameter, *T_Bid*, which is the value which that MAC will "bid" for TTRT in its Claim frames. A MAC receiving a Claim frame from upstream compares the received bid with its own T_Bid value. If its own T_Bid is a smaller time, then the MAC strips the Claim and sends its own Claim, otherwise it repeats the Claim received from upstream. Ties are resolved by comparing MAC addresses. When a MAC receives its own Claim frame, it has "won" Claim and issues a token. On the first rotation of the token every MAC sets its TTRT to the value bid in the winning Claim. On the second rotation the ring is said to be operational and MACs begin transmitting frames.

When the ring is operational the token may become damaged or "lost." MACs detect late tokens or long periods of ring inactivity and return to Claim to restart the ring. If Claim does not complete within a specified period (*T_Max*), then a MAC will enter Beacon to determine if a closed ring exists. If Beacon completes, then Claim is re-entered. Whenever the ring configuration changes (typically a station enters or leaves the ring or changes its own internal configuration), SMT recognizes this and forces a Beacon.

2.6 Station Management (SMT)

Figure 13 depicts FDDI SMT management model. In every FDDI node there is one SMT entity, which manages overall operation of the FDDI node. The principle parts of SMT are SMT Frame Services, Ring Management (RMT) and Connection Management (CMT). The Managed Object object classes managed by SMT are SMT objects, MAC objects, Path objects and Port objects. Each object includes a number of attributes which may be managed. A Port consists of a PHY entity and a PMD entity which together form one end of a physical connection. A Path represents a a segment of a ring that passes though the station. Paths and Ports are illustrated in Figure 3 a) and the elements of a path are illustrated in more detail in Figure 9. There may be multiple instances of the MAC, Port and Path objects in a station, but there is precisely one SMT object in a station.

The Management Agent Process itself is not defined in SMT. It may simply be a process that allows a station operator to manage the FDDI station from his keyboard and display, or it may be an agent for remote management of FDDI using a remote management protocol (see 2.6.3). SMT provides for the management of objects by Notifications and Operations. Notifications are unsolicited reports generated by a managed object. Operations are functions performed by a Management Agent Process on a managed Object. Get and Replace operations are defined in the standard

2.6.1 Management Information Base (MIB)

The *Management Information Base (MIB)* is defined in SMT clause 6. The definition uses a format defined by the *Guidelines for the Definition of Managed Objects* [GDMO]. Within a single node there may be multiple instances of MAC, PATH, and PORT objects.

Attributes, Behaviors, Notifications and Actions are specified for each object class in SMT clause 6 and given a formal name. MAC attributes frequently correspond to parameters defined in the MAC standard For example, the MAC standard defines a parameter T_Req. This is the value that a MAC bids in a Claim frame and the smallest value of bid during the Claim process determines the parameter T_Neg, which is the value a station understands to be the negotiated current TTRT. The MAC parameter T_Neg corresponds to the MIB attribute *fddiMACT-Neg*, while the MAC parameter T_Req corresponds to the MIB attribute *fddiMACT-Req*

Literally hundreds of MIB attributes are defined in SMT clause 6, some of which, like the examples given above correspond to parameters or counters defined in MAC or the CMT and RMT clauses of SMT. A detailed enumeration of the MIB attributes is beyond the scope of this summary.

In general, the attributes for SMT, MAC and PORT objects allow the monitoring and control of SMT, MAC, PHY and PMD objects. The PATH and ATTACHMENT attributes allow a remote manager to the configuration of nodes and to determine a detailed physical map of an FDDI network.

2.6.2 Frame Management Services

The frame management services of SMT are defined in SMT clause 7, which defines the SMT management frames and clause 8, which specifies several frame-based management protocols. Three basic types of frames are defined:

- Announcement frames, which announce some condition for the use of other stations, particularly a network manager. For example an SMT detecting a suspicious connection periodically announces that fact, and every MAC periodically announces its address.
- Request frames, which are used to request an SMT entity to take some action, for example to get or change a parameter or counter value, or to echo a test frame.
- Response frames, which are used in response to a request frame, for example to report a parameter value as a result of a request frame.

Several SMT Frame Classes are defined:

- *Neighbor Information Frames (NIF)* are used with with the Neighbor Notification protocol to allow a station to determine the MAC addresses of its upstream and downstream neighbors and to detect duplicate addresses. They also provide a periodic frame handshake that verifies the operation of the local MAC transmit and receive paths, in the absence of other traffic. Each MAC periodically (the default period is every 30 s) issues NIF Request Frames to its neighbor and responds with NSF Response frames to its upstream neighbor. The frames contain the station type and the address of its upstream neighbor.

  A special Frame Control field value, called *Next Station Address (NSA)* is Recognized by MACs. When a MAC receives a frame with an NSA Frame Control field addressed to the Broadcast Address with the A trailing indicator reset it knows that it has received a frame from its immediate upstream neighbor; repeats the frame with the A indicator set and reports its source address to SMT as the address of the upstream neighbor. A remote manager can also use NSF Request frames addressed to specific MAC addresses to build a logical ring map.

  Implementation of NIF Request and Response frames is mandatory. An optional NIF Announcement Frame is defined for compatibility with early draft versions of SMT.
- *Status Information Frames (SIF)* are used to request and provide a station's configuration and operating information. SIF Configuration response frames report on the station's configuration, the paths through it, its neighbors MAC addresses, and the version of SMT supported; the information supplied can be used to build both logical and physical ring maps. SIF Operation response frames reports on operational parameters such as MAC status, MAC frame counters, and PORT Link Error Monitor status; the information can be used for fault isolation and performance monitoring. Stations are required to receive SIF request frames and generate response frames.
- *Echo Frames (ECF)* are used to generate an echo response from a station for test purposes. Stations are required to respond to an ECF request frame as an ECF response frame.
- *Resource Allocation Frames (RAF)* are used to allocate resources. The only resource which currently is allocated by these frames in FDDI is synchronous bandwidth. They provide a means for a station to request a synchronous bandwidth allocation from a bandwidth management process. The process for managing synchronous bandwidth is not specified, however SMT clause 8.7.3 does provide guidelines for a such a process. Support of the RAF request and response frames is optional.

– *Request Denied Frames (RDF)* are response frames used to indicate that a request is denied. It is expected that new revisions of SMT will add features which are not supported in the current version, and that both new and old versions of SMT will exist in the same ring. When a station receives a request for a protocol or function it does not support it responds with an RDF frame. Support for RDF frames is mandatory.

– *Status Report Frames (SRF)* are announcement frames transmitted to the SRF Multicast Address that are used to report certain conditions or events that may be of interest to a manager. Examples include undesirable or illegal connections, neighbor changes, configuration changes and excessive error rates. SRF frames are mandatory.

– *Parameter Management Frames (PMF)* are used to provide remote management of station attributes. Support for Get PMF frames is mandatory (that is stations must receive Get PMF request frames and must respond with Get PMF response frames), but support of other PMF Set Response frames is optional (that is stations are not required to accept and respond to PMF Set Request frames).

– *Extended Service Frames (ESF)* are provided to support optional vendor enhancements to SMT. Their use is vendor specific and specific ESF frames are identified by a unique 48-bit number obtained from the IEEE universal address space.

A protocol is defined in SMT clause 8.2 for Neighbor Notification, using NIF frames. A Status Reporting protocol using SRF frames is defined in SMT clause 8.3. This clause specifies the conditions which may result in or require a status report. A Parameter Management protocol is defined in SMT clause 8.4.

Implementation of the Status Reporting and Parameter Management Protocols may be a matter of some concern. The earliest drafts of SMT did not include status reporting and some early SMT implementations did not provide it. Moreover many of the working drafts of SMT carried a “grandfather clause” which stated that implementations designed prior to the adoption of the standard would not be deemed nonconforment if they did not implement the SIF frames and protocol. While that grandfather clause has been removed from the current SMT document, there is a possibility that some older FDDI implementations may not implement this feature. It may be advisable to specifically require the implementation of SIF frames in FDDI procurements for this reason.

2.6.3 Remote Management

Figure 13 shows a Management Agent Process that controls the management of the station. This manager may be a agent for a remote manager using a remote management protocol such as the OSI *Common Management Information Protocol (CMIP)* [ISO9596] which is an application at the top of an OSI stack; in this model remote management then takes place through the CMIP protocol. However, at the present time, few FDDI products are available which offer CMIP management. The most commonly available network management protocol in FDDI equipment, particularly bridges, routers and concentrators, is the *Simple Network Management Protocol (SNMP)* [DAVI].

In addition, the PMF frames described in 2.6.2 provide a complete set of data link layer-management frames and can be used for remote management of FDDI nodes without requiring a higher level management process. As noted in 2.6.2 above, while implementation of PMF Get Response frames is mandatory in all FDDI stations, implementation of PMF Set response Frames is optional. This means that all FDDI stations must respond to a PMF Get Request, but need not act on PMF SET Requests. Therefore, an FDDI network management application that relies on PMF frames is assured that it can get the value or status of all attributes in every station (except for "MACless concentrators"), but cannot be assured that it can change parameters simply because all stations conform to the FDDI standard. The PMF Get Response and the required periodic transmission of NIF frames allows a remote manager to create a map of the active ring and determine the value of every managed attribute of every station on an FDDI ring.

Users must decide the management approach they favor and act accordingly. Users who desire "full" layer management capabilities should ensure that their FDDI networks are composed entirely of stations that implement the optional Set PMF Response frames. Users who favor SNMP or another higher level or out of band management strategy should ensure that the desired management protocol or mechanism is supported by at least the nodes that require remote management; that would ordinarily include all dual attachment nodes on a trunk dual ring and any single attachment stations which are bridges, routers, gateways or concentrators with many subordinate stations. Note that "MACless" concentrators cannot be remotely managed at all, unless it is by some "out of band" mechanism not defined in the standard. See 7.3 for information about the *Government Network Management Profile (GNMP).*

2.6.4 Connection Management (CMT)

Connection Management (CMT) is responsible for managing FDDI connections and the configuration of stations. When an FDDI port is turned on it first performs a self test called *Path Test.* The standard does not define what is done in Path Test, which depends on the design of the station, but a complete self test can take a relatively long period of time (seconds to minutes). CMT then starts a basic signaling protocol through which the CMT entities in two connected nodes determine the remote port type (A, B, S or M) and performs a *Link Confidence Test*, to provide assurance that the link quality is adequate to include the connection in an FDDI ring. Since the test is relatively brief (a few seconds, while a test of nearly about 10 min would be needed to provide a 90% confidence that the link error rate is less than 2.5×10^{-10}), the assurance is limited, but the Link Confidence Test does preclude inserting a very bad link into a network. Finally, if the port types are compatible (only M-to-M is forbidden under the default connection policy) the link is inserted into a ring. The CMT protocol uses a robust line state oriented mechanism which operates even with very high error rates.

While a connection is active the CMT Link *Error Monitor (LEM)* function monitors errors on the incoming link, by noting code violations and code values which are invalid in the current line state (for example, the only valid symbols which can follow an I symbol are a J symbol or another I symbol). From this the LEM maintains a running estimate of the current link error rate. The current link error rate estimate may be interrogated by a remote manager, and if the estimate exceeds a preset threshold (the default value is a BER estimate of 10^{-8}) an SRF frame transmitted to announce a link error rate alarm. If the estimate exceeds a further threshold (the default value is 10^{-7}) the link is disabled automatically.

CMT provides *Trace*, a special line state sequence which propagates upstream until a node with a MAC is encountered and then back downstream to the initiator. In some cases Trace may pass through several nodes before coming to the first upstream node with a MAC in its path. All connections involved in the Trace are then broken and every path goes to the station's internal Path Test. If the Path Tests complete successfully, connections are re-established by CMT. Path Test may however detect a failure, allowing the ring to be reconfigured around the fault. Since these internal self tests may take a long time (seconds or even minutes), Trace is a last resort invoked by RMT (see 2.6.5 below) when a ring cannot be initialized despite apparently good connections.

CMT controls the configuration of stations. For example it controls a bypass switch if one is present in dual attachment stations and concentrators. It causes dual attachment stations to wrap or go to through depending on the state of the connections on the A and B ports. It controls the M ports in concentrators and configures them in paths when a connection is established.

2.6.5 Ring Management (RMT)

Ring Management (RMT) monitors the operation of the ring. There is an instance or RMT for every MAC in a station. RMT detects and resolves duplicate MAC addresses, which may cause several problems in FDDI networks. It also causes the initiation of the *Trace* process when the logical ring is apparently broken immediately upstream of the MAC.

While the ring is operational duplicate MAC addresses cause problems because each MAC will strip frames sent by the other. This may mean that the frames are not received by their destination. The NIF request and response frames provide a means for duplicate address detection when the ring is operational.

Duplicate addresses may cause more serious problems during ring initialization (Beacon and Claim) and can, in some cases, prevent initialization from completing. This usually occurs when one of the duplicate address MACs is bidding the winning value of T_Req, which the other strips. When RMT detects that its MAC is a duplicate that prevents ring initialization it takes one of the following actions:

- RMT may remove the MAC from the FDDI ring;
- RMT may configure the MAC to lose the claim process; this allows ring initialization to proceed;
- RMT may change the MAC address to its unique universal address. Universal addresses should never be duplicates.

When a MAC Beacons for a long period of time but receives no Beacon (and the upstream connection is apparently still active to CMT), then the logical ring is apparently broken, probably due to a fault in the beaconing station or its upstream neighbor. Since CMT still indicates an active connection, the physical link is apparently not broken, but Beacon is not getting through. RMT then causes a special *Directed Beacon* to be sent, with the destination address value of its upstream neighbor. The primary purpose of the Directed Beacon is to announce where RMT suspects the problem lies. RMT then causes CMT to initiate the Trace function.

3. Overview of FDDI-II

FDDI-II adds an 8 kHz synchronized *isochronous* service to FDDI while it maintains the media defined for FDDI and the 100 Mbps data rate, as well as the token ring packet service. The word isochronous means having equal or uniform intervals of time. Digital telephone and video services are conventionally provided in an isochronous manner. The FDDI-II Isochronous service is a *Time Division Multiplexing (TDM)* circuit switched service compatible with North American and European telephone trunks services, such as the *Integrated Services Digital Network* (ISDN) B channel (64 kbps), the ISDN primary rate channel (1.536 Mbps), switched 56 kbps services, the DS-1 (1.536 Mbps), DS-2 (6.144 Mbps), and DS-3 (43.008 Mbps) services. Almost any data channel or carrier, up to the DS-3 rate, that can pass through the time division multiplexed (TDM) wide area public network maintained by the long distance and local telephone companies can also be carried through FDDI-II.

FDDI-II can be integrated with existing wide area network channels and a signal can originate on an FDDI-II node, then pass through the public network trunks to any normal terminal which is attached to the public network, at any rate up to DS-3. FDDI-II plays naturally with current video technology to facilitate multimedia applications. FDDI-II provides an integrated service Local Area Network which is well suited for video multimedia applications, using current equipment and technologies. B-ISDN and ATM LANs (see 9.5) intend to provide similar capabilities by means of ATM cell switching technology, however it may not be as easy to adopt cell switching to current video devices that expect a strictly isochronous channel. FDDI-II stations are backwards compatible with FDDI rings, that is they can be mixed with FDDI stations, however there can be no isochronous services in such a mixed ring.

3.1 Structure of the FDDI-II Standards

Figure 14 illustrates the architecture of FDDI-II. The various PMD alternatives are retained unchanged from FDDI. Four new standards are being created for FDDI-II.

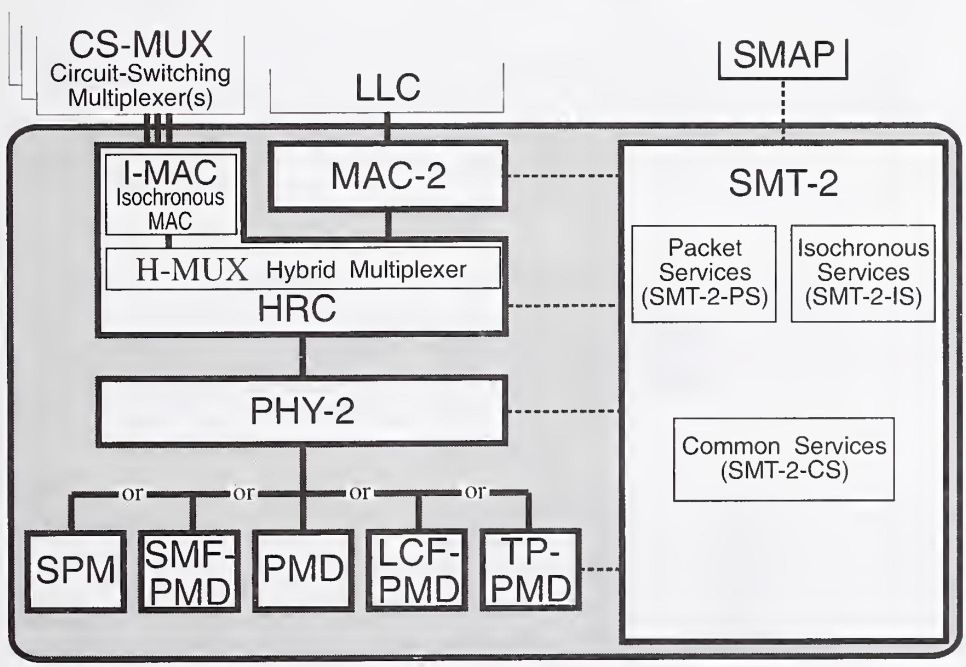

Figure 14 -FDDI-II Architecture.

3.1.1 Hybrid Ring Control

Hybrid Ring Control (HRC) [HRC] is the only entirely new sublayer in FDDI-II and it implements most of the new functionality in FDDI-II. The basic function of HRC is to create a cycle structure, illustrated in Figure 15, that allows the 100 Mbps data stream to be divided between the token ring packet service and the isochronous service. A *Cycle* has a duration of 125 μs for 8 kHz compatibility with the public network TDM carrriers. Each cycle includes a *Preamble* of five code symbols, a *Cycle Header* of 24 code symbols, a *Dedicated Packet Group* of 24 symbols and 96 *Cyclic Groups* of 32 symbols each. The symbols in the Dedicated Packet Group are always reserved for the packet service, providing a minimum packet service of .768 Mbps. This minimum packet service is needed for the management of the network even when all other available bandwidth is allocated to the isochronous service.

Each of the 16 bytes in each of the Cyclic Groups is allocated to one of 16 *Wide Band Channels (WBC)*. With 8,000 Cycles per second, 96 Cyclic Groups per second, and 8 bits per byte, the data rate of a WBC is 6.144 MHz. This is the DS-2 data rate, is four times the DS-1 or North American ISDN Primary (H11) payload rate, and three times the European CEPT1 or ISDN Primary (H12) carrier rate. Seven WBCs together equal the North American DS-3 (H32) payload rate of 43.008 Mbps. Thus the FDDI-II WBC structure maps conveniently into either the North American or European telephone trunk hierarchy.

Each of the 16 WBCs may be allocated to either the packet service or the isochronous service. For example, if seven WBCs were assigned to the isochronous service, then its total rate would be 43.008 Mbps, while the remaining nine WBCs plus the Dedicated Packet Group would provide an FDDI token ring packet service with a rate of 56.064 Mbps. The remaining .828 Mbps are lost to the cycle overhead.

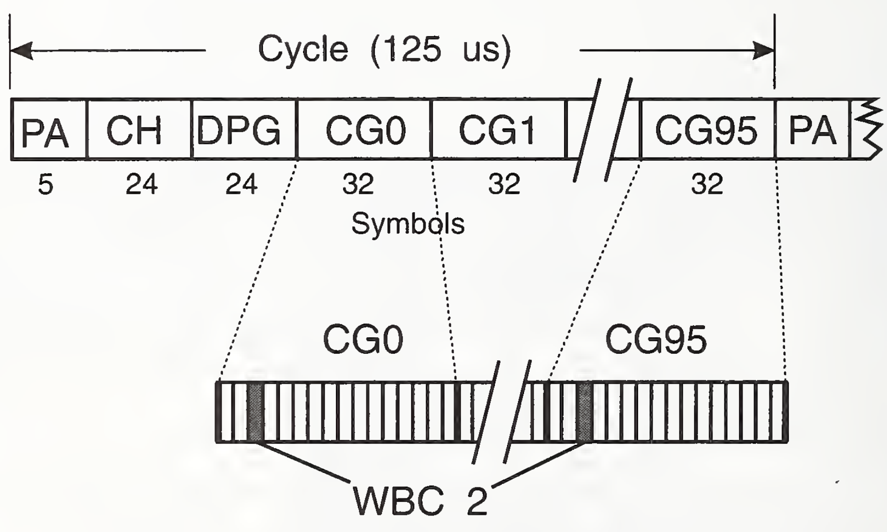

PA: Preamble
CH: Cycle Header
DPG: Dedicated Packet Group
CGx: Cyclic Group x (0 to 95)
WBC: Wide Band Channel (0 to 15)

Figure 15 - FDDI-II Cycle Structure.

Table 4 - FDDI-II Symbol Set

Binary Value	Notation	Assignment	Binary Value	Notation	Assignment
00000	Q	Quiet	10000	V	Violation
00001	V	Violation	10001	K	Start delimiter 2
00010	V	Violation	10010	8	Data
00011	V	Violation	10011	9	Data
00100	H	Halt	10100	2	Data
00101	L	Embedded Delimiter	10101	3	Data
			10110	A	Data
00110	V	Violation	10111	B	Data
00111	R	Reset	11000	J	Start delimiter 1
01000	V	Violation	11001	S	Set
01001	1	Data	11010	C	Data
01010	4	Data	11011	D	Data
01011	5	Data	11100	E	Data
01100	V	Violation	11101	F	Data
01101	T	Ending Delimiter	11110	0	Data
01110	6	Data	11111	I	Idle
01111	7	Data			

In FDDI-II there are two types of stations: *Monitor* and *Non-Monitor*. In a Monitor station HRC contains a Latency Adjustment Buffer, which is used to ensure that the delay around the ring is an exact multiple of 125 us, so that there are a fixed number whole cycles circulating around the ring. Non-Monitor stations do not have a Latency Adjustment Buffer.

When an FDDI-II ring starts up, it starts in FDDI or *basic* mode. In this mode the cycles do not exist and the protocol is compatible with FDDI. In fact FDDI and FDDI-II stations can be mixed on the same ring, but such a mixed ring can never leave the basic mode, that is it remains an FDDI packet only ring. In a ring of all FDDI-II stations, a single *Cycle Master* station is chosen from among the Monitor stations by a protocol defined in SMT-2. After the Cycle Master is chosen it generates the Programming Template for cycles and inserts its Latency Adjustment Buffer into the ring to cause the ring to enter the hybrid mode.

3.1.2 PHY-2

PHY-2 [PHY2] is an extended version of FDDI PHY. Provision is made in PHY-2 to frame *cycles* which are described below and new symbol, L is defined. The FDDI-II symbol set is shown in table 4. PHY-2 contains additional provisions for smoothing which are required for the isochronous service. PHY-2 is backward compatible with PHY and can replace PHY even in FDDI stations that do not implement the FDDI-II.

3.1.3 MAC-2

MAC-2 [MAC2] is a slightly extended version of FDDI MAC. It clarifies several open issues in MAC, such as the setting of the A & C bits with bridges and stripping for bridges (see 8.2). The MAC-2 frame format is illustrated in Figure 16. It is the same as the MAC format except for the inclusion of the Routing Information (RI) field, which is present if the first bit of a 48-bit source address field is set to one. If present, the RI field contains routing information for source routing bridges and the first octet of the field defines its length.

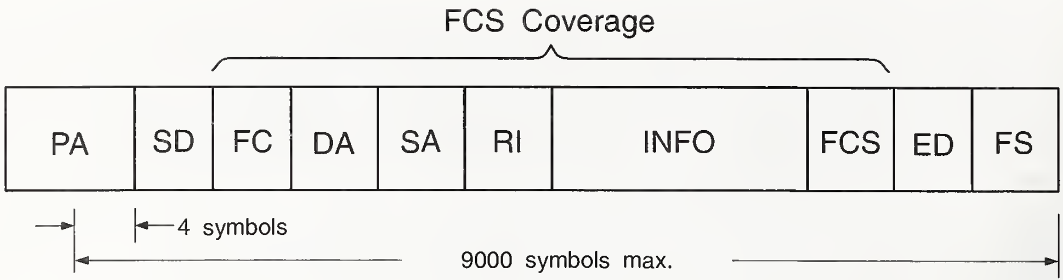

PA: Preamble (16 or more symbols)
SD: Starting Delimiter (2 symbols)
FC: Frame Control (2 symbols)
DA: Destination Address (12 symbols)
SA: Source Adress (12 symbols)
RI: optional Routing Information (4 to 30 symbols)
INFO: Information (0 or more symbol pairs)
FCS: Frame Check Sequence (8 symbols)
ED: Ending Delimiter (1 symbol)
FS: Frame Status (3 or more symbols)

One FDDI symbol is 4 data bits and 5 code bits

Figure 16 - MAC-2 Frame Format.

MAC-2, in effect, cleans up several problems and limitations with MAC several of which were noted as footnotes in the MAC standard, because the decision to change was made just after the approval of MAC. For example, MAC makes 16-bit addresses mandatory and 48-bit addresses optional, while MAC-2 makes 48-bit addresses mandatory and 16-bit addresses optional. In fact 16-bit addresses are nearly unused in nearly all FDDI stations, whether they implement MAC-2 or MAC. MAC-2 also strengthens the criteria for a valid ending delimiter and some MAC implementations use this stronger criteria.

MAC-2 is backward compatible with MAC and can be substituted for MAC in FDDI stations. In effect, most or all FDDI bridges and all stations that participate in source routing bridging implement at least some of the new features defined MAC-2. MAC-2 is required in FDDI-II stations.

3.1.4 SMT-2

SMT-2 will be an enhanced version of FDDI-SMT and will be subdivided into three parts:

- FDDI Station Management Packet Services (SMT-2-PS) [SMT2PS], which manage the packet services for an FDDI-II ring.
- FDDI Station Management Isochronous Services (SMT-2-IS) [SMT2IS], which will manage the isochronous services in an FDDI-II ring.
- FDDI Station Management Common Services (SMT-2-CS) [SMT2CS], which will provide management services common to both the packet and isochronous services.

3.1.5 Standards Status

At the present time MAC-2, PHY-2 and HRC are not yet approved as ISO or ANSI standards, but all are essentially completed and in the standards approval process. While SMT-2 will be able to adopt much from SMT, work is beginning on SMT-2. It therefore seems unlikely that the complete FDDI-II standard will be available before 1994. In particular, SMT-2 must define the protocols for allocating the WBCs between the isochronous and packet data services. Prototype ICs to implement MAC-2 and PHY-2 are available and a few early FDDI-II products are starting to emerge.

3.2 Applications and Planning for FDDI-II

The primary application for FDDI-II is expected to be multimedia computer systems, that is computer systems which integrate computers with voice and motion video. The FDDI-II isochronous service will simplify the use of conventional video and voice technology developed for use over existing TDM wide area networks, and the integration of LAN and TDM network services. A great deal of effort is being expended to develop practical video standards for low rate (56 kbps to 1536 kbps) isochronous wide area network services. Many early stand alone computer multimedia applications are based on CD-ROM storage, which provides a constant rate 704 kbps service. Even if good solutions are possible for packet video, the simplest solution may be to take advantage of these standards by using an isochronous medium such as FDDI-II, which will also facilitate integration with the telephone network, and rapid development of widespread videoconferencing via narrowband ISDN. Fibre Channel (see 9.2) and HIPPI (see 9.1) are other local network technologies that may provide circuit switching services at very high rates, probably through space or time division multiplexing in switches.

However, multimedia systems are still in their infancy, and the pace of the development of the market for multimedia systems is uncertain. The need for truly isochronous services to implement multimedia applications is debatable. Practical motion video for multimedia systems will probably use compression. Video compression relies on redundancy in the transmitted information and the fact that there is usually little change in the image from one frame to the next. However, motion or sudden changes in the image require transmission of large amounts of information when they occur. When a video scene change occurs, or when a video image has large amounts of motion (e.g., a sports telecast) and a constant rate limited bandwidth channel is used to transmit the video, the result is loss of image quality. In this sense, compressed motion video is rather like computer data: traffic is inherently bursty. It follows that there may be advantages to packet oriented video services, which are intended for bursty traffic. These kinds of compressed video may be well suited to ordinary packet LANs, or to a B-ISDN variable bit rate service, and not well suited to an isochronous service.

Eventually, the B-ISDN ATM technology (see 9.5) may provide a superior solution encompassing voice, video, and data in both the local and wide area networks. ATM technology will use 53 byte ATM cells (generally with a payload of either 48 or 44 bytes) to carry all types of data, including voice, data, and video. The *Synchronous Optical Network (SONET)* (see 9.3) standard provides an optical wide area trunking and multiplexing service which may serve as the basis for the delivery of ATM service, as well carrying the existing conventional telecommunications trunk services. SONET may provide a means of delivering *High Definition TV (HDTV)* services.

The FDDI-II standard does not cover isochronous multiplexors. It is possible that an FDDI-II backbone may not ordinarily be extended all the way to workstations, rather FDDI-II hubs may evolve which provide both an 802.3 packet service and a primary rate ISDN service to the workstation over twisted pair from the FDDI-II backbone. A single twisted pair 10-Base-T 802.3 link can be included in a workstation today for a very small cost, and similar technology can accommodate the ISDN primary rate. Very good quality video is now available over the primary rate, and existing public network switches are being upgraded to provide circuit switched primary rate services.

It is too soon and there are too many uncertainties to predict how multimedia will evolve and whether FDDI-II will be a key ingredient in its success. FDDI-II may be destined to simply be a curiosity, a technology in search of an application that never develops, or develops differently than the FDDI-II designers anticipate. Multimedia applications may prove so difficult to develop that B-ISDN will arrive before a large multimedia market develops. Business may not find a compelling need for the multimedia applications that are developed. However, if multimedia develops rapidly and stand alone CD-ROM multimedia applications and video conferencing are rapidly extended to LANs, then it may well be that FDDI-II provides the key enabling communications medium to facilitate the extension of multimedia from isolated systems to LANs and distributed multi-user systems.

Accordingly, it is difficult to provide guidance about planning for FDDI-II. It appears wise to wait for the development of specific multimedia applications and products using FDDI-II, before making any commitment to implement FDDI-II, and to make a decision to use FDDI-II in light of a specific intended application. In the future it may be possible to buy FDDI-II compatible packet only equipment (that is stations that use MAC-2, and PHY-2 and include HRC, but provide no isochronous services) and use it on an FDDI network. Since any FDDI station on a ring prevents the transition to hybrid mode, if the extra cost is small, it may be prudent to buy FDDI-II compatible packet only equipment when it becomes available, so that such equipment can eventually be included in an FDDI-II network. The extra cost of the FDDI-II H-mux need not be large, so vendors may find it prudent to simply offer all products as FDDI-II compatible, as soon as FDDI-II achieves some reasonable market share.

There are now commercial efforts to develop low-cost multimedia interfaces for PCs and workstations. These efforts typically combine an 802.3 equivalent service and an isochronous channel on twisted pair media. FDDI-II then would provide a backbone linking multimedia hubs, which would communicate with PCs and workstations over the low-cost twisted pair interface. This approach is similar in concept to the now widely used strategy of using FDDI as a backbone connecting 802.3 bridges and hubs.

The wiring infrastructure for FDDI and FDDI-II is identical; there is no need to change the wiring method for FDDI-II. However, there is a strong possibility that FDDI and FDDI-II may coexist in the same cable plant. Therefore it may be prudent allow for extra trunk fiber pairs for FDDI-II when installing installing a fiber cable plant for FDDI. If only a single fiber pair is pulled from a building wire center to a wire closet, and both FDDI-II and FDDI workstations are supported from that closet, then the fiber link into the closet must be FDDI-II and an FDDI to FDDI-II bridge must be provided in the closet. If two fiber pairs are pulled between each closet and the wire center, then there is flexibility to locate the bridging function in the wire center. There generally is little need to install extra horizontal wiring to the desktop for FDDI-II, because an individual workstation will generally need FDDI or FDDI-II, but not both at the same time.

Two pairs between each wiring closet and building wire center are sufficient to support a single dual ring, or two separate master/slave links, one FDDI and the other FDDI-II. If parallel FDDI and FDDI-II dual rings are contemplated, then four fiber pairs are needed. When several building wire centers are tied to a single campus or facility wire center, then four fiber pairs will allow for parallel FDDI and FDDI-II backbones. If the buildings in a campus or facility are ring wired, rather than star wired back to the wire center, then two pair on each link are sufficient to support both an FDDI and an FDDI-II dual ring.

Careful partitioning of a large FDDI network may facilitate a later transition to FDDI-II. If a large FDDI network is organized into a dual ring backbone with a small number of concentrators and bridges, and a number of cleanly partitioned FDDI trees, or bridged FDDI subnetworks, then it will be relatively simple to convert the backbone to FDDI-II and connect the FDDI trees to the FDDI-II backbone through a bridge rather than a concentrator. A well organized dual ring of trees, with a dual ring backbone limited to concentrators, and perhaps bridges and major servers, is a good idea in any large FDDI network, and will make for a more robust and manageable network even if there is never to be a transition to FDDI-II.

Network organization and partitioning will be important in FDDI-II, even if there is no transition from FDDI to manage. In an FDDI or an FDDI-II ring, a network reconfiguration (for example to include a station in the ring which had previously been switched off) causes the ring to be reinitialized, and may cause the loss of packets. The time needed to reinitialize the ring is generally short (a few ms) and the normal packet protocols quickly recover from the service interruption caused by the reconfiguration. Higher layer protocols and users will generally not be aware that the ring has been reconfigured.

However, an interruption of an isochronous service is likely to be much more apparent, since the essential nature of the service is its unvarying delivery of data. A process using an isochronous service expects a certain quantity of data every 125 μs. The consequence of an FDDI-II reconfiguration might, for example, be the loss of several video image frames, which might be quite apparent to a viewer. The service interruption might also mean the loss of some sort of subchannel synchronization, or the loss of important data, since some isochronous applications (e.g., telemetry) are real time applications which have no way to back up and recover lost data.

Therefore attention to FDDI-II network organization and partitioning is needed to maintain an acceptably low rate of network reconfiguration. It is likely that some sort of bridge will be used for FDDI-II in cases where a concentrator would suffice for FDDI. The bridges will break a large FDDI-II LAN into several smaller rings, so that reconfigurations can be limited to one small subnetwork rather than affect a larger network. It also appears prudent to design FDDI-II stations so that power is maintained in the FDDI-II port when the station is turned off, to avoid a reconfiguration every time the station is turned on or off.

4. FDDI Network Configuration Considerations.

In this section we will discuss some general considerations which apply to the configuration of FDDI rings. These include:

- How the size of the network and the number of stations and the *Target Token Rotation Time* (*TTRT*), largely determine the network performance characteristics;
- The limits on the maximum possible network size;
- Partitioning FDDI networks and determining the desirable ring size.

4.1 TTRT and FDDI Performance

The *TTRT* is the key parameter parameter affecting FDDI ring performance which a network manager may set. Increasing the ring's size, and therefore the ring latency (the time taken for a frame to circulate around the ring) causes packet queuing delay to increase and lowers the ring efficiency. Increasing the *TTRT* improves the ring efficiency at the expense of increasing the maximum *access delay*, the maximum delay which a MAC may experience before it receives a token which is not late and therefore may be used to transmit asynchronous data. The choice of a *TTRT* value is affected by the ring size and the delay constraints of applications and protocols, and there is a tradeoff between maximum efficiency and maximum access delay.

4.1.1 Analytical Models of FDDI Performance

The behavior of FDDI and other timed token protocols is well understood and analytical models allow determination of maximum network utilization as a function of ring latency, actual token delay and *TTRT* [WAIN, JAIN]. No known, broadly applicable analytical model provides accurate packet delay versus load estimates for heavily loaded FDDI networks [LAMA] but discrete event simulations have been used [JOHN87, DYK87, DYK88, BUX, BURR] to investigate FDDI delay performance under loads approaching the capacity of the network. The *TTRT* is a key parameter in controlling FDDI performance. Jain [JAIN] derives the following equations:

$$(1)\quad E = \frac{n \times (TTRT - L)}{n \times TTRT + L}$$

$$(2)\quad AD_{max} = (n-1) \times TTRT + 2 \times L$$

where E is the efficiency (ratio of bandwidth used to carry packets to channel capacity), n is the number of stations, L is the ring latency and $ADmax$ is the maximum access delay, when the ring is overloaded. For large n, the maximum efficiency is, $Emax$ is given by :

$$(3)\quad E_{max} = 1 - \frac{L}{TTRT}$$

$ADmax$ is linear with n and *TTRT*, but $Emax$ approaches 1 as *TTRT* becomes large. Therefore, as long as TTRT is at least several times L, increasing TTRT improves $Emax$ only a little, but causes a linear increase in $ADmax$.

4.1.2 Setting *TTRT*

Under the Implementation Agreements that will be specified by GOSIP version 3 [GOSIP] for FDDI, the default value of T_Req, the parameter which each MAC bids in the claim process, is T_Max, which is at least 165 ms. Thus if T_Req is not changed in any station, the operational

TTRT will be at least 165 ms. T_Max is the largest value for *TTRT*; it produces the highest possible utilization under very heavy loads and making it the default value of T_Req facilitates the explicit management of *TTRT*, since any one station bidding a smaller T_Req value during the Claim process wins Claim and sets *TTRT*. Therefore, a network manager can set *TTRT* to any value less than T_Max by setting the value of T_Req for any one MAC, while leaving T_Req unchanged in other MACs. However, as Jain points out, with large networks, this can result in large maximum access delays (i.e., that is the time between receiving tokens that are not late and can be used by the station to send traffic).

The behavior of traffic FDDI in a "saturated" FDDI ring, where every station has an asynchronous load which exceeds its share of the ring bandwidth, is simple. Each MAC periodically gets a transmission opportunity and then transmits frames for a period of *TTRT* before its token holding timer expires. That MAC then waits until every other MAC in succession has also had an opportunity to transmit for a period equal to the *TTRT*. The results is fair round robin sharing of the ring bandwidth in *TTRT* size increments. Since token rotation is minimized, and MACs are sending data whenever the token is not rotating, the efficiency of the network is maximized. But the maximum access delay is n times *TTRT*, where n is the number of MACs in the network. For example, with a 50 station network and the default *TTRT* of 165 ms, the interval between transmission opportunities would be 8.25 s and a station could send 16.5 Mbits of data at each opportunity. This gives an average bandwidth per station of 2 Mbps (less the bandwidth lost while the token was actually rotating). Note that the access delay in this saturated condition is nearly independent of the amount of cable in the ring, and depends only upon the *TTRT* and the number of MACs.

Long maximum access delays are unacceptable for some applications and higher layer protocols. The long intervals between transmission opportunities may cause sources to "time-out" and retransmit packets which had in fact been received, making congestion worse. Often protocols could more gracefully handle an overload situation with smaller access delays, even if the total available bandwidth were less.

Jain recommends setting *TTRT* to 8 ms. With a 50 MAC network, a *TTRT* of 8 ms limits the maximum delay between transmission opportunities to less than 0.4 s. Figure 17, which is based on a discrete event simulation, shows the delay versus load characteristics of a 50 station FDDI single ring for 1000 byte frames, at lengths of 10, 50, 100, and 200 km, when all packets have 1000 byte information fields. The load is uniformly distributed with exponential packet arrivals. Even with a 200 km ring, the maximum load is still more than 75 Mbps and the maximum efficiency is more than 75% with a *TTRT* of 8 ms. Figure 18 shows the effect of various TTRT settings on delay versus load in a moderately large (50 stations and 50 km of fiber) FDDI ring, with similar loads. We see here that there is almost no difference in delay between a TTRT value of 8 ms and one of 165 ms, until the load exceeds 90 Mbps and the maximum capacity of the ring (in Information field bps) increases from 93.8 to 97.2 Mbps. We can see from this that there is very little performance advantage to a *TTRT* greater than 8 ms, until an FDDI ring becomes quite large. In most cases a *TTRT* of 8 ms will provide a good compromise between maximizing network capacity and limiting the access delay in heavily loaded networks.

While the 8 ms rule of thumb may serve for most cases, larger *TTRT* values may be required for very large rings. Smaller *TTRT* values may be required for rings with very delay sensitive applications, particularly if priority timers are to be used (see [DYK88] for a full discussion of the characteristics of FDDI priorities). If the maximum delay that can be tolerated between

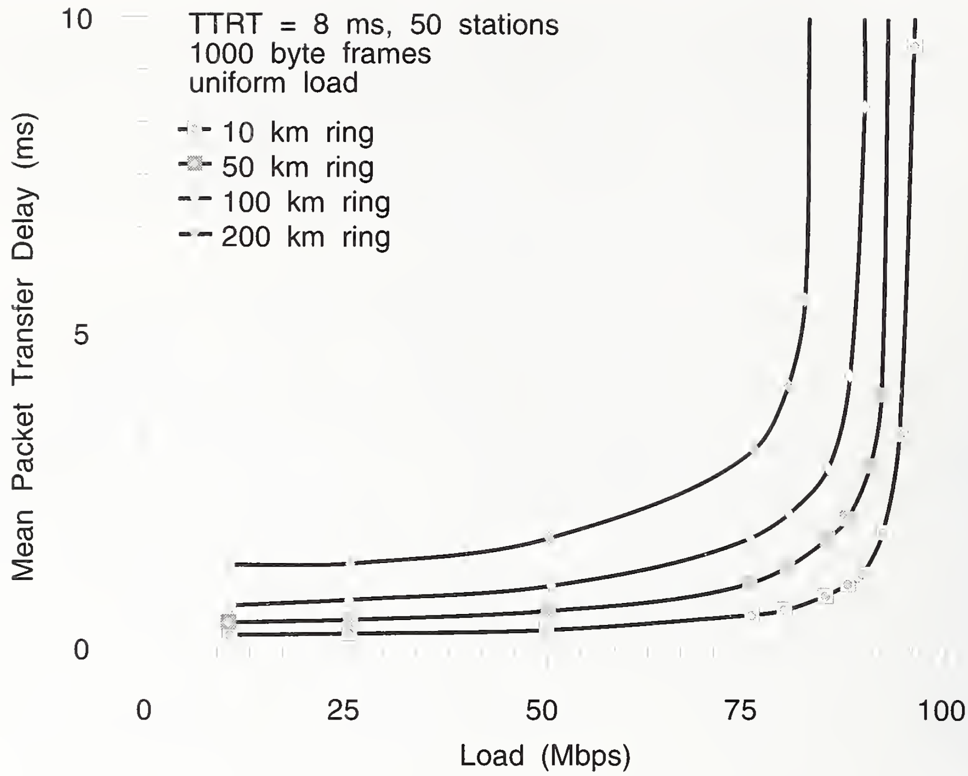

Figure 17 - Effect of Ring Length on FDDI Delay.

usable tokens is known, then *TTRT* should be set to a value not more than that maximum access delay divided by the number of stations in the ring.

4.1.3 Configuration Limits

Worst case latency is important because if it exceeds a maximum value, D_Max, ring initialization may not complete. FDDI is promoted as supporting networks of up to 1000 ports and 100 km (200 km in an unwrapped dual ring) of cable. Every FDDI standard contains the following sentence: "Default values for FDDI were calculated on the basis of 1000 physical connections and a total fiber path length of 200 km," or a very similar statement, in its Scope clause. This is a true but somewhat misleading statement, and the actual configuration limits are often more constraining.

The first reason that this is misleading is that one tends to read ports as stations or nodes. But every FDDI node potentially requires two port delays. A single attachment station is connected to an M port in a concentrator as well as its own S port. And whenever a dual ring is wrapped the ring passes twice through every dual attachment node. Finally, the light travels two ways on a master to slave link and also travels two ways on a wrapped dual ring. Therefore a less confusing statement of the FDDI configuration limit, which is sometimes given is: "500 nodes and 100 km of dual fiber cable."

This too, as a practical matter, is an overstatement of the actual FDDI configuration limits. Some delay is required for an incoming symbol to traverse a path through a node and be repeated. Because there is a path for every port, and ports are externally visible, we will refer to this as the port delay. PHY and PHY-2 specify maximum allowed port delays. The confusion arises in part because the original PHY standard specified a maximum port delay of 756 ns. Other default timer values were calculated from this to allow for 1000 ports and 200 km of fiber

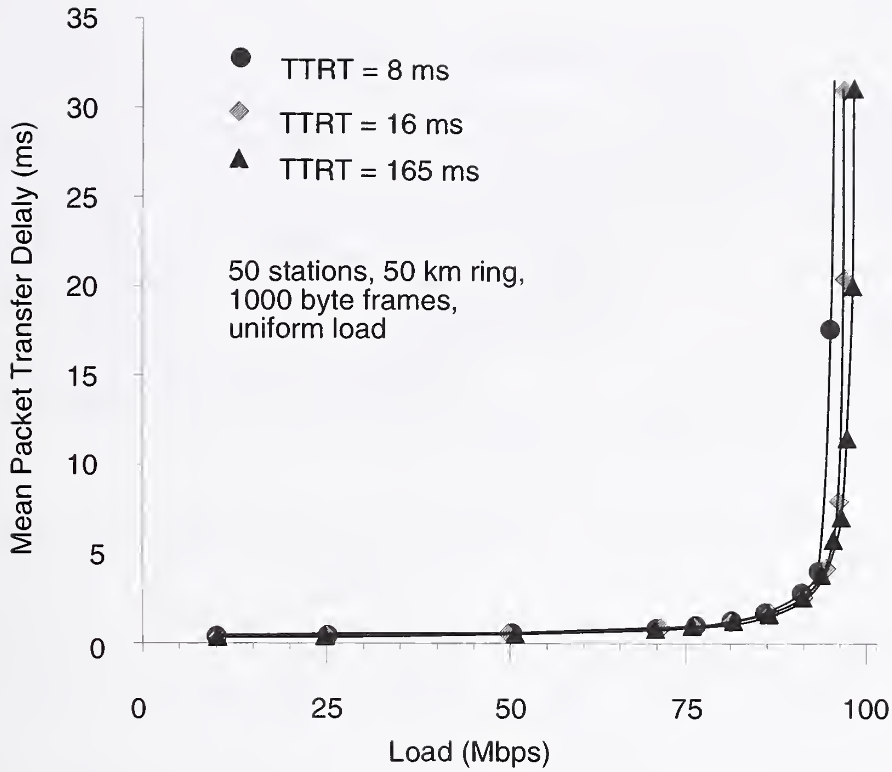

Figure 18 - Effect of TTRT on FDDI Delay.

propagation delay. However PHY was subsequently modified in ways that slightly increase the theoretical minimum delay per port and real implementations have proved to have somewhat greater delays than is theoretically necessary, but default times in MAC and SMT were not changed to allow for this.

The OSI Implementors' Workshop Stable Agreements for FDDI [STABLE], which are referenced by GOSIP [GOSIP], do allow for this increase in the port delay and PHY-2 gives maximum delay values consistent with the implementation agreements. SMT states default values for T_Max and TVX which are consistent with 500 nodes and 100 km of dual fiber fiber cable only if the port delay is 756 ns. It is likely that actual FDDI products have port delays which exceed the 756 ns Figure of PHY, and default timer values will not necessarily support network configurations as large as these statements indicate.

We can calculate the worst case network delay (the time it takes for a frame Starting Delimiter to circulate completely around the ring) from the length of each type of cable in the ring and the number and type of ports in the network. We assume that the signal propagates each way in the cable, which is always true in a M-to-S connection and is true in an A-to-B connection when the ring is wrapped. The signal propagation delay in optical fiber cable is approximately 5.08 µs per km, so allowing for propagation in each direction gives a total delay of 10.16 µs per km. Signal propagation in twisted pair media is slower than for fiber. We will use a typical propagation rate of 6.06 for twisted pair, yielding a round trip delay of 12.12 µs per km.

Each network port also adds delay. This port delay depends on whether HRC is present or not in the path, and, for FDDI-II networks, whether the network is in basic mode or hybrid mode. We take the following maxim per port delays from Annex A of the draft PHY-2 standard:

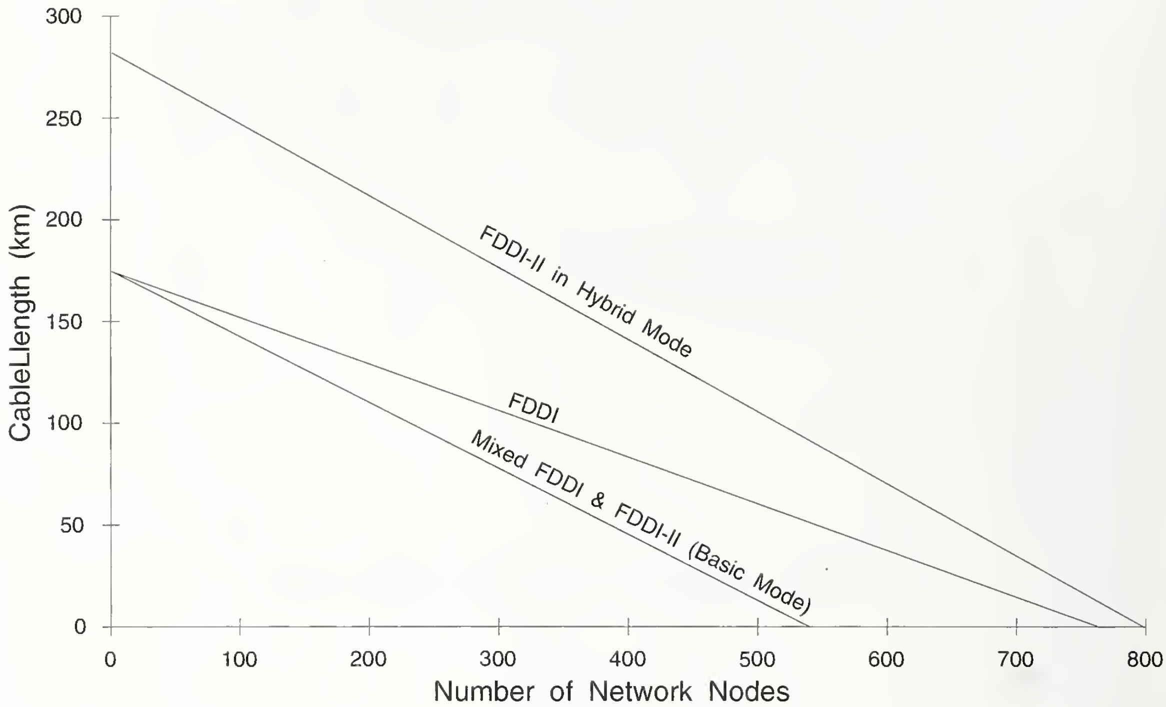

Figure 19 - Maximum Configuration with Default T_Max.

- The maximum delay for FDDI ports without HRC in the path (this will be all FDDI ports except for FDDI-II ports) is 1.164 μs. If there are any ports without HRC in the path in a ring then it can only operate in basic mode.
- The maximum delay for FDDI ports with HRC in the path (that is FDDI-II ports) operating in basic mode is 1.644 μs.
- The maximum delay for FDDI ports operating in hybrid mode and no latency adjustment buffer is configured in the repeat path is 1.804 μs.
- In addition, for FDDI-II hybrid mode operation, each latency adjustment buffer can contribute up to 11.52 μs in the repeat path of the packet channel. There is not necessarily more than one latency adjustment buffer in an FDDI-II network.

For FDDI or FDDI-II in basic mode, the default timer values mean that D_Max, the maximum delay that the ring initialization process allows, is 1777.45 μs. For FDDI-II in hybrid mode, the default timer values mean that D_Max is 2833 μs. While it may be possible in principle to change default values (T_Max and TVX) in every network MAC to allow larger configurations, this would be tedious and complex, since there would be a large number of MACs to change in most networks where this is a concern. We will consider only the maximum configurations without changing default values.

To calculate network delay assume that each node contributes two port delays (due to a wrapped dual ring and one S port plus one D port per single attachment node). Figure 19 provides a graph of maximum fiber cable length versus number of network nodes for three alternatives: (a) FDDI ports with no HRC in any paths, (b) FDDI-II ports with HRC in every path but in basic mode, and (c) FDDI-II ports in hybrid mode. Figure 19 assumes an entirely

Table 5 - D_Max Computation Worksheet

	Quantity	Times	Delay (μs)	Total Delay (μs)
length of fiber cable:	_________ km	10.16 μs/km		
length of TP cable	_________ km	12.12 μs/km		
Total delay due to signal propagation through cable:				
number nodes w/o HRC:	_________ nodes	2.328 μs/node		
number nodes with HRC:	_________ nodes	3.328 μs/port		
Total delay due to repeating through ports:				
Grand total delay due to ports and cable:				

fiber optic network. Table 5 provides a worksheet for calculating the maximum propagation delay of an FDDI network in basic mode.

We can see that the default configuration limits of FDDI are not actually quite as liberal as the often quoted 500 nodes and 100 km of fiber. Still they allow very large rings, and it is doubtful that many actual rings will stretch the limits.

4.2 Partitioning FDDI Rings

We considered the maximum configuration limits of FDDI networks in 4.1.3. However, even if we can have very large rings with several hundred nodes on one token ring, is it desirable? Clearly, even if the FDDI protocol allowed rings of unlimited size there would be practical reasons for limiting their size. There are a variety of means for connecting FDDI rings together, including bridges and routers. They are discussed in section 8. The question we will consider here is how big do we wish to let an FDDI ring get before we use a bridge or router to partition it into two separate networks.

This is a performance and a maintainability issue. If a single ring gets too big then the load on that ring may exceed its capacity and the ring latency will increase packet queuing delays. Broadcast and management traffic will consume a larger fraction of the bandwidth if there are too many nodes. If there are 500 nodes on a ring, then the average bandwidth available to each node is only 0.2 Mbps. When does it make sense to pay for a 100 Mbps network to get only 0.2 Mbps of bandwidth? While there may be special applications where most stations are primarily listeners, and while computer traffic is usually bursty it's also true that there are usually relatively long peak periods when the network load in computer networks is much higher than the average over a 24 hour period. Clearly, there are not going to be many situations when a 500 node FDDI ring makes good economic sense. Hardware that runs at 100 Mbps is too expensive to use when the peak per station load is relatively light. There will nearly always be more economical solutions to the need to attach very many such stations to one network.

It may also be be difficult to maintain reliable operation on large rings since a fault on any one node affects the entire ring. Devices such as bridges and routers (discussed in 8.) allow a network to be partitioned into several subnetworks and provide a "firewall" which may isolate the network as a whole from problems in a subnetwork. In many cases, traffic in a partitioned network will remain confined to only one subnetwork and will not affect the rest of the network, improving overall performance. But bridges and routers introduce additional traffic of their own, cause processing and additional MAC queuing delays at every bridge or router a frame traverses, and add extra nodes, cost, protocols and complexity to the network.

Philosophies vary on how large a single ring should be. The best answer depends upon the particular situation. For example, if most traffic is within work groups, this favors the partitioning of the network by work groups. Traffic between work group members would not affect the rest of the network. If, however, those work groups were spread around several buildings or locations, it might be impractical to cable the network into separate rings by work group.

Many organizations have a traffic pattern where most traffic flows between individual workstations and a few servers, rather than between individual workstations. Often these servers may be mainframe computers. In many cases the workstations will be primarily on 802.3 or 802.5 LANs rather than FDDI. In these cases the logical structure is a relatively small FDDI dual ring "backbone" to which only bridges, routers and major servers are connected.

Some vendors recommend that FDDI rings be limited to 100 or fewer nodes. This is probably sound general advice, although there must be exceptions. It is, however, clear that only the most exceptional cases will strain the generous FDDI configuration limits.

4.2.1 Dual Rings versus Trees

The FDDI dual ring of trees topology is discussed in 2.2. Dual rings provide a powerful fault isolation capability, the ability to wrap around faults including cable faults as and disconnected or failed nodes. But, although powerful, the capability is also limited. Two faults on a dual ring usually split it into two disjoint rings. It is important when cabling a dual ring for survivability to ensure that multiple dual ring cable segments do not share the same cableway or cable to prevent one accident from splitting the network into two or more disjoint parts.

Users will disconnect equipment from the network and this will cause wraps if the equipment is on a dual ring. It is therefore generally not advisable to run dual rings directly to user workstations. Some FDDI vendors take the relatively extreme position that nothing but concentrators should ever be on a dual ring, and those concentrators must be locked away from users. A less extreme position allows bridges, routers, major servers, and concentrators, located in restricted access areas, to be attached to dual ring backbones. For the protection of the FDDI network, FDDI equipment accessible to ordinary users should be a single attachment station connected through a concentrator

5. FDDI Media

FDDI can be used with a variety of physical media. A single ring can pass over different media types and the ports on dual attachment stations or concentrators may support different media. FDDI PMD standards have been defined or are under development for the following Media types:

- 62.5/125 µm multimode fiber
- single mode fiber
- shielded twisted pair wire
- unshielded twisted pair wire
- SONET (Synchronous Optical NETwork) carriers.

5.1 Optical fiber media

FDDI was designed to use fiber optic media. This section provides a brief introduction to optical fiber its limitations when used with FDDI.

5.1.1 Introduction to Optical Fiber

While optical fiber wave guides can be made of other materials, commercial data communications fibers today are made from very pure Silicon dioxide (SiO_2) glass, the same basic material as ordinary window glass. Small, precisely controlled amounts of Germanium (Ge) are added to the fiber, in a process called *doping* to control the index of refraction of the glass. Adding Ge increases the index of refraction. Figure 20 illustrates the three kinds of optical fiber in common use:

a) *Step index* multimode fiber which has an abrupt step change of refractive index, causing light conducted down the core of the fiber to be sharply reflected by the cladding of the fiber. Since there are many paths (modes) which a photon may take through the fiber of many different lengths, a pulse of light is spread as it travels down the fiber. For this reason step index fiber has a relatively low bandwidth and no standard variant of FDDI is intended to use step index fiber. However, it may be possible to successfully connect two FDDI stations by means of step index fiber, if the length of fiber is short.

b) *Graded index* multimode fiber which has a core with a carefully controlled index of refraction profile, so that the individual fiber modes are not abruptly reflected, rather they appear to be gradually bent. The greater the index of refraction, the slower the speed of light (the speed of light is a function of the dielectric constant of the medium). The index of refraction is greatest at the center of the fiber core and decreases as the distance from the center increases. For this reason, photons in the "low order" modes, which usually travel near the center of the core, move more slowly than photons in the "high order" modes which travel more of their path in the outer, lower index of refraction, region of the core. This effect compensates somewhat for the greater distance that photons in the high order modes travel, reducing the pulse spreading or *modal dispersion*, and correspondingly increasing the bandwidth of the fiber. This is the general type of fiber specified in FDDI PMD. A graded index multimode fiber is usually characterized in terms of its inner (core) diameter, outer diameter and *numerical aperture (NA)*. The numerical aperture is a measure of the "cone of acceptance" for light coupled into or emitted from a multimode fiber.

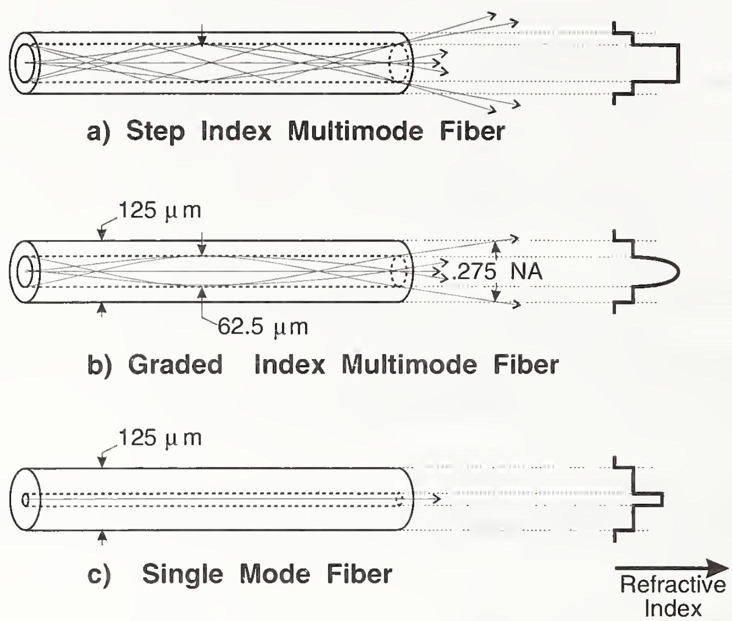

Figure 20 -Optical Fiber Type.

The fiber used in FDDI PMD is referred to as 62.5/125, NA 0.275 fiber. That is it has a core diameter of 62.5 µm, an outer diameter of 125 µm and a numerical aperture of 0.275. Reference [FIPS 159] provides a convenient specification for this type of fiber.

C) Single mode fiber which has a small core able to sustain only one propagation mode. This entirely eliminates modal dispersion, although other dispersion effects remain. Since the core of the fiber is very small, connectors must maintain very precise alignment, and laser transmitters are usually required to couple enough light into the fiber. SMF-PMD employs single mode fiber and laser transmitters to allow maximum link distances of up to 60 km.

There are two principle loss phenomena in optical fiber. One, which dominates the shape of the loss versus wavelength curve shown in Figure 21, is known as *Rayleigh scattering*. Rayleigh scattering is produced by localized variation in the index of refraction in the glass, due to the irregular arrangement of atoms of different atomic weight (Si, O and Ge) in the glass. It is an inherent property of the material and cannot be eliminated. Rayleigh scattering varies inversely with the fourth power of wavelength, so at short wavelengths it increases dramatically. The loss of the fiber decreases as wavelength increases until at about 1600 nm the photons excite thermal activity in the glass and are adsorbed.

The other affect is adsorption, caused by impurities in the glass. The most difficult impurity to entirely eliminate is typically water (in some processes the fiber is made with torches which burn hydrogen, producing, of course, water). The water results in hydroxyl (OH^-) ions, which exhibit distinct adsorption peaks in the glass. Figure 21 illustrates the resulting adsorption peaks, often called water peaks. It is not easy to precisely control the hydroxyl ions, so the height of the peaks will vary somewhat from fiber to fiber. For this reason communications systems avoid operation in the water peaks. These, and the semiconductor technologies used to

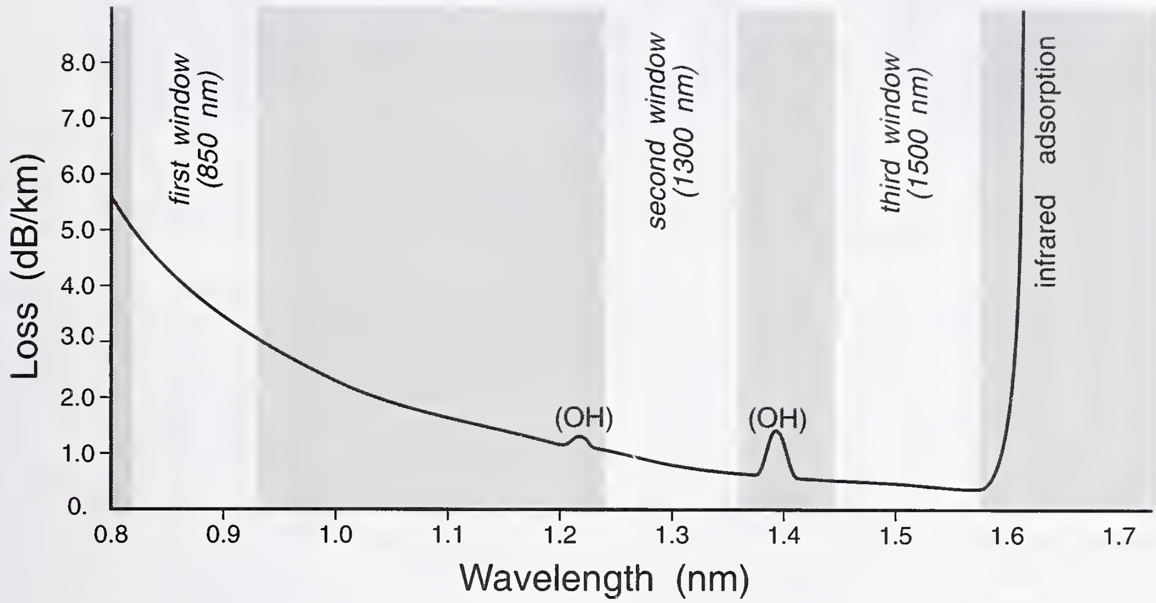

Figure 21 -Loss versus Wavelength in Silica Glass Fiber.

make transmitters and detectors divide the spectrum into three *windows* or regions of operation. They are called the *first window* (850 nm window), the *second window* (1300 nm window) and the *third window* (1500 nm window).

Losses in the first window are much higher than in the second window, and losses in the third window are lower still. However, relatively inexpensive AlGaAs LED and laser devices can transmit light in the first window. More expensive InGaAsP devices are generally needed for second and third window operation. Inexpensive Silicon detectors can be used in the first window, while more expensive InGaAs detectors are needed for the second and third windows. In general first window devices are the least expensive, while third window devices are the most expensive.

If cost were no object, it would seem best to operate in the third window. However, in addition to the modal dispersion described above, fiber bandwidth is also affected by *chromatic dispersion*. Chromatic dispersion is produced because the speed of light is affected by the wavelength of the light as well as the dielectric constant of the medium. The chromatic dispersion curve in most Ge doped SiO_2 glass fiber becomes zero at a wavelength of about 1330 mn. For this reason all three fiber FDDI PMD variants operate in the second window.

Still, for PMD and LCF-PMD, which use LED sources which have a relatively broad chromatic spectrum, chromatic dispersion is the factor which usually limits their maximum link distance. While PMD has a restrictive maximum chromatic width requirement for transmitters, there is no such requirement for LCF-PMD. This eliminates the cost required to test the chromatic width of LED transmitters for LCF-PMD.

5.1.2 Fiber Cables

Optical fiber is normally procured and installed in cables that have one or morel glass fibers, strength members and various sheaths, buffers and jackets protecting the fibers.

A wide variety of cables are available for different purposes. Cables are generally divided into two types:

- *loose buffer* cables in which several optical fibers are placed in a rigid plastic tube whose diameter is large enough to include a water resistant gel that protects the optical fiber. Loose buffer cables are normally used outdoors and the gel is a fire hazard which prevents their use inside buildings;
- *tight buffer* cables made by tightly wrapping the fibers in a protective jacket that is suitable for use in plenums. In general, tight buffered cables are used inside buildings, are thinner than loose buffer cables and allow smaller bend radiuses than loose buffer equivalents. Some tight buffered cables are also suitable for burial and aerial applications.

Indoor tight buffer cables are available in a wide variety of configurations. Simplex cables, illustrated in Figure 22 a), contain a single fiber protected by a coating and a buffer, with a strength member and an outer coating. Duplex zip cable, illustrated in Figure 22 b), is similar, but two cables are joined by the outer jacket material. Distribution cables, illustrated in Figure 22 c), enclose a number of buffered fibers in one jacket for trunking applications. Breakout cables, illustrated in Figure 22 d), have a number of separately jacketed fiber cables are enclosed in one sheath to facilitate attaching connectors and installing the fiber. Cables must be plenum rated for use in plenum areas in buildings; this is primarily a matter of a flame retardant low-smoke jacket. Riser cables are designed to be suspended in vertical risers in buildings.

Fibers are terminated with connectors. The strength member of the cable is connected to a strain relief device on the shell of the connector, to prevent the fiber from carrying any force and the fiber itself is stripped down to the cladding and inserted in the ferrule of the connector. Epoxy is generally used to bond the ferrule to the fiber and, after the epoxy has cured, the ends are polished to make a flat surface. Connector vendors offer kits for terminating connectors. The FDDI MIC is illustrated in Figure 8; in building patch panels and cross connects simplex "ST" or "SC" are normally used. Adapters are available to couple two simplex connectors to an FDDI MIC.

5.1.3 Link Dispersion and Power Budgets

A FDDI fiber optic link is constrained by two limits: its power budget and its dispersion budget. The two multimode PMD variants are most strongly constrained by dispersion. There are two main components to dispersion, modal and chromatic dispersion. A source with a narrow spectrum (often called a narrow *line width*) reduces the effect of chromatic dispersion. FDDI PMD was designed to allow the use of surface-emitting LEDs which have a rather broad spectrum. FDDI LCF-PMD limits only the center wavelength of the transmitter. So beyond specifying the PMD type, users have little control over the spectral width of FDDI stations.

FDDI PMD and LCF-PMD was designed to use 62.5/125 μm multimode fiber with a numerical aperture of 0.275. This is the standard multimode fiber recommended for use in office building wiring in the United States (see 6.). This type of fiber is available with over a range of modal or "laser" bandwidths. FDDI PMD specifies a minimum laser bandwidth of 500 MHz·km. This ensures that a link can extend for two km, and a fiber with a higher bandwidth is not guaranteed to allow the link to be extended to a significantly greater distance, because chromatic dispersion is often the dominant effect. However 62.5/125 μm fiber fiber that exceeds 500 MHz·km is

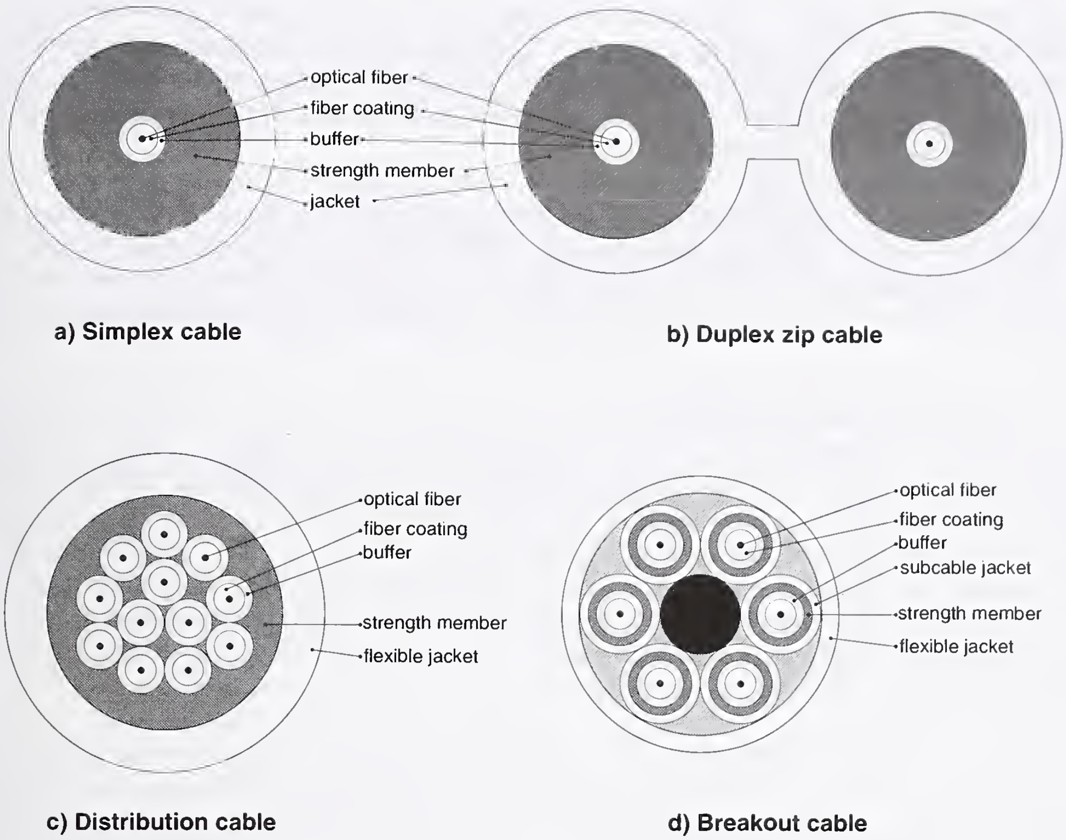

Figure 22 - Fiber Cables.

readily available and its use may be prudent in backbone and trunk cabling to allow for the use of that fiber with future, higher bandwidth systems.

The power budget of an LCF-PMD link is 7 dB and the power budget of a PMD link is 11 dB. A margin of 1 dB is typically reserved for component aging, leaving a remaining budget of 6 dB and 10 dB respectively. Link loss is estimated by summing the loss for the connectors, splices, bypass switches and the fiber itself along the link. Cabled fiber used in building wiring should be specified to have a loss of not more than 2.5 dB per km at 1300 nm, while cabled fiber with a loss of 1.0 dB per km at 1300 nm is available. Connector losses depend on the type of connector and the skill with which the termination is made. A nominal loss of 0.5 dB is commonly assigned for each connector in patch panels, cross connects, and wall plugs between the two stations, however the end connectors that plug directly into FDDI stations do not add extra loss; they are already allowed for in the FDDI output and input port specifications. Splice losses must also be accounted for in the budget and a nominal loss of 0.5 dB per slice is usually assumed.

MAC specifies a maximum bypass switch loss of 2.5 dB in the bypass state, although switches with less loss are available. If optical bypass protection is desired, the 10 (or 6) dB link power budget for the bypass condition must include all the connectors, bypass switches and cable loss between operating stations. As a practical matter it will rarely be practical to bypass more than two or three successive stations.

When the cable plant is installed link loss can be measured directly with an optical power meter and a suitable light source. The estimate obtained by summing the individual losses respective loss elements is often higher than the total loss, even when the individual losses of the elements measured individually are accurately known. This typically is because loss elements such as connectors and switches frequently attenuate high order modes more than lower order modes. The first such loss element effectively removes most higher mode energy successive loss elements therefore cause less attenuation.

Although specified for use with 62.4/125 µm fiber, FDDI links can use other fibers, and links that mix fiber links can also sometimes be used. While 62.5/125 mm fiber should be installed when fiber is installed for FDDI, other fiber that has already been installed can sometimes be used. PMD Annex C provides some information on estimating the effects of using other fiber types.

5.2 Twisted Pair Media for FDDI

FDDI may be used with both unshielded and shielded twisted pair media. Variants of FDDI using twisted pair media are sometimes informally called "CDDI" for Copper Distributed Data Interface. A number of vendors have developed products for FDDI using twisted pair cables. These earliest products are not necessarily interoperable with each other, however the FDDI committee is producing a standard for the use of FDDI over twisted pair media, and that effort is nearing technical completion. Some vendors are anticipating the final standard and have announced products that are interoperable with each other (although it is early to tell if they will conform to the letter of the final standard).

The greatest challenge in sending data over unshielded twisted pair wiring is meeting the limitations established by the FCC [FCC] governing the electrical emissions allowed for computer equipment. To accomplish this, the FDDI TP-PMD standard will require the use of very high quality twisted pair wiring, described below, as well as the use of MLT-3 a three level code (see Fig. 10) and scrambling of the signal. The scrambling spreads the spectrum of the signal so that there are no peaks that would violate the FCC regulations.

5.2.1 Unshielded Twisted Pair

While the current FED-STD 1090 (based on EIA/TIA 568) specifies some transmission characteristics for unshielded twisted pair (UTP) cables, it makes no distinction between grades of such cabling. The cable characteristics are specified in this standard up to a frequency of 16 MHz. Recent developments indicate that UTP cables can be used for data at higher rates if the cables are carefully constructed, hence the term "data grade UTP cables" has become common in the industry, though until recently there has not been a clear definition of what the term means. The EIA/TIA has addressed this problem by issuing Technical Systems Bulletin TSB-36, *Additional Cable Specifications for Unshielded Twisted Pair Cables.* The Bulletin defines five categories of UTP cables as follows:

a) Categories 1 and 2: For voice and low speed data. Not defined by EIA/TIA 568;

b) Category 3: Characteristics are specified for frequencies up to 16 MHz for data on LANs like IEEE 802.3 10BASE-T and IEEE 802.5 4 Mbps token ring;

c) Category 4: Characteristics are specified up to 20 MHz. This category can be used for IEEE 802.5 the 16 Mbps token ring;

d) Category 5: Characteristics are specified up to 100 MHz. This category of cable will be specified by the forthcoming FDDI twisted pair PMD standard for links of up to 100 m.

The transmission characteristics described in the TSB apply to cables consisting of four pairs of 24 AWG conductors, as the 568 standard recommends for horizontal wiring. Cables of four-pair 22 AWG wires may also be used if they meet or exceed all the transmission requirements.

Cables in categories 4 and 5 are referred to as Enhanced Unshielded Twisted Pair (EUTP) cables; they provide improved Signal-to-Crosstalk noise margins that are required for LANs that operate at 16 Mbps or higher over distances of up to 100 meters. This 100 meters includes the 90 meters of the horizontal wiring run, plus the 10 meters allowed for patch cords, cross-connects, and cords from outlets to terminals.

The transmission characteristics that must be carefully specified to qualify cables as EUTP are mutual capacitance, characteristic impedance, attenuation, and near end crosstalk.

- **Mutual capacitance** of any pair shall not exceed 20 nF per 304 m (1000 ft) for category 3 cable, and shall not exceed 17 nF per 304.8 m (1000 ft) for categories 4 and 5.
- **Characteristic impedance** must be 100 Ω +/- 15% in the frequency range from 1 MHz up to the highest referenced frequency. As a result of structural non-uniformities, the measured input impedance for an electrically long length of cable will fluctuate as a function of frequency. These random fluctuations are superimposed on the curve for characteristic impedance which approaches a fixed value at frequencies above 1 MHz. The characteristic impedance is obtained from these measurements by using a smoothing function over the bandwidth of interest.
- **Attenuation** for each category is given in a table which gives a Figure in dB per 304.8 m (1000 ft) at frequencies ranging from 64 kHz to the maximum applicable to each category. As an example, the maximum attenuation at 16 MHz for category 3 is 40 dB per 304.8 m (1000 ft); for category 4 it is 27 dB, and for category 5 it is 25 dB. Attenuation Figures for category 5 cables go up to 100 MHz where 67 dB per 304.8 m (1000 ft) is specified.
- **Near End Crosstalk** (NEXT) is specified in terms of NEXT Loss, dB per 304.8 m (1000 ft), at specific frequencies in the band of interest for each cable category. The higher the loss, the better the cable. At 16 MHz, the Figure is 23 dB for category 3, 38 dB for category 4, and 44 dB for category 5. NEXT in a cable is measured by applying a balanced signal to a disturbing pair while measuring the output of a disturbed pair at the near end of the cable. At the far end of the cable, both the disturbed and disturbing pair are terminated with a 100 Ω resistor.

The characteristics given are based on measurements performed on cables removed from the reel and stretched along a nonconducting surface or supported in an aerial span. It should be noted that TSB-36 does not claim that these characteristics are sufficient in themselves to insure adequate system performance; no attempt is made to cover all critical system design parameters. However, the categorization provided by this Bulletin will be an aid to the equipment designer and systems integrator in determining the suitability of UTP for data transmission rates previously thought to be beyond the capabilities of this media.

The category 4 and category 5 cables are currently not specified in FED-STD 1090 (EIA/TIA 568) because their performance levels were not firmly established when the standard was published. When EIA/TIA revises 568 the contents of TSB-36 will be integrated with it. Cable manufacturers are already responding to this categorization standard by referencing TSB-36 in their product literature.

Table 6 - FDDI Media

Medium	PMD type	Maximum Distance	Loss Budget	Primary Use
62.5/125 µm multimode optical fiber, NA: 0.275	PMD	2 km	11 dB*	backbone; dual ring with bypass switches
	LCF-PMD	500 m	7 dB	concentrator to workstation
125 µm cladding single mode optical fiber	SMT-PMD	20 to 64 km	10 to 32 dB	links between separate facilities in one metropolitan area
SONET STS-3c	SPM	only limit is ring delay	N/A	links between separate facilities
100 Ω category 5 unshielded TP	TP-PMD	100 m	dispersion & EMI limited	concentrator to workstation
150 Ω type 1 shielded TP				

* note: The 11 dB PMD budget is for a 2.5 x 10^{-10} bit error rate. The PMD budget for a 10^{-12} bit error rate is 9 dB. LCF and TP loss budgets are based on a 10^{-12} bit error rate.

5.2.2 Shielded Twisted Pair

Shielded twisted pair cables with an impedance of 150 Ω are specified for use with the IEEE 802.5 token Ring standard at rates of 4 Mbps and 16 Mbps. The revised EIA/TIA 568 will contain a specification for "type 1" 150 Ω shielded twisted pair cable, suitable for use with 802.5 and FDDI twisted pair PMD. The FDDI twisted pair PMD standard now under development will support links using this medium for distances of up to 100 m.

5.2.3 Conclusion

It is possible to use twisted pair horizontal wiring for FDDI. Several proprietary commercial FDDI twisted pair products now exist. Early twisted pair FDDI adapter cards appear to be somewhat less expensive than multimode fiber cards and significant savings may be possible. The FDDI committee is developing a standard which will provide for interoperability between products. Until products that implement this standard are available, twisted pair solutions are proprietary, and a decision to use twisted pair media will generally require that both a concentrator and station adapters be procured from the same vendor, or from vendors whose products are known to be compatible.

In many cases, existing unshielded twisted pair wiring, even "datagrade" will not be satisfactory for FDDI, however existing "type1" shielded twisted pair will generally be acceptable. Where unshielded twisted pair is to be installed for FDDI, category 5 twisted pair should be specified. Use of FDDI twisted pair for FDDI should be confined to horizontal wiring and the maximum distance should be limited to 100 m. FDDI adapters for twisted pair may apparently work when used with existing twisted pair wiring, particularly for short distances, however there is a significant risk that FCC regulations [FCC] governing emissions from computer equipment will be violated. For that reason it is recommended that only category 5 unshielded twisted pair or type 1 shielded twisted pair cables be used with FDDI.

5.3 Selecting FDDI Media

Table 6 summarizes the FDDI media alternatives and their applications. The medium used for an FDDI link depends upon the use of that link and, in particular the distance the link must traverse. In the case of multimode and single mode fiber, there are PMD alternatives that use the same fiber.

When FDDI links of more than 2 km are required, generally to connect different buildings or facilities in the same metropolitan area, there are two choices:

- single mode fiber (SMF-PMD)
- SONET (SPM).

The choice of SONET implies that an STS-3 SONET channel is available between the end-points.

For building and campus backbones and in building risers 62.5/125 multimode fiber is the usual choice, provided links are of 2 km or less. Optical bypass switches can be used with dual rings to improve reliability. When backbone links are limited to 500 m or less, and involve no more than two patch panels or cross connects, and have no bypass switch, then LCF-PMD may be employed for the link, while if distances exceed 500 m or an optical bypass is to be used, the PMD must be selected. In either case the fiber is the same, only the components in stations and concentrators change.

Wiring closet to desktop links can be either:

- 62.5/125 multimode fiber
- category 5 unshielded twisted pair
- shielded twisted pair

When multimode fiber is used to the desktop, the maximum distance from a concentrator to the desktop should not exceed 500 m, and not more than two patch panels or cross connects should be employed in the link. This will allow the use of LCF-PMD components, which are limited to a maximum link distance of 500 m and a power budget of 7 dB.

When twisted pair is used, the maximum link distance is 100 m. This means that, in many cases, it will be necessary to locate a concentrator in each wiring closet to prevent link distances from exceeding 100 m.

At the present time, early twisted pair FDDI equipment enjoys a significant cost advantage over multimode fiber PMD products. LCF-PMD products are not yet available and these may significantly reduce the cost differential between twisted pair and fiber. While satisfactory operation of FDDI can be accomplished over twisted pair media, fiber is undoubtedly more robust, and offers nearly complete immunity from radiated noise problems. If networks or links are to be installed in electrically very noisy environments, then fiber is the clear choice.

6. FDDI and Standards for Wiring of Office Buildings

In general, FDDI is consistent with the standards which define structured wiring plans for office buildings. However, FDDI does place some additional constraints on the types of media and distances for which they can be used.

6.1 Introduction to Building Wiring Standards

Three Federal Information Processing Standards adopting EIA/TIA standards have been issued to be used in the wiring of Federal office buildings. The document specifies minimum requirements for telecommunications wiring within a building and between buildings on a campus. The specifications are built around a recommended topology and specify allowable distances with selected transmission media. The transmission media may be twisted pair, coaxial cable, or optical cable, but the allowable variation within these basic types is restricted.

FIPS 174 [FIPS174] is intended for use by all departments and agencies of the Federal Government, both for new buildings and for major upgrades of existing plant. The standard states that it applies to sites with a geographical extent up to 3000 m (9,840 ft.), up to 1,000,000 square meters (10 million sq. ft.) of office space, and with a population of up to 50,000 users. The wiring systems so defined are expected to have a useful life in excess of 10 years. FIPS 176 [FIPS176] is a companion standard that is aimed at smaller installations.

FIPS 175 [FIPS175] addresses design and construction practices that have a significant impact on the design of a building wiring system. This document provides design guidelines for telecommunications closets, ducts, raceways, service,entrances, etc. Grounding and binding requirements will be addressed in another FIPS to be issued in 1993 [FIPSXXX].

The telecommunications infrastructure in commercial buildings has become so complex it is imperative that the documentation and system administration be approached in a structured manner. Hence, to complete the series on building wiring, the TIA is developing TIA-PN-2290, *Telecommunications Administration Standard for Commercial Buildings,* to expedite the management of documentation of all building wiring and the related pathways and spaces that contain wiring.

6.2 Summary of FIPS 174 Requirements

Basic Structure. The primary functional elements that comprise the generic wiring system are the horizontal wiring, backbone wiring, telecommunications closet, and the work area. The standard establishes the performance criteria for each of these elements so that a multitude of telecommunication products and services can be accommodated and for ease of making changes as usage evolves.

Horizontal Wiring. The horizontal wiring is that portion of the system that extends from the closet to the work area. It is referred to as horizontal because it usually runs along the floor or ceiling. It is usually less accessible than the backbone wiring, hence it must be designed so that changes to configuration can be made in the wiring closet. Therefore, the horizontal wiring always has a star topology, i.e., each work area outlet is connected directly to the closet. It should be noted that bus and ring topologies can be accommodated with this scheme, but this is accomplished by connections inside the closet and are considered part of the backbone wiring.

The distance from closet to work area outlet is specified to be no more than 90 meters (295 ft) independent of media type. In establishing this distance, an allowance was made for 3 additional meters from the outlet to the terminal equipment. There are also length limitations on the cross-connect jumpers and patch cords in the wiring closet.

Only four types of cables are recognized as acceptable in the horizontal wiring system. These are:

(1) four-pair 100-ohm unshielded twisted pair (UTP);

(2) two-pair 150-ohm shielded twisted pair (STP);

(3) 50-ohm coaxial cable;

(4) 62.5/125 m optical fiber cable.

A minimum of two telecommunications outlets shall be provided for each individual work area; one outlet may be associated with voice and the other with data. (It is not necessary that they be on separate face plates.)

One outlet shall be supported by a four-pair UTP cable; the second outlet shall be supported by any one of the other copper media. If optical fiber is to be installed it is to be in addition to these two outlets. The optical fiber cable may consist of only fiber or it can be a hybrid cable with one or more of the recognized media.

FIPS 174 adopts the industry wiring standard entirely with only one important change: for the unshielded twisted pair work-area outlet, the EIA/TIA standard allows an optional pin/pair assignment on the eight-position jack. The FIPS deletes that option and allows only one configuration, the one that is compatible with the majority of equipment manufactured in North America. This is a simple but very important step in the direction of achieving equipment interoperability when transporting terminal equipment from one building to another.

Backbone Wiring. The backbone wiring provides the interconnections between wiring closets, equipment rooms, and entrance facilities in the wiring system structure. The closets, equipment rooms, and network interfaces may be located in different buildings. The backbone wiring includes not only the transmission media, but also the group termination blocks (cross-connects) that are used to administer the building wiring. All backbone wiring must use a hierarchical star topology, i.e., each closet is wired to a main cross-connect, or to an intermediate cross-connect then to a main cross-connect. There can be no more than two hierarchical levels of cross-connects in the backbone wiring; this limitation is imposed to limit signal degradation in passive systems and to simplify administration. Systems which are designed for non-star configurations such as ring, bus, or tree, can be accommodated by the star topology through the use of appropriate interconnections, electronics, or adapters in the cross connects in the wiring closets. Where bus or ring configurations are anticipated, direct connections between closets are allowed. Such connections are additional to the connections of the basic star topology.

The cables recognized for backbone transmission media are the same as for the horizontal wiring except that there is no limitation on the number of twisted pairs within a cable. The user is warned, however, that crosstalk between individual, unshielded pairs may affect the transmission performance of multipair cables. Engineering to satisfy specific requirements is considered beyond the scope of the standard. The standard recognizes that it may be advantageous to use single mode optical fiber in the backbone for some applications. This is left open as a subject for future study.

Maximum distances for cable runs in the backbone depend on the media type used. The maximum distance between the main cross connect and the wiring closet for unshielded twisted pair or 50-ohm coax is 500 meters (1640 ft); for shielded twisted pair the distance is 700 meters (2296 ft), and for 62.5/125 optical fiber it is 2000 meters (6560 ft). The distance allowed between a closet and an intermediate cross connect is similar to the above, but for optical fiber the total distance from closet through the intermediate to main cross-connect shall not exceed the maximum of 2000 meters. It is, of course, generally advantageous to locate the main cross-connect near the center of the site.

Cable Specifications. The cable specifications given in the standard are for essential media transmission characteristics; these are in addition to requirements to conform to the National Electric Code (NEC) or any other applicable building codes.

For 100-ohm unshielded twisted pair, 24 AWG (American Wire Gauge) wires with thermoplastic insulation is the recommended type, although 22 AWG may be used if it meets the other transmission requirements. These other requirements include DC resistance, mutual capacitance, attenuation, characteristic impedance, and near-end crosstalk (NEXT). Attenuation, characteristic impedance, and NEXT are specified at several frequencies from 64 kHz to 16 MHz. Note that these specifications are not sufficient for FDDI over unshielded twisted pair, which requires a higher grade of cable (see 46).

For detailed specifications on 150-ohm shielded twisted pair, the user is referred to an EIA Interim Standard Omnibus Specification, NQ-EIA/IS-43. Similarly, the requirements for the 50-ohm coaxial cable are found in IEEE 802.3, 10BASE2 (horizontal wiring) or 10BASE5 (backbone).

The optical fiber to be used is multimode, graded index waveguide with a nominal core/cladding diameter of 62.5/125 μm. This is the fiber specified in FIPS PUB 159 (supersedes FED-STD 1070); this document adopts the EIA/TIA-492AAAA specification. The essential transmission characteristics are maximum attenuation and minimum information transmission capacity; these are specified for two different wavelengths. For a wavelength of 850 nm the maximum attenuation is 3.75 dB/km and the minimum information capacity 160 MHz-km; for 1300 nm these Figures are 1.5 dB/km and 500 MHz-km, respectively.

Connecting Hardware. All the four-pair UTP cables at the work area are to be terminated with the 8-pin modular jack (RJ-45) that is common in the telephone industry. Pin/pair assignments and color codes for the individual wires are given in the standard. As mentioned above, this is the one area where FED-STD 1090 differs from EIA/TIA-568. The latter allows an optional pair assignment at the connector while the FED-STD allows only one arrangement.

The connector used for terminating the shielded twisted pair at the outlet is that specified by ANSI/IEEE 802.5 (token ring). This connector is hermaphroditic in design so that two identical units will mate when oriented 180 degrees with respect to each other. The connector also provides the means to preserve the integrity of the shielding.

The 50-ohm coaxial cable is terminated with a female BNC connector at the workstation outlet, as specified in IEEE 802.3 10BASE2. The coaxial cable in the backbone wiring is terminated with the male Type N connector per IEEE 802.3 10BASE5.

No connector type is specified in the standard for the optical fiber, either at the workstation or in the wiring closet. However, it does state that "all backbone fiber wiring shall be terminated with permanently installed rematable connectors." The maximum optical attenuation per each mated connector pair shall not exceed 1.0 dB and the connectors shall sustain a minimum of 200 mating cycles without violating specifications.

6.3 FDDI with structured wiring.

FDDI fits well into the structured wiring plan of FIPS 174, with a few additional constraints. The major additional constraint is that category 5 twisted pair, or shielded twisted pair should be used in horizontal wiring where operation of FDDI over twisted pair is contemplated. The wire distance between a wire closet containing an FDDI concentrator and a workstation must be kept to 100 m or less if twisted pair is to be employed for FDDI.

Figure 23 illustrates the application of structured wiring to FDDI. A campus network is illustrated. Each building contains a dual attachment concentrator in the building wire center. The links between these concentrators use either multimode PMD or single mode PMD and are either 62.5/125 µm multimode fiber or single mode fiber. Bypass switches may be used at the concentrators to improve availability.

Within a building, 62.5/125 mm multimode fiber is used in the trunks and risers between the

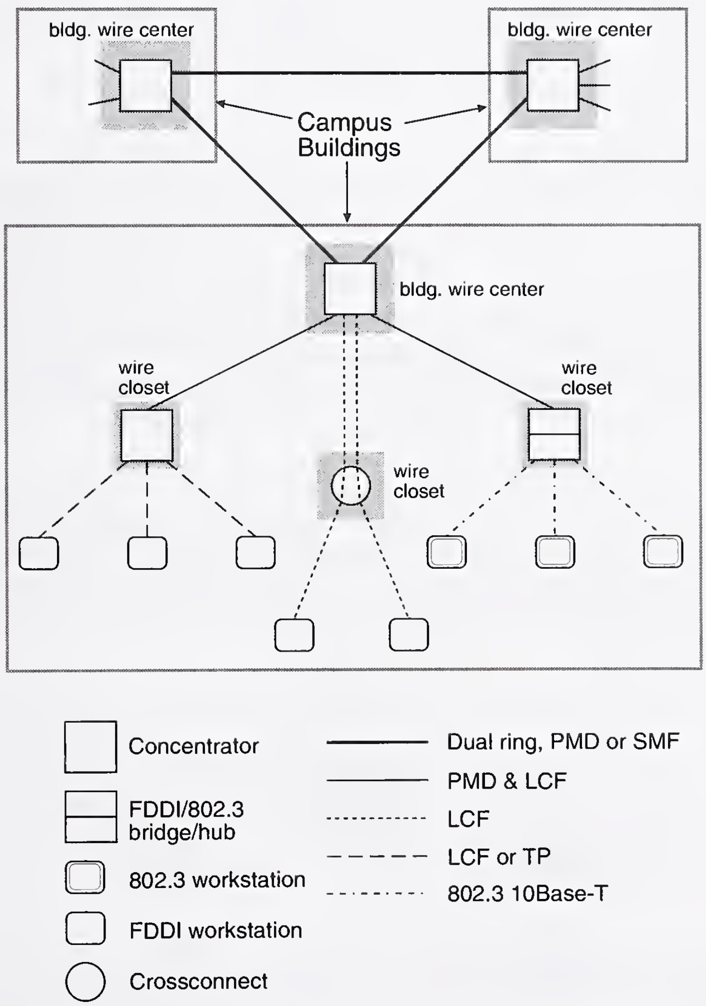

Figure 23 - FDDI and Structured Wiring.

concentrator in the building wire center and concentrators in wire closets throughout the building. Either PMD or LCF PMD FDDI ports may be used, depending upon distance, and the two may be mixed on the same link, provided the LCF distance and loss budgets are not exceeded.

Horizontal wiring to workstations may be category 5 twisted pair, shielded twisted pair or 62.5/125 mm multimode fiber. In most cases, where workstations are attached via optical fiber, they will use the LCF-PMD. While twisted pair should be limited to short (less than 100 m) horizontal runs, it is practical to run FDDI links from a concentrator, through raceways or risers, to a wire closet, and then, through an optical cross connect, on in the horizontal wiring to a workstation. In this case, assuming the use of LCF-PMD FDDI ports, the maximum distance for a link should not exceed 500 m and the maximum loss should not exceed 7 dB.

Other standards LANs, such as 802.3 and 802.5 can be integrated with FDDI and fit into the structured wiring approach, however they have somewhat less stringent media requirements. Figure 23 illustrates the use of 10Base-T connected to the FDDI backbone. While category 5 twisted pair that can be used with FDDI can also be used for 10Base-T, less expensive category 3 twisted pair will suffice for 10Base-T. If category 5 twisted pair is pulled in all horizontal runs, then any cable can be used for FDDI, 802.3, 802.5 or telephone service. However, if, the predominant workstation interface is to be 802.3, then it may be more cost effective to pull category 3 twisted pair in ordinary horizontal wiring, and pull category 5 twisted pair, shielded twisted pair or multimode fiber on an as needed basis.

6.4 Wiring Crossovers

All FDDI connector receptacles for the same media type are polarized identically. That is, they are polarized so that the same ferrule or electrical contact carries the transmit and receive signals. It is necessary that the cabling between two ports connect transmitters to a receivers. The transmitter contact or ferrule on the connector plug on one end of an FDDI cable link must be connected to the receiver contact or ferrule at the other end. This is informally called a "crossover." Nearly all FDDI fiber optic "patch cables," with FDDI connectors on the end, will have such a crossover, as illustrated in Figure 24. This is not necessarily the case with twisted pair patch cables.

According to FIPS 174, both optical and electrical horizontal are "straight through" from wiring closet to wiring closed and from closet to outlet. Crossovers may occur in the patch panels in wiring closets and in the wall outlet to FDDI equipment cables. There are too many possible configurations (including FDDI concentrators in wiring closets) to discuss or illustrate them each here, but it is necessary to ensure that there are an odd number of crossovers in the wiring between any two FDDI ports.

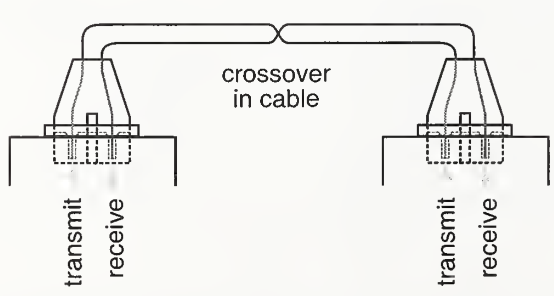

Figure 24 - Crossover in FDDI Patch Cable.

7. GOSIP, GNMP, Functional Standards and FDDI

This section explains how FDDI fits into the Government Open Systems Interconnection Profile (GOSIP) [FIPS146] and Federal Information Processing Standard 179 Government Network Management Profile (GNMP) [FIPS179]. Boland [BOLA] provides an introduction to GOSIP.

7.1 The GOSIP Architecture and FDDI

The OSI Reference Model [OSI] defines the familiar seven layer protocol stack illustrated in Figure 1. The OSI reference model is a very general structure and a large family of OSI standards has evolved. In fact, so many standards and options exist that simple conformance to OSI standards does not ensure interoperability. To promote interoperability and competition, GOSIP provides a common selection of OSI protocols and options for Federal use.

The current version of GOSIP is Version 2, which does not yet include FDDI, but does state the intention to include FDDI in a future version. GOSIP is a procurement specification. Use of GOSIP is mandatory when agencies acquire, "computer networking products and services and communications systems or services that provide equivalent functionality to the protocols defined in GOSIP." Note, however, that GOSIP recommends rather than mandates physical interface specifications, and allows the use of other non-proprietary interface standards.

GOSIP Version 3, which is in preparation, will be based on the *Industry/Government Open Systems Specification (IGOSS)*[IGOSS] which will include FDDI. The primary source of protocol specifications for the IGOSS is the *Stable Implementation Agreements for Open Systems Interconnection Protocols* [STABLE] called the "Workshop Agreements." There are now implementation agreements for FDDI in the Workshop Agreements and the IGOSS will include protocol specifications for FDDI LANs.

IGOSS defines several "subprofiles" as building blocks that may be selected and combined to define a particular procurement. Subprofiles define specific multi-layer protocol requirements for the provision of a particular service or subnetwork technology. The Application Subprofiles, Lower Layer Subprofiles and Subnetwork Profiles specified in IGOSS are illustrated in Figure 25.

7.1.1 Application Subprofiles

Application subprofiles are defined in the IGOSS for the services shown in Figure 25. These application subprofiles include specifications for the three OSI layers directly associated with applications:

- Application Layer: This layer consists of protocols and services to support user-defined application processes. Functions and services that are common to several application protocols are contained in "service elements" in this layer. The *Association Control Service Element (ACSE)* [ISO8650] and the *Remote Operation Service Element (ROSE)* [ISO9072] are examples of service elements specified in IGOSS. Application level protocols specified in the IGOSS include the *Message Handling System* [X.400], and the *File Transport, Access and Management (FTAM)* standard [ISO8571];
- Presentation Layer: This layer specifies or, optionally, negotiates the way information is represented for exchange by application entities. The Presentation layer is concerned only with the syntax of the transferred data;

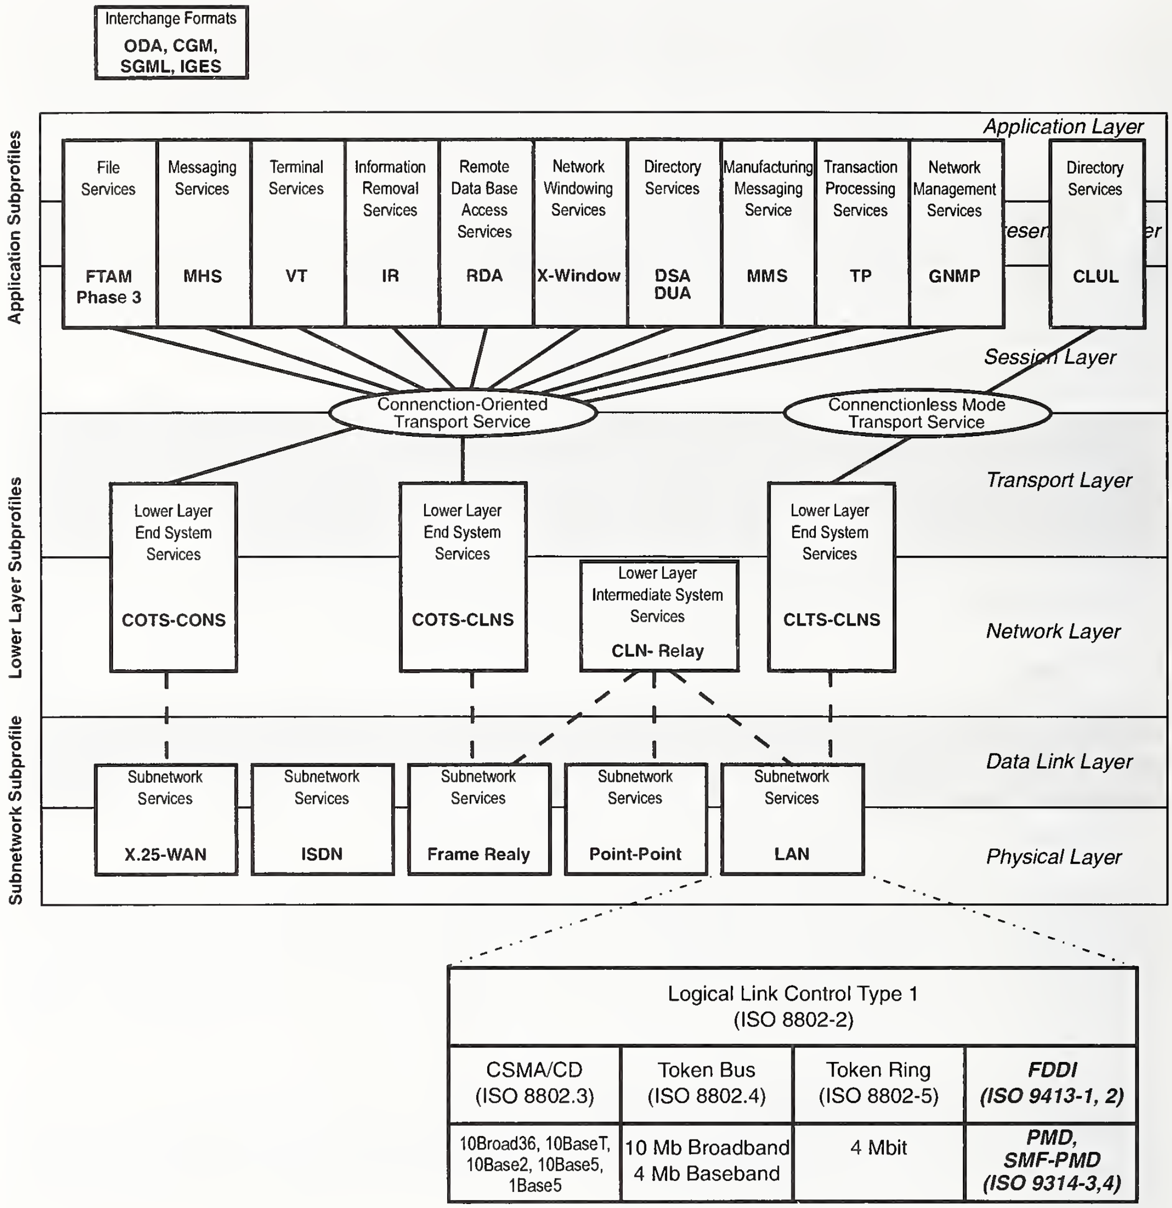

Figure 25 - IGOSS Subprofiles and FDDI.

- Session Layer: This layer aids in the orderly and reliable flow of data between users in cooperating end systems. Services include synchronization and check pointing.

7.1.2 Lower Layer Subprofiles

Lower Layer subprofiles provide reliable end-to-end data transfer. To assure interoperable data transfer for a variety of applications across a variety of subnetwork technologies, IGOSS requires support of Transport Protocol Class 4 (TP4) [ISO8073] and the Connectionless Mode Network Protocol (CLNP) [ISO8473]. Additional, optional, choices are provided for other lower level protocol combinations. The OSI layers included in the lower layer subprofiles are:

- Transport Layer: The connection-oriented TP4 protocol provides a reliable connection oriented end-to-end data service over possibly unreliable media. It can be used over either connectionless or connection oriented subnetworks. Specifically, this protocol ensures that data packets are received uncorrupted, in the correct order and that no packets are lost. TP4 provides for the retransmission of lost or corrupted packets and for flow control between end systems. The IGOSS also optionally provides a simple connectionless transport service that may be used over connection oriented subnetworks;

- Network Layer: The principle function of the network layer is routing and the interconnection of subnetworks, which may be of different types. In addition to routing, the Network Layer provides for the segmentation and reassembly of packets which may be required when packets cross from one subnetwork to another of a different type. Two Network Layer services are defined in GOSIP, the *Connectionless Network Service (CLNS)* [ISO8473], which is mandatory and a *Connection-oriented Network Service (CONS)* [ISO9574], which may be supported as an additional option. The CLNS, which is provided by the CLNP defined in ISO 8473, is specified by GOSIP for use with LANs such as FDDI.

7.1.3 Subnetwork Subprofiles

Subnetworks provide a data transport service between directly attached systems. As can be seen in Figure 25, IGOSS provides several subnetwork subprofiles. One of these, the LAN subprofile, includes FDDI along with other standardized LANs. The Subnetwork Subprofiles include specifications for the following two OSI layers:

- Data Link Layer: The Data Link layer is the lowest level at which frames or *Protocol Data Units (PDUs)* are apparent. It uses the raw facilities of the physical layer to provide nearly error free transmission of data frames, providing error detection and, optionally, correction. For LANs, such as FDDI, the Data Link layer is conventionally subdivided into two sublayers:

    - Logical Link Control (LLC): The LLC sublayer provides a consistent service interface for all standard LANs. The IEE 8802-2 LLC [ISO88022] protocol provides for two types of service. GOSIP selects LLC type 1 which provided a connectionless, unacknowledged (or “datagram”) service for use with the CLNS.

    - Media Access Control (MAC): The MAC sublayer provides data frames, error detection, address recognition and arbitration of access to a shared communications medium. The operation of the FDDI MAC is described in 2.5. Note that in Figure 25 FDDI PHY is included with MAC rather than with PMD. FDDI PHY is described in 2.4. Other standard LANs do not separate PHY from MAC, and the same PHY standard is applicable to all the FDDI PMD alternatives.

- Physical Layer: The physical layer provides for the transmission of data bits over a physical interconnection medium. The Physical layer defines the physical, electrical, mechanical and optical means used to transport data. In FDDI the physical layer is divided into two sublayers, PHY and PMD. FDDI PMD is described in 2.3. IGOSS recognizes that there are situations where LAN interfaces are required to support non-standard media and, while discouraging their use, does not prohibit nonstandard physical layer interfaces. This is significant for FDDI because, while the original multimode fiber PMD is currently included in IGOSS, the emerging twisted-pair and low-cost fiber media options (see 2.3.4 and 2.3.2) are not yet fully standardized and included in IGOSS. When standards are finalized, they will probably be included in later revisions of IGOSS, and their use is not now precluded by IGOSS.

7.2 FDDI's Role in GOSIP

As can be seen above in Figure 25, FDDI is simply one of several LAN alternatives to be included in the IGOSS and GOSIP Version 3, and LANs are only one of several Subnetwork Subprofiles. A decision to use FDDI does not mean that GOSIP is necessarily required; it is the functionality of the network or communications services to be procured which determines the applicability of GOSIP.

7.2.1 FDDI Testing

The IGOSS organizations are cooperating to establish a testing program, based on the GOSIP testing program, that is designed to verify that implementations conform to standards and can interoperate with other implementations of the same standards. A formal methodology for conformance testing is defined in [ISO9646]. The formal methodology requires the preparation of a *Protocol Implementation Conformance Statement (PICS)* for each standard. The PICS enumerates each of the mandatory and optional features of a standard. Using the PICS an *Abstract Test Suite (ATS)* is developed to verify conformance to each of the requirements of the standard. An ATS for protocols at the MAC layer and above are specified in a formal notation, the *Tree and Tabular Combined Notation (TTCN)*. TTCN is not applicable to an ATS at the Physical Level, which require measurement of physical parameters rather than the performance of a protocol based on PDUs.

A PICS and an ATS are being developed in the FDDI committee (X3T9.5) for MAC, PHY, SMT, and PMD. When complete, they will provide a conformance test for FDDI implementations. A great part of the tests for MAC and SMT are devoted to verifying the exception handling and error recovery protocol specifications. Many of the tests require specialized test equipment to produce fault conditions that normal FDDI equipment cannot ordinarily generate.

Interoperability tests verify that implementations interoperate successfully with each other. Conformant implementations may not interoperate and interoperable implementations may not be conformant. In an interoperability test it is not necessarily practical to produce and verify the response of implementations for every condition, state and parameter specified in the standard.

7.3 FDDI and the GNMP

FIPS Publication 179, *Government Network Management Profile (GNMP)* [FIPS179] supports the monitoring and control of the network and system components. The GNMP specifies:

- Use of the the OSI *Common Management Information Protocol (CMIP)* [ISO9596] and the *Common Management Information Services (CMIS)* [ISO9595] to exchange management information;
- Seven specific management functions and services:
  - Object Management Function;
  - State Management Function;
  - Attributes for Reporting Relationships;Alarm Reporting Function;
  - Alarm Reporting Function;

- Event Report Management Function;
- Log Control Function;
- Security Alarm Function.

– The syntax and semantics of management information. Managed objects are defined in standards documents and in the OIW Stable Agreements [STABLE]. The documents currently included in the GNMP are:

- FDDI Station Management [SMT];
- Definition of Management Information [ISO10164]];
- IEEE 802.1B LAN/MAN Management [LMAN];
- IEEE 802.3 Repeater Management [HUB];
- CCITT Generic Network Information Model[M.3100];
- ISO/IEC JTC1/SC6 Management Information Related to OSI Network Layer Standards [ISO10733];
- ISO/IEC JTC1/SC6 Management Information Related to Intermediate Systems to Intermediate Systems Intra-Domain Routing Information Exchange Protocol [ISO10589];
- ISO/IEC JTC1/SC6 Management Information Related to OSI Transport Layer Standards [ISO10733];
- Annexes A and B of part 18 (Network Management Implementors Agreements) of the OIW Stable Implementors Agreements [STABLE];
- The Network Management Forum Management Information Library [NMFMIL]

The GNMP is effective June 14, 1993 and its use will become compulsory for Federal agencies for use in solicitations and contracts for new network management functions and services 18 months after the effective date. A users guide on GNMP is also being written at NIST.

GNMP builds on GOSIP and, like GOSIP, the primary source of technical specifications for GNMP is the OIW Stable Implementation Agreements [STABLE]. The GNMP includes by reference the managed objects specified in FDDI SMT. The GNMP, however, applies primarily to the integrated management of overall enterprise networks, as illustrated in Figure 26. FDDI LANs will be only one subnetwork of the overall enterprise network.

While GNMP specified the use of CMIP to communicate management information and commands between open systems, little FDDI equipment now exists that implements CMIP management agents; SNMP or FDDI layer management (see 2.6.3) are more commonly implemented in FDDI nodes. It may therefore in some cases be necessary to use a CMIP agent that uses SNMP or FDDI layer management as the LAN manager to integrate FDDI LANs into the GNMP enterprise network management.

7.4 Functional Standards

The concept of profiles developed in GOSIP and the OIW has inspired the concept of an *International Standardized Profiles (ISP)*, otherwise known as a *functional standard.* ISO/IEC Joint Technical Committee 1 (JTC1) on Information Technology has established a Special Group on Functional Standardization, to coordinate the activities of the three workshops responsible for

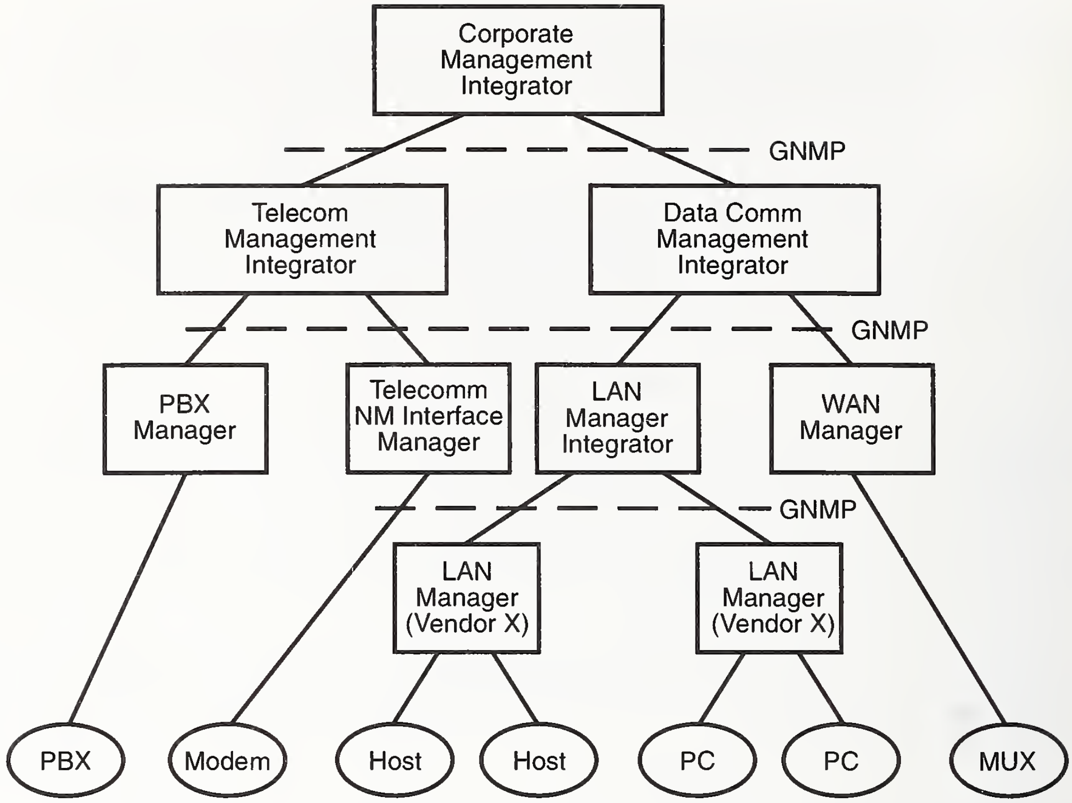

Figure 26 - Integrated Network Management Using GNMP.

the development of functional standards: the Asian-Oceana Workshop (AOW), the European Workshop on Open Systems (EWOS) and the OSI Implementors Workshop (OIW). ISPs are based upon the PICS provided for the base standard and the optional functionality required for particular profiles. An ISP for FDDI is under development by the AOW [ISO10608]. The IGOSS currently references the OIW implementation agreements for FDDI, which are similar to the draft FDDI ISP, but much less detailed, rather than the draft FDDI ISP.

8. Connecting an FDDI LAN to another Network

The first major application of FDDI has been to serve as a backbone for other LAN Networks such as the "802.3" *Carrier Sense Multiple Access with Collision Detection (CSMA/CD)* [ISO88023] and "802.5" *Token Ring* [ISO88025] LAN standards. Planning for an FDDI network will usually involve connecting it with other networks, often both LANs and *Wide Area Networks (WANs).*

It follows that users planning to install an FDDI LAN should, as a part of their process, plan how the FDDI LAN will be connected to other networks. The ad hoc model for LAN configuration has been the norm. An isolated LAN was added here for one purpose, another application and protocol was later added to that LAN for some other purpose, and a second LAN of another type was added for some other purpose. Three disparate noncommunicating groups of stations existed on two separate LANs. Then a need was perceived for communication between these groups and an ad hoc solution implemented. Soon other LANs, other applications and other protocols were added and so it went; ad hoc solutions improvised to communicate between them.

This has been more the rule than the exception. The reason for installing an FDDI LANs will often be to provide a systematic, manageable, high bandwidth replacement for an existing hodgepodge of improvised LAN interconnections. Performance on the existing system may be unsatisfactory and it may also be barely manageable. But FDDI is not a panacea. The problem of planning for such an FDDI network becomes a problem of organizing and systematizing a large heterogeneous network that just grew into its present shape. A significant concern is understanding the physical layout of the facility, and including a comprehensive wiring plan to account for the cable plant needs of all communications applications including voice, data and video. If the network designer is fortunate enough to be starting an entirely new network, he will still need to address those concerns, but will not be so constrained by history.

In this section we discuss the devices used to connect an FDDI LAN to other networks or to connect two separate FDDI LANs together. The following discussion will provide a brief general introduction to the kinds of products available to connect an FDDI LAN to other networks. Since this market is developing very rapidly, the introduction will be conceptual, and will not try to characterize the detailed characteristics of particular devices which are changing rapidly.

8.1 Some Definitions

Figure 27 provides a conceptual framework for describing the general categories of interconnection devices, in terms of the OSI levels at which network interconnection devices operate. Note that the terms used here are broadly used, but not always with precisely the same meaning, and the market for network interconnection devices is evolving very rapidly. The following terms are used in this chapter:

- *Bridge*. A LAN bridge originally was a MAC level device which joins two similar LAN networks. This concept has been extended to allow bridges to join different LANs whose main similarity is only that they share a common MAC level address space. There are several major classes of bridge protocols which are described in section 8.2.1 below.
- *Brouter*. A brouter is a single device which combines both bridge and router functionality. Brouters usually route PDUs that they recognize as addressed to their router function and which follow the appropriate protocol, while for other packets they function as bridges.

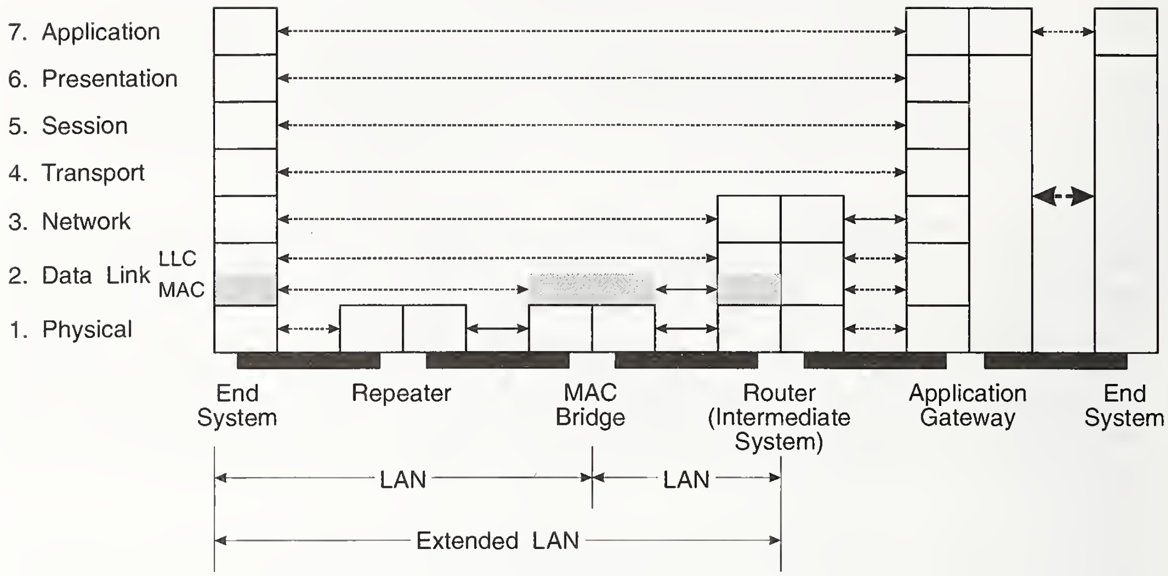

Figure 27 - OSI Layers and Network Connection Devices.

- *Concentrator.* An FDDI repeater node which serves as a hub for star or tree wired FDDI stations. See "hub" and "repeater" below and "concentrator"on page 4.

- *Extended LAN.* Two or more LANs joined together by bridges. Each MAC entity within the extended LAN has a unique address. In standard LANs such as FDDI, 802.3 and 802.5, all MACs have a 48-bit *universally administered address* (see Fig. 12). The IEEE provides a registration service that allows LAN vendors to ship each MAC with a 48-bit address which is unique in the world. FDDI also allows locally administered addresses, and, if they are used, they must be unique within the extended LAN.

- *Filtering Rate.* The number of packets per second which a bridge can examine to see if it must forward them to another network.

- *Forwarding Rate.* The number of packets per second which a bridge or a router can forward to another network, assuming that only the bridge or router is loading the output network. In very fast bridges or routers this may be limited by the speed of the slower network.

- *Gateway.* In the general sense, a gateway is any device which does a protocol conversion. By this definition a bridge which bridges an FDDI LAN to an 802.3 LAN is a gateway. By convention, however, such a device is simply called a "translating bridge," although it might be as sensible to call it a "MAC level gateway." In principle gateways can work at any level, however, since the architectures of various protocols do not not necessarily provide interface levels with comparable sets of services, one common kind of gateway is an *application level gateway*, which is shown in Figure 27. Such a gateway might, for example link a proprietary LAN mail protocol to the X.400 message handling protocol.

Another kind of common gateway is used with terminal emulation software in a workstation. In this case a number of PCs or workstations are tied to the gateway on a LAN, and the gateway is connected into mainframe network with a centralized control architecture. A special terminal emulation program provides the PC with the functionality of the "dumb" terminals normally used with the mainframe network. The PC terminals communicate over the LAN with the gateway, which appears to the Mainframe to be a normal communications controller device. The PCs are then able to serve as mainframe network terminals as well as PCs.

- *Hub.* This term is used in a variety of contexts. Here we use it to mean a device which allows a number of network links, of similar or dissimilar kinds, to be star wired to the hub. Hubs provide for systematic star or tree wiring of networks, support standard building wiring plans and often provide key network management facilities. Hubs may support the functions of bridges, repeaters or routers. By this definition an FDDI concentrator is a kind of hub.
- *Repeater.* A repeater is a physical level device that retimes and repeats a signal. It does not change or alter frames or MAC protocols, but it may change media. For example, an FDDI concentrator is a kind of repeater (and also a hub). The various FDDI links to the concentrator may use different PMD types, unshielded twisted pair, shielded twisted pair, multimode fiber, etc. Repeaters extend the size of a LAN by eliminating signal loss and dispersion, but two segments connected by a repeater remain one LAN subnetwork, with the total configuration limits of that type of LAN. In the context of FDDI or the 802.5 token ring, a single LAN subnetwork has one token path. Note that some commercially available 802.3 hub repeater devices may have functions which stretch the limits of this definition, for example they may allow the selective blocking of packets from some inputs to other outputs.
- *Router.* A router is a device that interconnects subnetworks at the network level, with network level addressing. This is a higher level of addressing than the MAC address and, while the structure differs between different Network level protocols, network level addressing is is hierarchical while MAC addressing (see 2.5.2) is "flat." That is network addressing identifies a structure of domains or of networks, systems and users, while a LAN MAC addresses is simply a unique 48-bit number, which is associated not with a network, but with a MAC entity of a station. In OSI terminology a router is an *intermediate system.*

8.2 Bridges and Routers

Gateways are too specialized and too particular to be considered here. We will consider the lower level network interconnection devices, bridges and routers. Typically network interconnection strategies will involve some combination of bridges and routers.

8.2.1 Bridges

Network interconnection devices which operate at the MAC level are considered to be bridges. The basic function of a MAC bridge connecting subnetworks together using MAC addresses to decide which packets to forward from one network to another. Today the concept of bridges has been expanded to allow for the conversion of frames from one subnetwork format to another and the connected networks need not share a common MAC level protocol. FDDI uses the IEEE 48-bit MAC addresses, and is commonly bridged to the 802.3 CSMA/CD and 802.5 token ring LANs, which also use the same addresses, although they have different frames and protocols.

Local bridges connect to LANs together directly at one bridge. *Remote bridges* connect two distant LANs through a long distance circuit (which is invisible to the stations on the LANs). In the past remote bridges have used a leased communications circuit between two similar bridges from the same vendor. The protocol between the bridges has been proprietary. The advent of new common carrier wide area communications protocols such as Frame Relay (see sec. 9.4) and the Switched Multimegabit Data Service (SMDS) (see sec. 9.6) offer alternatives to connecting remote bridges by leased line circuits.

Local bridges tie two or more LAN subnetworks directly together without an intervening long distance communications link. A local bridge or a pair of remote bridges may participate as a single element in other bridging protocols such as the spanning tree transparent bridging protocol discussed below.

Bridges serve the following functions:

- to interconnect LANs of similar or dissimilar kinds which are separated by long distances;
- to locally interconnect LANs of dissimilar type (i.e., 802.3, 802.5 and FDDI);
- to filter traffic between LAN subnetworks, increasing the capacity of the extended LAN network;
- to prevent faults on one LAN subnetwork from affecting the entire extended LAN.

Bridges can be divided into two general categories: *learning bridges* and *source routing* bridges. Learning bridges learn whether they must forward packets by observing the source addresses of packets on the networks to which they are connected. The bridge maintains a table of source addresses for each subnetwork. When a packet is received by the bridge on one subnetwork, with a destination address which matches a source address in the bridge's table for that subnetwork, then the bridge does not forward the packet because it knows that the destination is on the same subnetwork. All packets with destination addresses that do not match a source address the bridge has observed on the subnetwork are forwarded to all other bridged subnetworks. The bridge learns to forward only those packets which must be forwarded. An aging process deletes entries in the internal bridge address tables if a source address does not appear on the network for some period; this allows the bridge to remove or update entries for stations which have been removed or moved elsewhere on the network. FDDI requires MACs to announce their presence at regular intervals with an SMT NIF frame, so a learning bridge will quickly discover all the MACs on any FDDI ring to which it is attached.

When a learning bridge is first turned on, all packets are forwarded until the bridge learns which MAC addresses are on each subnetwork. If a destination is turned off, then its source address never appears on a packet and the bridge never learns where it is and therefore always forwards traffic addressed to it. Some LAN protocols involve periodic polling of stations to determine their status. If the polled destination is turned off, then there is never a response, and learning bridges will forward the polling messages throughout the extended LAN, contributing to network congestion. This means that protocols which rely on polling may cause problems on large extended LANs with learning bridges, even though the polled stations are on the same subnetwork as the polling station.

Learning bridges generally participate in a *spanning tree* algorithm, in which the bridges communicate with each other to establish a tree through the extended LAN, so that there is one and only one path between any two stations, preventing endlessly circulating packets. A root bridge node is determined by the spanning tree protocol. There may be redundant bridge links, however the spanning tree algorithm causes only one of the redundant bridge links to be active. If bridges fail or are turned off, the spanning tree is automatically reconfigured.

Learning bridges connecting dissimilar LANs may be either *translating bridges* or *encapsulating bridges*. The first FDDI to IEEE 802.3 bridges available were encapsulating bridges. They simply encapsulated the 802.3 frame in a proprietary frame format. Other encapsulating bridges on the FDDI ring recognized these encapsulated frames, removed the encapsulation, and forwarded them to their destination on an 802.3 LAN. The obvious limitation to this scheme is that all sources and destinations must be on the 802.3 LAN. A station on the 802.3 LAN may not communicate directly with a station on the FDDI ring. These bridges are only suitable for cases where all end stations were on the 802.3 LAN.

Translating bridges, as the name suggests, translate packets from the format of one kind of LAN to the format of another type of LAN. With a translating bridge, an 802.3 or an 802.5 packet is converted to an FDDI packet or vice versa, and an 802.3 or 802.5 station can communicate directly with an FDDI station. The IEEE 802.1D [802.1D] standard is a standard for translating, encapsulating, learning bridges and should ordinarily be specified when learning bridges are desired with FDDI LANs. This kind of bridge, sometimes called a *transparent bridge* because its operation is ordinarily not visible from normal stations, is the kind ordinarily used when bridging 802.3 LANs to each other (although translation is not needed in this case) and when bridging 802.3 to FDDI or 802.5. It can also be used between FDDI and 802.5 LANs, but, because some polling protocols are fairly common in 802.5 environments, polling of stations that are turned off may be a problem. The IEEE 802.D standard defines the operation of these transparent bridges.

Conversion of 802.5 to FDDI is straightforward, since the format of the frames is quite similar and since both have the same order of bit transmission and significance. Both are "big endian" networks which transmit the most significant bit of an octet first. The 802.3 LAN, however, uses a "little endian" approach in which the least significant bit of an octet is transmitted first. This causes no problem in converting the 48-bit MAC addresses in packet headers when bridging 802.3 to FDDI (or 802.5), because these are sent most significant bit first, even in 802.3. However problems can arise when 48-bit addresses are sent in the information field of packets in the "canonical notation" defined in the IEEE 802.1A draft standard, where the most significant bit (the I/G bit of the address - see Fig. 12), is defined in little-endian fashion as the least significant bit of the most significant octet. In some cases, 802.3 to FDDI bridges may inappropriately invert the bits in information field octets which are used to transport MAC addresses used by higher level protocols. Some bridges attempt to recognize these situations and compensate for them, but this too can cause problems. See reference [LATI] for a more complete discussion of this problem.

The style of bridge most commonly used in 802.5 LANs is the source routing bridge. With source routing bridges, the source and destination stations explicitly participate in the routing through the bridges. The source station inserts the route through the bridges to the information field of the packet. The bridge, in turn, just uses the routing information supplied by the source station to route packets. Bridges recognize source routed packets by examining the I/G bit of the Source Address. A source address I/G (see Fig. 12) value of 1 indicates a packet with bridge routing information, and bridges examine this information to make the decision to route the package to another LAN.

Source routing stations wishing to transmit to a MAC address whose location is unknown usually first send an LLC Test Command to the station. If a response is received, the destination is on the same subnetwork, and source routing through bridges is not required. If no response is received the source issues a search frame (sometimes called an explorer frame). The search frame is forwarded by all source routing bridges, which insert routing information as they do so. Redundant bridge paths are allowed and search frames are repeated on all of these paths, so that the destination will receive a copy of the search frame for every path between it and the source[1]. Destinations then send a response frame to the source indicating the available routes and the source selects the route to be used in further communications.

Source routing has advantages and disadvantages. It requires the stations to participate in the bridging and to know the routing for destinations. This frees bridges from maintaining this information. Source routing allows multiple active paths, while only one path between two nodes is allowed by transparent bridges. Just as packets addressed to unknown stations can flood networks with transparent bridges, the source routing search stations can do the same and this flooding causes practical limits to the size of either type of extended LAN.

Source routing and transparent bridging protocols are not incompatible, and both protocols can be used on one extended LAN. However, stations which do not use the source routing protocol (i.e., most 802.3 stations) cannot send packets to stations through source routing bridges. A revision has been proposed to the 802.5 standard which defines a *Source Routing Transparent (SRT)* bridging protocol. An SRT bridge can bridge a transparent extended LAN with a source routing extended LAN by recognizing packets from the transparent side which are addressed to stations on the source routing side and adding the necessary source routing information. The SRT bridge must build and maintain tables or routing information for the source routing side to make this transformation. A description of such a bridge is found in [LATI].

Two performance metrics are commonly given for bridges:

- *filtering rate,* which is the number of packets which a bridge can examine per second to make a forwarding decision. Some bridges can filter packets at the maximum possible rate for 802.3 or 802.5, and FDDI filtering rates of 400,000 or more packets per second are available.
- *forwarding rates* are the number of packets per second which a bridge can forward from on LAN to another

.

Bridges do not normally support packet segmentation and reassembly, when different subnetworks support different maximum size packets and a packet to be bridged is too large for the destination LAN, however, in some cases, some bridge products may now resegment packets, when they recognize higher level protocols (e.g., TCP) that can accept this. Since FDDI supports a larger packet size than 802.3 and 802.5 LANs, segmentation is only an issue when going from FDDI to the 802 LAN. Bridges do not participate in flow control; there is no protocol for one bridge to tell another bridge or an end station to stop sending because it is congested.

Bridges which connect more than two LANs together in a single bridge are often called *multiport* bridges. They may be a viable alternative to an FDDI backbone LAN in many cases; that is, instead of using an FDDI backbone network and a number of bridges or routers to connect 802.3 or 802.5 LANs to the backbone, it may be viable to connect the separate 802.3 and 802.5

---

1 The spanning tree algorithm may optionally be used to ensure that there is only one path between any two stations for search frames.

LANs together at one multiport bridge and dispense with the backbone altogether, particularly if the relatively long links supported by FDDI are not required. An FDDI backbone does however provide a high bandwidth port for those servers or workstations whose bandwidth requirements alone exceed the capacity of an 802.3 or 802.5 LAN.

At the present time there are few local FDDI to FDDI bridges in use, but, as FDDI LANs grow, they will be used to increase the capacity of networks which become overloaded, and to limit the effect of faults. Opinions of FDDI implementors vary concerning the maximum desirable number of nodes on a ring (see 4.2), but an FDDI ring which grows too big will be subject to unreliable operation and poor performance. FDDI to FDDI bridges will provide both a way to increase the traffic capacity of the FDDI network and to limit the effects of faults and other incidents, such as plugging and unplugging stations, to a small subset of the network

Current bridging techniques evolved after the development of the FDDI MAC standard. Therefore there is no discussion of bridges in the FDDI MAC standard. In general, this presents no problem for transparent bridges, except for the issue of setting the A and C trailing indicators in FDDI frames (see Frame Status on page 18). Since these bridges evolved in the 802.3 community, where there is no link level acknowledgment, transparent bridges generally do not set these indicators. The proposed MAC-2 [MAC-2] standard explicitly addresses bridges in its Annexes A on addressing and Annex C on bridging. Annex A states that the I/G bit in source addresses indicates source routing, as it does in 802.5 LANs. Annex C allows (but does not require) transparent bridges that intend to forward packets they have copied to set the C indicator without setting the A indicator (which is otherwise nonsensical). Annex C requires source routing bridges to set both the A and C indicator on packets they are forwarding.

Bridges forward frames without changing their source addresses, and may forward to an FDDI LAN many packets from different sources while the bridge is holding the FDDI token. The bridge must strip the frames it transmits on the FDDI LAN. It might be very burdensome to strip such frames by their source addresses, so Annex C allows bridges to use several other techniques for stripping, including sending a Void frame with the bridge's source address immediately before releasing the token, and stripping all frames until that void frame is received by the bridge.

8.2.2 Routers

In contrast to bridges, which operate the Data Link layer, routers operate at the Network Layer, which is the layer specifically intended to perform routing and to make protocol conversions. While the address space of bridges is a flat, universal 48-bit address space, router addresses are typically structured or hierarchical. For example the *Internet Protocol (IP)* address normally consists of a 32-bit number divided into a Network ID and a Host ID. OSI uses a hierarchical network level address coded as a character string. In both cases, the structure of the address simplifies routing decisions.

Network level entities of source stations explicitly address network PDUs to the MAC addresses of routers and the routers need not monitor all sub network traffic to discover the packets that they must route, as transparent bridges must. The most commonly used routing protocols are connectionless, that is packets are routed as datagrams, connections are not maintained in the routers, and each packet is independently routed. Although it is normally unusual, with connectionless services packets between the same endpoints may theoretically go by different routes and may therefore may not arrive in the order transmitted.

Routers convert between different data link protocols and resegment transport level PDUs as necessary to accomplish this. These PDUs are reassembled by the destination end point Transport protocol entity.

There are several routing protocols in common use. When routers are used to connect FDDI to other networks, it is important to be certain that the routers support the needed network level protocols. However, multi-protocol router products that simultaneously support several routing protocols are becoming common. Such routers simply recognize the type of network protocol used and process it accordingly.

8.2.3 Bridges vs. Routers

Bridges are generally considered to be faster than routers since the processing they perform is simpler. This is particularly true where the router is, as is often the case, simply an ordinary workstation computer with ports on two subnetworks, which performs routing in addition to other functions. Many workstations with two LAN ports can serve as routers. However, there are dedicated function, high performance routers with much greater performance than a general purpose workstation filling the role of a router. The availability of very high performance RISC processors to perform router functions allows routers to be built with with excellent performance. Special purpose integrated circuits are also being developed to improve the performance of layer three and four processing, and this may also improve router performance.

Routers are also limited to particular routing protocols, while bridges may be transparent to most routing protocols. Therefore a bridge may have an advantage where several routing protocols must be served, or where the protocols in use may not all be known. Multiprotocol routers are becoming common, however, so that one router may handle several routing protocols.

Bridging protocols are more or less automatic. Routers depend on routing tables which typically must be managed and maintained. It may be practical to simply plug a new terminal or a new bridge into a network where bridges are used, but it is usually necessary to explicitly update routing tables to add a new user, station or router. On the other hand, much effort is being devoted to management protocols for routers, to simplify and automate their management.

Bridge protocols limit the size of any extended LAN network. Spanning tree or source routing protocols both generate bridge management packets which, as the network grows, eventually consume too much of the network bandwidth just to manage the spanning tree or to find routes through the network. The spanning tree and its root limit the total bandwidth available in spanning tree bridges. Therefore, while the size of networks joined by routers is nearly unlimited, there are limits the the size of bridged networks. While an extended LAN can cross a continent and join thousands of stations, large wide area networks always rely on routers.

If a PDU on one medium is to large for its destination, routing protocols are usually able to resegment the PDU into smaller PDUs. In general, bridges cannot do this.

The particular application or network operating system that dominates LAN usage will often determine the preferred routing or bridging strategy. LAN operating systems may support several lower level protocol alternatives, but may work best with their "native" protocol. While there is a strong trend to make servers available through a variety of stacks, and, in particular, to support the IP protocol, performance may be better if workstations can reach their primary servers through bridges and use the server's "native" protocol without additional protocols.

There is no one universal choice. Many FDDI networks will use both bridges and routers. Many brouter products provide the functionality of both in a single package. Packets addressed to the router function are routed, while other packets are bridged. Bridges often provide better performance and protocol independence within restricted areas, while routers provide connection to wide area network's. Modular products, which add various bridging and and routing functions as modules to a general purpose frame are becoming common.

8.3 Hubs and Collapsed Backbones.

In the most general sense a "LAN hub" is simply a device that facilitates wiring LANs as stars or trees. As noted in section 6., standards for commercial wiring encourage such wiring, rather than bus or ring wiring within buildings. By this definition an FDDI concentrator is a kind of hub, as are 802.3 multiport repeaters. 802.5 LANs have always been star wired and require that all stations be connected to a *multistation access unit (MAU)*, the equivalent of an FDDI concentrator. With the emergence of the twisted pair 10-Base-T and the fiber optic 10-Base-F technologies, 802.3 LANs are also becoming predominantly star wired. So, while FDDI and 802.3 support dual rings and buses respectively, largely for backbones, most LANs today are predominantly star or tree wired. This facilitates management of the wiring plant and remote management of the network, since the hubs are natural points for remote management.

Simple "workgroup" hubs are often repeaters or concentrators with a modest number of ports that may or may not offer remote management features. There is now a trend to much larger "departmental" or "enterprise" hubs with much more flexibility and functionality. These typically consist of a chassis with a high bandwidth bus of some sort, into which a number of modules may be plugged. These modules may be bridges, concentrators, repeaters, routers, managers, or even, in some cases, LAN servers. One unit can be adapted as required to provide a very wide range of capabilities. There are now a number of vendors offering modular enterprise hub products that can route or bridge FDDI, 802.3 and 802.5 LANs. Products that also offer ATM LAN connection also becoming available.

With simple hub repeaters or with 802.3 bus networks all stations share the 10MHz network bandwidth. Users may contemplate converting to FDDI to provide more bandwidth. However, switching hubs can make the entire 10 Mbps bandwidth of each 10-Base-T link available individually to each attached workstation and provide a high bandwidth bridge to an FDDI backbone to reach station attached to other hubs. In some cases, then, the ability of hubs to switch 10-Base-T links will mean that a more expensive FDDI workstation port is not required. In other cases, the hubs may make it possible to dispense with even an FDDI backbone and the bus of the hub becomes what is sometimes called a "collapsed backbone." Enterprise hubs with backplane bandwidths of more than a Gbit/s also provide a growth path for the time when FDDI is used as a workstation interface and a dual FDDI ring no longer has adequate bandwidth to serve as the backbone.

This is a very fast moving area with new products emerging rapidly. Any user planning on installing an FDDI network should consider the capabilities of hub products and how they may affect the layout of the network, or perhaps even provide a less expensive solution that does not include FDDI.

9. FDDI and Newer Emerging Network Technologies

Although FDDI has been thought of as a backbone LAN technology, FDDI developers are expecting FDDI to soon be a common desktop interface as the rapid increase in the power of personal computers and workstations creates the need for more bandwidth on the desktop. When that happens, FDDI itself may not be fast enough to make a good backbone for other FDDI LANs. Partitioning of FDDI networks with bridges and routers will help to carry the load, just as it has with 802.3 networks, but eventually another, faster technology will be needed for the backbone. Moreover, while FDDI can span large campuses, and is even in use as a Metropolitan Area Network (MAN) in some areas, the token ring is not an ideal medium access method for spanning large metropolitan areas and is a hopeless technology for truly wide area networks. This section is a brief introduction to other networking developments which are now under way and may either compete with FDDI in the future, or provide backbones or wide area connectivity for FDDI in the future.

9.1 HIPPI

HIPPI is a high-performance point-to-point interface which supports very high performance computing. HIPPI may be simplex or duplex. For duplex operation HIPPI uses two 32-(or 64-) bit wide unidirectional, point-to-point electrical paths to obtain an 800 (or 1600) Mbps transfer rate. The maximum distance supported is 25 m with copper twisted-pair cables. A 32-(or 64-) bit *word* is transferred on every 40 ns cycle; words are grouped into *bursts* and bursts are grouped into *packets*. Four (or eight) parity lines provide byte parity for each word, while a vertical parity word protects each burst. Bursts contain 256 words (or fewer). Packets consist of 1 or more bursts; the first burst contains a 64-bit HIPPI-FP packet header.

There is a mapping for IEEE 802.2 LLC frames (including FDDI frames) on HIPPI. There is another mapping for IPI-3 command sets over HIPPI. A protocol defining memory read, write and lock operations is provided.

HIPPI supports both circuit and packet switching. Physical level switches provide circuit switched connections between several devices. A 32-bit I-field is provided during connection establishment which can be used to control the switch. Although existing HIPPI switches are implemented on a circuit switching basis, the switching can be on a packet by packet basis, therefore the performance is rather like a packet packet switch, provided at least that the distance between the switch and the station is short. That is necessarily the case with the 50 m maximum parallel interface defined in the standard. Although not described in the standard, HIPPI serial extenders are commercially available and can extend HIPPI links to several km.

HIPPI consists of the following standards as illustrated in Figure 28:

- X3.183-1991, *High-Performance Parallel Interface - Mechanical, Electrical and Signaling Protocol Specification (HIPPI-PH)*, DIS 11518-1, which specifies a 32- or 64-bit wide parallel electrical bus.
- X3.210-199X, *High-Performance Parallel Interface - Framing Protocol (HIPPI-FP)*, CD 11518-2, which specifies the framing of words into bursts and packets.;
- X3.218-199x, *High-Performance Parallel Interface - Encapsulation of ISO 8802-2 (IEEE Std. 802.2) Logical Link Control Protocol Data Units (HIPPI-LE)*,CD 11518-3, which specifies how normal LAN LLC packets are carried on HIPPI.;
- *High-Performance Parallel Interface - Mapping to IPI-3 Command Sets (HIPPI/IPI-3)*,11518-4, which specifies how the IPI-3 command sets for disk and tape drives are

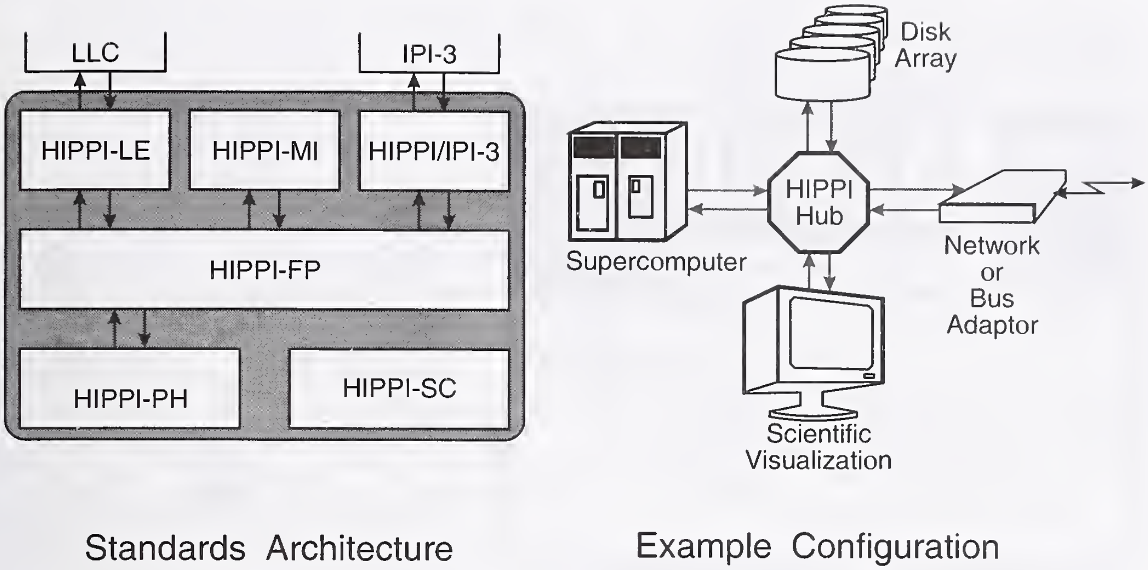

Figure 28 - HIPPI Overview .

used with HIPPI;

- X3.222-198x, *High-Performance Parallel Interface - Switch Control (HIPPI-SC)*, CD 11518-6, which specifies how HIPPI switches are controlled.

HIPPI is philosophically a hybrid between a LAN and a conventional computer I/O port. Although a parallel electrical interface, HIPPI supports a packet structure and can use the same LLC used by LANs, including FDDI. HIPPI also contains specific provisions for running peripheral device protocols and is used for attaching high performance peripherals to computer systems. HIPPI products are now available and HIPPI is coming into wide current use on very high performance scientific computer systems, particularly for applications which demand very high point to point bandwidth, such as parallel disk arrays and high resolution scientific visualization. HIPPI to FDDI bridges are available.

9.2 Fibre Channel

The Fibre Channel provides a transport vehicle for the Intelligent Peripheral Interface Generic Command Sets (IPI-3) and packetized Small Computer System Interface (SCSI) command sets for storage devices, as well as the HIPPI data framing. Fibre Channel defines an interface which is used with a Switching Matrix to interconnect many ports which may be spread over a campus or larger area. Figure 29 provides an overview of the structure of Fibre Channel.

The basic orientation of Fibre Channel is to provide a replacement for conventional I/O channels, which extend their distance and allows many host computers to share storage devices. Particular "hooks" are provided to facilitate this. In particular, I/O buses typically provide hardware signals to indicate whether a transfer is an I/O command, a device reporting its status, or an actual data transfer. The channel hardware then steers the transferred data to an appropriate host facility (e.g., a DMA controller for data transfers, an I/O driver for peripheral device status, etc.). Conventional LANs have no such provision for steering data in hardware and, when a packet is received, it is first parsed by network protocol software before it is delivered to the appropriate process in the host. Fibre Channel includes a one byte Type field in the frame which allow the Fibre Channel port hardware in the host to steer received data to the appropri-

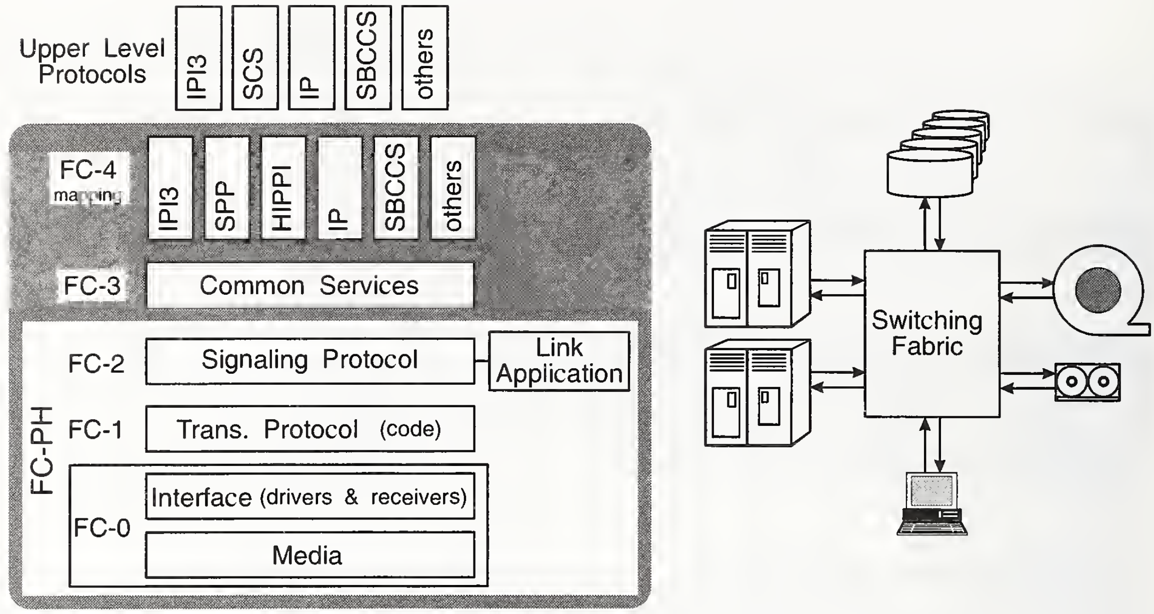

Fibre Channel Structure   Example Configuration

Figure 29 - Fibre Channel Overview.

ate process, without invoking general purpose network driver software. Code points are provided for specific peripheral device and network protocols, including the Intelligent Peripheral Interface (IPI) and Small Computer System Interface (SCSI) as well as the Internet Protocol. This facilitates including specialized processors to support efficient I/O without interrupting the host.

As depicted in Figure 29, the FC interface has four layers:

- FC-0 defines the media, connectors, transmitters and receivers. Rates of 12, 25, 50, and 100 Mbyte/s are supported. The various FC-0 media interfaces are summarized in Table 7.
- FC-1 defines the transmission protocol including a DC balanced 8 of 10 code;
- FC-2 defines variable length (up to 2112 information bytes) frames with 24-bit addresses. A signaling protocol provides connections through the fabric and flow control to ports;
- FC-3 defines common services which support FC-4;
- FC-4 defines the mapping of higher level device command sets or communications services (*e. g.* IPI or SCSI-3) onto FC.

Fibre Channel has three classes of service:

- Class 1, which provides a dedicated connection with guaranteed delivery and bandwidth. Frames are delivered in the order transmitted. This is effectively a circuit switched service and may be implemented with space division multiplex switches. It is expected to be useful for applications such as video, where a dedicated connection may exist for a long period of time. Flow control is end to end.
- Class 2, which is a packet switched service providing a virtual connection with guar-

Table 7 - Summary of Fibre Channel Physical Layer

Designation	Transmitter	Medium	Distance
850 Mbps (100 Mbytes/s nominal)			
100-SM-LL-L	1300 nm laser	single mode fibre	2 m - 10 km
100-SM-LL-I	1300 nm laser	single mode fibre	2 m - 2 km
100-TV-EL-S	electrical (ECL levels)	75 Ω video coax	0 - 25 m
100_MI-EL-S	electrical (ECL levels)	75 Ω mini coax	0 - 10 m
425 Mbps (50 Mbytes/s nominal)			
50-SM-LL-L	1300 nm laser	single mode fiber	2 m - 10 km
50-M5-SL-I	780 nm "CD type" laser	50/125 μm multimode fibre	2 m - 1 km
50-TV-EL-S	electrical (ECL levels)	75 Ω video coax	0 - 50 m
50_MI-EL-S	electrical (ECL levels)	75 Ω mini coax	0 - 25 m
212.5 Mbps (25 Mbytes/s nominal)			
25-SM-LL-L	1300 nm laser	single mode fibre	2 m - 10 km
25-SM-LL-I	1300 nm laser	single mode fibre	2 m - 2 km
25-M5-SL-I	780 nm "CD type" laser	50/125 μm multimode fibre	2 m - 2 km
25 -M6-LE-I	1300 nm LED	62.5/125 μm multimode fibre	0 m - 1 km
25-TV-EL-S	electrical (ECL levels)	75 Ω video coax	0 - 75 m
25-MI-EL-S	electrical (ECL levels)	75 Ω mini coax	0 - 30 m
25-TP-EL-S	electrical (ECL levels)	150 W shielded twisted pair	0 - 50 m
106.25 Mbps (12.5 Mbytes/s nominal)			
12-M6-LE-I	1300 nm LED	62.5/125 μm multimode fibre	0 - 1 km
12-TV-EL-S	electrical (ECL levels)	75 Ω video coax	0 - 100 m
12-MI-EL-S	electrical (ECL levels)	75 Ω mini coax	0 - 40 m
12-TP-EL-S	electrical (ECL levels)	150 Ω shielded twisted pair	0 - 100 m

anteed delivery. However frames may be delivered out of order. Bandwidth is not guaranteed. Flow control is provided between the switching fabric and ports. The intended use for this service is I/O transfers to intelligent storage devices.

- Class 3, which is datagram service. Neither delivery or sequencing of frames is guaranteed. Most higher level protocols used over LANs provide end to end recovery for out of sequence or lost packets, and this service is suitable for most conventional LAN applications.

The planned set of FC standards is:

- X3.230-199x, *Fibre Channel Physical and Signaling Interface (FC-PH)* (includes FC-0, FC-1 & FC-2). This is the only FC standard which is near completion at this time;
- *Fibre Channel Enhanced Physical and Signaling Interface (FC-EP)*, which will provide for an enhanced FC physical level;
- *Fibre Channel Fabric Generic Requirements (FC-FG)*, which will provide additional definition of the overall switching fabric;
- *Fibre Channel Cross Point Switched Fabric Requirements (FC-XS)*, which will de-

fine the characteristics of FC cross point switches;

- *Fibre Channel Single-Byte Command Code Sets (FC-SB)*, which will define how common storage peripheral bus command sets (e.g., SCSI & IPI3)are carried over the FC.;
- *Fibre Channel Implementation Guide (FC-IG)*, which will provide guidance on the implementation of Fibre Channel networks and the relationship between the physical and logical functions of the Fibre Channel;
- *Fibre Channel Mapping to HIPPI-FP (FC-FP)*, which will define the connection of the FC to HIPPI networks;
- *Fibre Channel Internet Protocol (FC-IP)*, which will define how Internet Packets are handled by the Fibre Channel;
- *Fibre Channel Low Cost Topologies (FC-LT)*, which will be based on FC-PH but provide a low cost loop (ring) topology.

Fibre Channel defines an almost bewildering array of rates and media choices, from 106.25 to 850 Mbps. Table 7 summarizes the FC media alternatives.

The basic FC-PH1 standard has been developed and products that use it are starting to appear in the marketplace. It apparently will be supported by major mainframe computers, storage peripherals and workstations. FC looks inward more than outward, that is it caters to the internal needs of distributed computer systems to communicate with each other and their storage peripherals, rather than to communicate with wide area networks and generalized network protocols. Nevertheless, Fibre Channel provides a general purpose datagram LAN, which supports very high bandwidths and could well function as a backbone for lower speed LANs, including FDDI, and does provide specific features for connection to LANs (including recognition of 48-bit universal addresses), wide area networks (including recognition of 60-bit CCITT addresses) as well as support for the Internet Protocol. It remains to be seen if it will be largely relegated to a "backend" storage oriented role, of it will widely encompass LAN applications as well.

In some ways the approach of FC and ATM LANs (see 9.5)are similar; that is they both rely on a fabric of switches. In the case of FC, these switches support variable length packets, while the ATM LANs will switch small, fixed size cells. The ATM network simplifies the switching somewhat by using a fixed size cell, which puts the burden of disassembling and reassembling larger packets on the stations. Where ATM LANs look outward to Broadband ISDN and wide area networks, FC looks inward, to the LAN operating system and its storage devices.

There is another similarity between FC and ATM LANs. While both are emerging technologies, and products for both are starting to appear, the station to network interface are their best defined parts at present. At the current state of the evolution of FC, it may be that network station adapters from different vendors will work with a switching fabric from other vendors, but it is unlikely that a working switching fabric itself can be assembled in mix and match fashion from switches of different vendors. This is in contrast to the more mature distributed LAN technologies such as FDDI, where the station is not distinguished from the switching fabric, and stations, bridges, concentrators, and routers from different vendors generally will interoperate.

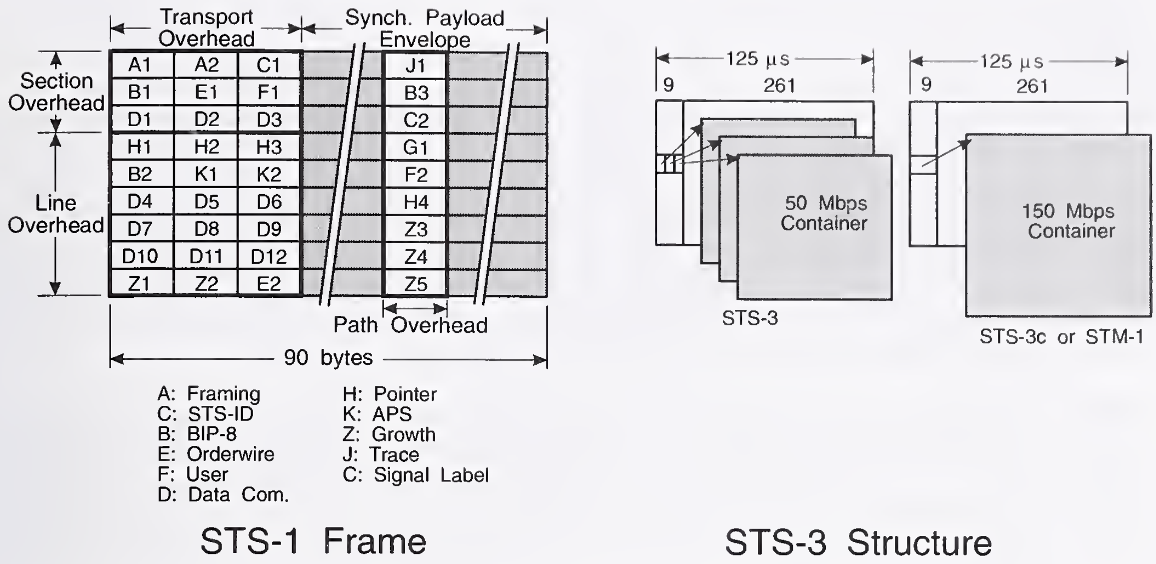

Figure 30 - SONET Overview.

9.3 SONET

The *Synchronous Optical NETwork (SONET)* [T1.105] standard defines a *Time Division Multiplexing (TDM)* standard for future public network trunk circuits. The CCITT has adopted SONET internationally, while changing the name to the *Synchronous Digital Hierarchy (SDH)* [G.707, G.708]. SONET (or SDH) is becoming the principle long distance optical transmission standard worldwide. In addition, the physical interface provided by public networks to the future B_ISDN services (See 9.5) will be based on SONET.

The basic building block for SONET is the *Synchronous Transport Signal (STS)* frame illustrated in Figure 30. The frame can be viewed as a nine by ninety byte bloc. The first three columns of 9 bytes each is the *Transport Overhead.* The remaining 87 columns of 9 bytes each are the *Synchronous Payload Envelope (SPE).* In some cases another column of 9 bytes of overhead, called the *Path Overhead,* is carried in the Synchronous Payload.

SONET is synchronized to an 8 kHz clock. The terminology *STS-N* is used to refer to the signal that transmits N times 8000 frames per second, while the the terminology *Optical Carrier N (OC-N)* designates the corresponding optical signal. Another way of expressing this is to observe that an STS-N or an OC-N signal has N 9 by 90 byte frames, each 125 μs. An STS-1 signal has a line rate of one times 8000 frames per second, times 90 columns per frame, times 9 rows per frame, times 8 bits per byte, giving a rate of 51.840 Mbps or a payload of about 50 Mbps. An STS-3 signal is three times that or a line rate of 155.520 Mbps and a payload of about 150 Mbps, as illustrated in Figure 30. Table 8 shows the rates that have been defined.

Table 8 - Standard OC Rates

OC level	Line Rate (Mbps)
OC-1	51.840
OC-3	155.520
OC-9	466.560
OC-12	622.080
OC-18	933.120
OC-24	1244.160
OC-36	1866.240
OC-48	2488.420

As illustrated in Figure 30, the payload of an STS-N signal can be packaged as an STS-N carrier with N separate 50 M containers, or as an STS-Nc (STS-3c is also called a *Synchronous*

Transport Module (STM) and STS-3c is equivalent to STM-1) carrier with a single N times 50 Mbps payload.

The synchronous containers are plesiochronous to the frames, that is, while they are also driven by an 8 kHz clock, it is not necessarily the same clock as the frame clock, and the containers may drift over long periods of time with respect to the frame boundaries. Pointers in the Transport Overhead point to the start of containers. The Path Overhead also provides for positive and negative "stuff bytes" which allows SONET to carry a signal that originates in a TDM network with one clock, independent of the SONET clock, and transport it to another TDM network with yet another independent clock, without introducing unnecessary clock conversions (which can result in either the insertion or loss of data due to the different clock rates).

SONET and SDH define an extensive multiplexing hierarchy of "tributaries" to make up the containers. This allows lower rate and existing carriers, such as the North American Digital Hierarchy of DS-1 (1.536 Mbps), DS-2 (6.144 Mbps) and DS-3 (43.008 Mbps) to be carried efficiently in the SONET payload. While SONET OC-N connections will eventually be available to users, there will be a long transition where existing DS trunks are multiplexed over SONET channels in the network.

The overhead bytes of SONET, shown in Figure 30, is divided into the Transport Overhead, which is in turn subdivided into the Section Overhead and Line Overhead, and the Path Overhead. Figure 31 illustrates the meaning of these terms. A Section is a segment between any two network equipments including repeaters. Among other functions for maintenance and management, it contains the framing information. The Line Overhead allows communication of necessary information between network equipments, such as terminals, multiplexors and switches, that are more complex than simple repeaters. It includes the pointers for frequency justification. The path overhead is contained only once in each Synchronous Payload Envelope and is terminal to terminal

There is a project in X3T9.5 to define a SONET Physical Mapping (SPM)14 standard for mapping the 125 Mbps FDDI line code bit rate onto an STS-3c payload envelope (see 2.3.5). This project has been delayed somewhat by difficulty in understanding the complexities of jitter in SONET, particularly concerning its use with the FDDI-II isochronous services. When completed, SPM will provide a means for carrying FDDI links over SONET, which may allow the use of services available from network service providers, when it would otherwise be necessary to run dedicated fiber for this purpose.

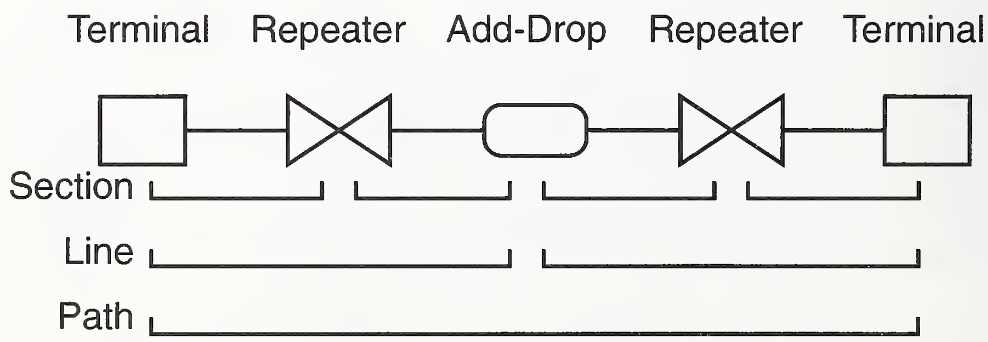

Figure 31 - The SONET Path.

9.4 Frame Relay

Frame relay is a new, standardized packet switching service now being offered by communications carriers. In many ways it is a streamlined X.25 packet service, the major difference being that X.25 provides hop by hop flow control and and recovery from transmission errors, while frame relay has simplified flow control and does not recover transmission errors. Error recovery with frame relay is end to end, which fits well with modern transport layer protocols. The reduced processing in switches allows frame relay switches to be faster than their X.25 predecessors.

Figure 32 illustrates the format of the frame relay packet which is similar to the format of the LAPB and LAPD data link frames used with X.25. Like them it uses a unique flag character and "bit stuffing" to delineate variable length frames. Frame relay is defined as an ISDN bearer service but is also available by dedicated connection to Frame Relay networks through 56 kbps and DS1 circuits, without the ISDN service. DS3 based services may eventually be offered.

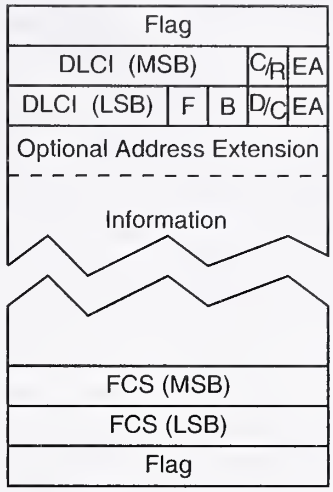

Flag: "01111110" pattern frame delimiter
DLCI: Data Link Connection ID (address)
F: Forward congestionilndication
B: Backward congestion indication
EA: Address field extension bit
C/R: Command/Response
D/C: DLCI or Control Indication
FCS: Frame Control Sequence

Figure 32 - Frame Relay Packet.

At the present time frame relay is usually offered in the form of fixed virtual circuits. Dynamically switched frame relay services may also become available in time. The major market for frame relay is expected to be connecting LANs together. In many cases it will be less expensive to use a frame relay virtual circuit to connect two LANs than to lease a dedicated circuit for that purpose. The use of frame relay virtual circuits with LAN bridges or routers is illustrated in Figure 33 below.

The major competitor of frame relay for interconnecting LANs is likely to be the SMDS service described in 9.6. Some carriers have both SMDS and frame relay service offerings and choice of the most advantageous service may simply depend on the tariffs.

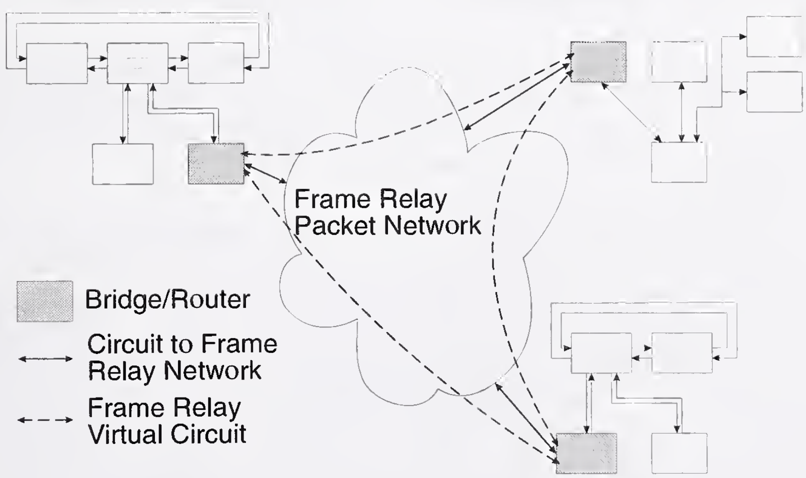

Figure 33 - Frame Relay Networks.

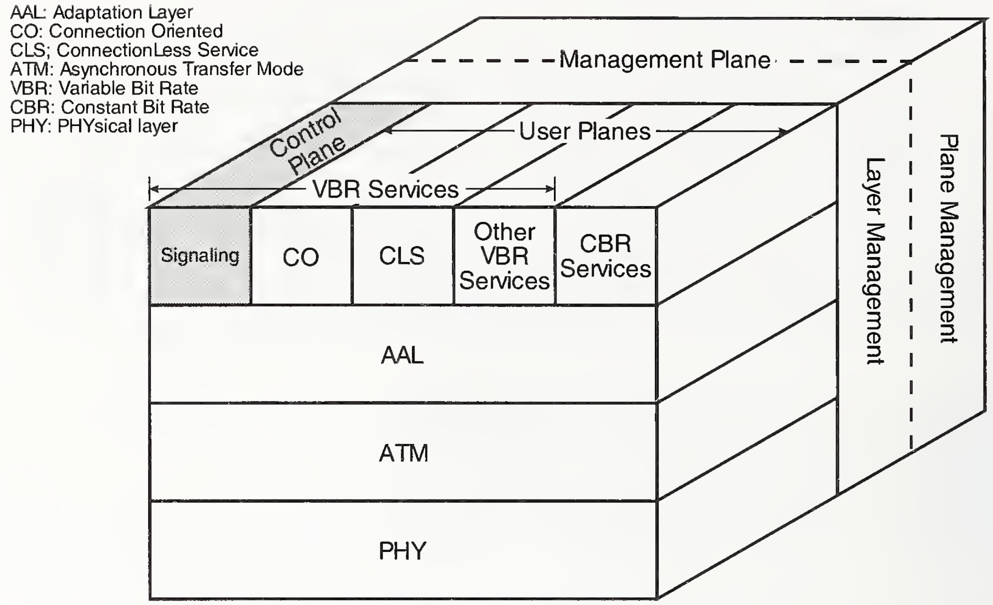

Figure 34 - B-ISDN Protocol Reference Model.

9.5 ATM LANs and B-ISDN

The concept of cell switching was proposed in the mid 1980s as a technology for providing high bandwidth integrated services switching, that is able to switch both bursty data traffic and constant rate traffic such as conventional voice or video [TURN]. The concept was rapidly adopted for Broadband-ISDN (B-ISDN), and the term *Asynchronous Transfer Mode (ATM)* came to be used for the cell switching concept of B-ISDN. The cells used are called *ATM cells*, and are 53 bytes each.

Figure 34 illustrates the protocol reference model used by B-ISDN [T1S1]. Unlike the OSI model, it is a three dimensional model, with three "planes:" a User Plane, a Control Plane, and Management Plane, which is subdivided into Layer Management and Plane Management. Somewhat confusingly, the planes are not depicted entirely as all parallel to one another, rather the Control Plane and the User Plane are horizontal slices, which ride atop the basic communications protocols, while the Management Plane is a vertical plane behind them.

Data transfer is provided by three layers. At the bottom is the Physical layer, which generally will be based on SONET. Above this is the ATM layer, which deals with the transmission and switching of 53 octet ATM cells. Above that is an *Adaptation Layer (AAL),* which adapts the basic ATM transport mechanism to the services above them.

The services are broken down into *Variable Bit Rate (VBR)* services, *Constant Bit Rate (CBR)* services and *signaling*. The VBR services are the normal data services and both *connectionless (CLS)* and *connection oriented (CO)* services are included. The CBR services are provided to support applications like conventional voice and video, which operate a a constant rate. *Signaling* refers to communications between the terminal and the network which request and control the services of the network. The analogy here is that of a telephone network, where a call is dialed to establish a circuit, rather than a LAN, where all stations simply address a packet to a destination address and send it. In the signaling protocol a B-ISDN terminal requests from the network a call setup with another terminal giving, in effect, the directory number of the destination, using the telephone network addressing system of ISDN [I.330]. A call is established by the network and both the "calling" and "called" parties are given a kind of temporary address, which will be described in the discussion of the ATM layer below, to be used for the duration of the call.

In general, the B-ISDN data services will provide higher layers with PDUs that apparently are conventional variable length packets. The AAL converts between these packets and fixed ATM cells. Four AALs have been adopted by CCITT. A fifth AAL, called AAL 5 has been proposed for packet data applications, particularly ATM LANs, and it is likely that it also will be adopted by the CCITT. AAL 5 is illustrated in Figure 35. It differs from the other packet data AALs which use an additional 2 octets of header and 2 octets of trailer from the payload of every cell, leaving a net payload of only 44 octets. AAL 5, has 48 octets of payload in each ATM cell. AAL 5 accepts PDUs of up to 65535 octets from higher layers, adds a trailer to the end which includes a length field and the same 32-bit CRC used in FDDI, and sends this segmented into 48 octet cell payloads. At the destination end AAL 5 reassembles the higher layer PDU from the separate cell payloads.

It has been shown that the 32-bit CRC will detect any cell misordering [WANG], and there are no cell sequence numbers, and no payload checksum in individual cells. Each cell has a header with a *Virtual Path Indicator (VPI)* and a *Virtual Circuit Indicator (VCI)*. In combination these fields are used by switches to route the packet through the switch, and they may be changed in each switch. The ATM LAN or the wideband B-ISDN network will be a virtual circuit network. The signaling protocol is used to set up a virtual connection between two stations and the cells will be routed along this circuit, which is identified by the VPI/VCI.

The ATM layer is primarily concerned with getting the 53 octet cells through the network. During call setup, a path through the network switches is determined for the call. A switch receives a cell on an input port and, from its VPI/VCI determines the output port (usually a fast hardware lookup). If we ignore the possible complexities caused by congestion at switches, output blocking, priorities and bandwidth reservation to accommodate constant rate services, ATM switching is simple and fast, and a wide variety of switch architectures has been discussed in the literature [NEWM], [ENG].

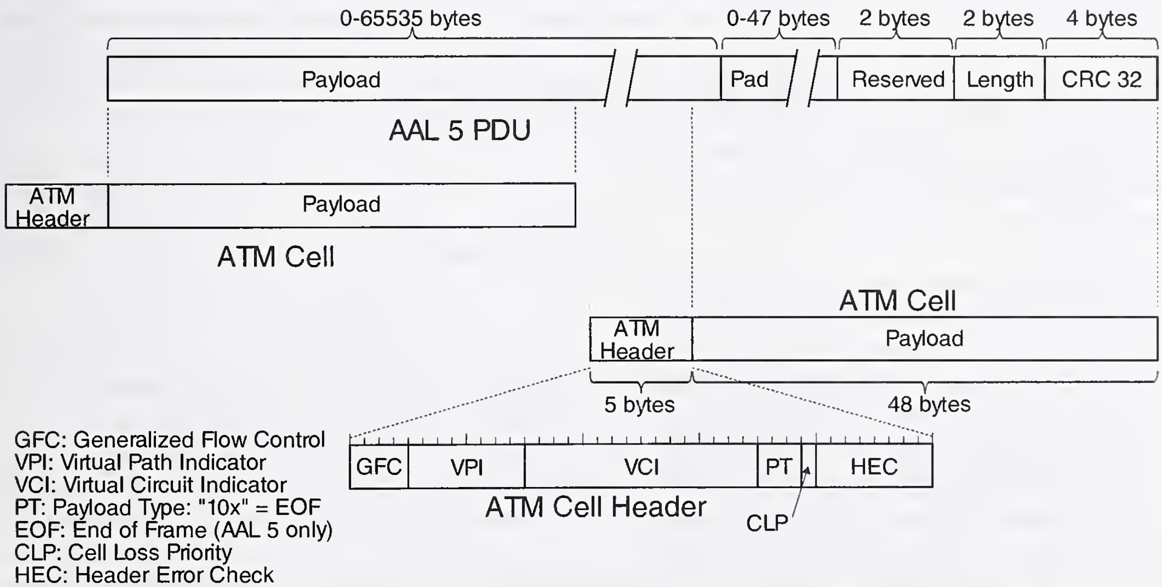

Figure 35 - Segmentation of AAL5 Packets into ATM Cells.

At the bottom of the stack is the physical layer, probably, in many cases, based on the use of SONET (see 9.3). At the next layer up, the ATM layer, data is transferred and switched in 53 byte ATM cells. The AAL converts higher layer PDUs to 53 octet ATM cells and back to higher layer PDUs. The higher layers might include data services such as the OSI stack above the data link layer or the Internet stack, as well as "constant bit rate" services.

B-ISDN was conceived as a wide area network. There has recently been much commercial activity, and informal standards efforts to adapt the ATM technology of B-ISDN to what have come to be called ATM LANs. This has four principle motivations:

- The switch based ATM architecture is inherently scalable, by replacing or augmenting switches, without changing stations (terminals);
- Stations do not share the link bandwidth as they do in FDDI in other LANs;
- It is believed that ATM technology can provide more effective constant bit rate services than conventional LANs, enabling multimedia applications that may require them;
- Integration with expected B-ISDN wide area networks will be simplified.

There are some complications. The telephone network like B-ISDN call setup procedure, and the ISDN I.30 addressing are foreign to LANs. Similarly, the basic broadcast paradigm of LANs and the ability to broadcast messages to Group Addresses, or to the entire network, which are integral to some LAN protocols, do not come naturally to ATM networks. If ATM LANs are to succeed, then they must transparently run the vast body of software and protocols which are now used over conventional LANs, particularly the important LAN Operating Systems.

Figure 36 (a) illustrates the data transport part of the B-ISDN reference model, which, approximately corresponds to layers 1 and 2 of the OSI model. Figure 36 (b) illustrates a LAN adaptation of this model, which some developers of ATM LANs are following. In this model an ATM MAC layer is inserted, which is responsible for providing a normal LAN LLC layer, with the services it expects from a LAN. The ATM LAN then looks like FDDI or one of the IEEE 802 LANs to the LLC. In this scheme every ATM LAN station has a 48-bit universal LAN type MAC address and a CCITT I.330 type address, and the ATM MAC layer maintains a table associating universal MAC addresses, I.330 address and the VPI/VCI of active connections. Protocols and directories are required to find the I.330 address for a new MAC address and to identify "virtual LANs," which support broadcast messages to subsets of the network stations.

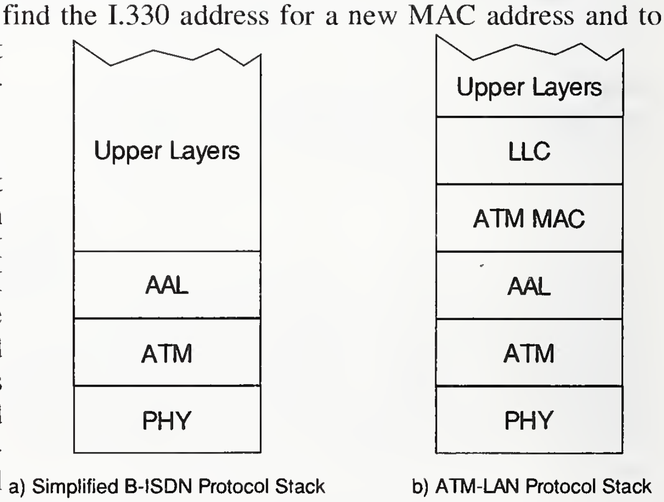

Figure 36 Reference Model for ATM LANs.

Figure 37 illustrates the path of a packet or packets which originate at station 1 on an ATM LAN, passes through ATM switch 2, a router, 3, to a wide area ATM network and switches 4, 5, and 6. At the destination end is another ATM LAN and ATM switch 7 before the packet reaches its destination, station 8. This is a hybrid arrangement, which is likely to be common, where conventional routers as well as ATM networks are involved in the end-to-end transfer. In this particular case, router 3 is used to tie both an ATM

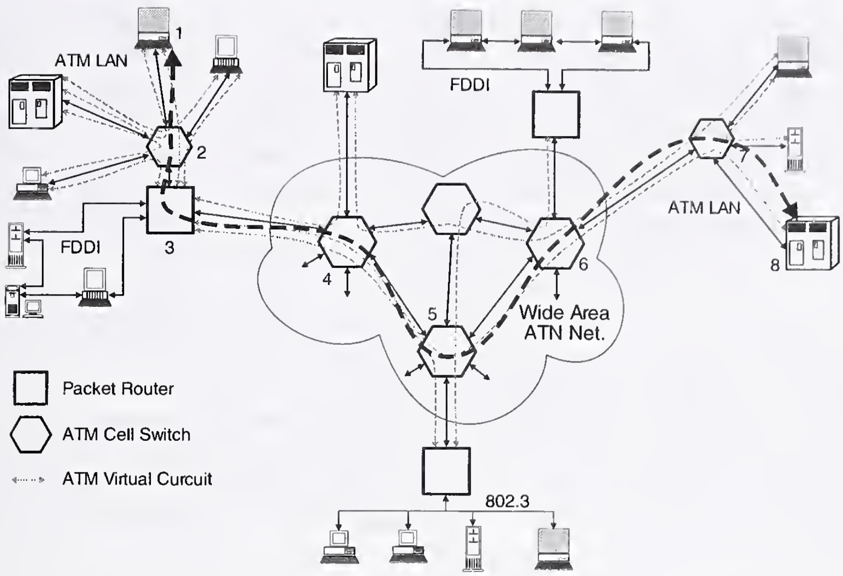

Figure 37 - A Virtual Circuit Through an ATM Network.

LAN and a conventional LAN to each other and to the wide area ATM network. Two ATM virtual circuits are traversed, one from station 1 to router 3, and one from router 3 to station 8.

Figure 38 illustrates the end-to-end protocol stacks. Note that the LLC PDUs are segmented into cells at station 1 and these cells are reassembled and then resegmented at router 3. Obviously this involves extra delay, moreover router processing is comparatively complex, and usually done in software. But, even if end-to-end ATM eventually becomes the rule, there will be a long transition period where hybrids like the one illustrated remain. The short stack of the ATM switches is indicative of the relatively minimal processing performed at each ATM switch. The cell processing in ATM switches will be done entirely in hardware and be very fast. Congestion, however, can still cause either long cell queuing delays or lost packets in switches and congestion control, particularly for constant bit rate traffic, will be a great challenge in ATM networks [HONG].

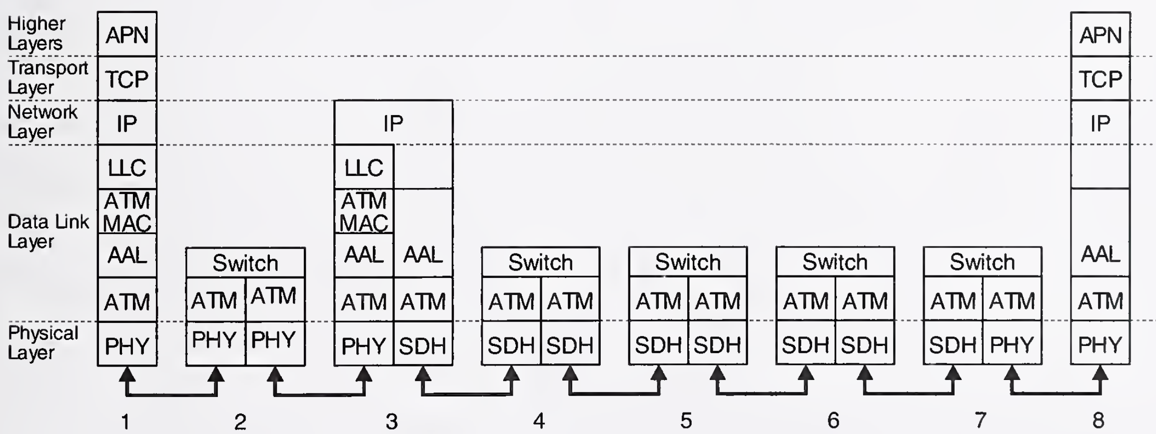

Figure 38 - End-to-End ATM Protocol Layers.

B-ISDN and ATM LAN standards for signaling are not yet complete. Nor are there yet formal standards for the protocols for an ATM MAC, although products are becoming available that implement this approach. B-ISDN standards for the physical layer will be based on SONET. In the ATM LAN arena there may be other alternatives, since the market is quite sensitive to component cost. Some early vendors of ATM LAN products are using the FDDI PHY and PMD components to build ATM data links at the FDDI rate of 100 Mbps. This takes advantage of the comparatively mature FDDI component market, and, as FDDI Twisted Pair is standardized, will more or less automatically provide a low cost physical layer.

ATM LANs are sometimes viewed as a competitor to FDDI. The geometric growth of the processing power of desktop computers makes it certain that a higher bandwidth solution than FDDI will ultimately be needed for backbone networks, it is doubtful that there now exist many overloaded FDDI backbones. FDDI has matured rapidly and there now exist only a few applications or computer systems whose normal network requirements exceed FDDI bandwidths. FDDI vendors have aggressive cost goals aimed at making FDDI a reality for workstations and high end PCs in the relatively near term. FDDI now allows interoperable multivendor solutions and the FDDI standards are now all but complete.

Much remains to be done to complete the standards for ATM LANs and B_ISDN, particularly in the areas of signaling, terminal interfaces and equivalents to LAN MAC functionality. FDDI probably has a fundamental cost advantage; an FDDI concentrator is a simpler device than an ATM switch, and there are costs and inefficiencies in the segmentation of packets into ATM cells. But ATM stations do not share the bandwidth of links with other stations. It seems likely that FDDI will be well established and not easily displaced, before interoperable, standardized, multivendor ATM LANs and B-ISDN become a widespread marketplace reality, similar to FDDI today. Just as the existence of the lower bandwidth 802.3 and 802.5 technologies provided a natural market for FDDI as a backbone and upgrade path, existing FDDI LANs will do the same for ATM based networks. FDDI has strengthened the growth of 802.3 and 802.5 LANs, by providing a backbone and a growth path. ATM LANs may do the same for FDDI, and further strengthen the LAN market in general by providing a natural link between the local environment and the wide area network.

9.6 DQDB and SMDS

The *Distributed Queue Dual Bus (DQDB)* standard is specified in [802.6]. DQDB is a Metropolitan Area Network (MAN) standard, and its functionality somewhat overlaps that of FDDI, in that FDDI has also been adopted to MAN applications. DQDB provides the customer interface for the *Switched Multimegabit Data Service (SMDS)* provided by several North American public carriers. DQDB is also of interest because the slots are generally compatible (there are differences in some of the control fields, but the overall slot size and payload are the same) with the emerging B-ISDN standards (see 9.5 above) which use 53-byte ATM cells, and SMDS can be viewed as the precursor or the first implementation of B-ISDN, now becoming available from North American Public networks service providers.

DQDB provides a means of attaching users to the SMDS service. While ATM LANs stress cell switches, DQDB provides a bus network for cells, which can be used to attach a number of user stations to a cell switching network, or it can serve as a stand alone network. Somewhat similar concepts are under consideration in B-ISDN, under the name of *generic flow control*. Figure 39 (a) illustrates the DQDB MAN open bus topology. Two buses, x and y, propagate through stations in opposite directions. At the head of each bus is a head-end station which originates empty *slots* (equivalent to ATM cells) down the bus. The simplest protocol one might imagine on such a dual bus is to allow stations to simply fill any empty slot passing through it. Such a protocol is obviously quite unfair, giving an advantage to stations immediately downstream of each head-end. The queue arbitrated DQDB medium access method is a method of improving the fairness of the use the first empty slot protocol.

Each DQDB slot has a header with a Busy bit and 3 Request bits. When the Busy bit is set, the

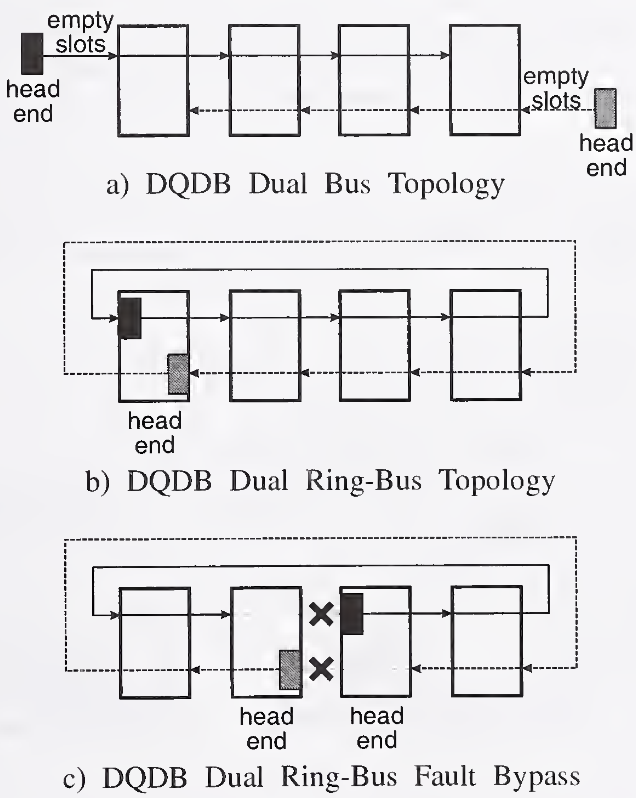

Figure 39 - DQDB topology.

slot is used, and each of the Request bits corresponds to a different priority level. DQDB stations request slots on bus x by setting a request bit of the appropriate priority on bus y, and vice-versa. Stations maintain counters for each bus which are incremented by requests of an equal or higher priority on the opposite bus and decremented by empty slots on the original bus. When the counter is zero (indicating no pending downstream requests) a station may use an empty slot and mark it busy. Transmission on bus y is the mirror image. A station monitors Busy bits on bus y, and Request bits on bus x.

While the DQDB protocol depends on the arrangement of stations in a nonbranching bus, it is possible to close that bus back on itself, with both head-ends at the same station, forming an open dual ring as shown in Figure 39 (b). This will probably be a common way of configuring a DQDB MAN, since, as illustrated in Figure 39 (c), it allows a broken cable or station to be bypassed, with stations on either side of the break assuming the role of the head-ends. The bandwidth and delay of the network are not affected by the reconfiguration and its fault recovery properties are similar to the FDDI.

DQDB provides for three distinct classes of services: an *Isochronous* Data service, a *Connection Oriented* Data Service, and *MAC* to *Logical Link Control (LLC)* Data Service. The connection oriented and MAC to LLC data services depend upon the *Queue Arbitrated (QA)* functions of DQDB. The isochronous service, which is not fully defined in the existing standard, will use the *Pre-Arbitrated (PA)* functions. The pre-arbitrated and queue arbitrated functions use slots which are marked by the network head end station as either PA or QA.

The DQDB MAC to LLC data service is the only one for which a convergence function has been fully defined in the current standard. The MAC to LLC data service is similar to the LLC service of LANs and closely corresponds to the service provided by FDDI to LLC. For compatibility with existing LANs, the MAC to LLC data service defines a multislot protocol data unit structure, similar to the packets of conventional LANs such as FDDI. This is an analog of

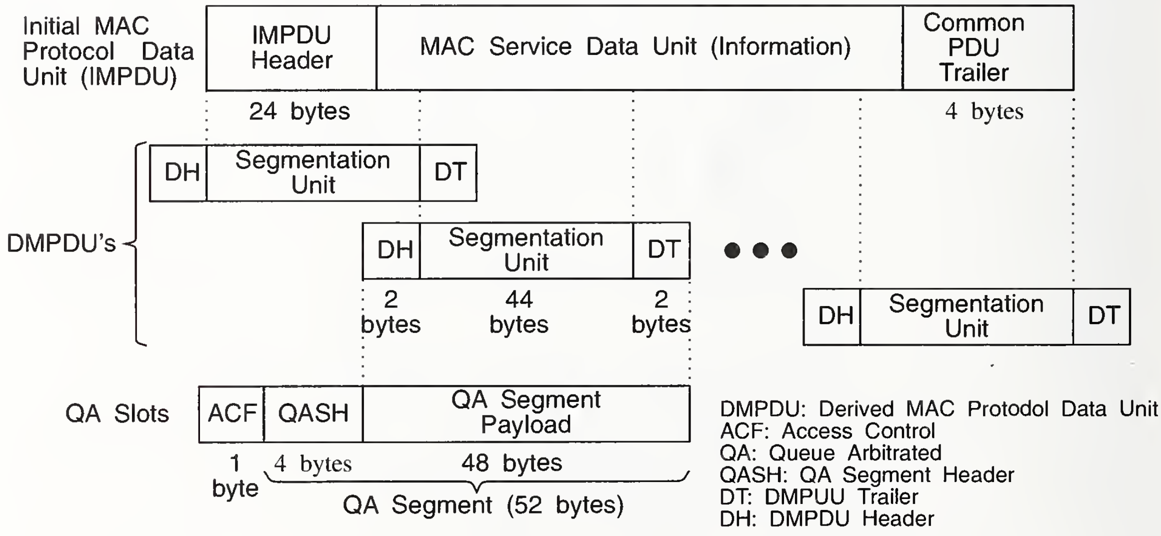

Figure 40 - Mapping Frames to Cells in DQDB.

the B-ISDN Adaptation Layer. The DQDB *Initial MAC Protocol Data Unit (IMPDU)*, which we will call a packet, is illustrated in Figure 40 and is similar to the frames of FDDI LANs. The IMPDU includes a header, trailer and Information field. The Information field may be from 0 to 9188 bytes long. The header is 28 bytes long and contains both a source and destination address, as well as other control fields. An optional 32-bit Cyclic Redundancy Check (CRC) may be contained in the trailer. With the CRC the trailer is 8 bytes, without it the trailer is 4 bytes. Since only 44 of 53 bytes in each slot carry payload there is a significant overhead and the maximum load carried cannot exceed 83.02 % of the channel capacity.

Tariffed SMDS/DQDB services are becoming available using DS-1 and DS3 carriers for access. Services may eventually be eventually introduced at SONET OC3 SONET (155 Mbps). The DS-1 access service provides an effective user bandwidth of 1.17 Mbps, while DS-3 access services are tariffed at user bandwidths from 4 to 34 Mbps.

While some independent service providers are offering FDDI based MAN services, and the performance envelopes of FDDI and DQDB overlap [BURR], SMDS/DQDB is only marginally a competitor of FDDI. Rather SMDS/DQDB will provide an effective vehicle for connecting LANs in a metropolitan or larger area together. SMDS/DQDB compete more directly with Frame Relay based services, also offered by public carriers as another means of connecting LANs together.

From a user perspective the greatest difference between current SMDS and frame relay offerings is that SMDS offers a connectionless service, while frame relay is conceptually connection oriented and current offerings are largely confined to predefined "virtual circuits." If SMDS becomes widely installed, then it may offer an advantage for providing wide area access to those FDDI LANs whose external communications needs include connectivity to many different remote locations.

9.7 100 Mbps Ethernet

Several vendors have recently announced plans to produce variants of Ethernet that operate at 100 Mbps using twisted pair wiring. The IEEE 802.3 standards committee has decided to write standards for two incompatible variants, one that preserves the CSMA-CD protocol of 802.3, and one that does not.

Many of the characteristics of these prospective standards are not yet clear. They apparently are directed at wiring closet to desktop connections, where they may provide an alternative to FDDI for those workstations that require the 100 Mbps bandwidth. Vendors apparently plan to produce hub products that support both 10 Mbps and 100 Mbps on different ports of the same hub. While very low cost goals have been announced, it is not clear that these 100 Mbps Ethernet products will have a fundamental cost advantage over FDDI on twisted pair. The distance limitations of the twisted pair medium and the distance limitations caused by the CSMA-CD medium access protocol may limit their ability to serve as backbone networks.

Appendix A: Sources for Documents

ANSI Standards and International Standards for FDDI are available from:

American National Standards Institute
11 West 42nd Street, 13th Floor
New York, NY 10036
Sales Department: (212) 642-4900
Fax: (212) 398-0023

Draft FDDI standards are available from:

Global Engineering Documents
15 Inverness Way East
Englewood, CO 80112-5704
(800) 854-7179; outside the United States and Canada (303) 792-2181

The mailings of X3T9.5, the FDDI committee, are available from the X3 Secretariat by subscription. For specific information about mailings contact:

Katrina Gray
X3 Secretariat, CBEMA
1250 Eye Street NW
Washington, DC 20005
Phone: (202) 6265741.

Federal Information Processing Standards Publications (FIPS PUB), NIST Special Publications (NIST SP) and NIST Internal Reports (NIST IR) are produced by the National Institute of Standards and Technology in Gaithersburg, MD. These documents can be ordered from:

National Technical Information Service (NTIS)
United States Department of Commerce
5285 Port Royal Rd.
Springfield, VA 22161
Phone: (703) 487-4650
FTS: 737-4650

Federal Standards and Specifications can be ordered from:

General Services Administration (GSA)
Federal Supply Service Bureau
Specifications Branch
490 East L'Enfant Plaza, SW
Suite 8100
Washington, DC 20407
Telephone: (202) 755-0325 or 775-0326
Fax: (202) 205-3720

IEEE Standards can be ordered from:

IEEE Computer Society Press
Order Department
10662 Los Vaqueros Circle
Los Alamitos, CA. 90270
(800) 272-6657

Current versions of the *Stable Implementation Agreements for Open System Interconnection Protocols* can be ordered from:

Superindentent of Documents
U.S. government Printing Office (GPO)
Washington, DC 20402
(202) 783-3238

Network Management Forum documents can be obtained from:

Network Management Forum
40 Morristown Rd.
Benardsville, NJ 07924
(908) 766-1544
FAX: (908)-766-5741

Many draft standards documents are available over the Internet from anonymous FTP servers. These documents may include meeting minutes, working papers, meeting information and working drafts of standards. Drafts are usually available as a text file or a Postscript print file. Published ANSI standards are not available. Table A.1 lists FTP servers and their addresses.

Table A.1 - List of Anonymous FTP Servers

Subject	Internet Address	Directory	Sub-dir.	Mail to
FDDI	nis.nsf.net (35.1.1.48)	working.groups	fddi	brian.cashman @um.cc.umich.edu
FFOL	ffol.lbl.gov	ffol	docs, papers, minutes, project	RLFink@lbl.gov
HIPPI	nsco.network.com	hippi		
GOSIP	osi.ncsl.nist.gov (129.6.48.100)	pub	gosip, oiw, gug gnmp	

Appendix B: List of Acronyms

The following acronyms are used in this document or are often encountered in discussions and papers about FDDI and other LANs.

AAL	Adaptation Layer - B-ISDN.
ACSE	Association Control Service Element - OSI.
ANS	American National Standard - ANSI.
ANSI	American National Standards Institute - standards body.
AOW	Asian-Oceana Workshop - regional functional standards workshop.
ASC	Accredited Standards Committee - accredited by ANSI.
ASN.1	A Syntax Notation 1 - ISO standard notation for defining data elements in directories and MIBs (ISO 8824 & ISO 8825).
ATM	Asynchronous Transfer Mode - B-ISDN.
ATS	Abstract Test Suite - OSI conformance test jargon.
B-ISDN	Broadband ISDN - new CCITT/T1 standards super family, based om ATM concept.
CASE	Common Application Service Element - OSI.
CGM	Color Graphics Metafile, interchange format.
CBR	Constant Bit Rate service - B-ISDN.
CCITT	Consultative Committee on International Telephony and Telegraphy - treaty standards body.
CD	Committee Draft - ISO.
CLNP	ConnectionLess Network Protocol - ISO 8473 standard.
CLNS	Connectionless network Service - OSI.
CLS	Connectionless - OSI jargon for services.
CLTS	Connectionless Transport Service - OSI.
CO	Connection Oriented - OSI.
COTS	Connection Oriented Transport Service - OSI.
CMIP	Common Management Information Protocol - OSI.
CMT	Connenction ManagemenT - FDDI SMT function.
COS	Corporation for Open Systems - U. S. OSI testing organization.
CSMA	Carrier Sense Multiple Access - LAN media access protocol.
CSMA/CD	Carrier Sense Multiple Access with Collision Detection - LAN media access protocol / IEEE 802.3 standard.
CRC	Cyclic Redundancy Code - industry jargon.
DAC	Dual Attachment Concentrator - FDDI.
DAS	Dual Attachment Station - FDDI.

DCE Data Communications Equipment - X.25.

DIS Draft International Standard - ISO.

DISP Draft ISP - OSI jargon, see ISP.

dpANS draft proposed American National Standard; - ANSI.

DQDB Dual Queue Dual Bus - IEEE 802.6 MAN standard.

DS-1 1.536 Mbps telephone trunk circuit - U.S. telephone industry jargon; equivalent to T1.

DS-3 43.008 Mbps telephone trunk circuit - U.S. telephone industry jargon; equivalent to T3.

DTE Data Terminal Equipment - X.25.

ECC Error Correction Code - industry jargon.

ECF Echo Frame - FDDI management frame.

EIA Electronics Industry Association - standards body.

EMI Electronic EMIssions - industry jargon.

ESF Extended Service Frame - FDDI management frame.

FC Fibre Channel - interface standard.

FCC Federal Communications Commission; often used in reference to regulations for RFI - *FCC Regulations Part 15, Subpart J* - government agency.

FCS Frame Check Sequence - industry jargon.

FDDI Fiber Distributed Data Interface; - LAN standard.

FDDI-II Enhancement of FDDI - LAN standard.

FIPS Federal Information Processing Standard - issued by NIST.

FTAM File Transfer, Access and Management - X.500/ISO 8571-4; OSI application;

FTP File Transfer Protocol - TCP/IP application.

GDMO Guidelines for the Definition of Managed Objects - OSI standard; FDDI SMT defines the FDDI MIB according to the GDMO.

GFC Generalized Flow Control - B-ISDN.

GNMP Government Network Management Profile - network management FIPS.

GOSIP Government Open Systems Interconnection Profile - FIPS PUB 146.

HIPPI High-Performance Parallel Interface - interface standard.

I/O Input/Output - industry jargon.

IP Internet Protocol - a Network layer protocol.

IEC International Electrotechnical Commission - standards body.

IEEE P802 Institute for Electrical and Electronics Engineers Project 802 - LAN network standards body.

IGOSS Industry/Government Open System Specification - GOSIP.

IS International Standard - ISO.

ISP	International Standardized Profile - OSI jargon; a "functional standard."
ISDN	Integrated Services Digital Network - CCITT/T1 telephone & data communications standards super family.
ISO	International Standards Organization; - standards body.
LAN	Local Area Network; - industry jargon.
LCF	Low Cost Fiber - one of the FDDI PMD standards.
LED	Light Emitting Diode - industry jargon.
LEM	Link Error Monitor - FDDI SMT function.
LLC	Logical Link Control - IEEE 802.2/ISO 8802 LAN standard.
MAC	Medium Access Control; - FDDI & general LAN jargon.
MAC-2	enhancement of FDDI MAC for FDDI-II standard.
MAN	Metropolitan Area Network; - industry jargon.
MAP	Manufacturing Automation Protocol - OSI profile for manufacturing.
MHS	Message Handling System - X.400; OSI application; CCITT standard;
MIB	Management Information Base - OSI.
MIC	Media Interface Connector; pronounced mic - FDDI PMD.
NBS	National Bureau of Standards - U.S. government agency, now NIST.
NIF	Neighbor Information Frame - FDDI management frame.
NIST	National Institute of Standards and Technology; pronounced nist - U.S. government agency; formerly NBS.
NRZ	NonReturn to Zero - data transmission coding method; industry jargon.
NRZI	NonReturn to Zero Inverted; - data transmission coding method; industry jargon.
ODA	Open Document Architecture - interchange format (ISO 8613:1989).
OIW	OSI Implementors Workshop - North American regional functional standards workshop.
OSI	Open System Interconnection - plan for families of standards.
PC	Personal Computer - industry jargon.
PDU	Protocol Data Unit - OSI.
PHY	Physical Layer Protocol, - FDDI, FDDI-II & FFOL standards.
PHY-2	enhancement of FDDI PHY for FDDI-II standard.
PICS	Protocol Implementation Conformance Statement - OSI.
PMD	Physical Medium Dependent - FDDI, & FFOL standards.
PMF	Parameter Management Frame - FDDI management frame.
RAF	Resource Allocation Frame - FDDI management frame.
RDF	Request Denied Frame - FDDI Management Frame.
RFI	Radio Frequency Interference - industry jargon.

RISC	Reduced Instruction-Set Computer; - industry jargon.
RMT	Ring ManagemenT - FDDI SMT function.
ROSE	Remote Operations Service Element - OSI.
SAC	Single Attachment Concentrator - FDDI.
SAP	Service Access Point; - OSI.
SAS	Single Attachment Station - FDDI.
SASE	Specific Application Service Element - OSI.
SDH	Synchronous Digital Hierarchy - CCITT standard telephony transmission system; equivalent to SONET.
SIF	Status Information Frame - FDDI management frame.
SGML	Standard Generalized Mark-up Language, interchange format (ISO 8613:1989).
SMAP	System Management Application Process - OSI.
SMF	Single Mode Fiber - one of the FDDI PMD standards.
SMT	Station Management - FDDI, FDDI-II and FFOL standards.
SMT-2-CS	FDDI Station Management Common Services - FDDI-II standard.
SMT-2-IS	FDDI Station Management Isochronous Services - FDDI-II standard.
SMT-2-PS	FDDI Station Management Packet Services - FDDI-II standard.
SMUX	Synchronous Multiplexer - FDDI-II.
SNMP	Simple Network Management Protocol - TCP/IP.
SONET	Synchronous Optical Network - family of standards for public network trunks.
SPM	SONET Physical Mapping - for FDDI; X3T9 standards project.
SRF	Status Report Frame - FDDI Management frame.
STP	Shielded Twisted Pair - FDDI physical medium.
T1	ANSI accredited standards body for telecommunications.
T1	1.536 Mbps telephone trunk circuit - U.S. telephone industry jargon; equivalent to DS-1.
T3	43.008 Mbps telephone trunk circuit - U.S. telephone industry jargon; equivalent to DS-3.
TCP	Transport Control Protocol - a Transport layer protocol.
TIA	Telecommunications Industry Association - standards body.
TP	Twisted Pair - one of the FDDI PMD standards.
TP4	Connection-oriented Transport Protocol Class 4 - ISO.
TRT	Token Rotation Trimer - FDDI MAC timer.
TTRT	Target Token Rotation Time - key FDDI parameter.
UTP	Unshielded Twisted Pair - FDDI physical medium.
VBR	Variable Bit Rate service - B-ISDN.

VCI	Virtual Circuit Indicator - B-ISDN.
VPI	Virtual Path Indicator - B-ISDN.
WAN	Wide Area Network; - industry jargon.
WBC	Wide Band Channel - FDDI-II HRC.
X.25	DTE to DCE packet interface; - CCITT Standard.
X.400	Message Handling System; - CCITT standard & OSI application.
X.500	File Transfer, Access and Management; pronounced exx dot 500 - CCITT standard & OSI application.
X3	ANSI accredited standards body for computers & information systems.
X3T9	domestic X3 technical committee on I/O Interfaces, under which FDDI falls.
X3T9.5	domestic standards committee which is the primary focus of worldwide FDDI standards development. A subgroup of X3T9.

Appendix C: List of References

[802.1D] *Draft IEEE Standard 802.1: Part D, MAC Bridges*, document P802.1D/D8, March 1989.

[802.6] ANSI/IEEE Standard 802.61991, *Distributed Queue Dual Bus (DQDB) Metropolitan Area Network (MAN).*

[BOLA] Tim Boland, NIST Special Publication 500 192, *Government Open Systems Interconnection Profile Users' Guide, Version 2*, National Institute of Standards and Technology, Gaithersburg, MD, October 1991.

[BURR] W. Burr, Xiamei Quan, S. Wakid, and D. Vaman, "A Comparison of FDDI Asynchronous Mode and DQDB Queue Arbitrated Mode Data Transmission for Metropolitan Area Networks," to be published in *IEEE Transactions on Communications.*

[BUX] Werner Bux, "Token ring networks and their performance," *Proceedings of the IEEE*, vol. 77, no. 2, Feb. 1989.

[DAVI] J. Davin et al., RFC 1157, *Simple Network Management Protocol (SNMP),* 1990.

[DQDB] Project 802 Local & Metropolitan Area Networks, Proposed Standard: *Distributed Queue Dual Bus (DQDB) Metropolitan Area Network (MAN)*, October 1, 1990.

[DYK87] Doug Dykeman and Werner Bux, "An investigation of the FDDI media access control protocol," *Proceedings of EFOC/LAN 87*, Information Gatekeepers, Boston, June 1987, pp. 229-236.

[DYK88] Doug Dykeman and Werner Bux, "Analysis and tuning of the FDDI Media Access Control protocol," *IEEE Journal on Selected areas in Communications*, vol. 6., no. 6, July 1988.

[ENG] Kai Y. Eng, Mark J. Karol, and YuShuan Yeh, "A Growable Packet (ATM) Switch Architecture: design Principles and Applications," *IEEE Transactions on Communications*, vol. 40, no. 2, February 1992, pp. 423-430.

[FCC] *FCC Rules and Regulations,* Part 15, Federal Communications Commission.

[FIPS146] FIPS PUB 146-1, *Government Open Systems Interconnection Profile (GOSIP).*

[FIPS159] FIPS PUB 159, *Detail Specification for 62.5 μm Core Diameter/ 125 μm Cladding Diameter class 1A Multimedia Graded-Index Optical Waveguide Fibers*, 1990.

[FIPS174] FIPS PUB 174, *Federal Building Telecommunication Wiring Standard*; adopts EIA/TIA 568, 1991.

[FIPS175] FIPS PUB 175, *Building Standard for Telecommunications Pathways & Spaces*; adopts EIA 569, 1991.*[FIPS176]FIPS PUB 176, Residential & Light Commercial Telecommunications Wiring Standard*; adopts EIA/TIA 570-1991.

[FIPS179] FIPS PUB 179, *Government Network Management Profile,* Version 1.

[FIPSXXX] Draft FIPS PUB, *Grounding and Bonding Requirements for Telecommunications in*

Commercial Buildings. This standard will be based on an EIA/TIA draft standard).

[G.707] .CCITT Recommendation G.707, "Synchronous Digital Hierarchy Bit Rates," *Blue Book*, 1988.

[G.708] .CCITT Recommendation G.707, "Network Node Interface for the Synchronous Digital Hierarchy," *Blue Book*, 1988.

[GDMO] draft International Standard 10165/4, Structure of Management Information Part 4: Guidelines for the Definition of Managed Objects, June 15, 1990.

[HONG] Duke Hong and Tatsuya Suda, "Congestion Control and Prevention in ATM Networks," *IEEE Network*, vol. 5, no. 4, July 1991, pp. 1016.

[HUB] IEEE P802.3.K/D10, *Draft Supplement to ANSI/IEEE Standard 802.3 - 1992 Edition, Repeater Management,* July 11, 1992.

[HRC] ANSI X3.186-1992, *FDDI Hybrid Ring Control (HRC)*; ISO standard will be 9314-5.

[I.330] CCITT Recommendation I.330 *ISDN numbering and addressing principles.*

[I.121] CCITT Draft Recommendation I.121, *Broadband Aspects of ISDN*, Study Group XVIII, Geneva, May 1990.

[IGOSS] *Industry/Government Open Systems Specification*, to be published by NIST.

[ISO7498] ISO 7498:1984, *Information Processing Systems- Open Systems Interconnection- Basic Reference Model.*

[ISO8073] ISO/IEC 8073:1992, *Information Technology- Telecommunications and Information Exchange Between Systems- Protocol for Providing the Connection-mode Transport Service.*

[ISO8327] ISO 8327:1987, *Information Processing Systems - Open Systems Interconnection - Basic Connection Oriented Session Protocol Specification.*

[ISO8473] ISO 8473:1988, *Information Processing Systems- Data Communications- Protocol for Providing the Connectionless-mode Network Service.*

[ISO8571] ISO 85711-1:1988, *Information Processing Systems- Open Systems Interconnection- File Transfer, Access and Management Part 1: General Introduction.*

[ISO8650] ISO 8650:1988, *Information Processing Systems- Open Systems Interconnection- Protocol Specification for Association Control Service Element: Association Control.*

[ISO8822] ISO 8822:1988, *Information Processing Systems- Open System Interconnection- Connection Oriented Presentation Service Definition.*

[ISO88022] ISO 8802-2:1989, *Information Processing System- Local Area Networks- Part 2: Logical Link Control.*

[ISO88023] ISO/IEC 8802-3:1992, *Information Processing System- Local and Metropolitan Area Networks- Part 3: Carrier Sense Multiple Access Method and Physical Layer Specifications.*

[ISO88025] ISO/IEC 8802-5:1992, *Information Processing System- Local and Metropolitan Area Networks- Part : Token Ring Access Method and Physical Layer Specifications.*

[ISO9041] ISO 9041:1990, *Information Technology - Open Systems Interconnection - Virtual Terminal Protocol Basic Class.*

[ISO9072] ISO/IEC 9072-1:1989, *Information Processing Systems- Text Communication- Remote Operations - Part 1: Model, Notation and Service Definition.*

ISO/IEC 9072-2:1989, *Information Processing Systems- Text Communication- Remote Operations - Part 2: Model, Notation and Service Definition.*

[ISO9574] ISO/IEC 99574:1992, *Information Technology- Provision of OSI Connection-Mode Network Service by Packet Mode Terminal Equipment Connected to an Integrates Services Digital Network (ISDN).*

ISO9595] ISO/IEC 9595:1991, *Information Technology- Open Systems Interconnection- Common Management Information Service Definition.*

[ISO9596] ISO/IEC 9596-1:1991, *Information Technology- Open Systems Interconnection- Common Management Information Protocol- Part 1: Specification.*

[ISO9646] ISO/IEC 9646-1:1991, *Information Technology- Open Systems Interconnection- Conformance Testing Methodology and Framework - Part 1: General Concepts.*

ISO/IEC 9646-2:1991, *Information Technology- Open Systems Interconnection- Conformance Testing Methodology and Framework - Part 2: Abstract Test Suite Specification.*

ISO/IEC 9646-3:1992, *Information Technology- Open Systems Interconnection- Conformance Testing Methodology and Framework - Part 3: Tree and Tabular Combined Notation (TTCN).*

[ISO10164] ISO/IEC 10164-4:1992, *Information Technology- Open Systems Interconnection- Systems management: Alarm Reporting Function.*

[ISO10608] ISO/IEC DISP 10608-14, *Information Processing System - International Standardized Profile TAnnnn - Connection-mode Transport Service over Connectionless-mode Network Service- Part 14: MAC, PHY and PMD Sublayer Dependent and Station Management Requirements over an FDDI LAN Subnetwork.*

[ISO10589] ISO/IEC 10589, *Information Technology- Telecommunications and Information Exchange Between Systems- Intermediate System to Intermediate System Intra-Domain Routing Information Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO8473).*

[ISO10733] ISO/IEC 10733, *Information Technology- Telecommunications and Information Ex-*

change Between Systems - Elements of Management Information Related to OSI Network Layer Standards.

[ISO10737] ISO/IEC 10737, *Information Technology- Telecommunications and Information Exchange Between Systems- Elements of Management Information Related to OSI Transport Layer Standards.*

[JAIN] Raj Jain, "Performance Analysis of FDDI Token Ring Networks: Effect of Parameters and Guidelines for Setting TTRT," *Proceedings ACM SIGCOMM'90 Symposium on Architectures and Protocols,* September 2427, 1990, Philadelphia, PA, pp. 264275.

[JOHN87] Marjory J. Johnson, "Proof that timing requirements of the FDDI token ring protocol are satisfied," *IEEE Transactions on Communications,* vol. COM35, no. 6, June 1987, pp. 620625.

[JOHN88] M. J. Johnson, "Performance Analysis of FDDI," *Proc. of EFOC/LAN 88,* held June 29 July 1 1988 in Amsterdam, Information Gatekeepers, Boston, 1988, pp. 295300.

[LAMA] R. O. LaMaire "An M/G/1 Vacation Model of an FDDI Station," *IEEE Journal on Selected Areas in Communications,* vol. 9, no. 2, February 1991, pp. 257264.

[LATI] Aamer Latif, Edward J Rowlance, and R. Hole Adams, "The IBM 8029 LAN Bridge," *IEEE Network,* vol. 6, no. 3, May 1992, pp. 2837.

[LCF] X3.237-199x, draft proposed American National Standard, *FDDI Low-Cost Fibre Physical Layer Medium Dependent (LCF-PMD),* 1 September 1992 (can be obtained from Global Engineering Documents - see Annex A); ISO reference is CD 8314-9

[LMAN] IEEE P802.1B/D20, *Draft Standard 802.1B: LAN/MAN Management,* January 27, 1992.

[M.3100] CCITT Draft Recommendation M.3100, *Generic Information Model,* November 1991.

[MAC] ISO 9314-2:1989, *Information Processing Systems- Fiber Distributed Data Interface (FDDI)- Part 2: Token Ring Media Access Control (MAC).*; Equivalent to ANSI X3.139-1987.

[MAC2] draft proposed American National Standard X3.239-199X, *FDDI Medium Access Control (MAC-2)*; ISO standard will be 9314-8 (draft can be obtained from Global Engineering Documents - see Appendix A).

[MIL1777] MIL-STD 1777, *Military Standard Internet Protocol,* 1983.

[MIL1778] MIL-STD 1778, *Military Standard Transport Protocol,* 1983.

[NEWMN] Peter Newman, "ATM Technology for Corporate Networks," *IEEE Communications Magazine,* vol. 30, no. 4, April 1992, pp. 90101.

[NMFMIL] Network Management Forum: Forum 006, *Forum Library - volume 4: ONMIPoint Definitions,* Issue 1.0, August 1992.

[OSI] ISO 7498, *Information Processing Systems Open Systems Interconnection Basic Reference Model,* International Standards Organization, Geneva, 1984.

[PHY] ISO 9314-1:1989, *Information Processing Systems- Fiber Distributed Data Interface (FDDI)- Part 1: Token Ring Physical Layer Protocol (PHY)*; Equivalent to ANSI X3.1481988.

[PHY2] draft proposed American National Standard X3.231-199X, *FDDI Physical Layer Protocol (PHY-2)*; ISO standard will be 9314-7 (draft can be obtained from Global Engineering Documents - see Appendix A)

[PMD] ISO 9314-3:1990, *Information Processing systems- Fiber Distributed Data Interface (FDDI)- Part 3: Token Ring Physical Medium Dependent (PMD)*; Equivalent to ANSI X3.1661989.

[POST] J. Postel, RFC 768, *User Datagram Protocol,* 1980.

[SEVI] Kenneth C. Sevick and Marjory J. Johnson, "Cycle time properties of the FDDI token ring protocol," *IEEE Transactions on Software Engineering,* vol. SE13, no. 3, March 1987, pp. 376385.

[SMF-PMD] draft proposed American National Standard X3.184-199x, *Fibre Distributed Data Interface (FDDI) - Single-Mode Fiber Physical Layer Medium Dependent (SMF-PMD)*; ISO standard will be 9314-4 (can be obtained from Global Engineering Documents - see Appendix A).

[SMT] draft proposed American National Standard X3.231-199X, *FDDI Station Management (SMT),* June 25, 1992; ISO standard will be 8314-6 (can be obtained from Global Engineering Documents - see Appendix A).

[SMT2CS] draft proposed American National Standard, *FDDI Station Management-2 Common Services (SMT-2-CS)*, (draft can be obtained from Global Engineering Documents - see Appendix A)

[SMT2IS] draft proposed American National Standard, *FDDI Station Management-2 Common Services (SMT-2-IS)*, (draft can be obtained from Global Engineering Documents - see Appendix A)

[SMT2PS] draft proposed American National Standard, *FDDI Station Management-2 Packet Services (SMT-2-PS)*, (draft can be obtained from Global Engineering Documents - see Appendix A)

[STABLE] NIST Special Publication 500-206, *Stable Implementation Agreements for Open Systems Interconnection,* Version 6 Edition 1.

[T1.618] ANSI T1.618, *DSS1Core Aspects of Frame Protocol for use with Frame Relay Bearer Service*, American National Standards Institute, Inc., New York, 1991.

[T1.617] ANSI T1.617, *DSS1 Signaling Specification for Frame Relay Bearer Service,* American National Standards Institute, Inc., New York, 1991.

[T1.602] ANSI T1.602, *Frame Relay Bearer Service Architectural Framework and Service De-*

scription, American National Standards Institute, Inc., New York, 1990.

[T1.105] ANSI-T1.105-1988, *American National Standard for Telecommunications Digital Hierarchy Optical Interface Rates and Formats Specifications*, 1988.

[T1S1] T1S1/92302, *Draft American National Standard for Telecommunications Broadband ISDN ATM Layer Functionality and Specification.*

[T1S1606] Draft Addendum to T1606, *Frame Relay Bearer Service Architectural Framework and Service Description,* American National Standards Institute, Inc., New York, July 1990.

[TSB36] EIA/TIA TSB36, *Additional Cable Specifications for Unshielded Twisted Pair Cables,* Electronics Industries Association, Washington, DC, November 1991.

[TURN] Jonathan S. Turner, "Design of an Integrated Services Packet Network," *IEEE Journal on Selected Areas in Communications,* vol. SAC4, no. 8, November 1966, pp. 1373-1380.

[ULM] J. M. Ulm, "A timed token ring protocol and its performance characteristics," *Proc. IEEE 7th Local Computer Networks Conference,* Minneapolis, MN, Feb. 1982, pp. 5056.

[WAIN] Nick Wainright and Andrew Myles, "A Comparison of the Delay Characteristics of the FDDI and IEEE 802.6 MAC Layer Protocols," Technical Memo, Hewlett Packard Laboratories Filton Road, Stoke Gifford, Bristol BS12 6QZ, England, March 1989.

[WANG] Zheng Wang and Jon Crowcroft, "SEAL Detects Cell Misordering," *IEEE Network,* vol. 6, no. 4, July 1992, pp. 89.

[X.400] CCITT Recommendation X.400 (Red Book, 1985), *Message Handling System: System Model Service Elements.*

Index

10-Base-F 69
10-Base-T 69

A

Abstract Test Suite (ATS) 58
analytical models of FDDI 34
announcement frames 22
Asian-Oceana Workshop (AOW) 60
asynchronous transfer mode (ATM) 79 - 82, 84
  ATM LAN 27, 75, 79 - 81, 84
attachment
  See FDDI attachment

B

Beacon 8, 20, 26
  directed 26
bit error rate (BER) 11
bridge 24, 61, 63, 68
  encapsulating 65
  learning 64
  local bridge 64
  remote bridge 64
  source routing 64 - 65
  spanning tree 64
  translating 65
  transparent 64 - 65
broadband ISDN (B-ISDN) 27, 31, 76, 79, 81 - 82
  adaptation layer (AAL) 79 - 80, 86
brouter 61
building wiring standards 50
  backbone wiring 51
  horizontal wiring 50
bypass 8
bypass switch 4

C

Claim 8, 20, 22, 26, 35
collapsed backbone 69
common management information protocol (CMIP) 24, 58
common management information services (CMIS) 58
concentrator 62
connection
  See FDDI connection
connection management (CMT) 8, 10, 25 - 26
connectionless mode network protocol (CLNP) 56
cycle master
  See FDDI-II, cycle master
cyclic redundancy check (CRC) 80, 86

D

D_Max 36
data link layer
  See OSI reference model
datagram service 57
distributed queue dual bus (DQDB) 84 - 86
dual attachment concentrator (DAC) 4
dual attachment station (DAS) 2, 10
dual homing 9
dual ring 4, 6 - 7
dual ring of trees 2, 7 - 8

E

end system 2
european workshop on open systems (EWOS) 60
extended LAN 62

F

fault detection 25
fault recovery 6 - 7
FDDI asynchronous mode priorities 20
FDDI asynchronous traffic 19 - 20
FDDI attachment 4 - 5
FDDI concentrator 4, 24, 63, 69
FDDI configuration limits 36 - 38
FDDI connection 4 - 5, 9
  A-to-A 4
  A-to-B 4, 7
  B-to-B 4
  M-to-A 4
  M-to-B 4
  M-to-M 4, 25
  M-to-S 4, 7
FDDI follow-on lan (FFOL) 31, 70 - 71
FDDI link 5, 10
FDDI management information base (MIB) 22
FDDI path 4 - 5, 15, 21
FDDI port 5, 11 - 12, 21
  A port 4, 25
  B port 4, 25
  M port 4
  S port 4
FDDI ring initialization 20
FDDI synchronous traffic 19 - 20
FDDI token 19
FDDI wrap 6 - 7, 10
FDDI-II 20, 27, 31 - 33
  configuration limits 38

cycle 28
cycle header 28
cycle master 29
cyclic group 28
dedicated packet group 28
HRC 28, 30
MAC-2 29 - 30, 67
PHY-2 29 - 30
SMT-2 29 - 30
wide band channel (WBC) 28
fibre channel (FC) 72 - 75
filtering rate 62, 66
forwarding rate 62, 66
frame management services 22
frame relay 64, 78

G

gateway 62
GDMO 22
GNMP 25, 55, 58
GOSIP 34, 56
FDDI role in 58
government network management profile
See GNMP
government network management profile
See GNMP

H

high-performance parallel interface (HIPPI) 71 - 72
hub 63, 69
hybrid ring control (HRC)
See FDDI-II, hybrid ring control

I

IEEE standard
802.1d 65
802.2 2
802.3 2, 19, 61, 63, 65
802.5 2, 19, 61, 63, 65
IGOSS 55
Industry/Government Open System Specification
See IGOSS
integraded services digital network (ISDN) 27
intelligent peripheral interface (IPI) 72 - 73
intermediate system 2, 63
international standardized profile (ISP) 59
internet protocol (IP) 2, 67, 73

J

JTC1 59

L

LCF-PMD 11 - 12, 43, 49
link
See FDDI link
link confidence test 25
link error monitor (LEM) 25
link error rate 25
LLC test command 65
local area network (LAN) 2
logical link control (LLC) 2, 85
low cost fiber PMD
See LCF- PMD

M

MAC 2, 7, 19, 21, 25, 61
MAC address 19, 25, 63
MAC bridge
See bridge
management agent process 22
master port
See FDDI M port
mater/slave loop
See topology errors
media acess control
See MAC
media interface connector (MIC) 11 - 12
MLT-3 13
multistation acces unit (MAU) 69

N

network utilization 34
NRZI 13

O

open systems interconnection
See OSI reference model
optical fiber
50/125 multimode fiber 74
62.5/125 multimode fiber 11 - 13, 41 - 42, 48 - 49, 51 - 54, 74
loss in 43
modal dispersion 41
multimode fiber 41, 44, 48 - 49
numerical aperture 41
power budget 49
single mode 13, 41 - 42, 49, 74
optical fiber cables 44
OSI implementors workshop (OIW) 60
OSI reference model 2 - 3, 55
and LAN interconnection devices 61
application layer 55
data link layer 2, 57, 67
logical link control (LLC) sublayer 57
media access control (MAC) sublayer 57
network layer 57, 63, 67
physical layer 2, 57

presentation layer 55
session layer 55
transport layer 56

P

path
See FDDI path
path test 25
PHY 2, 36
PHY-2 36
physical layer
See OSI reference model
physical medium dependent
See PMD
PICS 58, 60
PMD 2, 10 - 12, 27, 43 - 44, 49
port
See FDDI port
port delay 36
primary ring 4, 7
protocol implementation conformance statement
See PICS

R

remote management of FDDI 24
repeater 63
request frames 22
response frames 22
ring management (RMT) 25 - 26
router 2, 24, 63, 67 - 68, 81

S

secondary ring 4, 7
simple network management protocol (SMAP) 59
simple network management protocol (SNMP) 24
single attachment concentrator (SAC) 4, 7
single attachment station (SAS) 4, 7 - 9
single mode fiber PMD
See SMF-PMD
slave port
See FDDI S port
small computer system interface (SCSI) 72 - 73
SMDS 64
SMF PMD 12
SMF-PMD 49
SMT 2, 8, 20 - 21
SMT frames
echo frame (ECF) 23
extended service frame (ESF) 24
neighbor information frame (NIF) 22, 64
parameter management frame (PMF) 23 - 24
request denied frame (RDF) 23
resource allocation frame (RAF) 23
status information frame (SIF) 23, 25
status report frame (SRF) 23
SONET 13, 49
spanning tree algorithm 64
SPM 49
station management
See SMT
status information frame (SIF) 10
subnetwork subprofiles 57
switched multimegabit data service (SMDS) 84 - 86
synchronous digital hierarchy (SDH) 76
synchronous optical network (SONET) 2, 31, 76 - 77, 79, 81
optical carrier (OC-n) rates 76
path structure 77
synchronous payload Envelope (SPE) 76
synchronous transport module (STM) 77
synchronous transport signal (STS) 76

T

T_Neg 22
T_Req 22, 34
target token rotation time (TTRT) 19 - 20, 34 - 37
TCP/IP 2
time division multiplexing (TDM) 27, 31, 71, 77
timed token protocol 19
token holding timer (THT) 19
token rotation timer (TRT) 19
topology errors 9
A-to-A link 9
B-to-B link 9
master/slave loop 9
twisted ring 4, 9
TP-PMD 13
trace 25 - 26
transport control protocol (TCP) 2, 66
transport protocol class 4 (TP4) 56
tree and tabular combined notation (TTCN) 58
twisted pair 48 - 50
category 5 unshielded 13, 48

shielded in horizontal wiring 51
type 1 shielded 48
unshielded 51
unshielded categories 46
unshielded in horizontal wiring 51
twisted pair PMD
See TP-PMD

U

universal MAC address 19

W

wide band channel (WBC)
See FDDI-II WBC
Workshop Agreements 55
wrap 7
See FDDI wrap

☆ U.S. GOVERNMENT PRINTING OFFICE : 1993 O - 361-460

ANNOUNCEMENT OF NEW PUBLICATIONS ON COMPUTER SYSTEMS TECHNOLOGY

Superintendent of Documents
Government Printing Office
Washington, DC 20402

Dear Sir:

Please add my name to the announcement list of new publications to be issued in the series: National Institute of Standards and Technology Special Publication 500-.

Name ____

Company ____

Address ____

City ____________________ State ______________ Zip Code ____________

(Notification key N-503)

NIST *Technical Publications*

Periodical

Journal of Research of the National Institute of Standards and Technology—Reports NIST research and development in those disciplines of the physical and engineering sciences in which the Institute is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Institute's technical and scientific programs. Issued six times a year.

Nonperiodicals

Monographs—Major contributions to the technical literature on various subjects related to the Institute's scientific and technical activities.

Handbooks—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

Special Publications—Include proceedings of conferences sponsored by NIST, NIST annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

Applied Mathematics Series—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

National Standard Reference Data Series—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a worldwide program coordinated by NIST under the authority of the National Standard Data Act (Public Law 90-396). NOTE: The Journal of Physical and Chemical Reference Data (JPCRD) is published bimonthly for NIST by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements are available from ACS, 1155 Sixteenth St., NW, Washington, DC 20056.

Building Science Series—Disseminates technical information developed at the Institute on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

Technical Notes—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NIST under the sponsorship of other government agencies.

Voluntary Product Standards—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The standards establish nationally recognized requirements for products, and provide all concerned interests with a basis for common understanding of the characteristics of the products. NIST administers this program in support of the efforts of private-sector standardizing organizations.

Consumer Information Series—Practical information, based on NIST research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

Order the **above** *NIST publications from: Superintendent of Documents, Government Printing Office, Washington, DC 20402.*

Order the **following** *NIST publications—FIPS and NISTIRs—from the National Technical Information Service, Springfield, VA 22161.*

Federal Information Processing Standards Publications (FIPS PUB)—Publications in this series collectively constitute the Federal Information Processing Standards Register. The Register serves as the official source of information in the Federal Government regarding standards issued by NIST pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

NIST Interagency Reports (NISTIR)—A special series of interim or final reports on work performed by NIST for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Service, Springfield, VA 22161, in paper copy or microfiche form.