

ANSI Fiber Distributed Data Interface (FDDI) Standards

In this report:

Applications	2
FDDI Topologies	5
FDDI Specifications	7
Physical Medium Dependent (PMD) Specification	8
Physical Protocol (PHY) Sublayer Specification	9

Datapro Summary

In October 1982, the American National Standards Institute (ANSI), Committee X3T9.5, was chartered to develop a high-speed data networking standard to create high-speed backbones between supercomputers. This work yielded a suite of standards cumulatively known as Fiber Distributed Data Interface (FDDI)—a packet switched, token-passing, ring-based backbone networking technology that transports data at 100M bps over a variety of fiber- and copper-based media. FDDI does not handle realtime, interactive voice and video applications well, and its speed is inadequate for some high-performance broadband networking applications. FDDI, nevertheless, has since found widespread acceptance interconnecting distributed LANs, servers, and high-performance workstations. Delays in the evolution of ATM and the development of switched LAN technology have breathed new life into the FDDI market making FDDI a viable interim networking solution. FDD II, an enhancement and superset of FDDI, is evolving to add the capability to carry voice and video more efficiently. This effort, however, has limited industry support since ATM promises to deliver solutions within the same time frame. Also under consideration is FDDI Follow-On (FFOL), which plans higher speed transport and connection of multiple FDDI networks. It has similar limited support due to competing technologies.

Analysis

FDDI addresses the bottom two layers of the OSI model. The spirit of the FDDI standard is to use ISO protocols for the other layers, where possible. The optical-based FDDI LAN was designed to offer the same type of serial interconnection provided by traditional Ethernet and token-ring LANs while adding the high bandwidth, inherent noise immunity, and security of fiber connections (see Figure "A Sample FDDI Network").

At FDDI's inception in 1982, fiber was used mostly for point-to-point applications and not for the many configurations allowed by LANs. In this sense, FDDI was a breakthrough. Although it is possible to achieve considerably higher data rates over fiber (up to 4G bps with current point-to-point technology and 500M bps

on rings), higher rates result in significantly increased costs and shorter transmission distances between repeaters. Its designers, intending that FDDI should provide relatively inexpensive connectivity, focused on 100M bps transmission. (In sustained transmissions, a maximum data transfer rate of 80M bps is achievable. The remaining bandwidth is reserved for various overhead functions.)

FDDI is a token-passing, dual-ring based network accommodating synchronous and asynchronous data transmission, as well as isochronous channels for realtime digitized voice and compressed video (see Figure "The FDDI Environment"). *Class A* nodes are those that reside on the ring backbone. *Class B* nodes are non-ring stations; they may be standalone devices or elements within tree or star subnetworks.

Unlike existing open standards for LANs, where fiber optic variants have followed copper implementations, FDDI has been designed from the start as a fiber optic network. This design approach required establishment of standards in areas such as duplex optical connectors, fiber characteristics, optical bandwidth, bypass relays,

—By M. Scott Kingsley

Glossary of FDDI Terms

Asynchronous—A class of data transmission service whereby all requests for service contend for a pool of dynamically allocated ring bandwidth and response time.

Attachment—A Port or pair of Ports, optionally including an associated optical bypass, that are managed as a functional unit. A dual attachment includes two ports: a Port A and a Port B. A single attachment includes a Port S.

Bypass—The capability of a node to optically isolate itself from the FDDI network while maintaining the continuity of the cable plant.

Capture—The act of removing a token from the ring for the purpose of frame transmission.

Claim Token—A process whereby one or more stations bid for the right to initialize the ring.

Code Bit—The smallest signaling element used by the Physical Layer for transmission on the medium.

Code Group—The specific sequence of five code bits representing a DDL symbol.

Concentrator—An FDDI node that has additional ports beyond those required for its own attachment to an FDDI network. These additional ports (type M) are for attaching other FDDI nodes (including other concentrators) in a tree topology.

Connection Management (CMT)—That portion of the Station Management (SMT) function that controls network insertion, removal, and connection of PHY and MAC entities within a station.

Counter-rotating—An arrangement whereby two signal paths, one in each direction, exist in a ring topology.

Dual Attachment Concentrator—A concentrator that offers a dual attachment to the FDDI network and is capable of accommodating a dual (counter-rotating) ring.

Dual Ring (FDDI dual ring)—A pair of counter-rotating logical rings.

Entity—An active service or management element within an Open Systems Interconnection (OSI) layer or sub-layer.

Fiber Optic Cable—A cable containing one or more optical fibers.

Frame—A PDU transmitted between cooperating MAC entities on a ring, consisting of a variable number of octets and control symbols.

Jitter, Random—The probabilistic offsets of pulse transition edges from the expected time. Includes both Duty Cycle Distortion and Data Dependent Jitter.

Jitter, Systematic—The deterministic offsets of pulse transition edges from the expected time. Some sources of systematic jitter are differences in rise and fall times and propagation delays.

Logical Ring—The set of MACs serially connected to form a single ring. A fault-free FDDI network provides two logical rings.

Media Access Control (MAC)—The Data Link Layer responsible for scheduling and routing data transmissions on a shared-medium local area network (e.g., an FDDI ring).

Media Interface Connector (MIC)—A mated connector pair that provides an attachment between an FDDI node and a cable plant. The MIC consists of two parts: an MIC plug and an MIC receptacle.

MIC Plug—The male part of the MIC which terminates a fiber optic cable.

MIC Receptacle—The female part of the MIC which is contained in an FDDI node.

Network (FDDI Network)—A collection of FDDI nodes interconnected to form a trunk, a tree, or a trunk with multiple trees. This topology is sometimes called a dual ring of trees.

Node—A generic term applying to an active element in an FDDI network (station or concentrator).

NRZ—Non Return to Zero, a technique where a polarity level (+ or -) represents a logical "1" (one) or "0" (zero).

NRZI—Non Return to Zero Invert on Ones, a technique where a polarity transition represents a logical "1" (one). The absence of a polarity transition denotes a logical "0" (zero).

Parameter Management Frames (PMF)—PMFs provide remote access to the SMT MIB.

and cable assemblies. The network can support a total cable distance of 100 km. per ring, tolerate up to 2 km. of fiber between stations, and serve up to 500 attached devices or 1,000 physical connections (2 per device). As shown in Figure "The FDDI Environment," FDDI employs dual counter-rotating, token-passing rings—an approach that gives FDDI some of the features of conventional token-ring LANs. FDDI is not part of the well-established IEEE 802 family of LAN standards, however.

FDDI operates at 1300 nanometers (nm.). Current transmitter/receiver fiber technology operates at 850 nm., 1300 nm., or 1550 nm. While performance increases with wavelength, so does cost. For local data communications, both in LANs and in point-to-point applications employing fiber optic modems, the 850-nm. light sources are typically employed; however, this technology becomes unfeasible for 100M bps beyond a couple of miles. At the other end of the range (1550 nm.), the system becomes expensive and may provide unnecessary bandwidth.

The committee designing FDDI also investigated short-wavelength implementations. It became evident that to meet all the requirements, particularly the two-kilometer, station-to-station spacing, the system would require 1300-nm. wavelength. Using

1300-nm. technology, less expensive light-emitting diodes (LEDs) provide distance and data rates within the range desired for LANs and LAN backbones. The issue of fiber size was settled sometime after the wavelength decision.

The FDDI standard directly addresses the need for reliability, since the failure of a backbone system providing data transport for many user sessions would be serious. FDDI incorporates three reliability-enhancing features.

1. A failed or unpowered station can be bypassed by an optional automatic optical bypass switch.
2. Wiring concentrators are used in a star-wiring strategy to facilitate fault isolation and correction.
3. Two rings interconnect stations so that failure of a repeater or cable link results in the automatic reconfiguration of the ring using a loopback, or "wrapping" mechanism.

Applications

FDDI was originally designed to provide a high bandwidth backbone between high-performance computers. Rapid growth in

Physical Connection—The full-duplex Physical Layer association between adjacent PHY entities (in concentrators or stations) in an FDDI network (i.e., a pair of Physical Links).

Physical Layer (PHY)—The Physical Layer responsible for delivering a symbol stream produced by an upstream MAC Transmitter to the logically adjacent downstream MAC Receiver in an FDDI ring.

Physical Link—The simplex path (via PMD and attached medium) from the transmit function of one PHY entity to the receive function of an adjacent PHY entity (in concentrators or stations) in an FDDI network.

Physical Media Dependent (PMD)—PMD defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connector, the functions of the optical receiver, and (as an option) the bypass switch that can be incorporated into the station.

Port—A PHY entity and a PMD entity in a node, together creating a PHY/PMD pair, that may connect to the fiber media and provide one end of a physical connection with another node.

Primitive—An element of the services provided by one entity to another.

Protocol Data Unit (PDU)—Information delivered as a unit between peer entities which may contain control information, address information, and data (e.g., a Service Data Unit from a higher layer).

Receive—The action of a station in accepting a token, frame, or other signal sequence from the incoming medium.

Receiver (optical)—An electro-optical circuit that converts an optical signal to an electrical logic signal.

Repeat—The action of a station in receiving a token or frame from the adjacent upstream station and simultaneously sending it to the adjacent downstream station. The FDDI MAC may repeat received PDUs (tokens and frames) but it does not repeat the received signal stream between PDUs. When repeating a frame, MAC may copy the data contents and modify the control indicators as appropriate.

Repeater—A Physical Layer relay in an FDDI network.

Ring—A set of stations arranged in a closed loop, enabling information to pass sequentially from one station to

the next until it arrives at the intended destination station. In FDDI usage, the term "ring" or "FDDI ring" refers to a dual (counter-rotating) ring.

Service Data Unit (SDU)—The unit of data transfer between a service user and a service provider. The MAC SDU is the data contents of a frame. The PHY SDU is a symbol.

Services—The functions performed by one entity for another. Data services are provided by a higher level entity; management services are provided to a management entity at the same or another level.

Single Attachment Concentrator—A concentrator that offers a single attachment to the FDDI network.

Single Attachment Station—A station that offers a single attachment to the FDDI network.

Station—An addressable logical and physical node on the FDDI ring capable of transmitting, repeating, and receiving information. A station has exactly one SMT, at least one MAC, at least one PHY, and at least one PMD entity.

Station Management (SMT)—The supervisory entity within an FDDI station that

monitors and controls the various FDDI entities including PMD, MAC, and PHY.

Symbol—The smallest signaling element used by MAC (i.e., the PHY SDU). The symbol set consists of 16 data symbols and 8 control symbols. Each maps to a specific sequence of five code bits as transmitted by the Physical Layer.

Synchronous—A class of data transmission service whereby each requester is preallocated a maximum bandwidth and guaranteed a response time, which is not to exceed a specific delay value.

Token—An explicit indication of the right to transmit on a shared medium. On a token-ring LAN, the token circulates sequentially through the stations on the ring. At any time, it might be held by no stations or by one station. FDDI uses two classes of tokens: restricted and unrestricted.

Transmit—The action of a station in generating a token, frame, or other symbol sequence and placing it on the outgoing medium.

Transmitter (optical)—An opto-electrical circuit that converts an electrical logic signal into an optical signal.

bandwidth intensive applications, however, soon dwarfed even FDDI's 100M-bps capability. FDDI, moreover, began to find competition from Asynchronous Transfer Mode (ATM) technology, which promised high-bandwidth multiplexing and switching for both WAN and LAN applications. ATM's inability to deliver timely solutions, however, and other developments have fueled an FDDI resurgence. FDDI found a new mission—creating a higher speed backbone interconnecting multiple lower speed LANs into a single heterogeneous campus network. As bandwidth intensive applications trickled down to the desktop, FDDI was there to provide direct user connection to a single higher speed LAN rather than the 4M bps—16M bps Ethernet or token ring.

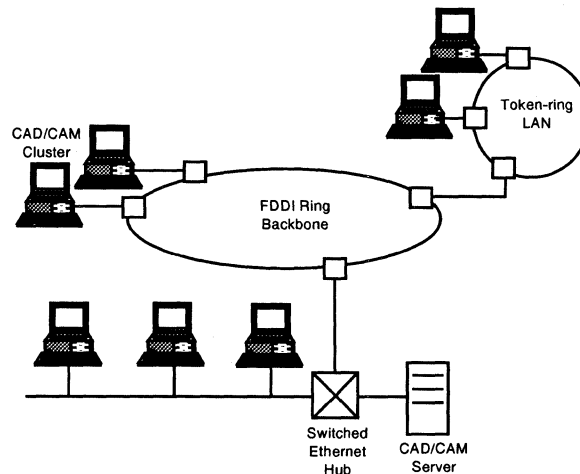
Ethernet has been expanded to 100M bps (100BASE-T and 100VG "Fast-Ethernet"). Much like ATM, however, it is an immature technology with arguable problems and limitations. A key strength of improved Ethernet technologies is that they can be implemented using voice-grade copper twisted-pair cabling. FDDI has responded to this development with support for lower cost fiber optic and twisted-pair copper cabling, thereby further encouraging use of FDDI.

This change seemed to mark the end of FDDI growth. High bandwidth and switching, however, are still needed but months or years away from practical implementation using ATM. The gap is presently being filled by switched LAN technology. Different LANs are connected to a centralized hub (e.g., switched Ethernet) for routing. This arrangement permits multiple LAN-based applications to access one or more large servers through shared high-speed pipes. FDDI is the natural technology choice for this pipe. LAN switching architectures are evolving toward FDDI switching—including hubs dedicated to FDDI connections alone to integrate different FDDI subnetworks into a single network. Vendors developing this type of implementation promise an increase in backbone bandwidth to more than 3G bps.

Security is one of FDDI's key benefits. Fiber optic cabling produces near zero emissions and prevents physical intrusion. Additionally, the dual-ring FDDI architecture can survive any single network failure, and it minimizes the damage of nodal failures.

The FDDI standard has been integrated into the Synchronous Optical Network (Sonet), the transport scheme developed and

Figure
A Sample FDDI Network




used by common carriers in long-haul fiber optic networks. Sonet extends the previously limited range of FDDI rings to very long distances.

FDDI designers envision improved performance of realtime interactive video and voice through FDDI-II and other FDDI follow-on (FFOL) technologies. In reality, most see ATM as taking the reins of network architecture evolution in the near future, even though it is currently an immature technology with interoperability problems and needing additional standards development. Other technology advancements, such as switched LANs, are also weakening future FDDI visions. In the short-term, however, FDDI is a proven and mature technology—one that has been documented, accepted as a national and international standard, and supported in a broad range of dependable products from more

than one hundred vendors. The following factors should signal a private network designer of the need to seriously consider FDDI:

1. When accessing remote LANs, network delays are too high or transmission throughput is unacceptable.
2. The computing environment must move toward a client/server, distributed database architecture, which will require higher performance.
3. The network requires a fault-tolerant networking solution.
4. Multimedia or other video-based applications are appearing.

Figure
The FDDI Environment



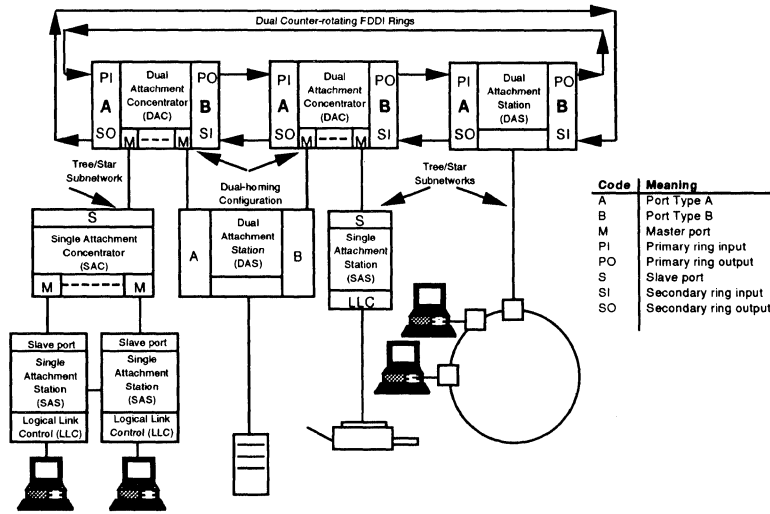


Figure FDDI Configuration Options

An FDDI network application illustrating node connectivity.

FDDI Topologies

FDDI networks can be configured using three topologies (see Figure "FDDI Configuration Options"). The most common topology is a ring with attached tree subnetworks. The ring is actually dual rings, called primary and secondary. Data on the two rings is counter-rotating, using one direction on the primary and the opposite on the secondary. Data is generally placed on only the primary ring, however, leaving the secondary one free for restoration in the event of a failure. The secondary ring can optionally be used for data, providing an aggregate throughput of 200M bps over the fiber backbone.

Subtending trees can be connected to the ring. The ring structure can be, but is not generally, extended into the trees, saving fiber and node costs. Standalone trees may be economical in some situations. A mesh, or dual-homing, topology can be implemented to supplement the ring restoration scheme.

FDDI Nodes

FDDI network nodes are categorized as either station nodes or concentrator nodes. Station nodes provide end-user device connections; concentrator nodes can accommodate connections to end-user devices or to other concentrators. A station or concentrator node can support either single or dual fiber connections, depending upon its function in the network. These attributes combine to create four different FDDI network node types:

1. Dual Attachment Station (DAS)
2. Dual Attachment Concentrator (DAC)
3. Single Attachment Station (SAS)
4. Single Attachment Concentrator (SAC)

A more recently defined fifth node type, called a Null Attachment Concentrator (NAC), permits a ringless tree in a star topology.

Dual-attachment nodes are used to establish connections to the primary and secondary FDDI rings; these connections are referred to as port types A and B, respectively. As seen in Figure

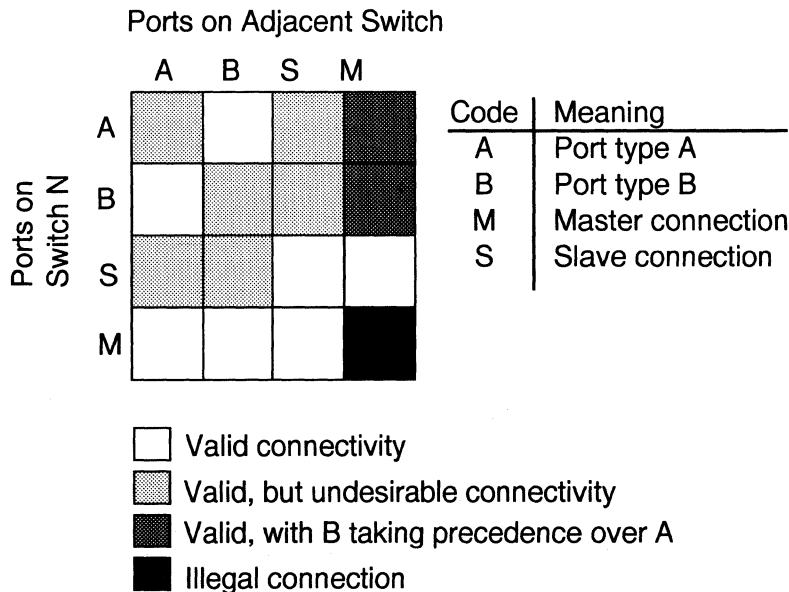



Figure Port Connections

Port connectivity matrix.

Figure
Failure Recovery

Ring rearrangement under failure.



"FDDI Configuration," Port A connects to the Primary Input (PI) and the Secondary Output (SO); port B connects to the Primary Output (PO) and Secondary Input (SI). In this way multiple FDDI nodes are daisy-chained together to form the FDDI primary and secondary ring.

Dual attachment nodes isolate the FDDI ring from all non-FDDI network elements. Their master (M) ports support connections to other FDDI concentrators and stations, which in turn support connections to non-FDDI devices and subnetworks through their master ports. A Logical Link Control (LLC) mechanism is involved in each connection to a non-FDDI device. Allowable attachments are summarized in Figure "Port Connections." The physical interfaces used, called Media Interface Connectors (MICs), are unique to each allowable configuration and are keyed to prevent inappropriate connections.

Restoration

Ring restoration is primarily accomplished using a wrapping mechanism. (See Figure "Failure Recovery.") As already explained, the primary ring carries data traffic in one direction around the ring. Originating traffic leaves the DAS or DAC on the ring via the Port B Primary Output (PO), and it is received at the destination node Port A Primary Input (PI). In the event that a severed fiber blocks communications between nodes, the wrapping mechanism becomes enabled. The traffic leaving the originating node Port B Primary Out is wrapped back to the Port A Secondary Out (SO), and it is received at the destination on the secondary fiber path. All data traffic is thus restored to the participating ring nodes.

In the FDDI ring, data enters an FDDI node, which either routes the data out of the ring or passes it to another FDDI node on the ring. If the node fails, it might destroy data destined for other nodes. Therefore, optical switches are installed in the FDDI ring to automatically cause such traffic to bypass the failed node.

In the event a fiber optic link in a tree subnetwork (outside of the FDDI ring) is severed, the tree is isolated. (See Figure "Failure of Class B Stations.") A dual-homing node configuration, shown in Figure "FDDI Configuration Options," provides an alternate route. In a dual-homing configuration, two Master ports—each from a different Dual Attachment Concentrator on the main ring—are connected to the A and B ports of a second-tier FDDI node. One port is used as the primary connection to the ring, while the other is in a hot-standby mode in case of failure of the

primary. This approach consumes a minimum of ring connections while providing a restoration path for mission-critical nodes.

Relationship to the OSI Reference Model

To fully understand the FDDI standard, some knowledge of the OSI Reference Model and how it applies to the LAN environment is necessary. (See Figure "Standards Relationship Matrix.") A brief discussion of the OSI model is followed by a more detailed description of FDDI specifications. Readers familiar with this material may skip to the section entitled **FDDI Specifications**.


The OSI Reference Model imposes order and structure on data communication, which often involves multiple, concurrent processes. It is natural to group these functions into layers which share task affinity and logical proximity. OSI layers are hierarchical; each layer calls for the services of the layer immediately beneath it. No layer can ask for the services of a higher level layer, skip the layer beneath it to directly reach an even lower layer, or jump into the middle of another layer. The OSI model precisely defines both the services each layer provides to the next higher layer and the request procedures.

Services defined for a given layer are utilized by the layer immediately above it. For outgoing transmissions, each layer passes down to the next lower layer blocks of data requiring processing for transmission, manipulation, or service. At the receiving end, this information is passed from each layer to the next higher layer. These layers normally attach a characteristic header that contains appropriate information (such as the real network address, block number, etc.). The headers are physically nested, with lower layer headers being outermost and higher layer headers being innermost. It is through the use of these well-defined headers that the protocols between the remote open systems are executed.

To effectuate the OSI model, the International Organization for Standardization has formulated standards—specifications for how information is to be coded and passed between communicating partners. The prospective equipment vendor needs only to implement the protocols, but it must employ all seven layers of the architecture.

The Reference Model and the service definitions are only structures for discussing the tasks involved in communicating between open systems. The OSI Reference Model is described by document ISO 7498, which was adopted as a standard in 1984.

Figure Class B Failure



Failure of Class B stations.

LAN Standards

Local Area Network (LAN) standards, defined by the Institute of Electrical and Electronic Engineers (IEEE), are connected with the lower two layers of the OSI model. Requirements are outlined for many types of LANs—including the popular Ethernet and token-ring networks. The IEEE standards divide the OSI Data Link layer into two sub-layers. The lower one is called the *Media Access Control (MAC)* sub-layer; the upper one is the *Logical Link Control (LLC)* sub-layer. This division facilitates moving across LAN sub-networks without leaving OSI layer two. From the LLC sub-layer and upwards, all LANs are generally compatible with the specifications of other protocol stacks.

FDDI Specifications

The IEEE does not define FDDI. This task has been left to the American National Standards Institute (ANSI), whose ANSI X3T12 standard follows both the OSI and IEEE recommendations. This standard accepts the MAC sub-layer of IEEE, and it further divides the OSI Physical layer into two sub-layers: the *Physical Protocol (PHY)* and the *Physical Medium Dependent (PMD)*. This additional division adds support for multiple media

types. In addition, ANSI defines Station Management (SMT) responsibilities that encompass the PHY, PMD, and MAC sub-layers.

Most ANSI FDDI standards have been or are in the process of being adopted as International Standards by the ISO. These core standards can be referenced as:

- Physical Medium Dependent (PMD: ANSI X3.166-1990/ISO 9314-3:1990)
- Physical Layer Protocol (PHY: ANSI X3.148-1988/ISO 9314-1:1989)
- Media Access Control (MAC: ANSI X3.139-1987/ISO 9314-2:1989)
- Station Management (SMT: ANSI X3.229-1994, Revision 7.3)

Many other FDDI-related standards are in development; most of them are updates to existing standards or are related to an FDDI extension (FDDI II), however.

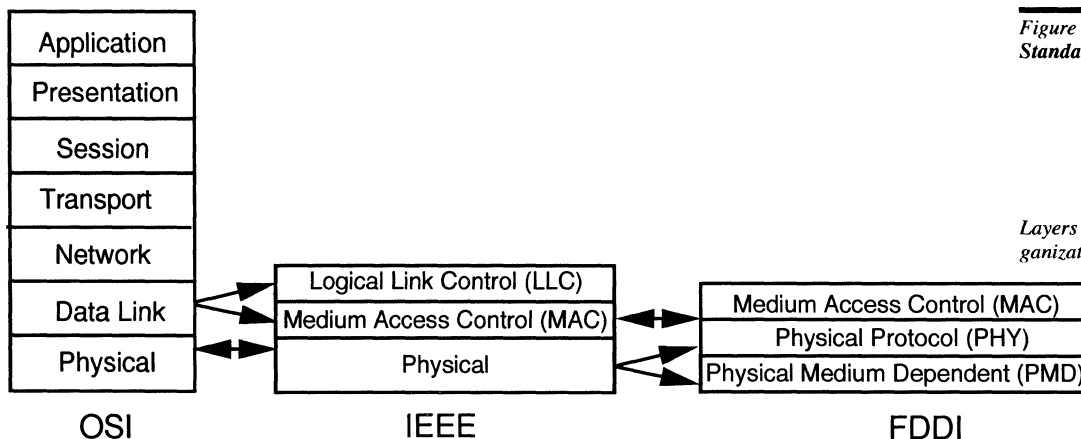



Figure Standards Relationship Matrix

Layers defined by standards organizations.

Figure
Media Interface Connector



An example of a Media Interface Connector (MIC) plug.

Physical Medium Dependent (PMD) Specification

Physical Medium Dependent (PMD) specification defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connectors, the functions of the optical receiver, and optionally, the bypass switch that can be incorporated into an FDDI station. PMD specifies the optical channel at the bulkhead of a station. The source is defined to radiate at a wavelength of 1300-nm. PMD also describes the peak optimal power, optical rise and fall times, and jitter constraints. The standard includes the following specifications:


1. PMD to PHY and PMD to SMT Services
2. Media Types
3. Media Interface Connectors
4. Media Interface Signals

Media Interface Connectors

PMD defines the duplex connector used for FDDI access (see Figure "Media Interface Connector"). The primary and secondary ring connections to each Class A connection are made simultaneously using the duplex connector and a dual-fiber cable. The connectors can be used for both Class A-to-Class A as well as

Figure
Comparison Coding Schemes

Comparison Coding Schemes
used in LANs and/or fiber.



In Manchester code, there is a transition at the middle of each bit period. The mid-bit transition serves as a clock and also as data. A high-to-low transition represents a 1, and a low-to-high transition represents a 0.

In Differential Manchester code, the mid-bit transition is used only to provide clocking. The coding of a 0 (1) is represented by the presence (absence) of a transition at the beginning of the bit period.

Class B-to-Class A (wiring concentrator) links. Four physical media connectors (one for each port type) are defined:

1. Dual Access Station/Concentrator Primary In/Secondary Out
2. Dual Access Station/Concentrator Primary Out/Secondary In
3. Single Access Station (SAS) inputs from concentrators
4. Connectors for SAS stations attached to concentrators

The bypass relay connects the optical inputs (at the primary and secondary rings) directly to the optical output in case of a station or link failure, allowing the ring to maintain continuity.

Media

FDDI was originally designed to use multimode fiber. It was enhanced, however, to accommodate other types of fiber as well as conventional copper twisted pair. This enhancement improved performance and reduced costs.

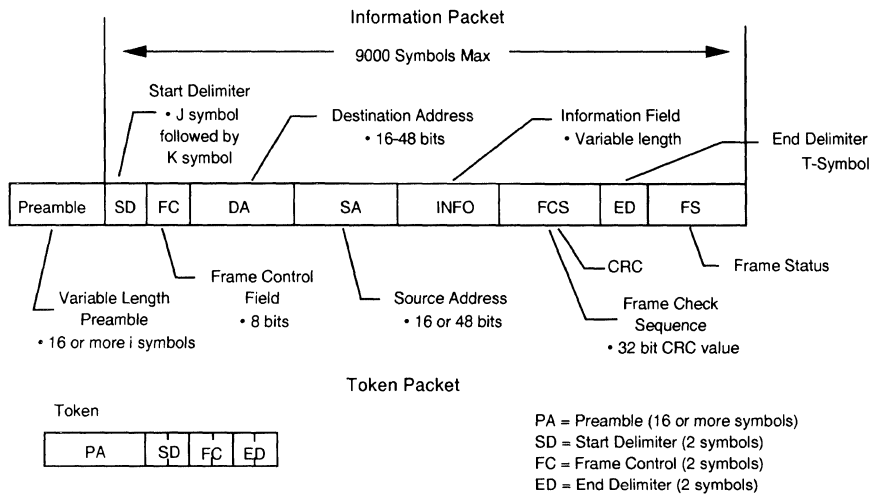
Fiber

Multimode fiber operates in the 1300 nanometer light region, with defined core/cladding dimensions of 62.5/125 micrometers and a maximum distance of two kilometers between repeaters. As an alternative, more economical fiber cabling with dimensions of 50/125 and 85/125 can be used. Single-mode fiber, with core/cladding dimensions of 8/125 to 10/125 micrometers, was introduced to support distances of 40 to 60 kilometers without signal regeneration. Low-cost fiber with 62.5/125 to 200/230 micrometer dimensions was defined for distances not exceeding 500 meters. The cost savings over multimode fiber systems—up to 30%—are somewhat misleading. They come mostly from decreased transmitter electronics expenses rather than from fiber cost savings.

Twisted Pair Copper Cable

Shielded twisted pair (STP) (type 1/2 150 Ohm) and unshielded twisted pair (UTP), (Category 5 100 Ohm) cable are defined for use with FDDI. These non-fiber connections are less expensive. To conform to strict electrical and electromagnetic requirements, however, it is necessary to modify the encoding scheme. Consideration is also being given to supporting Category 3 UTP (much

Figure
FDDI Packet Format.



cheaper than Category 5), which is used extensively in everyday telephone installations as well as in popular Ethernet 10BASE-T LAN installations.

FDDI Over Sonet

FDDI can be mapped into Synchronous Optical Networks (Sonet) transmission systems, which carriers use to transport information over long distances via single-mode fiber. Network designers, however, must carefully take into account distance delays and other considerations. See *ANSI Standard T1.105-02-1994: Sonet Mapping* for an explanation of these issues.

Media Interface Signals

In the 1300-nm. region, the dispersion due to multimode interference is minimal. PMD assumes that if LEDs are being used, they are either the surface-emitting or edge emitting type. PMD does not require, however, that the emitter be an LED. It permits lasers, as long as the optical interface parameters are within recommended values. In the future, some FDDI equipment manufacturers may begin to offer lower cost local-loop-grade laser emitters or long-haul-grade LEDs whose optical output conforms to FDDI standards. For the foreseeable future, however, all manufacturers pursuing FDDI products will use LEDs.

FDDI standards define the input/output signal characteristics for interfaces—including parameters for bit error rate, mean output, rise and fall time, and signal detection time.

Physical Protocol (PHY) Sublayer Specification

The Physical Protocol (PHY) sub-layer, which corresponds to the upper region of OSI's Physical Layer, defines the encoding scheme used to represent data and control information. PHY also describes the method for retiming transmission within the node. The standard includes the following specifications:

- Services
 - PHY to MAC Services
 - PHY to PMD Services
 - PHY to SMT Services
- Facilities
 - Coding
 - Symbol Set

—Line States

Digital data must be encoded in some form for proper transmission. The type of encoding depends on the nature of the transmission medium; the data rate; and other factors such as noise, reliability, and cost. Given the fact that fiber is inherently an analog medium, a digital-to-analog conversion technique is required. Intensity modulation is the norm for fiber; a binary 1 is represented by a pulse of light and a binary 0 by absence of optical power. This simple approach cannot reliably maintain synchronization between the clocks of the sending and receiving devices when long strings of zeros are being transmitted. The solution is to first encode the binary data in such a way as to guarantee the presence of signal transitions—even if there are no transitions in the incoming digital signal. After this encoding step has been performed, the signal is presented to the optical source for transmission using intensity modulation. A typical encoding scheme is Manchester encoding (see Figure "Comparison Coding Schemes").

Differential Manchester encoding is only 50% efficient, since each data bit is represented by transitions in signal. Two transitions allow a degree of robustness in the presence of noise, as would be the case in coaxial cable. Since fiber is less susceptible to noise, two transitions are not required to identify a bit with a good degree of confidence.

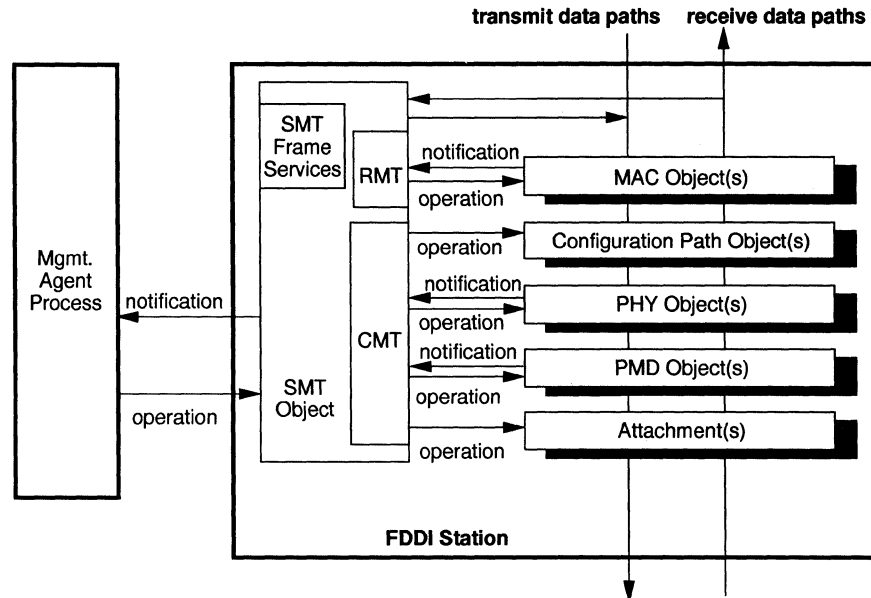
To avoid having to use a 200MHz signal, FDDI specifies a code referred to as 4B/5B group encoding. The result is that the 100M bps throughput is achieved in FDDI with a 125MHz clock rate, rather than the 200MHz clock rate needed in differential Manchester. This encoding helps keep down the cost and complexity of equipment.

One drawback of the group encoding concerns clock recovery. Since differential Manchester has more pulses in its stream, it is easy to extract the clock in that scheme. One of the key responsibilities of the PHY sublayer is to decode the 4B/5B nonreturn to zero inverted (NRZI) signal from the network into symbols which can be recognized by the FDDI station, and vice versa.

The synchronization clock is derived from the incoming signal. The data is then retimed to an internal clock through an elasticity buffer. In this scheme, 4 bits of data are translated into a 5 baud value transmitted over the network, giving an 80% efficiency factor. This group-encoding scheme employed in FDDI is a departure from differential Manchester codes normally specified in LAN standards.

Figure
SMT Architectural Model

The relationships between managed objects in an FDDI station and the SMT object.



To understand how the FDDI scheme achieves synchronization, the user must realize that there are two stages of encoding. In 4B/5B, the encoding is performed four bits at a time. Each group of four data bits is encoded into a symbol with five cells, each containing a single signal element (presence or absence of light). In effect, each set of four bits is encoded as five bits. Each element of the 4B/5B stream is then treated as a binary value and encoded using NRZI. In this code, a binary 1 is represented with a transition at the beginning of the bit interval; there are no other transitions.

The advantage of NRZI is that it employs differential encoding: the signal is decoded by comparing the polarity of adjacent signal elements rather than by interpreting the absolute value of a signal element. This scheme is relatively robust in detecting transitions in the presence of noise or other distortions; NRZI encoding, therefore, improves reception reliability.

This scheme encodes 4 bits (16 possible combinations) with 5 bit patterns (32 possible combinations); not all possible patterns are needed. The rules for processing the 4-bit patterns require that a transition occur at least twice for each 5-bit code. In the NRZI format, no more than three zeroes in a row can be allowed, since the absence of a transition would indicate a zero.

PHY also provides information about line states for stations establishing links with neighboring (upstream and downstream) services, enabling them to detect the integrity of those links. These line-state information transfers are also used for handshaking. A node receiving a line-state message on its primary input can respond by sending the proper line-state message on the secondary output. The line states are composed of a repetition of one or more "I" symbols.

Another issue resolved in this sublayer is that of timing jitter. Jitter is the timing deviation that can occur as the receiver attempts to recover the clocking signal as well as data from the incoming signal. The clock recovery can deviate randomly from the transitions of the received signal. If no measures are taken to resolve this problem, jitter will accumulate around the entire ring.

The IEEE 802 LAN standard requires that in local area networks, only one master clock should be used on a ring, and that the station owning the clock is responsible for eliminating jitter through an elastic buffer. As data on the ring runs ahead or behind the master clock, the elastic buffer expands or contracts accordingly. This centralized clocking method is not practical for a

100M bps ring, where a transmission interval of only 10 ns. per bit (compared to an interval of 250 ns. per bit on a 4M-bps token-ring LAN) makes the effect of distortion more severe.

FDDI utilizes a distributed clocking scheme. Each station has an elastic buffer into which incoming data traffic arrives at the clock rate recovered from the incoming bit stream itself. Data leaves the buffer, however, at the station's own clock rate. This distributed approach, which gives each FDDI station its own autonomous clock, is considered superior to the centralized clocking method of token-ring LANs. Letting each station perform its own clocking, moreover, permits an unlimited number of repeaters around the ring.

Media Access Control Specification

In the context of LANs, Layer 2 (the Data Link Layer) of the OSI Reference Model is traditionally divided into two sublayers: Logical Link Control (LLC) and Media Access Control (MAC). FDDI standards deal only with MAC, which is concerned with data flow over the ring. The token-passing protocol incorporated into FDDI controls transmission over the network. MAC defines packet formation (headers, trailers, etc.), addressing, and cyclic redundancy checking (CRC). It also defines the recovery mechanisms. This standard defines the following specifications:

- Services:
 - MAC-to-LLC Services
 - PHY-to-MAC Services
 - MAC-to-SMT Services
- Facilities:
 - Symbol set
 - Protocol data units
 - Fields
 - Timers
 - Frame counts
 - Frame Check Sum

The FDDI packet format is shown in Figure "FDDI Packet Format." Packets are preceded by a minimum of 16 IDLE control symbols. The packet itself is characterized by a Start Delimiter

composed of the J and K control symbols. This is followed by a Frame Control field that identifies the type of packet. The Destination Address, which follows, identifies the frame recipient. The Source Address is also included to identify which station originated the packet. The address field can be 26 or 48 bits in length. The variable information field follows, along with a Frame Check Sequence field of 32 bits. The check sequence covers the Frame Control Field, the two addresses, and the information field. An End Delimiter, which consists of the T symbol, is transmitted. The maximum packet length is limited by the size of the elastic buffer in the Physical Sublayer and by the worst-case frequency difference between two nodes, the upper bound in 9,000 octets. Figure "FDDI Packet Format" also shows the format of the token.

Flow control is the other major function of the MAC sublayer. In an idle condition, MAC connects to an internal source of IDLE control symbols to be transmitted over the ring. When a Start Delimiter is detected from the ring, MAC switches to a repeat path; the packet is monitored and copied if it is meant for this destination. The packet is simultaneously repeated onto the ring for relaying. The MAC can also inject its own packet or issue a token. Packets are removed only by the originating station. The MAC repeats the packet only until the Sender Address field is detected. If the destination recognizes the Sender Address field as its own station, it will insert IDLE control symbols back onto the ring (the fragmented packet is ignored and removed by any station holding a token for transmission). Stations wishing to transmit must first obtain the token.

The procedures for obtaining the token and the amount of time allowed for data transmission (to retain fairness) are specified in the Timed Token Protocol (TTP). A station obtains the token by performing the stripping function on the incoming token. Only the Start Delimiter field is repeated onto the ring; the station will inject its own information at this juncture. When the packet is sent, the station immediately issues a new token. TTP guarantees a maximum token rotation time.

TTP allows synchronous and asynchronous transmission modes. In the synchronous mode, stations obtain a predefined amount of transmission bandwidth on each token rotation. The balance of the bandwidth is shared among stations using the asynchronous service. These stations can send data when the token arrives earlier than expected. Any unused capacity left over from synchronous capacity is available to asynchronous traffic, which may be subdivided into up to eight levels of priority.

The amount of time allowed for asynchronous transmission is bounded by the difference of the token's actual arrival time and the expected arrival time. In essence, each station keeps track of how long it has been since it last saw the token. When it next sees the token, it can send synchronous traffic and/or any asynchronous traffic for which time remains available.

Station Management (SMT)

The FDDI Station Management (SMT) specification describes software-based, low-level data link management and integrated network control functions of all stations attached to an FDDI LAN and of the LAN itself. Each FDDI station contains only one

SMT entity (see Figure "SMT Architectural Model"). SMT initializes the network, monitors error rates and fault conditions in each network segment, and automatically reconfigures the network to isolate problem links. SMT components are Connection Management (CMT), which includes Entity Coordination Management, Physical Connection Management (PCM); and Ring Management (RMT). SMT is intended to operate regardless of equipment type, vendor, protocols, or applications.

SMT types (managed objects) have specific attributes indicating state, capabilities, and operation.

SMT managed objects are:

- Station or concentrator (SMT)
- MAC object(s)
- Path object(s)
- PHY object(s)
- PMD object(s)
- Attachment(s)

Attributes are:

- Attribute Identification (ID)
- Configuration
- Operational (status, counters, etc.)

Each attribute is defined in terms of Access Rights, and whether it is Mandatory or Optional. Each attribute also carries an FDDI specification reference or a specific definition if it is not defined in the FDDI specification.

SMT Connection Management (CMT) operates at the logical level, controlling the interface between PHY and MAC entities in a given station and controlling SMT-to-SMT communications across the ring. When a session requires a connection to another station, CMT causes PHY to send a stream of symbols to the targeted station. Upon receiving the symbols, the receiving station's PHY returns a continuous stream of symbols (primitives) indicating line state, station status, and willingness to carry out the requested action (establish a link). QUIET symbols indicate a disabled link. HALT or alternating HALT and QUIET (MASTER) symbols indicate an operating link and the receiving station's status (master, slave, or peer). A stream of IDLE symbols indicates willingness to connect. Once the link is established, CMT configures the PHY and MAC.

Entity Coordination Management (ECM) controls the optional optical bypass switch and signals the Physical Connection Management (PCM) entity when the bypass is complete. ECM also performs the Path Test to determine a fault's location.

Physical Connection Management (PCM) initializes the adjacent station's PHYs and manages signaling. Maintenance support functions are also part of PCM.

Configuration Management (CFM) interconnects PHYs and MACs. It automatically configures these connections according to PCM flags. CFM is defined differently for stations and concentrators.

Ring Management (RMT) relays MAC and CFM status information. It detects stuck signaling beacons, initiates the trace function, detects duplicate addresses and resolves them to allow continued ring operation, and notifies SMT of MAC status. ■

M. Scott Kingsley is a Systems Engineer and member of the technical staff at Alcatel Network Systems in Richardson, Texas. Mr. Kingsley has 15 years of experience in telecommunications, including jobs at DSC Communications and MCI Telecommunications. Mr. Kingsley holds a Masters of Science degree in Telecommunications Systems Management from Southern Methodist University.

ANSI FDDI Standards

In this report:

Alternative Standards.....	2
Technology Overview.....	3
FDDI Specifications.....	5
Sponsoring Organizations..	13

Datapro Summary

In October 1982, the American National Standards Institute (ANSI), Committee X3T9.5, was chartered to develop a high-speed data networking standard. Known as Fiber Distributed Data Interface (FDDI), the standard specifies a packet switched LAN-to-LAN backbone that transports data at high throughput rates over a variety of multimode fibers. FDDI is a token-passing network employing two fiber pairs operating at 100M bps. Since the release of Version 7.2 of the Station Management (SMT) portion in June 1992, the core FDDI standard is complete.

Analysis

FDDI addresses the bottom two layers of the OSI model. It is the spirit of the FDDI standard to use ISO protocols for the other layers, where possible. The optical-based FDDI LAN was designed to provide the same type of serial interconnection provided by LANs while providing the high bandwidth, inherent noise immunity, and security offered by fiber.

At FDDI's inception in 1982, fiber was used mostly for point-to-point applications, and not for the many configurations allowed by LANs. In this sense, FDDI was a breakthrough. Although it is possible to achieve considerably higher data rates over fiber (up to 3.7G bps with current point-to-point technology, and 500M bps on rings), higher rates result in significantly increased costs and shorter transmission distances between repeaters. Its designers intend FDDI to provide relatively inexpensive connectivity and, therefore, focused on the 100M bps rate. FDDI can be configured to support a sustained transfer rate of approximately 80M bps. The remaining bandwidth is reserved for various overhead functions.

—By Timothy McElgunn
Associate Analyst

Applications

Although optical fiber is widely deployed in the telecommunications environment (long hauls, interoffice, feeder plant, and so on), it has not done as well in the LAN arena. Three reasons can be identified: 1) increased technical complexity compared to passive copper and coaxial cable; 2) cost considerations; and 3) lack of a workable standard. FDDI solves the third problem and in the process begins to resolve the second.

FDDI allows designers to 1) build larger capacity LANs or LAN backbones to serve new data needs (file transfer, graphics, and so on) and some voice needs; or 2) interconnect LANs in metropolitan area networks (MANs). Thus, FDDI can be used directly as a LAN or as a backbone to interconnect slower-speed LANs into a single network over relatively large geographies.

The initial application of FDDI as a "back-end" interconnect for high-powered computing devices and peripherals required a high degree of fault tolerance and data integrity. As developers proceeded, it became obvious that FDDI could also serve high-speed "front-end" applications. Front-end applications include terminal-to-terminal and terminal-to-server communications typical of a LAN. In large networks (3,000 to 10,000 terminals, particularly when workstations are involved), the aggregate demand for network resources can overwhelm a 4M bps or 10M bps LAN. At that point, FDDI's bandwidth becomes important.

A high-speed FDDI ring is ideal as a backbone for other "departmental" LANs, typically operating at lower speeds. High capacity on LANs can be achieved in two ways by using multiple channels at low speeds, or by using one channel at a relatively high speed. The multiple-channel approach can be used with a broadband bus LAN. A drawback of this approach is that bridges must be provided between channels, and the architecture must be designed to avoid high rates of interchannel traffic and bottlenecks at the bridges.

Desktop systems that are likely to be linked via high-speed networks are now applied in realtime simulations, graphic transfers in a CAD/CAM environment, supercomputer terminals, medical imaging, and video.

The one area where FDDI is not well suited is broadband LANs carrying full analog video (6MHz). A broadband LAN can easily carry multiple channels of video as well as data. Digitizing a 6MHz TV channel results in a 45M bps to 90M bps data rate (45M bps is a slightly compressed version), which can easily swamp an FDDI backbone. FDDI II, discussed later in this report, has been designed to handle video efficiently.

FDDI is not the end, however. Even larger bandwidths are envisioned in the next five years under the thrust of Broadband ISDN (B-ISDN). BISDN for video needs, particularly in a High Definition TV (HDTV) environment, will require approximately 150M bps of dedicated bandwidth per channel. Planners are now discussing delivery of up to three channels per domicile, reaching into the 600M bps range. FDDI is a shared-medium technology; uncompressed video applications and/or local loop applications may not be capable of using networks built on the original FDDI standards.

Alternative Standards

High-speed LAN transport, until recently, was a question with a single answer—FDDI. In the past year, however, new technologies have proliferated, at least in the pages of the trade weeklies and in users' consciousness, if not on distributors' loading docks. The grand vision of all the workstations in the world linked by gigabit-per-second ATM transport is doubtless several years from becoming a reality, and several vendors have rushed into the gap with proposals for alternative high-speed LAN technologies.

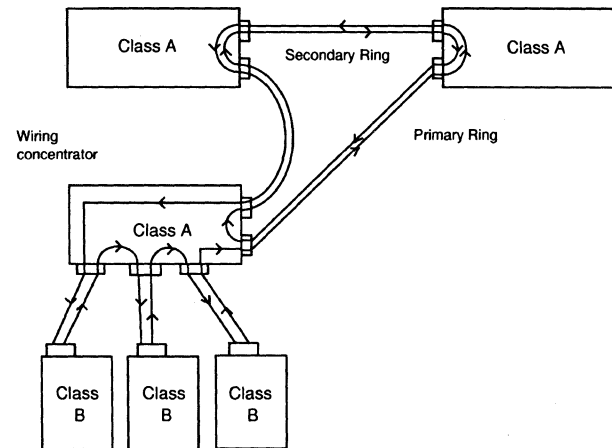
FDDI Variations

A number of vendors and vendor consortia have developed FDDI implementations operating on unshielded twisted pair wire. This is envisioned as a "drop" medium from the wiring closet to the work area. While UTP is considerably less expensive than fiber, distances will be limited. These implementations are called twisted-pair DDI (TPDDI) and copper DDI (CDDI). The vendor's rationale for developing FDDI-over-STP technology is to allow users to employ existing cabling when migrating from existing LANs to FDDI without installing fiber optic cabling.

FDDI proponents have advanced three separate proposals—two for shielded twisted pair and a third for unshielded twisted pair. The initial "Green Book" recommendation for STP, endorsed by Digital Equipment Corp., Advanced Micro Devices, Motorola, Chipcom, and SynOptics, among others, was never implemented. IBM's SDDI (Shielded twisted-pair Distributed Data Interface) was championed by ten other vendors including Cabletron, National Semiconductor, Motorola, Chipcom, and SynOptics.

SDDI, with the backing of the main industry proponent of shielded twisted pair, IBM, took longer to die than the Green Book scheme, but ultimately succumbed to the ANSI X3T9 committee's desire to field a single proposal that covered both unshielded and shielded twisted pair. When the committee made that preference known, there was only one company with a UTP

Figure 1.
The FDDI Environment



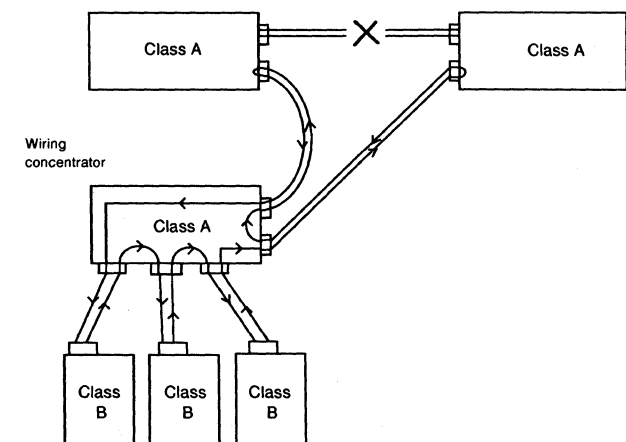
product ready to ship—start-up Crescendo Communications. To address the problems experienced in trying to transmit 100M bps signals over copper, Crescendo developed a new encoding and scrambling scheme. While the vendor claims that the required coding translation is transparent to the user, the scheme requires a new chipset design.

Several vendors have announced their support for CDDI, including Hewlett-Packard; AT&T; Ungermann-Bass; Silicon Graphics; Cabletron; and SynOptics. In September 1993 Crescendo was acquired by router heavyweight Cisco Systems.

Codenoll Technologies introduced a nonstandard version of the FDDI interface board for Extended Industry Standard Architecture (EISA) computers. While using an 850-nm. LED in the Codenoll design does reduce the cost of FDDI connectivity, users will be unable to operate with standard-based equipment containing 1300-nm. transmitters and receptors.

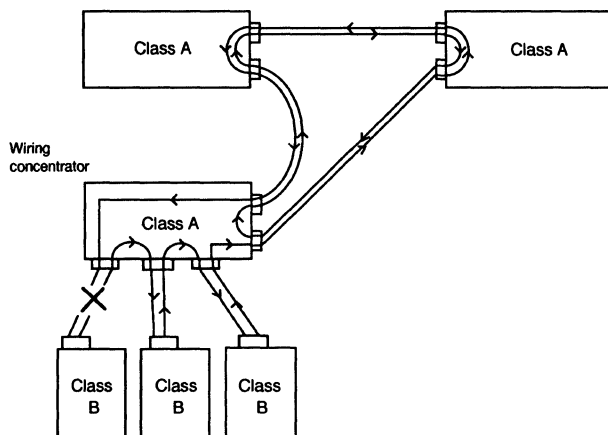
Finally, a low-cost alternative fiber PMD provides a cost enhancement for networks that do not require the standard's full distance.

Figure 2.
Failure Recovery



Ring rearrangement under failure.

Figure 3.
Class B Failure



Failure of Class B stations.

Alternative 100M bps Proposals

The slow acceptance of FDDI has been attributed to its high cost per seat on the network, as well as the perceived technical difficulty of implementing the technology. Until recently it required fiber optic cabling. Bridging between FDDI and lower speed networks was problematic, especially for token-ring with its source route bridging scheme.

Grand Junction, a start-up company, proposes a faster carrier sense multiple access with collision detection (CSMA/CD) network—that is, one that retains the Ethernet media access control (MAC) layer. This makes 10M bps-to-100M bps bridging relatively simple and could, in fact, deliver the lowest cost and quickest time to market for new 100M bps transports. This is borne out by the fact that Grand Junction is already shipping its FastSwitch products. Other Fast Ethernet vendors, including 3Com, are talking up their intention to market switchable 10M/100M adapter cards that will supposedly give users greater flexibility in upgrading individual workstations connections to the higher speed. 3Com has also announced that its fast Ethernet products will run on various grades of unshielded twisted-pair wire—including the most common type—Grade 3. Grand Junction's initial products require Grade 5 UTP.

Late in 1993, Hewlett-Packard, in association with IBM, brought token-ring into the 100M bps fold. In fact, that is what prompted the name 100VG-AnyLAN—Hewlett-Packard's design was originally dubbed 100BASE-VG. In either case, the VG stands for "voice grade" indicating that Hewlett-Packard's product will deliver 100M bps over Category 3 cabling. Category 3 is the type installed in most sites that are wired with twisted-pair today.

100VG-AnyLAN drives local area networking closer to a connection-oriented paradigm. Its Demand Priority access protocol uses a handshaking sequence in which the workstation sends a Request-To-Send query to the hub, and the hub, if all is well, responds with an Acknowledge-To-Send. This is in stark contrast to both traditional Ethernet, in which workstations start sending as soon as they sense a clear line, and Ethernet switching, which in most cases is a store-and-forward technology. In the first case, the hub plays no role in transmission at all; in the second, the Ethernet switch acts on packets as they are received, but no direct interaction with the transmitting station takes place.

Asynchronous Transfer Mode (ATM)

ATM is a high-bandwidth, low-delay *switching and multiplexing* packet technology; ATM communication is a connection-oriented process, although it is designed as a basis for supporting both

connectionless and connection-oriented services. Asynchronous transfer mode packetizes data into 53-byte cells, including a 5-byte destination header. The use of these small, fixed-length cells gives ATM the multiplexing efficiency of X.25 packet switching with the low delay and high throughput of Time-Division Multiplexing (TDM). Under ATM, a number of the layer 2 data link protocol functions are removed to the edge of the network. Core layer 2 capabilities are supported, however, in addition to layer 1 functions (clocking, bit encoding, physical medium connection).

ATM switches' high speed and guaranteed bandwidth availability allow users to transmit local area voice and video in addition to data traffic. Contention-based local area networks like Ethernet cannot transport time-dependent traffic efficiently, requiring users to dedicate the LAN bandwidth or risk losing portions of the audio or video signal, which is generally insupportable. ATM permits bit rate allocation on demand, so the bit rate per connection can vary according to demand or design. In addition, the actual "channel mix" at the broadband interface can dynamically change. ATM supports labeled channels operating at any fixed rate ranging from $N \times K$ bps up to the total payload capacity of the interface—potentially gigabits per second.

FDDI-II and Other Future Efforts

Vendors have also pursued a scheme where available network bandwidth is divided between voice and data using time-division multiplexing. FDDI-II, a superset of FDDI, is an upward-compatible, fiber-based LAN incorporating the current data capabilities in addition to the ability to handle voice and T1-compressed video traffic.

The FDDI-II standard follows the original FDDI standard, adding a fifth document to the present standard. The new document, Hybrid Ring Control, describes a hybrid operating mode comprising the packet-switching scheme used in FDDI and an isochronous transport mode similar to that used in public switched networks. Adding the isochronous mode will enable networks based on the expanded standard to transfer pixel data—key to video and computer graphics transport—in addition to circuit switched voice signals.

Some people are already thinking of even higher rates, particularly when considering voice and data. Development of an FDDI follow-on, or FFOL, is under consideration at ANSI. The FFOL project proposal covers the capability to operate as a backbone for multiple FDDI networks; interconnection to wide area networks, including Broadband Integrated Services Digital Networks (B-ISDN); a data rate between 600M bps and 1.25G bps initially, with intermediate data rates matched to the Synchronous Digital Hierarchy (SDH) underlying Sonet and eventual support for data rates up to 2.4G bps; duplex links; support for existing FDDI cabling, where possible; and support for both single mode and multimode fiber. FFOL interest appears to be waning, however, as asynchronous transfer mode cell switching and other technologies capture increased interest.

Technology Overview

FDDI is a token-passing, dual-ring network accommodating synchronous and asynchronous data transmission, as well as isochronous channels for realtime digitized voice and compressed video.

Figure 4.
OSI Functions

APPLICATION LAYER (7)
–Application-oriented functions (example: file transfer) –Management functions
PRESENTATION LAYER (6)
Data formats, codes, and representation
SESSION LAYER (5)
Dialogue control and synchronization, initialization
TRANSPORT LAYER (4)
End-to-end control of transportation of data, optimization of usage of network resources
NETWORK LAYER (3)
Forwarding blocks of data over a network route towards their final destination
LINK LAYER (2)
Reliable transfer of blocks of data between adjacent nodes (error detection and recovery)
PHYSICAL LAYER (1)
Control of data circuits, transfer of bits
PHYSICAL MEDIA
Transmission of electrical signals

Functions performed by the layers of the OSI Reference Model.

Unlike existing open standards for LANs, where fiber optic variants have followed copper implementations, FDDI has been designed from the start as a fiber optic network. This involved standardizing such features as duplex optical connectors, fiber characteristics, optical bandwidth, bypass relays, and cable assemblies. The FDDI ring is designed for an overall bit error rate of less than 10^{-9} . The network can tolerate up to 2 km. of fiber between stations and can support a total cable distance of 100 km. around the ring with 500 attachments (1,000 physical connections and a total fiber path of 200 kilometers). FDDI topology is a counterrotating token-passing ring (note the arrows in Figure 1). FDDI, however, is not part of the well-established IEEE 802 family of LAN standards. In the case of FDDI, copper variants have followed the fiber versions, as designers attempt to reduce costs and ease installation.

FDDI operates at 1300 nanometers (nm.). Current transmitter/receiver fiber technology operates at 850 nm., 1300 nm., or 1550 nm. While performance increases with wavelength, so does cost. For local data communications, both in LANs and in point-to-point applications employing fiber optic modems, the 850-nm. light sources are typically employed; however, this technology becomes unfeasible for 100M bps beyond a couple of miles. At the other end of the range (1550 nm.), the system becomes expensive and may provide unnecessary bandwidth.

The committee designing FDDI also investigated short-wavelength implementations. It became evident that to meet all the

requirements, particularly the two-kilometer, station-to-station spacing, the system would require 1300 nm. wavelength. Using 1300 nm. technology, less expensive light-emitting diodes (LEDs) provide distance and data rates within the range desired for LANs and LAN backbones. The issue of fiber size was settled sometime after the wavelength decision.

The FDDI standard directly addresses the need for reliability. A backbone system transports a large number of user sessions, and its loss would be serious. FDDI incorporates three reliability-enhancing features. First, a failed or unpowered station can be bypassed by an optional automatic optical bypass switch; second, wiring concentrators are used in a star-wiring strategy to facilitate fault isolation and correction; third, two rings interconnect stations so that failure of a repeater or cable link results in the automatic reconfiguration of the network (loopback) (see Figures 2 and 3).

To fully understand the FDDI standard, some knowledge of the OSI Reference Model and how it applies to the LAN environment is necessary.

The OSI Reference Model imposes order and structure on data communication. To achieve orderly communication, many functions must be performed. It is natural to group these functions into layers which share task affinity and logical proximity (see Figure 4). OSI layers are hierarchical in the sense that a given layer calls for the services of the layer immediately beneath it. The given layer cannot ask for the services of a layer at a higher level, nor can it skip the layer beneath it to directly reach a lower layer, or even jump into the middle of another layer. The model includes precise definitions of the services provided by each layer to its next higher layer and request procedures.

Services defined for a given layer are, in turn, employed by the layer immediately above it. Each transmitting layer passes down to the lower layer (and up to the next layer, at the receiving end) blocks of data requiring processing for transmission, manipulation, or service. These layers normally attach a characteristic header that contains appropriate information (such as the real network address, block number, and so on). The headers are physically nested, with lower layer headers being outermost and higher layer headers being innermost. It is through the use of these well-defined headers that the protocols between the remote open systems are accomplished.

To effectuate the OSI model, the International Organization for Standardization has formulated standards—specifications for how information is coded and passed between communicating partners. Only the protocols need to be implemented by a prospective vendor. Users should note that when working at a terminal connected to a host, or a PC connected to a LAN, all seven layers of the architecture must be employed.

The Reference Model and the service definitions are only structures for discussing the tasks involved in communicating between open systems. The OSI Reference Model is described by document ISO 7498, which was adopted as a standard in 1984.

LAN Protocol Suites

In a LAN environment, one typically defines Layer 1 and Layer 2 standards specific to local area networks. Layers 1 and 2 are defined by the IEEE 802 standards. With some "internetworking" protocols defined at Layer 3 (such as ISO 8473), and typically some connection-oriented Transport protocols (for instance, FTAM and MHS), sessions can then employ the normal protocol suite up to Layer 7 described above. Prior to the establishment of ISO-based standards, however, the use of Transmission Control Protocol/Internet Protocol (TCP/IP) was common for the upper layers. Below, one finds descriptions of the standards at Layer 1 and Layer 2, which are specific to LANs.

Logical Link Control

IEEE Standard 802.2-1985 (ISO 8802/2) describes the peer-to-peer protocol procedures for the transfer of information and control between any pair of Data Link Layer Service Access Points (SAPs) on a local area network.

Medium Access Control: Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

IEEE Standard 802.3-1985 (ISO 8802/3) provides a medium access method by which two or more stations share a common-bus transmission medium. The standard applies to several media types and provides the necessary specifications for a 10M bps baseband local area network.

Medium Access Control

Token-passing bus access method IEEE Standard 802.4-1985 (ISO 8802/4) deals with all elements of the token-passing bus access method and its associated physical signaling and media technologies. The goal is to achieve compatible interconnection of stations in a LAN. The access method coordinates the use of the shared medium among the attached stations. It specifies the electrical and physical characteristics of the transmission medium, the electrical signaling used, the frame formats of the transmitted data, the actions of a station upon receipt of a data frame, and the services provided at the conceptual interface between the Medium Access Control sublayer and the Logical Link Control sublayer above it.

Medium Access Control: Token-Passing Ring Access Method

IEEE Standard 802.5-1985 (ISO 8802/5) specifies the formats and protocols used by the token-passing ring medium at the control sublayer and Physical Layer. It also specifies the means of attachment to the token-passing ring access method. The protocol defines the frame format including delimiters, addressing, and frame check sequence and includes timers, frame counts, and priority stacks. It also defines the medium access control protocol and provides finite-state machines and state tables supplemented with descriptions of the algorithms. It identifies the services provided by the Medium Access Control sublayer to the Logical Link Control sublayer, and the services provided by the Physical Layer to the Medium Access Control sublayer. These services are defined in terms of service primitives and associated parameters. Also defined are the Physical Layer functions of symbol encoding and decoding, symbol timing and latency buffering, and the 1M bps and 4M bps shielded twisted-pair attachments of the station to the medium.

Figure 5, modeled after documentation of the German Commission for Computer-Integrated Manufacturing, depicts typical OSI protocol suites for LANs, particularly in a MAP/TOP environment. Figure 5 also depicts some additional details at the lower layers, including the positioning of FDDI juxtaposed to other LAN standards. Once again, note that when a user works at a PC connected to a LAN or FDDI ring, all seven layers of the protocol architecture must be employed, as shown in Figure 5. Thus, while the FDDI standard concentrates on Layers 1 and 2, the upper layers are required to accomplish end-to-end communication.

FDDI Specifications

FDDI is defined according to the OSI Reference Model and LAN protocol architecture. Layer 1 (Physical Layer) is specified in two documents: the FDDI Physical Medium Dependent (PMD) and

the FDDI Physical Sublayer (PHY). (See the following sections.) The Physical Layer provides the medium, connectors, optical bypassing, and driver/receiver requirements. It also defines encode/decode and clock requirements for framing data for transmission on the medium or to the higher layers of FDDI.

When the FDDI committee realized there would be considerable discussion on fibers, connectors, and other hardware, it decided to break the standardization of the OSI Physical Layer into two pieces. In this way, the relatively noncontroversial issues—such as coding and other matters that IC chip manufacturers need to know to begin design—could be put in a formal document and approved independently of other items pertaining to the standard.

The Data Link Layer is also divided into two sublayers:

1. A Media Access Control portion that provides fair and deterministic access to the medium, address recognition, and generation, and verification of frame check sequences. Its primary function is the delivery of frames, including frame insertion, repetition, and removal.
2. A Logical Link Control portion that provides a common protocol for data assurance services between the Media Access Control and the Network Layer.

Physical Medium Dependent (PMD) Specification

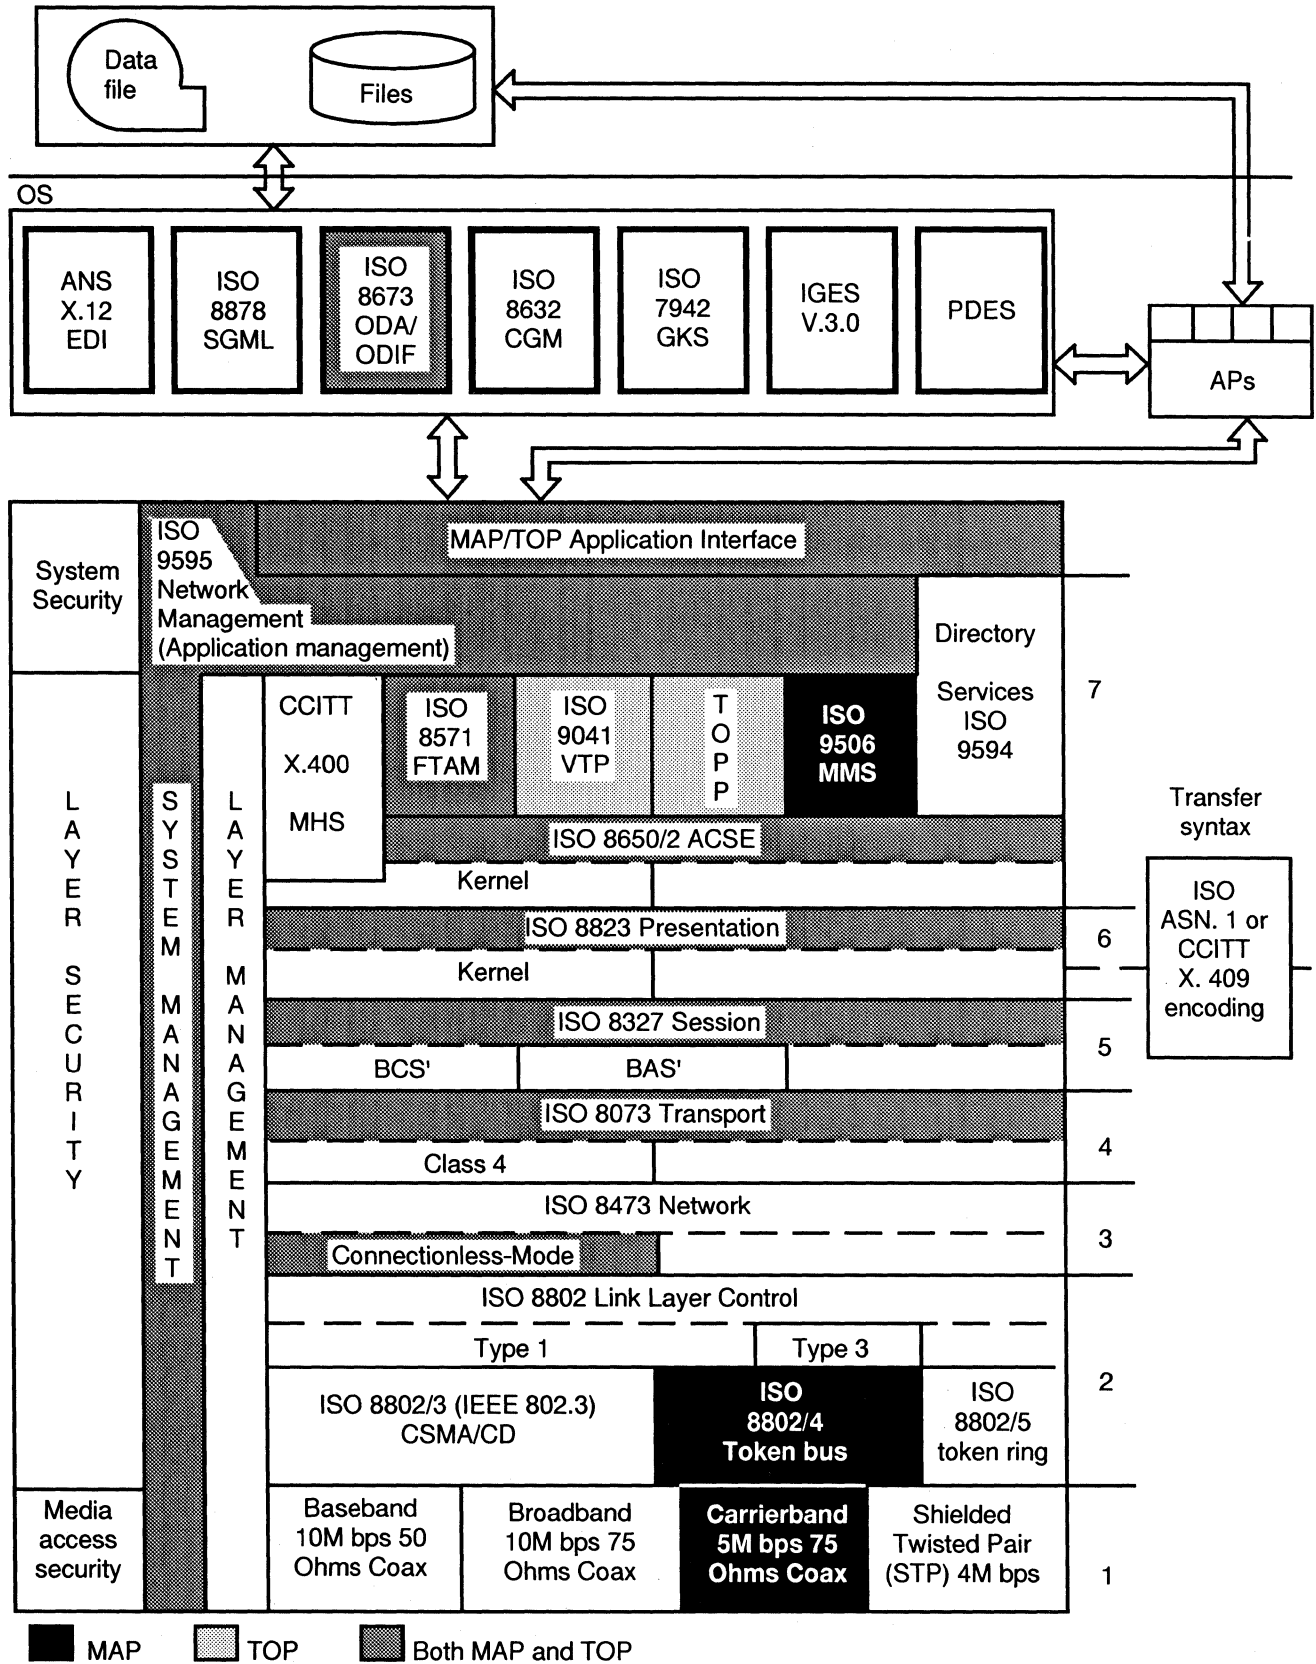
PMD defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connector, the functions of the optical receiver, and (as an option) the bypass switch that can be incorporated into the station. It specifies the optical channel at the bulkhead of a station. The source is defined to radiate in the 1300-nm. wavelength. PMD also describes the peak optimal power, optical rise and fall times, and jitter constraints. The minimum rise/fall time is 0.5 nanosecond. The standard includes the following specifications:

1. Services: PMD to PHY Services; PMD to SMT Services
2. Media attachment
3. Media Signal Interface
4. Interface signals
5. Cable Plant Interface Specification

Multimode fibers are employed (at least initially) up to a distance of two kilometers. Optical fiber dimensions are specified in terms of its core diameter and the outer diameter of the cladding layer. Fiber specifications are 62.5/125 micron (core diameter/cladding diameter) and 85/125 micron. The nominal numerical aperture is around 0.26. Applicable standards for the fiber itself are EIA-455-48 (core), EIA-455-27 or EIA-455-48 (cladding), and EIA-455-57 (aperture).

Listed in the appendix of the standard are two other fibers: a 50/125 micron fiber and a 100/140 micron fiber. These two fibers have not been extensively studied; the maximum achievable distance of 2 km. specified in the standard may not be possible. Thus, these two fibers are not officially part of the standard but are listed as "alternatives." Smaller diameters offer higher bandwidths but also more expensive, higher loss connectors. The 50/125 fiber has been used primarily for military applications, and is not likely to be as widely available as the other fibers. The 100/140 fiber has been added primarily because it is used in IBM's cabling system, and a number of customers have already installed it. On the other hand, the 62.5/125 fiber has been in production for some time and has become common for local applications, such as AT&T's Premises Distribution System (PDS). Component costs for this type of fiber are dropping most rapidly.

Figure 5.
ISO Suites



ISO suites in a LAN environment.

Some observers believe it is unfortunate that the FDDI standard could not specify a single fiber type, since this might have lowered costs and made it easier for the customer to start small and expand later. The FDDI committee settled on a 62.5-micrometer core, with advisory information about 50-, 85-, and 100-micrometer fiber sizes. While a debate remains about these other sizes, the issue is not really a critical one: as long as a fiber meets the optical power, channel bandwidth, and distance requirements, it can conform to the FDDI standard at the interface between the FDDI box and the network, independent of the fiber type (the charts and data presentations in the PMD have been written so that they can be applied to various fiber sizes).

Like all Layer 1 specifications, PMD defines the duplex connector used for FDDI access (see Figure 6). The primary and secondary ring connections to each Class A station are attached simultaneously using the duplex connector and a dual-fiber cable. The connectors can be used for both Class A-to-Class A as well as Class B-to-Class A (wiring concentrator) links. The bypass relay connects the optical inputs (at the primary and secondary rings) directly to the optical output in case of a station or link failure, allowing the ring to maintain continuity.

In the 1300-nm. region, the dispersion due to multimode interference is at a minimum. The combination of physical parameters selected ensures the desired 10^{-9} bit error rate. LEDs (either surface emitting or edge emitting) are implicitly assumed in PMD; however, PMD does not specify the emitter must be an LED. It could also be a laser, as long as the optical interface parameters at the optical port are met. At some future point, some manufacturers may include lower cost local-loop-grade laser emitters, or even long-haul-grade LDs in an FDDI package, by adjusting the optical output at the optical port to conform with the standard. For the foreseeable future, however, all manufacturers pursuing FDDI products will use LEDs.

At least one vendor, Codenoll, has elected to use an 850-nm. LED. While this wavelength does not meet the standard, Codenoll claims that the only effect of the change is to shorten the maximum allowable distance between stations from 2,000 meters to 1,000 meters. Because Codenoll's products are microcomputer adapter boards, the distance limitation is less significant than it is in internetworking or mainframe attachment environments. Such nonstandard implementations, however, cannot communicate with devices using standard 1300-nm. components, requiring users to install a nonstandard device at both ends of the link.


An FDDI subcommittee is now developing a twisted pair PMD specification describing a 100-meter STP or data grade UTP connection. The committee was also investigating a standard for standard voice-grade UTP, but has deferred that effort for now.

Physical Sublayer Specification

PHY represents the upper sublayer within OSI Layer 1. It defines the encoding scheme used to represent data and control symbols. It also describes the method for retiming transmission within the node. The standard includes the following specifications:

1. Services
 - PHY to MAC Services
 - PHY to PMD Services
 - PHY to SMT Services
2. Facilities
 - Coding
 - Symbol Set
 - Line States

Figure 6.
Media Interface Connector



Example of Media Interface Connector (MIC) plug.


Digital data must be encoded in some form for proper transmission. The type of encoding depends on the nature of the transmission medium; the data rate; and other factors such as noise, reliability, and cost. Given the fact that fiber is inherently an analog medium, a digital-to-analog conversion technique is required. Intensity modulation is the norm for fiber: a binary 1 is represented by a pulse of light and a binary 0 by absence of optical power. The disadvantage of using this method, in simplest form, is its lack of synchronization. Long strings of ones or zeroes create a situation where the receiver is incapable of synchronizing its clock to that of the transmitter. The solution is to first encode the binary data in such a way as to guarantee the presence of signal transitions, even if there are no transitions in the incoming digital signal; after this encoding is performed, the signal can be presented to the optical source for transmission using intensity modulation. A typical encoding scheme is Manchester encoding (see Figure 7).

Differential Manchester is only 50% efficient since each data bit is represented by transitions in signal. Two transitions allow a degree of robustness in the presence of noise, as would be the case in coaxial cable. Since fiber is less susceptible to noise, two transitions are not required to identify a bit with a good degree of confidence.

To avoid having to use a 200MHz signal, FDDI specifies a code referred to as 4B/5B group encoding. The result is that the 100M bps throughput is achieved in FDDI with a 125MHz rate, rather than the 200MHz rate needed in differential Manchester. This helps keep down the cost and complexity of equipment.

One drawback of the group encoding pertains to clock recovery. Since differential Manchester has more pulses in its stream, it

Figure 7.
Comparison Coding Schemes



In Manchester code, there is a transition at the middle of each bit period. The mid-bit transition serves as a clock and also as data. A high-to-low transition represents a 1, and a low-to-high transition represents a 0.

In Differential Manchester code, the mid-bit transition is used only to provide clocking. The coding of a 0 (1) is represented by the presence (absence) of a transition at the beginning of the bit period.

Comparison Coding schemes used in LANs and/or fiber.

Table 1. 4B/5B Codes Used in FDDI

Function or 4-bit group (4B)	Group Code (5B)	Symbol
<i>Starting Delimiter</i>		
First symbol of sequential SD pair	11000	J
Second symbol of sequential SD pair	10001	K
<i>Data Symbols</i>		
0000	11110	0
0001	01001	1
0010	10100	2
0011	10101	3
0100	01010	4
0101	01011	5
0110	01110	6
0111	01111	7
1000	10010	8
1001	10011	9
1010	10110	A
1011	10111	B
1100	11010	C
1101	11011	D
1110	11100	E
1111	11101	F
<i>Ending Delimiter</i>		
Used to terminate datastream	01101	T
<i>Control Indicators</i>		
Logical ZERO (reset)	00111	R
Logical ONE (set)	11001	S
<i>Line Status Symbols</i>		
Quiet	00000	Q
Idle	11111	I
Halt	00100	H
<i>Invalid Code Assignment</i>		
These patterns shall not be transmitted because they violate consecutive code-bit zeroes or duty cycle requirements. Some of the codes shown shall nonetheless be interpreted as a Halt if received.	00001	Void or Halt
	00010	Void or Halt
	00011	Void
	00101	Void
	00110	Void
	01000	Void or Halt
	01100	Void
	10000	Void or Halt




Figure 8.
FDDI Packets

FDDI packet format.

is easy to extract the clock in that scheme. One of the key responsibilities of this FDDI sublayer is to decode the 4B/5B nonreturn to zero inverted (NRZI) signal from the network into symbols which can be recognized by the station, and vice versa.

The synchronization clock is derived from the incoming signal. The data is then retimed to an internal clock through an elasticity buffer. In this scheme, 4 bits of data are translated into a 5 baud value transmitted over the network, giving an 80% efficiency factor. This group-encoding scheme employed in FDDI is a departure from differential Manchester codes normally specified in LAN standards.

To understand how the FDDI scheme achieves synchronization, one must realize that there are two stages of encoding. In 4B/5B, the encoding is performed four bits at a time. Each four bits of data are encoded into a symbol with five cells such that each cell contains a single signal element (presence or absence of light). In effect, each set of four bits is encoded as five bits. Then, each element of the 4B/5B stream is treated as a binary value and encoded using NRZI. In this code, a binary 1 is represented with a transition at the beginning of the bit interval; there are no other transitions.

The advantage of NRZI is that it employs differential encoding: the signal is decoded by comparing the polarity of adjacent signal elements rather than the absolute value of a signal element. This scheme is relatively robust in detecting transitions in the presence of noise or other distortions; therefore, the NRZI encoding will improve reception reliability. Table 1 shows the symbol encoding used in FDDI.

Since this scheme is encoding 4 bits (16 combinations) with 5 bit patterns (32 combinations), there will be patterns that are not needed. The codes selected to represent the sixteen 4-bit patterns are such that a transition is present at least twice for each 5-bit code. Given an NRZI format, no more than three zeroes in a row can be allowed, since the absence of a transition indicates a zero. The remaining symbols are either declared invalid or are assigned special meaning as control symbols as shown in Table 1.

PHY also provides line states for establishing the station's links with its neighbors (upstream and downstream) and to detect the integrity of the station's links to these neighbors. These line states are used to exchange a handshake with a neighbor. A node receiving a line state on its primary input can respond by sending the proper line state on the secondary output. The line states are composed of a repetition of one or more "I" symbols.

Another item that must be resolved in this sublayer is the issue of timing jitter. Jitter is the deviation of clock recovery that can occur when the receiver attempts to recover clocking as well as data from the received signal. The clock recovery will deviate in a

random fashion from the transitions of the received signal. If no countermeasures are used, the jitter will accumulate around the ring.

In LANs, the IEEE 802 standard specifies that only one master clock will be used on the ring, and that the station with the clock will be responsible for eliminating jitter using an elastic buffer. If the ring as a whole runs ahead or behind the master clock, the elastic buffer expands or contracts accordingly.

This centralized clocking method, however, is not practical for a 100M bps ring. At this speed the bit time is only 10 ns. compared to a bit time of 250 ns. at 4M bps, making the effect of distortion more severe. Consequently, FDDI specifies the distributed clocking scheme.

In this environment, each station uses its own autonomous clock to transmit or repeat information onto the ring. Each station has an elastic buffer where data is clocked in at the clock rate recovered from the incoming stream, but is clocked out at the station's own clock rate. This distributed system is considered stronger than the centralized method and minimizes jitter. As a consequence of relocking at each station, jitter does not limit the number of repeaters in the ring, as is the case in LANs where a master clock is used.

Media Access Control Specification

Layer 2 (Data Link Layer) of the OSI Reference Model is traditionally divided into two sublayers in a LAN context: Link Layer Control (LLC) and Media Access Control (MAC). FDDI only defines MAC, which controls data flow over the ring. The token-passing protocol incorporated into FDDI controls transmission over the network. MAC defines packet formation (headers, trailers, etc.), addressing, and cyclic redundancy checking (CRC). It also defines the recovery mechanisms. This standard defines the following specifications:

- Services:
 - MAC to LLC Services
 - PHY to MAC Services
 - MAC to SMT Services
- Facilities:
 - Symbol set
 - Protocol Data Units
 - Fields
 - Timers
 - Frame Counts
 - Frame Check Sum

Glossary of FDDI Terms

Asynchronous—A class of data transmission service whereby all requests for service contend for a pool of dynamically allocated ring bandwidth and response time.

Attachment—A Port or pair of Ports, optionally including an associated optical bypass, that are managed as a functional unit. A dual attachment includes two ports: a Port A and a Port B. A single attachment includes a Port S.

Bypass—The capability of a node to optically isolate itself from the FDDI network while maintaining the continuity of the cable plant.

Capture—The act of removing a token from the ring for the purpose of frame transmission.

Claim Token—A process whereby one or more stations bid for the right to initialize the ring.

Code Bit—The smallest signaling element used by the Physical Layer for transmission on the medium.

Code Group—The specific sequence of five code bits representing a DDL symbol.

Concentrator—An FDDI node that has additional ports beyond those required for its own attachment to an FDDI network. These additional ports (type M) are for attaching other FDDI nodes (including other concentrators) in a tree topology.

Connection Management (CMT)—That portion of the Station Management (SMT) function that controls network insertion, removal, and connection of PHY and MAC entities within a station.

Counterrotating—An arrangement whereby two signal paths, one in each direction, exist in a ring topology.

Dual Attachment Concentrator—A concentrator that offers a dual attachment to the FDDI network and is capable of accommodating a dual (counterrotating) ring.

Dual Ring (FDDI dual ring)—A pair of counterrotating logical rings.

Entity—An active service or management element within an Open Systems Interconnection (OSI) layer or sub-layer.

Fiber Optic Cable—A cable containing one or more optical fibers.

Frame—A PDU transmitted between cooperating MAC entities on a ring, consisting of a variable number of octets and control symbols.

Jitter, Random—The probabilistic offsets of pulse transition edges from the expected time. Includes both Duty Cycle Distortion and Data Dependent Jitter.

Jitter, Systematic—The deterministic offsets of pulse transition edges from the expected time. Some sources of systematic jitter are differences in rise and fall times and propagation delays.

Logical Ring—The set of MACs serially connected to form a single ring. A fault-free FDDI network provides two logical rings.

Media Access Control (MAC)—The Data Link Layer responsible for scheduling and routing data transmissions on a shared-medium local area network (e.g., an FDDI ring).

Media Interface Connector (MIC)—A mated connector pair that provides an attachment between an FDDI node and a cable plant. The MIC consists of two parts: an MIC plug and an MIC receptacle.

MIC Plug—The male part of the MIC which terminates a fiber optic cable.

MIC Receptacle—The female part of the MIC which is contained in an FDDI node.

Network (FDDI Network)—A collection of FDDI nodes interconnected to form a trunk, tree, or a trunk with multiple trees. This topology is sometimes called a dual ring of trees.

Node—A generic term applying to an active element in an FDDI network (station or concentrator).

NRZ—Non Return to Zero, a technique where a polarity level (+ or -) represents a logical "1" (one) or "0" (zero).

NRZI—Non Return to Zero Invert on Ones, a technique where a polarity transition represents a logical "1" (one). The absence of a polarity transition denotes a logical "0" (zero).

Parameter Management Frames (PMF)—PMF Frames provide remote access to the SMT MIB.

The FDDI packet format is shown in Figure 8. Packets are preceded by a minimum of 16 IDLE control symbols. The packet itself is characterized by a Start Delimiter composed of the J and K control symbols. This is followed by a Frame Control field that identifies the type of packet. The Destination Address, which follows, identifies the frame recipient. The Source Address is also included to identify which station originated the packet. The address field can be 26 or 48 bits in length. The variable information field follows, along with a Frame Check Sequence field of 32 bits. The check sequence covers the Frame Control Field, the two addresses, and the information field. An End Delimiter, which consists of the T symbol, is transmitted. The maximum packet length is limited by the size of the elastic buffer in the Physical Sublayer and by the worst case frequency difference between two nodes, the upper bound in 9,000 octets. Figure 8 also shows the format of the token.

Flow control is the other major function of the MAC. In an idle condition, MAC connects to an internal source of IDLE control symbols to be transmitted over the ring. When a Start Delimiter is detected from the ring, MAC switches to a repeat path; the packet is monitored and copied if it is meant for this destination. The packet is simultaneously repeated onto the ring for relaying.

The MAC can also inject its own packet or issue a token. Packets are removed only by the originating station. The MAC repeats the packet only until the Sender Address field is detected. If the destination recognizes the Sender Address field as its own station, it will insert IDLE control symbols back onto the ring (the fragmented packet is ignored and removed by any station holding a token for transmission). Stations wishing to transmit must first obtain a token.

The procedures for obtaining the token and the amount of time allowed for data transmission (to retain fairness) are specified in the Timed Token Protocol (TTP). A station obtains the token by performing the stripping function on the incoming token. Only the Start Delimiter field is repeated onto the ring; the station will inject its own information at this juncture. When the packet is sent, the station immediately issues a new token. TTP guarantees a maximum token rotation time.

TTP allows two types of transmission: synchronous and asynchronous. In the synchronous mode, stations obtain a predefined amount of transmission bandwidth on each token rotation. The balance of the bandwidth is shared among stations using the asynchronous service. These stations can send data when the token arrives earlier than expected. Any unused capacity left over from

Physical Connection—The full-duplex Physical Layer association between adjacent PHY entities (in concentrators or stations) in an FDDI network, i.e., a pair of Physical Links.

Physical Layer (PHY)—The Physical Layer responsible for delivering a symbol stream produced by an upstream MAC Transmitter to the logically adjacent downstream MAC Receiver in an FDDI ring.

Physical Link—The simplex path (via PMD and attached medium) from the transmit function of one PHY entity to the receive function of an adjacent PHY entity (in concentrators or stations) in an FDDI network.

Physical Media Dependent (PMD)—PMD defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connector, the functions of the optical receiver, and (as an option) the bypass switch that can be incorporated into the station.

Port—A PHY entity and a PMD entity in a node, together creating a PHY/PMD pair, that may connect to the fiber media and provide one end of a physical connection with another node.

Primitive—An element of the services provided by one entity to another.

Protocol Data Unit (PDU)—Information delivered as a unit between peer entities which may contain control information, address information, and data (e.g., a Service Data Unit from a higher layer).

Receive—The action in a station in accepting a token, frame, or other signal sequence from the incoming medium.

Receiver (optical)—An electro-optical circuit that converts an optical signal to an electrical logic signal.

Repeat—The action of a station in receiving a token or frame from the adjacent upstream station and simultaneously sending to the adjacent downstream station. The FDDI MAC may repeat received PDUs (tokens and frames) but does not repeat the received signal stream between PDUs. While repeating a frame, MAC may copy the data contents and modify the control indicators as appropriate.

Repeater—A Physical Layer relay in an FDDI network.

Ring—A set of stations wherein information is passed sequentially between stations,

each station, in turn, examining or copying the information, finally returning it to the originating station. In FDDI usage, the term "ring" or "FDDI ring" refers to a dual (counterrotating) ring.

Service Data Unit (SDU)—The unit of data transfer between a service user and a service provider. The MAC SDU is the data contents of a frame. The PHY SDU is a symbol.

Services—The services provided by one entity to another. Data services are provided by a higher level entity; management services are provided to a management entity in the same or another level.

Single Attachment Concentrator—A concentrator that offers a single attachment to the FDDI network.

Single Attachment Station—A station that offers a single attachment to the FDDI network.

Station—An addressable logical and physical node on a ring capable of transmitting, repeating, and receiving information. A station has exactly one SMT, at least one MAC, at least one PHY, and at least one PMD entity.

Station Management (SMT)—The supervisory entity within an FDDI station that

monitors and controls the various FDDI entities including PMD, MAC, and PHY.

Symbol—The smallest signaling element used by MAC, i.e., the PHY SDU. The symbol set consists of 16 data symbols and 8 control symbols. Each maps to a specific sequence of five code bits as transmitted by the Physical Layer.

Synchronous—A class of data transmission service whereby each requestor is preallocated a maximum bandwidth and guaranteed a response time not to exceed a specific delay.

Token—An explicit indication of the right to transmit on a shared medium. On a tokenring, the Token circulates sequentially through the stations on the ring. At any time it may be held by zero or one stations. FDDI uses two classes of Tokens: restricted and unrestricted.

Transmit—The action of a station in generating a token, frame, or other symbol sequence and placing it on the outgoing medium.

Transmitter (optical)—An opto-electrical circuit that converts an electrical logic signal into an optical signal.

synchronous capacity is available to asynchronous traffic, which may be subdivided into up to eight levels of priority.

The amount of time allowed for asynchronous transmission is bounded by the difference of the token's actual arrival time and the expected arrival time. In essence, each station keeps track of how long it has been since it last saw the token. When it next sees the token, it can send synchronous traffic and/or any asynchronous traffic for which time remains available.

Station Management (SMT)

The FDDI Station Management (SMT) specification describes software-based, low-level data link management and integrated network control functions of all stations attached to an FDDI LAN and of the LAN itself. Each FDDI station contains only one SMT entity. SMT initializes the network, monitors error rates and fault conditions in each network segment, and automatically reconfigures the network to isolate problem links. SMT components are Connection Management (CMT), which includes Entity Coordination Management, Physical Connection Management (PCM), and Ring Management (RMT). SMT is intended to operate regardless of equipment type, vendor, protocols, or applications. Figure 9 presents the SMT architectural model.

SMT types (managed objects) have specific attributes indicating state, capabilities, and operation.

SMT managed objects are the following:

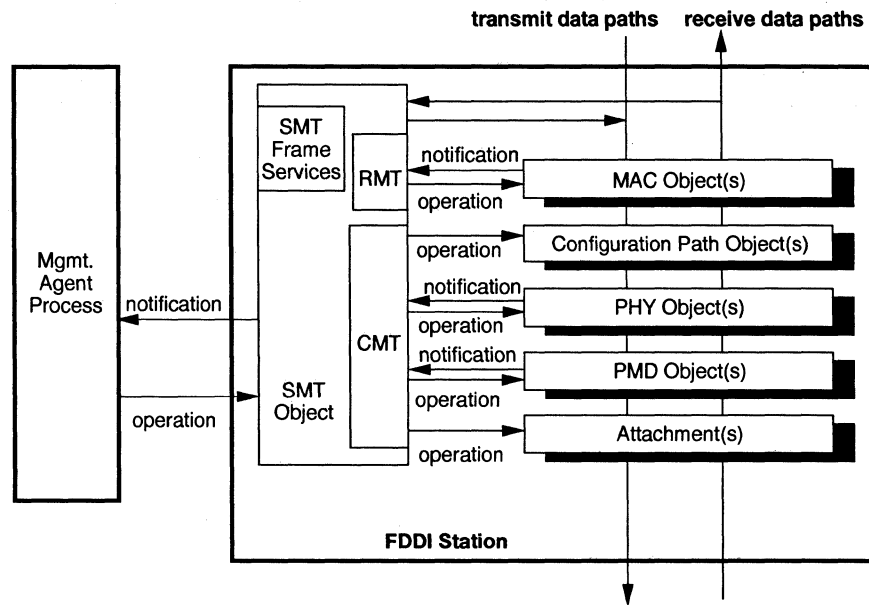
- Station or concentrator (SMT)
- MAC object(s)
- Path object(s)
- PHY object(s)
- PMD object(s)
- Attachment(s)

Attributes are the following:

- Attribute Identification (ID)
- Configuration
- Operational (Status, Counters, etc.)

Figure 9.
SMT Architectural Model

The relationships between managed objects in an FDDI station and the SMT object.



Each attribute is defined in terms of Access Rights, and whether it is Mandatory or Optional. Each attribute also carries an FDDI specification reference or a specific definition if it is not defined in the FDDI specification.

SMT Connection Management (CMT) operates at the logical level, controlling the interface between PHY and MAC entities in a given station and controlling SMT-to-SMT communications across the ring. When a session requires a connection to another station, CMT causes PHY to send a stream of symbols to the targeted station. Upon receiving the symbols, the receiving station's PHY returns a continuous stream of symbols (primitives) indicating line state, station status, and willingness to carry out the requested action (establish a link). QUIET symbols indicate a disabled link. HALT or alternating HALT and QUIET (MASTER) symbols indicate an operating link and the receiving station's status (master, slave, or peer). A stream of IDLE symbols indicates willingness to connect. Once the link is established, CMT configures the PHY and MAC.

Entity Coordination Management (ECM) controls the optional optical bypass switch, and signals the Physical Connection Management (PCM) entity when the bypass is complete. ECM also performs the Path Test to determine a fault's location.

Physical Connection Management (PCM) initializes the adjacent station's PHYs, and manages signaling. Maintenance support functions are also part of PCM.

Configuration Management (CFM) interconnects PHYs and MACs. It automatically configures these connections according to PCM flags. CFM is defined differently for stations and concentrators.

Ring Management (RMT) relays MAC and CFM status information. It detects stuck signaling beacons, initiates the trace function, detects duplicate addresses and resolves them to allow continued ring operation, and notifies SMT of MAC status.

Parameter Management Frames (PMF) are a newly required feature of FDDI SMT, implemented in Version 7.2. PMF frames provide remote FDDI stations with access to the SMT MIB. SMT Version 7.3, describing hybrid ring control (HRC) was presented to the X3 parent committee in late 1993 and should be released in early 1994. SMT Version 7.2 interoperates with Version 6.2, but earlier versions of SMT do not support HRC.

In August 1992, Digital Equipment Corp., AT&T/NCR, Motorola, and Distributed Systems International announced a successful SMT interoperability test. The tests demonstrated interoperability between SMT Versions 6.2 and 7.2, as well as between FDDI products from different vendors. Part of the tests involved the operation of the PMF.

Higher Layers

From LLC upwards, ANSI intends FDDI to generally fit traditional protocol stacks. The ANSI FDDI committee has not yet formally drafted protocols for Layer 3 (network) and 4 (transport), which are needed for any type of internetworking. Many suppliers believe the TCP/IP protocol suite can serve this purpose. TCP/IP software was originally developed by the U.S. government for Arpanet, the worldwide packet switched network, but has been used with great success in commercial applications—particularly for internetworking among different LANs. TCP and IP follow layered networking concepts, occupying Layers 4 and 3, respectively, of the OSI model. Vendors have adopted the protocols for Ethernet and other local area networks. Most commercial implementations of the TCP/IP protocol suite include three standardized upper layer protocols: Telnet (virtual network terminal), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP).

Although more versatile than TCP/IP, the OSI internetworking protocols—spanning Layers 3 through 5—are less practical to implement in the real world. Much effort, however, has been expended to migrate to ISO-based standards. These market demands are forcing investigations into upgrading the third and fourth layers. The OSI stack also has a well-developed suite of upper layer applications, including X.400 (electronic mail), File Transfer Access Management (FTAM), and X.500 Directory System Protocol.

Certification

FDDI vendors conduct interoperability evaluations at Advanced Micro Device's Advanced Networking Center (ANC) in Sunnyvale, California, and the University of New Hampshire's Interoperability Lab. Although the actual results of these interoperability tests are restricted to vendors, network managers can insist on seeing the Seal of Approvals, product interoperation matrixes, and testing procedures before purchasing specific products.



Sponsoring Organizations

Advanced Networking Test Center (ANTC)

Advanced Micro Devices
P.O. Box 3453, Mail Stop 17
Sunnyvale, CA 94088
Tel: +1 408 749 4235
Fax: +1 408 749 5008

ANSI

1430 Broadway
New York, NY 10018
Tel: +1 212 642 4900

FDDI SMT Development Forum

Synermetics Inc.
800 W. El Camino Real
Mountain View, CA 94040
Tel: +1 415 903 2225
Fax: +1 508 366 8901

Global Engineering Documents

2805 McGaw Avenue, P.O. Box 19539
Irvine, CA 92714
Tel: +1 714 261 1455

Interop, Inc.

480 San Antonio Road, Suite 100
Mountain View, CA 94040
Tel: +1 415 941 3399
Fax: +1 415 949 1779

University of New Hampshire Interoperability Laboratory (IOL)

Science and Engineering
Research Building, Room 336
Durham, NH 03824
Tel: +1 603 862 4532
Fax: +1 603 862 1915 ■

ANSI Fiber Distributed Data Interface (FDDI) Standards

In this report:

Technology Overview	2
Applications.....	2
FDDI Specifications.....	5

Background

In October 1982, the American National Standards Institute (ANSI), Committee X3T9.5, was chartered to develop a high-speed data networking standard. Known as Fiber Distributed Data Interface (FDDI), the standard specifies a packet switched LAN-to-LAN backbone that transports data at high throughput rates over a variety of multimode fibers. FDDI is a token-passing network employing two fiber pairs operating at 100M bps. With the release of Version 7.2 of the Station Management (SMT) portion in June 1992, the core FDDI standard is now complete. Work continues in ANSI X3T9.5 to develop an alternative physical media dependent (PMD) standard for copper and for lower-cost fiber.

Since the initial standard was proposed, ANSI has begun work on enhancements to FDDI. FDDI II, which is considered a superset of FDDI, will add the capability of carrying voice and video signals over FDDI-based networks. FDDI follow-on, or FFOL, will allow interconnection of multiple FDDI networks, connection to wide-band facilities, and SONET compatibility. FDDI II should be complete by the end of 1992; FFOL, by 1995.

This report details the FDDI draft standard, which covers the first two layers of the OSI Reference Model, and discusses future directions for the standard.

—By *Timothy McElgunn*
Associate Analyst

Analysis

Technology Overview

FDDI addresses the bottom two layers of the OSI model. It is the spirit of the FDDI standard to use ISO protocols for the other layers, where possible. The optical-based FDDI LAN was designed to provide the same type of serial interconnection provided by LANs while providing the high bandwidth, inherent noise immunity, and security offered by fiber.

At FDDI's inception in 1982, fiber was used mostly for point-to-point applications, and not for the many configurations allowed by LANs. In this sense, FDDI was a breakthrough. Although it is possible to achieve considerably higher data rates over fiber (up to 3.7G bps with current point-to-point technology, and 500M bps on rings), higher rates result in significantly increased costs and shorter transmission distances between repeaters. Its designers intend FDDI to provide relatively inexpensive connectivity and, therefore, focused on the 100M bps rate.

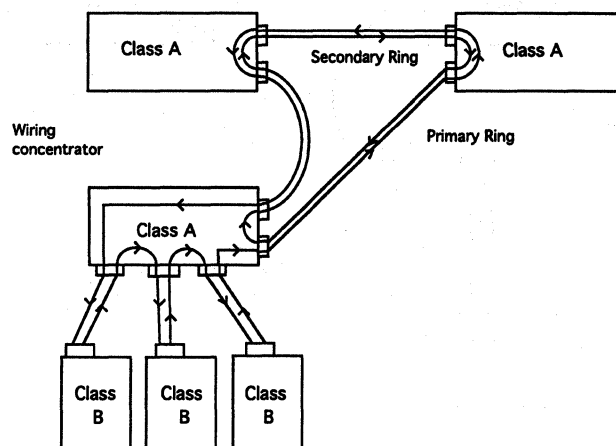
FDDI can be configured to support a sustained transfer rate of approximately 80M bps. The remaining bandwidth is reserved for various overhead functions.

FDDI is a token-passing, dual-ring network accommodating synchronous and asynchronous data transmission, as well as isochronous channels for realtime digitized voice and compressed video. Unlike existing open standards for LANs, where fiber optic variants have followed copper implementations, FDDI has been designed from the start as a fiber optic network. This has involved standardization issues in such areas as duplex optical connectors, fiber characteristics, optical bandwidth, bypass relays, and cable assemblies. The FDDI ring is designed for an overall bit error rate of less than 10^{-9} . The network can tolerate up to 2 km. of fiber between stations and can support a total cable distance of 100 km. around the ring with 500 attachments (1,000 physical connections and a total fiber path of 200 kilometers). FDDI topology is a counterrotating token-passing ring (note the arrows in Figure 1). FDDI, however, is not part of the well-established IEEE 802 family of LAN standards.

FDDI operates at 1300 nanometers (nm.). Current transmitter/receiver fiber technology operates at 850 nm., 1300 nm., or 1550 nm. While performance increases with wavelength, so does cost. For local data communications, both in LANs and in point-to-point applications employing fiber optic modems, the 850-nm. light sources are typically employed; however, this technology becomes unfeasible for 100M bps beyond a couple of miles. At the other end of the range (1550 nm.), the system becomes expensive and may provide unnecessary bandwidth.

The committee designing FDDI also investigated short-wavelength implementations. It became evident that to meet all the requirements, particularly the two-kilometer, station-to-station spacing, the system would require 1300-nm. wavelength. Using 1300-nm. technology, less expensive light-emitting diodes (LEDs) provide distance and data rates within the range desired for LANs and LAN backbones. The issue of fiber size was settled sometime after the wavelength decision.

Figure 1.
The FDDI Environment

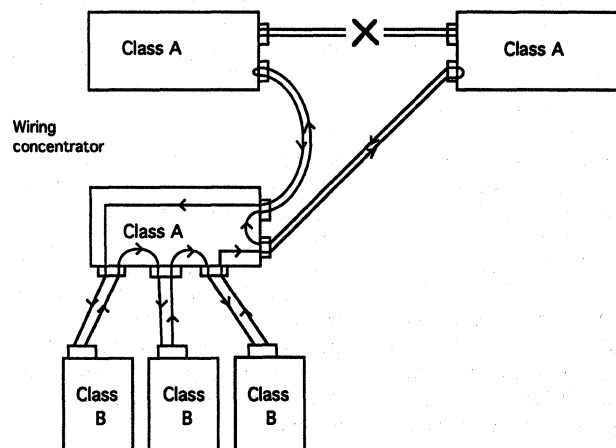


The FDDI standard directly addresses the need for reliability. A backbone system transports a large number of user sessions and its loss would be serious. FDDI incorporates three reliable-enhancing features. First, a failed or unpowered station can be bypassed by an optional automatic optical bypass switch; second, wiring concentrators are used in a star-wiring strategy to facilitate fault isolation and correction; third, two rings interconnect stations so that failure of a repeater or cable link results in the automatic reconfiguration of the network (loopback) (see Figures 2 and 3).

Applications


Although optical fiber is widely deployed in the telecommunications environment (long hauls, interoffice, feeder plant, and so on), it has not done as well in the LAN arena. Three reasons can be identified: 1) increased technical complexity compared to passive copper and coaxial cable; 2) cost considerations; and 3) lack of a workable standard.

Figure 2.
Failure Recovery



Ring rearrangement under failure.

Figure 3.
Class B Failure



Failure of Class B stations.

FDDI solves the third problem and in the process also begins to resolve the second.

FDDI allows designers to 1) build larger capacity LANs or LAN backbones to serve new data needs (file transfer, graphics, and so on), and some voice needs; or 2) interconnect LANs in metropolitan area networks (MANs). Thus, FDDI can be used directly as a LAN or as a backbone to interconnect slower-speed LANs into a single network over relatively large geographies.

The initial application of FDDI as a "back-end" interconnect for high-powered computing devices and peripherals required a high degree of fault tolerance and data integrity. As developers proceeded, it became obvious that FDDI could also serve high-speed "front-end" applications. Front-end applications include terminal-to-terminal and terminal-to-server communications typical of a LAN. In large networks (3,000 to 10,000 terminals, particularly when workstations are involved), the aggregate demand for network resources can overwhelm a 4M bps or 10M bps LAN. At that point, FDDI's bandwidth becomes important.

A high-speed FDDI ring is ideal as a backbone for other "departmental" LANs, typically operating at lower speeds. High capacity on LANs can be achieved in two ways by using multiple channels at low speeds, or by using one channel at a relatively high speed. The multiple-channel approach can be used with a broadband bus LAN. A drawback of this approach is that bridges must be provided between channels, and the architecture must be designed to avoid high rates of interchannel traffic and bottlenecks at the bridges.

Desktop systems that are likely to be linked via high-speed networks are now applied in realtime simulations, graphic transfers in a CAD/CAM environment, supercomputer terminals, medical imaging, and video.

The one area where FDDI is not well suited is broadband LANs carrying full analog video (6MHz). A broadband LAN can easily carry multiple channels of video as well as data. Digitizing a 6MHz TV channel results in a 45M bps to 90M bps data rate (45M bps is a slightly compressed version), which can easily swamp an FDDI backbone. FDDI II, discussed later in this report, has been designed to handle video efficiently.

FDDI is not the end, however. Even larger bandwidths are envisioned in the next five years under the thrust of Broadband ISDN (B-ISDN). BISDN for video needs, particularly in a High Definition TV (HDTV) environment, will require approximately 150M bps of dedicated bandwidth per channel. Planners are now discussing delivery of up to three channels per domicile, reaching into the 600M bps range. FDDI is a shared-medium technology; uncompressed video applications and/or local loop applications may not be capable of using networks built on the original FDDI standards.

Relationship to Other Standards

To fully understand the FDDI standard, some knowledge of the OSI Reference Model and how it applies to the LAN environment is necessary. A brief discussion of the OSI model is followed by a more detailed description of FDDI specifications and emerging vendor products. Readers intimately familiar with this material can skip to the section entitled "FDDI Specifications."

The OSI Reference Model imposes order and structure on data communication. To achieve orderly communication, many functions must be performed. It is natural to group these functions into layers which share task affinity and logical proximity (see Figure 4). OSI layers are hierarchical in the sense that a given layer calls for the services of the layer immediately beneath it. The given layer cannot ask for the services of a layer at a higher level, nor can it skip the layer beneath it to directly reach a lower layer, or even jump into the middle of another layer. The model includes precise definitions of the services provided by each layer to its next higher layer and request procedures.

Services defined for a given layer are, in turn, employed by the layer immediately above it. Each transmitting layer passes down to the lower layer (and up to the next layer, at the receiving end) blocks of data requiring processing for transmission, manipulation, or service. These layers normally attach a characteristic header that contains appropriate information (such as the real network address, block number, and so on). The headers are physically nested, with lower layer headers being outermost and higher layer headers being innermost. It is through the use of these well-defined headers that the protocols between the remote open systems are accomplished.

To effectuate the OSI model, the International Organization for Standardization has formulated standards—specifications for how information is coded and passed between communicating partners. Only the protocols need to be implemented by a prospective vendor. Users should note that, when working at a terminal connected to a host or a PC connected to a LAN, all seven layers of the architecture must be employed.

The Reference Model and the service definitions are only structures for discussing the tasks involved in communicating between open systems. The OSI Reference Model is described by document ISO 7498, which was adopted as a standard in 1984.

LAN Protocol Suites

In a LAN environment, one typically defines Layer 1 and Layer 2 standards specific to local area networks. Layers 1 and 2 are defined by the IEEE 802 standards. With some "internetworking" protocols defined at Layer 3 (such as ISO 8473), and typically some connection-oriented Transport protocols (for instance, FTAM and MHS), sessions

Glossary of FDDI Terms

Asynchronous—A class of data transmission service whereby all requests for service contend for a pool of dynamically allocated ring bandwidth and response time.

Attachment—A Port or pair of Ports, optionally including an associated optical bypass, that are managed as a functional unit. A dual attachment includes two ports: a Port A and a Port B. A single attachment includes a Port S.

Bypass—The capability of a node to optically isolate itself from the FDDI network while maintaining the continuity of the cable plant.

Capture—The act of removing a token from the ring for the purpose of frame transmission.

Claim Token—A process whereby one or more stations bid for the right to initialize the ring.

Code Bit—The smallest signaling element used by the Physical Layer for transmission on the medium.

Code Group—The specific sequence of five code bits representing a DDL symbol.

Concentrator—An FDDI node that has additional ports beyond those required for its own attachment to an FDDI network. These additional ports (type M) are for attaching other FDDI nodes (including other concentrators) in a tree topology.

Connection Management (CMT)—That portion of the Station Management (SMT) function that controls network insertion, removal, and connection of PHY and MAC entities within a station.

Counterrotating—An arrangement whereby two signal paths, one in each direction, exist in a ring topology.

Dual Attachment Concentrator—A concentrator that offers a dual attachment to the FDDI network and is capable of accommodating a dual (counterrotating) ring.

Dual Ring (FDDI dual ring)—A pair of counterrotating logical rings.

Entity—An active service or management element within an Open Systems Interconnection (OSI) layer or sub-layer.

Fiber Optic Cable—A cable containing one or more optical fibers.

Frame—A PDU transmitted between cooperating MAC entities on a ring, consisting of a variable number of octets and control symbols.

Jitter, Random—The probabilistic offsets of pulse transition edges from the expected time. Includes both Duty Cycle Distortion and Data Dependent Jitter.

Jitter, Systematic—The deterministic offsets of pulse transition edges from the expected time. Some sources of systematic jitter are differences in rise and fall times and propagation delays.

Logical Ring—The set of MACs serially connected to form a single ring. A fault-free FDDI network provides two logical rings.

Media Access Control (MAC)—The Data Link Layer responsible for scheduling and routing data transmissions on a shared-medium local area network (e.g., an FDDI ring).

Media Interface Connector (MIC)—A mated connector pair that provides an attachment between an FDDI node and a cable plant. The MIC consists of two parts: an MIC plug and an MIC receptacle.

MIC Plug—The male part of the MIC which terminates a fiber optic cable.

MIC Receptacle—The female part of the MIC which is contained in an FDDI node.

Network (FDDI Network)—A collection of FDDI nodes interconnected to form a trunk, tree, or a trunk with multiple trees. This topology is sometimes called a dual ring of trees.

Node—A generic term applying to an active element in an FDDI network (station or concentrator).

NRZ—Non Return to Zero, a technique where a polarity level (+ or -) represents a logical "1" (one) or "0" (zero).

NRZI—Non Return to Zero Invert on Ones, a technique where a polarity transition represents a logical "1" (one). The absence of a polarity transition denotes a logical "0" (zero).

can then employ the normal protocol suite up to Layer 7 described above. Prior to the establishment of ISO-based standards, however, the use of Transmission Control Protocol/Internet Protocol (TCP/IP) was common for the upper layers. Below, one finds descriptions of the standards at Layer 1 and Layer 2, which are specific to LANs.

Logical Link Control

IEEE Standard 802.2-1985 (ISO 8802/2) describes the peer-to-peer protocol procedures for the transfer of information and control between any pair of Data Link Layer Service Access Points (SAPs) on a local area network.

Medium Access Control: Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

IEEE Standard 802.3-1985 (ISO 8802/3) provides a medium access method by which two or more stations share a common-bus transmission medium. The standard applies to several media types and provides the necessary specifications for a 10M bps baseband local area network.

Medium Access Control

Token-passing bus access method IEEE Standard 802.4-1985 (ISO 8802/4) deals with all elements of the token-passing bus access method and its associated physical signaling and media technologies. The goal is to achieve compatible interconnection of stations in a LAN. The access method coordinates the use of the shared medium among the attached stations. It specifies the electrical and physical characteristics of the transmission medium, the electrical signaling used, the frame formats of the transmitted data, the actions of a station upon receipt of a data frame, and the services provided at the conceptual interface between the Medium Access Control sublayer and the Logical Link Control sublayer above it.

Medium Access Control: Token-Passing Ring Access Method

IEEE Standard 802.5-1985 (ISO 8802/5) specifies the formats and protocols used by the token-passing ring medium at the control sublayer and Physical Layer. It also specifies the means of attachment to the token-passing ring access method. The protocol defines the frame format including

Parameter Management Frames (PMF)—PMF Frames provide remote access to the SMT MIB.

Physical Connection—The full-duplex Physical Layer association between adjacent PHY entities (in concentrators or stations) in an FDDI network, i.e., a pair of Physical Links.

Physical Layer (PHY)—The Physical Layer responsible for delivering a symbol stream produced by an upstream MAC Transmitter to the logically adjacent downstream MAC Receiver in an FDDI ring.

Physical Link—The simplex path (via PMD and attached medium) from the transmit function of one PHY entity to the receive function of an adjacent PHY entity (in concentrators or stations) in an FDDI network.

Physical Media Dependent (PMD)—PMD defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connector, the functions of the optical receiver, and (as an option) the bypass switch that can be incorporated into the station.

Port—A PHY entity and a PMD entity in a node, together creating a PHY/PMD

pair, that may connect to the fiber media and provide one end of a physical connection with another node.

Primitive—An element of the services provided by one entity to another.

Protocol Data Unit (PDU)—Information delivered as a unit between peer entities which may contain control information, address information, and data (e.g., a Service Data Unit from a higher layer).

Receive—The action in a station in accepting a token, frame, or other signal sequence from the incoming medium.

Receiver (optical)—An electro-optical circuit that converts an optical signal to an electrical logic signal.

Repeat—The action of a station in receiving a token or frame from the adjacent upstream station and simultaneously sending to the adjacent downstream station. The FDDI MAC may repeat received PDUs (tokens and frames) but does not repeat the received signal stream between PDUs. While repeating a frame, MAC may copy the data contents and modify the control indicators as appropriate.

Repeater—A Physical Layer relay in an FDDI network.

Ring—A set of stations wherein information is passed sequentially between stations, each station, in turn, examining or copying the information, finally returning it to the originating station. In FDDI usage, the term "ring" or "FDDI ring" refers to a dual (counterrotating) ring.

Service Data Unit (SDU)—The unit of data transfer between a service user and a service provider. The MAC SDU is the data contents of a frame. The PHY SDU is a symbol.

Services—The services provided by one entity to another. Data services are provided by a higher-level entity; management services are provided to a management entity in the same or another level.

Single Attachment Concentrator—A concentrator that offers a single attachment to the FDDI network.

Single Attachment Station—A station that offers a single attachment to the FDDI network.

Station—An addressable logical and physical node on a ring capable of transmitting, repeating, and receiving information. A station has exactly one SMT, at least one MAC, at least one PHY, and at least one PMD entity.

Station Management (SMT)—The supervisory entity within an FDDI station that monitors and controls the various FDDI entities including PMD, MAC, and PHY.

Symbol—The smallest signaling element used by MAC, i.e., the PHY SDU. The symbol set consists of 16 data symbols and 8 control symbols. Each maps to a specific sequence of five code bits as transmitted by the Physical Layer.

Synchronous—A class of data transmission service whereby each requester is preallocated a maximum bandwidth and guaranteed a response time not to exceed a specific delay.

Token—An explicit indication of the right to transmit on a shared medium. On a token-ring, the Token circulates sequentially through the stations on the ring. At any time it may be held by zero or one stations. FDDI uses two classes of Tokens restricted and unrestricted.

Transmit—The action of a station in generating a token, frame, or other symbol sequence and placing it on the outgoing medium.

Transmitter (optical)—An optoelectrical circuit that converts an electrical logic signal into an optical signal.

delimiters, addressing, and frame check sequence and includes timers, frame counts, and priority stacks. It also defines the medium access control protocol and provides finite-state machines and state tables supplemented with descriptions of the algorithms. It identifies the services provided by the Medium Access Control sublayer to the Logical Link Control sublayer, and the services provided by the Physical Layer to the Medium Access Control sublayer. These services are defined in terms of service primitives and associated parameters. Also defined are the Physical Layer functions of symbol encoding and decoding, symbol timing and latency buffering, and the 1M bps and 4M bps shielded twisted-pair attachments of the station to the medium.

Figure 5, modeled after documentation of the German Commission for Computer-Integrated Manufacturing, depicts typical OSI protocol suites for LANs, particularly in a MAP/TOP environment. Figure 6 depicts some additional details at the lower layers, including the positioning of FDDI juxtaposed to other LAN standards. Once again, note that when a user works at a PC connected to a LAN or FDDI ring, all seven layers of the protocol architecture

must be employed, as shown in Figure 5. Thus, while the FDDI standard concentrates on Layers 1 and 2, the upper layers are required to accomplish end-to-end communication.

Specifications

FDDI Specifications

FDDI is defined according to the OSI Reference Model and LAN protocol architecture. Layer 1 (Physical Layer) is specified in two documents: the FDDI Physical Medium Dependent (PMD), and the FDDI Physical Sublayer (PHY). (See Figure 7.) The Physical Layer provides the

Figure 4.
OSI Functions

<p style="text-align: center;">Application Layer (7)</p> <ul style="list-style-type: none"> - Application oriented functions (example: file transfer) - Management functions
<p style="text-align: center;">Presentation Layer (6)</p> <p>Data formats, codes and representation</p>
<p style="text-align: center;">Session Layer (5)</p> <p>Dialogue control and synchronization, initialization</p>
<p style="text-align: center;">Transport Layer (4)</p> <p>End to end control of transportation of data, optimization of usage of network resources</p>
<p style="text-align: center;">Network Layer (3)</p> <p>Forwarding blocks of data over a network route towards their final destination</p>
<p style="text-align: center;">Link Layer (2)</p> <p>Reliable transfer of blocks of data between adjacent nodes (error detection and recovery)</p>
<p style="text-align: center;">Physical Layer (1)</p> <p>Control of data circuits, transfer of bits</p>
<p style="text-align: center;">Physical Media</p> <p>Transmission of electrical signals</p>

Functions performed by the layers of the OSI Reference Model.

medium, connectors, optical bypassing, and driver/receiver requirements. It also defines encode/decode and clock requirements for framing data for transmission on the medium or to the higher layers of FDDI.

When the FDDI committee realized there would be considerable discussion on fibers, connectors, and other hardware, it decided to break the standardization of the OSI Physical Layer into two pieces. In this way, the relatively noncontroversial issues—such as coding and other matters that IC chip manufacturers need to know to begin design—could be put in a formal document and approved independently of other items pertaining to the standard.

The Data Link Layer is also divided into two sublayers:

1. A Media Access Control portion that provides fair and deterministic access to the medium, address recognition, and generation, and verification of frame check

sequences. Its primary function is the delivery of frames, including frame insertion, repetition, and removal.

2. A Logical Link Control portion that provides a common protocol for data assurance services between the Media Access Control and the Network Layer.

Physical Medium Dependent (PMD) Specification

PMD defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connector, the functions of the optical receiver, and (as an option) the bypass switch that can be incorporated into the station. It specifies the optical channel at the bulkhead of a station. The source is defined to radiate in the 1300-nm. wavelength. PMD also describes the peak optimal power, optical rise and fall times, and jitter constraints. The minimum rise/fall time is 0.5 nanosecond. The standard includes the following specifications:

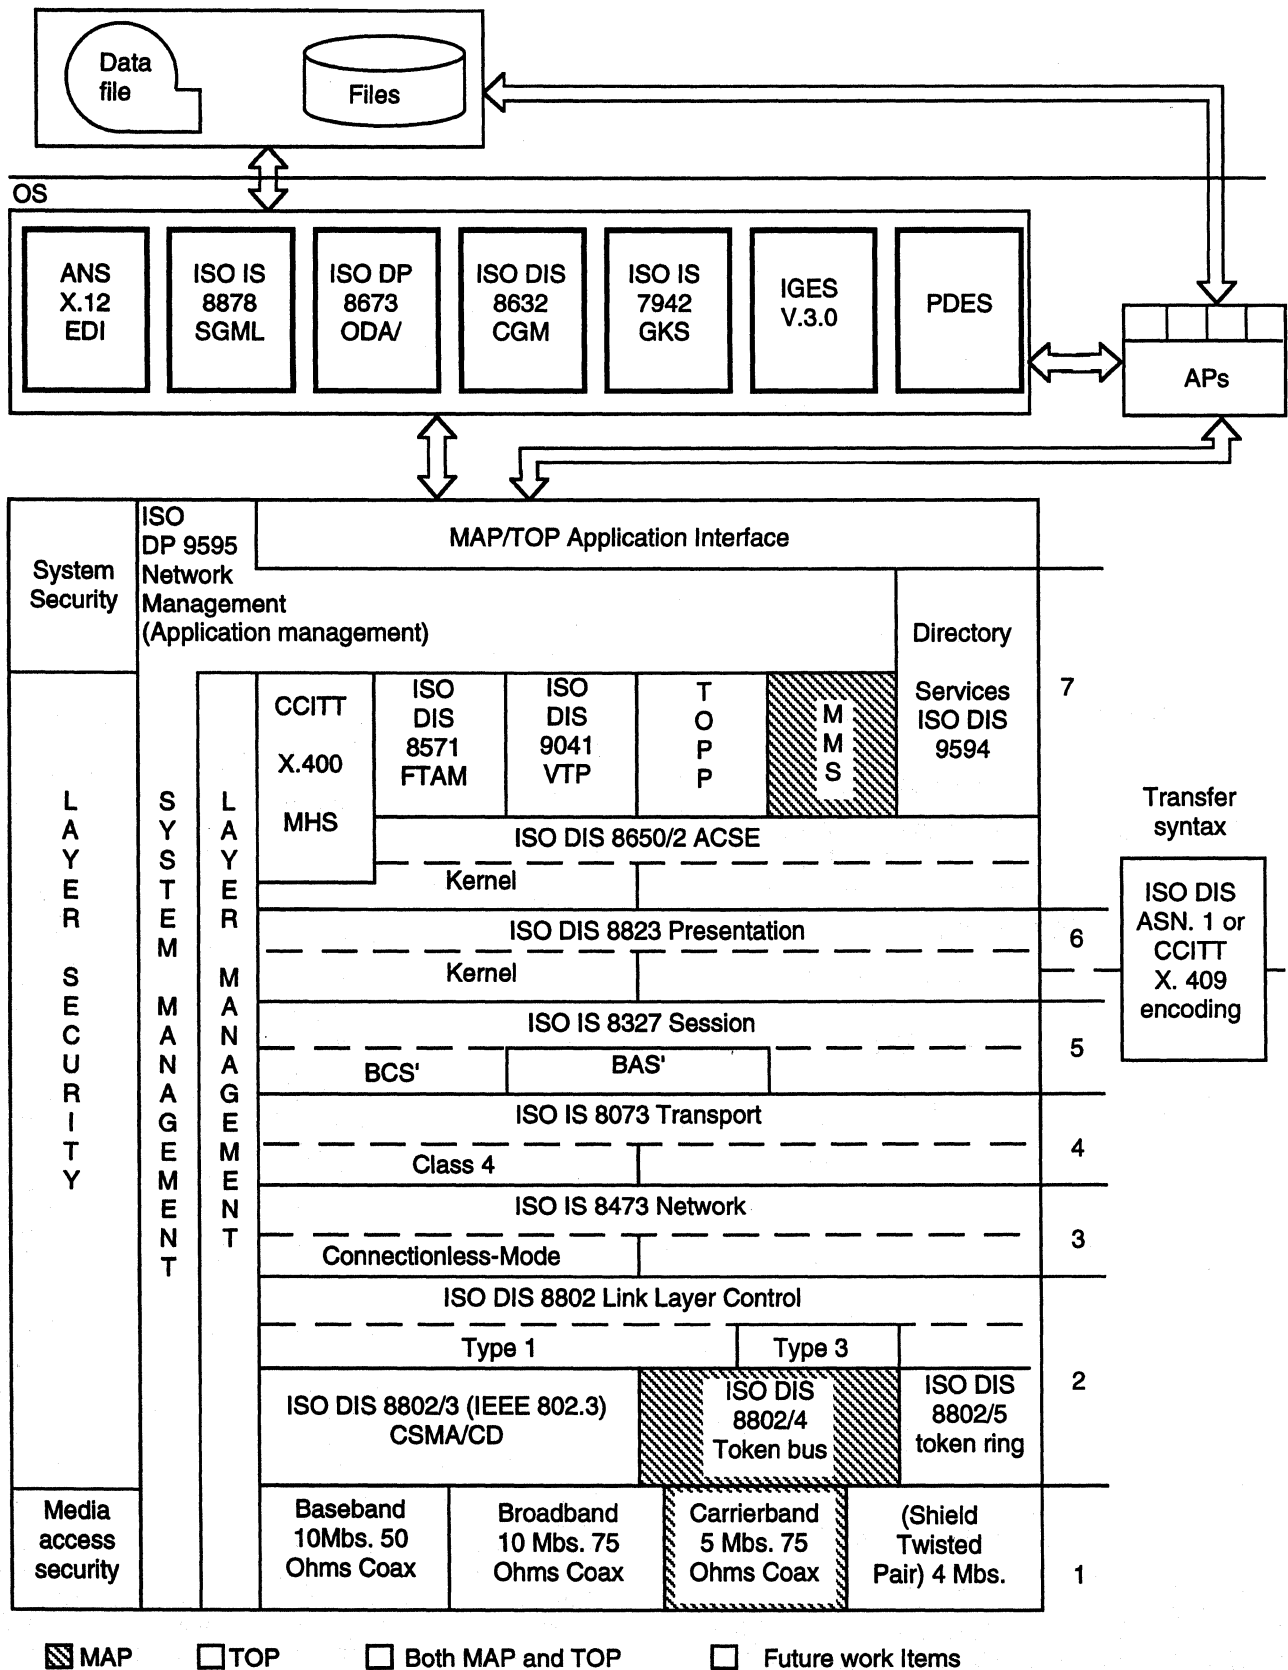
1. Services: PMD to PHY Services; PMD to SMT Services
2. Media attachment
3. Media Signal Interface
4. Interface signals
5. Cable Plant Interface Specification

Multimode fibers are employed (at least initially) up to a distance of two kilometers. Optical fiber dimensions are specified in terms of its core diameter and the outer diameter of the cladding layer. Fiber specifications are 62.5/125 micron (core diameter/cladding diameter) and 85/125 micron. The nominal numerical aperture is around 0.26. Applicable standards for the fiber itself are EIA-455-48 (core), EIA-455-27 or EIA-455-48 (cladding), and EIA-455-57 (aperture).

Listed in the appendix of the standard are two other fibers: a 50/125-micron fiber and a 100/140-micron fiber. These two fibers have not been extensively studied; the maximum achievable distance of 2 km. specified in the standard may not be possible. Thus, these two fibers are not officially part of the standard but are listed as "alternatives." Smaller diameters offer higher bandwidths but also more expensive, higher loss connectors. The 50/125 fiber has been used primarily for military applications, and is not likely to be as widely available as the other fibers. The 100/140 fiber has been added primarily because it is used in IBM's cabling system and a number of customers have already installed it. On the other hand, the 62.5/125 fiber has been in production for some time and has become common for local applications, such as AT&T's Premises Distribution System (PDS). Component costs for this type of fiber are dropping most rapidly.

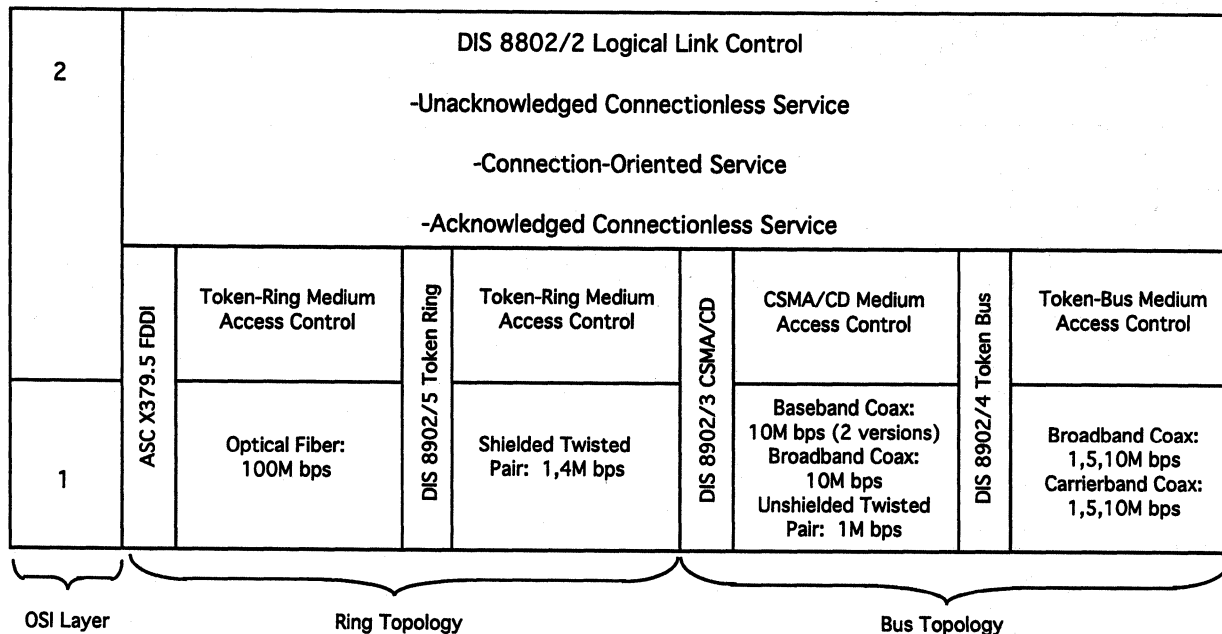
Some observers believe it is unfortunate that the FDDI standard could not specify a single fiber type, since this might have lowered costs and made it easier for the customer to start small and expand later. The FDDI committee settled on a 62.5-micrometer core, with advisory information about 50-, 85-, and 100-micrometer fiber sizes. While a debate remains about these other sizes, the issue is not really a critical one: as long as a fiber meets the optical power, channel bandwidth, and distance requirements, it

Figure 5. ISO Suites



ISO suites in a LAN environment.

Figure 6. ISO Lower Layers



Lower layers in a LAN environment.

can conform to the FDDI standard at the interface between the FDDI box and the network, independent of the fiber type (the charts and data presentations in the PMD have been written so that they can be applied to various fiber sizes).

Like all Layer 1 specifications, PMD defines the duplex connector used for FDDI access (see Figure 8). The primary and secondary ring connections to each Class A station are attached simultaneously using the duplex connector and a dual-fiber cable. The connectors can be used for both Class A-to-Class A as well as Class B-to-Class A (wiring concentrator) links. The bypass relay connects the optical inputs (at the primary and secondary rings) directly to the optical output in case of a station or link failure, allowing the ring to maintain continuity.

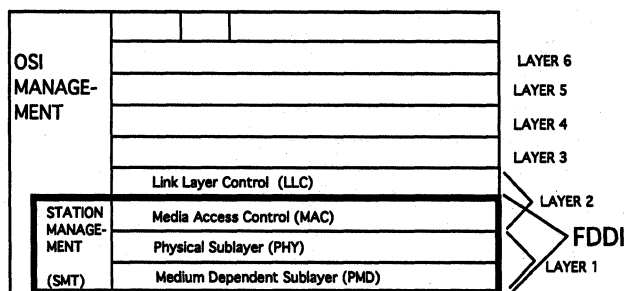
In the 1300-nm. region, the dispersion due to multi-mode interference is at a minimum. The combination of physical parameters selected ensures the desired 10⁻⁹ bit error rate. LEDs (either surface emitting or edge emitting)

are implicitly assumed in PMD; however, PMD does not specify the emitter must be an LED. It could also be a laser, as long as the optical interface parameters at the optical port are met. At some future point, some manufacturers may include lower cost local-loop-grade laser emitters, or even long-haul-grade LDs in an FDDI package, by adjusting the optical output at the optical port to conform with the standard. For the foreseeable future, however, all manufacturers pursuing FDDI products will use LEDs.

At least one vendor, Codenoll, has elected to use an 850-nm. LED. While this wavelength does not meet the standard, Codenoll claims that the only effect of the change is to shorten the maximum allowable distance between stations from 2,000 meters to 1,000 meters. Because Codenoll's products are microcomputer adapter boards, the distance limitation is less significant than it is in inter-networking or mainframe attachment environments. Such nonstandard implementations, however, cannot communicate with devices using standard 1300-nm. components, requiring users to install a nonstandard device at both ends of the link.

An FDDI subcommittee is now developing a twisted-pair PMD specification describing a 100-meter STP or data grade UTP connection. The committee was also investigating a standard for standard voice grade UTP, but has deferred that effort for now.

Figure 7. FDDI Protocols




Physical Sublayer Specification

PHY represents the upper sublayer within OSI Layer 1. It defines the encoding scheme used to represent data and control symbols. It also describes the method for retiming transmission within the node. The standard includes the following specifications:

1. Services
 - PHY to MAC Services

Figure 8.
Media Interface Connector



Example of Media Interface Connector (MIC) plug.

- PHY to PMD Services
- PHY to SMT Services

2. Facilities

- Coding
- Symbol Set
- Line States

Digital data must be encoded in some form for proper transmission. The type of encoding depends on the nature of the transmission medium; the data rate; and other factors such as noise, reliability, and cost. Given the fact that fiber is inherently an analog medium, a digital-to-analog conversion technique is required. Intensity modulation is the norm for fiber: a binary 1 is represented by a pulse of light and a binary 0 by absence of optical power. The disadvantage of using this method, in simplest form, is its lack of synchronization. Long strings of ones or zeroes create a situation where the receiver is incapable of synchronizing its clock to that of the transmitter. The solution is to first encode the binary data in such a way as to guarantee the presence of signal transitions, even if there are no transitions in the incoming digital signal; after this encoding is performed, the signal can be presented to the optical source for transmission using intensity modulation. A typical encoding scheme is Manchester encoding (see Figure 9).

Differential Manchester is only 50% efficient since each data bit is represented by transitions in signal. Two transitions allow a degree of robustness in the presence of noise, as would be the case in coaxial cable. Since fiber is less susceptible to noise, two transitions are not required to identify a bit with a good degree of confidence.

To avoid having to use a 200MHz signal, FDDI specifies a code referred to as 4B/5B group encoding. The result is that the 100M bps throughput is achieved in FDDI with a 125MHz rate, rather than the 200MHz rate needed in differential Manchester. This helps keep down the cost and complexity of equipment.

One drawback of the group encoding pertains to clock recovery. Since differential Manchester has more pulses in its stream, it is easy to extract the clock in that scheme. One of the key responsibilities of this FDDI sublayer is to decode the 4B/5B nonreturn to zero inverted (NRZI) signal from the network into symbols which can be recognized by the station, and vice versa.

The synchronization clock is derived from the incoming signal. The data is then retimed to an internal clock through an elasticity buffer. In this scheme, 4 bits of data are translated into a 5 baud value transmitted over the network, giving an 80% efficiency factor. This group-encoding scheme employed in FDDI is a departure from differential Manchester codes normally specified in LAN standards.

To understand how the FDDI scheme achieves synchronization, one must realize that there are two stages of encoding. In 4B/5B, the encoding is performed four bits at a time. Each four bits of data are encoded into a symbol with five cells such that each cell contains a single signal element (presence or absence of light). In effect, each set of four bits is encoded as five bits. Then, each element of the 4B/5B stream is treated as a binary value and encoded using NRZI. In this code, a binary 1 is represented with a transition at the beginning of the bit interval; there are no other transitions.

The advantage of NRZI is that it employs differential encoding: the signal is decoded by comparing the polarity of adjacent signal elements rather than the absolute value of a signal element. This scheme is relatively robust in detecting transitions in the presence of noise or other distortions; therefore, the NRZI encoding will improve reception reliability. Table 1 shows the symbol encoding used in FDDI.

Since this scheme is encoding 4 bits (16 combinations) with 5 bit patterns (32 combinations), there will be patterns that are not needed. The codes selected to represent the sixteen 4-bit patterns are such that a transition is present at least twice for each 5-bit code. Given an NRZI format, no more than three zeroes in a row can be allowed, since the absence of a transition indicates a zero. The remaining symbols are either declared invalid or are assigned special meaning as control symbols as shown in Table 1.

PHY also provides line states for establishing the station's links with its neighbors (upstream and downstream) and to detect the integrity of the station's links to these neighbors. These line states are used to exchange a handshake with a neighbor. A node receiving a line state on its primary input can respond by sending the proper line state on the secondary output. The line states are composed of a repetition of one or more "I" symbols.


Another item that must be resolved in this sublayer is the issue of timing jitter. Jitter is the deviation of clock recovery that can occur when the receiver attempts to recover clocking as well as data from the received signal. The clock recovery will deviate in a random fashion from the transitions of the received signal. If no countermeasures are used, the jitter will accumulate around the ring.

In LANs, the IEEE 802 standard specifies that only one master clock will be used on the ring, and that the station with the clock will be responsible for eliminating jitter using an elastic buffer. If the ring as a whole runs ahead or behind the master clock, the elastic buffer expands or contracts accordingly.

This centralized clocking method, however, is not practical for a 100M bps ring. At this speed the bit time is only 10 ns, compared to a bit time of 250 ns, at 4M bps, making the effect of distortion more severe. Consequently, FDDI specifies the distributed clocking scheme.

In this environment, each station uses its own autonomous clock to transmit or repeat information onto the ring. Each station has an elastic buffer where data is clocked in at the clock rate recovered from the incoming stream, but is clocked out at the station's own clock rate. This distributed system is considered stronger than the centralized method and minimizes jitter. As a consequence of relocking at each station, jitter does not limit the number of repeaters in the ring, as is the case in LANs where a master clock is used.

Figure 9.
Comparison Coding Schemes



In Manchester code, there is a transition at the middle of each bit period. The mid-bit transition serves as a clock and also as data. A high-to-low transition represents a 1, and a low-to-high transition represents a 0.

In Differential Manchester code, the mid-bit transition is used only to provide clocking. The coding of a 0 (1) is represented by the presence (absence) of a transition at the beginning of the bit period.

Comparison coding schemes used in LANs and/or fiber.

Media Access Control Specification

Layer 2 (Data Link Layer) of the OSI Reference Model is traditionally divided into two sublayers in a LAN context: Link Layer Control (LLC) and Media Access Control (MAC). FDDI only defines MAC, which controls data flow over the ring. The token-passing protocol incorporated into FDDI controls transmission over the network. MAC defines packet formation (headers, trailers, etc.), addressing, and cyclic redundancy checking (CRC). It also defines the recovery mechanisms. This standard defines the following specifications:

1. Services:
 - MAC to LLC Services
 - PHY to MAC Services
 - MAC to SMT Services
2. Facilities:
 - Symbol set
 - Protocol Data Units
 - Fields
 - Timers
 - Frame Counts
 - Frame Check Sum

The FDDI packet format is shown in Figure 10. Packets are preceded by a minimum of 16 IDLE control symbols. The packet itself is characterized by a Start Delimiter composed of the J and K control symbols. This is followed by a Frame Control field that identifies the type of packet. The Destination Address, which follows, identifies the frame recipient. The Source Address is also included to identify which station originated the packet. The address field can be 26 or 48 bits in length. The variable information field follows, along with a Frame Check Sequence field of 32 bits. The check sequence covers the Frame Control Field, the two addresses, and the information field. An End Delimiter, which consists of the T symbol, is transmitted. The

maximum packet length is limited by the size of the elastic buffer in the Physical Sublayer and by the worst case frequency difference between two nodes, the upper bound in 9,000 octets. Figure 10 also shows the format of the token.

Flow control is the other major function of the MAC. In an idle condition, MAC connects to an internal source of IDLE control symbols to be transmitted over the ring. When a Start Delimiter is detected from the ring, MAC switches to a repeat path; the packet is monitored and copied if it is meant for this destination. The packet is simultaneously repeated onto the ring for relaying. The MAC can also inject its own packet or issue a token. Packets are removed only by the originating station. The MAC repeats the packet only until the Sender Address field is detected. If the destination recognizes the Sender Address field as its own station, it will insert IDLE control symbols back onto the ring (the fragmented packet is ignored and removed by any station holding a token for transmission). Stations wishing to transmit must first obtain a token (this is the unique six-symbol packet shown in Figure 10).

The procedures for obtaining the token and the amount of time allowed for data transmission (to retain fairness) are specified in the Timed Token Protocol (TTP). A station obtains the token by performing the stripping function on the incoming token. Only the Start Delimiter field is repeated onto the ring; the station will inject its own information at this juncture. When the packet is sent, the station immediately issues a new token. TTP guarantees a maximum token rotation time.

TTP allows two types of transmission synchronous and asynchronous. In the synchronous mode, stations obtain a predefined amount of transmission bandwidth on each token rotation. The balance of the bandwidth is shared among stations using the asynchronous service. These stations can send data when the token arrives earlier than expected. Any unused capacity left over from synchronous capacity is available to asynchronous traffic, which may be subdivided into up to eight levels of priority.

The amount of time allowed for asynchronous transmission is bounded by the difference of the token's actual arrival time and the expected arrival time. In essence, each station keeps track of how long it has been since it last saw the token. When it next sees the token, it can send synchronous traffic and/or any asynchronous traffic for which time remains available.

Station Management (SMT)

The FDDI Station Management (SMT) specification describes software-based, low-level data link management and integrated network control functions of all stations attached to an FDDI LAN and of the LAN itself. Each FDDI station contains only one SMT entity. SMT initializes the network, monitors error rates and fault conditions in each network segment, and automatically reconfigures the network to isolate problem links. SMT components are Connection Management (CMT), which includes Entity Coordination Management, Physical Connection Management (PCM); and Ring Management (RMT). SMT is intended to operate regardless of equipment type, vendor, protocols, or applications. Figure 11 presents the SMT architectural model.

SMT types (managed objects) have specific attributes indicating state, capabilities, and operation.

SMT managed objects are:

- Station or concentrator (SMT)

Table 1. 4B/5B Codes Used in FDDI

Function or 4-bit group (4B)	Group Code (5B)	Symbol
<i>Starting Delimiter</i>		
First symbol of sequential SD pair	11000	J
Second symbol of sequential SD pair	10001	K
<i>Data Symbols</i>		
0000	11110	0
0001	01001	1
0010	10100	2
0011	10101	3
0100	01010	4
0101	01011	5
0110	01110	6
0111	01111	7
1000	10010	8
1001	10011	9
1010	10110	A
1011	10111	B
1100	11010	C
1101	11011	D
1110	11100	E
1111	11101	F
<i>Ending Delimiter</i>		
Used to terminate datastream	01101	T
<i>Control Indicators</i>		
Logical ZERO (reset)	00111	R
Logical ONE (set)	11001	S
<i>Line Status Symbols</i>		
Quiet	00000	Q
Idle	11111	I
Halt	00100	H
<i>Invalid Code Assignment</i>		
These patterns shall not be transmitted	00001	Void or Halt
because they violate consecutive	00010	Void or Halt
code-bit zeroes or duty cycle requirements.	00011	Void
Some of the codes shown shall	00101	Void
nonetheless be interpreted as a Halt if	00110	Void
received.	01000	Void or Halt
	01100	Void
	10000	Void or Halt

- MAC object(s)
- Path object(s)
- PHY object(s)
- PMD object(s)
- Attachment(s)

Attributes are:

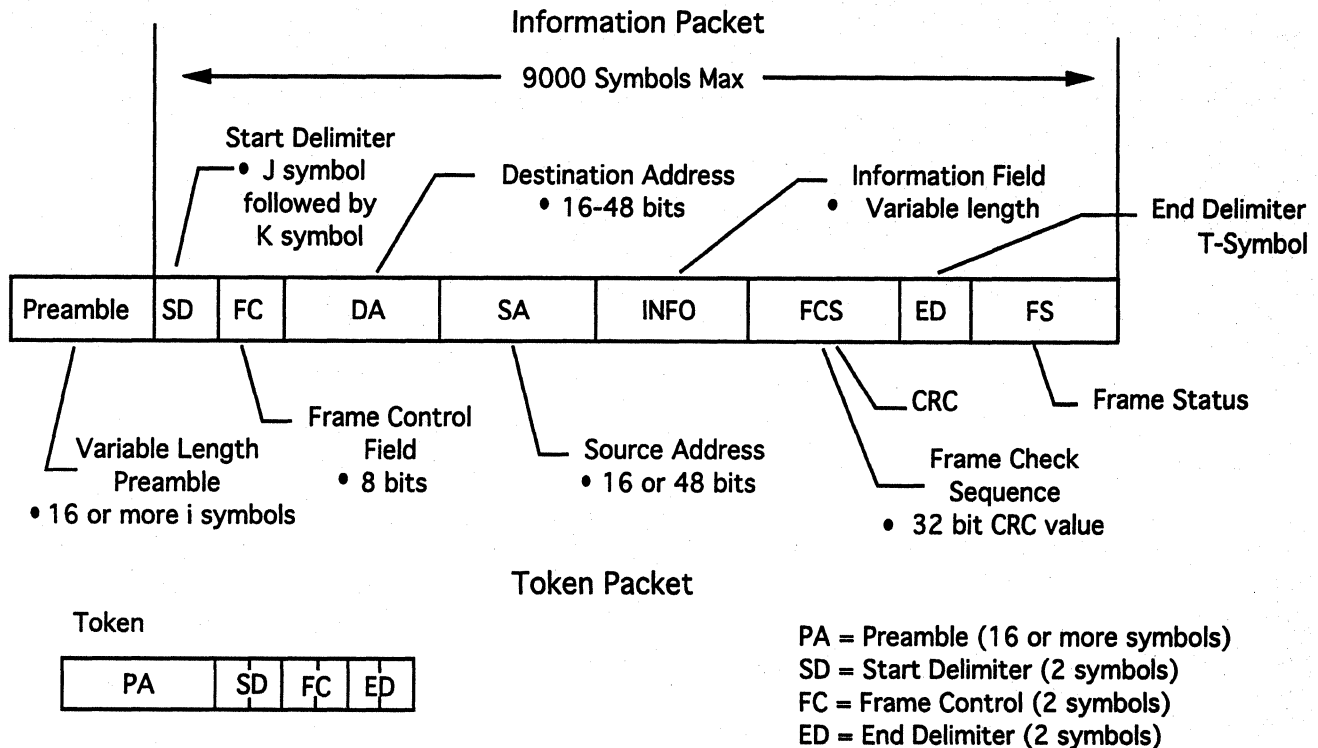
- Attribute Identification (ID)

- Configuration
- Operational (Status, Counters, etc.)

Each attribute is defined in terms of Access Rights, and whether it is Mandatory or Optional. Each attribute also carries an FDDI specification reference or a specific definition if it is not defined in the FDDI specification.

SMT Connection Management (CMT) operates at the logical level, controlling the interface between PHY and

Figure 10.
FDDI Packets



FDDI packet format.

MAC entities in a given station and controlling SMT-to-SMT communications across the ring. When a session requires a connection to another station, CMT causes PHY to send a stream of symbols to the targeted station. Upon receiving the symbols, the receiving station's PHY returns a continuous stream of symbols (primitives) indicating line state, station status, and willingness to carry out the requested action (establish a link). QUIET symbols indicate a disabled link. HALT or alternating HALT and QUIET (MASTER) symbols indicate an operating link and the receiving station's status (master, slave, or peer). A stream of IDLE symbols indicates willingness to connect. Once the link is established, CMT configures the PHY and MAC.

Entity Coordination Management (ECM) controls the optional optical bypass switch, and signals the Physical Connection Management (PCM) entity when the bypass is complete. ECM also performs the Path Test to determine a fault's location.

Physical Connection Management (PCM) initializes the adjacent station's PHYs, and manages signaling. Maintenance support functions are also part of PCM.

Configuration Management (CFM) interconnects PHYs and MACs. It automatically configures these connections according to PCM flags. CFM is defined differently for stations and concentrators.

Ring Management (RMT) relays MAC and CFM status information. It detects stuck signaling beacons, initiates the trace function, detects duplicate addresses and resolves them to allow continued ring operation, and notifies SMT of MAC status.

Parameter Management Frames (PMF) are a newly required feature of FDDI SMT, implemented in Version 7.2.

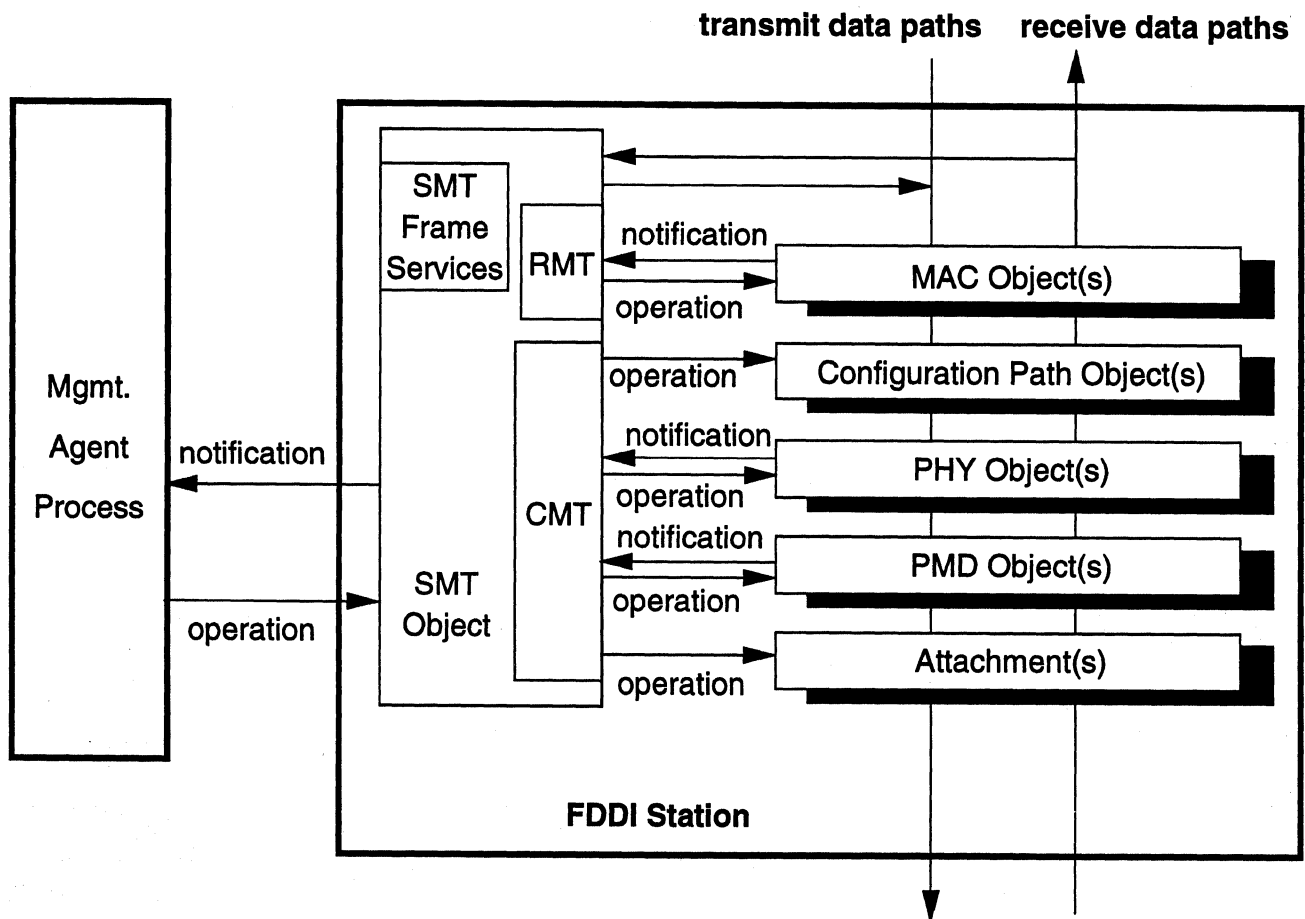
PMF frames provide remote FDDI stations with access to the SMT MIB. SMT Version 7.2 is complete and scheduled for release by the first quarter of 1993. SMT Version 7.2 will interoperate with Version 6.2. In August 1992, Digital Equipment, AT&T/NCR, Motorola, and Distributed Systems Intl. announced a successful SMT interoperability test. The tests demonstrated interoperability between SMT Versions 6.2 and 7.2, as well as between FDDI products from different vendors. Part of the tests involved the operation of the PMF.

Higher Layers

From LLC upwards, ANSI intends FDDI to generally fit traditional protocol stacks. The ANSI FDDI committee has not yet formally drafted protocols for Layer 3 (network) and 4 (transport), which are needed for any type of internetworking. Many suppliers believe the TCP/IP protocol suite can serve this purpose. TCP/IP software was originally developed by the U.S. government for Arpanet, the worldwide packet switched network, but has been used with great success in commercial applications—particularly for internetworking among different LANs. TCP and IP follow layered networking concepts, occupying Layers 4 and 3, respectively, of the OSI model. Vendors have adopted the protocols for Ethernet and other local area networks. Most commercial implementations of the TCP/IP protocol suite include three standardized upper layer protocols: Telnet (virtual network terminal), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP).

Although more versatile than TCP/IP, the OSI internetworking protocols—spanning Layers 3 through 5—are less

Figure 11.
SMT Architectural Model



The relationships between managed objects in an FDDI station and the SMT object.

practical to implement in the real world. Much effort, however, has been expended to migrate to ISO-based standards, particularly under the MAP/TOP thrust, as shown in Figure 5. These market demands are forcing investigations into upgrading the third and fourth layers. The OSI stack also has a well-developed suite of upper-layer applications, including X.400 (electronic mail), File Transfer Access Management (FTAM), and X.500 Directory System Protocol. The federal government mandates its own version of the OSI Reference Model (GOSIP) in government procurements.

Some believe that middle-layer OSI protocols will be changed to look more like TCP/IP. In today's commercial networking applications, however, vendors are blending different protocols' stacks from various sources to match user needs. One vendor's network protocol, for instance, might blend different layers from OSI, TCP/IP, or IBM's SNA. In reality, networking protocols are still evolving, and prospective users must evaluate them with an eye toward future standards.

Open Issues

In May 1991, five networking equipment and semiconductor manufacturers announced that they had joined together to define and publish an open, interoperable solution for transmitting FDDI over shielded twisted-pair (STP) cabling. Advanced Micro Devices, Chipcom, Digital

Equipment, Motorola, and SynOptics published the proposed specification in hopes that additional companies would adhere to it, helping to lower the cost of the needed components, thus accelerating the installation of FDDI to the desktop. The vendor's rationale for developing FDDI-over-STP technology is to allow users to employ existing cabling when migrating from existing LANs to FDDI without installing fiber optic cabling.

The first ANSI-chartered Twisted Pair-Physical Medium Dependent (TP-PMD) ad hoc meeting was held in August 1990. Chipcom, Digital Equipment, and SynOptics shared the results of their independent research into implementing FDDI over STP. The companies' proposals were so sufficiently similar that the companies agreed to work together to develop an open, interoperable proposal.

The proposed solution specifies a Physical Medium Dependent sublayer that uses 150-ohm shielded twisted-pair cables. While the document does not specify a particular design, it does call for the complete replacement of the existing optical PMDs with electrical counterparts. Because all other aspects of the FDDI standard remain the same, no changes are required beyond the PMD level.

A proposal from Crescendo Communications, a Sunnyvale, CA, start-up, provides yet another alternative. According to the company, Crescendo's Copper Distributed

Data Interface (CDDI) technology allows fully FDDI-compliant signals to run over shielded or unshielded twisted-pair wire (UTP), as well as fiber. To address the problems experienced in trying to transmit 100M bps signals over copper, Crescendo developed a new encoding and scrambling scheme. While the vendor claims that the required coding translation is transparent to the user, the scheme requires a new chipset design.

Other vendors working on their own FDDI-over-UTP designs are AT&T and British Telecom. All interested vendors have submitted their designs to ANSI and, until a formal standard for FDDI over copper is set, users are unlikely to implement any of the competing technologies in significant volume. The FDDI over copper is expected to be a stable working draft by the end of 1992.

As mentioned earlier, Codenoll Technologies introduced a nonstandard version of the FDDI interface board for Extended Industry Standard Architecture (EISA) computers. While using an 850-nm. LED in the Codenoll design does reduce the cost of FDDI connectivity, users will be unable to operate with standard-based equipment containing 1300-nm. transmitters and receptors.

Finally, a low-cost alternative fiber PMD forwarded to ANSI in June 1992 provides a cost enhancement for networks that do not require the standard's full distance.

FDDI-II and Other Future Efforts

Some discussion is under way to define a scheme where available network bandwidth is divided between voice and data using time-division multiplexing. FDDI-II, a superset of FDDI, is being defined as an upward-compatible, fiber-based LAN incorporating the current data capabilities in addition to the ability to handle voice and T1-compressed video traffic.

The FDDI-II proposal will follow the original FDDI standard, adding a fifth document to the present standard. The new document, Hybrid Ring Control, describes a hybrid operating mode comprising the packet-switching scheme used in FDDI and an isochronous transport mode similar to that used in public switched networks. Adding the isochronous mode will enable networks based on the expanded standard to transfer pixel data—key to video and computer graphics transport—in addition to circuit switched voice signals.

Because it builds on the existing FDDI framework, FDDI II should be completed and approved very quickly. The Hybrid Ring Control document should be an approved standard by the end of 1992.

Some people are already thinking of even higher rates, particularly when considering voice and data. The FDDI follow-on, or FFOL, is under development at ANSI. The FFOL project proposal covers the capability to operate as

a backbone for multiple FDDI networks; interconnection to wide area networks, including Broadband Integrated Services Digital Networks (B-ISDN); a data rate between 600M bps and 1.25G bps initially, with intermediate data rates matched to the Synchronous Digital Hierarchy (SDH) underlying SONET and eventual support for data rates up to 2.4G bps; duplex links; support for existing FDDI cabling, where possible; and support for both single mode and multimode fiber. If the ANSI effort keeps on schedule, the family of standards comprising FFOL should be complete by late 1995.

There is also a separate effort called High-Speed Channel being undertaken in X3T9.3. This is a point-to-point system for digital data interface (channel extension rather than a network configuration intrinsic in FDDI). These systems are designed to carry 400M, 800M, or 1600M bps on very short copper links, but they are being developed to be compatible with future fiber optic links operating over longer distances.

Sponsoring Organizations

FDDI Information Sources

The chairperson of the ASC X3T9.5 task group is Gene Milligan of Seagate (Oklahoma City, OK).

ANSI is involved in the form of the X3 Secretariat, as is the Computer Business Equipment Manufacturers Assn. (CBEMA).

ANSI's address is: 1430 Broadway, New York, NY 10018. Telephone (212) 642-4900.

CBEMA's address is: Suite 500, 311 First Street NW, Washington, DC. Telephone (202) 737-8888.

Specification X3.139 can be obtained from Global Engineering Documents.

Specification X3.148 can be obtained from Global Engineering Documents.

Specification X3.166 can be obtained from Global Engineering Documents.

Draft Specification X3T9.5/84-49 is available from ANSI for review.

The address of Global Engineering Documents is: 2805 McGraw Avenue, P.O. Box 19539, Irvine, CA 92714. Telephone (714) 261-1455. ■

ANSI Fiber Distributed Data Interface (FDDI) Standards

In this report:

Applications	4
OSI Reference Model	6
FDDI Specifications	10
FDDI-II and Other Enhancement Efforts	17
FDDI Information Sources	18

Synopsis

Editors Note

In October 1982, the American National Standards Institute (ANSI), Committee X3T9.5, was chartered to develop a high-speed data networking standard. Known as FDDI (Fiber Distributed Data Interface), the standard specifies a packet switched LAN-to-LAN backbone that transports data at high throughput rates over a variety of multimode fibers. FDDI is a token-passing network employing two fiber pairs operating at 100M bps.

Report Highlights

The FDDI standard describes an optical fiber-based local area network operating at 100M bps. The standard covers the first two layers of the OSI Reference Model. The LAN is configured as dual redundant fiber rings to protect against disruptions caused by station failures. The standard is essentially complete.

The last of the four FDDI components, Station Management (SMT), was forwarded to the ANSI X3T9 committee to be prepared for a letter ballot. The draft standard is technically stable and should be ratified by the end of 1991.

Since the initial standard was proposed, ANSI has begun work on enhancements to FDDI. FDDI II, which is considered a superset of FDDI, will add the capability of carrying voice and video signals over FDDI-based networks. FDDI follow-on, or FFOL, will allow interconnection of multiple FDDI networks, connection to wideband facilities, and Synchronous Optical Network (SONET) compatibility. FDDI II should be complete by the end of 1991, FFOL by 1995.

Other efforts focus on implementing FDDI over twisted-pair cabling and specifying a single-mode fiber for use in FDDI networks. This report details the FDDI draft standard and discusses future directions for the standard.

Analysis

Technology Overview

FDDI addresses the bottom two layers of the OSI model. FDDI's designers expect users will use ISO protocols for the other layers, where possible. The optical-based FDDI LAN was designed to provide the same type of serial interconnection provided by LANs while providing the high bandwidth, inherent noise immunity, and security offered by fiber. At FDDI's inception in 1982, fiber was used mostly for point-to-point applications, and not for the many configurations supported by LANs. In this sense FDDI was a breakthrough. Although it is possible to achieve considerably higher data rates over fiber (up to 3.7G bps with current point-to-point technology, and 500M bps on rings), higher rates result in significantly increased costs and shorter transmission distances between repeaters. Its designers intend FDDI to provide relatively inexpensive connectivity and therefore focused on the 100M bps rate. FDDI can be configured to support a sustained transfer rate of approximately 80M bps. The remaining bandwidth is reserved for various overhead functions.


FDDI is a token-passing, dual-ring network accommodating synchronous and asynchronous data transmission as well as isochronous channels for realtime digitized voice and compressed video. Unlike existing open standards for LANs, where fiber optic variants have followed copper implementations, FDDI has been designed from the start as a fiber optic network. This has involved standardization issues in such areas as duplex optical connectors, fiber characteristics, optical bandwidth, bypass relays, and cable assemblies. The FDDI ring is designed for an overall bit error rate of less than 10^{-9} . The network can tolerate up to 2 km. of fiber between stations and can support a total cable distance of 100 km. around the ring with 500 attachments (1,000 physical connections and a total fiber path of 200 kilometers). FDDI topology is a counterrotating, token-passing ring

(note the arrows in Figure 1). FDDI, however, is not part of the well-established IEEE 802 family of LAN standards.

FDDI operates at 1300 nanometers (nm.). Current transmitter/receiver fiber technology operates at 850 nm., 1,300 nm., or 1,550 nm. While performance increases with wavelength, so does cost. For local data communications, both in LANs and in point-to-point applications employing fiber optic modems, 850-nm. light sources are typically employed; however, this technology becomes unfeasible for 100M bps beyond a couple of miles. At the other end of the range (1,550 nm.), the system becomes expensive and may provide unnecessary bandwidth. The committee designing FDDI also investigated short-wavelength implementations. It became evident that to meet all the requirements, particularly the two-kilometer station-to-station spacing, the system would require 1,300-nm. wavelength. Using 1,300-nm. technology, less expensive light-emitting diodes (LEDs) provide distance and data rates within the range desired for LANs and LAN backbones. The issue of fiber size was settled after the wavelength decision.

The FDDI standard directly addresses the need for reliability. A backbone system transports a large number of user sessions, and its loss would be serious. FDDI incorporates three reliability-enhancing features. First, a failed or unpowered station can be bypassed by an optional automatic optical bypass switch; second, wiring concentrators are used in a star wiring strategy to facilitate fault

Figure 1.
The FDDI Environment



isolation and correction; and third, two rings interconnect stations so that failure of a repeater or cable link results in the automatic reconfiguration of the network (loopback).


FDDI consists of four adopted or draft proposals. The Media Access Control (MAC) is specified in X3.139-1987, approved on November 5, 1986. The Physical layer standard (PHY) is contained in X3.148-1988. The Physical Medium Dependent (PMD) is specified in X3.166-1988. Station Management (SMT) is described in draft proposal X3T9.5/84-49.

MAC was the first FDDI standard completed. MAC specifies access to the medium, addressing, data checking, and frame generation/reception. PMD specifies the optical fiber link and related optical components. PHY specifies encode/decode, clocking, and data framing. A Station Management standard specifies the control required for proper operation of stations on the ring. Services include station management, configuration management, fault isolation and recovery, and scheduling procedures.

The token-passing protocol used in FDDI is based on the IEEE 802.5 standard for a 4M or 16M bps twisted-pair ring. When no station is transmitting, the token (a small control packet) circulates around the ring. When a station needs to transmit data, it waits until it detects the token. The station then removes the token from the ring and transmits the data in packetized form. The data packet will circulate around the ring until it reaches its intended recipient. That station copies the packet and returns it to the ring. The packet continues to the transmission station, which removes it from the network and places the token back on the ring for the next station to use.

The choice of FDDI's ring topology is based on optical communications characteristics. Bus and passive star topologies require the optical transmission to be detected at several sources simultaneously. Although practical fiber optic taps are available, attenuation is still such that the number of nodes is relatively limited. Since fiber optic transmission is best suited to a point-to-point configuration, FDDI included this aspect in its definition. Two topologies are possible with point-to-point links: the active star and the ring. The active star has a single point of failure. A defective hub

Figure 2.
Failure Recovery



Ring rearrangement under failure.

incapacitates the entire network. A similar problem affects a single ring. Consequently, the FDDI specification calls for a dual ring.

FDDI Station Types

Two FDDI station types are allowed: Class A and Class B. Class A stations are dual-attachment stations. They connect to both the primary and secondary rings of the network. Data flows in opposite directions on the two rings. A Class A station can act as a wiring concentrator to interconnect several Class B (single attachment) stations. Wiring concentrators give the network administrator a common maintenance point for a number of stations (see Figure 1). Ring reconfiguration following a failure is shown in Figure 2. Here, the link between two Class A devices is broken. The two stations detect the failure and patch the network by looping the data path onto the secondary ring, thus creating a single ring.

Figure 3 illustrates a break in the cable between a Class A and a Class B device. Here, communication continues over the primary link, as the Class A device detects the failure and makes appropriate internal modifications. The Class B device, however, remains detached. Class B devices trade lower cost against the fault tolerance of the more sophisticated Class A device.

The secondary ring can also be used to carry traffic. This gives the fully configured FDDI system 200M bps of effective throughput. When the

Glossary of FDDI Terms

The definitions given here apply to the ANSI FDDI standards. As the final section of the standard (Station Management) is completed, additional definitions may be added to this list.

Asynchronous—A class of data transmission service whereby all requests for service contend for a pool of dynamically allocated ring bandwidth and response time.

Attachment—A Port or pair of Ports, optionally including an associated optical bypass, that are managed as a functional unit. A dual attachment includes two ports: a Port A and a Port B. A single attachment includes a Port S.

Bypass—The capability of a node to optically isolate itself from the FDDI network while maintaining the continuity of the cable plant.

Capture—The act of removing a Token from the ring for the purpose of Frame transmission.

Claim Token—A process whereby one or more stations bid for the right to initialize the ring

Code Bit—The smallest signaling element used by the Physical Layer for transmission on the medium.

Code Group—The specific sequence of five code bits representing a DDL symbol.

Concentrator—An FDDI node that has additional ports beyond those required for its own attachment to an FDDI network. These additional ports (type M) are for attaching other FDDI nodes (including other concentrators) in a tree topology.

Connection Management (CMT)—That portion of the Station Management (SMT) function that controls network insertion, removal, and connection of PHY and MAC entities within a station.

Counterrotating—An arrangement whereby two signal paths, one in each direction, exist in a ring topology.

Dual Attachment Concentrator—A concentrator that offers a dual attachment to the FDDI network and is capable of accommodating a dual (counterrotating) ring.

Dual Ring (FDDI dual ring)—A pair of counter-rotating logical rings.

Entity—An active service or management element within an Open Systems Interconnection (OSI) layer or sublayer.

Fiber Optic Cable—A cable containing one or more optical fibers.

Frame—A PDU transmitted between cooperating MAC entities on a ring, consisting of a variable number of octets and control symbols.

Jitter, Random—The probabilistic offsets of pulse transition edges from the expected time. Includes both Duty Cycle Distortion and Data Dependent Jitter.

Jitter, Systematic—The deterministic offsets of pulse transition edges from the expected time. Some sources of systematic jitter are differences in rise and fall times and propagation delays.

Logical Ring—The set of MACs serially connected to form a single ring. A fault-free FDDI network provides two logical rings.

Media Access Control (MAC)—The Data Link Layer responsible for scheduling and routing data transmissions on a

shared-medium local area network (e.g., an FDDI ring).

Media Interface Connector (MIC)—A mated connector pair that provides an attachment between an FDDI node and a cable plant. The MIC consists of two parts: an MIC plug and an MIC receptacle.

MIC Plug—The male part of the MIC which terminates a fiber optic cable.

MIC Receptacle—The female part of the MIC which is contained in an FDDI node.

Network (FDDI Network)—A collection of FDDI nodes interconnected to form a trunk, tree, or a trunk with multiple trees. This topology is sometimes called a dual ring of trees.

Node—A generic term applying to an active element in an FDDI network (station or concentrator).

NRZ—Non Return to Zero, a technique where a polarity level (+ or -) represents a logical "1" (one) or "0" (zero).

NRZ1—Non Return to Zero Invert on Ones, a technique where a polarity transition represents a logical "1" (one). The absence of a polarity transition denotes a logical "0" (zero).

two rings merge to support a backup configuration, the network data rate drops to 100M bps.

Applications

Although optical fiber is widely deployed in the telecommunications environment (long-hauls, interoffice, feeder plant, and so on), it still has not exploded in the LAN environment. (This is not true in Japan, where fiber optic LANs dominate.)

Physical Connection—The full-duplex physical layer association between adjacent PHY entities (in concentrators or stations) in an FDDI network, i.e., a pair of Physical Links.

Physical Layer (PHY)—The Physical Layer responsible for delivering a symbol stream produced by an upstream MAC Transmitter to the logically adjacent downstream MAC Receiver in an FDDI ring.

Physical Link—The simplex path (via PMD and attached medium) from the transmit function of one PHY entity to the receive function of an adjacent PHY entity (in concentrators or stations) in an FDDI network.

Port—A PHY entity and a PMD entity in a node, together creating a PHY/PMD pair, that may connect to the fiber media and provide one end of a physical connection with another node.

Primitive—An element of the services provided by one entity to another.

Protocol Data Unit (PDU)—Information delivered as a unit between peer entities which may contain control information, address information,

and data (e.g., a Service Data Unit from a higher layer).

Receive—The action in a station in accepting a Token, frame, or other signal sequence from the incoming medium.

Receiver (optical)—An electro-optical circuit that converts an optical signal to an electrical logic signal.

Repeat—The action of a station in receiving a Token or Frame from the adjacent upstream station and simultaneously sending to the adjacent downstream station. The FDDI MAC may repeat received PDUs (Tokens and Frames) but does not repeat the received signal stream between PDUs. While repeating a Frame, MAC may copy the data contents and modify the control indicators as appropriate.

Repeater—A Physical Layer relay in an FDDI network.

Ring—A set of stations wherein information is passed sequentially between stations, each station in turn examining or copying the information, finally returning it to the originating station. In FDDI usage, the term

“ring” or “FDDI ring” refers to a dual (counterrotating) ring.

Service Data Unit (SDU)—The unit of data transfer between a service user and a service provider. The MAC SDU is the data contents of a Frame. The PHY SDU is a symbol.

Services—The services provided by one entity to another. Data services are provided by a higher level entity; management services are provided to a management entity in the same or another level.

Single Attachment Concentrator—A concentrator that offers a single attachment to the FDDI network.

Single Attachment Station—A station that offers a single attachment to the FDDI network.

Station—An addressable logical and physical node on a ring capable of transmitting, repeating, and receiving information. A station has exactly one SMT, at least one MAC, at least one PHY, and at least one PMD entity.

Station Management (SMT)—The supervisory entity within an FDDI station that monitors and controls the various FDDI entities including PMD, MAC, and PHY.

Symbol—The smallest signaling element used by MAC, i.e., the PHY SDU. The symbol set consists of 16 data symbols and 8 control symbols. Each maps to a specific sequence of five code bits as transmitted by the Physical Layer.

Synchronous—A class of data transmission service whereby each requestor is preallocated a maximum bandwidth and guaranteed a response time not to exceed a specific delay.

Token—An explicit indication of the right to transmit on a shared medium. On a token-ring, the Token circulates sequentially through the stations on the ring. At any time it may be held by zero or one stations. FDDI uses two classes of Tokens: restricted and unrestricted.

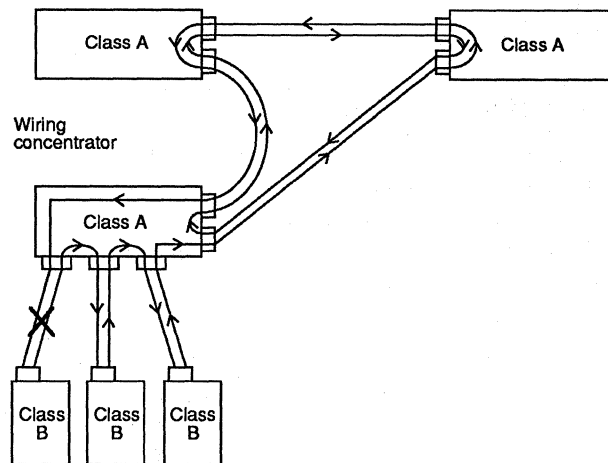
Transmit—The action of a station in generating a Token, Frame, or other symbol sequence and placing it on the outgoing medium.

Transmitter (optical)—An opto-electrical circuit that converts an electrical logic signal into an optical signal.

Three reasons can be identified: 1) increased technical complexity compared to passive copper and coaxial cable; 2) cost considerations—especially the cost of writing off huge existing cable investments; and 3) lack of a workable standard. FDDI solves the third problem, and in the process it also begins to resolve the second.

FDDI allows designers to 1) build larger capacity LANs or LAN backbones to serve new data needs (file transfer, graphics, and so on) and some voice needs; or 2) interconnect LANs in metropolitan area networks (MANs). Thus, FDDI can be

Figure 3.
Class B Failure



Failure of Class B stations.

used directly as a LAN or as a backbone to interconnect slower speed LANs into a single network over relatively large geographies.

The initial application of FDDI as a “backend” interconnect for high-powered computing devices and peripherals required a high degree of fault tolerance and data integrity. As developers proceeded, it became obvious that FDDI could also serve high-speed “front-end” applications. Front-end applications include terminal-to-terminal and terminal-to-server communications typical of a LAN. In large networks (3,000 to 10,000 terminals, particularly when workstations are involved), the aggregate demand for network resources can overwhelm a 4M/16M bps token-ring or 10M bps Ethernet LAN. At that point, FDDI’s bandwidth becomes important.

A high-speed FDDI ring is ideal as a backbone for other “departmental” LANs, typically operating at lower speeds. High capacity on LANs can be achieved in two ways: by using multiple channels at low speeds or by using one channel at a relatively high speed. The multiple-channel approach can be used with a broadband bus LAN. A drawback of this approach is that bridges must be provided between channels, and the architecture must be designed to avoid high rates of interchannel traffic and bottlenecks at the bridges.

The one area where FDDI is not well suited is broadband LANs carrying full analog video (6MHz). A coaxial-based broadband LAN can easily carry multiple channels of video, as well as data, in analog form. Digitizing a 6MHz TV channel,

however, results in a 45M to 90M bps data rate (45M bps is a slightly compressed version), which can easily swamp an FDDI backbone. FDDI-II, discussed later in this report, has been designed to handle video efficiently.

FDDI is not the end, however. Even larger bandwidths are envisioned in the next five years under the thrust of Broadband ISDN (BISDN). BISDN for video needs, particularly in a (High Definition TV (HDTV)) environment, will require approximately 150M bps of dedicated bandwidth per channel. Planners are now discussing delivery of up to three channels per domicile, reaching into the 600M bps range. FDDI is a shared-medium technology; uncompressed video applications and/or local loop applications may not be capable of using networks built on the original FDDI standards.

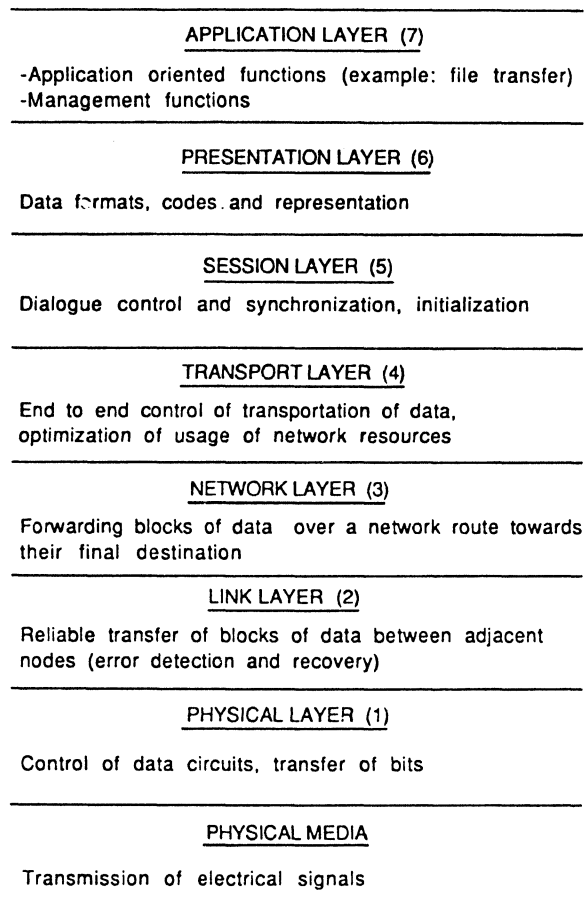
OSI Reference Model

To fully understand the FDDI standard, some knowledge of the OSI model and how it applies to the LAN environment is necessary. A brief discussion of the OSI model is followed by a more detailed description of FDDI specifications and emerging vendor products. Readers intimately familiar with this material can skip to the section entitled FDDI Specifications.

The OSI Reference Model imposes order and structure on data communications. To achieve orderly communications, many functions must be performed. It is natural to group these functions into layers that share task affinity and logical proximity. OSI layers are hierarchical in the sense that a given layer calls for the services of the layer immediately beneath it. The given layer cannot ask for the services of a layer at a higher level, nor can it skip the layer beneath it to directly reach a lower layer, or even jump into the middle of another layer. The model includes precise definitions of the services provided by each layer to its next higher layer and request procedures.

Services defined for a given layer are, in turn, employed by the layer immediately above it. Each transmitting layer passes down to the lower layer (and up to the next layer, at the receiving end) blocks of data requiring processing for transmission, manipulation, or service. These layers normally attach a characteristic header that contains appropriate information (such as the real network

Figure 4.
OSI Functions



Functions performed by the layers of the OSI Reference Model.

address, block number, and so on). The headers are physically nested, with lower layer headers being outermost and higher layer headers being innermost. It is through the use of these well-defined headers that the protocols between the remote open systems are accomplished. To effectuate the OSI model, the International Organization for Standardization has formulated standards—specifications for how information is coded and passed between communicating partners. Only the protocols need to be implemented by a prospective vendor. Users should note that, when working at a terminal connected to a host or a PC connected to a LAN, all seven layers of the architecture must be employed.

The Reference Model and the service definitions are only structures for discussing the tasks involved in communicating between open systems.

The OSI Reference Model is described by document ISO 7498, which was adopted as a standard in 1984.

LAN Protocol Suites

In a LAN environment, one typically defines Layer 1 and Layer 2 standards specific to local area networks. Layers 1 and 2 are defined by the IEEE 802 standards. With some “internetworking” protocols defined at Layer 3 (such as ISO 8473), and typically some connection-oriented Transport protocols (for instance, FTAM and MHS), sessions can then employ the normal protocol suite up to Layer 7 described above. Prior to the establishment of ISO-based standards, however, the use of Transmission Control Protocol/Internet Protocol (TCP/IP) was common for the upper layers. Below, one finds descriptions of the standards at Layer 1 and Layer 2, which are specific to LANs.

Logical Link Control

IEEE Standard 802.2-1985 (ISO 8802/2) describes the peer-to-peer protocol procedures for the transfer of information and control between any pair of Data Link Layer Service Access Points (SAPs) on a local area network.

Medium Access Control

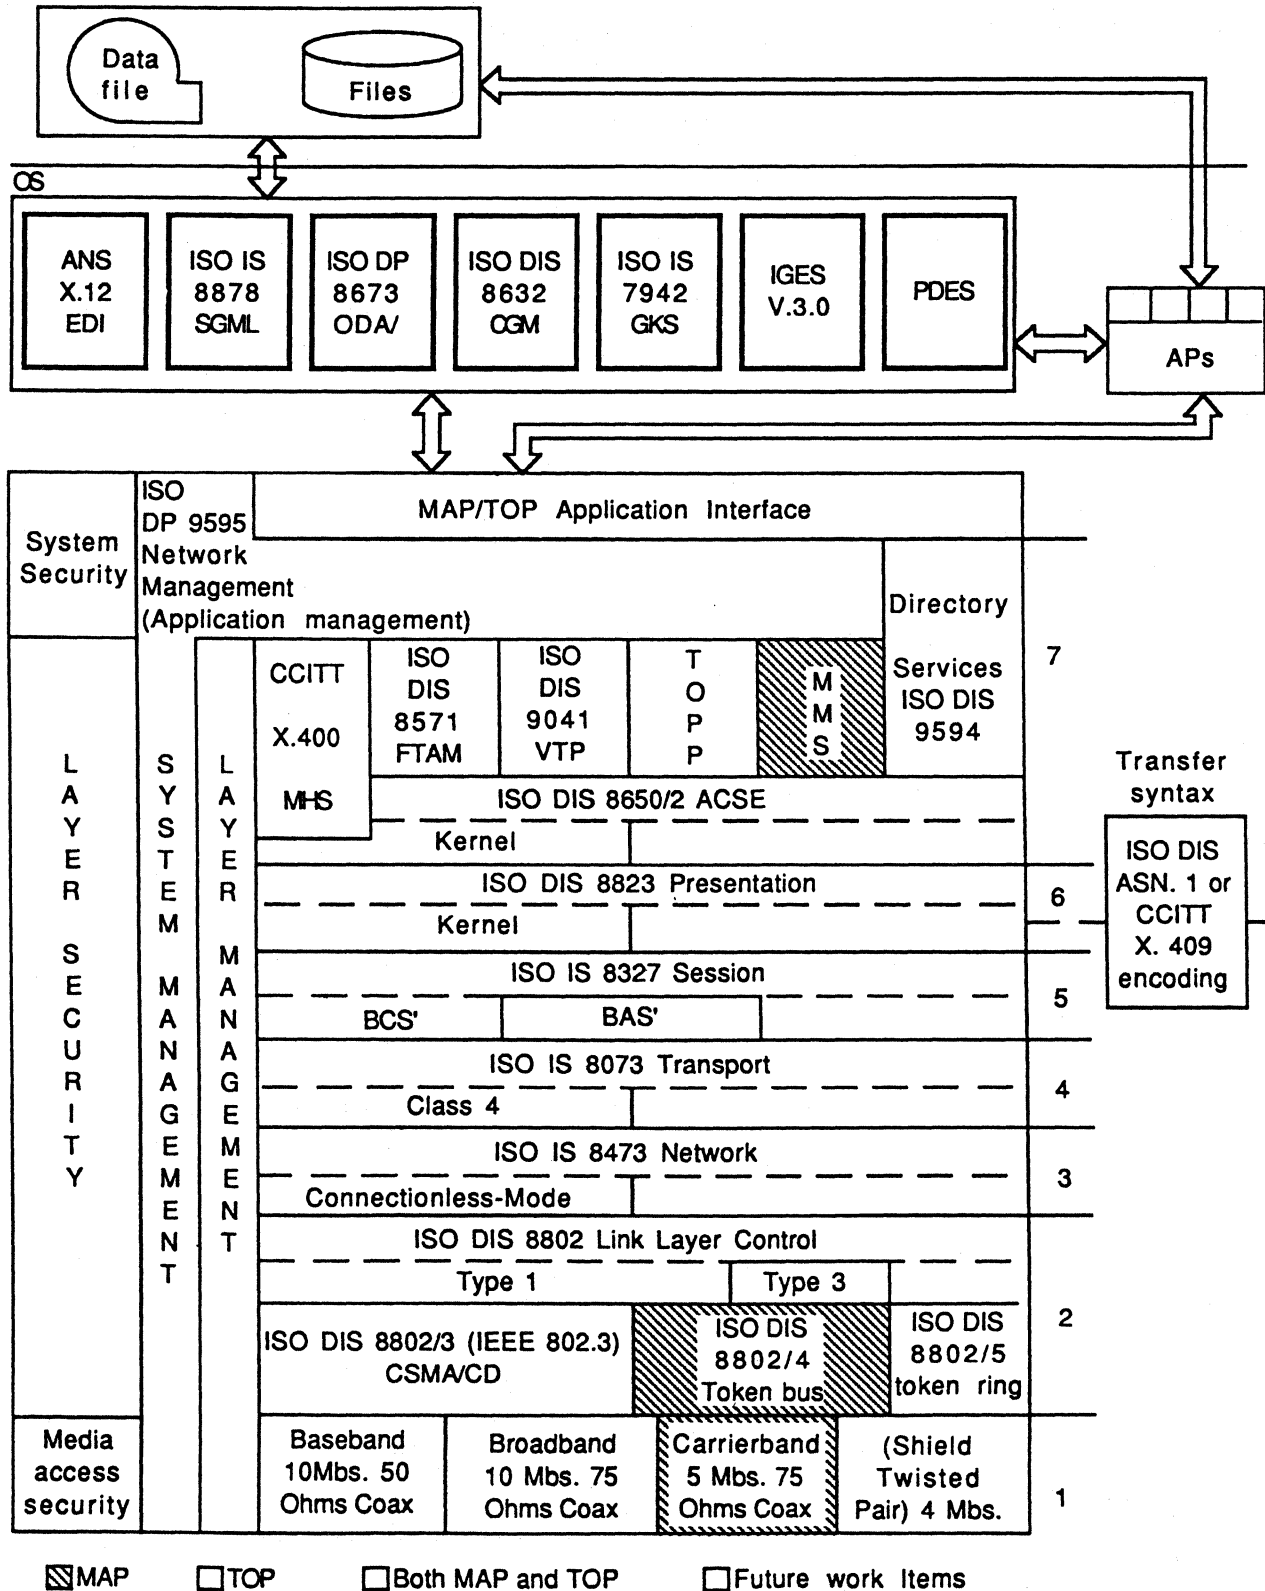
Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

IEEE Standard 802.3-1985 (ISO 8802/3) provides a medium access method by which two or more stations share a common-bus transmission medium. The standard applies to several media types and provides the necessary specifications for a 10M bps baseband local area network.

Token-Passing Bus Access Method

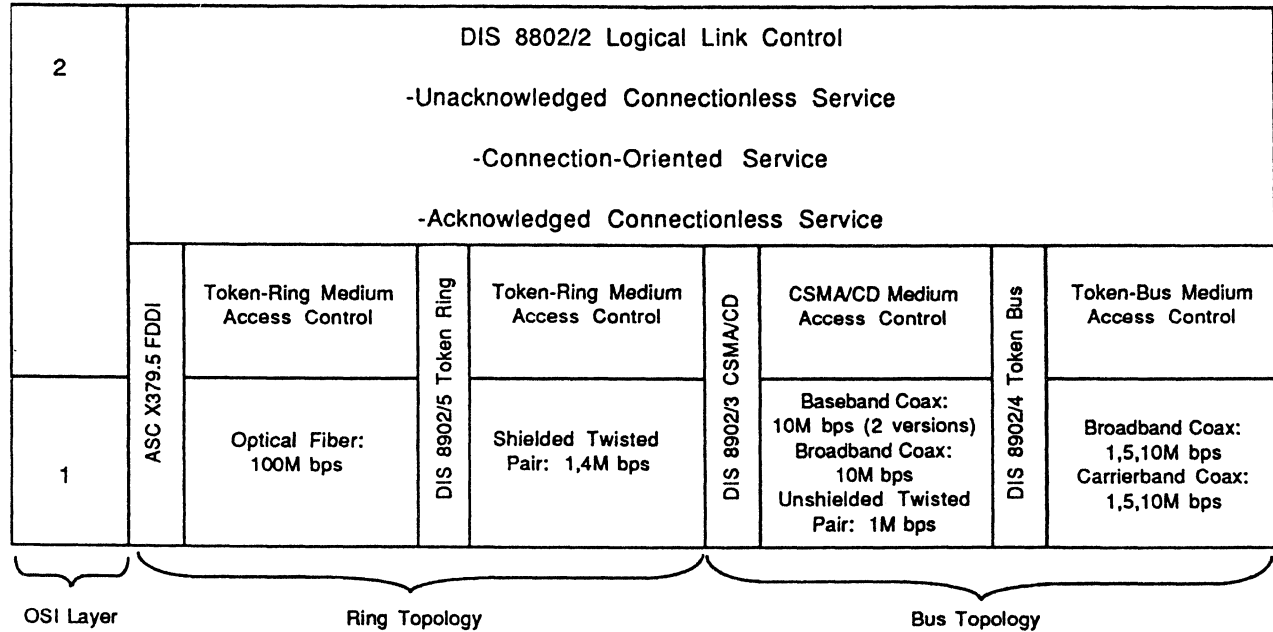
IEEE Standard 802.4-1985 (ISO 8802/4) deals with all elements of the token-passing bus access method and its associated physical signaling and media technologies. The goal is to achieve compatible interconnection of stations in a LAN. The access method coordinates the use of the shared medium among the attached stations. It specifies the electrical and physical characteristics of the transmission medium, the electrical signaling used, the frame formats of the transmitted data, the actions of a station upon receipt of a data frame, and

Figure 5.
ISO Suites



ISO suites in a LAN environment.

Figure 6.
ISO Lower Layers



Lower layers in a LAN environment.

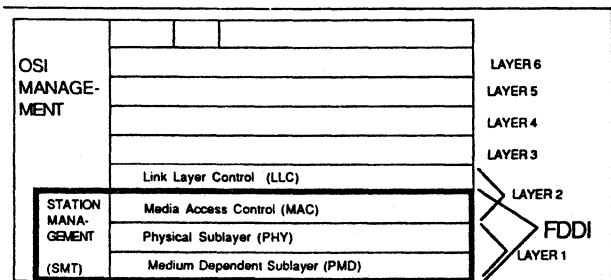
the services provided at the conceptual interface between the Medium Access Control sublayer and the Logical Link Control sublayer above it.

Token-Passing Ring Access Method

IEEE Standard 802.5-1985 (ISO 8802/5) specifies the formats and protocols used by the token-passing ring medium at the control sublayer and physical layer. It also specifies the means of attachment for the Token-Passing Ring Access Method. The protocol defines the frame format including delimiters, addressing, and frame check sequence and includes timers, frame counts, and priority stacks. It also defines the medium access control protocol and provides finite-state machines and state tables supplemented with descriptions of the algorithms. It identifies the services provided by the Medium Access Control sublayer to the Logical Link Control sublayer and the services provided by the physical layer to the Medium Access Control sublayer. These services are defined in terms of service primitives and associated parameters. Also defined are the physical layer functions of symbol encoding and decoding, symbol timing and latency buffering, and the 1M bps and 4M bps, shielded twisted-pair attachments of the station to the medium.

Figure 5, modeled after documentation of the German Commission for Computer-Integrated Manufacturing, depicts typical OSI protocol suites for LANs, particularly in a MAP/TOP environment. Figure 6 depicts some additional details at the lower layers, including the positioning of FDDI as it relates to other LAN standards. Once again, note that when a user works at a PC connected to a LAN or FDDI ring, all seven layers of the protocol architecture must be employed, as shown in Figure 5. Thus, while the FDDI standard concentrates on Layers 1 and 2, the upper layers are required to accomplish end-to-end communications.

Figure 7.
FDDI Protocols



FDDI Specifications

FDDI is defined according to the OSI Reference Model and LAN protocol architecture. Layer 1 (physical layer) is specified in two documents: the FDDI Physical Medium Dependent (PMD) and the FDDI Physical Sublayer (PHY). (See Figure 7.) The Physical Layer provides the medium, connectors, optical bypassing, and driver/receiver requirements. It also defines encode/decode and clock requirements for framing data for transmission on the medium or to the higher layers of FDDI.

When the FDDI committee realized there would be considerable discussion on fibers, connectors, and other hardware, they decided to break the standardization of the OSI Physical Layer into two pieces. In this way, the relatively noncontroversial issues—like coding and other matters that IC chip manufacturers need to know to begin design—could be put in a formal document and approved independently of other items pertaining to the standard.

The Data Link Layer is also divided into two sublayers:

1. A Media Access Control portion that provides fair and deterministic access to the medium, address recognition, and generation and verification of frame check sequences. Its primary function is the delivery of frames, including frame insertion, repetition, and removal.
2. A Logical Link Control portion that provides a common protocol for data assurance services between the Media Access Control and the Network Layer.

Physical Medium Dependent (PMD) Specification

PMD defines the optical interconnecting components used to form a link. It describes the wavelengths for optical transmission, the fiber optic connector, the functions of the optical receiver, and (as an option) the bypass switch that can be incorporated into the station. It specifies the optical channel at the bulkhead of a station. The source is defined to radiate in the 1,300-nm. wavelength. PMD also describes the peak optimal power, optical rise and fall times, and jitter constraints. The minimum rise/fall time is 0.5 nanosecond. The standard includes the following specifications:

1. Services: PMD to PHY Services; PMD to SMT Services
2. Media attachment
3. Media Signal Interface
4. Interface signals
5. Cable Plant Interface Specification.

Multimode fibers are employed (at least initially) up to a distance of two km. Optical fiber dimensions are specified in terms of its core diameter and the outer diameter of the cladding layer. Fiber specifications are 62.5/125 micron (core diameter/cladding diameter) and 85/125 micron. The nominal numerical aperture is around 0.26. Applicable standards for the fiber itself are EIA-455-48 (core), EIA-455-27 or EIA-455-48 (cladding), and EIA-455-57 (aperture).

Listed in the appendix of the draft standard are two other fibers: a 50-/125-micron fiber and a 100-/140-micron fiber. These two fibers have not been extensively studied; the maximum achievable distance of two km. specified in the standard may not be possible. Thus, these two fibers are not officially part of the standard but are listed as "alternatives." Smaller diameters offer higher bandwidths but also more expensive, higher loss connectors. The 50/125 fiber has been used primarily for military applications and is not likely to be as widely available as the other fibers. The 100/140 fiber has been added primarily because it is used in IBM's cabling system and a number of customers have already installed it. On the other hand, the 62.5/125 fiber has been in production for some time and has become common for local applications, such as AT&T's Premises Distribution System (PDS). Component costs for this type of fiber are dropping most rapidly.

Some observers believe it is unfortunate that the FDDI standard could not specify a single fiber type, since this might have lowered costs and made it easier for the customer to start small and expand later. The FDDI committee settled on a 62.5-micrometer core, with advisory information about 50-, 85-, and 100-micrometer fiber sizes. While a debate remains about these other sizes, the issue is not really a critical one: as long as the fiber can meet the optical power, channel bandwidth, and distance requirements, it can conform to the FDDI standard at the interface between the FDDI box and the network, independent of the fiber type (the

charts and data presentations in the PMD have been written so that they can be applied to various fiber sizes).


Like all Layer 1 specifications, PMD defines the duplex connector used for FDDI access. The primary and secondary ring connections to each Class A station are attached simultaneously using the duplex connector and a dual-fiber cable. The connectors can be used for both Class A-to-Class A as well as Class B-to-Class A (wiring concentrator) links. The bypass relay connects the optical inputs (at the primary and secondary rings) directly to the optical output in case of a station or link failure, allowing the ring to maintain continuity.

In the 1,300-nm. region, the dispersion due to multimode interference is at a minimum. The combination of physical parameters selected ensure the desired 10^{-9} bit error rate. LEDs (either surface emitting or edge emitting) are implicitly assumed in PMD; however, PMD does not specify the emitter must be an LED. It could also be a laser, as long as the optical interface parameters at the optical port are met. At some future point, some manufacturers may include lower cost local-loop-grade laser emitters, or even long-haul-grade LDs in an FDDI package, by adjusting the optical output at the optical port to conform with the standard. For the foreseeable future, however, all manufacturers pursuing FDDI products will use LEDs. At least one vendor, Codenoll, has elected to use an 850-nm. LED. While this wavelength does not meet the standard, Codenoll claims that the only effect of the change is to shorten the maximum allowable distance between stations from 2,000 meters to 1,000 meters. Because Codenoll's products are microcomputer adapter boards, the distance limitation is less significant than it is in internetworking or mainframe attachment environments. Such nonstandard implementations, however, cannot communicate with devices using standard 1,300-nm. components, requiring users to install a nonstandard device at both ends of the link.

Physical Sublayer Specification

PHY represents the upper sublayer within OSI Layer 1. It defines the encoding scheme used to

Figure 8.
Comparison Coding Schemes



In Manchester code, there is a transition at the middle of each bit period. The mid-bit transition serves as a clock and also as data. A high-to-low transition represents a 1, and a low-to-high transition represents a 0.

In Differential Manchester code, the mid-bit transition is used only to provide clocking. The coding of a 0 (1) is represented by the presence (absence) of a transition at the beginning of the bit period.

Comparison coding schemes used in LANs and/or fiber.

represent data and control symbols. It also describes the method for retiming transmission within the node. The standard includes the following specifications:

1. Services
 - PHY to MAC Services
 - PHY to PMD Services
 - PHY to SMT Services
2. Facilities
 - Coding
 - Symbol Set
 - Line States

Digital data must be encoded in some form for proper transmission. The type of encoding depends on the nature of the transmission medium, the data rate, and other factors such as noise, reliability, and cost. Intensity modulation is the normal method of representing digital data for transmission over fiber: a binary 1 is represented by a pulse of light and a binary 0 by absence of optical power. The disadvantage of using this method, in its simplest form, is its lack of synchronization. Long strings of ones or zeroes create a situation where the receiver is unable to synchronize its clock to

that of the transmitter. The solution is to first encode the binary data in such a way as to guarantee the presence of signal transitions, even if there are no transitions in the incoming digital signal; after this encoding is performed, the signal can be presented to the optical source for transmission using intensity modulation. A typical encoding scheme is Manchester encoding (see Figure 8).

Differential Manchester is only 50% efficient since each data bit is represented by transitions in signal. Two transitions allow a degree of robustness in the presence of noise, as would be the case in coaxial cable. Since fiber is less susceptible to noise, two transitions are not required to identify a bit with a good degree of confidence. To avoid having to use a 200MHz signal, FDDI specifies a code referred to as 4B/5B group encoding. The result is that the 100M bps throughput is achieved in FDDI with a 125MHz rate, rather than the 200MHz rate needed in differential Manchester. This helps keep down the cost and complexity of equipment. One drawback of the group encoding pertains to clock recovery. Since differential Manchester has more pulses in its stream, it is easy to extract the clock in that scheme. One of the key responsibilities of this FDDI sublayer is to decode the 4B/5B nonreturn to zero inverted (NRZI) signal from the network into symbols that can be recognized by the station, and vice versa.

The synchronization clock is derived from the incoming signal. The data is then retimed to an internal clock through an elasticity buffer. In this scheme, four bits of data are translated into a five-baud value transmitted over the network, giving an 80% efficiency factor. This group-encoding scheme employed in FDDI is a departure from differential Manchester codes normally specified in LAN standards.

To understand how the FDDI scheme achieves synchronization, one must realize that there are two stages of encoding. In 4B/5B, the encoding is performed four bits at a time. Each four bits of data are encoded into a symbol with five cells such that each cell contains a single signal element (presence or absence of light). In effect, each set of four bits is encoded as five bits. Then, each element of the 4B/5B stream is treated as a binary value and encoded using NRZI. In this code, a binary 1 is represented with a transition at the beginning of the bit interval; there are no other

transitions. The advantage of NRZI is that it employs differential encoding: the signal is decoded by comparing the polarity of adjacent signal elements rather than the absolute value of a signal element. This scheme is relatively robust in detecting transitions in the presence of noise or other distortions; therefore, the NRZI encoding will improve reception reliability. Table 1 shows the symbol encoding used in FDDI.

Since this scheme is encoding 4 bits (16 combinations) with 5 bit patterns (32 combinations), there will be patterns that are not needed. The codes selected to represent the sixteen four-bit patterns are such that a transition is present at least twice for each five-bit code. Given an NRZI format, no more than three zeroes in a row can be allowed, since the absence of a transition indicates a zero. The remaining symbols are either declared invalid or are assigned special meaning as control symbols as shown in Table 1.

PHY also provides line states for establishing the station's links with its neighbors (upstream and downstream) and to detect the integrity of the station's links to these neighbors. These line states are used to exchange a handshake with a neighbor. A node receiving a line state on its primary input can respond by sending the proper line state on the secondary output. The line states are composed of a repetition of one or more "I" symbols.

Another item that must be resolved in this sublayer is the issue of timing jitter. Jitter is the deviation of clock recovery that can occur when the receiver attempts to recover clocking as well as data from the received signal. The clock recovery will deviate in a random fashion from the transitions of the received signal. If no countermeasures are used, the jitter will accumulate around the ring. In LANs, the IEEE 802 standard specifies that only one master clock will be used on the ring and that the station with the clock will be responsible for eliminating jitter using an elastic buffer. If the ring as a whole runs ahead or behind the master clock, the elastic buffer expands or contracts accordingly. This centralized clocking method, however, is not practical for a 100M bps ring. At this speed, the bit time is only 10 ns, compared to a bit time of 250 ns. at 4M bps, making the effect of distortion more severe. Consequently, FDDI specifies the distributed clocking scheme.

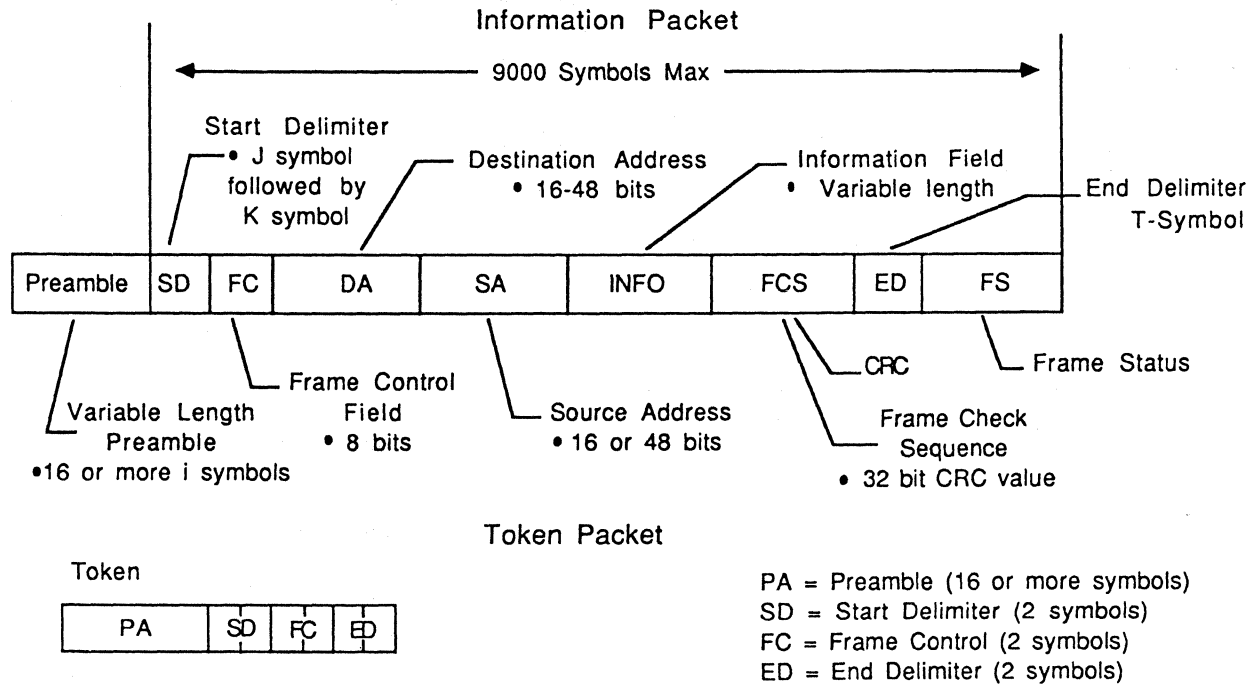
Table 1. 4B/5B Codes Used in FDDI

Function or 4-bit group (4B)	Group Code (5B)	Symbol
Starting Delimiter		
First symbol of sequential SD pair	11000	J
Second symbol of sequential SD pair	10001	K
Data Symbols		
0000	11110	0
0001	01001	1
0010	10100	2
0011	10101	3
0100	01010	4
0101	01011	5
0110	01110	6
0111	01111	7
1000	10010	8
1001	10011	9
1010	10110	A
1011	10111	B
1100	11010	C
1101	11011	D
1110	11100	E
1111	11101	F
Ending Delimiter		
Used to terminate datastream	01101	T
Control Indicators		
Logical Zero (reset)	00111	R
Logical One (set)	11001	S
Line Status Symbols		
Quiet	00000	Q
Idle	11111	I
Halt	00100	H
Invalid Code Assignment		
These patterns shall not be transmitted because they violate consecutive code-bit zeroes or duty cycle requirements. Some of the codes shown shall nonetheless be interpreted as a Halt if received.	00001	Void or Halt
	00010	Void or Halt
	00011	Void
	00101	Void
	00110	Void
	01000	Void or Halt
	01100	Void
	10000	Void or Halt

In this environment, each station uses its own autonomous clock to transmit or repeat information onto the ring. Each station has an elastic

buffer where data is clocked in at the clock rate recovered from the incoming stream, but it is clocked out at the station's own clock rate. This

Figure 9.
FDDI Packets



FDDI packet format.

distributed system is considered stronger than the centralized method and minimizes jitter. As a consequence of reclocking at each station, jitter does not limit the number of repeaters in the ring, as is the case in LANs where a master clock is used.

Media Access Control Specification

Layer 2 (Data Link Layer) of the OSI Reference Model is traditionally divided into two sublayers in a LAN context: Link Layer Control (LLC) and Media Access Control (MAC). FDDI only defines MAC, which controls data flow over the ring. The token-passing protocol incorporated in FDDI controls transmission over the network. MAC defines packet formation (headers, trailers, etc.), addressing, and cyclic redundancy checking (CRC). It also defines the recovery mechanisms. This standard defines the following specifications:

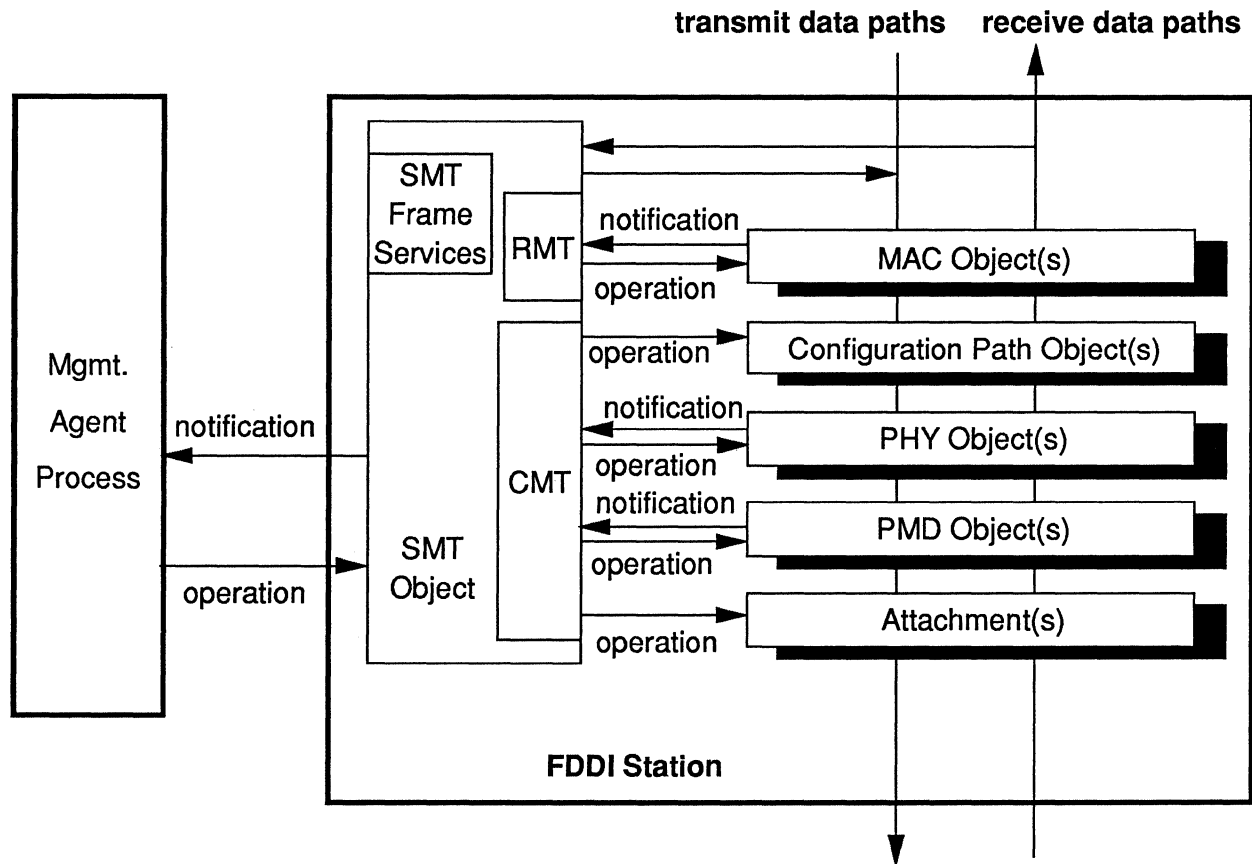
1. Services:
 - MAC to LLC Services
 - PHY to MAC Services
 - MAC to SMT Services
2. Facilities:
 - Symbol Set
 - Protocol Data Units
 - Fields

- Timers
- Frame Counts
- Frame Check Sum

The FDDI packet format is shown in Figure 9. Packets are preceded by a minimum of 16 IDLE control symbols. The packet itself is characterized by a Start Delimiter composed of the J and K control symbols. This is followed by a Frame Control field that identifies the type of packet. The Destination Address, which follows, identifies the frame recipient. The Source Address is also included to identify which station originated the packet. The address field can be 26 or 48 bits in length. The variable information field follows, along with a Frame Check Sequence field of 32 bits. The check sequence covers the Frame Control Field, the two addresses, and the information field. An End Delimiter, which consists of the T symbol, is transmitted. The maximum packet length is limited by the size of the elastic buffer in the Physical Sublayer and by the worst case frequency difference between two nodes, the upper bound in 9,000 octets. Figure 9 also shows the format of the token.

Flow control is the other major function of the MAC. In an idle condition, MAC connects to

Figure 10.
SMT Architectural Model




The relationships between managed objects in an FDDI station and the SMT object.

an internal source of IDLE control symbols to be transmitted over the ring. When a Start Delimiter is detected from the ring, MAC switches to a repeat path; the packet is monitored and copied if it is meant for this destination. The packet is simultaneously repeated onto the ring for relaying. The MAC can also inject its own packet or issue a token. Packets are removed only by the originating station. The MAC repeats the packet only until the Sender Address field is detected. If the destination recognizes the Sender Address field as its own station, it will insert IDLE control symbols back onto the ring (the fragmented packet is ignored and removed by any station holding a token for transmission). Stations that wish to transmit must first obtain a token (this is the unique six-symbol packet shown in Figure 9).

The procedures for obtaining the token and the amount of time allowed for data transmission (to retain fairness) are specified in the Timed Token Protocol (TTP). A station obtains the token by performing the stripping function on the incoming

token. Only the Start Delimiter field is repeated onto the ring; the station will inject its own information at this juncture. When the packet is sent, the station immediately issues a new token. TTP guarantees a maximum token rotation time. TTP allows two types of transmission: synchronous and asynchronous. In the synchronous mode, stations obtain a predefined amount of transmission bandwidth on each token rotation. The balance of the bandwidth is shared among stations using the asynchronous service. These stations can send data when the token arrives earlier than expected. Any unused capacity left over from synchronous capacity is available to asynchronous traffic, which may be subdivided into up to eight levels of priority. The amount of time allowed for asynchronous transmission is bounded by the difference of the token's actual arrival time and the expected arrival time. In essence, each station keeps track of how long it has been since it last saw the token. When it

Figure 11.
Media Interface Connector



Example of Media Interface Connector (MIC) plug.

next sees the token, it can send synchronous traffic and/or any asynchronous traffic for which time remains available.

Station Management (SMT)

The FDDI Station Management (SMT) specification describes software-based, low-level data link management and integrated network control functions of all stations attached to an FDDI LAN and of the LAN itself. Each FDDI station contains only one SMT entity. SMT initializes the network, monitors error rates and fault conditions in each network segment, and automatically reconfigures the network to isolate problem links. SMT components are Connection Management (CMT), which includes Entity Coordination Management, Physical Connection Management, (PCM), and Ring Management (RMT). SMT is intended to operate regardless of equipment type, vendor, protocols, or applications. Figure 10 presents the SMT architectural model.

SMT types (managed objects) have specific attributes indicating state, capabilities, and operation.

SMT managed objects are:

- Station or concentrator (SMT)
- MAC object(s)
- Path object(s)
- PHY object(s)
- PMD object(s)
- Attachment(s)

Attributes are:

- Attribute Identification (ID)

- Configuration
- Operational (Status, Counters, etc.)

Each attribute is defined in terms of Access Rights and whether it is Mandatory or Optional. Each attribute also carries an FDDI specification reference or a specific definition if it is not defined in the FDDI specification.

SMT Connection Management (CMT) operates at the logical level, controlling the interface between PHY and MAC entities in a given station and controlling SMT-to-SMT communications across the ring. When a session requires a connection to another station, CMT causes PHY to send a stream of symbols to the targeted station. Upon receiving the symbols, the receiving station's PHY returns a continuous stream of symbols (primitives) indicating line state, station status, and willingness to carry out the requested action (establish a link). QUIET symbols indicates a disabled link. HALT or alternating HALT and QUIET (MASTER) symbols indicate an operating link and the receiving station's status (master, slave, or peer). A stream of IDLE symbols indicates willingness to connect. Once the link is established, CMT configures the PHY and MAC.

Entity Coordination Management (ECM) controls the optional optical bypass switch and signals the Physical Connection Management (PCM) entity when the bypass is complete. ECM also performs the Path Test to determine a fault's location.

Physical Connection Management (PCM) initializes adjacent stations' PHYs and manages signaling. Maintenance support functions are also part of PCM.

Configuration Management (CFM) interconnects PHYs and MACs. It automatically configures these connections according to PCM flags. CFM is defined differently for stations and concentrators.

Ring Management (RMT) relays MAC and CFM status information. It detects stuck signaling beacons, initiates the trace function, detects duplicate addresses and resolves them to allow continued ring operation, and notifies SMT of MAC status.

The draft SMT standard was forwarded out of the X3T9.5 committee in April 1990. Letter ballots returned in midsummer 1990 included comments that required some additional work. The draft is now technically stable, and approval is expected by late 1991.

Higher Layers

From LLC upward, ANSI intends FDDI to generally fit traditional protocol stacks. The ANSI FDDI committee has not yet formally drafted protocols for Layers 3 (network) and 4 (transport), which are needed for any type of internetworking. Many suppliers believe the TCP/IP protocol suite can serve this purpose. TCP/IP software was originally developed by the U.S. government for Arpanet, the worldwide packet switched network, but has been used with great success in commercial applications—particularly for internetworking among different LANs. TCP and IP follow layered networking concepts, occupying Layers 4 and 3, respectively, of the OSI Model. Vendors have adopted the protocols for Ethernet and other local area networks. Most commercial implementations of the TCP/IP protocol suite include three standardized upper layer protocols: Telnet (virtual network terminal), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP).

Although more versatile than TCP/IP, the OSI internetworking protocols—spanning Layers 3 through 5—are less practical to implement in the real world. Much effort, however, has been expended to migrate to ISO-based standards, particularly under the MAP/TOP thrust, as shown in Figure 5. These market demands are forcing investigations into upgrading the third and fourth layers. The OSI stack also has a well-developed suite of upper layer applications, including X.400 (electronic mail), File Transfer Access Management (FTAM), and X.500 Directory System Protocol. Moreover, the federal government has mandated its own version of the OSI Reference Model—called GOSIP—must replace TCP/IP in government procurements after August 1990.

Some believe that middle-layer OSI protocols will be changed to look more like TCP/IP. In today's commercial networking applications, however, vendors are blending different protocol stacks from various sources to match user needs. One vendor's network protocol, for instance, might blend different layers from OSI, TCP/IP, or IBM's SNA. In reality, networking protocols are still evolving, and prospective users must evaluate them with an eye toward future standards.

Open Issues

In May 1991, five networking equipment and semiconductor manufacturers announced that they

had joined together to define and publish an open, interoperable solution for transmitting FDDI over shielded twisted-pair (STP) cabling. Advanced Micro Devices, Chipcom, Digital Equipment Corp., Motorola, and SynOptics published the proposed specification in hopes that additional companies would adhere to it—helping to lower the cost of the needed components—thus accelerating the installation of FDDI to the desktop. The vendors' rationale for developing FDDI-over-STP technology is to allow users to employ existing cabling when migrating from existing LANs to FDDI without installing fiber optic cabling.

The first ANSI-chartered Twisted Pair-Physical Medium Dependent (TP-PMD) ad hoc meeting was held in August 1990. Chipcom, Digital Equipment, and SynOptics shared the results of their independent research into implementing FDDI over STP. The companies' proposals were sufficiently similar that the companies agreed to work together to develop an open, interoperable proposal.

The proposed solution specifies a Physical Medium Dependent sublayer that uses 150-ohm, shielded twisted-pair cables. While the document does not specify a particular design, it does call for the complete replacement of the existing optical PMDs with electrical counterparts. Because all other aspects of the FDDI standard remain the same, no changes are required beyond the PMD level.

As mentioned earlier, Codenoll Technologies introduced a nonstandard version of the FDDI interface board for Extended Industry Standard Architecture (EISA) computers. While using an 850-nm. LED in the Codenoll design does reduce the cost of FDDI connectivity, users will be unable to operate with standard-based equipment containing 1,300-nm. transmitters and receptors.

FDDI-II and Other Enhancement Efforts

Work is nearly complete on defining a scheme where available network bandwidth is divided between voice and data using time-division multiplexing. FDDI-II, a superset of FDDI, is being defined as an upward-compatible, fiber-based LAN incorporating the current data capabilities in addition to the capability to handle voice and T1-compressed video traffic. The FDDI-II proposal will follow the original FDDI standard, adding a

fifth document to the present standard. The new document, Hybrid Ring Control, describes a hybrid operating mode comprising the packet-switching scheme used in FDDI and an isochronous transport mode similar to that used in public switched networks. Adding the isochronous mode will enable networks based on the expanded standard to transfer pixel data—key to video and computer graphics transport—in addition to circuit switched voice signals. Because it builds on the existing FDDI framework, FDDI II should be completed and approved very quickly. According to Gene Milligan, chairman of the ANSI X3T9.5 committee, the Hybrid Ring Control document should be an approved standard by the end of 1991.

Some people are already thinking of even higher rates, particularly when considering voice and data. The FDDI follow-on, or FFOL, is under development at ANSI. The FFOL project proposal covers the capability to operate as a backbone for multiple FDDI networks; interconnection to wide area networks, including Broadband Integrated Services Digital Networks (BISDN); a data rate between 600M bps and 1.25G bps initially, with intermediate data rates matched to the Synchronous Digital Hierarchy (SDH) underlying SONET and eventual support for data rates up to 2.4G bps; duplex links; support for existing FDDI cabling, where possible; and support for both single-mode and multimode fiber. If the ANSI effort keeps on schedule, the family of standards comprising FFOL should be complete by late 1995.

There is also a separate effort called High Speed Channel being undertaken in X3T9.3. This is a point-to-point system for digital data interface (channel extension rather than a network configuration intrinsic in FDDI). These systems are designed to carry 400M-, 800M-, or 1600M bps on very short copper links, but they are being developed to be compatible with future fiber optic links operating over longer distances.

FDDI Information Sources

The chairperson of the ASC X3T9.5 task group is Gene Milligan of Seagate (Oklahoma City, OK).

ANSI is involved in the form of the X3 Secretariat, as is CBEMA (Computer Business Equipment Manufacturers Association).

Specification X3.139, X3.148, and X3.166 can be obtained from Global Engineering Documents.

Draft Specification X3T9.5/84-49 is available from ANSI for review.

ANSI
1430 Broadway
New York, NY, 10018 (212) 642-4900.

CBEMA
311 First Street, Suite 500
Washington D.C. (202) 737-8888.

Global Engineering Documents
2805 McGraw Avenue, P.O. Box 19539
Irvine, CA 92714 (714) 261-1455

Information regarding the FDDI/STP effort is available from:

Advanced Micro Devices, Inc.
901 Thompson Place
Sunnyvale, CA 94088 (408) 982-7880

Chipcom Corp.
118 Turnpike Road
Southborough, MA 01772 (508) 460-8900

Digital Equipment Corp. (DEC)
550 King Street
Littleton, MA 01460 (508) 486-5096

Motorola
6501 William Cannon Drive West
Austin, TX 78735 (512) 891-2140

SynOptics Communications
4401 Great America Parkway
Santa Clara, CA 95052 (408) 764-1013 ■