

Fiber Distributed Data Interface

Background

The Fiber Distributed Data Interface (FDDI) standard was produced by the ANSI X3T9.5 standards committee in the mid-1980s. During this period, high-speed engineering workstations were beginning to tax the capabilities of existing local-area networks (LANs) (primarily Ethernet and Token Ring). A new LAN was needed that could easily support these workstations and their new distributed applications. At the same time, network reliability was becoming an increasingly important issue as system managers began to migrate mission-critical applications from large computers to networks. FDDI was developed to fill these needs.

After completing the FDDI specification, ANSI submitted FDDI to the International Organization for Standardization (ISO). ISO has created an international version of FDDI that is completely compatible with the ANSI standard version.

Today, although FDDI implementations are not as common as Ethernet or Token Ring, FDDI has gained a substantial following that continues to increase as the cost of FDDI interfaces diminishes. FDDI is frequently used as a backbone technology as well as a means to connect high-speed computers in a local area.

Technology Basics

FDDI specifies a 100-Mbps, token-passing, dual-ring LAN using a fiber-optic transmission medium. It defines the physical layer and media-access portion of the link layer, and so is roughly analogous to IEEE 802.3 and IEEE 802.5 in its relationship to the Open System Interconnection (OSI) reference model.

Although it operates at faster speeds, FDDI is similar in many ways to Token Ring. The two networks share many features, including topology (ring), media-access technique (token passing), reliability features (redundant rings, for example), and others. For more information on Token Ring and related technologies, refer to Chapter 6, “Token Ring/IEEE 802.5.”

One of the most important characteristics of FDDI is its use of optical fiber as a transmission medium. Optical fiber offers several advantages over traditional copper wiring, including security (fiber does not emit electrical signals that can be tapped), reliability (fiber is immune to electrical interference), and speed (optical fiber has much higher throughput potential than copper cable).

FDDI defines use of two types of fiber: *single mode* (sometimes called *monomode*) and *multimode*. Modes can be thought of as bundles of light rays entering the fiber at a particular angle. Single-mode fiber allows only one mode of light to propagate through the fiber, while multimode fiber allows multiple modes of light to propagate through the fiber. Because multiple modes of light propagating through the fiber may travel different distances (depending on the entry angles), causing them to arrive at the destination at different times (a phenomenon called *modal dispersion*), single-mode

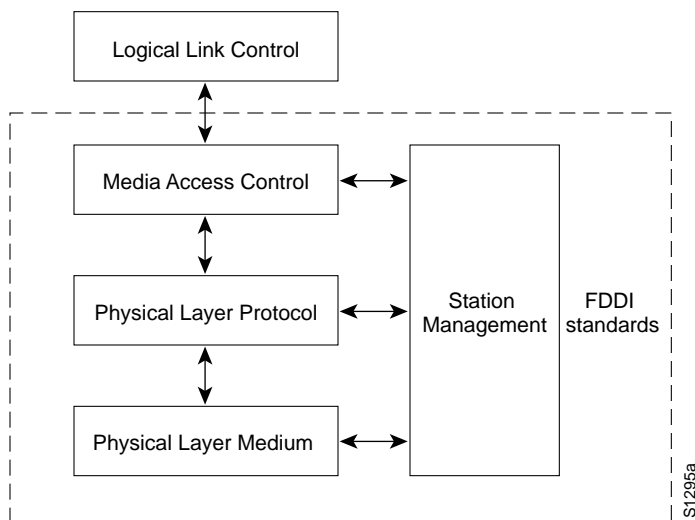
fiber is capable of higher bandwidth and greater cable run distances than multimode fiber. Due to these characteristics, single-mode fiber is often used for interbuilding connectivity, while multimode fiber is often used for intrabuilding connectivity. Multimode fiber uses light-emitting diodes (LEDs) as the light-generating devices, while single-mode fiber generally uses lasers.

FDDI Specifications

FDDI is defined by four separate specifications (see Figure 7-1):

- *Media Access Control (MAC)*—Defines how the medium is accessed, including frame format, token handling, addressing, algorithm for calculating a cyclic redundancy check value, and error recovery mechanisms.
- *Physical Layer Protocol (PHY)*—Defines data encoding/decoding procedures, clocking requirements, framing, and other functions.
- *Physical Layer Medium (PMD)*—Defines the characteristics of the transmission medium, including the fiber-optic link, power levels, bit error rates, optical components, and connectors.
- *Station Management (SMT)*—Defines the FDDI station configuration, ring configuration, and ring control features, including station insertion and removal, initialization, fault isolation and recovery, scheduling, and collection of statistics.

Figure 7-1 FDDI Standards



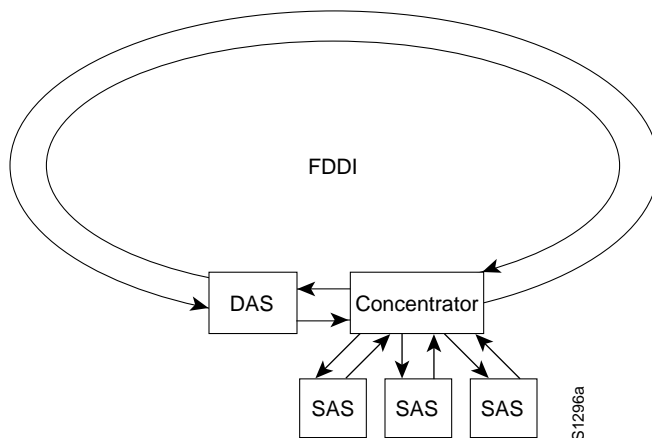
Physical Connections

FDDI specifies the use of dual rings. Traffic on these rings travels in opposite directions. Physically, the rings consist of two or more point-to-point connections between adjacent stations. One of the two FDDI rings is called the *primary* ring; the other is called the *secondary* ring. The primary ring is used for data transmission, while the secondary ring is generally used as a backup.

Class B or *single-attachment stations (SAS)* attach to one ring; *Class A* or *dual-attachment stations (DAS)* attach to both rings. SASs are attached to the primary ring through a *concentrator*, which provides connections for multiple SASs. The concentrator ensures that failure or power down of any given SAS does not interrupt the ring. This is particularly useful when PCs, or similar devices that frequently power on and off, connect to the ring.


A typical FDDI configuration with both DASs and SASs is shown in Figure 7-2.

Figure 7-2 FDDI Nodes: DAS, SAS, and Concentrator



Each FDDI DAS has two ports, designated A and B. These ports connect the station to the dual FDDI ring. Therefore, each port provides a connection for both the primary and the secondary ring, as shown in Figure 7-3.

Figure 7-3 FDDI DAS Ports



Traffic Types

FDDI supports real-time allocation of network bandwidth, making it ideal for a variety of different application types. FDDI provides this support by defining two types of traffic: *synchronous* and *asynchronous*. Synchronous traffic can consume a portion of the 100-Mbps total bandwidth of an FDDI network, while asynchronous traffic can consume the rest. Synchronous bandwidth is allocated to those stations requiring continuous transmission capability. Such capability is useful for transmitting voice and video information, for example. Other stations use the remaining bandwidth asynchronously. The FDDI SMT specification defines a distributed bidding scheme to allocate FDDI bandwidth.

Asynchronous bandwidth is allocated using an eight-level priority scheme. Each station is assigned an asynchronous priority level. FDDI also permits extended dialogues, where stations may temporarily use all asynchronous bandwidth. The FDDI priority mechanism can essentially lock out stations that cannot use synchronous bandwidth and have too low an asynchronous priority.

Fault-Tolerant Features

FDDI provides a number of fault-tolerant features. The primary fault-tolerant feature is the dual ring. If a station on the dual ring fails or is powered down or if the cable is damaged, the dual ring is automatically “wrapped” (doubled back onto itself) into a single ring, as shown in Figure 7-4. In this figure, when Station 3 fails, the dual ring is automatically wrapped in Stations 2 and 4, forming a single ring. Although Station 3 is no longer on the ring, network operation continues for the remaining stations.

Figure 7-4 Station Failure, Ring Recovery Configuration

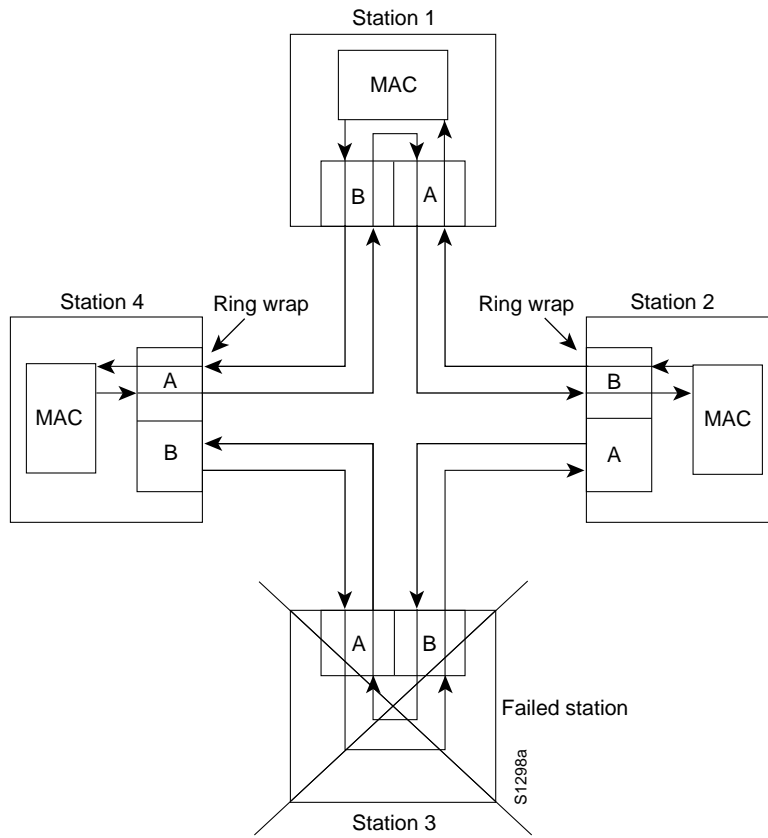
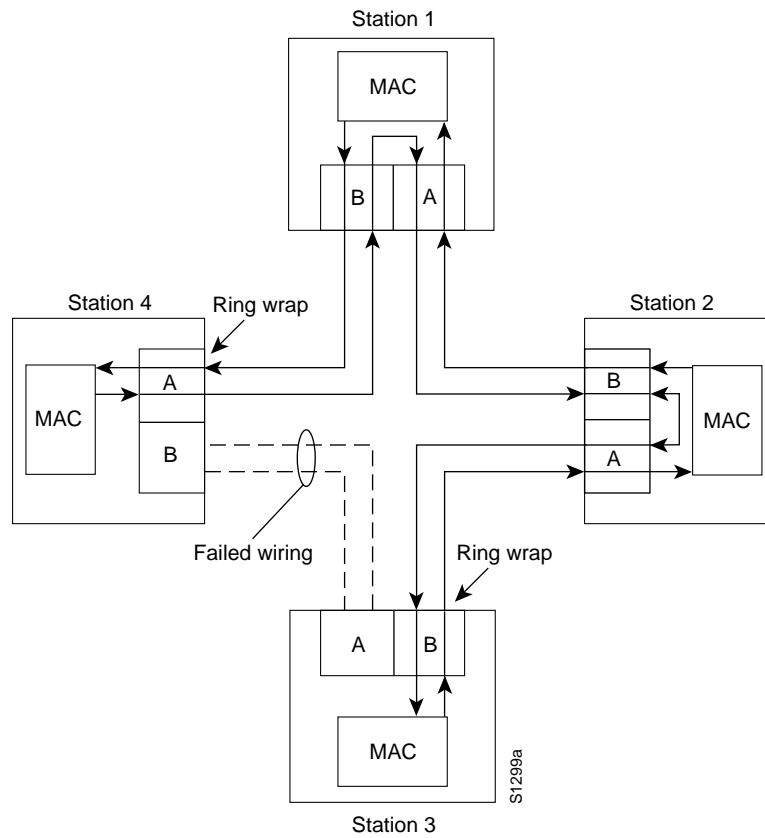


Figure 7-5 shows how FDDI compensates for a wiring failure. Stations 3 and 4 wrap the ring within themselves when wiring between them fails.

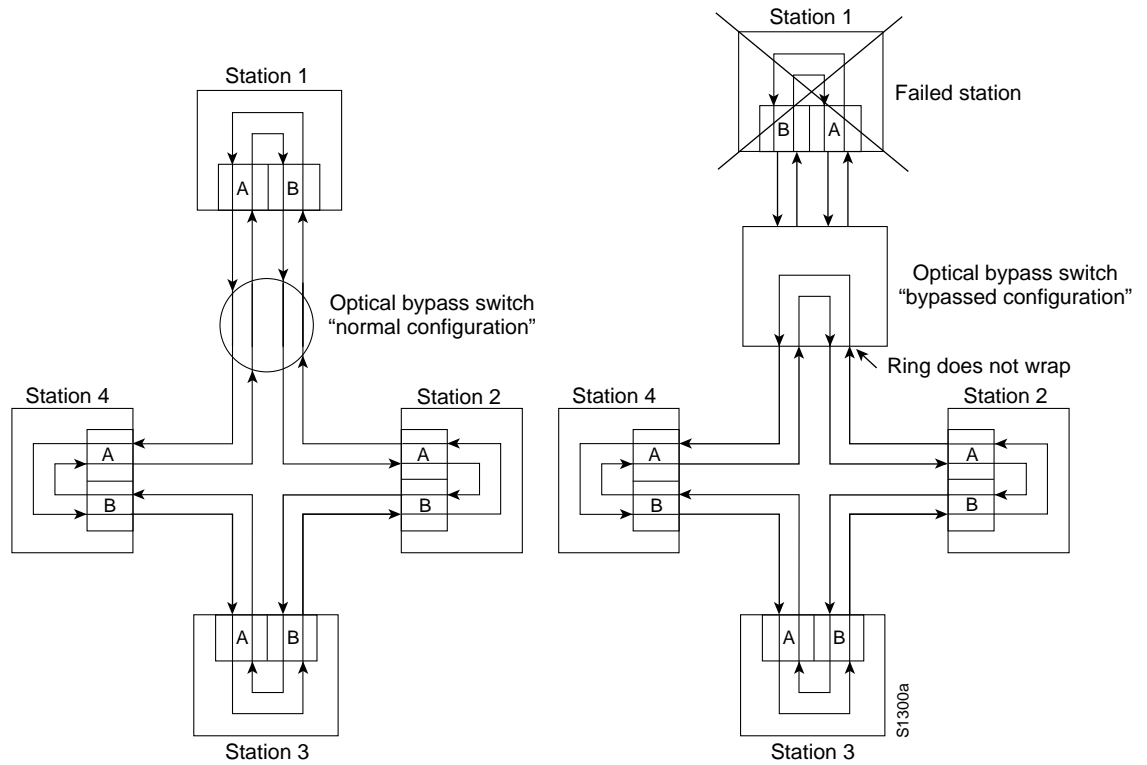
Figure 7-5 Failed Wiring, Ring Recovery Configuration



As FDDI networks grow, the possibility of multiple ring failures grows. When two ring failures occur, the ring will be wrapped in both cases, effectively segmenting the ring into two separate rings that cannot communicate with each other. Subsequent failures cause additional ring segmentation.

Optical bypass switches can be used to prevent ring segmentation by eliminating failed stations from the ring. This is shown in Figure 7-6.

Figure 7-6 Use of Optical Bypass Switch

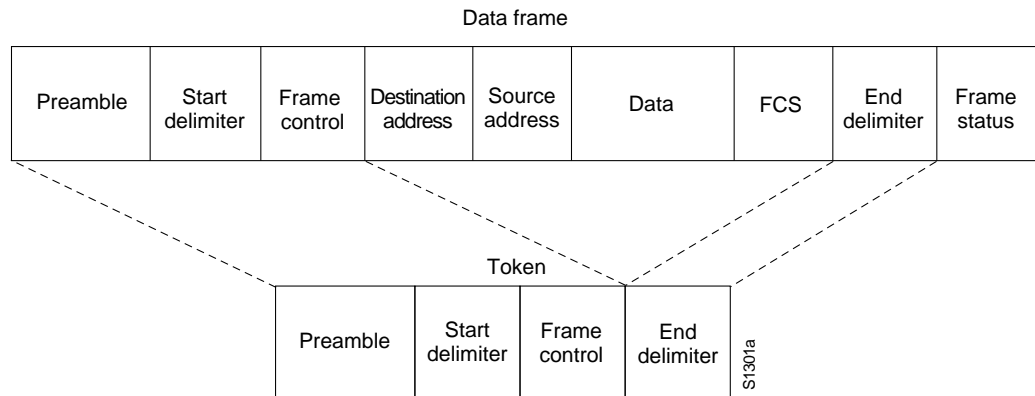


Critical devices such as routers or mainframe hosts can use another fault-tolerant technique called *dual homing* to provide additional redundancy and help guarantee operation. In dual-homing situations, the critical device is attached to two concentrators. One pair of concentrator links is declared the active link; the other pair is declared passive. The passive link stays in backup mode until the primary link (or the concentrator to which it is attached) is determined to have failed. When this occurs, the passive link is automatically activated.

Frame Format

FDDI frame formats (shown in Figure 7-7) are similar to those of Token Ring.

Figure 7-7 FDDI Frame Format



The fields of an FDDI frame are as follows:

- *Preamble*—Prepares each station for the upcoming frame.
- *Start delimiter*—Indicates the beginning of the frame. It consists of signaling patterns that differentiate it from the rest of the frame.
- *Frame control*—Indicates the size of the address fields, whether the frame contains asynchronous or synchronous data, and other control information.
- *Destination address*—Contains a unicast (singular), multicast (group), or broadcast (every station) address. As with Ethernet and Token Ring, FDDI destination addresses are 6 bytes.
- *Source address*—Identifies the single station that sent the frame. As with Ethernet and Token Ring, FDDI source addresses are 6bytes.
- *Data*—Contains either information destined for an upper-layer protocol or control information.
- *Frame check sequence (FCS)*—Filled by the source station with a calculated *cyclic redundancy check (CRC)* value dependent on the frame contents (as with Token Ring and Ethernet). The destination station recalculates the value to determine whether the frame may have been damaged in transit. If so, the frame is discarded.
- *End delimiter*—Contains nondata symbols that indicate the end of the frame.
- *Frame status*—Allows the source station to determine if an error occurred and if the frame was recognized and copied by a receiving station.

CDDI

The high cost of fiber-optic cable has been a major impediment to the widespread deployment of FDDI to desktop computers. At the same time, shielded twisted-pair (STP) and unshielded twisted-pair (UTP) copper wire is relatively inexpensive and has been widely deployed. The implementation of FDDI over copper wire is known as *Copper Distributed Data Interface (CDDI)*.

Before FDDI could be implemented over copper wire, a problem had to be solved. When signals strong enough to be reliably interpreted as data are transmitted over twisted-pair wire, the wire radiates electromagnetic interference (EMI). Any attempt to implement FDDI over twisted-pair wire had to ensure that the resulting energy radiation did not exceed the specifications set in the United States by the Federal Communications Commission (FCC) and in Europe by the European Economic Council (EEC). Three technologies reduce energy radiation:

- Scrambling—When no data is being sent, FDDI transmits an idle pattern that consists of a string of binary ones. When this signal is sent over twisted-pair wire, the EMI is concentrated at the fundamental frequency spectrum of the idle pattern, resulting in a peak in the frequency spectrum of the radiated interference. By scrambling FDDI data with a pseudo-random sequence prior to transmission, repetitive patterns are eliminated. The elimination of repetitive patterns results in a spectral peak that is distributed more evenly over the spectrum of the transmitted signal.
- Encoding—Signal strength is stronger, and EMI is lower when transmission occurs over twisted-pair wire at lower frequencies. MLT3 is an encoding scheme that reduces the frequency of the transmitted signal. MLT3 switches between three output voltage levels so that peak power is shifted to less than 20 MHz.
- Equalization—Equalization boosts the higher frequency signals for transmission over UTP. Equalization can be done on the transmitter (predistortion), or at the receiver (postcompensation), or both. One advantage of equalization at the receiver is the ability to adjust compensation as a function of cable length.

In June 1990, ANSI established a subgroup called the Twisted Pair-Physical Medium Dependent (TP-PMD) working group to develop a specification for implementing FDDI protocols over twisted-pair wire. ANSI approved the TP-PMD standard in February 1994. Approval of the standard is pending in Europe.

Note Of the many categories and types of twisted-pair wire, the ANSI standard only recognizes Category 5 UTP and Type 1 STP.
