

Fiber Distributed Data Interface (FDDI)

SOME OF THE MAIN TOPICS IN THIS CHAPTER ARE

FDDI Dual-Ring Topology	2
FDDI Protocol Standards	6
Transmitting Data on an FDDI Ring	7
Common Problems Using FDDI	10

FDDI stands for *Fiber Distributed Data Interface*. Work was started on FDDI in the early 1980s to provide a reliable, high-speed method of networking to connect the faster workstations and computers that were becoming available. Although Ethernet and Token-Ring networks were operating at that time at maximum speeds of 10Mbps, FDDI was designed to provide a 100Mbps bandwidth, which was a substantial increase over existing technologies at that time. Because of the more recent development of other high-speed networking technologies (such as Fast Ethernet and Gigabit Ethernet) and protocols that provide better control and quality of service (such as Frame Relay and ATM), and because of the usually higher cost, basic FDDI is not as popular as it once was.

However, when used for the backbone in a network, FDDI is still considered to be a viable option for the network administrator mainly because of the built-in redundancy provided by two rings—if one ring failed, the other could maintain the network. The technology is a mature one, and you can still find vendors and installers for FDDI. When a new building is being added to an existing network that already contains FDDI segments, it might be more cost-effective to use FDDI than to incorporate newer technologies. However, because most FDDI installations use fiber-optic cabling, you'll be better off if you can use a newer technology, such as Gigabit Ethernet, which can use the same cabling. Gigabit Ethernet might also be a better choice if you are planning a complete overhaul of your network or are in the process of building a new one from scratch.

Note

To compare FDDI with Fast and Gigabit Ethernet, see Chapter 14, "Ethernet: The Universal Standard." If you need a more deterministic local area network technology, see "Token-Ring Networks" located on the upgradingandrepairingpcs.com Web site.

An important feature that made FDDI an attractive networking solution is the security features inherent in using optical fiber as a network medium. Unlike topologies based on copper wire cables, optical fiber has no electrical properties and emits no signals that can be detected outside the cable. However, there are many alternatives today that also use fiber-optic cabling, which reduces this advantage considerably. For example, Ethernet has been adapted to fiber-optic cables. Fibre Channel, used mainly for Storage Area Networks (discussed in Chapter 11, "Network Attached Storage and Storage Area Networks"), also uses fiber-optic cabling for the most part. The advantage that FDDI once trumpeted is now available with other more common network standards.

In this chapter, we will look at the components that make up the FDDI protocol and the methods used to transmit data and perform maintenance functions on an FDDI ring, and examine some typical configurations. If you are still using FDDI in your network, and if it works for you, then that's okay. If you want to increase the speed of your network, consider an upgrade to newer technologies discussed in other chapters of this book. With the same fiber-optic cables, you can achieve a higher bandwidth and keep communications secure by using the latest encryption technologies.

For those users who are still using FDDI, the remainder of this chapter will acquaint you with the techniques used and assist you in managing your infrastructure.

FDDI Dual-Ring Topology

Two of the main goals for FDDI during its development were speed and reliability. Optical fiber was selected as the network transmission medium because of its capability to transmit data at high speeds. The topology chosen was the ring topology, similar to Token-Ring networks. However, to provide enhanced reliability, a dual-ring topology was developed that uses two rings that transmit data in opposite directions (counter-rotating rings). Figure 1 shows the layout of a simple FDDI dual ring.

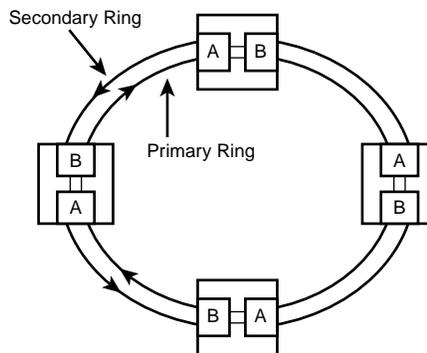


Figure 1 FDDI uses the dual counter-rotating ring topology.

By using a dual-ring topology, with one ring operating as the primary ring and the other as a secondary backup ring, a simple failure of one fiber cable segment is less likely to cause disruption of the entire network. You can see that each station on the ring in Figure 1 has two ports, labeled A and B, which connect it to the ring. At first glance it might appear that one port is used for each ring. However, that is not the case. Instead, on the primary ring the A port receives data from its neighbor, which is transmitting the signal on its B port. On the backup secondary ring, the A port transmits data to its neighbor in the opposite direction, and the neighbor receives the data on its B port.

Ports and Stations

The first example of a simple FDDI ring that we looked at consists of stations that have two ports, labeled A and B. This kind of station is called a *dual-attached station*. This topology is called a *dual-attached ring*. Other ports that you might find on an FDDI network include the following:

- M or *master* port
- S or *slave* port

These ports enable you to connect other kinds of stations to the FDDI ring, specifically a concentrator or a *single-attached station (SAS)*. Single-attached stations do not have an A or B port, but instead have a slave port. In this case, the SAS is not directly attached to the dual ring but is connected through a master port on a concentrator. In Figure 2, you can see an example of a concentrator that consists of A, B, and master ports. The master ports are used to connect multiple workstations to the ring through their slave ports.

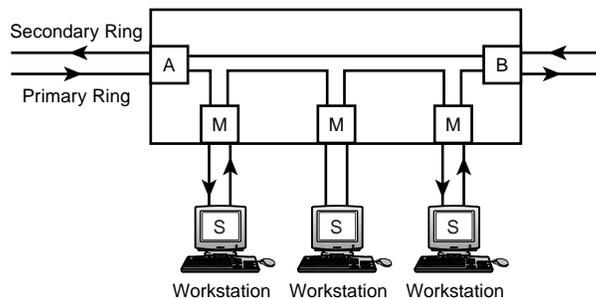


Figure 2 Concentrators can be used to connect single-attached stations to an FDDI ring.

Using this concentrator-based ring topology can aid in reducing network outages due to problems with workstations attached to the ring. The concentrator can act in a manner similar to the Ethernet hub or Token-Ring MSAU and isolate individual workstations that fail, preventing them from causing problems for other nodes in the ring. For dual-attached stations, however, even a simple power-down of the unit can cause problems because the node, which is powered off, no longer is capable of receiving or retransmitting data, causing a break in the ring.

Ring Wrap

The *dual-ring* topology allows FDDI to heal itself when faced with simple failures of one node. When a station fails or if a cable between two stations is damaged, other stations detect the failure and the ring is automatically wrapped back onto itself to form a single ring. Figure 3 shows four network nodes (dual-attached stations) operating normally in a dual-ring topology.

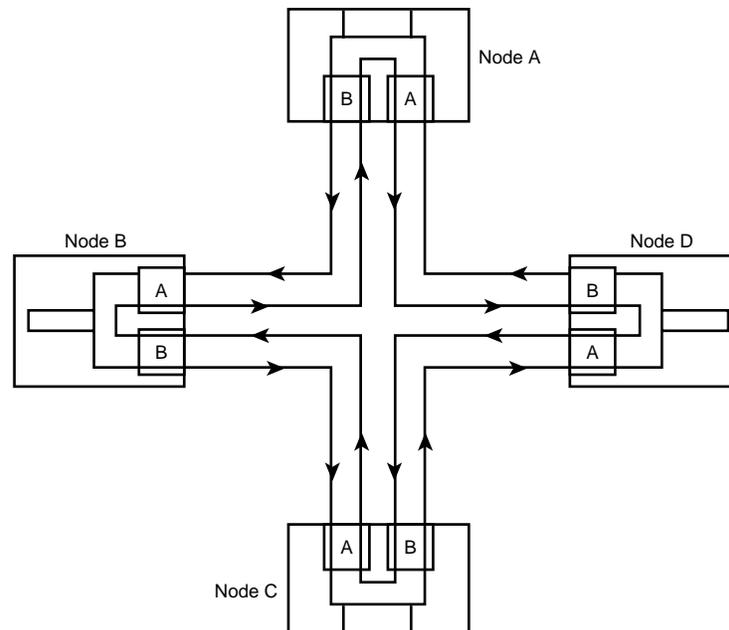


Figure 3 During normal operation, FDDI uses a dual-ring topology with one ring serving as the primary network medium.

In Figure 4, you can see what happens when a single node on the ring fails. Here Node C has failed and no longer participates in ring communications. Instead, Nodes B and D have internally wrapped the ring back onto itself. That is, the cable that was used for the secondary ring has been joined in these nodes to the primary ring, creating a single ring consisting of Nodes A, B, and D.

This self-repair method works well when only one node on the ring has failed. If more than one node fails, ring wrapping can result in two or more separate rings. For this reason, it is preferable to use concentrators to join individual workstations to a ring rather than make each workstation or PC a dual-attached station. The concentrator itself becomes a single point of failure, and if it fails it can take down all workstations connected to it. However, because even something as simple as powering down a workstation that is dual-attached can cause a break in the ring, it is generally more likely that a workstation or PC will be the point of failure rather than a concentrator that is safely tucked away in a wiring closet.

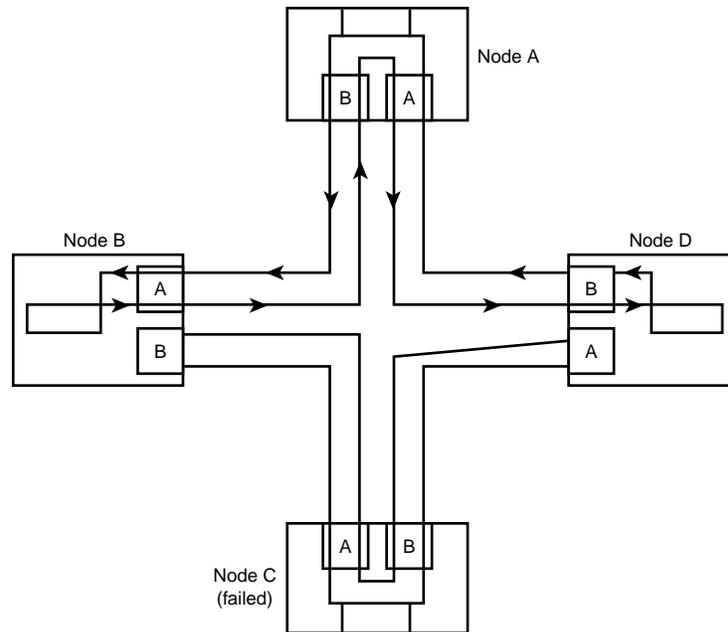


Figure 4 When a node on the ring fails, the Media Access Control layer forms a single ring to isolate the malfunctioning node so that the other nodes can continue to operate.

Another consideration to think about is that, because ring wrapping increases the actual distance it takes for information to travel around the ring, you must plan the maximum size of the ring accordingly. In a large LAN it would be impractical to size the ring to be able to accommodate a large number of failures. However, for a ring that is rated as a highly available network, you should plan to limit the maximum size to account for at least two to four simultaneous failures.

Optical Bypass Switches

Another device you can use to help prevent ring failures is called an *optical bypass switch*. This device can prevent ring wrap from happening if the station that is attached to it fails. This is accomplished in the bypass switch by an optical relay that causes the signal to bypass the failed station and continue on to its neighbor in the ring. However, these kinds of switches are expensive and result in a small amount of signal loss. Using them for each station on a large ring can cause signal degradation that can make network performance suffer noticeably. You might balance the cost of this kind of switch against the amount of allowable downtime you can have in your network. For networks that can tolerate very little downtime (such as a financial institution), the cost of the switch will be minimal compared to the cost of the downtime necessary to repair the ring.

Dual-Homing Important Devices

Some network nodes usually are considered to be more important than others. For example, a user might think that his or her workstation is the most important device on the network, but the failure of a router or file server that also is attached to the local ring can cause more than one user to be affected. Important devices can be made less susceptible to network failure by a technique called *dual-homing*. In this situation, an important device is attached to two separate concentrators that are then attached to the ring. Only one of the links is active, and the other acts as a standby. If the active connection made to the first concentrator becomes unavailable, the device can then use the secondary connection to continue operations.

FDDI Protocol Standards

FDDI is defined by standards from the American National Standards Institute (ANSI) and the International Organization for Standardization (ISO). The four key components of FDDI are as listed here:

- Media Access Control (MAC) layer
- Physical (PHY) layer
- Physical Media Dependent (PMD) layer
- Station Management (SMT) protocol

The MAC layer interfaces with higher-level protocols, such as TCP/IP, and passes their Protocol Data Units (PDUs) to the PHY layer after repackaging them into packets of as many as 4,500 bytes. Other functions performed by the MAC layer are addressing, scheduling, and routing. The MAC specification defines the frame format and takes care of error recovery. It also is responsible for calculating the CRC value and token-handling procedures.

The PHY layer is responsible for actually encoding and decoding the packet data into the format used by the transmission media. The method used for encoding is called *4B/5B* encoding, which means that 4 bits of information are encoded into 5-bit groups. This encoding technique is used to ensure that under normal circumstances the bit stream that is transmitted by FDDI will never contain 4 zero bits in a row, which is important from a timing standpoint. FDDI has no active monitor like a Token-Ring network, and each station has its own clock and must be able to synchronize with other stations.

Although we normally think of 8-bit values (a byte) when encoding data, the 4B/5B technique concentrates on units of only 4 bits in length. The 4B/5B technique uses data symbols to represent the actual data that is being transmitted (labeled 0–9 and A–F), and eight control symbols used for things such as indicating the state of the link (symbols I, H, and Q) or frame delimiters (symbols J, K, and T).

Table 1 shows the symbols used for data, along with the actual binary value each symbol represents. The last column shows the actual bits transmitted for this symbol on the FDDI ring.

Table 1 4B/5B Encoding Values for Data Symbols

Actual Binary Value	Symbol Name	Symbol Bits
0000	0	11110
0001	1	01001
0010	2	10100
0011	3	10101
0100	4	01010
0101	5	01011
0110	6	01110
0111	7	01111
1000	8	10010
1001	9	10011
1010	A	10110
1011	B	10111

Table 1 Continued

Actual Binary Value	Symbol Name	Symbol Bits
1100	C	11010
1101	D	11011
1110	E	11100
1111	F	11101

As you can see, a binary value of 0000 would be transmitted as 11110, and the binary value of 1111 would be transmitted as 11101. In the actual bits transmitted, no symbol contains more than two zeros in a row, so when combined with other symbols, the bit stream will not contain more than four zeros in a row.

The PMD layer does the actual physical signaling on the transmission media, which is fiber-optic cable for FDDI. To make it easier to attach common devices such as PCs and workstations to an FDDI network, the CDDI (Copper Distributed Data Interface) specification allows the PMD layer to also transmit data using copper wires as the network medium.

The SMT protocol is the component responsible for managing the ring. Similar to Token-Ring networks, SMT functions include neighbor identification, detection of faults, and reconfiguration of the ring due to faults or insertion or removal of a station from the ring.

FDDI can extend for up to 100km when using multimode fiber, with stations being up to 2km from each other. When using single-mode fiber, stations can be up to 20km in distance from each other. The maximum number of stations on the ring is 500.

Transmitting Data on an FDDI Ring

The method used by FDDI to transmit information across optical fiber is light. Two kinds of fiber-optic cables can be used; they are classified as either *Single mode* or *Multimode*. Single-mode fiber uses a laser as its source of light and can be used over longer distances than Multimode fiber. Multimode fiber cables allow multiple rays of light, entering the cable at different angles, to carry signals through the cable and use a light-emitting diode (LED) as their light source.

Caution

Remember that when you're dealing with lasers, don't look directly at the light emitted from the end of a cable! Doing this can cause permanent damage to your eyesight!

Using Light to Encode Bits

A station on the ring looks at the state of the light beam on the fiber about every eight nanoseconds. The only possibilities are that the station will sense that light is present or sense that it is not. To determine whether a zero or a one is being sent, the station will compare the current sampling with the one immediately before it. If the state of the fiber has changed—that is, it has gone from no light detected to light detected, or vice versa—it is determined that a bit representing a one has been transmitted. If no change between this sampling period and the one immediately preceding it is detected, a zero bit has been transmitted. This technique for signaling is known as *nonreturn to zero with inversion (NRZI)* modulation.

FDDI Frames

FDDI frames are similar to Token-Ring frames. A token frame (shown in Figure 5) is passed from one station in the ring to the next. When a station on the ring has data that it needs to transmit, it seizes the token, adds addressing information and data to the frame, and then transmits the data frame (see Figure 6).

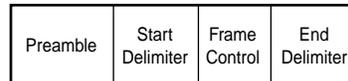


Figure 5 The token frame is used to mediate access to the ring.

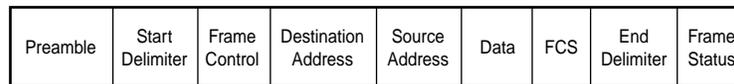


Figure 6 The data frame contains the same fields as the token frame, but addressing information and data also are included.

The fields that appear in Figure 5 and Figure 6 are used for the following functions:

- **Preamble**—Indicates that a frame is arriving at the station so that the adapter can prepare to receive it.
- **Start Delimiter**—Designates the start of the frame.
- **Frame Control**—Specifies the size of the address fields and other control information.
- **Destination Address**—Indicates the 6-byte address of the destination of the frame. This also can be a multicast address such as a broadcast or group address.
- **Source Address**—Contains the 6-byte address of the station that originated the data frame.
- **Data**—Contains control information or data that is destined for a higher-level protocol, such as TCP/IP.
- **Frame Check Sequence (FCS)**—Contains a 32-bit CRC value used to check the integrity of the frame when it arrives at its destination.
- **End Delimiter**—Designates the end of the frame.
- **Frame Status**—Contains bits used to indicate errors, address recognized, and frame copied.

The token frame is passed around the ring in an orderly fashion. When a node on the ring wants to transmit data, it grabs the token. That is, instead of retransmitting the token frame to its neighbor, it transmits one or more data frames. Data frames can be used to send information that is destined for a higher-level protocol, or it can control information used in ring maintenance procedures.

Some of the more important FDDI management frames are as listed here:

- **Neighbor Information Frame (NIF)**—This frame is sent around the ring periodically (usually from 2 to 30 seconds) and is used to exchange information about stations and their neighbors on the ring.
- **Status Information Frame (SIF)**—This frame gives stations configuration information and operation information stored in a *MIB (management information base)*.

- **Parameter Management Frame**—This frame is used to read information from, or write information to, the MIB.
- **Status Report Frame**—This frame uses a multicast address and is used to give status information to management applications.

The Target Token Rotation Timer and the Claim Process

When the ring is first initialized, the Target Token Rotation Time (TTRT) is determined. After this is done, the nodes participating in the ring must decide which one will generate the first token frame. The selection process is done using the following steps:

1. The Token Rotation Timer (TRT) is initially set to the TTRT value.
2. If the TRT expires before a station receives a token, it assumes that some kind of error has occurred and it starts the claim process.
3. During this process, the stations also negotiate to determine which one will be responsible for generating the first token on the ring. Each station is assigned a value for TTRT. During the claim process, each station compares the value for TTRT that was set by the previous station in the ring.
4. If the TTRT it receives from its neighbor is longer than its own value, it substitutes its own value and retransmits the claim frame.
5. If the received TTRT value is shorter than its own, it retransmits the frame. If its value is the same as that which it receives from its neighbor, the station with the highest address wins.
6. When a station receives the claim frame back and recognizes that it is the winner, it sets the TTRT for the ring and generates the first token.

After normal operations begin, each node on the ring monitors the ring and frames passing around to check for errors.

When a station transmits a data frame, the frame travels around the ring and is read by each node and retransmitted to the next node on the ring. When a node recognizes itself as the destination address of the frame, it copies the address to a buffer and sets the address recognized and frame copied bits in the Frame Status field before retransmitting it. When the frame arrives back at the station from which it was originally generated, the station can determine whether the frame was received by checking these bits. If the error bit is set, it can take actions to resend the data.

When a station begins to transmit frames on the ring, it can continue to do so for as long as the rules negotiated with other stations allow. This time monitored at each station by a timer is called the *Token Holding Timer*.

Beaconing

When a station on the ring fails to receive a frame (either token or data) from its neighbor, it begins the beaconing process by transmitting a beacon frame. Each station in the ring retransmits the beacon frame to the next station in the ring. If the beacon frame travels around the ring back to the station where it originated, the station stops transmitting beacon frames because it assumes that the fault has been repaired.

If it does not receive the beacon frame back after about 10 seconds, it will start a trace process. During this process, the station uses the secondary ring instead of the primary ring to communicate with its upstream neighbor. Both nodes remove themselves from the ring and test the connections between them. If no fault is found, both stations rejoin the ring. If one of the stations encounters an error, it stays out of the ring and the other station performs the ring wrap function so that other nodes on the ring can continue to operate.

Common Problems Using FDDI

Although FDDI rings perform some basic maintenance functions to help take care of problems, it is still necessary to monitor the LAN periodically to ensure that the network is operating optimally. Also, many problems can't be corrected by software, such as faulty network adapters or network cables. Tools you can use for monitoring and troubleshooting efforts include a cable tester (one intended for use with fiber-optic cable) and a standard LAN protocol analyzer. Most FDDI vendors also will provide a station management application that can be used to examine ring functionality and gather statistics and error information. It is a good idea to get a thorough understanding of station management software so that you are better prepared when problems occur.

Ring Wrapping

This process allows for a malfunctioning node to be isolated from the other nodes in the ring. To restore the ring to normal functioning, it is necessary to track down the offending node and determine the cause of the failure. Station management software can provide the information you need in order to determine which station has left the ring. Perhaps the most common reason that ring wrapping occurs is a simple power failure. This can result from a faulty power supply in one of the attached workstations, or perhaps in a concentrator on the ring. It also can result from human error when someone who doesn't understand how the ring operates mistakenly powers down a station.

Other possible causes include all the associated hardware, from the cable to the connectors to the interface card that is installed in the workstation. Be sure to check that all connectors are correctly fastened and are not loose. Check for crushed or otherwise damaged cables. Fiber-optic cable is not as forgiving as twisted-pair can be, so it should be handled with care.

If all else fails, check the interface card on the computer. Use the vendor-supplied diagnostic software to determine whether the card is in good working order. If the card passes all the vendor's tests, try using it in a different slot on the computer bus or swapping it out for a card that is known to be working at this time.

Ring Initializations and Frame Check Sequence (FCS) Errors

Errors in the transmission of bits (causing FCS errors) or a high number of ring initializations (beacon or claim frames) can indicate problems with the light signal strength. It is easy to introduce dust, fingerprints, or other obstacles to the light signal when handling connectors. Connectors that are not tightly fastened also can lead to problems with the quality of the light signal and promote these kinds of errors. A cable tester designed for fiber-optic cabling can be used to test the signal strength for each segment.

Exceeding the maximum distance specifications for FDDI also can cause problems resulting in FCS errors or frequent ring initializations. Although a ring might be able to continue functioning if a single node is removed from the ring through ring wrapping, the removal of more than one can cause the total distance through the ring to be in excess of the specifications, resulting in a poor signal. If you are using optical bypass switches to help prevent ring wrapping, you can still experience these problems, because each optical bypass switch can introduce a two-decibel reduction in the signal strength.

Although not as common, a faulty network interface card or port on a concentrator might be the cause of these problems. Other possible sources that can affect signal quality include low-quality fiber-optic cable and imperfect splices in the cable. You can track down a faulty network adapter using the station management software or by using a LAN analyzer and looking for the station that is starting the claim process. Look at this station's upstream neighbor to find out whether a faulty card is causing the problem.

Making Repairs

Fiber-optic cabling is more expensive than copper wire cabling and requires a skilled technician for installation and repairs. A network administrator can easily perform tasks such as replacing a network adapter or moving a node to a different port on a concentrator. However, someone who is well trained in the techniques should splice fiber-optic cables or attach connectors. For all but large shops this usually means using an outside vendor or contractor.

- ◀◀ In Chapter 53, "Network Testing and Analysis Tools," you will find more information about instruments that can be used to test fiber-optic cables and to perform LAN analysis functions on an FDDI network.

In addition, with the introduction of obstructions such as dust or fingerprints, the light signal also can become degraded due to incorrect polishing when the cable is cut and attached to a connector. This doesn't mean that you cannot make such repairs yourself. If the size of the network warrants the expense, purchasing the tools and training an employee in the necessary techniques might be appropriate.

