Token Ring/IEEE 802.5

Background

The Token Ring network was originally developed by IBM in the 1970s. It is still IBM's primary local area network (LAN) technology and is second only to Ethernet/IEEE 802.3 in general LAN popularity. The related IEEE 802.5 specification is almost identical to and completely compatible with IBM's Token Ring network. In fact, the IEEE 802.5 specification was modeled after IBM Token Ring, and it continues to shadow IBM's Token Ring development. The term *Token Ring* generally is used to refer to both IBM's Token Ring network and IEEE 802.5 networks. This chapter addresses both Token Ring and IEEE 802.5.

Token Ring and IEEE 802.5 networks are basically compatible, although the specifications differ in minor ways. IBM's Token Ring network specifies a star, with all end stations attached to a device called a *multistation access unit* (MSAU). In contrast, IEEE 802.5 does not specify a topology, although virtually all IEEE 802.5 implementations are based on a star. Other differences exist, including media type (IEEE 802.5 does not specify a media type, although IBM Token Ring networks use twisted-pair wire) and routing information field size. Figure 9-1 summarizes IBM Token Ring network and IEEE 802.5 specifications.

Figure 9-1   Although dissimilar in some respects, IBM's Token Ring Network and IEEE 802.5 are generally compatible.
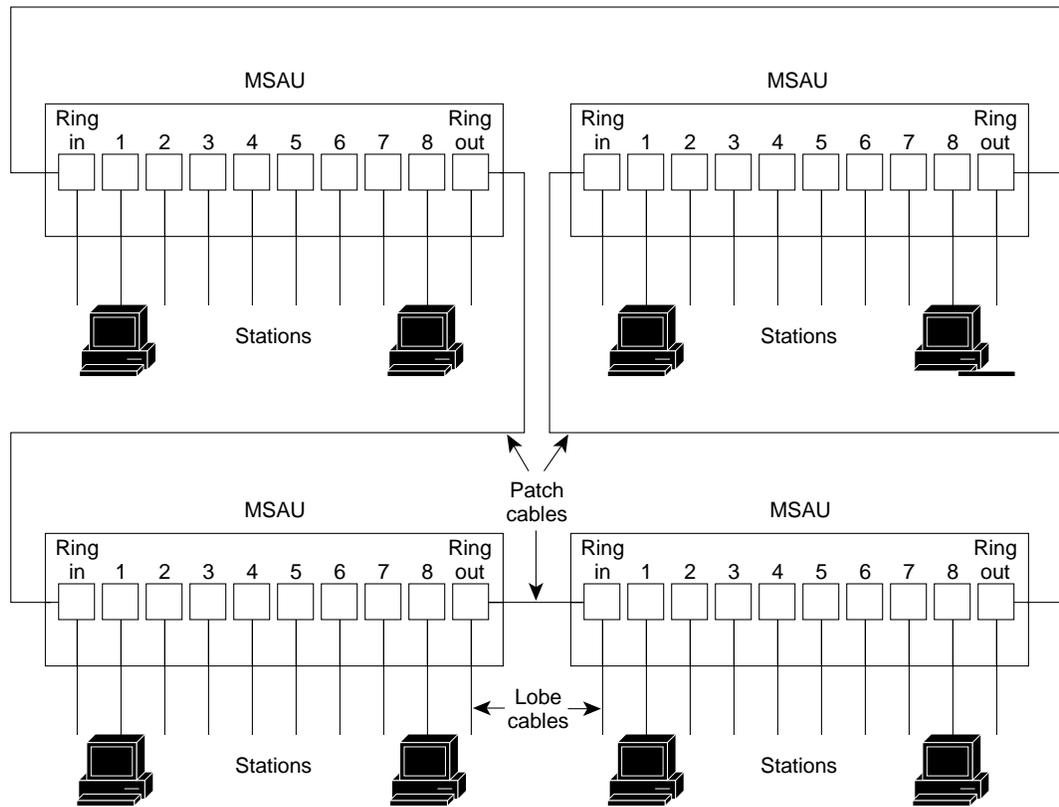
	IBM Token Ring network	IEEE 802.5
Data rates	4 or 16 Mbps	4 or 16 Mbps
Stations/segment	260 (shielded twisted pair) 72 (unshielded twisted pair)	250
Topology	Star	Not specified
Media	Twisted pair	Not specified
Signaling	Baseband	Baseband
Access method	Token passing	Token passing
Encoding	Differential Manchester	Differential Manchester

S1292a

Physical Connections

IBM Token Ring network stations are directly connected to MSAUs, which can be wired together to form one large ring (see Figure 9-2). Patch cables connect MSAUs to adjacent MSAUs, while lobe cables connect MSAUs to stations. MSAUs include bypass relays for removing stations from the ring.

Figure 9-2   MSAUs can be wired together to form one large ring in an IBM Token Ring network.

Token Ring Operation

Token Ring and IEEE 802.5 are two principal examples of token-passing networks (FDDI being the other). Token-passing networks move a small frame, called a *token*, around the network. Possession of the token grants the right to transmit. If a node receiving the token has no information to send, it passes the token to the next end station. Each station can hold the token for a maximum period of time.

If a station possessing the token does have information to transmit, it seizes the token, alters one bit of the token, which turns the token into a start-of-frame sequence, appends the information it wants to transmit, and sends this information to the next station on the ring. While the information frame is circling the ring, no token is on the network (unless the ring supports *early token release*), which means that other stations wanting to transmit must wait. Therefore, collisions cannot occur in Token Ring networks. If early token release is supported, a new token can be released when frame transmission is complete.

The information frame circulates the ring until it reaches the intended destination station, which copies the information for further processing. The information frame continues to circle the ring and is finally removed when it reaches the sending station. The sending station can check the returning frame to see whether the frame was seen and subsequently copied by the destination.

Unlike CSMA/CD networks (such as Ethernet), token-passing networks are *deterministic*, which means that it is possible to calculate the maximum time that will pass before any end station will be able to transmit. This feature and several reliability features, which are discussed in the section

"Fault-Management Mechanisms" later in this chapter, make Token Ring networks ideal for applications where delay must be predictable and robust network operation is important. Factory automation environments are examples of such applications.

Priority System

Token Ring networks use a sophisticated priority system that permits certain user-designated, high-priority stations to use the network more frequently. Token Ring frames have two fields that control priority: the *priority* field and the *reservation* field.

Only stations with a priority equal to or higher than the priority value contained in a token can seize that token. After the token is seized and changed to an information frame, only stations with a priority value higher than that of the transmitting station can reserve the token for the next pass around the network. When the next token is generated, it includes the higher priority of the reserving station. Stations that raise a token's priority level must reinstate the previous priority after their transmission is complete.

Fault-Management Mechanisms

Token Ring networks employ several mechanisms for detecting and compensating for network faults. One station in the Token Ring network, for example, is selected to be the *active monitor.* This station, which potentially can be any station on the network, acts as a centralized source of timing information for other ring stations and performs a variety of ring- maintenance functions. One of these functions is the removal of continuously circulating frames from the ring. When a sending device fails, its frame may continue to circle the ring. This can prevent other stations from transmitting their own frames and essentially can lock up the network. The active monitor can detect such frames, remove them from the ring, and generate a new token.

The IBM Token Ring network's star topology also contributes to overall network reliability. Because all information in a Token Ring network is seen by active MSAUs, these devices can be programmed to check for problems and selectively remove stations from the ring if necessary.

A Token Ring algorithm called beaconing detects and tries to repair certain network faults. Whenever a station detects a serious problem with the network (such as a cable break), it sends a beacon frame, which defines a failure domain. This domain includes the station reporting the failure, its nearest active upstream neighbor (NAUN), and everything in between. Beaconing initiates a process called autoreconfiguration, where nodes within the failure domain automatically perform diagnostics in an attempt to reconfigure the network around the failed areas. Physically, the MSAU can accomplish this through electrical reconfiguration.

Frame Format

Token Ring and IEEE 802.5 support two basic frame types: tokens and data/command frames. Tokens are 3 bytes in length and consist of a start delimiter, an access control byte, and an end delimiter. Data/command frames vary in size, depending on the size of the Information field. Data frames carry information for upper-layer protocols, while command frames contain control information and have no data for upper-layer protocols. Both formats are shown in Figure 9-3.

Figure 9-3   IEEE 802.5 and Token Ring specify tokens and data/command frames.
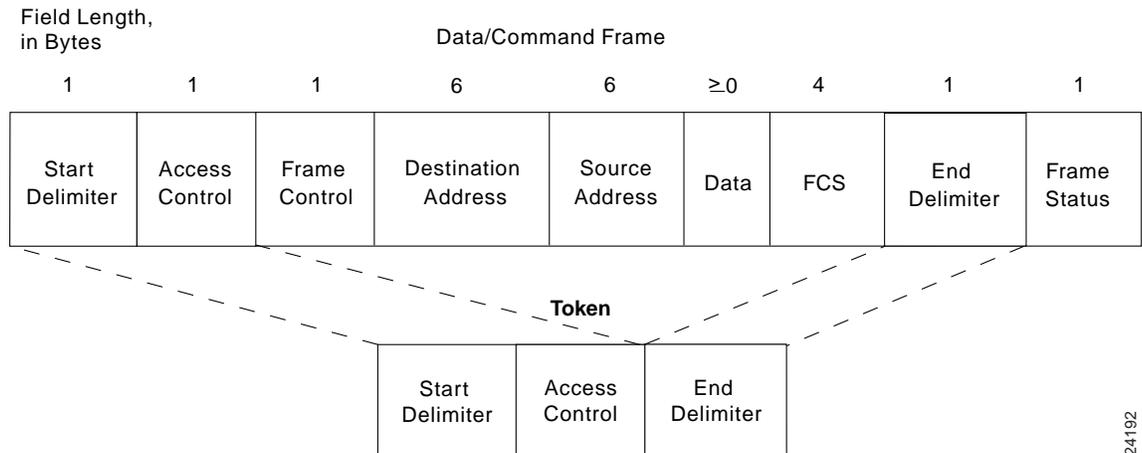
Token Frame Fields

The three token frame fields illustrated in Figure 9-3 are summarized in the descriptions that follow:

- *Start Delimiter*—Alerts each station of the arrival of a token (or data/command frame). This field includes signals that distinguish the byte from the rest of the frame by violating the encoding scheme used elsewhere in the frame.

- *Access-Control Byte*—Contains the Priority field (the most significant 3 bits) and Reservation field (the least significant 3 bits), as well as a token bit (used to differentiate a token from a data/command frame) and a monitor bit (used by the active monitor to determine whether a frame is circling the ring endlessly).

- *End Delimiter*—Signals the end of the token or data/command frame. This field also contains bits to indicate a damaged frame and identify the frame that is the last in a logical sequence.

Data/Command Frame Fields

Data/Command frames have the same three fields as Token Frames, plus several others. The Data/Command frame fields illustrated in Figure 9-3 are described in the following summaries:

- *Start Delimiter*—Alerts each station of the arrival of a token (or data/command frame). This field includes signals that distinguish the byte from the rest of the frame by violating the encoding scheme used elsewhere in the frame.

- *Access-Control Byte*—Contains the Priority field (the most significant 3 bits) and Reservation field (the least significant 3 bits), as well as a token bit (used to differentiate a token from a data/command frame) and a monitor bit (used by the active monitor to determine whether a frame is circling the ring endlessly).

- *Frame-Control Bytes*—Indicates whether the frame contains data or control information. In control frames, this byte specifies the type of control information.

- *Destination and Source Addresses*—Two 6-byte address fields identify the destination and source station addresses.

- *Data*—Length of field is limited by the ring token holding time, which defines the maximum time a station can hold the token.

- *Frame-Check Sequence (FCS)*—Filed by the source station with a calculated value dependent on the frame contents. The destination station recalculates the value to determine whether the frame was damaged in transit. If so, the frame is discarded.

- *End Delimiter*—Signals the end of the token or data/command frame. The end delimiter also contains bits to indicate a damaged frame and identify the frame that is the last in a logical sequence.

- *Frame Status*—A 1-byte field terminating a command/data frame. The Frame Status field includes the address-recognized indicator and frame-copied indicator.