An Introduction to Token Ring

National Semiconductor
Application Note 857
John von Voros

CONTENTS

INTRODUCTION

Token Ring and its associated hardware was first introduced by IBM® in the Fall of 1985. The token passing protocol was chosen because it would perform better under heavy network loading and be more easily integrated into existing IBM synchronous networks. In late 1985, the International Organization for Standardization (ISO) adopted the Token Ring Standard (ANSI/IEEE Std. 802.5–1985). A newer revision of the specification was approved by ISO in 1992 and has been published as International Standard **ISO/IEC 8802-5:1992.**

Initially, the 4 Mbps data rate was considered sufficient for most office automation and networking tasks. Realizing the need for increased bandwidth, IBM released the 16 Mbps version of Token Ring in 1989. Since its inception, Token Ring has become the most popular implementation of the token passing protocol and is second only to Ethernet® in nodes shipped.

In the token passing protocol, a token circulates around the ring until a station requests permission to transmit (*Figure 1*). Once granted, the transmitting station sends its information over the network. Each station along the ring checks the destination address of the data frame to determine if the packet is to be copied into a local buffer or simply repeated onto the ring. The data frame is removed from the ring by the originator of that frame and the token is released.

PHYSICAL CONNECTION

Each station is linked to the ring via a concentrator (*Figure 2*). Although this appears as a star-based topology, closer inspection reveals that it is indeed a ring. Most of the current concentrators are based on a passive technique utilizing simple relay logic. Each station has transformer-coupled receive and transmit twisted-pair wires (either shielded or unshielded) which attach the station to the concentrator using a Media Interface Connector (MIC).

When inserting a new station into the ring, the station first completes some initialization routines to insure proper operation. The station then "impresses" a D.C. voltage onto the Media Interface Cable to power the relay in the concentrator (*Figure 3*). This is called a *phantom* drive voltage because it is transparent to symbols being transmitted (due to AC coupling). The relay closes the switches so that the concentrator changes from bypass to insertion mode.

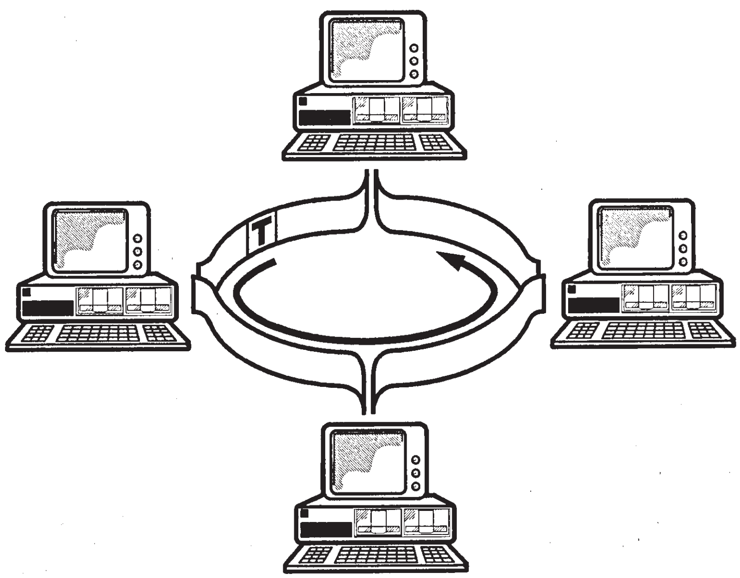

TL/F/11716–1

FIGURE 1. Token Ring

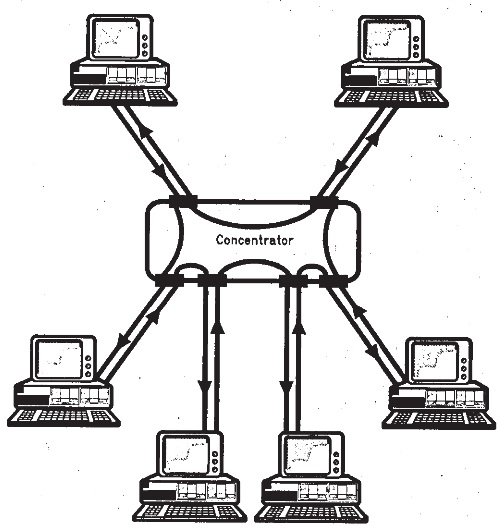

TL/F/11716-2

FIGURE 2. Station Connection to a Concentrator

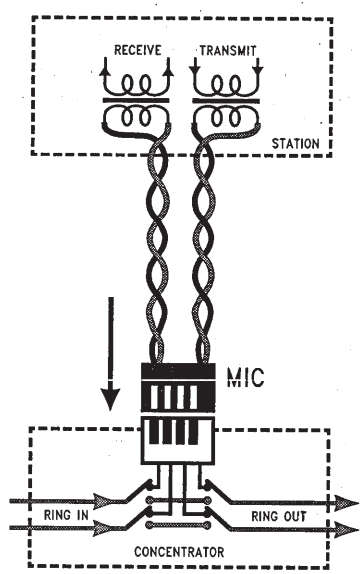

TL/F/11716-3

FIGURE 3. Physical Connection

RING SYNCHRONIZATION

Each ring must have a station (the Active Monitor) which is responsible for maintaining the master clock for all workstations. This is always the first station to get onto the ring. All other stations are considered Standby Monitors which constantly check the ring to make sure an Active Monitor is present. If a problem with the Active Monitor station develops, the other stations will negotiate to become the new Active Monitor.

The Active Monitor is the only station on a ring to synchronize to a clock oscillator *(Figure 4)*. All other stations synchronize to the embedded clock of the incoming signal using a phase-locked loop (PLL). One disadvantage to this approach is that each PLL along the ring introduces jitter into the signal. Jitter is the time varying difference between the phase of the master clock and the clock recovered from the signal. If this accumulated jitter reaches a certain threshold, the next station will not be able to recover a clock and the entire ring will fail. This is why the standards committees and manufacturers are sensitive to the amount of jitter introduced by token ring equipment. The current I.E.E.E. specification only allows for 250 stations in a ring to limit the potential for this catastrophic error.

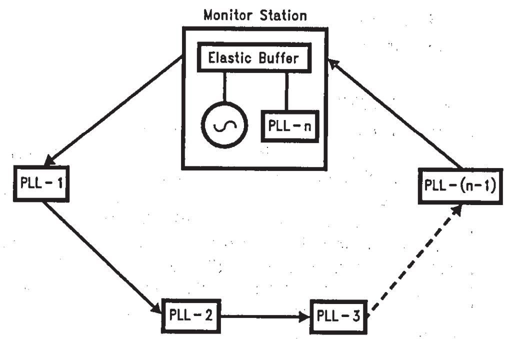

TL/F/11716-4

FIGURE 4. Ring Synchronization

Other responsibilities of the Active Monitor include recovery from error situations, Neighbor Notification, and maintaining a latency buffer. Neighbor Notification allows each station to know both its upstream and downstream neighbors on the ring. This is important for fault isolation. The latency buffer allows the Active Monitor to introduce a 24-bit delay (length of a token) to allow for a token to completely circulate around the ring without overlapping the originating station.

FRAME TYPES

The token ring standard currently supports two kinds of transmission units: tokens, and frames. As previously mentioned, tokens are responsible for controlling access to the network. The frames transfer data as well as control instructions for the network.

Token Format

The token is responsible for granting permission for an individual station to transmit onto the ring. This frame contains three fields: Starting Delimiter, Access Control, and Ending Delimiter *(Figure 5)*.

SD	AC	ED
1 Byte	1 Byte	1 Byte

SD = Starting Delimiter
AC = Access Control
ED = Ending Delimeter

FIGURE 5. Token Format

The Starting Delimiter field is a unique sequence of symbols indicating the beginning of a frame or token (See TROPIC™-Front End Description, AN-850). This field is the same for all frame types.

As shown in *Figure 6*, the Access Control field contains 3 bits indicating the priority of the frame or token (PPP), 1-bit indicating whether this is a frame (T=1) or a token (T=0), 1-bit for monitoring the ring to prevent the recirculation of errant tokens or frames (M), and 3 bits for requesting the needed priority setting for the next available token (RRR). Each station on the ring is assigned a maximum priority with which it may transmit.

P	P	P	T	M	R	R	R

PPP = Priority Bits
T = Token Status
M = Monitor
RRR = Reservation Bits

FIGURE 6. Access Control Field

The Ending Delimiter field contains a unique pattern of symbols indicating the end of a frame or token. One bit, the Intermediate Frame Bit, is used to indicate if the current frame is the end of a station's series of transmitted frames. The Error Detected Bit, the last bit in this field, is set by any station on the ring that detects an error. All tokens and frames must contain the Ending Delimiter sequence to be considered valid.

Frame Format

The frame is the basic data transmission unit of Token Ring networks. The IEEE standard for Token Ring (802.5) currently supports two frame types: MAC frames and LLC frames. The Medium Access Control (MAC) frames are used to control the operation of the ring. Examples of MAC frames include Claim Token, Duplicate Address Test, and Lobe Test which are explained in the Ring Insertion section. Logical Link Control (LLC) frames are responsible for the transmission of data and the establishment of communications between two or more nodes.

The basic frame format consists of the data to be sent encapsulated by a physical header and trailer *(Figure 7)*. The SD, AC, and ED fields are the same as those in the token.

The Frame Control field (FC) defines the frame to be either a MAC or LLC frame. This field also indicates how a MAC frame is to be buffered.

The Destination Address field (DA) is a 6 byte ID Standardized by I.E.E.E. which identifies the address of the receiving station on the ring. Bit 0 of byte 0 (the first bit of the DA transmitted) indicates whether the address is an individual or group address. All bits are set to a "1" for a broadcast message.

The Source Address field (SA) is also a 6 byte ID which identifies the station that originated the frame. Bit 0 of byte 0 is used to specify whether the data portion of the frame contains a routing information field (see section on internetworking). Routing information is only needed for frames that will leave the source ring.

The data field is of variable length and ranges from 0 bytes to a limit defined by the maximum time a station is allowed to hold the token. The current I.E.E.E. specification for the Token Hold Timer (THT) states that each station must support an 8.9 ms hold time which corresponds to a maximum information field length of 17800 bytes.

The Frame Check Sequence field (FCS) is a 32-bit check word that is generated during transmission of the FC, DA, SA, and INFO fields and then appended to the outgoing frame. During reception, the receiving station will calculate this check word on these same fields and compare it with the FCS byte transmitted to determine if the frame was sent correctly. Each station on the ring checks every frame for correct transmission and sets the appropriate bit in the Ending Delimiter field.

The Frame Status field (FS) indicates whether the current frame's address was recognized and if the information was copied by the receiving station *(Figure 8)*. The reserved bits are not currently used. The Address Recognized and Frame Copied bits are duplicated within the Frame Status byte because the error checking routine does not cover this field. These flags allow the network to determine three distinct conditions:

1. The station does not exist or is inactive.
2. The station exists, but the frame was not copied (this could be caused by network congestion).
3. The frame was copied properly by the receiving station.

A	C	r	r	A	C	r	r

A = Address Recognized
C = Frame Copied
r = Reserved

FIGURE 8. Frame Status Field

RING OPERATION

The following section covers the procedures for ring insertion, transmission of a frame, and frame reception. A more detailed explanation can be found in the Token Ring Network Architecture Reference by IBM (see references).

Ring Insertion

Step A: Check the physical connection and receive logic.

1. Transmit *lobe test MAC frames* onto the station's cable only (not onto the ring). This is done to verify cable attachment.
2. Transmit *duplicate address test MAC frames* (refer to Step C). These are only used for testing the receive logic at this point (not for actual address duplications).
3. After both tests are passed, the station is inserted into the ring by exerting the phantom voltage which causes the concentrator to switch the station into the ring.

Physical Header						Physical Trailer		
SD	AC	FC	DESTINATION ADDRESS	SOURCE ADDRESS	INFO	FCS	ED	FS
1 BYTE	1 BYTE	1 BYTE	6 BYTES	6 BYTES	VARIABLE	1 BYTE	1 BYTE	1 BYTE

TL/F/11716–5

SD = Starting Delimiter
AC = Access Control
FC = Frame Control
FCS = Frame Check Sequence
ED = Ending Delimiter
FS = Frame Status

FIGURE 7. Frame Format

Step B: Check for an active monitor on the ring.

1. Initiate countdown timer. (Join Ring Timer = 18 second maximum)
2. Within this time, the station must detect an active monitor by receiving any of three special MAC frame types: ring purge frame **(PRG)**, active monitor present frame **(AMP)**, or standby monitor present frame **(SMP).** If this occurs, the initialization routine proceeds to Step C.
3. If not, the station will assume one of three possibilities; This is the first station on the ring, no active monitor is present, or insertion of this station has broken the ring. The token claim process is initiated to contend to become the active monitor.

Step C: Check for a duplicate address on the ring.

1. Transmit a duplicate address test frame onto the ring to see if another station has the same physical address. This is done by setting the destination address to be the same as the address of your station. If the frame returns with the address recognized bit clear (address not recognized), proceed to Step D.
2. If the frame returns with the address recognized bit set, a duplicate physical address has been found. The station notifies the network manager and removes itself from the ring.

Step D: Neighbor notification (used for network management).

1. Determine the nearest available upstream neighbor **(NAUN).**
2. Notify the downstream neighbor of your station's identity.

Step E: Request initialization.

1. Identify your station to the network manager including the address, the address of the NAUN, the product instance identification, and the revision of the controller's microcode.
2. The station is now active on the ring and enters the normal repeat mode.

Transmitting on the Ring

Step A: Capture token.

1. The station transmitting the frame must have a priority greater than or equal to the priority of the token in order to use it.
2. If a frame or a token with higher priority passes, the reservation bits RRR in the Access Control (AC) field are set to request a token of the appropriate priority.

Step B: Transmit Frame.

1. Upon receipt of a usable token, the station converts the token to a start of frame sequence (SFS) by setting the "T" bit in the AC field to a "1".
2. The station now changes from repeat mode to transmit mode and transmits its frame onto the ring.
3. The frame check sequence (FCS) is calculated and appended to the information field of the outgoing frame.
4. The address recognized (A) and frame copied (C) bits are cleared within the frame status (FS) field which is the final field transmitted in a frame.

Step C: Transmit token and strip frames from the ring.

1. Unless early token release is enabled (see note), the sending station will delay the transmission of the token until it receives the source address field of the last transmitted frame. This delay is accomplished by transmitting idles (any combination of 1's or 0's; a.k.a. fill bits) and may be of any length within the constraints of the maximum amount of time a token can be held by a station. These idles are required so that the PLL's of the stations maintain synchronization with the active monitor.

Note: Early Token Release (ETR) was introduced to make more efficient use of the available ring bandwidth. The token is immediately transmitted following the Frame Status byte of the outgoing frame. Without ETR, the sending station would wait until it had received the physical header back before releasing the token. If the header returns to the sending station before the token is released, the token will be transmitted with a priority equal to the reservation bits in the Access Control field. If the header has not yet been received, the token is released with a priority equal to the transmitted frame. Early token release still only allows one token on the ring at any one time.

2. While continuing to strip the frame, the sending station transmits a token with the appropriate priority as determined by the reservation bits in the Access Control field.
3. The sending station will continue to transmit idles after the token until the stripping process is completed.

Step D: Station returns to normal repeat mode.

Frame Reception

Step A: Check the destination address of the incoming frame to determine if it matches your station's individual or group address, or is a broadcast address.

Step B: If no match occurs, the station repeats the frame onto the ring. This is known as normal repeat mode. The station will check the data in the tokens and frames received, and will set the appropriate A, C, and E (address recognized, frame copied, and error detected) bits within the Ending Delimiter and Frame Status fields at the end of the frame. If an error is detected by any station, the E bit will be set and the frame will not be copied by the destination station.

Step C: If the destination address matches, the station begins copying the frame into its buffer while simultaneously repeating the frame onto the ring. The station will then set the A, C, and E bits appropriately.

TOKEN RING BRIDGING

Multiple rings are connected by bridges which are responsible for copying frames from the source ring and forwarding that information to the destination ring *(Figure 9)*. Bridges are characterized by ring numbers, bridge numbers, hop count limits, largest frame sizes, and single route broadcast indicators. Ring numbers identify the numbers of the rings on either side of the bridge. Each bridge between two rings is assigned a unique bridge number. The hop count limits the number of rings an All Routes Broadcast frame can travel before the bridge discards it. The largest frame size indicates the length of the frame that is either supported by the bridge, or the rings on either side. The single route broadcast indicator determines if the bridge is capable of supporting single route broadcast frames.

Frames are routed through Token Ring networks by using a procedure called source routing in which frames carry the information on the route to be taken. In contrast, transparent bridging requires that the bridges keep track of all routes between stations.

When first determining a route, the source station sends a test frame on the local ring with the desired destination address. If the proper test frame response is not received, the station must look elsewhere on the network. Two methods are used to determine the routes to destinations on other rings: All Routes Broadcast Route Determination and Single Route Broadcast Determination.

When the All Routes Broadcast technique is chosen, the source station sends a test frame to all rings. As this frame traverses the network, it records information on the bridges it encounters on the way to the destination. Since multiple paths may exist between source and destination, many copies of this test frame may reach the destination with unique routing information. As each copy is received by the destination, a test frame response is sent back, via the same route, to the source station with the acquired routing information. The source station then decides which route is preferred. This decision could be based on cost, bandwidth, network congestion, etc. Once the route has been determined, all dialogue between the two stations proceeds along the chosen path.

The other possible method of route determination, Single Route Broadcast, sends the test frame so that only one copy of the frame reaches the destination. This technique utilizes the single route broadcast method that is an inherent feature of source routing bridges. When the destination station receives the test frame, it answers with an All Routes Broadcast test response frame which could result in multiple copies appearing at the source station, each with its own routing information. Once again, the source station determines the preferred route of transmission.

Bridges scrutinize the routing information to make sure that a frame is not continuously circulating around the network. If a bridge discovers the number of the next ring included in the routing information field, it discards the frame. The bridge also discards frames if the hop count limit has been exceeded.

CONCLUSION

Token Ring is gaining popularity due to standardization efforts and the decreasing cost of hardware. In mid-1992, IBM introduced the capability of using standard category 3, 4, or 5 unshielded twisted-pair for 16 Mbps Token Ring which should further reduce the cable plant cost. IEEE is expected to approve this standard by early 1993. In February of 1992, IBM licensed its Token Ring technology to National Semiconductor thus making IBM silicon available to O.E.M.'s for integration into their products.

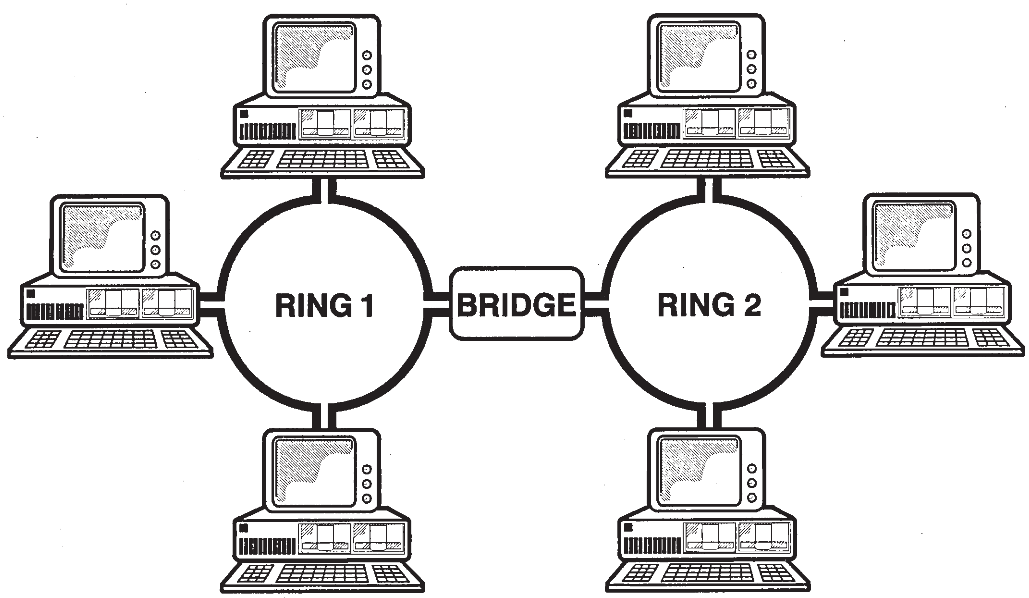

TL/F/11716-6

FIGURE 9. Token Ring Bridging

REFERENCES

International Standard ISO/IEC 8802-5: 1992 (E) IEEE Std 802.5—1992

Naugle, Matthew G., "Local Area Networking", McGraw-Hill, Inc., San Francisco, CA, 1991.

"Token Ring Network Architecture Reference", 3rd ed., IBM Corporation. Research Triangle Park, N.C., 1989.

DEFINITIONS
(From IEEE Specification)

Abort Sequence. A sequence that terminates the transmission of a frame prematurely.

Accumulated Jitter. The jitter measured against the clock of the active monitor. Like alignment jitter, this is not a type of jitter but a way to measure total jitter growth throughout the ring. It is normally used to determine the required size of the elastic buffer.

Alignment Jitter. The jitter measured against the clock of the upstream adapter. This is not a type of jitter per se; rather, it is a way to measure jitter. When "zero transferred jitter" is specified, the jitter measured is alignment jitter.

Broadcast Transmission. A transmission addressed to all stations.

Channel. The channel is the transmission path from the MIC at the transmitter to the first MIC at the receiver. It may include TCUs and connectors in addition to transmission line.

Configuration Report Server (CRS). A function that controls the configuration of the ring. It receives configuration information from the stations on the ring and either forwards them to the network manager or uses them to maintain a configuration of the ring. It can also, when requested by the network manager, check the status of stations on the ring, change operational parameters of stations on the ring, and remove stations from the ring.

Correlated Jitter. The portion of the total jitter that is related to the data pattern. Since every adapter receives the same *pattern*, this jitter is *correlated* among all adapters and therefore grows in a *systematic* way along the ring. Correlated jitter is also called *pattern jitter* or *systematic jitter*.

Differential Manchester Encoding. A signaling method used to encode clock and data bit information into bit symbols. Each bit symbol is split into two halves, where the second half is the inverse symbol of the first half. A 0 bit is represented by a polarity change at the start of the bit time. A 1 bit is represented by no polarity change at the start of the bit time. Differential Manchester encoding is polarity-independent.

Fill. A bit sequence that may be either 0 bits or 1 bits or any combination thereof.

Frame. A transmission unit that carries a protocol data unit (PDU) on the ring.

Jitter. The time-varying difference between the phase of the recovered clock and the phase of the source clock. Jitter is measured in fractions of a clock cycle, or unit interval (UI).

Logical Link Control (Sublayer) (LLC). That part of the data link layer that supports media-independent data link functions, and uses the services of the MAC to provide services to the network layer.

Medium. The material on which the data may be represented. Twisted pairs, coaxial cables, and optical fibers are examples of media.

Medium Access Control (Sublayer) (MAC). The portion of the data station that controls and mediates the access to the ring.

Medium Interface Connector (MIC). The connector between the station and TCU at which all transmitted and received signals are specified.

Monitor. The monitor is that function that recovers from various error situations. It is contained in each ring station; however, only the monitor in one of the stations on a ring is the *active monitor* at any point in time. The monitor function in all other stations on the ring is in standby mode.

Multiple Frame Transmission. A transmission where more than one frame is transmitted when a token is captured.

Physical (Layer) (PHY). The layer responsible for interfacing with the medium, detecting and generating signals on the medium, and converting and processing signals received from the medium and the MAC.

Protocol Data Unit (PDU). Information delivered as a unit between peer entities that contains control information, and optionally, data.

Protocol Implementation Conformance Statement (PICS). A statement of which capabilities and options have been implemented for a given Open Systems Interconnection protocol.

Repeat. The action of a station in receiving a bit stream (for example, frame, token, or fill) from the previous station and placing it on the medium to the next station. The station repeating the bit stream may copy it into a buffer or modify control bits as appropriate.

Repeater. A device used to extend the length, topology, or interconnectivity of the transmission medium beyond that imposed by a single transmission segment.

Ring Error Monitor (REM). A function that collects ring error data from ring stations. The REM may log the received errors, or analyze this data and record statistics on the errors.

Ring Latency. In a token ring MAC system, the time (measured in bit times at the data transmission rate) required for a signal to propagate once around the ring. The ring latency time includes the signal propagation delay through the ring medium plus the sum of the propagation delays through each station connected to the token ring.

Ring Parameter Server (RPS). That function that is responsible for initializing a set of operational parameters in ring stations on a particular ring.

Routing Information. A field, carried in a frame, used by source routing transparent bridges that provides source routing operation in a bridged LAN.

Service Data Unit (SDU). Information delivered as a unit between adjacent entities that may also contain a PDU of the upper layer.

Source Routing. A mechanism to route frames, through a bridged LAN. Within the source routed frame, the station specifies the route that the frame will traverse.

Station (or Data Station). A physical device that may be attached to a shared medium LAN for the purpose of transmitting and receiving information on that shared medium. A data station is identified by a destination address (DA).

Static Phase Offset. The constant difference between the phase of the recovered clock and the optimal sampling position of the received data.

Station Management (SMT). The conceptual control element of a station that interfaces with all of the layers of the station and is responsible for the setting and resetting of control parameters, obtaining reports of error conditions, and determining if the station should be connected to or disconnected from the medium.

Token. The symbol of authority that is passed between stations using a token access method to indicate which station is currently in control of the medium.

Transferred Jitter. The amount of jitter in the recovered clock of the upstream adapter. Transferred jitter is important because each adapter must both limit the amount of jitter it generates, and track the jitter delivered by the upstream adapter.

Transmit. The action of a station generating a frame, token, abort sequence, or fill and placing it on the medium to the next station. In use, this term contrasts with *repeat*.

Transparent Bridging. A bridging mechanism, in a bridged LAN, that is transparent to the end stations.

Trunk Cable. The transmission cable that interconnects two TCUs.

Trunk Coupling Unit (TCU). A physical device that enables a station to connect to a trunk cable. The TCU contains the means for inserting the station into the ring, or conversely, bypassing the station.

Uncorrelated Jitter. The portion of the total jitter that is independent of the data pattern. This jitter is generally caused by noise that is uncorrelated among adapters and therefore grows in a nonsystematic way along the ring. Uncorrelated jitter is also called *noise jitter* or *nonsystematic jitter*.

Unit Interval (UI). One half of a bit time. 125 ns for 4 Mb/s transmission and 31.25 ns for 16 Mb/s transmission. UI is used in the specification of jitter.

Upstream Neighbor's Address (UNA). The address of the station functioning upstream from a specific station.