Network Troubleshooting

by Othmar Kyas

8 Token Ring

An Agilent Technologies Publication

Agilent Technologies

8 Token Ring

"If the odds are a million to one against something occurring, the chances are 50-50 that it will."

UNKNOWN

8.1 Token Ring: Specification and Implementation

Unlike Ethernet, Token-Ring networks are based on deterministic data transmission techniques. Each station in the ring is periodically given an opportunity to transmit data, and no two stations can transmit simultaneously. Token-Ring technologies offer data speeds of 4 and 16 Mbit/s. Special frames, known as "tokens", control access to the LAN medium. A token is a sequence of bits that

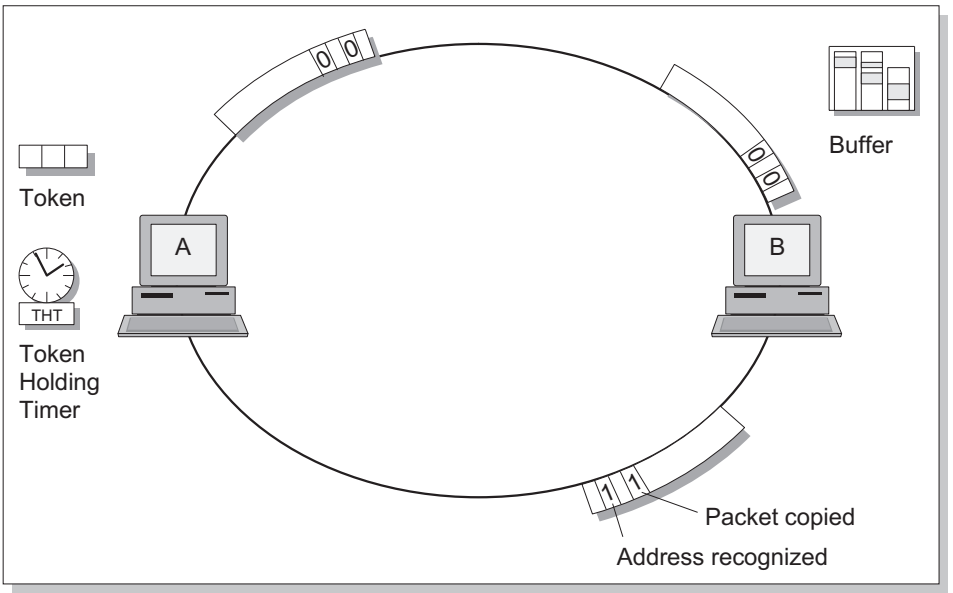

Figure 8.1 The token-passing principle (Token Ring, FDDI)

circulates continuously through the ring. When a station receives a token it can transmit data. The token must be released after a specified period called the token holding time, at which point it becomes available to the next station. When no data needs to be transmitted, the token circulates unused through the ring. If a station that receives the token has data waiting for transmission, it immediately begins transmitting its own data rather than passing the token to the next station.

When a station recognizes its own address as the destination for a data packet, it copies the packet to its receive buffer, sets the address recognized and frame copied bits in the frame header to 1, and returns the modified frame to the source station. When the station that originally sent the data receives the modified frames, it evaluates the information in the frame status field (address recognized, frame copied) and removes the frame from the ring. Stations not participating in a given data exchange act as repeaters, simply forwarding frames downstream to the next station in the direction of ring traffic (Figure 8.1).

Data packets cannot collide in Token-Ring networks because only the station that has the token can transmit data. Consequently, theoretical total bandwidth is used more efficiently in Token Ring than in Ethernet networks. Only when capacity use rises above 80 percent, the time a station must wait for the token doubles, and a decrease in network performance becomes noticeable.

8.1.1 The Physical Layer in Token-Ring Networks

Token Ring distinguishes between four types of signal on the physical layer:

- 0 binary 0
- 1 binary 1
- J non-data J
- K non-data K

Signals are transmitted in Differential Manchester encoding (as opposed to the simple Manchester encoding used in Ethernet). In this process, voltage transitions can occur at the beginning and in the middle of each bit time. A logical 1 is identified by the fact that there is no voltage transition at the beginning of the bit; for a logical 0 there is a transition at the beginning. A voltage transition always occurs in the middle of the bit interval. This ensures that the resulting signal is DC-balanced and can be inductively or capacitively coupled. Only the J and K symbols deviate from the rule described for splitting signals. A J signal begins with the same polarity as the signal that preceded it, and a K signal begins with the opposite polarity of its preceding signal. To avoid a residual

DC component in the ring, J and K symbols are transmitted in pairs. One disadvantage of Differential Manchester coding is that, as in Ethernet, the effective bit rate of 4 or 16 Mbit/s is only half the actual baud rate of 8 or 32 MHz.

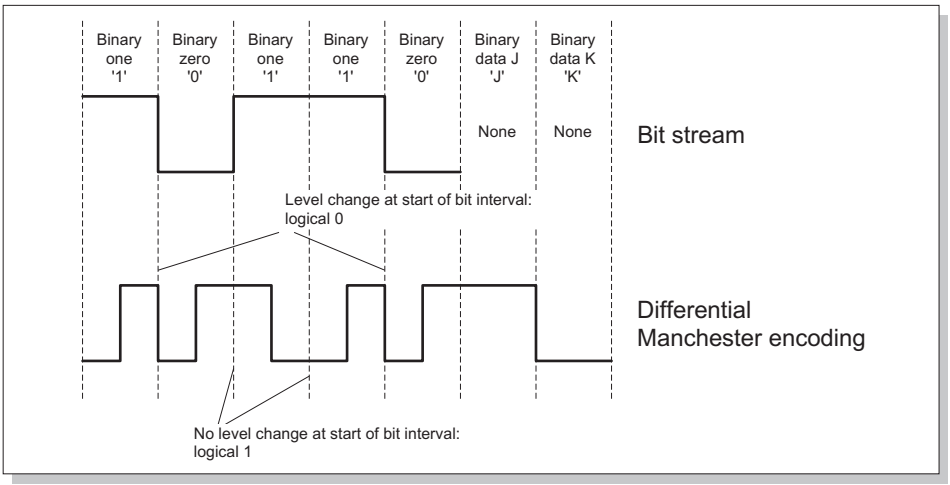

Figure 8.2 Differential Manchester encoding

8.1.1.1 Physical and Logical Ring Lengths

From the physical point of view, a Token-Ring network consists of stations daisy-chained together in a circle. Each station has a one-bit buffer into which each bit that is received is copied before being passed on along the ring. While this bit is in the buffer it can be analyzed and changed if necessary. Each station increases the ring length by exactly one bit time. At a typical signal propagation speed of 200 m/µs (approximately 0.7 c), each bit in a 4 Mbit/s Token-Ring network travels some 50 meters (or 12.5 meters in a 16 Mbit/s ring). Consequently, each additional station increases the apparent ring length by some 50 meters (or 12 meters in the case of a 16 Mbit/s ring). Furthermore, by this calculation, a ring that is 1,000 meters long can only hold 20 bits at a time! Yet a Token Ring must be capable of carrying at least 24 bits at a time, which is the length of a token. This is why the active monitor, the station charged with monitoring the ring, operates a delay buffer to guarantee a logical ring length of at least 24 bits at a time.

8.1.1.2 Synchronization and Pulse

The physical layer of each station regenerates the clock information and minimizes phase jitter in the signal received. During normal operation, a Token Ring has one active monitor that sets the ring clock rate to its own local oscillator. All

other stations are synchronized with the frequency and phase of this clock by a phase-locked loop with the following minimum precision requirements:

- The maximum dynamic jitter permitted in a station is 3 sigma (= 10 degrees).
- When a station enters the ring or has lost synchronization, it must be able to resynchronize in phase with the active monitor within 1.5 ms.
- The timing of the implementation must be precise enough to permit at least 250 active stations on the ring at one time.

Jitter and synchronization problems are among the most common sources of malfunctions in Token-Ring networks. If the signal deviates too far from the clock rate during data transmission, the ring's stations may get out of sync with the data stream. Technologically advanced Token-Ring components can stabilize themselves without the help of the active monitor using their own local oscillators.

8.1.1.3 Phase Jitter Compensation

A slight variation in data speeds in the ring can result from signal jitter. With 250 stations in a ring, this can lead to differences in data speed of up to ±3 bits. If the round trip time is decreased, bits are lost. To prevent this, a dynamic buffer of up to 6 bits (corresponding to 12 signal elements or 12 voltage transitions) is added to the 24-bit circulation buffer in the active monitor. The total buffer is initialized with a capacity of 27 bits. If the speed of the data received by the active monitor is slightly higher than that of the master oscillator, the buffer capacity can be expanded to 28, 29 or 30 bits as required. If the data speed is lower, buffering can be reduced to 24 bits.

8.1.1.4 Connecting Ring Stations

Each station is connected to the Token-Ring network by means of a concentrator, also called a trunk coupling unit (TCU) or multi-station access unit (MSAU). The cable that connects the station to the concentrator, called the lobe cable, consists of shielded, four-lead copper cable with an impedance of 150 ± 15 Ω. The connector at the concentrator end (medium interface connector or MIC) specified for Token Ring is an IBM Type 1 connector. When a station is inserted in the ring, two DC voltages are applied to the concentrator (one over each wire pair), and the two DC circuits are inductively coupled with both the station's and concentrator's actual transmit and receive leads. In this way the AC signals, or data streams, are transported through the inductive coupling, while the DC current in the two wire pairs can be monitored to detect open or short circuits. The DC voltage also controls the relay in the concentrator, which creates the actual mechanical connection between the station and the ring. When the

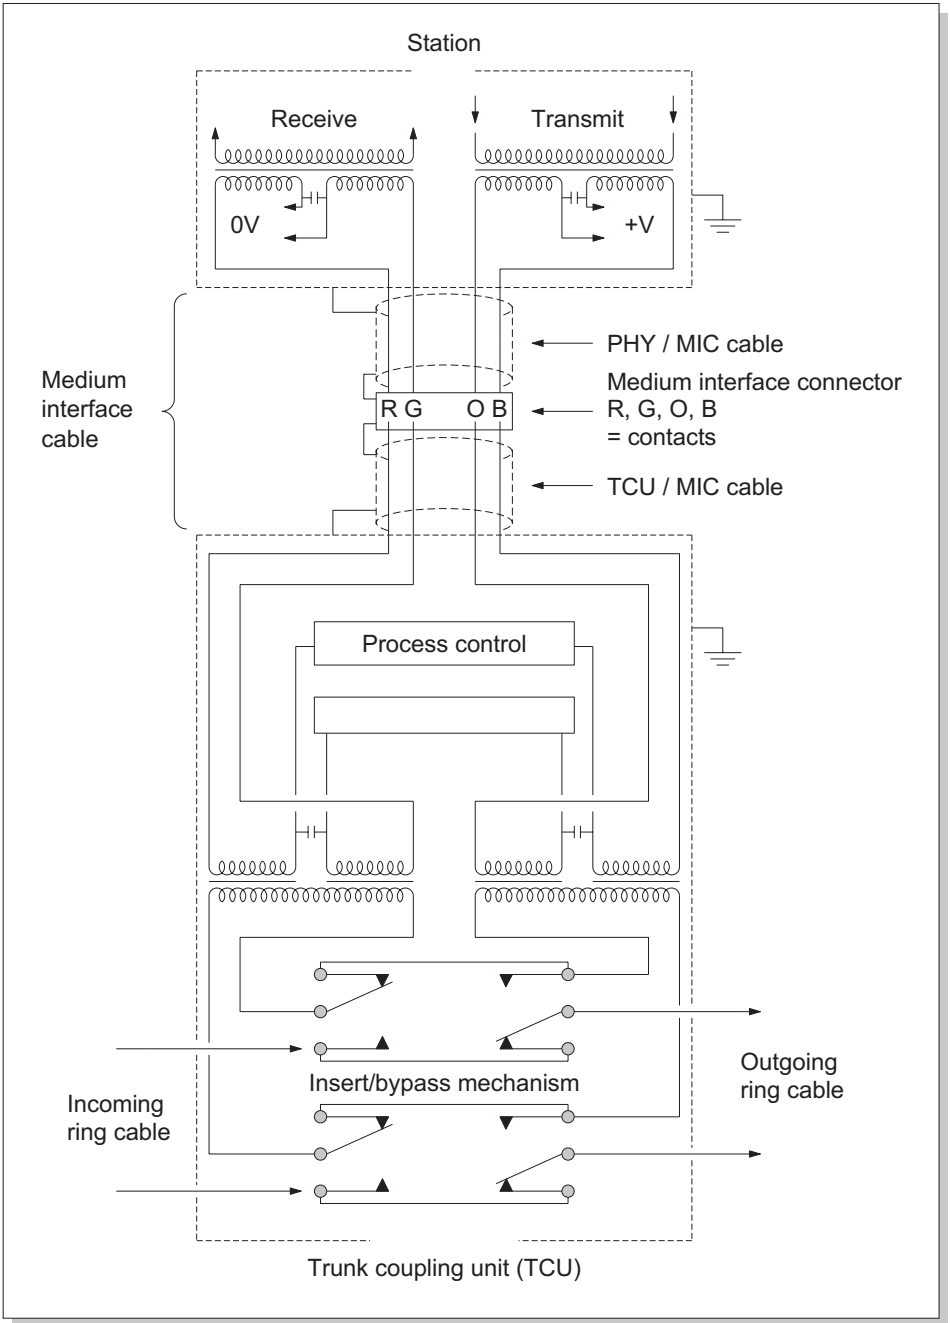

Figure 8.3 Station and concentrator

station is included in the ring, 4.1 to 7.0 volts are applied across points B and O (Figure 8.3), resulting in a current of 0.65 to 2.0 mA between points G and R. In the bypass state, that is when the station is not participating in the ring, the voltage is below 1 volt.

8.1.2 The Token-Ring Data Format

There are two types of frames in Token-Ring networks: tokens and data packets. A token is 3 bytes long and consists of a starting delimiter, an access control

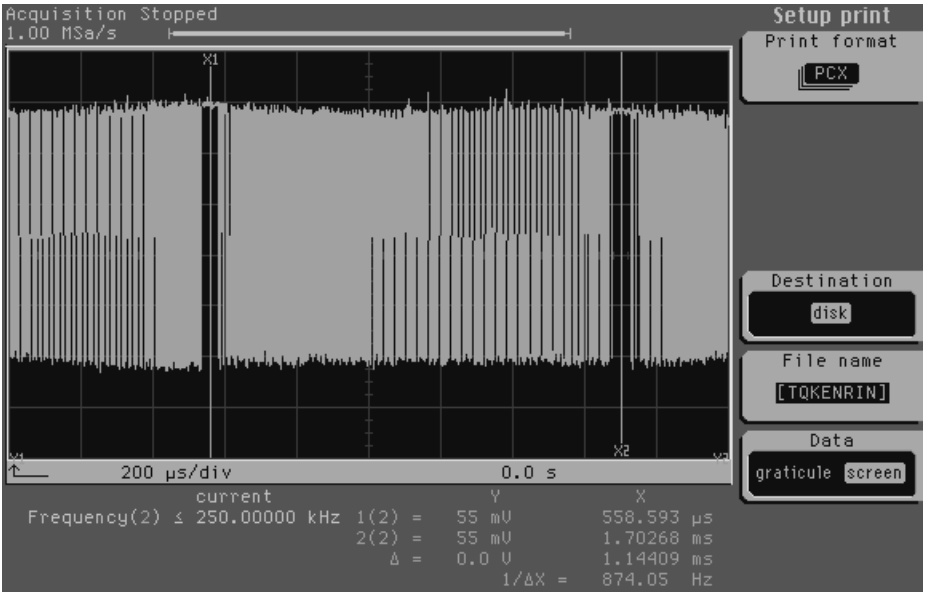

Figure 8.4 Bit sequence of a token recorded with a digital oscilloscope

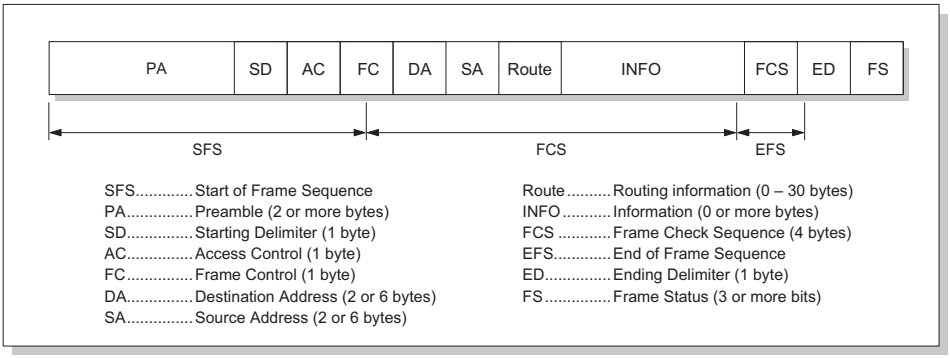

Figure 8.5 Data packet format in Token Ring

field, and an ending delimiter. Figure 8.4 shows a token sequence recorded with a digital oscilloscope, circulating in an empty ring at a frequency of 250 kHz.

All other data packets can have lengths from 13 to 4,500 bytes in a 4 Mbit/s ring, or 13 to 17,800 bytes in a 16 Mbit/s ring, and consist of the following fields: start of frame sequence, preamble, starting delimiter, access control, frame control, destination address, source address, information, frame check sequence, end of frame sequence, ending delimiter and frame status.

8.1.2.1 Token and Data Packet Fields

Starting Delimiter (SD)

Every frame, including tokens, begins with this field. Any sequence of bits that does not begin with this field is discarded.

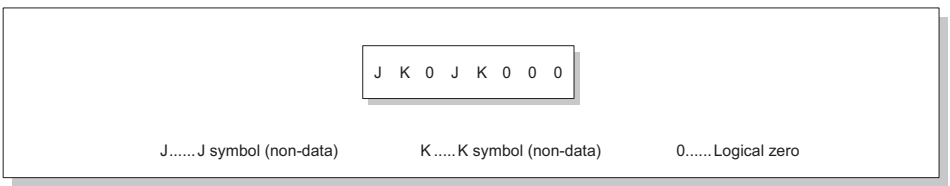

Figure 8.6 The starting delimiter field in Token Ring

Access Control (AC)

The access control field is responsible for controlling access to the network medium in Token Ring. The priority bits (PPP) specify the priority level of the token. There are eight levels of priority, increasing from 000 to 111 (for example, priority 110 is higher than 011). In a token, the token bit T is set to 1; in all other frames it is set to 0. When a station receives a token with a priority level equal to or lower than that of the data packet in its send queue, the station can transform the token into a start of frame sequence and transmit its data. The monitor bit M prevents data packets and tokens with a priority greater than 0 from circulating continuously in the ring. The default value for this bit in all frames, including tokens, is 0. The active monitor, the station acting as ring manager, sets this bit to 1 before forwarding any frame. Whenever a frame has the monitor bit set to 1,

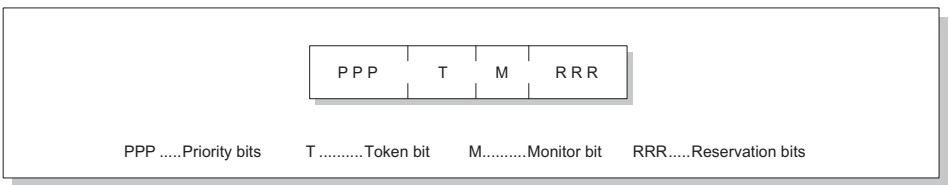

Figure 8.7 The access control field in Token Ring

this indicates that the frame has passed through the active monitor since its generation.

By setting the reservation bits RRR, stations can request a token with a specific priority level. This may be necessary when a station has data packets that need to be sent urgently and cannot wait for a normal token.

Frame Control (FC)

The frame control field identifies the frame type. It is used to distinguish between MAC frames and LLC frames. MAC frames are used for ring management, while LLC frames contain user data. The information about the frame type is contained in the two format bits (FF):

00 MAC frame
01 LLC frame
1X Undefined

How stations react to a MAC frame is determined by the control bits the frame contains. In an LLC frame, the first three control bits ZZZ are set to 0, and the data packet's priority is indicated in the remaining three bits.

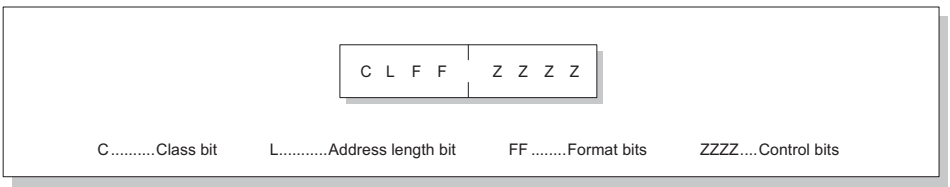

Figure 8.8 The frame control field in Token Ring

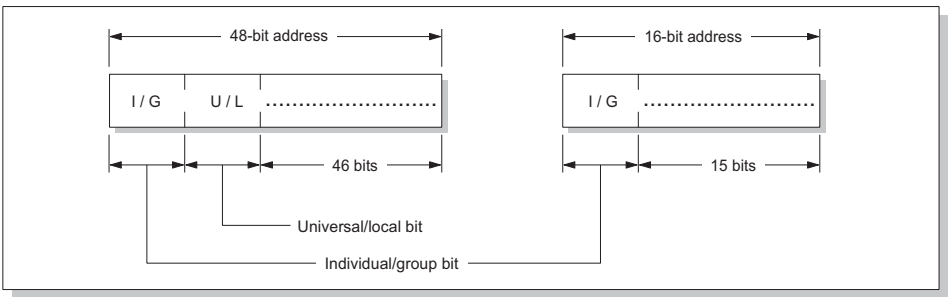

Figure 8.9 The destination address field in Token Ring

Destination Address (DA)

The destination address identifies the station that is the intended recipient of a given data packet. The first bit indicates whether the destination is an individual

address (bit 1=0) or a group address (bit 1=1); the second bit (only in 48-bit addresses) shows whether the address is locally administered (bit 2=1) or universally administered (bit 2=0). If the destination address field contains all 1s, then the destination is a broadcast address, which means the packet is intended for all stations on the ring. An address consisting entirely of 0s is called a "zero address"; packets with this address are not intended for any station.

Source Address (SA)

The source address identifies the station from which a data packet originates. The length and format are the same as in the destination address field, but the first bit is always a 0.

Routing Information

The routing information field, with a length of up to 30 bytes, is required only when source routing is implemented to determine transmission paths in the network. When this field is present, it consists of a 2-byte routing control field and a variable number of route designator fields.

Information (INFO)

The information or data field can contain zero, one or more bytes, which may be addressed to the MAC layer, the LLC layer, or to network management functions—the total packet length must not exceed 4,500 bytes (4 Mbit/s)/17,800 bytes (16 Mbit/s). If the information field is part of a MAC frame it is directly processed by the Token-Ring protocol. In an LLC frame, the information field contains user data, which is passed to the higher protocol layers at the destination node. Each byte in the information field is transmitted starting with the most significant bit. The information field of an LLC frame begins with the following three fields: destination service access point (DSAP, 1 byte), source service access point (SSAP, 1 byte) and the LLC control field (1 or 2 bytes). The DSAP and SSAP designate the protocol in the subsequent information fields. These fields can specify any of 128 different protocols. The Sub-Network Access Protocol (SNAP), which is designated by the values DSAP=AA, SSAP=AA and LLC Control=03, is a kind of meta-protocol: it suspends the limit of 128 protocols by adding another field to allow manufacturers practically unlimited use of proprietary protocols.

Frame Check Sequence (FCS)

The FCS field contains a 32-bit checksum calculated from the FC, DA, SA and INFO fields. This field allows the receiving node to detect bit errors in transmission.

Ending Delimiter (ED)

A receiving station considers the ending delimiter to be valid if the first six symbols, JK1JK1, are received correctly. The intermediate frame bit indicates whether other frames follow in the sequence. If the value of the 1 bit is 0, the frame is the last in its sequence. The error detection bit of a data packet, token, or abort sequence is always set to 0 by the transmitting node. If any station detects an error (FCS error, non-data signal, etc.) when forwarding a frame, it sets the error detection bit to logical 1. Otherwise, the frame is passed along with the error detection bit unchanged.

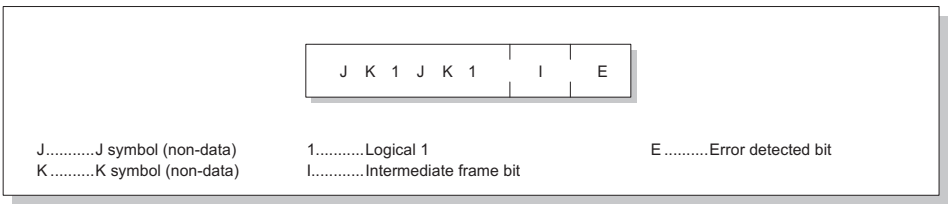

Figure 8.10 Ending delimiter

Frame Status

The destination node uses the frame status field to provide the source node with confirmation that the frame was received. When a frame is first sent, the transmitting node sets the A (address recognized) and C (frame copied) bits to 0. When a receiving station identifies the destination address of the data packet as its own, it sets the address recognized bit to 1. If the receiving station is able to copy the data packet to its receive buffer, the frame copied bit is also set to 1. This allows the transmitting node to distinguish between three cases upon receiving the returning frame:

- The destination node does not exist in the ring or is inactive,
- The station exists but the data packet was not received, or
- The data packet was copied into the receive buffer of the destination node.

The A and C bits are used regardless of the status of the error detection bit as long as the frame is found to be valid.

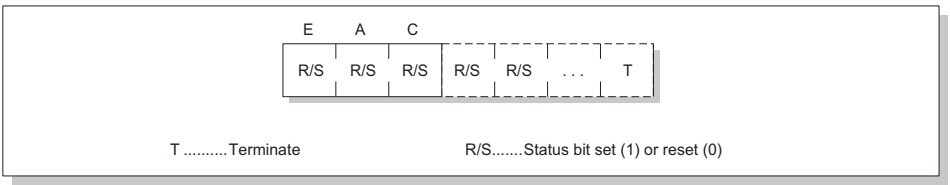

Figure 8.11 The frame status field

Abort Sequence

An abort sequence consists of a starting delimiter and an ending delimiter. This sequence can be transmitted anywhere in the bit stream without regard to byte alignment and stops all transmission.

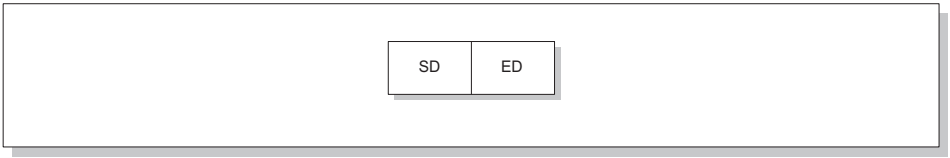

Figure 8.12 The abort sequence

8.1.2.2 MAC Layer Token-Ring Frames

MAC layer Token-Ring frames are used by the Token-Ring control protocol to conduct ring management operations, including error handling, in various operating states. The value in the control field of a MAC frame indicates one of three priority levels, which determine whether the packet is copied into the buffer of a receiving station. If this value is 00, the packet is copied into the receive buffer only if sufficient memory is available. Packets with 01 in the control field are always copied into the receive buffer, even if this means discarding existing data. All other MAC frames with a control field value greater than 01 are addressed to all stations, but are only copied into a station's receive buffer if enough memory is available. There are five different types of MAC frames.

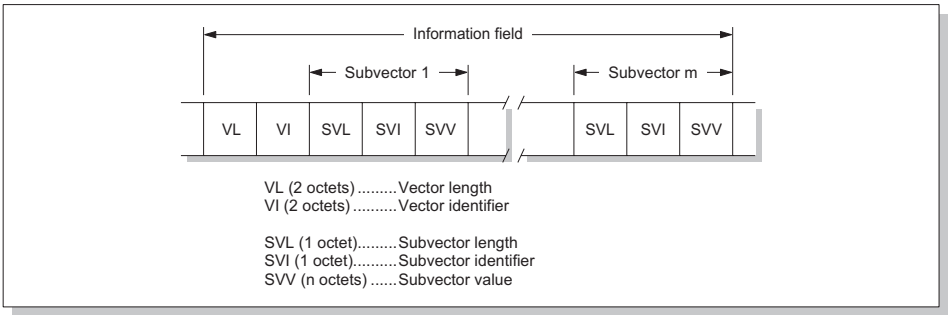

Figure 8.13 The information field of a MAC frame

Claim Token

If a station recognizes that there is no active monitor because the Timer Standby Monitor (TSM) has run out, it initiates a monitor contention process to select a new active monitor by sending a claim token frame. When a new active monitor

has been selected, the selected station will return the ring to operational status by issuing a new token.

The station initiating the process begins transmitting claim token frames while noting the source address of any claim token frame it receives. Other stations can choose whether or not to take part in the monitor contention process. A station that does not take part in this process simply repeats the claim token frames and starts its own claim token timer. Stations that actively participate in the claim token process compare the source address of each claim token frame received with their own address. If the value of the source address in the claim token frame is greater than that of the station's own address, the station withdraws from the process. If the station's address is higher, it replaces the source address of the claim token frame with its own address before passing the frame on. Once a station receives a claim token frame with its own address as the source, it becomes the active monitor. It clears all data circulating in the ring (that is, performs a ring purge) and issues a new token.

Priority	0	
Frame Control field	0000 0011	Claim token
Destination address	1111 1111 1111 1111	Broadcast
Info field VI	0003 hex	Claim token
Info field SVI-1	02 hex	Received upstream neighbor address
Info field SVV-1	XXXX ...	(The upstream neighbor's address)

Figure 8.14 The fields in a claim token MAC frame

Duplicate Address Test (DAT)

When a new station enters the ring it transmits a DAT frame. The new station sends a frame with its own address as the destination. If the address recognized bit is set to 1 when the frame returns, then another station already exists with the same address. In this case, the new station removes itself from the ring.

Priority	0	
Frame Control field	0000 0000	
Destination address		(= sender's address)
Info field VI	0007 hex	Duplicate address test

Figure 8.15 The fields in a duplicate address test frame

Active Monitor Present (AMP)

An AMP frame is transmitted by the active monitor after a ring purge, or when the Timer Active Monitor (TAM) times out, to inform the standby monitor stations of its presence. Each station that receives an AMP frame resets its TSM to 0.

Priority	7	
Frame Control field	0000 0101	
Destination address		Broadcast
Info field VI	0005 hex	Active monitor present
Info field SVI-1	02 hex	Received upstream neighbor address
Info field SVV-1	XXXX ...	(The upstream neighbor's address)

Figure 8.16 The fields in an active monitor present frame

Standby Monitor Present (SMP)

Standby monitors transmit SMP frames to announce their presence. (A standby monitor is any station in the ring that is not currently serving as active monitor.) When a station receives an SMP or AMP frame from its upstream neighbor with the A and C bits set to 0, it saves the source address as its current next active upstream neighbor (NAUN) and sets the A and C bits to 1. The station also starts its Timer Queue PDU (TQP) and transmits its own SMP frame when the timer expires. The TQP ensures that the process of circulating SMP and AMP frames—the neighbor notification or "ring poll" process—does not take up more than 1 percent of the ring's bandwidth.

Priority	0	
Frame Control field	0000 0110	
Destination address		Broadcast
Info field VI	0006 hex	Standby monitor present
Info field SVI-1	02 hex	Received upstream neighbor address
Info field SVV-1		(The upstream neighbor's address)

Figure 8.17 The fields in a standby monitor present frame

Beacon (BCN)

A BCN frame is transmitted when there is a problem in the ring that the Token-Ring protocol cannot solve otherwise. Such a problem may involve the cabling, a

network interface card (NIC), or a concentrator. The failure domain begins with the station directly upstream from the one transmitting the beacon frame. Its address is placed in the SVV-1 information field of the beacon frame.

Priority	0	
Frame Control field	0000 0010	
Destination address		Broadcast
Info field VI	0002 hex	Beacon
Info field SVI-1	02 hex	Received upstream neighbor address
Info field SVV-1		(The upstream neighbor's address)
Info field SVI-2	01 hex	(Reserved for future use)
Info field SVV-2	0002 hex	Continuous non-data J waveform; see physical layer
	0003 hex	TNT (Timer No Token) expired and no claim token frame detected
	0004 hex	TNT (Timer No Token) expired during claim token process

Figure 8.18 The fields in a beacon frame

Token-Ring Detailed Decode

Control | Config | Actions | Format | Other displays | Help

Frame: 2 Time: Sep 22@16:17:48.6817660 Length: 36

Field	Value	Description
Access Control:		
Frame Priority	000.-....	Non-Priority
Token Bit	...1-....	Frame
Monitor Bit	....-0...	Not Passed Monitor
Reservation	....-.011	Reserved
Frame Control:		
Frame Type	00..-....	MAC Frame
Reserved	..00-....	Reserved
Control	....-0100	Express Buffered
MAC Destination Address	Broadcast	Broadcast
MAC Source Address	4000F602661A	No source routing
Destination Class	0000-....	Ring Station
Source Class	....-0000	Ring Station
MAC Frame Type	04	Ring Purge
Subvector Type	00000000	Physical Location
Subvector Type	IBM---26E062	NAUN
Frame check sequence	6B-C4-D4-C7	
> Intermediate bit	....-..0.	Single Frame
> Error bit	....-...0	No Errors Detected
> AC bits		Recognized and Copied
> Data size	0	

Figure 8.19 Ring purge recorded with a protocol analyzer

Purge (PRG)

A PRG frame initializes the ring, deleting all data in it. A PRG frame is transmitted in the following situations:

- After a token claiming process, but before the new token is transmitted
- After the Timer, Valid Transmission (TVX) expires
- Following the appearance of a monitor bit set to 1 in an access control field

Priority	0	
Frame Control field	0000 0100	
Destination address		Broadcast
Info field VI	0004 hex	Purge
Info field SVI-1	02 hex	Received upstream neighbor address
Info field SVV-1		(The upstream neighbor's address)

Figure 8.20 The fields in a purge frame

8.1.2.3 Timers in Token Ring

All processes in the Token-Ring protocol are time-controlled. Analysis of timer activity can yield valuable information about the state of a ring. The main timers are described here.

Timer Return to Repeat (TRR)

Every station has a TRR. This timer ensures that the station returns from transmission to repeat state within a defined period. The timer setting must be higher than the transmission delay in the ring plus the sum of the delay times of all stations. The default setting for the TRR is 2.5 ms.

Timer Holding Token (THT)

The THT limits the time during which a station can transmit once it has received a token. The default setting is 10 ms.

Timer Queue PDU (TQP)

The TQP defines how long a station waits after receiving the upstream neighbor's AMP or SMP frame before it transmits its own SMP frame. This limits the overhead traffic on the ring due to the AMP/SMP process to less than 1 percent. The default value is 10 ms.

Timer Valid Transmission (TVX)

Every station has a TVX. A token error occurs when the active monitor does not receive a valid signal before its TVX expires. The setting for this timer is the sum of the settings for the THT and TRR. The default setting is 12.5 ms.

Timer No Token (TNT)

The TNT allows the ring to recover from a variety of token errors. The value for this timer is $n \cdot$ THT + TRR, where n equals the number of ring stations. The default setting for the TNT is 1 second.

Timer Active Monitor (TAM)

Every station has a TAM. When the active monitor's TAM expires, it transmits an AMP frame. The default setting is 3 seconds.

Timer Standby Monitor (TSM)

Every station has a TSM. This timer ensures that there is always an active monitor in the ring (see the previous description of the claim token frame). The default setting for this timer is 7 seconds.

8.1.3 Process Control in Token-Ring Networks

As described at the beginning of this chapter, access to the transmission medium in Token Ring is controlled by a token: a station must receive a token before it can transmit data over the ring. The data to be transmitted must also have a priority level that is equal to or higher than that of the token. If the token cannot be used, or has a higher priority level than the data packets to be transmitted by a given station, the station can request a lower priority token by entering the corresponding priority level in the token's RRR bit in the access control field. Once a suitable token is received, this token is transformed into a start of frame sequence by setting the token bit. The station then begins transmitting its own outgoing data rather than repeating received data.

8.1.3.1 Beaconing and Neighbor Notification

When persistent problems occur in the ring, the Token-Ring protocol is often able to recover or at least indicate the failure domain by means of the beacon and neighbor notification processes. Neighbor notification, or ring polling, is an essential process for automatic recovery because it ensures that all active stations and their relative positions are known at all times. The active monitor initiates neighbor notification at regular intervals by broadcasting an AMP frame. The nearest downstream neighbor reacts as follows when it receives this frame:

- It resets its TSM.
- It copies the AMP broadcast frame into its receive buffer and saves the NAUN address (or UNA).
- It sets the A and C bits of the AMP frame to 1 and forwards it.
- It transmits an SMP frame.

One after another, each station receives an SMP frame with the A and C bits set to 0, which indicates that the frame was sent by the NAUN. At the conclusion of the neighbor notification process, each station knows the address of its NAUN.

The beacon process allows the ring to recover after a failure. The station downstream from the location of the failure is the first to notice a ring error because it no longer receives valid signals. This node's first reaction to the absence of valid signals is to start a claim token (or monitor contention) process. If this does not succeed before the station's claim token timer expires, it enters the beacon state and begins broadcasting beacon frames. Beacon frames contain both the address of the station that generated them and that of its NAUN. No new station can enter the ring while beaconing is occurring. If the NAUN of the station that initiated the beacon process receives eight beacon frames that contain its address in the NAUN field, it assumes that it is the cause of the error, removes itself from the ring and performs a self-test. If no error is found, it re-enters the ring. After a defined period (usually 26 seconds), the station that initiated beaconing also removes itself from the ring, assuming that it might be the cause of the error, and performs a self-test. If this does not pinpoint the source of the error, the ring goes into the "streaming beacon" state. At this point, the ring can no longer recover automatically.

There are four types of beacon frames: Type 1 beacon frames are only used by stations that implement the Dual Ring protocol (IEEE 802.5c). Type 2 beacon frames indicate a complete loss of valid signals. Type 3 beacon frames indicate bit-streaming (the NAUN is transmitting a continuous stream of padding bits) and Type 4 beacon frames indicate "claim streaming" (the NAUN is continuously transmitting claim token frames).

8.1.3.2 Optional Token-Ring Services

In addition to the functions described previously, Token Ring can include a number of other functional elements to optimize management: these include the ring parameter server, ring error monitor, configuration report server and LAN bridge server. Ring stations report the required data to these services regardless of whether the corresponding components are present in the ring. In practice, this data is analyzed by special Token-Ring management software packages or monitored using a protocol analyzer. The ring parameter server saves all ring operating data, such as timer settings, active monitor stations and standby monitor stations. The ring error monitor keeps track of any errors that occur. The configuration report server informs network management services on the configuration of individual ring stations. And the LAN bridge server can be used to analyze the performance of bridges in the network.

8.1.4 Design Guidelines for Token-Ring Networks

The most important guidelines for designing Token-Ring networks are those governing the maximum number of stations (250) and the maximum distance between adjacent nodes on the ring. In 4 Mbit/s rings, this distance must not exceed 240 meters when using IBM type 1 cabling, or 100 meters when using Cat. 3 UTP cabling. The corresponding limitations for 16 Mbit/s rings are 100 meters for IBM type 1 and 45 meters for Cat. 3 UTP cabling.

8.1.4.1 Verifying Your Token-Ring Design

You can check your Token-Ring network against the tables 8.21 through 8.24 to make sure it conforms to the design guidelines. All you need to ascertain is the number of concentrators (MAUs), the number of wiring closets that contain the concentrators, and the cable lengths between the wiring closets. Begin by using these numbers to calculate the adjusted ring length (ARL): add the length of the cables between wiring closets and subtract the length of the shortest cable connecting two wiring closets. Tables 8.21 through 8.24 contain the values for

Token Ring (4 Mbit/s) over Type 1 and Type 2 cabling: Maximum ARL (in meters)
(When Type 6 or Type 9 cabling is used, divide values by 1.33.
For Type 8 cabling, divide by 2.)

Concentrators	Number of wiring closets										
	2	3	4	5	6	7	8	9	10	11	12
2	363										
3	354	350									
4	346	341	336								
5	337	332	328	323							
6	329	324	319	316	310						
7	320	315	311	306	301	297					
8	311	306	302	297	293	288	283				
9	302	298	293	289	284	279	274	270			
10	294	289	284	280	275	271	266	262	257		
11	285	280	276	271	266	262	257	253	248	244	
12	276	272	267	262	258	253	249	244	240	235	230
13	268	263	258	254	249	244	240	235	231	226	222
14	259	254	250	245	240	236	231	227	222	217	213
15	250	246	241	236	232	227	223	218	213	209	204

Figure 8.21 Token-Ring design table: 4 Mbit/s, Type 1, Type 2 cabling

Token Ring (16 Mbit/s) over Type 1 and Type 2 cabling: Maximum ARL (in meters) (When Type 6 or Type 9 cabling is used, divide values by 1.33. For Type 8 cabling, divide by 2.)									
Concentrators	Number of wiring closets								
	2	3	4	5	6	7	8	9	10
2	162								
3	155	150							
4	148	144	138						
5	142	137	132	127					
6	135	130	125	120	115				
7	129	123	19	113	109	103			
8	122	350	112	197	105	97	92		
9	115	110	105	100	95	90	85	80	
10	108	104	98	93.6	88	84	79	73	69
11	102	97	92	87	82	77	72	67	62
12	95	90	85	80	75	70	65	60	55
13	247	77	72	67	62	57	52	47	42
14	69	64	59	54	49	44	39	34	29
15	56	51	139	41	36	31	26	21	16

Figure 8.22 Token-Ring design table: 16 Mbit/s, Type 1, Type 2 cabling

Token Ring (4 Mbit/s) over UTP cabling				
Concentrators	Number of wiring closets			
	1	2	3	4
1	223			
2	217	206		
3	211	201	196	
4	205	195	190	185
5	199	189	184	179
6	194	183	178	173
7	188	178	173	167
8	182	172	167	162
9	176	166	161	156
10	170	160	160	150

Figure 8.23 Token-Ring design table: 4 Mbit/s, UTP cabling

the stations' maximum drive distances when using passive concentrators, which equal the maximum ARL plus the maximum lobe cable length. If active concentrators are used, the table values correspond directly to the maximum ARL.

Token Ring (16 Mbit/s) over UTP cabling				
Concentrators	Number of wiring closets			
	1	2	3	
1	55			
2	45	39		
3	35	29	23	
4	26	20		
5	16			

Figure 8.24 Token-Ring design table: 16 Mbit/s, UTP cabling

If a 16 Mbit/s Token Ring is made up of the three wiring closets A, B and C, of which A and B contain one concentrator each and C houses two concentrators, and the distances between them are AB = 34 meters, BC = 56 meters and CA = 64 meters, then the ARL calculation yields:

ARL = 34 + 56 + 64 - 34 = 120 meters

The corresponding field in table 8.22 (four concentrators, three wiring closets) provides the value for the maximum allowable ARL, which is 144 meters. If passive concentrators are used, the maximum allowable length of lobe cables is 144 - 120 = 24 meters. If the lobe cables do not exceed this length, the ring design is suitable for use with passive or active concentrators. If the lobe cables are longer, active concentrators must be used.

8.1.5 Token-Ring Standards

IEEE Std 802.5c-1991, Supplement to IEEE Std 802.5-1989, Local and Metropolitan Area Networks: Recommended Practice for Dual Ring Operation with Wrapback Reconfiguration.

IEEE 802.5e, Token Ring Station Management Entity Specifications.

IEEE Std 802.5j-1997, Supplement to Information Technology–Telecommunications and Information Exchange Between Systems–Local and Metropolitan Area Networks–Specific Requirements–Part 5: Token Ring Access Method and Physical Layer Specifications–Fiber Optic Media Requirements.

IEEE Std 802.5r-1997, Standard for Information Technology–Telecommunications and Information Exchange Between Systems–Local and Metropolitan Area Networks–Specific Requirements–Part 5: Token Ring Access Method and Physical Layer Specifications–Dedicated Token Ring Operation.

IEEE 802.5t, Supplement to ISO/IEC 8802-5: 1995 Specific Requirements–Part 5: Token Ring Access Method and Physical Layer Specifications–100 Mbit/s Dedicated Token Ring Operation Over 2-Pair Cabling.

IEEE 802.5u, Supplement to ISO/IEC 8802-5: 1995 Specific Requirements–Part 5: Token Ring Access Method and Physical Layer Specifications–100 Mbit/s Dedicated Token Ring Operation Over Multi-mode Fiber.

IEEE 802.5v, Supplement to ISO/IEC 8802-5: 1995, Specific Requirements–Part 5: Token Ring Access Method and Physical Layer Specifications. Media Access Control Parameters, Physical Layers, and Management Parameters for 1000 Mbit/s Operation or Above.

RFC 1231 IEEE 802.5 Token Ring MIB.

The IEEE 802.5 Working Group can be reached on the Internet at

http://www.8025.org/802.5/documents/

8.2 Troubleshooting in Token-Ring Networks

8.2.1 Gathering Information on Symptoms and Recent Changes

The first step in any troubleshooting process is to gather information. The more information you have about the symptoms and characteristics of a problem–including *when* it first occurred–the better your chances of solving the problem quickly and efficiently. Typical questions you might ask at this stage include:

- Do the symptoms occur regularly or intermittently?
- Are the symptoms related to certain applications, or do they affect all network operations?
- Do the symptoms correlate to other activities in the network?
- When was the first occurrence of the symptom?
- Was there any change in any hardware or software network component?
- Has anyone connected or disconnected a PC (laptop or desktop) or any other component to or from the network?
- Has anyone installed an interface card in a computer?
- Has anyone stepped on a cable?

- Has any maintenance work been performed in the building recently (by a telephone company or building maintenance personnel, for example)?
- Has anyone (including cleaning personnel) moved any equipment or furniture?

8.2.2 Starting the Troubleshooting Procedure

Troubleshooting in Token-Ring LANs is primarily performed using protocol analyzers, network management software and cable testers, and special Token-Ring management software that tracks and displays Token-Ring operating messages. Unlike Ethernet, Token Ring is second-generation network technology and has a number of self-diagnosis functions that enable it to resolve certain critical operation states on its own. Key requirements for successful troubleshooting in Token-Ring networks include a detailed understanding of its operational processes and the availability of a protocol analyzer and specialized Token-Ring management software to monitor these processes.

The first step in the troubleshooting procedure involves using a protocol analyzer to determine the main operating statistics of the network. These statistics

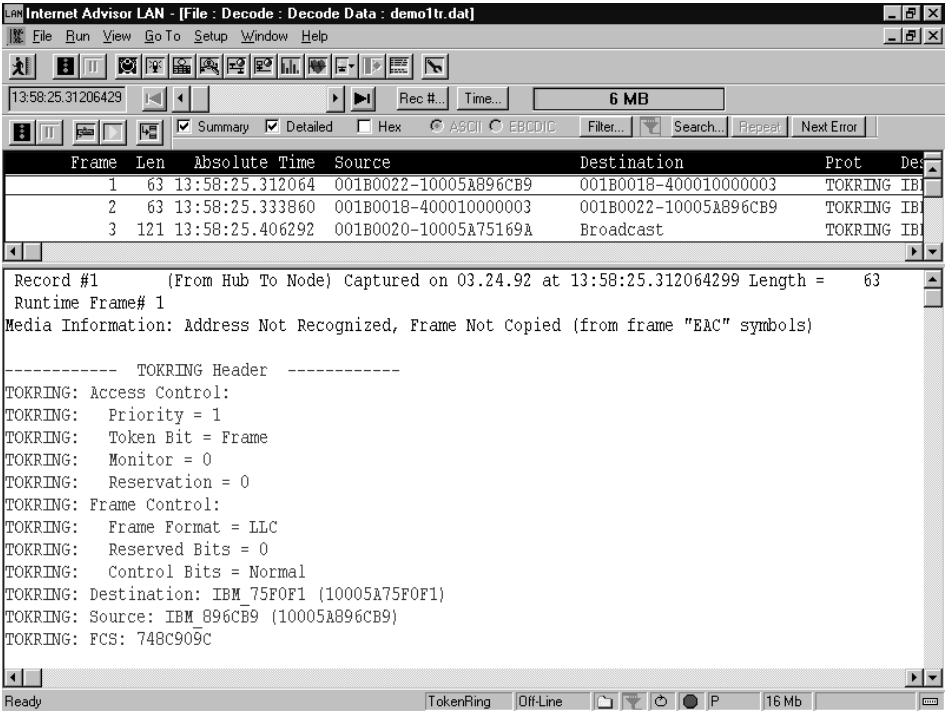

Figure 8.25 Display of characteristic Token-Ring operating data on the Agilent Advisor

include ring capacity use as a percentage; throughput in frames per second; the number of receiver congestion messages, burst errors, line errors, beacons, monitor contentions and ring purges; and the number of transmitting stations. The analysis of these statistics often points to possible causes of the problem. Unlike the defective frames found in Ethernet, the various Token-Ring MAC frames that report and handle errors generally give a fairly precise indication of how the errors came about. If the problem cannot be located in this way, however, additional trend measurements are necessary. This involves recording the main operating parameters over a period of hours, or even days, and analyzing the results for correlations. By charting the network load together with the number of active nodes and the number of ring purges, for example, you can tell at a glance whether there is a correlation between the occurrence of ring purges and stations entering or leaving the ring. If this is the case, these ring purges are probably normal operating behavior. Otherwise, defective NICs or concentrator ports may be the cause. Measurement results from different network segments connected by bridges can be similarly correlated. In this way, possible causes can be systematically eliminated until the source of the problem is limited to a small area.

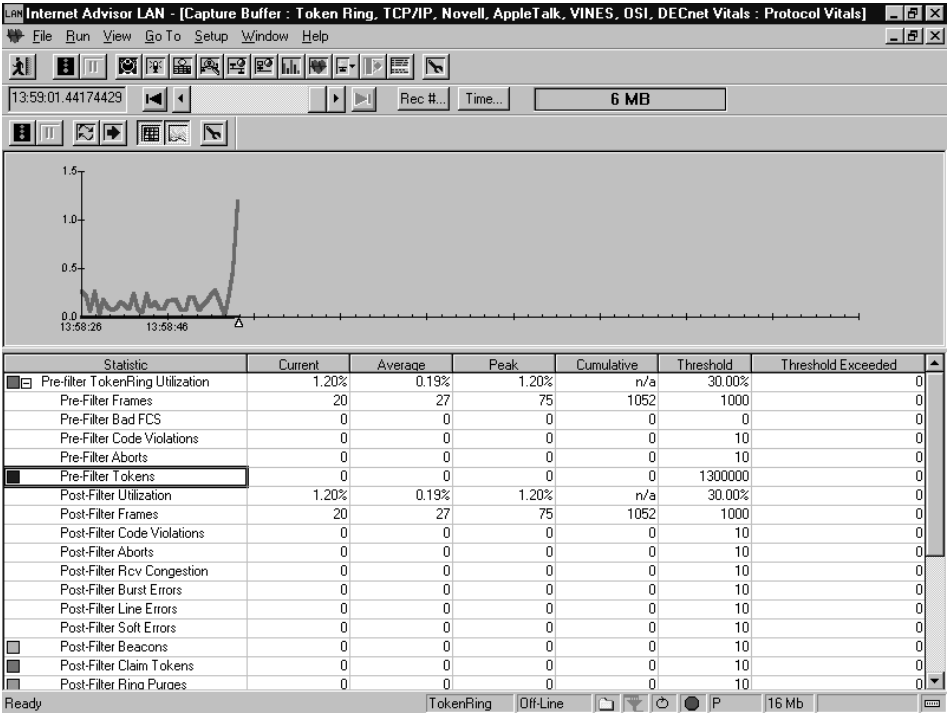

Figure 8.26 Trend measurements and correlation analysis in Token-Ring networks

The steps to take after a protocol analyzer performed the basic measurements depend on the nature of the symptoms. If the symptoms can be localized, occur periodically, or can at least be reproduced, then the troubleshooting process continues with the network component nearest to the problem. If the problem source cannot be detected there, the range of analysis is successively expanded. For example, if the problems are found to be related to a single network node, the next step is to analyze the station's software and hardware components. If no fault is found, the examination progresses to the lobe cable, the connector, the wall jack, the cable to the concentrator, the concentrator itself, the cable to the server, and so on.

If the problem cannot be localized at all, or if problems that were thought to have been localized cannot be pinpointed, the only way to find the source of the problem is through systematic segmentation of the network. To do this, check each of the concentrators in turn and replace any defective units found. If the error persists, physically remove each concentrator in turn until the source of the problem becomes apparent. This method causes considerable disruption in network operation, and is therefore applied only as a last resort, when the problem itself severely impairs normal network operation.

If the symptoms occur intermittently, long-term measurements are necessary. These must be performed continuously until the basic network operating statistics have been measured during the occurrence of the fault. This information usually provides the first clue to the error source. Furthermore, it is essential to log the exact time of intermittent error events. Later this information can be used to find temporal correlations with other events in the network or on a given node, such as backups, the start of specific applications, connections through routers, access to the Internet, users' working hours, or other possible factors. If this does not help to track down the error, you may have to resort to the segmentation method. Depending on which causes the least inconvenience to users, you can either systematically disable network functions and applications, or physically separate concentrators. These methods usually lead to the error source.

8.2.3 Error Symptoms in Token Ring

The Token-Ring protocol has built-in operation management functions that generate highly specific messages in most cases of error. Errors in Token Ring can be divided into two categories: isolating errors, which can be traced to a specific node, and non-isolating errors. Error messages that report isolating errors contain both the address of the station reporting the error and the address of its upstream neighbor; this information points to the probable failure

domain. Non-isolating errors, by contrast, such as token errors or lost frames, do not provide information about the location of the failure. The following error types are reported by the Token-Ring protocol.

8.2.3.1 Error Types

Abort Error (Isolating)

Token-Ring nodes transmit abort error messages if they have recovered from an internal operating error, if recovery from an internal error failed, or after having received a corrupt token.

A/C Error (Isolating)

An A/C error is reported if a station receives more than one AMP frame, or if it receives an SMP frame without having received an AMP frame to initiate a neighbor notification process. This is often caused by problems in the upstream neighbor or by errors involving bridges or routers.

Beacon Error (Token-Ring Protocol)

Beacons are transmitted if a station is no longer able to receive frames of any type, and its attempt to initiate a claim token process has failed. Common causes are defective concentrators, interface cards or cabling.

Burst Error (Isolating)

Burst errors are reported if a station does not receive a valid signal for more than five half-bit times between a starting delimiter and an ending delimiter.

Claim Token (Token-Ring Protocol)

Claim token frames are circulated in the monitor contention process to select a new active monitor. There are a number of reasons why such a frame may be transmitted, including the following:

- A standby monitor is not receiving valid signals
- A standby monitor is not receiving tokens
- A standby monitor has noticed the lack of a neighbor notification process
- The active monitor is unable to perform a ring purge
- A new station is inserted into the ring
- A station detects a clock error
- A beaconing station receives its own beacon frame
- A station receives a given frame a second time
- A station receives beacon frames for an extended period

If the monitor contention process cannot be completed successfully, the ring goes into a beaconing state.

Frame Copy Error (Non-Isolating)

If a station receives a frame in which its own address is given as the destination but the address recognized bit is already set, that station reports a frame copy error. This error is often caused by a duplicate Token-Ring address in the ring.

Frequency Error (Non-Isolating)

A frequency error is reported if the clock rate of the signal received exceeds the frequency tolerance limits.

Internal Error (Isolating)

A Token-Ring node reports internal errors after recovering from an internal operating error.

Line Error (Isolating)

A line error is reported if a station detects frames that contain invalid checksums or coding errors.

Lost Frame (Non-Isolating)

If a station's transmitted frame does not return to it, the station reports a lost frame error.

Purge (Token-Ring Protocol)

An active monitor transmits a purge frame if it does not receive a valid token before its TNT expires. If the purge frame returns to the active monitor successfully, the ring is considered operational and a new token is released.

Receiver Congestion Error (Non-Isolating)

If a destination node is not able to copy a frame into its receive buffer due to memory overflow, it reports a receive congestion error. So long as such errors do not originate from bridges or routers, they are of minor importance and have little impact on ring performance.

Token Error (Non-Isolating)

A token error is reported if a token is detected with a monitor bit set to 1, if a corrupt token is received, or if a TNT timeout occurs.

8.2.3.2 Principle Error Conditions During Normal Ring Operation

As in Ethernet, there are error conditions in Token Ring that occur during normal operation and do not present a problem so long as they do not exceed certain levels. The following events trigger the error messages listed:

Station insertion:

Active monitor:	Ring purge
Inserted station:	2 Duplicate Address tests, 1 Report NAUN frame
Downstream neighbor:	1 Report NAUN frame
Inserted station:	2-4 Request Initialization frames
Active monitor:	Report Soft Error frame (1-4 lost tokens)
Downstream neighbor:	Report Soft Error frame (2-4 burst errors)
Other stations:	Report Soft Error frame (1 lost frame)

Station removal:

Active monitor:	Ring purge
Downstream neighbor:	1 Report NAUN frame
Active monitor:	1 Report Soft Error frame (1-4 lost tokens)
Downstream neighbor:	2-4 Burst errors
Other stations:	Report Soft Error frame (approximately 1 lost frame)

Station shutdown:

Downstream neighbor:	4 Claim Token frames
New active monitor:	Ring purge
Downstream neighbor:	Report Soft Error frame (reports loss of the old active monitor; 1-4 burst errors)
New active monitor:	Report New Monitor frame
Downstream neighbor of the former active monitor:	Report NAUN frame

8.2.3.3 Critical Error States Requiring Investigation

Operational states that indicate possible problems on the ring may be indicated by:

- Ring purges that are not related to station insertion or removal
- Incomplete neighbor notification processes
- Report active monitor errors
- Change of active monitor
- Beaconing or streaming beaconing

8.2.4 Cabling Problems

As in other networks, cabling problems are among the most common causes of errors in Token-Ring networks. Typical causes include defective or low-quality cables; incorrect characteristic impedance; and wiring mistakes or electromagnetic interference (noise) caused by air conditioning systems, photocopiers, pagers, elevators or production environments. These problems are discussed in detail in the chapter on cabling. One factor that must be mentioned with specific reference to Token Ring, however, is the limit on the distance between two active ring stations (see the section on design guidelines). If this limit is not observed, switching off ring stations may temporarily result in node distances that exceed these limits. Consequently, the Token-Ring signals may no longer be transported reliably between two nodes; this can lead to extraordinarily high error rates or even a complete failure of the ring. Another common source of problems is upgrading 4 Mbit/s rings to 16 Mbit/s. If you do not take into account the fact that the maximum distance allowed between two nodes on 16 Mbit/s rings is significantly less than that specified for 4 Mbit/s rings, the distance limitations may be exceeded.

8.2.5 Problems with Token-Ring Interface Cards

The first step in localizing a defective NIC is to identify suspicious nodes on the network. Begin by making a list of all network nodes that transmit defective frames. Most protocol analyzers provide this information with fully automatic test programs. If the source addresses of the defective frames are invalid and cannot be decoded, try the correlation method: begin by simultaneously charting the activity of the suspicious nodes and the error rate in the network. If you observe a correlation between the activity of a certain node and the error frequency, then you have probably found the defective interface card. Monitoring the states of the ring stations (active monitor, standby monitor) in correlation with error frames can provide additional information about the source of the problem. For example, if ring errors frequently occur when one particular station is the active monitor, there is a good chance that the station's NIC is defective.

Symptoms of Defective Token-Ring Interface Cards

Characteristic symptoms of defective Token-Ring interface cards include the occurrence of claim token frames in the absence of any ring activity, beacon frames, or receiver congestion errors. Common causes are station configuration errors, problems with power to the interface card, or a hardware failure on the interface card.

Defective Interface Card Hardware

If a defective interface card causes the concentrator relay to remain in the closed position, the ring goes into a beaconing state and cannot recover on its own.

Duplicate Token-Ring Addresses

When a new station is inserted into the ring, it may detect that it has the same address as another active node. In this case, the new station cancels the insertion process and removes itself from the ring. Duplicate addresses can result from typing errors during station configuration, from copying configuration files between stations, or from cloning a station (that is, copying a one-to-one disk drive image from one node to another).

Incorrect Ring Speed

An interface card or a bridge or router port configured for the wrong ring speed (for example, 4 Mbit/s instead of 16 Mbit/s) also cause beaconing.

8.2.6 Problems with Concentrators (MAUs/TCUs)

The nodes are connected to the ring by concentrators. Smooth functioning of the concentrators is, therefore, a key requirement for error-free ring operation. One defective concentrator can often cause complete failure of the entire ring. The following equipment is required for troubleshooting in concentrators:

- A backup concentrator known to be in working order
- A lobe cable in working order
- A MAU port reset connector (for resetting stuck concentrator ports)
- A Token-Ring mini-network (see Figure 8.27)

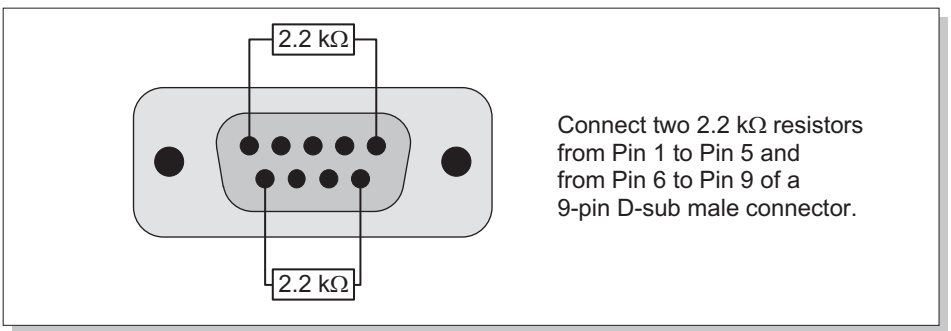

Figure 8.27 Mini-network for Token-Ring simulation

Symptoms of Defective Concentrators

The most common symptom of trunk concentrator unit (TCU) problems is the sporadic or continuous occurrence of beacon frames. This symptom is often caused by stuck TCU ports, defective lobe cables, loose or defective connectors, or general TCU hardware problems. To locate the source of the problem, use a protocol analyzer to capture and analyze the beacon frames. The addresses in the beacon frames, those of the beacon sender and its upstream neighbor, provide the information needed to locate the failure domain. The exact error source can be determined by systematically replacing the components in the failure domain with components known to be in working order (lobe cable, connector, interface card, TCU) and by testing components in the Token-Ring mini-network.

8.2.7 Problems with Bridges

Bridges are inter-network elements that connect network segments on OSI Layer 2 (the MAC layer). Bridges buffer and filter the frames they receive from connected segments and transmit them to their destination segments without regard to higher-layer protocols. The main functions of a MAC-layer bridge are preventing the spread of local traffic to neighboring segments and overcoming such limitations of the particular network topology as the maximum number of nodes per segment, the transmission delay or the distance limits.

8.2.7.1 Token-Ring Bridges

The basic functions of bridges are described in detail in the chapter on Ethernet networks; the following discussion deals only with the source-routing bridges developed specifically for Token-Ring topologies.

Source-Routing Bridges

Unlike the spanning-tree algorithm used by Ethernet bridges, the source-routing algorithm used to select routes in Token-Ring topologies lets the transmitting node rather than bridges determine the transmission path. The actual route information is contained in the optional routing information field of Token-Ring frames where route control information and route designators describe the exact route. Before actual communication occurs between two nodes in separate rings linked by one or more source-routing bridges, an explorer frame is sent to determine the optimum route. The explorer frame is duplicated by each bridge and copied into all linked rings until one of the explorers reaches the destination. Each crossed bridge writes its identification into the explorer's route designator field so that the entire transmission path is stored in the explorer frame on arrival. If more than one explorer frame reaches the destination, the

one that contains the shortest route is selected. The maximum allowable hop count is seven. The destination node then copies the route information from the explorer frame and returns it to the sender, which then starts the actual data transmission. The source-routing protocol is defined in the source-routing annex of IEEE 802.1 and is not part of the IEEE 802.5 standard. Today there are bridges available, known as source-routing transparent (SRT) bridges, that can link both source-routing and spanning-tree networks to one another.

Linking Token-Ring Networks to Other LAN Topologies

When linking Token-Ring networks to other network topologies, such as Ethernet or FDDI, the differences in data speeds, frame formats and access mechanisms can present a number of difficulties that necessitate certain special capabilities in the bridges used. These problems are discussed in detail in the section on Ethernet bridges.

8.2.7.2 Diagnosing Bridge Problems

The challenge when analyzing bridge problems is to correlate the occurrence of symptoms in several different network segments. Concurrent measurements in several segments—using probe-based monitoring systems, for example—can be very helpful. Less important are performance measurements of bridges by means of expensive, specialized multiport test systems. Most modern bridges are capable of forwarding frames at line speed anyway so that performance measurements in most cases just confirm the manufacturer's technical data. It is more efficient to request system specifications from the manufacturer based on standardized test methods, as specified in RFC 1242 and RFC 2544.

Most problems that affect bridges can best be located by a process of elimination that involves the correlation of specific measurements and an analysis of the network topology. Symptoms of bridge problems can include poor network performance in particular segments, intermittent or permanent loss of connection to particular stations, or the failure of certain protocols and services. The first phase of the troubleshooting process is, as always, a review of all configuration changes that were made in the network before the error occurred as well as the general information-gathering steps described previously. If the symptoms correlate to particular connections, begin by checking all bridges located along the corresponding transmission path. Otherwise, the next step is to prepare a list of all the stations, connections, protocols and services affected by the problems observed. To do this, measure the current parameters in the various network segments and compare the results with statistics gathered during normal operation. This involves recording and analyzing throughput and performance parameters of network nodes, protocols and services, as well as reviewing log files that contain the operating statistics on all bridges in the network.

The log files provide bridge statistics such as CPU capacity use, port capacity use, buffer capacity use and error rates. To measure the response times of connections across bridges, send loopback packets across the bridges from different network segments. Long-term response time measurement statistics can be gathered using dedicated response-time agents distributed throughout the network. This type of long-term measurement can be especially useful in diagnosing intermittent problems. Based on the results of these measurements, the potential sources of error can usually be narrowed down to specific components.

8.2.7.3 Symptoms and Causes of Bridge Problems

The symptoms for most bridge problems in Token-Ring networks differ only slightly from those in Ethernet or FDDI networks. As described in the section on Ethernet bridges, the most common difficulties are throughput problems, incorrectly configured filter settings, bridge buffer overflow, and faulty address tables. Problem characteristics of Token-Ring networks include incorrect ring speed settings, bridge ports configured with duplicate Token-Ring addresses, and incorrect frame length settings.

Incorrect Ring Speed

If the ring speed setting on the bridge port does not match the actual speed on the ring (4 Mbit/s versus 16 Mbit/s), beaconing begins and communication on the ring breaks down.

Bridge Port Configured with Duplicate Token-Ring Address

Because Token-Ring addresses are configured by software, the occurrence of duplicate Token-Ring addresses due to incorrect configuration (typing errors, copied configuration files) is not uncommon.

Inefficient Maximum Frame Length

Incorrectly configured bridge ports that restrict the maximum frame size can have a negative effect on performance. The throughput in 4 Mbit/s rings, for example, decreases significantly when the maximum frame size is under 256 bytes.

Installation and Configuration Errors

Among the leading causes of problems with bridges are incorrect installation or configuration of the equipment. Incorrectly configured ports (port not enabled; wrong operating mode, for example, 4 Mbit/s instead of 16 Mbit/s), bad connections (loose cables, connectors, or plug-in modules), and faulty connections to the back plane or the MAU are the most common error sources.

Hardware Problems

If you suspect hardware problems, check the power supply and connectors and run the bridge's self-test function.

8.2.8 Problems with Routers

Routers are internetworking components that connect network segments on OSI Layer 3. Because they operate on this layer, routers can link networks of any topology. Refer to the section on router problems in Chapter 7 for a detailed description of procedures for troubleshooting and diagnosing router errors.

8.2.9 Symptoms and Causes: Token Ring

Symptom: Active Monitor Error or Active Monitor Change

Cause (1): Active monitor detects a claim token frame, quits active monitor status, and sends a report active monitor frame.
This generally occurs when a standby monitor does not detect an active monitor in the ring. (Any station in the ring that is not the active monitor is a standby monitor.)

Cause (2): Active monitor detects an AMP frame that it did not generate
When this happens, the active monitor transmits a report active monitor frame with subvector 2 (duplicate monitor).

Cause (3): Station participating in the monitor contention process detects a claim token frame with its own address as the source but a NAUN address that does not match the NAUN address in its memory.
This station then transmits a report monitor error frame with subvector 3 (duplicate address during monitor contention).

Symptom: Address Recognized Error

Cause: Station detects more than one AMP frame or an SMP frame not preceded by an AMP frame.

Symptom: Burst Errors

Cause: Hardware problem such as a defective cable, NIC, MAU or concentrator.
A burst error frame is sent if no signal is received for five half-bit times between the starting and ending delimiters of a frame. Decode Token-Ring messages to locate the fault: determine which station reported the error and what stations are upstream from it (refer to the list of active stations). Analyze correlations between station activity and errors in the failure domain. Check the con-

centrator (run its self-test function). Check cables using a cable scanner.

Symptom: Beaconing, Streaming

Cause (1): Defective concentrator or NIC.

Cause (2): Loose or defective connectors (interface cards, wall jacks, concentrators, bridges, routers).
Analyze the beacon frames and trace the failure domain from the addresses for the sending station and its NAUN. The failure domain consists of the station transmitting the beacon frame and its incoming line, the sending station's NAUN and its outgoing line, and the concentrator between the two stations. All components within this domain (NICs, concentrators, cables, connectors, wall jacks) need to be inspected.

Symptom: Failed Insertion

Cause (1): Duplicate address.
During the duplicate address check (part of the station insertion process), the new station detects another station already in the ring with the same address.

Cause (2): Station unable to participate successfully in the neighbor notification process.

Cause (3): Station parameters not initialized correctly.

Symptom: Frame Copy Error

Cause: Station receives frame addressed to it, but detects that the address recognized/frame copied bits are not 0.
One likely reason for this is a duplicate MAC address in the ring. To locate another station with the same address, use a protocol analyzer and check for a failed insertion frame. Once you have identified the node with the duplicate MAC address, reconfigure it.

Symptom: Lost Frame Error

Cause: Failure to receive a transmitted frame.
This can happen when other stations enter or leave the ring.
This error is non-isolating and can't be assigned to any particular station.

Symptom: Frequency Error

Cause (1): Ring clock rate and NIC's internal clock rate differ significantly.

Cause (2): Poor cabling.

Cause (3): Defective NIC.
Frequency errors are non-isolating and can't be assigned to any particular station. Typical causes of frequency errors are poor-quality cabling, cabling that exceeds distance limitations, or defective NICs.

Symptom: Intermittent Errors and Connection Failures

Cause (1): Cabling exceeds the distance limitations between two ring stations.
If a station is removed from the ring, the distance between two ring nodes can become so great that the signals can no longer be transmitted reliably and serious connection problems can occur (non-isolating errors, token errors, etc.). Check the maximum allowable distance between two stations on the ring and redesign the ring if necessary.

Cause (2): Phase jitter, frequency errors, timeouts in Token-Ring protocol timers, or intermittent beaconing (see Figure 8.28).
Verify whether the maximum number of stations allowed in the ring has been exceeded (see the section "Network Design Guidelines for Token-Ring Networks" for details).

Cable type	Maximum number of nodes at 4 Mbit/s	Maximum number of nodes at 16 Mbit/s
IBM Type 1	260	140
Cat. 3 UTP	72	72
Cat. 5 UTP	132	132

Fig. 8.28 Maximum number of stations in a Token-Ring network

Symptom: Internal Error

Cause: Station detects an internal error and recovers on its own.
Internal errors are isolating errors and can be traced to the station where they originate. Capture and decode the internal error frame using a protocol analyzer and observe the node identified.

Symptom: No Connection to Server

Cause (1): Cable from the node to the concentrator is loose or disconnected, broken, short-circuited, or exposed to electromagnetic interference.

Cause (2): Defective network interface card.
Check cable, connectors and interface card and replace if necessary.

Cause (3): Address table of a bridge in the transmission path to the server missing the node's MAC address.
Addresses that are not used over a certain period are deleted by the bridge's aging function. If the bridge is in protected mode (that is, learning is deactivated), the transmitting node's address cannot be automatically added to the bridge table. Check the address tables and operating modes of the bridges in the transmission path to the server.

Cause (4): Bridge port deactivated or defective.
Check bridge ports, send ping packets to nodes beyond the bridge, and analyze the bridge logs.

Cause (5): Incorrectly configured bridge filter.
Examine the filter settings in bridges along the transmission path to the server.

Symptom: Intermittent Connection Failures

Cause: Duplicate MAC address.
If a station attempts to enter the ring with an address that is already in use, it is refused entry and receives a request to remove frame. To locate the other station with the same address, use a protocol analyzer to capture request to remove frames. When you have identified the node with the duplicate MAC address, reconfigure it.

Symptom: High Network Load

Cause: Overloaded or incorrectly configured router(s) and/or bridge(s).
Use a protocol analyzer to identify the most active stations in the ring and search for routing or bridging problems. If timeouts occur, measuring response times can provide clues to the source of the problem. Check the statistics on the routers and bridges involved. How many frames are discarded? Check the bridges' forwarding tables and filter settings. Deactivate optional bridge functions, such as the ring parameter monitor or configuration port server if they are not in use.

Symptom: Network Slow, Stations Locking Up

Cause: Line errors, burst errors, FCS errors, and superfluous ring purges. Burst and line errors are usually caused by defective station cables or hardware defects in the concentrator or the interface card. Check the network for line errors and burst errors. Then check the concentrators, cabling and connectors upstream from the station reporting the error.

Symptom: Neighbor Notification Error

Cause (1): Insertion or removal of a node.

Cause (2): Intermittent hardware problems in a NIC.

Symptom: Report Neighbor Notification Incomplete

Cause: Active monitor sends process incomplete frame to the ring error monitor and initiates a new AMP frame.

The neighbor notification process is initiated every 7 seconds, when the active monitor sends an AMP frame. When a station detects an AMP frame, it compares the address recognized bits and the frame copied bits in the AMP frame. If the frame has not yet been copied by any other station, the receiving station compares the source address of the AMP frame with its own NAUN address. If the addresses are different, the source address of the AMP frame is stored as the new NAUN address, and a report NAUN change frame is sent to the configuration report server. If the AMP frame is not returned to the active monitor before the neighbor notification timer expires, the active monitor sends a process incomplete frame to the ring error monitor and initiates a new AMP frame.

Use a protocol analyzer to check for request to remove frames and try to identify the stations with duplicate MAC addresses.

Symptom: Network Slow Despite Low Traffic

Cause (1): Poor configuration, inefficient protocols, or insufficient NIC memory.

Cause (2): Router or bridge port settings restrict the maximum allowable frame size.

The network load (as a percentage of its capacity) is not the only factor determining network performance. Other important factors include the size and type of frames being transported. LLC frames, for example, carry no user data but serve to set up and maintain connections. A high proportion of short LLC and MAC frames indicates an inefficient protocol. In the NetBIOS/SMB protocol, for example, the ratio of LLC frames to user data packets is about 1:1. The reason for this exceptionally poor ratio is that NetBIOS/SMB uses a connection-oriented protocol at the LLC level. NetWare IPXuses, a connectionless service, transfers user data without waiting for acknowledgement of receipt. LLC frames are rare in IPX, whereas connection-oriented protocols usually generate a huge number of management frames.

Small frame sizes can also have other causes, however. The data packet size that can be handled by a NIC depends on the card's memory. In a 4 Mbit/s ring the maximum frame size is 4,500 bytes, and in a 16 Mbit/s ring 17,800 bytes. Older NICs with 8 Kbytes of RAM can only process data packets of up to 1,000 bytes. State-of-the-art cards, however, usually support the maximum frame lengths of 4,500 and 17,800 bytes. Furthermore, certain network operating systems can restrict the maximum frame size. NetWare 3.11, for example, supports frames only up to 4,000 bytes.

Symptom: Ring Purges

Cause (1): Short-circuited cable.
Cause (2): Noise or crosstalk.
Cause (3): Token rotation time too long.
Cause (4): Defective NIC.

Ring purges are initiated by the active monitor to delete all signals on the ring in preparation for the release of a new token. They frequently occur when a station enters or leaves the ring. If ring purges occur when no station has been inserted or removed, this indicates hardware problems on the ring.

Symptom: Ring Resetting

Cause: Several consecutive claim token frames transmitted; ring recovers after beaconing.

See the previous section on beaconing as an error symptom.

Symptom: Receiver Congested

Cause: Insufficient buffer space to copy a frame.

If this error occurs frequently, you must replace or upgrade the interface cards of the affected nodes to increase card memory.

Symptom: Token Error

Cause (1): Station entering or leaving the ring.
Cause (2): Noise.
Cause (3): Defective NIC or cable.
Cause (4): Extremely high number of broadcasts.

A Token Error frame is transmitted in any of the following situations:

- A token with a priority greater than 0 and a monitor count of 1 is detected beyond the active monitor (indicating that the token is already on its second round).
- No token or frame is encountered before the Good Token timer expires (10 ms).
- Illegal coding is detected.
- Token errors are non-isolating and can't be assigned to any particular station.

Symptom: Request Station Removed

Cause: Duplicate Token-Ring MAC address.

Use a protocol analyzer to capture request to remove MAC frames and examine their source addresses.

Symptom: Token Direction Change

Cause (1): Insertion or removal of stations.
Cause (2): Problems with hardware or software components in the ring.

To determine the direction of rotation, use a protocol analyzer to analyze frames transmitted by the station that is the direct (physical) neighbor of the analyzer, making sure that neither the protocol analyzer nor the neighboring node is the active monitor at the time. If the monitor bit of these frames is set to 1, then the frames are moving from the analyzer to the selected neighbor node. If the value is 0, the frames are moving in the other direction. If the direction of token rotation changes, this indicates that at least one ring purge has occurred. This can be caused by normal operating events, such as the insertion or removal of stations, or by problems with hardware or software components in the ring.

Common Errors

The following list summarizes the most frequent sources of problems in Token-Ring networks (in alphabetical order):

- Bridge address list incorrectly configured; bridge in protected mode
- Bridge filter incorrectly configured
- Bridge overloaded
- Bridge's aging function deletes address entry
- Cable length between neighboring nodes exceeds specifications
- Connectors, loose or defective: interface cards, wall jacks, concentrators, bridges, routers
- Defective Trunk Concentrator Unit (TCU)
- Defective lobe cable
- Defective network interface card
- Duplicate MAC addresses
- Electromagnetic interference
- Faulty physical installation of router, bridge or concentrator (cable, connectors, plug-ins are loose; cable connections on the backplane are wrong)
- Frame length restrictions on router/bridge ports
- Frequency and jitter problems due to cabling, noise, too many stations
- Maximum frame length not supported by interface cards due to insufficient card memory
- NIC incorrectly configured
- Protocol inefficient, not well adapted to Token Ring (NetBIOS/SMB)
- Receive buffer on interface card insufficient
- Ring speed incorrectly set on bridge/router port: for example, 4 Mbit/s vs. 16 Mbit/s
- Router filter incorrectly configured
- Router overloaded
- Router protocol entries incorrectly configured (address tables, mapping tables, subnet masks, default gateways, routing tables, timers)
- Router settings incorrectly configured: port not active, protocol not active
- Short circuit in cable
- Source-routing problems
- Stations: too many on the ring
- WAN connections overloaded or of poor quality (high BER)

Figure 8.29 The most frequent sources of problems in Token-Ring networks

Index of chapter 8

M

N

P

R

S

T