Apple Token Ring Cards and Token Ring Timers (3/95)

Article Created: 24 May 1993
Article Reviewed/Updated: 29 March 1995

TOPIC ---

I'm running into a Token Ring timer problem with our Apple 4/16 Token Ring cards. I believe, for my environment, I have them set too low.

Here are my questions:

1) Can both the Apple Token Ring 4/16 NB card and TokenTalk card timers be adjusted? I believe the older TokenTalk card timer can't be adjusted.

2) If the TokenTalk timer can't be adjusted can you forward me what the default T timers are for the 4 MB token ring cards?

3) The defaults for the 4/16 cards are:

	Group 1	Group 2
T1(Response)	0.20	1.00
T2(Ack)	0.04	0.40
TI(Inactivity)	1.00	5.00

IBM terminology refers to T timer values for LOCAL and REMOTE (for example bridged) LANs. Does this correspond to GROUP1 and GROUP2?

DISCUSSION ---

1 & 2) The Group 1 and Group 2 timers are adjustable on both the 4/16 Token Ring and TokenTalk token ring cards if the TT 2.4.x software is being used.

NOTE: There is no way to adjust timer values with the newer Apple Token Ring NB/c card because this card does not use the Token Ring control panel.

The Token Ring Control panel will display timers with either of the supported cards. The old TT 2.2 did not offer this feature. The older MacDFT product allowed modification of these timers to account for routing delays in Token Ring nets for SNA traffic.

3) The Group 1 timers apply to link layer traffic which is passing on the local ring. The Group 2 timers apply to link layer traffic which is passing through bridges (routers) to other rings. SNA traffic uses these timers since it is a link layer protocol. For IBM PCs it's our understanding that if the destination TR address includes a route which specifies a bridge will be used, the Group 2 (larger timers) will be used. If a local address is used, Group 1 timers prevail.

How Macintosh Solutions Use These Timers

The Token Ring timer values to do NOT apply to any AppleTalk protocol traffic. The only types of data which use these values are Link Layer (DLC) protocols such as SNA where requests are sent with responses expected. The timers provide some basis for the DLC to probe and re-request potentially "lost" information. The Apple product which uses these timers is SNA•ps Token Ring gateways.

The Apple Token Ring 4/16 NB card does send out TEST frames with Non-Broadcast verses All-routes Broadcast, but the SNA•ps product does not use the information which may be returned about bridged versus non-bridged LANs. The SNA•ps Token Ring line driver instead uses the Group 1 and Group 2 timers in the Token Ring Control Panel as a reference for actual timer calculations. In the SNA•ps Token Ring Line configuration, there are the following timers:

- Response Timer - which equivalent to the T1 timer
- Receive Acknowledge Timer - which is equivalent to the T2 timer
- Inactivity Timer - which is equivalent to the TI timer

You will notice that these timers are not separately identified as Group 1 and Group 2. Instead, the values entered by the customer are limited by the SNA•ps Config application to a range of 1 to 10. Any numbers between 1 and 5 are used to multiply the actual Group 1 timers entered in the Token Ring Control Panel. If the numbers are between 6 and 10, 5 is subtracted from the entry and the resulting number is used to multiply the actual Group 2 timer.

SNA•ps Example using Default Values

Lets do some examples as a way to fully understand the Token Ring timers and how SNA•ps uses them on Token Ring gateways. If you use the default timer values SNA•ps Config does for a new Token Ring line, you get the following values:

- Response Timer: 5
- Receive Ack Timer: 2
- Inactivity Timer: 8

These "timer" values are not in seconds but are used to multiply the actual Group 1 and 2 Timers found in the Token Ring Control Panel. If the Control panel is still in its default values then:

	Group 1 (sec)	Group 2 (sec)
T1	0.20	1.00
T2	0.04	0.40

TI 1.00 5.00

So the SNA•ps gateway would use the following timers (in seconds):

- SNA•ps T1 = 5 x 0.2 = 1.0 seconds
- SNA•ps T2 = 2 x .04 = .08 seconds
- SNA•ps TI = (8-5) x 5.0 = 15 seconds (uses Group 2 timers)

Maximum Values for SNA•ps Timers

The Token Ring Control Panel values cannot be set higher than 10.2 seconds for T1. T2 and TI in Group 1 & 2. The maximum values in SNA•ps is a multiplier of 10 which actually multiplies the Group 2 timers by 5

So maximum values:

SNA•ps T1 = 5 x 10.2 = 51 seconds
SNA•ps T2 = 5 x 10.2 = 51 seconds
SNA•ps TI = 5 x 10.2 = 51 seconds

What happens if these timer values are changed?

The T1 Timer

Increasing the T1 timer from the default value (1.0 seconds) will not have any noticeable performance issues unless the LAN is exhibiting a large number of transmission errors and it would take longer for the sender to realize that the receiver never answered and tries it again. This value should only be increased to the 5 second range on very large Bridged Token Ring LANs where the sending node and the receiving node are separated by low speed WAN segments.

If the T1 timer is set too low, like .2 seconds, then we will be retransmitting too early to nodes which have not had a chance to reply.

The T2 Timer

Increasing the T2 timer from the default value (0.08 seconds (80 ms)) is not generally necessary. This will have a direct and noticeable effect on performance since the node which is desiring to respond to a poll will wait longer than necessary. There seems to be no strong reason in increase this timer beyond the (0.80 second (800 ms)) value in the largest bridged network. Never make the T2 timer > than the T1 timer since the responding node will wait more time than the sender will to wait to get an answer and a retransmit will be sent.

The TI Timer

Increasing the default value of the TI timer (15 seconds) is probably never necessary. This timer comes into play if the SNA•ps gateway has not been polled by the host for this period of time. It would be a very slow host that would require this value to be increased on a very large bridged LAN. The timer value does affect the delay in notification of down connections to end users who may be confused by "slow" response time.

Detailed explanation

There are three timers defined in the Token Ring Architecture Guide, T1, T2 and TI. The Group 1 and Group 2 designations are not common terms in IBM, we can apply the IBM terms "not bridged or not routed (Group 1) and bridged or routed (Group 2)" LANs.

To determine whether a link is bridged or not the IBM PC directs the LAN Support Program to send TEST frames with the broadcast indicator bits (in the routing information field) set to "non-broadcast". If a response to a frame with the "non-broadcast" bits set is received, then the PC assumes that the link is not bridged. If there is no response to the "non-broadcast" TEST frames, then the TEST frames are resent with the broadcast indicator bits set to "all-routes broadcast" to attempt to establish a bridged link. For token-ring connections (which use source-routing), this method of bridged/not-bridged detection works reliably.

I have shown below the IBM PC values (TRLI config values) which can be configured and compared that with the Macintosh values.

TIMER T1

TIMERT1 specifies the length of time that the PC waits for an acknowledgment before resending a frame that requires an acknowledgment. If TIMERT1 is too short, frames will be needlessly retransmitted, causing slower response time. If TIMERT1 is much too short, then the link will fail. If TIMERT1 is too long, then there will be excessive delay before a user is notified of a link failure.

The TIMERT1 value can be set to 0 though 10. The actual response timeout value falls into a range that depends upon the setting of the TIMERT1 parameter. The following table gives the actual response timeout for each possible setting.

Value Specified	Actual Value in Seconds
0	1.0 - 1.2
1	0.2 - 0.4
2	0.4 - 0.6
3	0.6 - 0.8
4	0.8 - 1.0
5	1.0 - 1.2
6	1 - 2
7	2 - 3

	8	3 - 4
	9	4 - 5
	10	5 - 6

When TIMERT1 is not specified in the TRLI entry, the response timeout will be:

Token-Ring, not bridged	1.0 to 1.2 seconds
Token-Ring, bridged	5.0 to 6.0 seconds
Ethernet or 802.3, either bridged or not bridged	1.0 to 1.2 seconds
PC Connected to Ethernet/802.3 LAN that is bridged to token-ring via 8209	1.0 to 1.2 seconds
PC Connected to token-ring that is bridged to Ethernet/802.3 LAN via 8209	5.0 to 6.0 seconds

TIMER T2

TIMERT2 specifies the length of time that the PC waits before sending an acknowledgment to a received frame. In general, the lower TIMERT2 is set, the better response time will be for this PC. However, setting TIMERT2 too low creates unnecessary traffic on the LAN which can, in turn, increase response time for all LAN users. To cause the PC to respond immediately to all traffic, TIMERT2 can be disabled by specifying a value of 11 or greater.

The TIMERT2 value can be set to 0 though 99. The actual acknowledgment timeout value falls into a range that depends upon the setting of the TIMERT2 parameter. The following table gives the actual acknowledgment timeout value for each possible setting.

Value Specified	Actual Value in milliseconds
0	80 - 120
1	40 - 80
2	80 - 120
3	120 - 160
4	160 - 200

5	200 - 240
6	400 - 800
7	800 - 1200
8	1200 - 1600
9	1600 - 2000
10	2000 - 2400
11 - 99	Timer disabled

When TIMERT2 is not specified in the TRLI entry, the acknowledgment timeout will be:

Token-Ring, not bridged	80 to 120 milliseconds
Token-Ring, bridged	400 to 800 milliseconds
Ethernet or 802.3, either bridged or not bridged	80 to 120 milliseconds
PC Connected to Ethernet/802.3 LAN that is bridged to token-ring via 8209	80 to 120 milliseconds
PC Connected to token-ring that is bridged to Ethernet/802.3 LAN via 8209	400 to 800 milliseconds

TIMER TI

TIMERTI specifies how often the PC should check to be sure that the link to the AS/400 is still operational. Smaller values of TIMERTI allow the user to be informed of a link failure more quickly but cause more traffic on the LAN. If TIMERTI is much too small, the PC will incorrectly determine that the link is not operational and will report a link failure even when the link is operational.

The TIMERTI value can be set to 0 though 10. The actual inactivity timeout value falls into a range that depends upon the setting of the TIMERTI parameter. The following table gives the actual inactivity timeout value for each possible setting.

Value Specified	Actual Value in Seconds

0	25 - 30
1	1 - 2
2	2 - 3
3	3 - 4
4	4 - 5
5	5 - 6
6	5 - 10
7	10 - 15
8	15 - 20
9	20 - 25
10	25 - 30

When TIMERTI is not specified in the TRLI entry, the inactivity timeout will be:

Token-Ring, not bridged	5 to 6 seconds
Token-Ring, bridged	25 to 30 seconds
Ethernet or 802.3, either bridged or not bridged	5 to 6 seconds
PC Connected to Ethernet/802.3 LAN that is bridged to token-ring via 8209	5 to 6 seconds
PC Connected to token-ring that is bridged to Ethernet/802.3 LAN via 8209	25 to 30 seconds

Article Change History:
29 Mar 1995 - Reviewed and updated format.
19 Sep 1994 - Added info on Token Ring NB/c card.
26 Jun 1993 - CORRECTED to show how Macintosh solutions use these timers.

Support Information Services

Keywords: supt,knts

==========

This information is from the Apple Technical Information Library.

19960215 11:05:19.00

Tech Info Library Article Number: 12188