FDDI Station Management with the National Chip Set

National Semiconductor
Application Note 728
Robert M. Grow, Jim Schuessler
January 1991

1.0 INTRODUCTION

The National DP83200 FDDI Chip Set includes special features that aid in the management of an FDDI station as well as the management of an FDDI ring. An attempt is made here to guide you through some of the details of Station Management (SMT) using National's DP83200 FDDI Chip Set with as little pain as possible. Special circumstances for non-standard applications are also discussed.

Due to the universally acknowledged complexity of the FDDI Standard, it is always a wise idea to have ready access to the original documents when reading collateral material—this is no exception! We recommend obtaining the ANSI X3T9.5 FDDI Standards set[1].

The BMAC™ device and PLAYER™ device are two of the devices comprising the National DP83200 FDDI Chip Set. They both provide a rich set of fully maskable interrupts. These interrupts are used to drive SMT protocols, including Frame Based Management, Connection Management and Ring Management. The BMAC device includes counters for fault isolation and station and ring performance monitoring that ease the implementation of SMT protocols. The PLAYER device includes multiplexing capability to implement the node configurations called out in the SMT Configuration Management state machine (the most popular being a Single Attach Station (SAS) and Dual Attach Station (DAS)), internal hardware for link error monitoring, and noise timing (see *Figure 1*). The full duplex architecture of the chip set provides for comprehensive testing and fault isolation.

DP83261 BMAC DEVICE (BASIC MEDIA ACCESS CONTROLLER)

- SMT multicast addressing on-chip
- Full duplex architecture
- Auto generation of Beacon and Claim frames
- Multiple transmit immediate modes
- Multiple diagnostic counters
- Ring latency timer
- 4-bit late counter
- Same information field detection for MAC frame filtering
- Duplicate address detection
- Multiple token detection

DP83251/55 PLAYER DEVICE (PHYSICAL LAYER CONTROLLER)

- On-chip configuration switch
- Line state history registers
- Link error detector
- Noise threshold timer
- Unique injection register
- Full duplex architecture

FIGURE 1. National DP83200 FDDI Chip Set SMT Features

[1] There are four documents comprising the FDDI Standard: PMD, PHY, MAC and SMT. The first three are published standards available through ANSI (phone: 212-642-4900), the last, SMT, is still in draft form at the publication time of this Application Note. It can be obtained through Global Engineering Documents, phone 800-854-7179, or 714-261-1455.

Four general areas of management will be discussed in this paper: how to select values of operational parameters, how to use the BMAC device for fault isolation, how to use the PLAYER device for implementing Connection Management (CMT), and how to monitor network status and performance.

2.0 OPERATIONAL PARAMETERS

FDDI supports a broad range of network configurations, and the FDDI standard specifies default parameters for operation of large configurations. The National FDDI Chip Set is designed to operate over an even larger range of configurations for special applications. The default values of these parameters must be changed for larger configurations. In a few systems, it is also valuable to optimize parameters in small ring configurations.

At the core of the timed token protocol implemented in the BMAC device are timers used to control the transmission of information on the network and to detect when ring recovery is required. These timers are loaded with values which ultimately determine the performance of the ring. Ring recovery/startup is a function which can be performed by the BMAC device automatically with default timer values. Other values for these timers can be loaded, but be warned: changing these values from the defaults may cause poor performance, (high usable token latency or low throughput) or worse yet, oscillation between Claim and operational states. For instance, a shorter Valid Transmission Timer (TVX) value might be used on a small ring to accelerate ring recovery, but if the shorter value was used on a large ring it could cause the above problems.

In most cases, the operation and performance of FDDI is determined by the most demanding (typically the shortest!) parameter in all stations on the ring. For example, the station with the shortest TVX value will frequently be the station starting a Claim Process. This is because any timed out TVX will cause the BMAC device to start the Claim Process (TVX timeout indicates no frame received within TVX time).

When powering up, or after a hardware reset, the BMAC device has been designed to revert to default values for many critical parameters; though, a station must initialize the Parameter RAM before participating in an FDDI ring. Parameter RAM values include the long and short addresses, for example.

The use of the BMAC device's programmable timers is discussed below to help illustrate the different alternatives for management of the parameters and selection of proper values for desired operational characteristics.

2.1 The Claim Process

Claim is a ring state which must be completed before the ring can become operational. The objective is to create an interoperable environment in which all stations may communicate with both asynchronous and synchronous traffic. The process does this by setting the ring's Target Token Rotation Time (TTRT) and determining who will issue the first token (the ''winner'' of claim). Following FDDI rules for the Claim Process, the station with the shortest Requested Target Rotation Time (TREQ) is the ''winner'', and will determine the Negotiated Token Rotation Time (TNEG) for the ring. TNEG is used as the operational value for the Token Rotation Time (TRT) in all stations on the ring.[2]

[2] In many cases, stations will have the same TREQ value. In this case other information in the Claim frame is used to select the ''winner''.

A simplified Claim Process flow is shown in Table I:

TABLE I. Timer Values in the Claim Process

Value	Becomes	Value
TREQ	→	Tbid in a station's Claim frames
Shortest Tbid	→	TNEG in all stations
TNEG	→	Token Rotation Timer (TRT) when the ring becomes operational

Note: See Recovery Required in MAC Standard

The BMAC device is capable of automatically generating Claim frames. It starts the Claim Process when TVX times out or Tlate = 2 (Tlate is a 4-bit counter which increments once each time TRT goes to zero). When the frames are transmitted, the BMAC device places the TREQ value stored in the BMAC device Parameter RAM into the Claim frame. This value then becomes Tbid to the next station on the ring. The receiving BMAC device saves the Tbid from the Claim frame as the potential TNEG, while comparing it to its TREQ. If the received Tbid time is shorter than its own TREQ, it stores the Tbid value as its new TNEG, stops transmitting Claim frames and repeats incoming Claim frames. If the received Tbid is larger than the current TREQ, the station keeps (or starts) transmitting its current Claim frame. The Claim Process completes when the BMAC device receives Claim frames with both source and destination address equal to its own, (its own Claim frame) as well as the Tbid value equal to its TREQ. This process insures that the shortest Tbid value of any node participating in the Claim process ends up as TNEG for all nodes. If an implementation externally generates Claim frames instead of letting the BMAC device generate them, then it is very important that TREQ in the BMAC device Parameter RAM be equal to the first four bytes of the Claim frame (Tbid). If this isn't done, the Claim Process may not complete, and a false duplicate address condition may be detected by the SMT Ring Management protocol.

2.2 Selection of TREQ

Two major factors are used in selecting a TREQ value. The first is the delay and segment size for synchronous traffic. The second is the desired queuing delay for all traffic, both asynchronous and synchronous.[3] The first factors: delay and segment size are just another way of specifying the throughput necessary for the synchronous application. The data source is usually isonchronous (periodic) and therefore must be serviced or ''disaster'' will strike. (An example might be voice data, where a delay would result in a noticeable blank spot in continuous speech.)

An application must be able to rely on the stability of TNEG, therefore, TREQ should not be changed frequently. It is generally a bad idea to change TREQ as application processes start and stop. In fact, there are really only two reasons for changing TREQ from the default TMAX: To create a synchronous service period, or to adjust the asynchronous maximum usable_token latency—both relate to token latency.

[3] For a discussion of Asynchronous versus Synchronous traffic, see the ANSI X3T9.5 FDDI MAC Standard.

2.2.1 Selection of TREQ for Synchronous Traffic

Changing TNEG, and thus TRT, has significant implications on synchronous traffic. Synchronous service is usually viewed as some number of bits per second; but in FDDI, synchronous service is provided in bytes per token rotation. Each time a token is received, a station may transmit X bytes based on its synchronous allocation. The total synchronous bandwidth allocation for the ring may not exceed TNEG less overhead.[4]

A synchronous application will generally determine the bandwidth requirement on application layer (OSI Layer 7) data; but the synchronous allocation requested in SMT protocols must include overhead for placing the application data in a frame. A change in TNEG changes the framing overhead. This is because the number of bytes of overhead is the same for 1 byte of application data or 4 kBytes of application data. Since an objective of synchronous service is to guarantee bandwidth to an application, a shorter TTRT will cause the application data to be segmented into smaller sizes. So as TNEG is lowered, response time decreases (faster token rotation) but overhead increases and therefore overall throughput decreases. This is the tradeoff between response time and throughput.

For example, the synchronous requirement for an application layer requiring 10,000 bytes per second above the MAC layer would increase 37% on MAC overhead alone when TNEG changes from 20 ms to 5 ms. A little math will illustrate this.

Remember that all bytes transmitted must be counted in the synchronous allocation. The fixed bytes of overhead per frame are shown in Table II.

TABLE II. Fixed Bytes of Overhead per Frame

Number of Bytes	Portion of Frame
8	Preamble
1	SD (JK Symbol)
1	FC
12	Addresses: SA, DA
. . .	DATA
4	FCS
2	ED
28	Total

If TRT is 20 ms that means the token will circulate 50 times per second (1/20 ms) at the maximum network utilization. Our requirement is for 10,000 bytes/s which means 200 bytes per frame (10,000 bytes/s)/(50 frames/s). If TRT is 5 ms, the other alternative, the token rotates 200 times per second (1/5 ms). The same 10,000 bytes are delivered in 50 bytes per frame (10,000 bytes/s)/(200 frames/s). Therefore:

5 ms Case:

28 bytes overhead + 50 bytes data = 78 bytes/frame → 36% overhead

78 bytes/frame * 200 frames = 15,600 total bytes (10,000 bytes data)

[4] Actual time is Target Token Rotation Time (TTRT) less the sum of maximum Ring Latency (D_MAX = 1.617 ms), maximum Frame Time (F_MAX = 0.361 ms), and Token Time (0.00088 ms).

20 ms Case:

28 bytes overhead + 200 bytes data = 228 bytes/frame → 12% overhead

228 bytes/frame * 50 frames = 11,400 total bytes (10,000 bytes data)

There is a 37% decrease in bandwidth in the 5 ms TRT case verses the 20 ms TRT.

If other protocol information like an LLC is transmitted on each frame, the increase would be even worse. In addition, a change in TNEG also creates ring stability problems for previously enqueued synchronous information. Frames which are queued at 20 bytes, for example, would cause the ring to go into the Claim Process if enough of them were queued when TRT changed to a lower value requiring 50 byte frames.

Synchronous service is not well specified in current FDDI documents. If each synchronous application is allowed to pick its own TREQ, then as applications start and stop, TNEG will increase or decrease. In most cases, it is better for an application to learn what the synchronous target time is and segment to that size, rather than dynamically changing it as applications start and stop. This simpler model of operation can be handled in the ring's synchronous bandwidth allocation algorithm.

All this is to say that applications (OSI Layer 7) should not specify or have control over TREQ. Synchronous allocation should be done by a process which has global knowledge of the throughput and latency requirements of all stations. SMT is uniquely qualified for this purpose.

2.2.2 Selection of TREQ for Asynchronous Traffic

Asynchronous traffic is not effected significantly by the value of TNEG in typical systems. Here, the desired TNEG is based on a tradeoff between network throughput and response time. For example, on a large ring of 150 stations (n) at 1 μs latency per station, the maximum throughput is 99% at 20 ms. TNEG, and 97% at 5 ms TNEG.

$$\frac{n(\text{TNEG} - \text{Ring Latency})}{n(\text{TNEG}) + \text{Ring Latency}} = \text{Percentage Throughput} \quad (1)$$

Therefore: $150\ (20\ \text{ms} - 150\ \mu\text{s})/(150(20\ \text{ms}) + 150\ \mu\text{s}) = 99\%$

and: $150\ (5\ \text{ms} - 150\ \mu\text{s})/(150(5\ \text{ms}) + 150\ \mu\text{s}) = 97\%$

Another important performance characteristic is the maximum usable_token latency. Usable_token latency is the time for a token to return to a particular station, and be usable for asynchronous traffic. This means the TRT has not timed out once since the station last saw the token. (This is opposed to maximum token latency which is 2 times TNEG, a much smaller time.)

It is possible, though highly improbable, that each station in an overloaded ring could use the token for TNEG time (minus ring latency) when the token is captured. This is a worst case scenario. The maximum usable_token latency would then be (n − 1) (TNEG) + 2 (Ring Latency). For this improbable event to occur each station on the ring must transmit for the maximum allowable time—in this case TNEG minus ring latency.

In the above configuration, the maximum usable_token latency is 3 seconds at 20 ms TNEG, and 0.75 seconds at 5 ms TNEG.

(n − 1) (TNEG) + 2 (Ring Latency) = Max. Usable Token Latency (2)

Therefore: (150 − 1) * 20 ms + 2(150 μs) = 3.0s

and: (150 − 1) * 5 ms + 2(150 μs) = 0.75s

This kind of tradeoff can best be made by a network management station as described later, since a station that attemps to minimize the usable_token latency can adversely effect network throughput in some configurations. For example, if the Ring_latency were 1 ms, instead of the 150 μs above, the change in TNEG from 20 ms to 5 ms would change the maximum throughput from 95% to 85%.

TABLE III. Example Configurations and TNEG Value

Stations	TNEG	Maximum Throughput	Maximum Usable_Token Latency
10	8	0.9986	0.072 ms
	64	0.9998	0.576 ms
	167	0.9999	1.503 s
150	8	0.9811	1.192 s
	64	0.9976	9.536 s
	167	0.9991	24.883 s
500	8	0.9374	3.993 s
	64	0.9922	31.937 s
	167	0.9970	83.334 s

Note: At 1 μs per station latency (including optical fiber propagation delay), actual latency may be different.

Note that a TNEG (TREQ) of 8 ms can significantly reduce the usable_token latency (a good thing) and still not decrease throughput significantly for rings of under about 150 stations.

A robust method of controlling the TNEG of a ring can be created using the operational characteristics described above, and the management features of the BMAC device. Normal stations operate with the default TMAX (approximately 167 ms) and TREQ equal to the TMAX. A network management station determines the desired TNEG for the ring based on knowledge of synchronous applications requirements and ring throughput implications. The network management station sets its TREQ to the desired TNEG value, thus determining the TNEG resulting from the Claim Process. See Table III.

If implementors allow stations to set TREQ independently, then it is advisable that a lower bound on TREQ be enforced to protect the network from serious denial of service problems. When TREQ is changed, the station can cause a Claim Process by setting the CLM bit of the Function Register (FR) in the BMAC device. If this is not done, then it may be a long time before the desired change in TNEG would occur.

2.3 Selection of Asynchronous Priority Thresholds

Asynchronous priorities are set in the THSH1, THSH2, and THSH3 Registers. These priorities are of greatest value when the ring latency is large. Two approaches can be used for setting the thresholds. The first sets the threshold at a load factor, for example 50% load. In the majority of systems, the latency will be small enough that all load factors default to the same effective level, namely transmit if no frames were transmitted on the last rotation. In the large 150 station ring described earlier, a 1 kByte frame would represent 50% network load to the MAC timers. A longer latency would require a larger frame to get the same 50% load factor. The latency timing feature of the BMAC device allows determination of the appropriate threshold for a load factor. (See Network Monitoring, Section 5, for an example load factor calculation).

The threshold may also be viewed as a reservation of time for transmission of frames at higher priority. In this case the thresholds can be set directly from the tables contained in the BMAC device datasheet (end of Section 5).

2.4 Selection of TMAX and TVX

Stations should only change the TVX and TMAX values when there are critical responsiveness requirements. No improvement in throughput will be achieved, and in most rings, no improvement in network availability will be achieved. However, a discussion of their selection criteria is provided her as guidance for tuning these parameters to the latency of the network.

2.4.1 Selection of TMAX

The BMAC device has been designed so that it can be applied to rings with extremely large latencies. In these cases, the value of TMAX will need to be longer than the default. Making TMAX shorter has a statistically insignificant impact on ring availability, therefore, most implementors need only use the default specified in the FDDI Standard. The BMAC device sets TMAX to default on reset. Special closed systems may benefit from a shorter TMAX value, since a very small set of ring failures is detected by timing for TMAX. In either the longer or shorter TMAX case, the BMAC device can be loaded with the desired value after reset.[5]

2.4.2 Selection of TVX

The Valid Transmission Time (TVX) register is used to hold the FDDI parameter TVX_valid. The token is assumed to be lost if the time between receiving valid frames or non-restricted tokens is longer than TVX. The BMAC device loads TVX with the default value recommended in the FDDI MAC Standard upon reset. TVX need only be changed in rings with latencies larger than 1.7 ms. If a station has a TVX value that is too small, the likely symptom will be a ring that oscillates between the Claim Process and being operational. The optional SMT Parameter Management Frame (PMF) capability can be used by a network manager to attempt to fix an oscillation condition. The National FDDI chip set has been designed with larger than default counters to allow large latency networks, other implementations may not be able to interoperate in one of these very large networks.

In stations which need a non-default TVX value, station implementors can provide a non-volatile storage location. This feature would avoid potential oscillation between the ring being operational and being in the Claim Process, when stations are powered off and on in a larger network than supported by the default value.

2.5 Adjustment of Value for Parameter Encoding

Two representations are used for timer values in the BMAC device. Where accuracy and resolution are important, the chip uses binary encoding. Where network manageability

[5] See the FDDI Standard for the calculation of these values. The value of DMAX should be computed from the equations in the PHY Standard. TVX and TMAX equations are given in the MAC Standard.

would not be compromised, an exponential representation was used. Tables for converting exponential values are included in the BMAC device datasheet. This optimization saved circuitry which allowed other functions to be included in the BMAC device. When desired values cannot be represented exactly in the chip, the guidelines shown in Table IV can be used.

TABLE IV. BMAC Device Parameter Encoding Guidelines

TREQ	Loaded Time ≤ Desired Time
TMAX	Loaded Time ≥ Desired Time
TVX	Loaded Time ≥ Desired Time
THSH	Loaded Time Is the Closest to Desired Time

2.6 Changing Addresses

The BMAC device has been designed to reduce the number of things that an implementor needs to worry about. The setting of the station addresses, unfortunately, is not one of those things. Addresses can be changed by the SMT processor through the control interface, and here lies potential danger. Dangers exist for both Group and Individual Addresses; but the more serious implications are in changing an Individual Address. If an Individual Address is changed while a frame is on the ring, or still enqueued, then a no owner frame can be created, since the address is changed one byte at a time (a no owner frame will eventually be stripped when it runs into a station which is transmitting). This can be avoided by waiting for the transmit queue to become empty, then disable the Individual Address with the BMAC Device Option Register, until the change is complete. If the implementation uses the optional external Claim or Beacon frames, the address must be changed in those frames also.

The BMAC device also includes an on-chip SMT group address recognition capability. The SMT committee has request addresses for use in SMT frame protocols. These reserved addresses are for the exclusive use of SMT processes. Changes to the base group addresses may result in frames being copied as the result of comparison against a partially changed address. If the group address capability is used for SMT addresses, individual addresses can be enabled and disabled without this problem.

2.7 Denial of Service Protection

Some of the discussion on individual parameters indicated how the ring can become unusable with the improper setting of that parameter. Improper use of parameters or frames can result in disruption of service to the entire ring, or in some cases, to a single station. These conditions can be grouped together as denial of service problems. In general, it is prudent that an implementor only allow operational changes by trusted software. This includes the setting of parameters, as well as the ability to source MAC frames (e.g., Claims and Beacons) and SMT frames. This is simplified by features in the BMAC device like internal Claim and Beacon generation, transmission of Source Address from the Parameter RAM, and reset to default values of important parameters.

3.0 USE OF THE BMAC DEVICE FOR FAULT ISOLATION

In some failure cases, communication on an FDDI ring will be impossible because of a low probability failure within some station on the network. Sometimes, this may result in the ring being non-operational, and other times the ring may be oscillating between operational and non-operational states. The BMAC device is designed to support network management applications that correct or isolate these rare faults. Described below are several methods which facilitate fault isolation.

3.1 How to Perform Transmit Immediate

Transmit Immediate is a BMAC device feature which allows the transmission of any frame without the ring becoming operational. In other words, no token needs to be received, the station just strips anything received and transmits its frame—thus: Transmit Immediate.

The transmit immediate capability can be used to isolate many faults. One tool that is useful is to allow a network manager to segment the ring. Using this capability, faults can be localized by monitoring the symptom of the failure. For example, if the ring cannot become operational, an application can segment the ring by forcing a configuration change in a remote station(s). If the symptom goes away in the segment containing the network management station, the fault is probably in the isolated segment of the network, if it doesn't, the fault is probably in the remaining segment of the network. The same procedure can then be used with a different remote station until the fault domain is located.

In the case of timer parameter faults, the problem may be corrected directly by performing a transmit immediate of an SMT PMF Request Frame to the station with the invalid parameter.

Applications using the transmit immediate capability must take into account three important items. Differing implementations of transmit immediate, transmission of MAC frames, and the effect of the Ring being Operational.

The BMAC device has the capability to perform transmit immediate under all ring conditions; other implementations do not. Therefore, the application cannot expect a response from other stations.

The fault isolation protocol must take into account that in a ring stuck at the Beacon Process, each repeating station will enter Claim every TMAX, destroying traffic being repeated at that time. As a result, the source or destination of the frame, or any station between may make the transition to Claim, causing an abort of the management frame. The probability of getting a frame to the destination is improved with a few techniques. Setting the Inhibit Recovery Required option (IRR bit in the BMAC device Option Register) will allow the station to transmit complete frames independent of the station's TRT expirations. Setting the IRPT option will stop MAC frames generated by other stations from aborting transmission. A short frame has a statistically smaller chance of being aborted by other stations; but retry of the frame may also be necessary.

3.2 Implementing SMT Events

The BMAC device and PLAYER device are designed to allow for reporting of significant network events. This includes timer expirations, received frame conditions, and counter increments and overflows. All of these conditions are implemented as maskable interrupts. Use of these interrupt conditions can eliminate any need for polling of status in the FDDI logic. SMT frames are transmitted as the result of some of these events. The ability for generation of interrupts on increment of a BMAC device statistical counter (e.g., Error Counter) allows for generation of event report frames directly from BMAC device interrupts.

4.0 USE OF THE PLAYER DEVICE FOR CONNECTION MANAGEMENT

The interface between the PLAYER device and the FDDI connection management (CMT) protocol is designed so that time critical operations are performed by the PLAYER device. The most time critical operation to be performed by the CMT software is PC_React. PC_React is equal to 3 ms and is defined by the standard as the maximum time for the Physical Connection Management (PCM) state machine (implemented by a combination of hardware and software) to make a transition from the active state to the break state in response to the Quiet Line state (QLS), a fault condition, or a request to start the PCM protocol (PC_Start).

Important fault isolation features are provided in the Connection Management (CMT) protocols specified in the SMT standard. The PLAYER device includes counter and interrupt logic to aid in implementing these protocols. The line state reporting of the PLAYER device includes individual reporting of each state transitioned, thus providing a history of all line states seen since the register was cleared. (Registers: Receive Condition A and B (RCRA, RCRB)). This is important since the Physical Connection Management (PCM) is intended to run, even when the optical link has an extremely high error rate. The line state history thus provides greater flexibility to the software implementing these protocols. In addition, individual masking of line state interrupts can eliminate interrupts for line states that are ignored in a specified operational state.

Another fault isolation capability supported directly by the PLAYER device is the noise timer, which detects jabber conditions (continuous transmission) and other faults that may occur in a system. This timer is referred to as the TNE timer in the standard. The PLAYER device has prescale and count registers which load countdown counters. Each Noise Line State, Active Line State or Line State Unknown symbol will decrement the noise counter. The noise counter is used by the Physical Connection Management (PCM) state machine to detect the length of the noise events.

When the TNE timer threshold is exceeded, an interrupt can be generated by the PLAYER device which causes the PCM state machine to transition from the active to the break state. The PLAYER device also includes hardware that implements the Link Error detector function of SMT. The LETR and CLECR registers are also used for performing the Link Confidence Test during link establishment.

5.0 NETWORK MONITORING

Maximum network throughput can be calculated as:

n(TTRT − Ring Latency)/(n(TTRT) + Ring Latency) (1)

Again

Where n is the number of stations in the network.

This basic equation can be used to determine the proper values for TREQ. This function determines the asymptote for network throughput. The actual utilization of the network can be calculated as:

(TRT − Ring Latency)/TRT (3)

Where TRT is the actual token rotation time.

This function produces the curve shown in *Figure 2*.

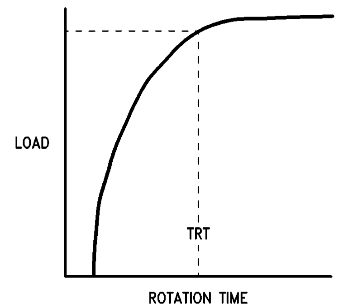

TL/F/11084–1

FIGURE 2. Timed Token Protocol

If the Token Rotation Time is equal to TNEG, then the network is at the maximum utilization. In a similar way, a T_pri value can be determined for loading into a THSH register by picking the desired load factor and determining the corresponding time value. Conversely, a time value can be picked to determine the maximum utilization at that T_pri.

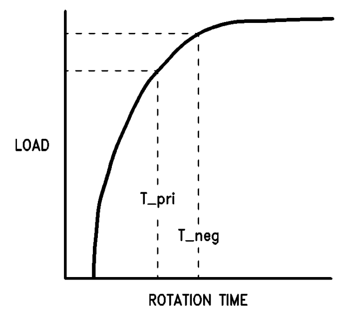

TL/F/11084–2

FIGURE 3. Asynchronous Priorities

The same basic equation (Eqn. 3) is used to monitor the throughput of the network. Either monitoring the throughput or setting timer parameters requires determination of the ring latency. The BMAC device includes hardware that performs a latency measurement of the ring. The latency of a ring will vary slightly through expansion and contraction of the elasticity buffers and smoothers required in PHY components like the PLAYER device, and more significantly through stations entering and exiting the network. Therefore periodic measurement may be necessary for network monitoring. The latency measurement, obtained through the RLCT counter in the BMAC device, can be used in conjunction with the token counter (TKCT) to determine average load over time. Average load is then represented by the equation.

(Time − (Token_Ct * Ring Latency))/Time (4)

For example:

(5 sec − (5000 Tokens * 150 μs))/5 sec = 85% Utilization

Over a measurement period of 5 seconds, 5000 tokens were counted in TKCT and the ring latency counter read 150 μs. Working the formula through results in an 85% load on the ring.

The BMAC and PLAYER devices also include required and optional statistical counters that can be used to evaluate traffic in terms of frames. One of the important counters added is the Frame Not Copied (FNCT) count of the BMAC device. This counter is very useful for evaluation of overload on stations in a network, since it indicates inability of the station to keep up with frames sent to it. Network managers need this type of information for proper administration of server stations. The BMAC device transmit and receive frame counters (FTCT and FRCT) provide valuable traffic information with virtually no software overhead. The error counters in the BMAC and PLAYER devices are useful indications of error rate problems on communication links in the ring. All of these counters are designed to simplify the implementation of station software.

REFERENCES

1. National Chip Specifications.
2. Fiber Distributed Data Interface (FDDI)—Token Ring Media Access Control (MAC), ANSI X3.139—1987.
3. Fiber Distributed Data Interface (FDDI)—Token Ring Physical Layer Protocol (PHY), ANSI X3.148—1988.
4. Fiber Distributed Data Interface (FDDI)—Token Ring Physical Layer Medium Dependent (PMD), ANSI X3.166—1990.
5. Fiber Distrubuted Data Interface (FDDI)—Station Management (SMT), X3T9.5/84-49, Rev. 6.2.
6. R. Jain, Performance Analysis of FDDI Token Ring Networks: Effect of Parameters and Guidelines for Setting TTRT, DEC-TR-655, September 1989.

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
2900 Semiconductor Drive
P.O. Box 58090
Santa Clara, CA 95052-8090
Tel: 1(800) 272-9959
TWX: (910) 339-9240

National Semiconductor GmbH
Livry-Gargan-Str. 10
D-82256 Fürstenfeldbruck
Germany
Tel: (81-41) 35-0
Telex: 527649
Fax: (81-41) 35-1

National Semiconductor Japan Ltd.
Sumitomo Chemical Engineering Center
Bldg. 7F
1-7-1, Nakase, Mihama-Ku
Chiba-City,
Ciba Prefecture 261
Tel: (043) 299-2300
Fax: (043) 299-2500

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semiconductores Do Brazil Ltda.
Rue Deputado Lacorda Franco
120-3A
Sao Paulo-SP
Brazil 05418-000
Tel: (55-11) 212-5066
Telex: 391-1131931 NSBR BR
Fax: (55-11) 212-1181

National Semiconductor (Australia) Pty, Ltd.
Building 16
Business Park Drive
Monash Business Park
Nottinghill, Melbourne
Victoria 3168 Australia
Tel: (3) 558-9999
Fax: (3) 558-9998