

Performance Analysis of FDDI

By Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
<http://www.cse.wustl.edu/~jain/>

This paper is a modified version of "Performance Analysis of FDDI Token Ring Networks: Effect of Parameters and Guidelines for Setting TTRT," by Raj Jain, published in the Proceedings of the SIGCOMM '90, September 1990. Copyright 1990, Association for Computing Machinery, Inc.]

Abstract

The performance of an FDDI LAN depends upon configuration and workload parameters such as the extent of the ring, the number of stations on the ring, the number of stations that are waiting to transmit, and the frame size. In addition, one key parameter that network managers can control to improve performance is the target token rotation time (TTRT). Analytical modeling

and simulation methods were used to investigate the effect of the TTRT on various performance metrics for different ring configurations. This analysis demonstrated that setting the TTRT at 8 milliseconds provides good performance over a wide range of configurations and workloads.

Institute (ANSI).[1,2] This standard allows as many as 500 stations to communicate by means of fiber-optic cables using a timed-token access protocol. Normal data traffic and time-constrained traffic, such as voice, video, and real-time applications, are supported. All major computer and communications vendors and integrated circuit manufacturers offer products that comply with this standard.

Unlike the token access protocol defined by the IEEE 802.5 standard, the timed-token access protocol used by FDDI allows synchronous and asynchronous traffic simultaneously.[3] The maximum access delay, i.e., the time between successive transmission opportunities for a station, is

Fiber distributed data interface (FDDI) is a 100-megabit-per-second (Mb/s) local area network (LAN) defined by the American National Standards

bounded for both types of traffic. Although this delay is short for synchronous traffic, that for asynchronous traffic varies with the

Performance Analysis of FDDI

network configuration and load and can be as long as 165 seconds. Long maximum access delays are undesirable and can be avoided by properly setting the network parameters and configurations.

This paper begins with a description of the timed-token access method used by FDDI stations and then proceeds to discuss how various parameters affect the performance of these systems. The target token rotation time (TTRT) is one of the key parameters. We investigated the effect of the TTRT on FDDI LAN performance and

developed guidelines for setting the value of this parameter. The results of our investigation constitute a significant portion of this paper.

Timed-Token Access Method

The token access method for network communication, as defined by the IEEE 802.5 standard, operates in the following manner. A token circulates around the ring network. A station that wants to transmit information waits for the arrival of the token. Upon receiving the token, the station can transmit for a fixed time interval called the token holding time (THT). The

the token by a station is called the token rotation time (TRT). Using this scheme, a station on an n-station ring may have to wait as long as n times the THT interval to receive a token. This maximum access delay may be unacceptable for some applications if the value of either n or THT is large. For example, voice traffic and real-time applications may require that this delay be limited to 10 to 20 milliseconds (ms). Consequently, using the token access method severely restricts the number of stations on a ring.

The timed-token access method, invented by Grow, solves this problem by ensuring that all stations on a ring agree to a target token rotation time (TTRT) and limit their transmissions to meet this target.[4]

There are two modes of transmission: synchronous and asynchronous. Time-constrained applications such as voice and real-time traffic use the synchronous mode. Traffic that does not have time constraints uses the asynchronous mode. A station can transmit synchronous traffic whenever it receives a token; however, the total transmission time for each opportunity is short.

station releases the token either immediately after completing transmission or after the arrival of all the transmitted frames. The time interval between two successive receptions of

This time is allocated at the ring initialization. A station can transmit asynchronous traffic only if the TRT is less than the TTRT.

The basic algorithm for asynchronous traffic is as follows. All stations on a ring agree on a target token rotation time. Each station measures the time elapsed since last receiving the token, i.e., the TRT. On token arrival, a station that wants to transmit computes a token holding time using the following formula:

$$\text{THT} = \text{TTRT} - \text{TRT}$$

If the value of THT is positive, the station can transmit for this time interval. After completing transmission, the station releases the token. If a station does not use its entire THT, other stations on the ring can use the remaining time through successive applications of the algorithm.

Note that even though the stations attempt to keep the TRT below the target, they do not always achieve this goal. The TRT can exceed the target by as much as the sum of all synchronous-transmission time allocations; however, these allocations are limited so that their sum is less than the TTRT. As a result, the TRT is always less than twice the TTRT.

In this paper, the performance was studied under asynchronous traffic

Performance Parameters

The performance of any system depends upon both system parameters and workload parameters as shown in Figure 1. There are two kinds of system parameters: fixed and user-settable. Fixed parameters cannot be controlled by

the network manager and vary from one ring to another. Cable length and the number of stations on a ring are examples of fixed parameters. It is important to study network performance with respect to these parameters; if performance is sensitive to them, each set of fixed parameters may require a different guideline. System parameters that can be set by the network manager or the individual station manager include various timer values. Most of these timers influence the reliability of the ring and the time it takes to detect a malfunction. The key settable parameters that affect system performance are the TTRT and the synchronous time allocations.

The workload also has a significant impact on performance. A guideline

conditions only. The presence of synchronous traffic will further restrict the choice of TTRT, as discussed later in the section Guidelines for Setting the Target Token Rotation Time.

or recommendation may be suitable for one workload but not for another. The key workload parameters are the number of active stations and the load per station. By active we mean stations on a ring that are either transmitting

Performance Analysis of FDDI

or waiting to transmit. A ring may contain a large number of stations, but generally only a few are active at any given time. Active stations include the currently transmitting station, if any, and stations that have frames to transmit and are waiting for the access right, i.e., for a usable token to arrive. The load per station varies with the number of stations, the interval between bursts, the number of frames per burst, and the frame size.

Performance Metrics

The quality of service a system provides is measured by its productivity and responsiveness.[5] For an FDDI LAN, productivity is measured by its throughput, and responsiveness is measured by the response time and maximum access delay.

The productivity metric of concern to the network manager is the total throughput measured in megabits per second. Over any reasonable time interval and for most loads, the throughput is equal to the load. For example, if the load on a ring is 40 Mb/s, then the throughput is also 40 Mb/s. But this does not hold if the load is high. For

key productivity metric is not the throughput under low load but the maximum obtainable throughput under high load. This latter quantity is also known as the usable bandwidth of the network. And the ratio of the usable bandwidth to the nominal bandwidth (e.g., 100 Mb/s for an FDDI LAN) is called the efficiency of the network. For instance, if we consider a set of configuration and workload parameters with a usable FDDI bandwidth of at most 90 Mb/s, the efficiency is 90 percent.

The response time is the time interval between the arrival of a frame and the completion of its transmission, including queuing delay, i.e., from the first bit in to the last bit out. This metric is meaningful only if a ring is not saturated, because at loads near or above capacity the response time approaches infinity. With such loads, the maximum access delay for a station, i.e., the time interval between wanting to transmit and receiving a token, has more significance.

Another metric that is of interest for a shared resource such as FDDI is the fairness with which the resource is allocated. Fairness is particularly

example, if there are three stations on a ring, each with a 100-Mb/s load, the total arrival rate is 300 Mb/s and the throughput is considerably less-close to 100 Mb/s. Thus, the

important under a heavy load. Given such a load, the asynchronous bandwidth is allocated equally to all active stations. However, the FDDI protocols are fair only if the priority

levels are not implemented. In the case of multiple priority implementation, it is possible for two stations with the same priority and the same load to have different throughput depending upon their location.[6] Low-priority stations that are close to high-priority stations may get better service than those farther downstream. We

assumed a single priority implementation to keep

the analysis simple. Since such implementations have no fairness problem, this metric will be discussed no further in this paper.

We used two methods to analyze performance: analytical modeling and simulation. We first present the analytical model used to compute the efficiency and maximum access delay of a network under a heavy load. Then we discuss the simulation model workload used to analyze the response time at loads below the usable bandwidth.

A Simple Analytical Model

The maximum access delay and efficiency are meaningful only under heavy load. Therefore, we assume

For an FDDI network with a ring latency D (i.e., the time it takes a bit to travel around the ring) and a TTRT value of T , the efficiency and maximum access delay are computed using the following formulas:

(1)

$$\text{Efficiency} = \frac{n(T-D)}{nT+D}$$

(2)

$$\text{Maximum access delay} = (n - 1)T + 2D$$

Equations (1) and (2) constitute the analytical model. Their derivation is simple and is presented in the next section.

Those readers who are not interested in the details can proceed to the section Application of the Model.

that there are n active stations, each generating enough frames to saturate the FDDI network.

Basic Equations

Digital Technical Journal Vol. 3 No. 3 Summer 1991

Performance Analysis of FDDI

Derivation

First consider a ring with three active stations, as shown in Figure 2. (Later, we will consider the general case of n active stations.) The figure shows the space-time diagram of various events on the ring. The space is shown horizontally, and the time t is shown vertically. The token reception is denoted by a thick horizontal line segment. The interval of time during which transmission of frames takes place is indicated by a thick vertical line segment.

Assume that all stations are idle until $t = D$, when the three active

stations suddenly get a large (infinite) burst of frames to transmit. The sequence of events shown in Figure 2 is as follows:

1. $t = 0$. Station S1 receives the token and resets its own token rotation timer to zero. Since the station has no frames to transmit, the token proceeds to the next station S2.
2. $t = t_{12}$. Station S2 receives the token and resets its token rotation timer to zero. t_{12} is equal to the signal propagation delay from S1 to S2.

4. $t = D$. Station S1 receives the token. Since S1 now has an infinite supply of frames to transmit, it captures the token and determines that the TRT is D . Thus, the time interval during which S1 can hold the token, the difference between TTRT and TRT, is $T - D$.
5. $t = T$. The THT at station S1 expires. S1 releases the token.
6. $t = T + t_{12}$. Station S2 receives the token. S2 last received the token at $t = t_{12}$; thus, the value of TRT is T . S2 cannot use the token at this time and releases it.
7. $t = T + t_{13}$. Station S3 receives the token. S3 last received the token at $t = t_{13}$; thus, its TRT is also T . S3 cannot use the token at this time and releases it.
8. $t = T + D$. Station S1 receives the token. S1 last received the token at $t = D$; its TRT is also T . (Note that the TRT is measured from the instant the token arrives at a station's receiver, i.e., event 4 for station S1, and not from the time it leaves a station's transmitter, i.e., event 5.) S1

3. $t = t_{13}$. Station S3
receives the token
and resets its token
rotation timer to zero.
 t_{13} is equal to the
signal propagation delay
from S1 to S3.

cannot use the token
and releases it.

9. $t = T + D + t_{12}$. Station S2 receives the token. Since TRT is only D, it sets the THT to the remaining time, i.e., $T - D$. S2 transmits for that interval and releases the token at $t = T + D + t_{12} + (T - D)$.
10. $t = 2T + t_{12}$. The THT at station S2 expires. S2 releases the token.
11. $t = 2T + t_{13}$. Station S3 receives the token. Since TRT is T, S3 releases the token.
12. $t = 2T + D$. Station S1 receives the token. Since TRT is T, S1 releases the token.
13. $t = 2T + D + t_{12}$. Station S2 receives the token. Since TRT is T, S2 releases the token.
14. $t = 2T + D + t_{13}$. Station S3 receives the token. Since TRT is only D, it transmits for the time interval $T - D$ and releases the token at $t = 2T + D + t_{13} + (T - D)$.
15. $t = 3T + t_{13}$. The THT at station S3 expires. S3 releases the token.
16. $t = 3T + D$. Station S1 receives the token, and the sequence of events begins to repeat. The

The above discussion illustrates that the system goes through a cycle of events and that the cycle time is $3T + D$. During every cycle, each of the three stations transmits for a time interval equal to $T - D$; the total transmission time is $3(T - D)$. The number of bits transmitted during this time is $3(T - D) \times 108$, and the throughput equals $3(T - D) \times 108 / (3T + D)$ bits per second. The efficiency, i.e., the ratio of the throughput to the FDDI bandwidth of 100 Mb/s, is $3(T - D) / (3T + D)$.

During the cycle, each station waits for a time interval of $2T + 2D$ after releasing the token for another opportunity to transmit. This interval is the maximum access delay. For lower loads, the access delay is shorter.

Thus, for a ring with three active stations,

$$\text{Efficiency} = \frac{3(T-D)}{3T+D}$$

$$\text{Maximum access delay} = (3 - 1)T + 2D = 2T + 2D$$

To generalize the above analysis for n active stations, substitute n for 3. Equations (1) and

token passes through
stations S1, S2, and
S3, all of which find it
unusable.

(2) are the results; the
derivation is complete.

Application of the Model

17. $t = 3T + 2D$. The cycle
continues with S1
capturing the token
as in event 4.

Digital Technical Journal Vol. 3 No. 3 Summer 1991

Performance Analysis of FDDI

Equations (1) and (2) can be used to compute the maximum access delay and the efficiency for any FDDI ring configuration. For example, consider a ring with 16 stations and a total fiber length of 20 kilometers (km). (Using a two-fiber cable, this corresponds to a cable

length of 10 km.) Light

waves travel along the fiber at a speed of 5.085 microseconds per kilometer ($\mu\text{s}/\text{km}$). The station delay between receiving and transmitting a bit is approximately 1 μs per station. The ring latency can be computed as follows:

$$\begin{aligned} \text{Ring latency } D &= (20 \text{ km}) \\ &\times (5.085 \mu\text{s}/\text{km}) + (16 \\ &\text{stations}) \times (1 \mu\text{s}/\text{station}) \\ &= 0.12 \text{ milliseconds (ms)} \end{aligned}$$

Assuming a TTRT of 5 ms and all 16 stations active,

$$\text{Efficiency} = \frac{16(5-0.12)}{16 \times 5 + 0.12} = 97.5\%$$

$$16 \times 5 + 0.12$$

$$\text{Maximum access delay} = (16$$

$$- 1) \times 5 + 2 \times 0.12 = 75.24 \text{ ms}$$

Thus, on this ring, the maximum possible throughput is 97.5 Mb/s. If the load

The key advantage of this model is its simplicity, which allows us to see immediately the effect of various parameters on network performance. With only one active station, which is usually the case, equation (1) becomes

$$\text{Efficiency}(n = 1) = \frac{T-D}{T+D}$$

$$T+D$$

As the number of active stations increases, the efficiency increases. With a very large number of stations,

$$\text{Maximum efficiency}(n = \infty) = \frac{1-D}{T}$$

This efficiency formula is easy to remember and permits "back-of-the-envelope" calculations of FDDI LAN performance. This special case of $n = \infty$ has already been studied.[7] Similarly, we can use

equation (2) to calculate

the maximum access delay with one active station as follows:

$$\text{Maximum access delay}(n = 1) = 2D$$

That is, a single active station may have to wait as long as twice the ring

is greater than this for any substantial length of time, the queues will build up, the response time will increase, and the stations may start to lose frames due to insufficient buffers. The maximum access delay is 75.24 ms; thus, asynchronous stations may have to wait as long as 75.24 ms to receive a usable token.

latency between successive transmissions because every alternate token that it receives would be unusable. For $n = \infty$, the maximum access delay approaches infinity.

Simulation Workload

One way to measure the

responsiveness of a system is to use simulation to analyze the response time. This metric depends upon the frame arrival pattern of the workload and is discussed further in the Response Time section. The workload we used in our simulations was based on an actual measurement of traffic at a customer site. The chief application at this site was the warehouse and inventory control (WIC). Hence, we named the workload WIC.

Previous network measurements show that when a station wants to transmit, it generally transmits not one frame, but a burst of frames. The WIC workload has this trait as well. Therefore, we used a bursty Poisson arrival pattern in our simulation model with an interburst time of 1 ms and five frames per burst.

We limited the frames to two sizes: 65 percent of

the frames were small (100 bytes), and 35 percent were large (512 bytes). This workload constitutes a total load per station of 1.22 Mb/s. Forty stations, each executing this load, would utilize 50 percent

Guidelines for Setting the Target Token Rotation Time

This section presents the rules specified by the ANSI FDDI media access control standard for setting the value of the TTRT. We also discuss efficiency, maximum access delay, and response time considerations, as well as reasons to limit the value of TTRT.
ANSI FDDI Standard

According to the ANSI FDDI standard, the following rules must be observed when setting the TTRT:

1. Since the TRT can be as long as twice the TTRT, a synchronous station may have to wait a time interval of up to $2T$ before receiving the token. Therefore, synchronous stations should request a TTRT value of one-half the required service interval. For example, a voice station that wants to receive a token every 20 ms or less should request a TTRT of 10 ms.
2. The TTRT must allow transmission of at least one maximum-size frame in addition to the synchronous time allocation, if any. That is,

$$\text{TTRT} > \text{ring latency}$$

of the FDDI bandwidth.
Higher load levels can be
obtained either by reducing
the interburst time or
increasing the number of
stations on the ring.

+ token time
+ maximum frame time

+ synchronous
time allocation

Performance Analysis of FDDI

The maximum-size frame on FDDI is 4500 bytes plus preamble and takes approximately 0.361 ms to transmit. The maximum ring latency is 1.773 ms. The token time (11 bytes including 8 bytes of preamble) is 0.00088 ms. This rule, therefore, requires that the TTRT be set at a value greater than or equal to 2.13 ms plus the synchronous time allocation. Violating this rule, for example, by overallocating the synchronous bandwidth, results in unfairness and starvation, i.e., some stations are unable to transmit.

3. A station must request a TTRT greater than or equal to the station parameter T_{min} . The default maximum value of T_{min} is 4 ms. Generally, most stations do not request a TTRT less than 4 ms.
4. A station must request a TTRT less than or equal to the station parameter T_{max} . The default minimum value of T_{max} is 165 ms. Assuming that there is at least one station with T_{max} equal to 165 ms, the TTRT on a ring cannot exceed this value. (In practice,

Efficiency and Maximum Access Delay Considerations

In addition to the rules specified by the standard, the TTRT values should be chosen to allow high-performance operation of a ring. This section discusses these performance considerations.

Figure 3 is a plot of efficiency as a function of the TTRT. Three configurations called "Typical," "Big," and "Largest" are shown.

The Typical configuration consists of 20 single attachment stations (SASs) on a 4-km fiber ring. The numbers used are based on an intuitive feeling of

what a typical ring would look like and not based on any survey of actual installations. Twenty offices located on a 50 m by 50 m floor require a 2-km cable or a 4-km fiber.

many stations will use a
value of $222 \times 40 \text{ ns} =$
 167.77216 ms , which can
be conveniently derived
from the symbol clock
using a 22-bit counter.)

The Big configuration consists of 100 SASs on a 200-km fiber. This configuration represents a reasonably large ring with acceptable reliability. Configuring a single ring with considerably more than this number of stations increases the probability of bit errors.[8]

The Largest configuration consists of 500 dual attachment stations (DASs) and a ring that has wrapped. A DAS can have one or two media access controllers (MACs). In this configuration, each DAS has two MACs. Thus, the LAN consists of 1000 MACs in a single logical ring. This is the largest number of MACs allowed on an FDDI LAN. Exceeding this number would require recomputation of all default parameters specified in the standard.

Figure 3 shows that for all three configurations, the efficiency increases as the TTRT increases. The efficiency is very low at TTRT values close to the ring latency but increases as the TTRT increases. Thus, to ensure a minimal efficiency, the minimum allowed TTRT on FDDI is 4 ms. This direct relationship between the efficiency and the TTRT may lead some to conclude that

the curve depends upon the ring configuration. For larger configurations, the knee occurs at larger TTRT values. Even for the Largest configuration, the knee occurs in the range of 6 to 10 ms. For the Typical configuration, the TTRT has little effect on efficiency as long as the TTRT is in the allowed range of 4 to 165 ms.

Figure 4 shows the maximum access delay as a function of the TTRT for the three configurations. To show the complete range of possibilities, we used a semilogarithmic scale on the graph. The vertical scale is logarithmic, while the horizontal scale is linear. The figure shows that increasing the TTRT brings about a corresponding increase in the maximum access delay for all three

configurations.

Table 1 presents the performance metrics for the maximum access delay and the efficiency as functions of the TTRT. As evidenced in the table, on the Largest ring, a TTRT of 165 ms causes a maximum access delay as long as 165 seconds. This means that in a worst-case situation, a station may have to wait several minutes to receive a usable token. For many

the largest possible TTRT be chosen. However, notice also that the efficiency gained by increasing the TTRT, i.e., the slope of the efficiency curve, decreases as the TTRT increases. The "knee" of

applications, this delay is unacceptable; therefore, a reduced number of stations or a smaller TTRT may be preferable.

Response Time

Performance Analysis of FDDI

Figure 5 shows the average response time as a function of the TTRT for a relatively large configuration, i.e., 100 stations and 10 km of fiber. The WIC workload was simulated at three load levels: 28, 58, and 90 percent. Two of the three curves are horizontal straight lines indicating that TTRT has no effect on the response times at these loads. The TTRT only affects the response time at a heavy load. In fact, it is only near the usable bandwidth that the TTRT has any effect on the response time.

Table 1
 Maximum Access Delay and Efficiency as Functions of the TTRT

TTRT	Maximum Access Delay (seconds)			Efficiency (percent)		
	Typical	Big 100	Largest	typical	Big 100	
Largest	20 SAS 4 km	SAS 200 km	500 DAS 20 km	20 SAS 4 km	SAS 20 km	50 DAS 200 km
4	0.08	0.40	4.00	98.94	71.87	49.55
8	0.15	0.79	8.00	99.47	85.92	74.77
12	0.23	1.19	11.99	99.65	90.61	83.18
16	0.30	1.59	15.99	99.74	92.95	87.38
20	0.38	1.98	19.98	99.79	94.36	89.91
165	3.14	16.34	164.84	99.97	99.32	98.78

To summarize the results presented so far, if the FDDI load is below saturation, the TTRT has little effect. At saturation, a larger value of TTRT gives a larger usable bandwidth and therefore increased efficiency. But a longer TTRT also results in longer maximum access delays. The selection of the TTRT requires a trade-off between these two requirements. To facilitate making this trade-off, the

the Largest ring is poor (50 percent). A very large value of TTRT, such as 165 ms, is also undesirable, because it results in long maximum access delays. The 8-ms value is the most desirable, since it yields 75 percent or more efficiency on all configurations and results in a maximum access delay of less than one second on Big rings. Eight milliseconds is, therefore, the recommended default TTRT.

two performance metrics for the three configurations are listed in Table 1. TTRT values in the allowable range of 4 to 165 ms are shown. The data shows that a very small value of TTRT, such as 4 ms, is undesirable, because the resulting efficiency on

Problems with a Large TTRT
There are three additional reasons for preferring an 8-ms TTRT over a large TTRT such as 165 ms. First, a large TTRT allows a station to receive a large number of frames back-to-back. To operate in such an environment, all

Performance Analysis of FDDI

adapters must be designed with large receive buffers. Although memory is not considered an expensive part of a computer, its cost is significant for low-cost components such as adapters. The board space for the additional memory required by choosing a larger TTRT is considerable as are the bus holding times required for such large back-to-back transfers.

Second, a very large TTRT results in an exhaustive service discipline (i.e., all frames are transmitted in one token capture), which has several known drawbacks. For example, exhaustive service is unfair. Frames coming to higher load stations have a greater chance of finding the token during the same transmission opportunity, whereas frames arriving at low load stations may have to wait. Thus, the response time is inversely dependent upon the load, i.e., higher-load stations yield lower response times and vice versa.[9].

Third, with exhaustive service, the response time of a station is dependent upon station location with respect to that of high-load stations. The station immediately downstream from a high-load station may

Parameters Other Than The TTRT That Affect Performance

Many parameters other than the TTRT affect the performance of a network. This section discusses four configuration and workload parameters: the extent of the ring, the total number of stations, the number of active stations, and the frame size.

Extent of the Ring

The total length of the fiber is called the extent of the ring. The maximum allowed extent on an FDDI LAN is 200 km. Figures 6 and 7 are graphs illustrating the efficiency and maximum access delay as functions of the extent. A star-shaped ring with all stations at a fixed radius from the wiring closet is assumed. The total cable length, shown along the horizontal axis, is twice the radius times the number of stations. As is evident from the figures, rings with a larger extent have a slightly lower efficiency and a longer maximum access delay than those with smaller extents.

Note that in Figure 7, the increase in maximum access delay for each configuration is not apparent due to the semilogarithmic scale.

obtain better throughput
than the one immediately
upstream.

Total Number of Stations

The total number of stations on a ring includes active and inactive stations. In general, increasing the number of stations adds to the ring latency because of the additional fiber length and additional station delays. Thus, the number of stations affects the efficiency and maximum access delay in

a way similar to that of the extent; a ring that contains a larger number of stations than another has a lower efficiency and a longer maximum access delay. In addition, a large number of stations on a ring increases the bit-error rate. Consequently, large rings are not desirable.

Number of Active Stations

As the number of active stations, i.e., MACs, increases, the total load on the ring increases. Figures 8 and 9 show the ring performance as a function of the number of active MACs on the ring. We considered a maximum-size ring with a TTRT value of 8 ms for the analysis. The figures show that increasing the number of active MACs has

Frame Size

Frame size does not appear in the simple models of efficiency and maximum access delays, because frame size has little impact on FDDI performance. In our analysis, we assumed that transmission stops at the instant the THT expires; however, the standard allows stations to complete the transmission of the last frame.

a slight positive effect on the efficiency, but considerably increases the maximum access delay. Therefore, it is preferable to keep a minimal number of active stations on each ring by segregating small groups on separate rings.

Digital Technical Journal Vol. 3 No. 3 Summer 1991

Performance Analysis of FDDI

The extra time used by a station after THT expiry is called asynchronous overflow. Assuming all frames are of fixed size, let F denote the frame transmission time. During every transmission opportunity, an active station can transmit as many as k frames:

$$k = \frac{T-D}{F}$$

Here, $\lceil \cdot \rceil$ is used to denote rounding up to the next integer value. The transmission time is equal to k times F , which is slightly more than T minus D . With asynchronous overflow, the modified efficiency and maximum access delay formulas become

$$\text{Efficiency} = \frac{nkF}{n(kF+D)+D}$$

Notice that substituting $kF = T - D$ in the above equations results in Equations (1) and (2).

Figures 10 and 11 show the efficiency and the maximum access delay as functions of the frame size. Frame size has only a slight effect on these metrics.

- o The time to process a frame increases only slightly with the size of the frame. A larger frame size results in fewer frames and, hence, in less processing at the host.

Overall, we recommend using as large a frame size as the reliability

considerations allow.

Summary

Although many parameters affect the performance of an FDDI ring network, the target token rotation time (TTRT) is the key parameter that network managers can control to optimize this performance. We analyzed the effect of other parameters such as the extent of the ring (the

length of the cable), the total number of stations, the number of active

stations, and frame size.

From our data we concluded the following:

- o Rings with a large extent and those with a

large number of stations are undesirable because they yield a longer maximum access delay and have only a slight positive effect on the

Larger frame sizes do have the following effects:

- o The probability of error is greater in a larger frame.
- o Since the size of protocol headers and trailers is fixed, larger frames require less protocol overhead.

efficiency of the ring.

- o It is preferable to minimize the number of active stations on a ring to avoid increasing the maximum access delay.

- o A large frame size is desirable, taking into consideration the acceptable probability of error.

The value of TTRT does not significantly affect the response time unless the load is near saturation.

Under very heavy load, response time is not a suitable metric. Instead, maximum access delay, i.e., the time between wanting to transmit and being able to do so, is more meaningful.

A larger value of TTRT

improves the efficiency, but it also increases the maximum access delay. A good trade-off is provided by setting TTRT at 8 ms. Since this value provides good performance for all ranges of configurations, we recommend that the default value of TTRT be set at 8 ms.

References

1. F. Ross, "An Overview of FDDI: The Fiber Distributed Data Interface," IEEE Journal on Selected Areas in Communications, vol. 7, no. 7 (September 1989): 1043-51.
2. Fiber Distributed Data Interface (FDDI)
3. Token Ring Access Method and Physical Layer Specifications, ANSI/IEEE Standard 802.5-1985, ISO/DIS 8802/5 (New York: The Institute of Electrical and Electronics Engineers, Inc., 1985).
4. R. Grow, "A Timed-token Protocol for Local Area Networks," Proceedings of the IEEE Electro '82 Conference on Token Access Protocols, Paper 17/3, Boston, MA (May 1982).
5. R. Jain, The Art of Computer Systems Performance Analysis, ISBN 0471-50336-3 (New York: John Wiley & Sons, 1991).
6. D. Dykeman and W. Bux, "Analysis and Tuning of the FDDI Media Access Control Protocol," IEEE Journal on Selected Areas in Communications, vol. 6, no. 6 (July 1988): 997-1010.
7. J. Ulm, "A Timed-token Ring Local Area Network and Its Performance Characteristics," Proceedings of the Seventh IEEE Conference on Local Computer Networks (February 1982): 50-56.
8. R. Jain, "Error

- Token Ring Media
Access Control (MAC),
ANSI X3.139-1987 (New
York: American National
Standards Institute,
1987).

Characteristics of
Fiber Distributed Data
Interface (FDDI),"
IEEE Transactions on
Communications, vol.
38, no. 8 (August 1990):

1244-1252.

Digital Technical Journal Vol. 3 No. 3 Summer 1991

Performance Analysis of FDDI

9. W. Bux and H. Truong, "Mean-delay Approximation for Cyclic-service Queueing Systems," Performance Evaluation, vol. 3 (Amsterdam: North-Holland, 1983): 187-196.

=====
Copyright 1991 Digital Equipment Corporation. Forwarding and copying of this
article is permitted for personal and educational purposes without fee
provided that Digital Equipment Corporation's copyright is retained with the
article and that the content is not modified. This article is not to be
distributed for commercial advantage. Abstracting with credit of Digital
Equipment Corporation's authorship is permitted. All rights reserved.
=====