Cycle Time Properties Of The FDDI Token Ring Protocol

Kenneth C. Sevcik
Marjory J. Johnson

August, 1985
(revised March, 1986)

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 85.10

(NASA-CR-187298) CYCLE TIME PROPERTIES OF THE FDDI TOKEN RING PROTOCOL (Research Inst. for Advanced Computer Science) 29 p

N90-71367

Unclas

00/61 0295385

RIACS

Research Institute for Advanced Computer Science

Cycle Time Properties Of
The FDDI Token Ring Protocol †

Kenneth C. Sevcik
Marjory J. Johnson

Research Institute for Advanced Computer Science
NASA Ames Research Center

ABSTRACT

The FDDI Token Ring Protocol controls communication over fiber optic rings with transmission rates in the range of 100 megabits per second. It is intended to give guaranteed response to time-critical messages by using a "timed token" protocol, in which non-critical messages may be transmitted only if recent movement of the token among stations has been sufficiently fast relative to a "target" token rotation time (TTRT).

In this paper, we prove two important properties of the protocol. The first is that the *average* token cycle time is bounded above by the TTRT, and the second is that the *maximum* token cycle time is at most twice the TTRT. Each property is treated first under the assumption that all overheads are negligible, and second with certain sources of overhead taken into account explicitly. It is found that the proposed standard protocol can be improved for situations of practical interest by a slight modification.

Current addresses and phone numbers of the authors:

Professor K. C. Sevcik
Computer Systems Research Institute
University of Toronto
Toronto, Ontario
CANADA M5S 1A4
(416) 978-6219

Dr. Marjory J. Johnson
RIACS – Mail Stop 230-5
NASA Ames Research Center
Moffett Field, California
U.S.A. 94035
(415) 694-6922

† This work was carried out while the first author was spending his sabbatical leave from the University of Toronto at the Research Institute for Advanced Computer Science at NASA Ames Research Center.

Cycle Time Properties Of The FDDI Token Ring Protocol

1. Introduction

Communication technology now makes it possible to support high data transmission rates at relatively low cost. In particular, optical fiber can be used as the medium in local area networks with data rates in the range of 100 megabits per second. Unfortunately, local area network topologies and communication protocols that work well with lower speed media are not necessarily appropriate when the data transmission rate is scaled up by approximately an order of magnitude. Recognizing this fact, an ANSI sub-committee (ANSI X3T9.5) has been working for the past two years on a proposed standard for a token ring protocol tailored to a transmission medium with transmission rate in the 100 megabits per second range. The protocol is referred to as the FDDI (Fiber Distributed Data Interface) Token Ring protocol. The proposal for the standard is now quite mature and nearly stable.

While numerous analyses of the performance of token ring protocols have been carried out and described in the literature, these have for the most part dealt with protocol variations of less complexity than FDDI. Konheim and Meister [17] provided an analysis of the case of a token ring joining symmetric stations. They obtained distributions of both delay and queue lengths using a discrete time model. Bux [5] studied several types of local area network protocols, including an exhaustive service symmetric token ring. He obtained the mean service time for a continuous time model as a limiting case of Konheim and Meister's results. A complete exact analysis of the two-station case was given by Boxma and Cohen [4, 9].

Berry and Chandy [3] treated the asymmetric case. They observed that, in the case of exhaustive service, the symmetric and asymmetric cases do not differ greatly. They analyzed the non-exhaustive case, using an M/G/1 model in which they adjusted the service time distribution to reflect the token passage times around the ring. An exact analysis of the asymmetric case was provided by Ferguson and Aminetzah [11].

Danthine and Vyncke [10] analyzed a slotted ring (Cambridge ring), a token ring and a token bus. They concluded that the token ring is best, particularly when the token is passed on immediately after transmission, rather than awaiting successful return of the message. Coffman and Gilbert [8] took a unique approach in which message arrivals occur continuously around the ring, rather than at discrete stations. This approach is a good approximation to the exhaustive service protocol with a very large number of stations on the ring. They obtained the distribution of the number of customers served on each cycle of the token. Takagi and Kleinrock [20] have provided an extensive overview of work on the analysis of polling systems (which includes the analysis of token rings).

The major feature that distinguishes the protocol of interest in this paper from token ring protocols that have been analyzed previously is the concept of a "timed token", which selectively allocates the right to transmit data among the stations depending in part on how rapidly the token progressed around the ring on the previous cycle. The "timed token" creates some dependencies among transmissions at various stations, and these dependencies complicate the analysis of the protocol's performance.

The basic ideas of the timed token protocol on which the FDDI protocol is based were first presented by Grow [12]. He distinguished two types of traffic. *Synchronous* traffic is a type of traffic that has delivery time constraints. Examples include voice and video transmissions, where delays in transmission can result in disruptions of the sound or picture signal. *Asynchronous* traffic has no such time constraints, or at least the time constraints are measured in units that are large relative to the token cycle time.

Here is a brief overview of how the "timed token" protocol works. The stations on the local area network choose, in a distributed fashion, a target token rotation time (TTRT). Basically, the TTRT is chosen to be sufficiently small that responsiveness requirements at every station will be met. The right to use network bandwidth for transmission of synchronous traffic is allocated among the stations in a manner such that it is guaranteed that network capacity is not exceeded. The token is then forced by the protocol to circulate with sufficient speed that all stations receive their allocated fractions of capacity for synchronous traffic. This is done by conditioning the right to transmit asynchronous data on the fact that the token has rotated sufficiently fast that it is "ahead of schedule" with respect to the target token rotation time. In essence, the TTRT value dictates a departure schedule for the token to pass from station to station, and asynchronous traffic can be transmitted only when doing so does not cause that schedule to be broken. (While synchronous transmissions can be initiated only if they will complete without making the token late, this is not the case for asynchronous traffic. For asynchronous traffic, the protocol requires only that the transmission be initiated before the token becomes late, and, once initiated, the transmission may run to completion. This creates an "asynchronous overrun" problem, which we shall ignore in Section 3, but then account for in Section 4.) Subsequently, Ulm [21] investigated the protocol described by Grow and determined its sensitivity to various parameters. He considered the effect of overheads and provided a number of graphs indicating the impact of various parameters on maximum transmission capacity.

The FDDI protocol has been developed from the ideas of Grow and Ulm. The formal description of the FDDI protocol is contained in several draft standard proposals under current development [1,2]. Additional papers have appeared dealing with FDDI directly. Ross and Moulton [19] give an overview of FDDI. Joshi and Iyer [16] describe the potential impact of FDDI. In two papers, Johnson describes the reliability mechanisms built into the FDDI protocol [14], and provides a proof of the robustness of the FDDI medium access scheme [15].

As well as describing the timed token protocol, Grow [12] and Ulm [21] included intuitive arguments supporting two fundamental properties of (a somewhat idealized version of) the protocol. These two properties are:

I. The average token cycle time in the absence of failures is at most the TTRT.

II. The maximum token cycle time in the absence of failures is at most twice the TTRT.

While Grow and Ulm assert that these properties hold for the timed-token protocol, neither formal proofs nor references are provided. Because the FDDI protocol is based on the same timed-token protocol, subsequent publications specifically describing the FDDI protocol also claim that the two properties hold [13,16]. A proof of the second property is included in Johnson's paper [15].

In this paper, we prove the first property formally, and state a proof of the second property in terms of the same notational framework used to prove the first property. From the derivations, it will become apparent that the protocol's restrictions on the transmission of asynchronous traffic can actually be relaxed somewhat while still guaranteeing properties (I) and (II). This leads to a variation on the FDDI protocol that is at least as easily implemented, that guarantees sufficient responsiveness and capacity for the transmission of synchronous traffic, and that may provide improved responsiveness to asynchronous transmissions in some situations where usage of synchronous allocations by stations is irregular. In section 2, we introduce our model and the notation to describe it. After treating an idealized situation in which several types of overhead are ignored in section 3, we will generalize the arguments in section 4 to take these overheads into account. Finally, we discuss the effect of the overheads on transmission capacity for some realistic values of network configuration parameters.

2. Model Definition

We will let N denote the number of stations in the network and we assume one network connection per station. Let T be the mutually agreed upon target token rotation time (which is chosen for reasons that will be made clear later to be one-half the minimum over all stations of the maximum tolerable time between visits of the token). Let f_i be the fraction of the network capacity allocated to station i for the transmission of its synchronous traffic. Clearly, the sum of these fractions cannot exceed one if the capacity guarantees are to be met. In fact, due to certain overheads, the maximum feasible sum of the allocated fractions is somewhat less than one.

Additional notation is needed to describe a particular behavior sequence in which traffic is transmitted under the FDDI protocol. We will index token visits to stations by a pair of subscripts, the first one, c, indicating the token cycle, and the second, i, indicating the station being visited. Our notation in the rest of the paper will often use the natural ordering of visits. Visit c,i is followed by visit $c,i+1$ if $1 \leq i < N$ and by visit $c+1,1$ if $i = N$. If $i=1$ when the subscript pair $c,i-1$ is used to denote the visit before c,i, then $c,i-1$ should be taken to be $c-1,N$. Similarly, if $i=N$, then the pair $c,i+1$ should be taken to be $c+1,1$. We will use these pairs to index visits even in summations. For example,

$$\sum_{j,k=w,x}^{y,z} Q_{j,k} = Q_{w,x} + \cdots + Q_{w,N} + Q_{w+1,1} + \cdots + Q_{y-1,N} + Q_{y,1} + \cdots + Q_{y,z}$$

is the sum of the quantity Q for all the visits starting with the cycle w visit to station x, and ending with the cycle y visit to station z.

Let $g_{c,i}$ and $a_{c,i}$ respectively be the times spent transmitting synchronous and asynchronous traffic on the cth visit to station i. The sum of $g_{c,i}$ and $a_{c,i}$ (plus some overhead quantities to be specified later) is the duration of the cth visit to station i, and will be denoted by $v_{c,i}$. We will assume (reasonably) that in each visit all synchronous transmission precedes any asynchronous transmission. The length of a full token cycle ending with the cth visit to station i is given by

$$C_{c,i} = \sum_{j,k=c-1,i+1}^{c,i} v_{j,k}$$

In the FDDI protocol, stations are allowed to transmit asynchronous messages only if the length of the cycle preceding the visit is less than the TTRT. In this case, where $T - C_{c,i-1} > 0$, we will say that the token is *early* on visit c,i. In fact, the specification of the FDDI protocol indicates that asynchronous transmission is allowed only if the token is not only early, but also any accumulated lateness on the previous cycle has been regained. We let $L_{c,i}$ denote the accumulated lateness of the token at its cth arrival to station i:

$$L_{c,i} = \max\,[0, C_{c,i-1} + L_{c-1,i} - T] \quad \text{with } L_{0,i} = 0 \text{ for all } i.$$

Then asynchronous transmission can be initiated by station c during the ith token visit if and only if $T - C_{c,i-1} - L_{c-1,i} > 0$. While Grow and Ulm and others have all suggested that it is necessary to work off accumulated lateness from the previous cycle, our results will show that that is not the case. By using the less restrictive transmission constraint for asynchronous traffic, $T - C_{c,i-1} > 0$, all stations are still guaranteed to receive their allocated bandwidth with acceptable frequency, while responsiveness to asynchronous requests may be improved in situations where stations make irregular use of their allocated bandwidth for synchronous traffic.

In order to compare the actual pace of the token's cycles with the target pace, we define some additional quantities. Let $P_{c,i}$ be the *pace* time (ignoring overheads once again) at the end of the cth visit to station i. It is defined by:

$$P_{c,i} = T \times [(c-1) + \frac{1}{A} \sum_{j=1}^{i} f_j] \quad \text{where } A = \sum_{j=1}^{N} f_j \tag{2.1}$$

(In section 4, we will generalize this definition of pace to take overheads into account.) Let $R_{c,i}$ be the actual realized time at the end of the cth visit to station i. Then

$$R_{c,i} = \sum_{j,k=1,1}^{c,i} v_{j,k}$$

The difference between the pace time and the realized time,

$$G_{c,i} = P_{c,i} - R_{c,i}$$

represents the time gained by the token over a steady pace of each cycle having length exactly T. It will be our intent to show that the FDDI protocol guarantees that this gain is generally increasing in time, indicating that the token indeed keeps up with the pace of at least one cycle per TTRT on average.

3. Negligible Overhead Case

In this section, we will establish the protocol properties stated earlier while ignoring any sources of overhead. In section 4, overhead considerations will be taken into account.

Here are the rules under which the protocol operates in the absence of overhead.

(1) The synchronous transmission in a visit may not exceed T times the station's allocated bandwidth fraction:

$$g_{c,i} \le F_i = f_i \times T$$

(2) The asynchronous transmission in a visit may not exceed the earliness of the token in its arrival to the station:

$$a_{c,i} \le max[0 \ , \ T - C_{c,i-1}]$$

(3) The total over all stations of the fraction of bandwidth allocated for synchronous traffic may not exceed one:

$$0 \leq A = \sum_{j=1}^{N} f_j \leq 1$$

(4) In the first cycle, there is no transmission, either synchronous or asynchronous ($g_{1,j} = a_{1,j} = 0$), and in the second cycle, there is no asynchronous transmission ($a_{2,j} = 0$):

$$R_{1,k} = 0 \quad \text{for } k = 1, 2, \cdots, N$$

$$R_{2,k} = \sum_{j=1}^{k} g_{2,j} \leq \sum_{j=1}^{k} F_j \quad \text{for } k = 1, 2, \cdots, N$$

Note that rule (2) above represents a weaker restriction on asynchronous transmission than the one suggested in the FDDI proposal, since "lateness" is not accumulated from cycle to cycle.

An important point to notice is that operation of the protocol depends only on local timings by local clocks. The amount of synchronous transmission on a visit to a station is constrained by a clock at that station. Similarly, the allowable amount of asynchronous transmission is calculated from the length of time between successive token arrivals to the station. The symbol $C_{c,i-1}$ is the time between the start of the $c-1$st visit and the cth visit to station i as observed by the clock at station i. Thus, our arguments will assume that the clocks at all stations run at precisely the same constant rate, but we need not assume that they are synchronized.

3.1. Bound On Average Token Cycle Time

In this section, we prove that the protocol rules guarantee that the average token cycle time is bounded above by the TTRT. The first step in the proof is the following lemma, which shows that the token never falls behind the TTRT driven pace.

Lemma 1:

With the token pace defined by eq. (2.1), the protocol based on rules (1), (2), (3), and (4) guarantees that

$$G_{c,i} \geq 0 \quad \text{for } i = 1, 2, \cdots, N, \text{ and } c = 1, 2, \cdots$$

Proof:

The proof is by contradiction. Assume that visit x,y is the first visit for which the gain is negative. Then $G_{x,y} < 0$, but $G_{j,k} \geq 0$ for $1,1 \leq j,k < x,y$. By protocol rule (4), x must be at least 3, because

$$G_{1,k} = P_{1,k} - R_{1,k} = \frac{1}{A}\sum_{j=1}^{k} F_j - 0 \geq 0$$

and

$$G_{2,k} = P_{2,k} - R_{2,k} = T + \left[\frac{1}{A}\sum_{j=1}^{k} F_j - \sum_{j=1}^{k} g_{2,j}\right] \geq T > 0$$

since protocol rules (1) and (3) together guarantee that the term in brackets is non-negative.

Now consider two cases.

Case 1: $g_{x,y} + a_{x,y} \leq F_y$

In this case,

$$G_{x,y} = G_{x,y-1} + [\frac{F_y}{A} - (g_{x,y} + a_{x,y})]$$

Since $A \leq 1$, the term in brackets is non-negative, and $G_{x,y} \geq G_{x,y-1}$.

Case 2: $g_{x,y} + a_{x,y} > F_y$

By protocol rule (1), $F_y - g_{x,y} \geq 0$, so $a_{x,y} > 0$. Considering protocol rule (2) also, we have

$$0 < a_{x,y} \leq T - C_{x,y-1}$$

Consider the relationship between $G_{x,y}$ and $G_{x-1,y-1}$.

$$G_{x,y} = G_{x-1,y-1} + (T + \frac{F_y}{A}) - (C_{x,y-1} + g_{x,y} + a_{x,y})$$

or, regrouping,

$$G_{x,y} = G_{x-1,y-1} + [(T - C_{x,y-1}) - a_{x,y}] + [\frac{F_y}{A} - g_{x,y}]$$

Since both terms in brackets are non-negative, $G_{x,y} \geq G_{x-1,y-1}$.

We assumed that visit x,y was the earliest for which gain was negative, yet in each of the two cases above, we showed that $G_{x,y}$ was no less than the gain at an earlier visit. This contradiction shows that our assumption that $G_{x,y} < 0$

for some x, y must be false.

QED

With the preceding lemma, we can establish an upper bound on the average token cycle time.

Theorem 1:

For the protocol based on rules (1), (2), (3), and (4), with negligible overheads, the average token cycle time in the absence of failures is no greater than the target token rotation time.

Proof:

By lemma 1,

$$G_{c,i} = P_{c,i} - R_{c,i} \geq 0 \quad \text{for all } c, i$$

Therefore,

$$T \times \left[(c-1) + \frac{1}{A} \sum_{j=1}^{i} f_j \right] - \sum_{j,k=1,1}^{c,i} v_{j,k} \geq 0$$

So,

$$T \geq \frac{\sum_{j,k=1,1}^{c,i} v_{j,k}}{\left[(c-1) + \frac{1}{A} \sum_{j=1}^{i} f_j \right]}$$

Because the right-hand side of the inequality above is total time taken divided by the number of completed cycles, we conclude that the average token cycle time is less than the target token rotation time.

QED

While we used $\frac{1}{A} \sum_{j=1}^{i} f'_j$ as the measure of a partially completed cycle (consisting of visits to only the first i stations) in the proof above, alternative measures leading to essentially the same result also exist.

3.2. Bound On Maximum Token Cycle Time

The originators of the timed token protocol upon which FDDI is based have asserted that the maximum token cycle time under the protocol is twice the target token rotation time (TTRT). This fact has been proven by Johnson in the context of demonstrating the robustness of the FDDI protocol [15]. In this section, we present a proof of the theorem in the notational framework of this paper. We begin with a lemma of broader utility in analyzing the FDDI protocol.

Lemma 2:

With the protocol based on rules (1), (2), (3), and (4), let j,k be the last visit at or before c,i for which the token arrived early (so that $T - C_{j,k} > 0$). If j,k is no later than $c-1,i$ then

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} g_{x,y}$$

If j,k is after $c-1,i$ then

$$C_{c,i} \le T + \sum_{x,y=j,k}^{c,i} g_{x,y} - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

Proof:

Consider first the case where there is no early visit in the cycle preceding c,i, so j,k is no later than $c-1,i$. If no visit in that cycle was early, then $a_{x,y} = 0$ for visits x,y from $c-1,i+1$ to c,i, and

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} g_{x,y}$$

Now consider the case where some visit in the cycle before c,i is early, so that j,k is after $c-1,i$. Then the second protocol rule guarantees that

$$a_{j,k} \le \max[0, T - C_{j,k-1}]$$

Because visit j,k was picked for being a visit on which the token arrived early, we know that

$$0 \le a_{j,k} \le T - C_{j,k-1} = T - \sum_{x,y=j-1,k}^{j,k-1} (g_{x,y} + a_{x,y})$$

Adding the quantity

$$g_{j,k} + \sum_{x,y=c-1,i+1}^{j,k-1} (g_{x,y} + a_{x,y})$$

to both $a_{j,k}$ and its upper bound given above, we obtain

$$\sum_{x,y=c-1,i+1}^{j,k} (g_{x,y} + a_{x,y}) \le T + g_{j,k} - \sum_{x,y=j-1,k}^{c-1,i} (g_{x,y} + a_{x,y})$$

Knowing that $a_{x,y} = 0$ for all visits after j,k since the token is not early on those visits, we add $\sum_{x,y=j,k+1}^{c,i} g_{x,y}$ to get

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + a_{x,y}) \le T + \sum_{x,y=j,k}^{c,i} g_{x,y} - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

QED

Lemma 2 above allows us to prove the following theorem.

Theorem 2:

In the absence of failures and with negligible overheads, token cycle time is bounded above by twice the target token rotation time.

Proof:

Consider visit c,i. If there was no early visit from $c-1,i+1$ to c,i, then, by lemma 2,

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} g_{x,y}$$

Since

$$\sum_{x,y=c-1,i+1}^{c,i} g_{x,y} \le \sum_{j=1}^{N} F_j = T \times \sum_{j=1}^{N} f_j$$

and

$$\sum_{j=1}^{N} f_j \le 1$$

then

$$C_{c,i} \le T < 2 \times T$$

On the other hand, if there was some early visit in the cycle ending with visit c,i, then by lemma 2,

$$C_{c,i} \le T + \sum_{x,y=j,k}^{c,i} g_{x,y} - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

where j,k is the index of the last early visit before c,i.

Since $v_{x,y} \geq 0$ for all x,y, and visit j,k is no earlier than visit $c-1,i+1$,

$$C_{c,i} \leq T + \sum_{x,y=j,k}^{c,i} g_{x,y} \leq T + \sum_{x,y=c-1,i+1}^{c,i} g_{x,y} \leq T + \sum_{j=1}^{N} F_j \leq 2 \times T$$

Therefore, in either case, the token cycle time is bounded above by twice the target token rotation time.

QED

4. The Effect of Overheads

4.1. Overhead Sources

In the previous section, we ignored all forms of overhead in obtaining the results that the average and maximum token rotation times are at most the TTRT and twice the TTRT, respectively. In this section, we will examine the kinds of overhead that must be taken into account, and we will see how the overheads influence the results obtained earlier. (These overheads and their impact on the protocol are discussed more thoroughly by Johnson [15].)

Here are the five types of overhead that we will consider.

(1) Medium propagation delay.
To the visit times at station i, we will add the propagation delay, p_i, between stations i and the next station on the ring. We will let $P = \sum_{i=1}^{N} p_i$ be the total propagation delay around the ring.

(2) Token transmission time.
Each visit must include the time required to transmit the token, X.

(3) Station latency.
At each station, messages pass through a buffer causing a delay, $e_{c,i}$, which is at most E.

(4) Capture delay.
After a station captures the token, there may be a delay, $d_{c,i}$, which is at most D, before transmission actually begins.

(5) Asynchronous overrun.
While synchronous message transmissions can only be started if they will

fit within the time allocated for synchronous transmission, asynchronous transmissions are not similarly constrained. An asynchronous transmission can be initiated right up until the time the token is due to be passed on, and these transmissions can then be completed. This overrun, $o_{c,i}$, can be as large as O, the time required to send a message of maximum allowable length.

Taking these forms of overhead into account, the length of visit c,i becomes

$$v_{c,i} = g_{c,i} + a_{c,i} + z_{c,i} + o_{c,i} + p_i$$

$$\text{where } z_{c,i} = d_{c,i} + e_{c,i} + X$$

Let $Z = D + E + X$ represent the upper bound on $z_{c,i}$. It is also necessary to generalize the definition of token pace in light of the overheads. We let

$$P_{c,i} = T \times [(c-1) + \frac{1}{A} \sum_{j=1}^{i} f_j] + \sum_{j=1}^{i} p_j + i \times Z \tag{2}$$

With these generalized definitions of $v_{c,i}$ and $P_{c,i}$, theorems 1 and 2 of the preceding section no longer hold. We will now investigate what changes to the protocol definition are required in order to obtain analogous theorems that take the overheads into account.

4.2. Asynchronous Overrun Problem

First, we point out an example in which the basic protocol fails to retain the bound on maximum token rotation time. For clarity in this example, we will assume that all the overheads are negligible except the asynchronous overrun. Consider a situation where no station has transmitted in the previous cycle. Then the token has arrived early by T. The first station can then do $F_{c,1}$ worth of synchronous transmission followed by as much as $(T - F_{c,1}) + o_{c,1}$ worth of asynchronous transmission (including the overrun). If all subsequent stations then use their full allocation of synchronous transmission, then the cycle length becomes

$$C_{c,N} = F_{c,1} + (T - F_{c,1}) + o_{c,1} + \sum_{j=2}^{N} F_{c,j}$$

Since $F_{c,1}$ may be as small as zero and $o_{c,1}$ may be as large as O, then the bound on maximum token rotation time of $2 \times T$ can be assured in the allocations of the $F_{c,j}$'s only if

$$\sum_{i=1}^{N} F_i \leq T - O$$

Such a reduction in the allocated synchronous bandwidth is the means suggested in the FDDI protocol to account not only for asynchronous overrun, but also for all other forms of overhead. The FDDI protocol does retain the property that average token rotation time is at most the TTRT, but it is able to do so only by retaining "lateness" (according to the stronger form of protocol rule (2)) from one cycle to the next so that any asynchronous overruns are compensated for by cycles less than TTRT in length before any additional asynchronous transmission is allowed. There is no upper bound on how long it might take to compensate for the overrun since all stations might require their full synchronous allocation for an unlimited number of successive cycles.

Rather than reducing the total amount of synchronous allocation in order to account for asynchronous overrun as is suggested by the FDDI protocol proposal, an alternative is to deal with the possibility of overrun directly in allowing asynchronous transmission. One approach is to allow the initiation of an asynchronous transmission only if the message is sufficiently short that transmission will be completed without any overrun. A less graceful alternative would be to allow asynchronous messages to be initiated right up until the time the token must be passed on, but to avoid overrun by aborting transmissions in the middle when necessary.

If implementation considerations make both of these alternatives undesirable, then another possibility is to account for the possibility of an overrun in determining the time for which a station is allowed to initiate the transmission of asynchronous messages. Specifically, if it is required that

$$a_{c,j} \leq \max\,[0, T - C_{c,j-1} - O]$$

then it is guaranteed that

$$a_{c,j} + o_{c,j} \leq T - C_{c,j-1}$$

and overruns cannot cause either of the two properties to be violated. In the remainder of this section, we will analyze a protocol based on this last approach.

The first rule by which the protocol operates is unchanged in the presence of overhead, but the other three must be replaced as follows.

(2') The duration of time in which the transmission of asynchronous messages may be initiated is constrained as follows:

$$0 \le a_{c,i} \le \max\,[0, T - C_{c,i-1} - O]$$

(3') The total over all stations of the fraction of the bandwidth allocated for synchronous traffic must satisfy

$$0 \le A = \sum_{j=1}^{N} f_j \le 1 - \frac{(N \times Z + P)}{T}$$

(4') In the first cycle, there is no transmission, either synchronous or asynchronous ($g_{1,i} = a_{1,i} = o_{1,i} = 0$), and in the second cycle, there is no asynchronous transmission ($a_{2,i} = o_{2,i} = 0$):

$$R_{1,i} = \sum_{j=1}^{i} (z_{1,j} + p_j) \le i \times Z + \sum_{j=1}^{i} p_j \qquad \text{for } i = 1, 2, \cdots, N$$

$$R_{2,i} = R_{1,N} + \sum_{j=1}^{i} (g_{2,j} + z_{2,j} + p_j) \le (N+i) \times Z + P + \sum_{j=1}^{i} (F_j + p_j)$$

$$\text{for } i = 1, 2, \cdots, N$$

4.3. Average Token Cycle Time

Taking overheads into account, we can obtain a result analogous to theorem 1, using the following lemma.

Lemma 1':

With the token pace defined by equation (2), the protocol based on rules (1), (2'), (3'), and (4'), guarantees that $G_{c,i} \ge 0$ for $i = 1, 2, \cdots, N$, and $c = 1, 2, \cdots$

The proof of the lemma is given in the appendix.

With lemma 1' replacing lemma 1, the proof of theorem 1 in the previous section suffices without change to prove also the following theorem.

Theorem 3:

For the protocol based on rules (1), (2'), (3'), and (4'), taking overheads into account, the average token cycle time in the absence of failures is no greater than the target token rotation time.

4.4. Maximum Token Cycle Time

Similarly, in the case of the bound on maximum token cycle time, generalization of Lemma 2 and Theorem 2 is straightforward given the new protocol characterization.

Lemma 2':

With the protocol based on rules (1), (2'), (3'), and (4'), let j,k be the last visit at or before c,i for which the token arrived early by at least O. If j,k is no later than $c-1,i$, then

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + p_y + z_{x,y})$$

and if j,k is after $c-1,i$, then

$$C_{c,i} \leq T + \sum_{x,y=j,k}^{c,i} (g_{x,y} + p_y + z_{x,y}) - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

With this lemma, we are able to prove theorem 4.

Theorem 4:

For a protocol based on constraints (1), (2'), (3'), and (4'), and taking overheads into account, the token cycle time in the absence of failures is bounded above by twice the target token rotation time.

The proofs of both the lemma and the theorem appear in the appendix.

4.5. Impact of Overheads on Capacity

In the previous subsections, we used a constraint that

$$0 \leq A = \sum_{j=1}^{N} f_j \leq 1 - \frac{(N \times Z + P)}{T}$$

in order to prove that the desired protocol properties would still hold in the presence of overhead. Here we will consider some reasonable values for system parameters and investigate how large the target token rotation time must be in order to make it possible to allocate various fractions of the ring's transmission capacity to synchronous traffic. (Ulm's paper [21] includes several graphs that relate attainable ring utilization to various overhead quantities.) In the equation above, the quantity, A, represents the fraction of the network's transmission capacity that is allocated for synchronous traffic. The inequality can be transformed into a lower bound on target token rotation time given a specified total allocation of synchronous bandwidth shares among the stations:

$$T \geq \frac{N \times Z + P}{1 - A}$$

According to the figures taken from the FDDI standards documents, estimates of the overhead quantities are as follows. The medium propagation delay is approximately .005 milliseconds per kilometer. The sum of station latency, capture delay, and the token transmission time is also very close to .005 milliseconds (with the majority being capture delay). Finally, given that the maximum length of any message is 36,000 bits, the asynchronous overrun can be at most 360 milliseconds.

Taking the values above, the quantity $N \times Z + P$ can be estimated as .005 milliseconds times the sum of the ring length in kilometers and the number of stations on the ring. For several ring configurations of possible interest, Table 1 indicates the minimum target token rotation time that will permit allocation of various fractions of the ring transmission capacity to synchronous traffic.

Stations	Kilo-meters	$N \times Z + P$	Total Fraction Allocated To Synchronous					
			.5	.9	.95	.99	.995	.999
19	1	.10	.2	1.0	2	10	20	100
40	10	.25	.5	2.5	5	25	50	250
490	10	2.50	5.0	25.0	50	250	500	2500
900	100	5.00	10.0	50.0	100	500	1000	5000
1000	200	6.00	12.0	60.0	120	600	1200	6000

Table 1. Minimum TTRT's (in milliseconds) to permit various fractions of total capacity to be allocated to synchronous traffic.

From the table, we observe that when the number of stations is small (less than 20) and the length of the ring is short (less than 1 km.), then the target token rotation time can be as small as four milliseconds and still allow more than 95% of the ring's capacity to be allocated to synchronous traffic. This is important since a four millisecond sampling interval suffices to give high-quality, real-time voice transmission.

Even with a very long non-local ring (about 200 km.) and many stations (about 1000), the target token rotation time can still be less than one second while allocating about 99% of the ring's capacity to synchronous traffic.

Next we consider the impact of the asynchronous overrun time, which may be as large as .360 milliseconds. With the original FDDI approach, the allocated synchronous bandwidth is reduced to account for this potential overhead. Thus, fraction O/T of the ring capacity cannot be allocated for synchronous traffic due to potential asynchronous overrun. When the target token rotation time is chosen to be small (in order to carry real-time voice traffic, for example), this can represent nearly 10% of the network's capacity.

With our more liberal variation of the FDDI protocol, the potential for asynchronous overrun does not decrease the allowable allocation to synchronous traffic at all. Instead, the initiation of asynchronous messages is merely prohibited unless it is guaranteed that the transmission can be completed without causing asynchronous overrun. The cost of insulating the amount of allocated bandwidth for synchronous traffic from the effect of asynchronous overrun is to postpone transmission of asynchronous traffic unless the token is early at a station by at least O. (In some situations, the original FDDI protocol can delay asynchronous traffic still more, however, because it requires that accumulated "lateness" be worked off in cycle times less that TTRT before any further asynchronous transmission is permitted.)

5. Conclusions

This study has formally demonstrated some properties of the FDDI protocol that were previously believed on the basis of intuitive arguments. If all overhead sources are assumed negligible, then the protocol as described in the standards documents [1,2] (which was derived from Grow [12] and Ulm [21]) has two properties:

I. The average token rotation time is less than or equal to TTRT.

II. The maximum token rotation time is less than or equal to twice the TTRT.

Both these properties are important to the successful operation of the protocol. The first one guarantees that the long run bandwidth provided to station i is at least fraction f_i of the network's capacity. The second property guarantees that, in the absence of component failures, the time between a station's successive opportunities to transmit synchronous traffic will never exceed twice the target token rotation time.

Our study actually treated a modification of the original FDDI protocol that allows asynchronous transmission without requiring that "lateness" be carried forward from cycle to cycle. Assuming overheads to be negligible, our proofs showed that the restraint on the transmission of asynchronous traffic can be relaxed from

$$0 \leq a_{c,j} \leq \max[0, T - C_{c,j-1} - L_{c-1,j}]$$

to

$$0 \leq a_{c,j} \leq \max[0, T - C_{c,j-1}]$$

The relaxed version retains the two desired properties, but allows asynchronous transmissions to occur in some situations where they would be disallowed by the original protocol.

When overheads were considered, it was found that the FDDI protocol as described satisfied the constraint on average token rotation time (relying on the retention of "lateness" from cycle to cycle), but not the one on maximum cycle time. According to the papers and the draft standards documents describing the FDDI protocol, all forms of overhead would be accounted for by reducing the total synchronous allocation permitted. Thus, the constraint corresponding to rule (3') in our proposed protocol is

$$0 \leq A \leq \sum_{j=1}^{N} f_j \leq 1 - \frac{N \times Z + P + O}{T}$$

so that a lower total bandwidth can be allocated to synchronous traffic. Without this reduction in allocated synchronous bandwidth, token rotation times could exceed the target token rotation time by as much as O, the maximum duration of asynchronous overrun.

The modified protocol proposed here, which is based on the rule

$$0 \leq a_{c,j} \leq \max[0, T - C_{c,j-1} - O]$$

has several advantages over the standard protocol, which uses the rule

$$0 \leq a_{c,j} \leq \max[0, T - C_{c,j-1} - L_{c-1,j}].$$

These advantages of the modified protocol include:

(1) It guarantees both property I and property II without having to retain "lateness" from one cycle to the next.

(2) Depending on the amount of synchronous bandwidth allocated and the regularity with which it is used, the modified protocol may provide better service to asynchronous requests in the case where O is small relative to the token rotation time.

(3) It is easier to implement since "lateness" need not be retained from cycle to cycle.

Work is underway on the task of quantifying the performance of the FDDI protocol by determining estimates of or tighter bounds on the average token rotation time and on the average delivery time of a submitted message. The properties established in this paper are required to form the basis of the quantitative analysis.

6. Acknowledgement

We wish to thank Terry Grant, who directs the Advanced Data Networks Concepts project at NASA Ames Research Center, for his support of this research work at RIACS. Also, discussions with Floyd Ross, Sunil Joshi, and Larry Green have been helpful in understanding the course of development of the FDDI protocol.

7. References

[1] ANSI X3T9.5,
FDDI Token Ring -- Physical Layer Protocol Standard,
draft proposed American National Standard, ANSI X3T9.5/83-15, revision 9, April, 1985.

[2] ANSI X3T9.5,
FDDI Token Ring -- Media Access Control,
draft proposed American National Standard, ANSI X3T9.5/83-16, revision 8, March, 1985.

[3] Robert F. Berry and K. Mani Chandy,
Performance Models of Token Ring Local Area Networks,
Performance Evaluation Review, (special issue), pp. 266-274, August, 1983.

[4] O. J. Boxma,
Two Symmetric Queues with Alternating Service and Switching Times,
in *PERFORMANCE '84: Proc. of the 10th Int. Conf. on Computer Performance*, ed. by E. Gelenbe, North-Holland Publ., 1984, pp. 409-431.

[5] Werner Bux,
Local Area Networks: A Performance Comparison,
IEEE Transactions on Communications, vol. COM-29, no. 10, pp. 1465-1473, October, 1981.

[6] Werner Bux, F. H. Closs, K. Kuemmerle, H. J. Keller, and H. R. Mueller,
Architecture and Design of a Reliable Token Ring Network,
IEEE Transactions on Selected Areas in Communications, vol. SAC-1, no. 5, pp. 756-765, November, 1983.

[7] Ralph T. Carsten, et al.,
A Simplified Analysis of Scan Times in an Asymmetrical Newhall Loop With Exhaustive Service,
IEEE Transactions on Communications, vol. COM-25, no. 9, pp. 951-957, September, 1977.

[8] E. G. Coffman and E. N. Gilbert,
A Continuous Polling System With Constant Service Times,
draft report, AT&T Bell Labs, 1984.

[9] J. W. Cohen and O. J. Boxma,
The *M/G/1* Queue with Alternating Service Formulated as a Riemann-Hilbert Problem,
in *PERFORMANCE '81: Proc. of the 8th Int. Conf. on Computer Performance*, ed. by F. J. Kylstra, North-Holland Publ., 1982, pp. 181-1[illegible]9.

[10] A. Danthine and E. Vyncke,
Broad Site Local Wideband Network,
in *Proc. of the First Esprit Technical Week*, Brussels, September, 1984.

[11] M. J. Ferguson and Y. J. Aminetzah,
Exact Results for Nonsymmetric Token Ring Systems,

IEEE Transactions on Communications, vol. COM-33, no. 3, pp. 223-231 March, 1985.

[12] Robert M. Grow,
A Timed-token Protocol For Local Area Networks,
presented at *Electro '82*, Token Access Protocols (17 / 3), May, 1982.

[13] Venkatraman Iyer and Sunil Joshi,
FDDI's 100Mbps Protocol Improves On 802.5 Spec's 4Mbps Limit,
Electronic Data Networks (EDN), pp. 151-160, May, 1985.

[14] Marjory J. Johnson,
Reliability Mechanisms of the FDDI High Bandwidth Token Ring Protocol,
in *Proc. of the 10th Conf. on Local Computer Networks*, IEEE, October, 1985, pp. 124-133.

[15] Marjory J. Johnson,
Proof That Timing Requirements of the FDDI Token Ring Protocol Are Satisfied,
Technical Report 85.8, Research Institute for Advanced Computer Science, NASA Ames Research Center, August, 1985.

[16] Sunil Joshi and Venkatraman Iyer,
New Standards for Local Networks Push Upper Limits For Lightwave Data,
Data Communications, pp. 127-136, July, 1984.

[17] A. G. Konheim and B. Meister,
Waiting Lines and Times in a System With Polling,
Journal of the ACM, vol. 21, no. 3, pp. 470-490, July, 1974.

[18] P. J. Kuehn,
Multi-queue Systems With Non-exhaustive Cyclic Service,
Bell System Technical Journal, vol. 58, no. 3, pp. 671-698, March, 1979.

[19] Floyd E. Ross and R. Kirk Moulton,
FDDI Overview -- A 100 Mbit per Second Solution,
presented at *Electro '84*, (2 / 1), May, 1984.

[20] H. Takagi and L. Kleinrock,
Analysis of Polling Systems,

Technical Report TR87, IBM Japan Science Institute, January, 1985.

[21] John N. Ulm,
A Timed Token Ring Local Area Network and Its Performance Characteristics,
in *Proc. of the 7th Conf. on Local Computer Networks*, IEEE, February, 1982. pp. 50-56.

[22] Ray Weiss,
Standards Bring the Future One Step Closer,
Electronic Design, pp. 99-103, October 13, 1984.

8. Appendix

Lemma 1':

With the token pace defined by equation (2), the protocol based on rules (1), (2'), (3'), and (4'), guarantees that $G_{c,i} \geq 0$ for $i = 1, 2, \cdots, N$, and $c = 1, 2, \cdots$

Proof:

The proof is by contradiction. Assume that visit x,y is the first visit for which the gain is negative. Then $G_{x,y} < 0$, but $G_{j,k} \geq 0$ for $1,1 \leq j,k < x,y$. By protocol rule (4'), x must be at least 3. It can't be one because, with $g_{1,i} = a_{1,i} = o_{1,i} = 0$ for $i = 1, 2, \cdots, N$,

$$R_{1,k} \leq k \times Z + \sum_{j=1}^{k} p_j$$

while

$$P_{1,k} = \frac{1}{A} \sum_{j=1}^{k} F_j + \sum_{j=1}^{k} p_j + k \times Z$$

So,

$$G_{1,k} = P_{1,k} - R_{1,k} \geq \frac{1}{A} \sum_{j=1}^{k} F_j \geq 0$$

Also, x cannot be two because

$$P_{2k} = T + \frac{1}{A}\sum_{j=1}^{k} F_j + \sum_{j=1}^{k} p_j + k \times Z$$

and

$$R_{2k} \le (N+k) \times Z + P + \sum_{j=1}^{k} (F_j + p_j)$$

So,

$$G_{2k} = P_{2k} - R_{2k} = [T - (N \times Z + P)] + \left[\frac{1}{A}\sum_{j=1}^{k} F_j - \sum_{j=1}^{k} F_j\right] \ge T > 0$$

since protocol rules (1) and (3') guarantee that both terms in brackets are non-negative.

Now consider two cases.

Case 1: $g_{x\,y} + a_{x\,y} + o_{x\,y} \le F_y$

In this case,

$$G_{x\,y} = G_{x\,y-1} + (\frac{F_y}{A} + p_y + Z) - (g_{x\,y} + a_{x\,y} + o_{x\,y} + p_y + z_{x\,y})$$

$$= G_{x\,y-1} + [\frac{F_y}{A} - (g_{x\,y} + a_{x\,y} + o_{x\,y})] + [Z - z_{x\,y}]$$

Since $A \le 1$, both terms in brackets are non-negative, and $G_{x\,y} \ge G_{x\,y-1}$.

Case 2: $g_{x\,y} + a_{x\,y} + o_{x\,y} > F_y$

By protocol rule (1), $F_y - g_{x\,y} \ge 0$, so $a_{x\,y} + o_{x\,y} > 0$. But $o_{x\,y} > 0$ only if $a_{x\,y} > 0$, so $a_{x\,y} > 0$. Considering protocol rule (2') also, we have

$$0 < a_{x\,y} \le T - C_{x\,y-1} - O$$

Consider the relationship between $G_{x\,y}$ and $G_{x-1\,y-1}$.

$$G_{x\,y} = G_{x-1\,y-1} + (T + \frac{F_y}{A} + p_y + Z) - (C_{x\,y-1} + g_{x\,y} + a_{x\,y} + o_{x\,y} + p_y + z_{x\,y})$$

or, regrouping,

$$G_{x\,y} = G_{x-1\,y-1} + [(T - C_{x\,y-1} - o_{x\,y}) - a_{x\,y}] + [\frac{F_y}{A} - g_{x\,y}] + [Z - z_{x\,y}]$$

Since all terms in brackets are non-negative, $G_{x\,y} \ge G_{x-1\,y-1}$.

We assumed that visit x,y was the earliest for which gain was negative, yet in each of the two cases above, we showed that $G_{x,y}$ was no less than the gain at an earlier visit. This contradiction shows that our assumption that $G_{x,y} < 0$ for some x,y must be false.

QED

Lemma 2':

With the protocol based on rules (1), (2'), (3'), and (4'), let j,k be the last visit at or before c,i for which the token arrived early by at least O. If j,k is no later than $c-1,i$, then

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + p_y + z_{x,y})$$

and if j,k is after $c-1,i$, then

$$C_{c,i} \le T + \sum_{x,y=j,k}^{c,i} (g_{x,y} + p_y + z_{x,y}) - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

Proof:

Consider first the case where there is no early visit in the cycle preceding c,i, so j,k is no later than $c-1,i$. If no visit in that cycle was early by at least O, then $a_{x,y} = 0$ and consequently $o_{x,y} = 0$ for visits x,y from $c-1,i+1$ to c,i, and

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + p_y + z_{x,y})$$

Now consider the case where some visit in the cycle before c,i is early by at least O so that j,k is after $c-1,i$. Then the second protocol rule guarantees that

$$a_{j,k} \le \max[0, T - C_{j,k-1} - O]$$

Because visit j,k was picked for being a visit on which the token arrived early by at least O, we know that

$$0 \le a_{j,k} + o_{j,k} \le a_{j,k} + O \le T - C_{j,k-1} = T - \sum_{x,y=j-1,k}^{j,k-1} v_{x,y}$$

Adding the quantity

$$g_{j,k} + p_k + z_{j,k} + \sum_{x,y=c-1,i+1}^{j,k-1} v_{x,y}$$

to both $a_{j,k} + o_{j,k}$ and its upper bound given above, we obtain

$$\sum_{x,y=c-1,i+1}^{j,k} v_{x,y} \le T + (g_{j,k} + p_k + z_{j,k}) - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

Knowing that $a_{x,y} = o_{x,y} = 0$ for all visits after j,k since the token was not early by at least O on those visits, we have

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} v_{x,y} \le T + \sum_{x,y=j,k}^{c,i} (g_{x,y} + p_y + z_{x,y}) - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

QED

Theorem 4:

For a protocol based on constraints (1), (2'), (3'), and (4'), and taking overheads into account, the token cycle time in the absence of failures is bounded above by twice the target token rotation time.

Proof:

Consider visit c,i. If there was no visit from $c-1,i+1$ to c,i at which the token was early by at least O, then, by lemma 2',

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + p_y + z_{x,y})$$

Since

$$\sum_{x,y=c-1,i+1}^{c,i} g_{x,y} \le \sum_{j=1}^{N} F_j = T \times \sum_{j=1}^{N} f_j$$

then

$$C_{c,i} = \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + p_y + z_{x,y}) \le T \times \sum_{j=1}^{N} f_j + N \times Z + P$$

Applying protocol rule (3'),

$$\sum_{j=1}^{N} f_j \leq 1 - \frac{(N \times Z + P)}{T}$$

we have

$$C_{c,i} \leq T \times \sum_{j=1}^{N} f_j + (N \times Z + P) \leq T < 2 \times T$$

On the other hand, if there was some visit in the cycle ending with visit c,i at which the token was early by at least O, then by lemma 2',

$$C_{c,i} \leq T + \sum_{x,y=j,k}^{c,i} (g_{x,y} + p_y + z_{x,y}) - \sum_{x,y=j-1,k}^{c-1,i} v_{x,y}$$

where j,k is the index of the last early visit before c,i.

Since $v_{x,y} \geq 0$ for all x,y, and visit j,k is no earlier than visit $c-1,i+1$,

$$C_{c,i} \leq T + \sum_{x,y=j,k}^{c,i} (g_{x,y} + p_y + z_{x,y}) \leq T + \sum_{x,y=c-1,i+1}^{c,i} (g_{x,y} + p_y + z_{x,y})$$

$$\leq T + \sum_{j=1}^{N} F_j + (N \times Z + P)$$

By protocol rule (3') again,

$$\sum_{j=1}^{N} F_j + N \times Z + P \leq T$$

so,

$$C_{c,i} \leq 2 \times T$$

Therefore, in either case, the token cycle time is bounded above by twice the target token rotation time.

QED