

Performance Analysis of Slotted Ring Protocols in HSLAN's

MIRJANA ZAFIROVIC-VUKOTIC, IGNAS G. NIEMEGERES, DURK S. VALK

Abstract—The performance of a number of slotted ring protocols supporting integration of synchronous and asynchronous traffic in HSLAN's is evaluated in this paper as follows: the Cambridge Fast Ring, a variant of the Cambridge Fast Ring, and Orwell. The performance of their basic access mechanisms is compared and contrasted to the multiple token ring. The effect of a uniframe scheme for supporting synchronous traffic is examined. The influence of system parameters and performance of the integrated services slotted ring protocols is evaluated.

I. INTRODUCTION

THE slotted ring protocols originate from the Cambridge Ring [8]. They are suitable for high speed LAN's (HSLAN's) because of their good performance at high transmission rates [20]. Furthermore, the synchronous transmission scheme makes the implementation at high transmission rates in fiber optics easier. In this paper, the performance of the following MAC layer protocols or access mechanisms (AM's) for slotted rings is evaluated: the Cambridge Fast Ring (CFR) [17], Orwell [4], and a variant of the CFR (CFRV) [25]. Similar protocols can be found in other slotted rings, e.g., FXNET [3] and Upperbus [6]. Our analysis is applicable to these systems as well. Transmission rates in excess of 100 Mbits/s are assumed.

The performance of the basic slotted ring AM's is analyzed. This includes an analysis of the sensitivity to the configuration and workload parameters, and a comparison to the multiple token ring protocol and the uniframe scheme. Furthermore, integration of synchronous (e.g., voice, video) and asynchronous traffic (e.g., file transfer) is studied. The influence of the system parameters under both symmetric and asymmetric loads is evaluated. The performance analysis is conducted using simulation and analytical models. In the analytical modeling, the MAC service provider is regarded as a multiple cyclic server system with a limited or exhaustive service discipline.

The performance of slotted ring protocols in high speed networks has been studied in [4], [5], [12], and [13]. However, except for [13], these analyses concern telephone traffic only. Furthermore, no comparison of the various slotted ring protocols except for a limited study

in [13] is available yet. In this paper, we present such a comparative analysis for both synchronous and asynchronous traffic. The sensitivity to configuration and load parameters, the influence of the system parameters, and the effect of asymmetric loads is also studied. The analytical models and the stability conditions for the slotted rings on which the analysis is mainly based have recently been developed by ourselves in [22], [24], and [25].

The paper is organized as follows. Protocols are described in Section II. Issues related to the integration of synchronous and asynchronous traffic are presented and an introduction to the performance analysis is given in Section III. Notation is introduced in Section IV. The workload model is presented in Section V. Configuration and system parameters are chosen in Section VI. A sensitivity analysis and a comparison to the token passing AM and a study of the uniframe scheme is done in Section VII. The influence of the system parameters is examined in Section VIII. A comparative performance analysis of the slotted ring protocols supporting integrated services is presented in Section IX. Section X contains the conclusions.

II. SLOTTED RING PROTOCOLS

The ring is partitioned into equal length slots (Fig. 1). We assume that this is achieved by introducing a latency register located in a monitor station to virtually lengthen the ring to a multiple of the slot length. Let us call a LLC_PDU a packet and a MAC_PDU a minipacket.

Slots circulate around the ring and can be empty or full. A full slot is occupied by a minipacket. Stations are actively coupled to the ring: they repeat or modify the slots. A sending station may fill an empty slot by a minipacket. A full slot circulating around the ring reaches the destination station which reads it and passes it on to a higher layer.

There are two classes of basic slotted ring AM's depending on which station empties a full slot: the source station or the destination station. Furthermore, after emptying the slot there are two possibilities: the slot can be used by the station that empties it or it must be passed to the next downstream station.

If the source station releases a slot which was full (Fig. 2), there are two variants of this AM depending on the maximum number of slots that can be used by one station at a time. Either at most one slot at a time can be occupied by a station or more than one slot at a time can be carrying

Manuscript received November 1, 1987; revised February 19, 1988.
M. Zafirovic-Vukotic and I. G. Niemegeers are with the University of Twente, 7500AE Enschede, The Netherlands.
D. S. Valk is with Logica BV Nederland, 9711HZ Groningen, The Netherlands.

IEEE Log Number 8821337.

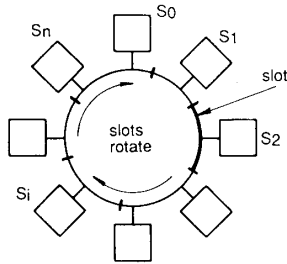
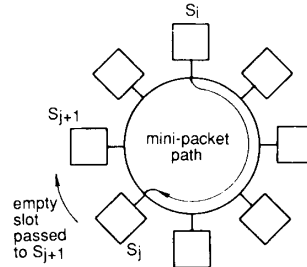
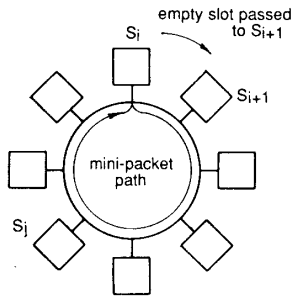


Fig. 1. The slotted ring structure.

Fig. 3. Path of a minipacket sent from S_i to S_j in the Orwell basic AM.Fig. 2. Path of a minipacket sent from S_i to S_j in the CFR and the CFRV basic AM.

minipackets from the same source. An example of the former protocol can be found in the CFR. Examples of the latter protocol are the CFRV, FXNET, and Upperbus. In all these protocols, the source station after emptying the slot has to release it and pass it on to the first downstream station. An exception are the channel slots in the CFR and in the CFRV which can be reused by the source.

If the destination station releases a slot which was full (Fig. 3) typically more than one slot at a time can be carrying minipackets from the same source. An example of such a protocol can be found in Orwell. After emptying the slot a destination station has to pass it on to the next station. Fairness is not guaranteed and precautions have to be taken to prevent a station from hogging the ring.

In the case of source release, an acknowledgment from the destination to the source station is provided using the same slot. If only one slot at a time can be used by a station an immediate retransmission of an erroneous minipacket is possible. In the case of destination release, no acknowledgment is provided at the MAC layer. In this paper, we do not study acknowledgment or retransmission mechanisms. We assume that each station is capable of using every empty slot that arrives and of reading every slot destined to itself.

Let us now describe the AM's of the CFR, the CFRV, and Orwell.

The CFR distinguishes two types of slots: normal and channel slots. Both kinds of slots are subject to source release. A normal slot must be passed by the source to the first downstream station. This way a periodical access of each station to this slot is ensured. A channel slot, on the

other hand, can be reused by the source. Channel slots are intended for large bandwidth or bursty sources. In the CFR, only one slot at a time can be used by a station. When we relax this condition and permit that more than one slot at a time can be used by a station, we obtain the CFRV. The CFR is a relatively inexpensive and simple LAN which has been designed for short ring lengths and moderate transmission rates, e.g., up to 100 Mbits/s. We compare it to the other slotted rings working at high transmission rates.

In Orwell, slots are released by the destination. It offers two service classes: class 1 for delay sensitive services, requiring a guaranteed bandwidth, e.g., synchronous traffic, and class 2 for delay tolerant services, e.g., asynchronous traffic. Access is organized in cycles to ensure fairness. Each station S_i in the ring has a counter D_i [4], indicating the number of packets allowed to be sent during a so called reset interval or a cycle. During a cycle each station can send up to $D_i - D_{\min}$ packets of class 1 (we denote this value by $D_{i,1}$) where D_{\min} provides a guaranteed (safeguard) bandwidth for class 2 packets (we denote this value by $D_{i,2}$). Since each station has independent counters, the throughput can be regulated according to the station's needs. After all the stations have used their counters or have nothing to send, a reset is issued which resets all the counters to their initial value and starts a new reset interval.

We distinguish three basic AM's which correspond to cases where in the CFR and in the CFRV only normal slots are used, and in Orwell it is assumed that reset intervals do not exist, i.e., $D_{i,1}$ and $D_{i,2}$ are infinite.

III. SCOPE OF THE PERFORMANCE ANALYSIS

The main characteristics of slotted rings in HSLAN's compared to their use in medium speed LAN's (transmission rate around 10 Mbits/s, e.g., IEEE P802.3, P802.4, and P802.5), which are relevant for modeling and performance analysis, are the following ones: 1) synchronous traffic sources and large bandwidth sources create most of the load [11]; such sources create a different type of load than the smaller bandwidth sources in medium speed LAN's, 2) the information field of a slot contains more bits in a HSLAN than in a medium speed LAN, so the number of slots per packet is typically smaller in HSLAN's than in medium speed LAN's for the same ap-

plication, and 3) typically, there are more slots in a HSLAN than in a medium speed LAN given the same configuration.

Synchronous traffic also demands a different quality of service (QoS) than asynchronous traffic. Small packet delays and a regular access to the medium have to be provided. There are two basic ways in which synchronous and asynchronous traffic can be integrated in a LAN: 1) by packetizing the synchronous traffic (e.g., a 64 kbit/s source produces 512 bit packets every 8 ms), and 2) by sending the synchronous traffic using a PCM-like transmission, i.e., one byte every 125 μ s. The former solution demands an upper bound on the delays for the synchronous traffic at the MAC layer and the latter one called a "uniframe scheme," requires a frame structure to support both PCM-like circuit switched channels and asynchronous traffic packet switching.

The CFR, CFRV, and Orwell implement the former way of integration. In Orwell, access is organized in cycles to provide fairness and ensure fulfillment of the QoS requirements for synchronous traffic. In the CFR and the CFRV, such a mechanism does not exist at the MAC layer. However, in these protocols there is a minimum bandwidth available to each station so the QoS given to the synchronous traffic can be guaranteed, provided synchronous traffic uses less than this minimum bandwidth and provided that it is given priority over asynchronous traffic at transmit time.

Examples of a uniframe scheme are FDDI II [14], Σ Network [7], and FXNET [3]. Let us now briefly describe it. In a frame of duration 125 μ s, there are a number of bytes used for synchronous transmission. Some are also used for (protocol) control information. Other bytes are available for asynchronous traffic using a LAN protocol: in our case, a slotted ring protocol or a token passing protocol. If a slotted ring AM is used, we assume as it has been done in FXNET that an integer number of slots fit into the remaining part of the frame.

In this paper, we first study the performance of the basic slotted ring AM's. Then we study the first type of integration of the two traffic classes in more detail and compare it to the uniframe method. A reason to study in detail the first type of integration is that the slotted ring protocols perform very well with short packets which are typical and desirable for synchronous traffic. Note that the merits of the two integration schemes depend not only on the performance of the MAC layer, but also on the technical and economical feasibility.

Only a few papers concerning performance analysis of slotted ring protocols in HSLAN's are available in the literature. In [4], it has been shown that the delays in Orwell are much smaller than in the token ring for voice only traffic. In [5], the performance of the Cambridge ring, used also only for telephony, has been evaluated. In [12], the blocking probability in Orwell for telephone traffic has been approximated. A scheme for integration of voice and data based on the Orwell basic AM has been proposed in [9]. The scheme has not been described in enough detail

to be used as a basis for further study. A limited performance comparison in [13] of the CFR, the CFRV basic AM's, and a simple (not a multiple) token passing ring operating at 100 Mbits/s shows that the slotted rings perform well with short packet lengths.

We define packet delay as the duration of the MAC layer service per packet, i.e., the packet delay from arrival at S_i till its complete delivery at the destination. We obtain the estimates of the expected packet delay at S_i . Simulations provide also the maximum packet delay which is the maximum packet delay observed during a simulation.

Let us now introduce the simulation and analytical models.

The simulation models are detailed models of the MAC protocols. They are written in SIMULA and are documented in [10] and [18]. Ninety-five percent level confidence intervals have been obtained on estimates of the expected delay except when the correlation between the batches was too large (see [18] and [2]).

The analytical models have been developed and tested by us in [22] for Orwell, in [24] for the CFR, and in [25] for the CFRV. An overview of analytical models is given in [21]. They provide exact stability conditions (though without a proof) and approximations of the expected packet delays. The accuracy of the models has been tested by comparing to simulation results for a number of cases thought to be representative of HSLAN's. The analytical models provide estimates that have a deviation up to 20 percent of the point estimates obtained by the simulations over a range of utilizations which go from 0.4 to 1.3 for Orwell, from 0.0 to 0.8 for the CFRV, and from 0.0 to 0.7 for the CFR (deviation even less than 10 percent for the CFR). At high utilizations the errors are larger, however, the approximations are sufficiently good and indicative for a comparative analysis. The exact solution of the token ring model as a cyclic server system with an exhaustive service discipline has been given in [16]. The same models are used for uniframe schemes. The load due to synchronous traffic and uniframe control information is accounted for by decreasing the available capacity for asynchronous traffic. This has been done by decreasing the transmission rate in the model. The error in delay estimates caused by this is less than a uniframe duration, i.e., 125 μ s.

IV. NOTATION

Let us introduce the following notation: n —number of stations minus one; S_i — i th station in the ring, $i = 0, 1, \dots, n$; w —transmission rate (Mbits/s); σ —duration of a slot (μ s); s —number of slots in the ring; λ_i —packet arrival rate at station S_i ($1/\mu$ s), $i = 0, \dots, n$; γ_i —the expected number of minipackets a packet is split into at S_i , $i = 0, \dots, n$; ρ —relative load, or the expected number of minipackets arriving in the system during σ time units; such that

$$\rho = \sum_{i=0}^n \lambda_i \gamma_i \sigma; \quad (1)$$

p_{ij} —an element of the packet communication source-to-destination matrix, $\|p_{ij}\|_{n+1 \times n+1}$ which represents the relative traffic intensity from source S_i to destination S_j , and

$$0 \leq p_{ij} \leq 1 \quad \text{and} \quad \sum_{j=0}^n p_{ij} = 1, \\ i, j = 0, \dots, n; \quad (2)$$

τ_{ij} —propagation time from S_i to S_j including the latency at station S_j (μs), $i, j = 0, \dots, n$; τ —slot rotation time or τ_{ii} for all i (μs), such that

$$\tau = \sigma\sigma. \quad (3)$$

V. WORKLOAD MODEL

Let us now specify a workload model representative for a HSLAN. It includes the arrival process of packets and of minipackets, the distribution of their lengths and the load specification.

Packets are segmented and MAC protocol control information (PCI) is added to form a number of minipackets. The information field of a minipacket in a slotted ring protocol has a constant length. The length of a minipacket is a constant and equals $\sigma \times w$ (bits) i.e., a slot length.

A study of the scenarios and workload models for a HSLAN conducted in [11] shows that it can be expected that in an office environment a great deal of the total load consists of synchronous traffic. We assume that synchronous traffic represents 80 percent and asynchronous traffic 20 percent of the total offered load at each station. Let us now specify the arrival process and the packet length distribution for synchronous and asynchronous traffic.

1) *Synchronous Traffic*: We denote this traffic by T_0 . If synchronous traffic is sent using a uniframe scheme, only a specification of the intensity is needed. If synchronous traffic is packetized, we use the following model. We assume that each packet has a constant length and fits into one slot. In our case, packet lengths are 512 bits and are equal to the slot information field length. Use of a Poisson arrival process for synchronous traffic cannot be justified (see [23]). We assume a bulk deterministic arrival process where the interarrival times are constant and determined according to the total load. In reality, the arrivals of synchronous traffic packets have a random but a periodic behavior, i.e., given a number of connections, the number of packets arriving in a packetization interval is a constant (possibly with small deviations). The deterministic arrival process has been chosen to model the regularity of the arrival process of synchronous traffic. The bulk arrival process represents in a conservative way the fact that not all the arrivals are equidistant in time. The bulk size is chosen to be 4 which is a realistic value but could be subject to discussion.

2) *Asynchronous Traffic*: We assume a Poisson arrival process and a bimodal packet length distribution which has been argued to be a good one for LAN's [19]. So we distinguish two packet lengths. We denote packets belonging to the group with shorter average packet length

by T_1 and those belonging to the group with longer average packet length by T_2 . T_2 represents, e.g., file transfers and T_1 , e.g., management messages, electronic mail, etc. We assume the following packet lengths: 1024 bits for T_1 , and 8192 bits for T_2 . The relative load (in terms of offered bits per time unit) is taken as follows: 13 percent of T_1 , and 87 percent of T_2 .

Two load types and traffic patterns have been used: a symmetric and an asymmetric one. In the symmetric case, the load intensity at each station S_i , $i = 0, \dots, n$, is the same and

$$p_{ij} = \begin{cases} \frac{1}{n}, & j \neq i \\ 0, & j = i \end{cases} \quad i, j = 0, \dots, n. \quad (4)$$

In the asymmetric case, one station S_0 is taken to be different from the others. In this case, the load at S_0 is $1/3$ of the total load. At all other stations S_j , $j = 1, \dots, n$ it is equal. p_{ij} take the following values:

$$p_{ij} = \begin{cases} \frac{1}{2}, & j = 0 \\ \frac{1}{2} \cdot \frac{1}{n-1}, & j \neq i, j \neq 0, \\ 0, & j = i \\ i = 1, \dots, n, j = 0, \dots, n, \end{cases} \quad (5)$$

and

$$p_{0j} = \begin{cases} \frac{1}{n}, & j \neq 0 \\ 0, & j = 0 \end{cases} \quad j = 0, \dots, n. \quad (6)$$

This asymmetric load has been taken as a representative of such type of load, not as a model of a particular application. However, S_0 could represent, e.g., a high throughput gateway.

VI. CONFIGURATION AND SYSTEM PARAMETERS

Configuration and system parameters typical for HSLAN's have been taken as follows: *configuration*: cable length = 5 km, number of stations = 40 or 10, *system parameters*: transmission rate = 140 Mbits/s, slot information field length = 512 bits, overhead in slot = 48 bits, and latency register at each station = 24 bits. An equal distance between the neighboring stations has been assumed, so $\tau_{i,i+1} = \tau_{j,j+1}$, $i, j = 0, \dots, n$. We assume that the preamble and the MAC_PCI length in a token ring add up to 200 bits. We assume that the length of the overhead in a uniframe scheme equals 84 bits.

A common slot information field length of 512 bits has been chosen for all the slotted rings. In the CFR [17], the slot information field is 256 bits and in Orwell [4] 128 bits. We do a comparative performance analysis of these protocols and take the same information field length for

all the protocols. The overhead (MAC_PCI) length is taken to be 48 as in the original protocols.

An analysis concerning the system parameters has been done for the case with 10 stations. The reason is that the protocols are more sensitive to the values of system parameters if the number of stations is smaller (see [22], [24], and [25]). However, the case with 40 stations is a more realistic one for the office environment. A delay analysis and a sensitivity analysis for the configuration and workload parameters of these protocols have been done for the case with 40 stations.

VII. A COMPARATIVE ANALYSIS OF THE BASIC SLOTTED RING AM'S AND THE TOKEN RING

The basic slotted ring AM's can be modeled as multiple cyclic server (Fig. 4), as has been done in [22], [24], and [25]. A slot corresponds to a server. The service discipline is limited, so only one minipacket can be served per visit to the queue. The switchover time is constant and equals $\tau_{i,i+1}$. The service time in the CFR and CFRV is also constant and equals τ . The Orwell basic AM has been studied in [22] under a symmetric load only, i.e., for

$$\lambda_i = \lambda_j, \quad \gamma_i = \gamma_j, \quad i, j = 0, \dots, n. \quad (7)$$

In Orwell, the service time is a random variable X_i such that

$$P\{X_i = \tau_{i+1,j+1}\} = \frac{1}{n}, \quad j \neq i, \quad j, i = 0, \dots, n.$$

In Orwell, the order of visiting the queues by the servers is random and in the CFRV it is a regular one. In the CFR, the service is performed only if there is no other server serving the same queue.

Let us present now the exact stability conditions which have been proposed in [22], [24], and [25]. We will use them to compare the capacity of the protocols. They are independent of the interarrival time and packet length distributions except for their first moments. For the CFR, it has been shown [24] that the system may have a few stable patterns when all the queues are instable. The necessary and the sufficient condition for all queues to be stable are different. The necessary condition is the following one:

$$\lambda_i \gamma_i \sigma < 1 - \rho, \quad i = 0, \dots, n, \quad (8)$$

and the sufficient one is

$$\lambda_i \gamma_i \tau < 1 - \rho, \quad i = 0, \dots, n. \quad (9)$$

For a symmetric load relations, (8) and (9) become, respectively,

$$\rho < 1 - \frac{1}{n+2} \quad (10)$$

and

$$\rho < 1 - \frac{s}{s+n+1}. \quad (11)$$

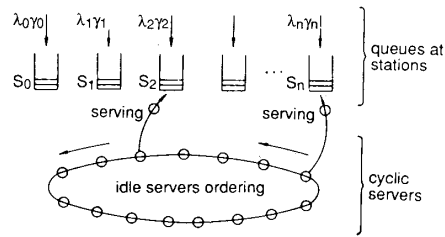


Fig. 4. A multiplequeue multiple cyclic server model of a slotted ring.

For the CFRV relations, (8) or (10) represent both a necessary and a sufficient stability condition. For the Orwell basic AM under a symmetric load, the following necessary and sufficient stability condition holds:

$$\rho < 2 - \frac{4}{n+3}. \quad (12)$$

Relation (12) shows that under a symmetric load the capacity of Orwell is close to twice the transmission rate. It is also about two times larger than the one of the CFRV [see (10)]. As expected, the CFRV is more efficient than the CFR. This is apparent from the fact that the sufficient stability conditions are different. Note that the stability conditions depend on ρ and n but not on s except for the sufficient condition of the CFR.

Let us now do a comparative sensitivity analysis of the delays of the basic slotted ring AM's by means of the analytical models developed in [22], [24], and [25]. The results are compared to the multiple token ring with an exhaustive service discipline, i.e., the target token rotation time is infinite. Exact results from [16] have been used for the token ring modeling. A comparative analysis of a uniframe integration scheme will also be presented.

The workload introduced in Section V yields an average packet length of 621 bits. If we choose the packet length of the synchronous traffic equal to the slot information field length, no bits are wasted. Synchronous traffic represents 80 percent of the total load. So, we determine the average packet length as the average packet length observed on the medium, i.e., it is calculated starting from the number of minipackets appearing in the ring. An exponential packet length distribution is assumed and a configuration with 40 stations is assumed.

The *packet delay versus load* is depicted in Fig. 5. The Orwell basic AM has the best performance. The CFRV basic AM performance is the next best and has very low delays at loads < 80 Mbits/s. The CFR basic AM has larger delays than the CFRV due to the fact that only one slot at a time can be used by a station which increases the waiting times for access to the medium. The token ring has the worst performance. We attribute this to its efficiency for short packet lengths, e.g., less than 1000 bits (see also Fig. 9).

The *packet delay versus the number of slots*, or the ring length, is shown in Fig. 6. The load is 80 Mbits/s. The delay increase for the CFRV and the Orwell basic AM is linear and equals the increase in the propagation delay.

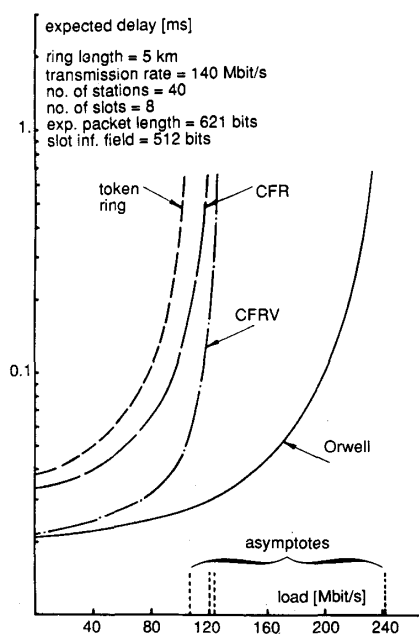


Fig. 5. Packet delay versus load in the basic AM's.

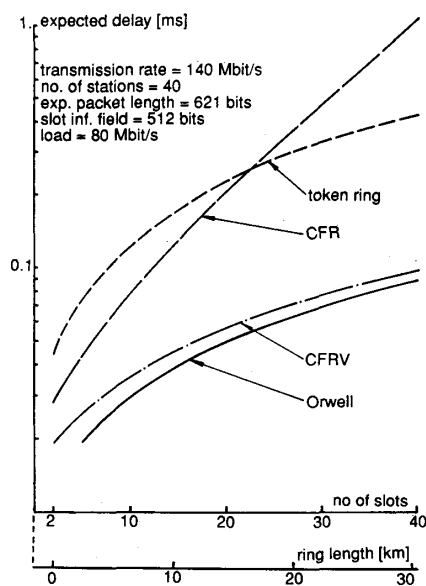


Fig. 6. Packet delay versus number of slots and ring length in the basic AM's.

So, these two AM's perform well with different numbers of slots. The token ring has a nonlinear and much larger increase in the delays than the CFRV and Orwell because, in addition, longer ring delays cause longer switchover times. In the slotted rings, these are compensated by introducing new slots which increase the capacity (i.e., the number of servers). The CFR has a degradation of performance if the number of slots increases (i.e., the ratio be-

tween the number of stations and the number of slots decreases) due to the fact that only one slot at a time can be used by a station.

The *packet delay versus the number of stations* is presented in Fig. 7. The load is 80 Mbits/s. The ring delay is kept constant. So, the addition of a new station does not introduce extra latency. The CFR is the most sensitive to the number of stations. This is because of the change of the ratio between the number of stations and the number of slots. The others are sensitive only for a small number of stations, e.g., less than 15. All delays converge to horizontal asymptotes. This is a property of cyclic servers systems (see [16], [22], [24], and [25]). In the slotted rings, the packet delays decrease with an increasing number of stations and in the token ring they increase. We attribute this to the different service disciplines, namely, the limited discipline in the slotted rings and the exhaustive discipline in the token ring. If the number of stations is smaller, the probability that the server is serving the queue where a packet arrives is larger and the expected waiting time is smaller in case of an exhaustive discipline. With a limited service discipline a more regular service is provided resulting in smaller delays for a larger number of stations.

The *packet delay versus transmission rate* is shown in Fig. 8. The relative load is held constant at 0.625, e.g., 80 Mbits/s load at 140 Mbits/s transmission rate. Above 200 Mbits/s, the CFR has increasing delays, mainly due to the increase of the number of slots. In the CFRV, Orwell and the token ring delays decline. The Orwell and the CFRV basic AM's have much smaller delays than the token ring in the case studied. The delay curves for the CFR and Orwell eventually have the same asymptote at the expected propagation time from source to the destination, i.e., $\tau/2$ in the limit. The token ring converges to an asymptote which equals the sum of the expected propagation time from source to destination and a nonzero waiting time for the token. In case of the slotted ring, increasing the transmission rate implies more slots in the ring, i.e., more servers. In a token ring, the service time (i.e., transmission time) of a packet tends to zero, however, due to a constant relative load the token is held for a nonvanishing time at each station.

The *packet delay versus average packet length* is shown in Fig. 9. The load is 80 Mbits/s. The increase in the delays of the CFRV and Orwell is caused by longer transmission times. The transmission time is dependent on the number of minipackets per packet (i.e., on the packet length and the slot information field length) and on the slot duration. Note that the queuing delays (packet delays excluding the transmission delay) are the same for short and long packets given the load intensity and the average packet length. The CFR has much larger delays than the CFRV if the packet lengths are larger because in addition to the effect of longer transmission times only one slot at a time can be used by a station. The delays in the token ring decrease with increasing packet length till about 1300 bits and then they increase. If the average

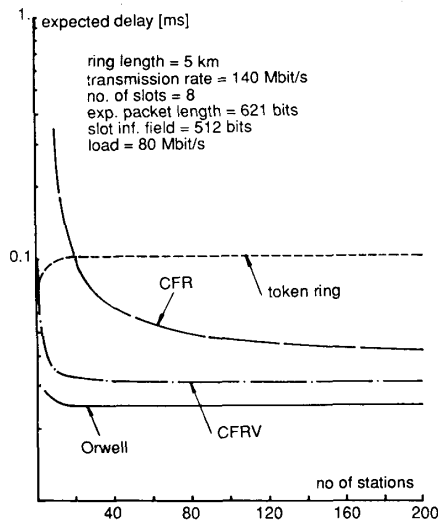


Fig. 7. Packet delay versus number of stations in the basic AM's.

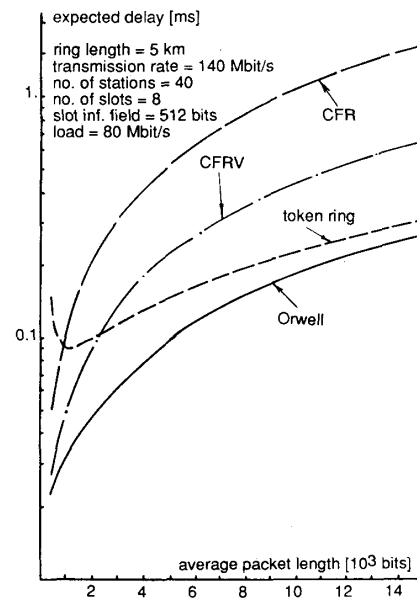


Fig. 9. Packet delay versus average packet length in the basic AM's.

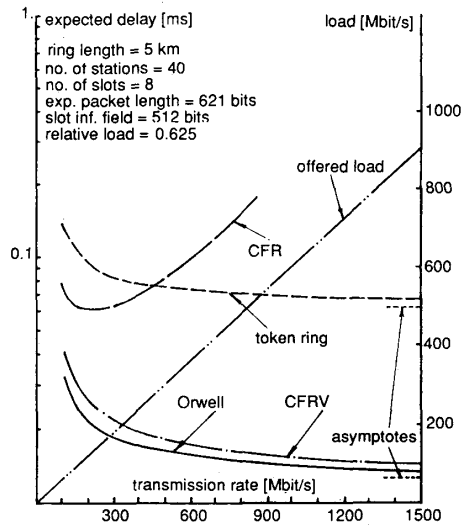


Fig. 8. Packet delay and offered load versus transmission rate in the basic AM's.

packet length is short then the token passing AM has a small throughput and a bad performance. This can be illustrated as follows. Assume that all the packets have a constant packet length and a layout which is the same as for the minipackets. The token passing AM with an exhaustive service discipline is worse than the CFRV provided that on the average the transmission at a station is shorter than the ring latency. In addition to this, if the packet length is short the ratio between the preamble and PCI length (assuming an asynchronous transmission scheme) and the data information field length in a MAC_PDU is large and the token passing AM inefficient. On the other hand, if the average packet length grows, the transmission time also grows and so do the delays. The

token ring has a worse performance than the CFRV if the average packet length is small relative to the slot information field length, e.g., in our case less than about 2300 bits. If the average packet length is large relative to the slot information field length, e.g., 15 000 bits in Fig. 9, the delays in the token ring are comparable to those in the Orwell basic AM. This demonstrates that in typical HSLAN's (without a uniframe) which are expected to carry a lot of synchronous traffic in short packets, slotted rings have an advantage in terms of throughput and expected delay over a token ring. This advantage is very pronounced for the CFRV and Orwell.

If a *uniframe scheme* is used, very small and constant delays are ensured for synchronous traffic. The *delays of the asynchronous traffic versus asynchronous traffic load* are depicted in Fig. 10. An average packet length of 4319 bits has been taken. This is according to the workload of Section V. An exponential packet length distribution is assumed. Again, we determine the average packet length as the average packet length observed on the medium, i.e., it is calculated starting from the number of minipackets appearing in the ring. The synchronous traffic load is 80 Mbits/s if full duplex connections are assumed. The delays in the CFR basic AM are much larger than in the other protocols. This is because the CFR performs badly for large packet lengths. The token ring performs better than the CFRV basic AM because of the relatively long average packet length. An additional reason is that some capacity is wasted in the slotted ring protocols since an integer number of slots must fit in a frame leaving some unused capacity. In token passing rings this does not occur. At low loads delays of the token ring are larger than in the CFRV due to waiting for access to the medium,

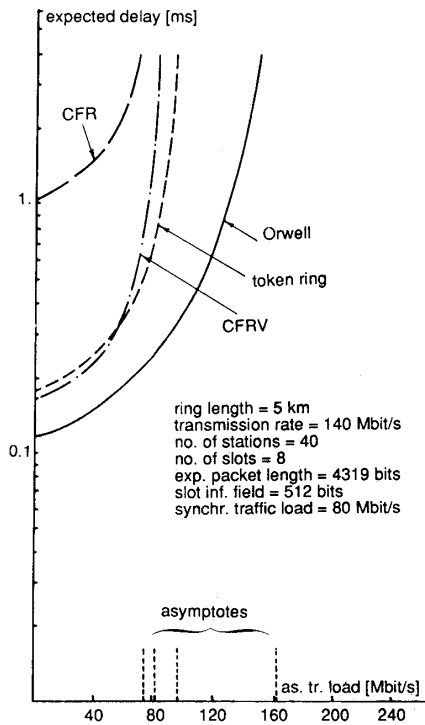


Fig. 10. Packet delay asynchronous traffic versus asynchronous traffic load in the basic AM's with a uniframe scheme.

i.e., waiting for a token. The maximum capacity of Orwell stays approximately the same, i.e., 242 Mbits/s. This is mainly because of the destination release of the slots in the Orwell basic AM which already almost provides a full duplex connection if the number of stations is large relative to the number of slots. Note that if the number of stations is small, e.g., 4 the usage of uniframe over Orwell may improve the performance significantly [26]. For all the other basic AM's when a uniframe is used the maximum carried load increases, e.g., 162 versus 124 Mbits/s for the CFRV and 175 versus 107 Mbits/s for the token ring (see also Fig. 5). Hence, usage of a uniframe does improve performance both for synchronous and asynchronous traffic for the token ring and the slotted ring protocols. It has the most effect on the token ring. The price for using a uniframe scheme is protocol complexity since it has to integrate both PCM-like circuit switched channels and a MAC protocol for data.

The slotted ring protocols are clearly advantageous if the average packet length is small and that happens when synchronous traffic is packetized. We present further on a performance analysis of these protocols for that case.

VIII. INFLUENCE OF THE SYSTEM PARAMETERS

In this section, we study by means of simulation the three slotted ring protocols under both symmetric and asymmetric loads in the cases when 1) the number of normal and channel slots changes in the CFR and the CFRV

and 2) the D_i allocation changes in Orwell. We use a configuration with 10 stations which implies seven slots in the ring. A queue prioritizing the synchronous traffic has been assumed.

A. Number of Normal and Channel Slots

Normal slots can be viewed as cyclic servers with a limited service discipline meaning that after carrying a minipacket from a particular station another station must be offered the service. Channel slots can be viewed as cyclic servers with an exhaustive service discipline. So, they serve a queue until it becomes empty. Only then the service is offered to another station. The fairness of the limited service discipline is provided at the expense of ceasing the service at a queue which is not empty. On the other hand, an exhaustive service discipline permits an unlimited possession of a server by a station and thus hogging of the ring. This property is undesirable. Since a queue is served exhaustively, this discipline, for a Poisson arrival process, provides smaller expected packet delays than the limited service discipline (see, e.g., [16] for a single cyclic server system). At low relative loads, smaller maximum packet delays are expected, but at larger relative loads, the unfairness of the service may result in too large maximum delays.

A symmetric load of 100 Mbits/s has been chosen. Results of the simulations for the CFR are presented in Table I and for the CFRV in Table II.

Let us analyze the CFR. With an increasing number of channel slots in the ring, the expected packet delays of the three traffic types (T_0 , T_1 , and T_2) decrease. The maximum packet delays also decrease for all traffic types except for T_1 if the number of channel slots equals 7 where it gets slightly larger again. Let us analyze the CFRV. With an increasing number of channel slots in the ring both the expected and the maximum packet delays for T_1 and T_2 decrease. The number of channel slots does not influence the performance of T_0 very much. The reason for the observed behavior is the same as was given in the introductory paragraph of this section.

All the expected and maximum packet delays are smaller in the CFRV than in the CFR. The only exceptions concern the maximum packet delays which are larger in the CFRV for T_1 and T_2 in the case of 0 channel slots, and for T_0 in the case of 7 channel slots. The ratio between the expected packet delay and its maximum can be used as an indication of the variance of the delays. For example, a larger expected value but a smaller maximum value of the packet delay in the CFR than in the CFRV, indicates that the variance of the delay in the CFR is much smaller. Both in the CFR and the CFRV under a symmetric load and if the number of channel slots equals 4, the performance is very good for all traffic types. Introducing more channel slots does not improve it. More channel slots increase the risks associated with unfairness.

The delays are much larger under an asymmetric load for the CFR than for the CFRV, so we take different load

TABLE I
PACKET DELAYS VERSUS NUMBER OF CHANNEL SLOTS IN THE CFR UNDER A SYMMETRIC LOAD OF 100 Mbits/s

traffic type	no of chan. slots	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
T0	0	0.145	0.003	1.052
	3	0.100	-	0.921
	4	0.097	0.001	0.842
	7	0.092	0.001	0.549
T1	0	0.289	-	4.131
	3	0.142	0.010	1.744
	4	0.132	0.009	1.488
	7	0.122	0.009	1.507
T2	0	0.825	-	3.481
	3	0.544	0.010	1.866
	4	0.526	0.008	1.843
	7	0.508	0.008	1.798

TABLE II
PACKET DELAYS VERSUS NUMBER OF CHANNEL SLOTS IN THE CFRV UNDER A SYMMETRIC LOAD OF 100 Mbits/s

traffic type	no of chan. slots	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
T0	0	0.087	-	0.938
	3	0.062	-	0.634
	4	0.067	-	0.715
	7	0.072	-	0.783
T1	0	0.273	-	4.840
	3	0.137	0.013	1.364
	4	0.119	0.008	1.257
	7	0.098	-	0.716
T2	0	0.633	-	5.291
	3	0.279	0.015	1.940
	4	0.241	0.012	1.595
	7	0.182	0.007	0.897

TABLE III
PACKET DELAYS VERSUS NUMBER OF CHANNEL SLOTS IN THE CFR UNDER AN ASYMMETRIC LOAD OF 40 Mbits/s

station	traffic type	no of chan. slots	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
S0	T0	0	0.129	-	0.768
		3	0.114	0.002	0.552
		4	0.112	-	0.548
		7	0.109	0.002	0.548
	T1	0	0.071	-	0.670
		3	0.065	-	0.515
		4	0.062	-	0.500
		7	0.058	-	0.493
	T2	0	0.569	-	1.107
		3	0.462	-	0.838
		4	0.457	-	0.850
		7	0.452	-	0.839
S5	T0	0	0.195	-	0.771
		3	0.144	-	0.545
		4	0.141	-	0.550
		7	0.139	-	0.536
	T1	0	0.764	-	6.171
		3	0.212	-	1.509
		4	0.211	-	1.513
		7	0.205	-	1.509
	T2	0	1.004	-	4.098
		3	0.611	-	1.919
		4	0.602	-	1.896
		7	0.595	-	1.770

TABLE IV
PACKET DELAYS VERSUS NUMBER OF CHANNEL SLOTS IN THE CFRV UNDER AN ASYMMETRIC LOAD OF 100 Mbits/s

station	traffic type	no of chan. slots	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
S0	T0	0	0.055	-	0.553
		3	0.058	0.003	1.087
		4	0.070	0.005	0.585
		7	0.055	-	0.621
	T1	0	0.067	-	0.805
		3	0.090	-	0.562
		4	0.092	-	0.923
		7	0.087	-	0.454
	T2	0	0.360	-	1.209
		3	0.073	-	1.136
		4	0.232	-	0.743
		7	0.178	-	0.750
S5	T0	0	0.160	0.015	0.862
		3	0.071	0.003	0.794
		4	0.073	-	0.804
		7	0.061	-	0.421
	T1	0	103.330	-	180.282
		3	0.125	-	0.930
		4	0.097	-	0.754
		7	0.088	-	0.794
	T2	0	59.751	-	476.384
		3	0.272	-	1.224
		4	0.223	-	1.088
		7	0.169	-	0.517

values. We take a 40 Mbit/s load for the CFR and 100 Mbits/s for the CFRV. Results of the simulations for the CFR are given in Table III and for the CFRV in Table IV. Note that the delays at the lightly loaded stations S_i , $i = 1, \dots, n$, do depend on the position of the station with respect to the heavily loaded station S_0 . The results are given for station S_0 which has a heavy load and for station S_5 which has a light load.

Let us examine the CFR. The delays are large already at a load of 40 Mbits/s. The different traffic types experience different delays, e.g., both the expected and the maximum delays are even 10 times larger for $T1$ at S_5 than at S_0 and the expected delays for $T2$ are about 1.1–1.5 times larger. Having 3–7 channel slots improves the performance.

Let us now analyze the CFRV. The delays are smaller at S_0 than at S_5 except for $T2$ in 4 and 7 channel slot configuration. For the same traffic types the performance is even better under the same total load (see also Table II) in the asymmetric case than in the symmetric case, except for the case with zero channel slots. The performance improves if the number of channel slots increases. The performance is much better already with 3 channel slots compared to zero.

Under an asymmetric load, the performance in both the CFR and in the CFRV is the best if the number of channel slots equals 7, but having, e.g., 4 channel slots already improves the performance significantly.

Both in the CFR and in the CFRV the delays are larger at a lightly loaded station than at a heavily loaded station, except for the case with 7 channel slots in the CFRV. In a single cyclic server case with an exhaustive service discipline, the delays are smaller at the heavily loaded sta-

tion [16]. This property is attributed to the large probability that customer arriving to the heavily loaded queue finds the server serving that very queue. It has been shown in [1] that this property also holds for a single cyclic server system with a limited service discipline and zero switchover times. For nonzero switchover times, no conclusions concerning this property could be deduced. The results of the simulations in Tables III and IV show that this property also holds with a large difference in the delays for the case with only normal slots. This property does not hold in the CFRV with only channel slots but the differences in the delays are not large.

B. D_i Allocation in Orwell

Let us now study the influence of D_i allocation. The analysis of the Orwell basic AM corresponds to a D_i value of infinity. The intent is not to give a complete analysis of the sensitivity to the D_i value, but to merely demonstrate that the D_i values have indeed a large impact on the performance of all traffic types. A number of experiments have been done. First a symmetric load of 140 Mbits/s has been used.

In the first set, the ratio $D_{i,1}/D_{i,2}$ is kept constant at 4, corresponding to the share of synchronous and asynchronous traffic, and $D_{i,1}$ is varied. The packet delays are shown in Table V and the reset intervals in Table VI. The performance of $T0$ improves if $D_{i,1} = 8$. Larger values do not improve the performance of $T0$ but $T1$ and $T2$ get better performance. The reset intervals increase with increasing D_i .

Next $D_{i,1}$ is kept constant at 8 and $D_{i,2}$ is changed. Results are shown in Tables VII and VIII. Table VII shows that the expected and the maximum delays of $T1$ and $T2$ decrease with increasing $D_{i,2}$. The delays of the synchronous traffic increase in that case. $D_{i,1} = D_{i,2} = 8$ provides a good compromise for all traffic types. The expected duration of the reset intervals does not change with $D_{i,2}$ contrary to Table VI, but the maximum does.

Results for $D_{i,1}/D_{i,2} = 1$ and different $D_{i,1}$ values are shown in Tables IX and X. $D_{i,1} > 5$ yields small delays. For $D_{i,1} = 4$ the performance for $T0$ is bad. The reset intervals also increase a lot if $D_{i,1}$ is changed from 4 to 5.

The performance for an asymmetric load of 120 Mbits/s is shown in Table XI and Table XII. $D_{i,1}/D_{i,2} = 1, i = 0, \dots, n$ and $D_{0,j}/D_{i,j} = 5, j = 1, 2, i = 1, \dots, n$. The expected and the maximum packet delays are smaller if $D_{i,1}$ and $D_{i,2}$ are larger, e.g., 20 at S_0 and 4 at other stations. If D_i is infinite, the delays at the heavily loaded station S_0 are smaller than at the lightly loaded station S_5 , which is the property also noticed for the CFR and the CFRV.

We can conclude that if the reset interval has to be small with respect to the total load, the Orwell protocol is very sensitive to the D_i values. If inadequately chosen they can degrade the performance strongly. In these cases, sudden changes of the loads can strongly influence the performance. On the other hand, if the reset intervals do not have to be small relative to the load, the performance is

TABLE V
PACKET DELAYS VERSUS $D_{i,1}$ IN ORWELL UNDER A SYMMETRIC LOAD OF 140 Mbits/s WHERE $D_{i,2} = 4 * D_{i,1}$

traffic type	$D_{i,1}$	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
T0	4	0.128	0.005	0.628
	8	0.086	0.001	0.404
	64	0.085	0.001	0.810
	∞	0.081	-	0.817
T1	4	26.623	-	86.005
	8	27.188	-	86.743
	64	11.381	3.272	54.381
	∞	0.078	0.004	0.989
T2	4	24.045	18.403	86.149
	8	24.044	19.191	87.121
	64	13.707	3.874	54.405
	∞	0.256	0.007	1.352

TABLE VI
RESET INTERVAL DURATION VERSUS $D_{i,1}$ IN ORWELL UNDER A SYMMETRIC LOAD OF 140 Mbits/s WHERE $D_{i,2} = 4 * D_{i,1}$

$D_{i,1}$	exp. reset int. [ms]	95% conf. int. [ms]	max reset int. [ms]
4	0.173	0.007	1.640
8	0.349	-	2.216
64	3.096	-	8.835
∞	16.399	-	67.185

TABLE VII
PACKET DELAYS VERSUS $D_{i,2}$ IN ORWELL UNDER A SYMMETRIC LOAD OF 140 Mbits/s WHERE $D_{i,1} = 8$

traffic type	$D_{i,2}$	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
T0	2	0.086	0.001	0.404
	4	0.092	0.002	0.458
	8	0.116	0.010	0.860
	10	0.123	0.010	0.892
	12	0.135	0.012	0.921
	100	0.176	0.021	1.344
T1	2	27.188	-	86.743
	4	0.702	0.138	7.526
	8	0.167	0.146	2.607
	10	0.112	0.009	1.798
	12	0.101	0.006	1.579
	100	0.076	0.003	0.383
T2	2	24.405	19.191	87.120
	4	1.806	-	9.097
	8	0.563	0.018	2.730
	10	0.494	0.013	2.104
	12	0.438	0.010	1.693
	100	0.177	-	0.544

TABLE VIII
RESET INTERVAL DURATION VERSUS $D_{i,2}$ IN ORWELL UNDER A SYMMETRIC LOAD OF 140 Mbits/s WHERE $D_{i,1} = 8$

$D_{i,2}$	exp. reset int. [ms]	95% conf. int. [ms]	max reset int. [ms]
2	0.349	-	0.443
4	0.398	-	2.153
8	0.387	-	1.850
10	0.394	-	1.910
12	0.393	-	1.528
100	0.393	-	2.153

TABLE IX
PACKET DELAYS VERSUS $D_{i,1}$ IN ORWELL UNDER A SYMMETRIC LOAD OF 140 Mbits/s WHERE $D_{i,2} = D_{i,1}$

traffic type	$D_{i,1}$	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
T0	4	6.125	-	14.171
	5	0.201	0.025	1.184
	8	0.116	0.010	0.860
	16	0.100	0.004	0.834
T1	4	0.188	0.017	3.174
	5	0.167	0.014	2.773
	8	0.167	0.146	2.607
	16	0.141	0.124	2.458
T2	4	0.779	0.020	3.139
	5	0.812	0.017	2.884
	8	0.563	0.018	2.730
	16	0.334	0.018	2.582

TABLE X
RESET INTERVAL DURATION VERSUS $D_{i,1}$ IN ORWELL UNDER A SYMMETRIC LOAD OF 140 Mbits/s WHERE $D_{i,2} = D_{i,1}$

$D_{i,1}$	exp. reset int. [ms]	95% conf. int. [ms]	max reset int. [ms]
4	0.173	0.007	1.640
5	0.257	-	1.401
8	0.349	-	2.216
16	0.739	-	2.970

TABLE XI
PACKET DELAYS VERSUS $D_{0,1}$ IN ORWELL UNDER AN ASYMMETRIC LOAD OF 120 Mbits/s WHERE $D_{i,2} = D_{i,1}$

$D_{0,1}$	exp. reset int. [ms]	95% conf. int. [ms]	max reset int. [ms]
10	0.156	-	0.244
15	0.193	-	2.478
20	0.235	0.007	0.442
30	0.245	0.009	0.556
∞	0.300	0.021	6.529

very good both under symmetric and asymmetric loads, and the Orwell basic AM also has a very good performance and a low sensitivity to the system and configuration parameters. This indicates that another way of integration of synchronous and asynchronous traffic may be needed to provide a better performance in these cases. Some attempts in that direction are proposed in [9]. The problems reported here concerning the D_i allocation are comparable to the problems of selecting values of certain system parameters which regulate the fairness of the access protocol in other LAN's. For example, it has been reported in [15] that the choice of TTRT in FDDI also has a large influence on the performance of FDDI.

IX. DELAY ANALYSIS OF THE INTEGRATED SERVICES SLOTTED RINGS

Let us now examine the integration of synchronous and asynchronous traffic. We assume a configuration with 40 stations implying 8 slots and we limit ourselves to symmetric cases.

TABLE XII
RESET INTERVAL DURATION VERSUS $D_{0,1}$ IN ORWELL UNDER AN ASYMMETRIC LOAD OF 120 Mbits/s WHERE $D_{i,2} = D_{i,1}$, AND $D_{0,j} = 5 * D_{k,j}, k \neq 0$

station	traffic type	$D_{i,1}$	exp. packet delay [ms]	95% conf. int. [ms]	max packet delay [ms]
S0	T0	10	45.707	-	96.884
		15	0.609	-	3.335
		20	0.170	-	1.961
		30	0.137	-	1.623
		∞	0.058	0.001	0.508
	T1	10	0.274	-	4.173
		15	0.204	-	3.839
		20	0.187	-	3.697
		30	0.196	-	3.377
		∞	0.053	-	0.396
	T2	10	1.413	0.044	4.632
		15	1.242	-	4.037
20		1.020	-	3.794	
30		1.059	0.036	3.512	
∞		0.155	0.005	0.635	
S5	T0	2	26.736	-	96.842
		3	0.376	-	3.278
		4	0.130	-	1.651
		5	0.103	-	1.242
		∞	0.067	0.002	0.437
	T1	2	0.177	-	2.412
		3	0.267	-	2.178
		4	0.140	-	1.948
		5	0.134	-	1.549
		∞	0.115	-	1.000
	T2	2	0.970	0.136	2.692
		3	0.866	0.108	2.479
4		0.708	0.087	2.287	
5		0.728	-	2.230	
∞		0.240	-	1.050	

The share of each traffic type is varied. In accordance with the results of Section VIII, 6 channel and 2 normal slots are used for the CFR and the CFRV. In Orwell, the total D_i allocation per station is 10. This ensures that the total D_i allocation is not too small, so reset intervals do not appear too often. $D_{i,1}/D_{i,2}$ is chosen equal to the ratio of synchronous and asynchronous traffic. Although this ratio is not the optimal one, it has been used since it provides a fair share of capacity to all traffic types and, on the other hand, obtaining the optimal ratio is difficult.

The delays of different traffic types versus asynchronous traffic load are shown in Fig. 11 for the CFR, in Fig. 12 for the CFRV, and in Fig. 13 for Orwell. The synchronous traffic load is used as a parameter. The results have been obtained by simulations.

In the CFR and the CFRV the delays increase rather slowly with the load until the utilization is about 0.8. From there on they rise sharply. Random fluctuations in load intensity may cause a sudden degradation of performance by significantly increasing the delays. If the load intensity of both synchronous and asynchronous traffic in the CFRV is controlled, bounded delays and a good QoS for both traffic classes can be provided. A guaranteed synchronous service can also be provided by merely controlling the amount of synchronous traffic itself (given queueing priority). However, this would not be sufficient when a large share, e.g., 80 percent of the load consists of synchronous traffic. The performance of the CFR is

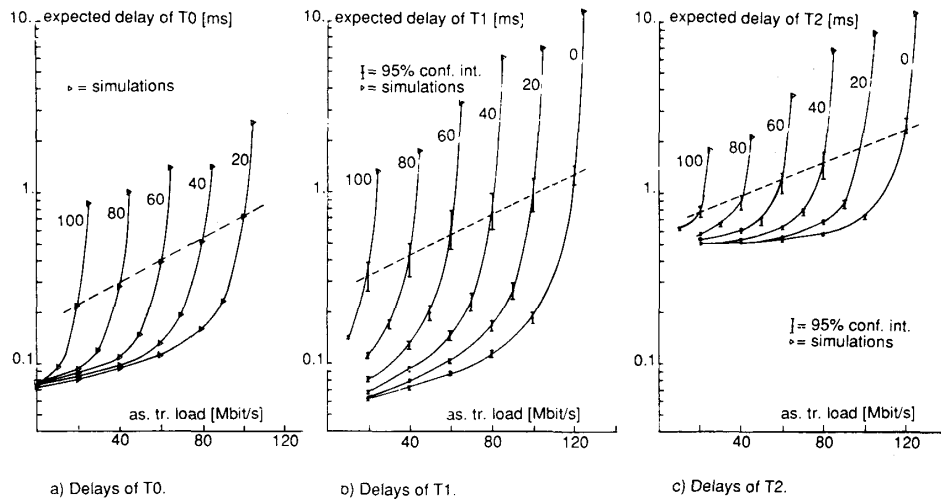


Fig. 11. Packet delays versus asynchronous traffic load given a synchronous traffic load in the CFR.

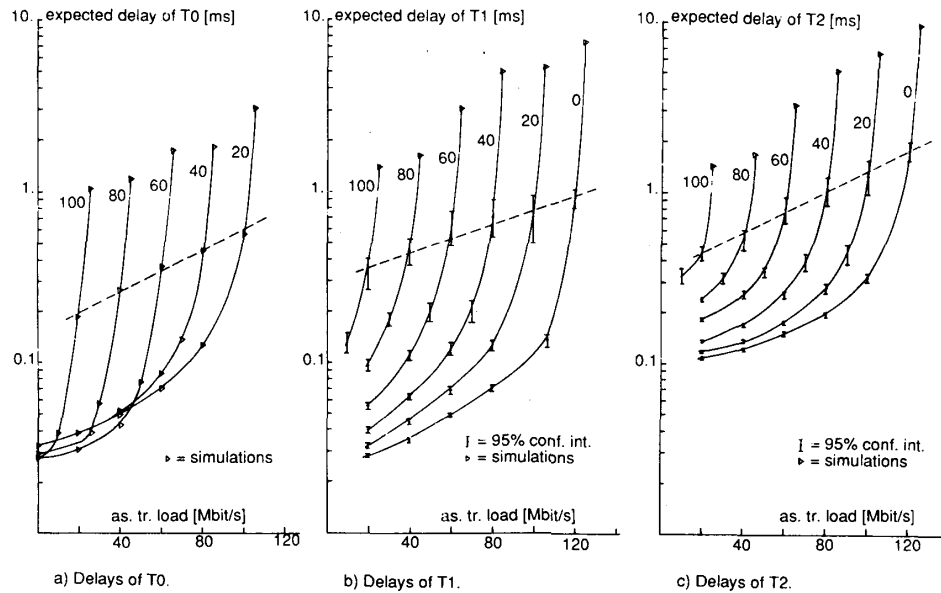


Fig. 12. Packet delays versus asynchronous traffic load given a synchronous traffic load in the CFRV.

worse than the CFRV. In the CFR and the CFRV, the delays increase if the share of asynchronous traffic increases given a constant total load. This is because of the following two reasons: the larger variance of the packet length distribution of asynchronous packets, and the larger variance of the asynchronous traffic interarrival time due to a Poisson arrival process. The dotted lines in Figs. 11 and 12 represent a constant total load curve of 120 Mbits/s. The delays of T_1 and T_2 vary by a factor 4 in the CFR and by a factor 3 in the CFRV when the share of asynchronous traffic changes from 20 to 120 Mbits/s.

In Fig. 12(a), a counterintuitive behavior of the CFRV appears: if there is no asynchronous traffic, the expected delay decreases when the load increases. This is due to the regularity of the deterministic arrival process of synchronous traffic. Note, however, that the difference in the delays due to this is small, less than $5 \mu\text{s}$.

Orwell exhibits a sharp increase of the asynchronous traffic delays already starting at a utilization of about 0.5 while the increase of the synchronous traffic delays is much slower. The difference is due to the D_i allocation. The delays in Orwell do not increase smoothly because of

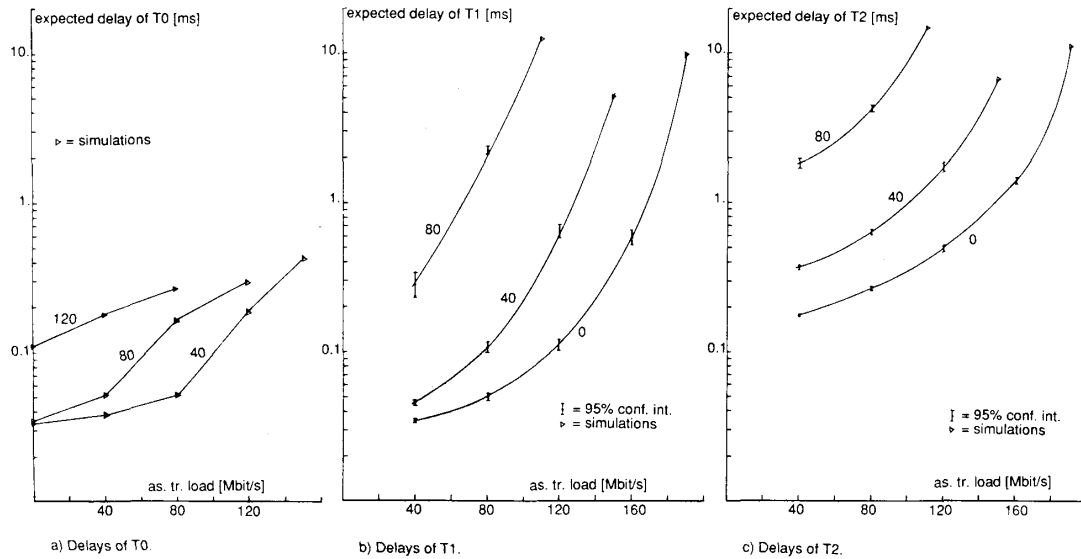


Fig. 13. Packet delays versus asynchronous traffic load given a synchronous traffic load in Orwell.

the high sensitivity to the choice of D_i : the $D_{i,1}/D_{i,2}$ ratio does not follow exactly the ratio of synchronous and asynchronous traffic. Note that the delays for synchronous traffic are small. This happens at the expense of the asynchronous traffic delays. This is in agreement with the results in Table V.

X. CONCLUSION

Slotted ring protocols are suitable for optical fiber HSLAN's because of their efficiency at high transmission rates and the synchronous transmission scheme. However, packet segmentation and reassembly may impose large processing demands if the packet lengths are large relative to the slot information field size. Let us now summarize the performance analysis results we obtained for these protocols. The multiple token passing ring has been used in this paper for comparison to the slotted ring basic AM's.

The Orwell and the CFRV basic AM's perform very well with respect to the number of slots and the number of stations. Their performance also improves if the transmission rate increases. The Orwell basic AM has the largest capacity due to the destination release of slots. Its capacity is almost twice the transmission rate under a symmetric load. The CFRV has about half of that capacity. The CFR basic AM has a much larger sensitivity to the system parameters and a much worse performance than the other two slotted ring protocols. Token passing is inferior to the CFRV and the Orwell basic AM's unless the average packet length is much larger than the slot length (e.g., >2300 bits). The performance of the token passing is improved most of all by the uniframe scheme. This scheme also improves the performance of the CFRV. It

does not improve the performance of the Orwell basic AM in a configuration with a large number of stations relative to the number of slots. The asynchronous traffic delays in the CFR with a uniframe are large, e.g., >1 ms.

The CFR has larger delays than the other slotted rings. So, the CFR has worse performance than the CFRV and Orwell with either long cable lengths, a small number of stations, or asymmetric loads. By increasing the transmission rate to more than, e.g., 200 Mbit/s the performance does not improve. This is the price for MAC layer acknowledgment and the possibility of immediate retransmission of erroneous minipackets.

The CFRV can be used successfully over a wide range of parameters. Combining the normal and the channel slots improves its performance especially under asymmetric loads. This corresponds to combining two service disciplines: limited and exhaustive. It performs very well under asymmetric loads. If the load intensity of both synchronous and asynchronous traffic on the network is controlled, bounded delays and a good QoS for both traffic classes can be provided. Furthermore, the CFRV allows for an acknowledgment at the MAC layer.

The good properties of the Orwell basic AM indicate that it can be used successfully over a wide range of parameters. Some improvement of the performance of Orwell is achieved if a larger $D_{i,2}$ allocation is taken for asynchronous traffic than what corresponds to the share of asynchronous traffic. The performance is good under asymmetric loads.

Because of the very good performance of Orwell and the CFRV protocol with long cable lengths and a large number of stations, they are suitable for metropolitan area networks, as well.

ACKNOWLEDGMENT

The authors would like to thank R. van den Brink for doing the simulation runs in Section VIII, which have been documented in [2].

REFERENCES

- [1] O. J. Boxma and B. Meister, "Waiting time approximations for cyclic-service systems with switch-over times," *Performance '86/ACM SIGMETRICS 1986*, 1986.
- [2] R. v. d. Brink, "Influence of system parameters on performance of the slotted rings: Simulations," University of Twente, Rep. 080-87-46, Oct. 1987.
- [3] L. M., Casey, R. C. Dittburner, and N. D. Gamage, "FXNET: A backbone ring for voice and data," *IEEE Commun. Mag.*, vol. 24, pp. 23-28, Dec. 1986.
- [4] R. M. Falconer and L. Adams, "Orwell: A protocol for an integrated services local network," *Brit. Telecom. Tech. J.*, vol. 3, no. 4, pp. 27-35, Oct. 1985.
- [5] R. M. Falconer, J. L. Adams, and G. M. Walley, "A simulation study of the Cambridge ring with voice traffic," *Brit. Telecom. Tech. J.*, vol. 3, no. 2, pp. 85-91, Apr. 1985.
- [6] W. K. Giloi, P. Behr, and G. Zuber, "Upperbus: A high speed backbone for metropolitan area networks," in *Proc. EFOC/LAN'86*, Amsterdam, The Netherlands, June 1986, pp. 286-289.
- [7] K. Hiyama, H. Narisawa, and H. Satou, "An integrated services optical fiber local area network 'Σ Network,'" *Hitachi Rev.*, vol. 32, no. 4, pp. 165-170, 1983.
- [8] A. Hopper, "The Cambridge ring—A local network," in *Advanced Techniques for Microprocessor Systems*, Hanna F. K., Ed. Stevenage, UK: Peter Pergrinus, Ltd., 1980, pp. 67-71.
- [9] J. O. Limb and S. Kandiyoor, "Improved scheduling of traffic for a high-speed slotted ring," *Electron. Lett.*, vol. 23, no. 15, pp. 806-765, July 1987.
- [10] K. v. d. Lande, "Design and implementation of a Cambridge fast ring variant simulation model, University of Twente, The Netherlands, Rep. 080-87-06, Feb. 1987.
- [11] L. Lamberelli, A. Luvison, D. Roffinella, and M. Sposini, "Service integration in wideband local area networks: Problems and system solutions," presented at Int. Tirrena Workshop Digital Commun., Tirrena (Pisa), North Holland, 1985.
- [12] I. Mitrani, J. L. Adams, and R. M. Falconer, "A modelling study of the Orwell ring protocol," in *Teletrafic Analysis and Computer Performance Evaluation*, O. Boxma, J. Cohen, and H. Tijms, Eds. New York: Elsevier, 1986, pp. 429-438.
- [13] N. N. Pedersen and R. Sharp, "Analysis of channel access schemes for high-speed LAN's," in *Proc. ACM SIGCOMM'84: Commun. Architect. Protocols*, Montreal, P. Q., Canada, June 1984, pp. 164-171.
- [14] F. E. Ross, "FDDI—A tutorial," *IEEE Commun. Mag.*, vol. 24, pp. 10-23, May 1986.
- [15] A. Schill and M. Zieher, "Performance analysis of the 100 Mbit/s optical token ring," in *IFIP WG 6.4 Workshop High Speed Local Area Networks*, A. Danthine and O. Spaniol, Eds. Aachen, W. Germany: North Holland, 1987, pp. 53-74.
- [16] H. Takagi, *Analysis of Polling Systems*. Cambridge, MA: The MIT Press, 1986.
- [17] S. Temple, "The design of the Cambridge fast ring," in *Ring Technology Local Area Networks*, I. N. Dallas and E. B. Spratt, Eds. New York: Elsevier, 1984, pp. 79-88.
- [18] D. Valk, "Performance analysis of integrated services slotted ring protocols in gathering HSLAN's," University of Twente, The Netherlands, Rep. 080-87-08, Mar. 1987.
- [19] E. Vyncke and A. Danthine, "A realistic simulation of a wideband backbone network," in *Proc. ICC'86*, North Holland, pp. 176-181, 1986.
- [20] M. Zafirovic-Vukotic and I. G. Niemegeers, "An evaluation of high speed local area network access mechanisms," in *Communication in Distributed Systems*, N. Gerner and O. Spaniol, Eds. Aachen, W. Germany: Informatik-Fachberichte no. 130, Springer-Verlag, Feb. 1987, pp. 426-440.
- [21] —, "Analytical models of the slotted ring protocols in HSLAN's," in *IFIP WG 6.4 Workshop High Speed Local Area Networks*, A. Danthine and O. Spaniol, Eds. Aachen, W. Germany: North Holland, 1987, pp. 115-134.
- [22] —, "Performance modelling of the Orwell basic access mechanism," in *Proc. ACM SIGCOMM'87: Frontiers Comput. Commun. Technol.*, Stowe, VT; also in *Comp. Commun. Rev.*, vol. 17, no. 5, Special Issue, Aug. 1987, pp. 35-48.
- [23] —, "Applications, workload models, benchmarks and configurations for HSLAN's," University of Twente, Rep. 080-86-33, Oct. 1987.
- [24] —, "Performance modelling of the Cambridge fast ring protocol," in *Proc. 1988 Int. Zurich Sem. Digital Commun.*, Zurich, Switzerland, Mar. 1988, pp. 163-170.
- [25] —, "Performance modelling of a HSLAN slotted ring protocol," presented at Proc. ACM SIGMETRICS'88, Santa Fe, NM, May 1988.
- [26] —, "HSLANs for communication within a gateway: A performance evaluation," in *Proc. IFIP WG 6.4 Workshop HSLAN'88*, A. Danthine, O. Spaniol, Eds. Liege, Belgium: North Holland, Apr. 1988, to be published.



Mirjana Zafirovic-Vukotic was born in Belgrade, Yugoslavia, in 1958. She received the B.S. degree in mathematics in 1981 and the M.S. degree in operations research in 1984 from the University of Belgrade. She is also a Ph.D. candidate. She worked as a Research Assistant in operations research at the University of Belgrade from 1982 to 1984. Since 1984 she has been working as a Research Assistant at the University of Twente, The Netherlands, on performance analysis of integrated services high speed LAN's. Her interests are in computer communication systems and performance analysis.



Ignas G. Niemegeers (S'70-M'78) was born in Gent, Belgium, in 1947. He received the M.S. degree in electrical engineering from the Rijksuniversiteit Gent in 1970. In 1972 he received the M.S. degree in computer engineering and in 1978, the Ph.D. degree from Purdue University, West Lafayette, IN.

From 1978 to 1981 he was a System Design Engineer at Bell Telephone Mfg. Co., Antwerp, Belgium, where he was involved in the design of the DPS 1500 system, a high capacity X.25 packet switching network. From 1981 to 1986 he was with the Department of Electrical Engineering, University of Twente, Enschede, The Netherlands. Since 1986 he has been a Professor of computer science at the same university. His present research interests are architectures for local integrated communication infrastructures, parallel processing in communication switches, performance analysis, and performability of fault tolerant systems.



Durk S. Valk was born in Leeuwarden, The Netherlands, in 1964. He received the Ir. degree in electrical engineering from the University of Twente, The Netherlands in 1987.

He joined Logica B.V., The Netherlands, in 1987. He is currently also working on performance analysis tools. His interests are in computer and communication performance analysis and communication systems.