A CSMA/CD Compatible MAC for Real-Time Transmissions Based on Varying Collision Intervals[F]

Oran Sharon[1], Michael Spratt
Internet Mobile Solutions Laboratory
HP Laboratories Bristol
HPL-2000-92
27th July, 2000*

E-mail: oran@mathcs11.haifa.ac.il
mps@hplb.hpl.hp.com

local area network,
LAN,
real-time,
CSMA/CD,
ethernet,
delay-sensitive

This paper suggests a CSMA/CD compatible MAC protocol for real-time transmissions in a shared Home or Small Office Local Area Network. The new MAC is based on the distributed assignment of special TAG numbers to the stations transmitting real-time traffic. These TAG numbers determine a Round Robin order of transmissions among the real-time stations. They also help in resolving collisions among real time stations by setting the length of the Jam signal, transmitted in case of a collision, to be a function of the TAG number. In a collision the station with the highest TAG number, and so with the longest Jam, is persisting with its Jam transmission for the longest until all the other stations defer. Thus, the collision terminates and the longest persisting station can transmit its packet successfully. The new protocol enables stations implementing the IEEE 802.3 MAC standard to transmit on the same network with stations implementing the new protocol.

After introducing the protocol, we compute an upper bound on the access delay that the protocol guarantees and prove the correctness of the distributed TAG assignment procedure. Finally, we simulate the protocol in a network consisting stations implementing the new protocol together with stations implementing the standard IEEE 802.3 MAC. We show that the access delays of the stations transmitting real-time traffic are indeed bounded as predicted.

[F] A preliminary version of this paper appeared in INFOCOM'98

[1] Department of Mathematics, University of Haifa, Haifa, Israel.

* Internal Accession Date Only Approved for External Publication
To be published in *Computer Networks, available from* http://www.elsevier.com

© Copyright Hewlett-Packard Company 2000

1 Introduction

In the last 20 years, Local Area Networks (LANs) have revolutionized the way in which computers have been used in the work place, and have allowed Personal Computers and the Client / Server paradigm to largely replace the mainframe / terminal technology prior to LANs. Following this example, it is anticipated that a Home wired or wireless LAN potentially can help to revolutionize home equipment by bringing together computing and electronic entertainment. Such a network might link the digital TVs, set-top box, games machines, Digital VCR, CD player producing digital audio, PCs, printer and ISDN routers. Such a network could also find use in the small office environment.

In this paper we propose a shared, half duplex bus LAN for the Home environment. This alternative is attractive due to its low cost, simplicity of installation and maintenance, and its wide use by Ethernet [1]. We also suggest that data devices shall continue to use the well known IEEE 802.3 CSMA/CD MAC [1]. Notice however, that due to collisions it is not clear if the CSMA/CD MAC can always guarantee the QoS requirements of real-time traffic such as bandwidth, bounded delay and jitter. For this reason we suggest a new MAC for the stations transmitting real-time traffic, which is compatible with the standard CSMA/CD MAC.

The issue of incorporating real-time transmissions over a CSMA/CD network was extensively investigated in the past and several approaches were suggested to solve this task. One approach is to use Ethernet switches [2] where stations are connected by point-to-point links to switches and transmit in a full duplex mode with-no collisions. We do not adopt such a solution for the Home environment because its installation and maintenance are more complicated and expensive than a passive shared bus. Also, in the IEEE 802.9 IsoEthernet standard [3] Ethernet is combined with a circuit switched ISDN channel to enable packet and real time transmissions respectively. Again, we see such a solution to be too complicated and expensive for use in the Home environment.

It is also important to mention that many works, such as [4], [5], [6], [7], [8], [9], [10] and [11] simulated Ethernet networks where all the stations, data and real-time, transmit by the standard CSMA/CD MAC. Various traffic models for data and real-time traffic were considered. Many of these works show that CSMA/CD is suitable to support both the data transmissions and the real-time traffic with its required QoS.

However, we are looking for a MAC that can guarantee the QoS requirements of real-time traffic in a deterministic way. First, it is not adequate to propose a network to a Home customer and saying that "simulation results show that the network can satisfy the requirements of real-time applications". Also, the data and real-time characteristics can change in the future [11] and be different than those that were simulated, e.g. future WWW data applications traffic and various alarm signals that may be required in the Home environment such as fire alarms. Second, when alarm signals are considered, it is clear that guaranteed access delay is required.

There were also suggestions to change the CSMA/CD MAC as it appears in the IEEE 802.3 standard [1], by replacing the Binary Exponential Backoff (BEB) scheme by that of the Binary Logarithmic Arbitration Method (BLAM) [12], [13], [14]. These schemes are used by stations to compute their backoff intervals after collisions. It is shown in [14] that BLAM achieves better fairness and supports a higher number of real-time stations in a network than BEB. However, BLAM is not considered by the IEEE 802.3 standard committee anymore [15] due to the wide trend to use switched Ethernet systems, with no collisions. Notice however, that we assume that a shared bus is an attractive solution for the Home LAN.

Finally, there were many proposals in the past to change the MAC of the stations that transmit real-time traffic over Ethernet [16], [17], [18], [19]. These proposals enable stations implementing the new MAC a higher priority in acquiring the bus for transmission, and also enable these stations to transmit in the same network with stations implementing the IEEE 802.3 MAC. This is also the approach we take in this paper but our new MAC is simpler for implementation and it imposes a minimal restriction on the real-time traffic.

In [17] an extension to the work of [16] is presented, where different real-time traffic rates can be supported. The MAC is based on trains of transmissions, with a fixed order among the transmitting stations. However, this MAC requires that the different rates of the real-time streams will all be a multiple of the same basic quantity and it requires that stations will keep information on the trains in order to cope with station failures or transmission terminations. The MAC proposed in [18] uses similar ideas but also requires some mechanisms that complicate its implementation such as the maintenance of active and idle station queues at the stations, the election of a special station that serves as a channel controller and the processing of several different control packets.

The MAC proposed in [19] is simpler and creates a TDM like system for real-time transmissions. In order to keep fixed size transmission slots, real-time packets contain an overflow field that is always transmitted, even if it is not required. The length of the overflow field is proportional to the transmission time of the longest possible data packet which is 1520 bytes in the IEEE 802.3 standard, and the rate of the real-time application. For example, if we consider a CBR video application of rate 1.5 Mb/s and packets are generated every $40ms$, then for every video frame of 7500 bytes there is an overflow field of 230 bytes, and as mentioned, it is not always used.

Also, the MAC requires a continuous transmission of the real-time stations, in a fixed rate, in order to keep their time slot. As we already mentioned, we want our MAC to support all kinds of real-time applications, e.g. such that also generate packets for transmission in a sporadic way, as alarm signals.

The main characteristics of our new MAC protocol are:

1. It is simpler than other proposals and it is closer in nature to the inherent simplicity of the CSMA/CD MAC - its main requirements are the setting of different jam lengths, the detection of collisions and their termination, and the examination of only two types of received packets. All these tasks are performed anyway in CSMA/CD. In addition, it requires the stations to monitor the bus in consecutive predefined time intervals and to implement one integer variable and one flag which signals a station if it is allowed to transmit. All these requirements are simpler than those in other proposals, such as [17] and [18] .

2. Our assumptions on the real-time traffic are less stringent. Stations can transmit real-time traffic in different rates and in a sporadic way. The only restriction we impose is the definition of a predefined time interval such that every real-time station must transmit at least one packet during this time interval. This restriction is needed for the assignment of TAG numbers to stations by a distributed procedure. It can be removed however, if these numbers are configured.

3. It is CSMA/CD compatible and so standard Ethernet stations can transmit on the same network with stations implementing the new MAC.

4. It guarantees a bounded access delay to real-time transmissions. This feature bounds the jitter on these transmissions and this in turn reduces the buffering requirements of the receivers used to compensate this jitter. Although memory costs have fallen, there will nevertheless in many cases be a requirement to minimize on buffer memory at the receiver. A stream of several Mbit/s across an Ethernet network without guarantees would require a not insignificant amount of buffering at the receiver to avoid packets being dropped.

   Another importance to bounded access delay is to ease the operation of the time-out computation mechanisms of transport protocols such as TCP.

5. It enables to configure real-time stations with priority to access the bus which is independent of the number of transmitting real-time stations. This feature has an importance for e.g. alarm signals which may be very important in the Home environment.

The main assumptions used when designing the protocol are:

1. The Home LAN will be restricted in length - we initialy assume that it will not be greater than 40 meters in length. We also assume that it will be based on 10BASE5 or 10BASE2 physical media which can expand to the above length with no need for repeaters. These physical media can also support the anticipated number of stations in the Home environment.
2. Compared to the larger office or work-group LAN, the home LAN will have a small number of end nodes. In particular, it will have say, 6 'high priority end nodes' transmitting real time traffic as video or digital audio at any one time.

The rest of the paper is organized as follows: in Section 2 we define the system model. In Sections 3 and 4 we describe the new MAC. In Section 5 we give some upper bounds on the performance of the protocol and in Section 6 we prove its correctness. In Section 7 we propose two possible methods to transmit real time traffic by the new MAC and present simulation results. Finally, Section 8 summarizes the paper.

2 Model and Definitions

As mentioned, we consider an Ethernet-like network for the Home or small office environments. We anticipate two kinds of stations that can transmit in such a network. Firstly, stations that transmit real-time traffic. Such stations may need to transmit a frame typically every few 10s of ms, or can generate sporadic frames. However, in both cases a strict constraint on the time from when a frame is generated until it is transmitted successfully is required. Secondly, we assume that there are stations that transmit other traffic streams which are not sensitive to delays and also have much more random characteristics. The 1-persistent CSMA/CD MAC used in Ethernet is defined in the IEEE 802.3 standard [1], and because of possible collisions, it cannot guarantee a bounded access delay to frames transmitted in the system. However, in order to enable efficient real-time transmissions, we would like to somehow bound the access delay. We achieve this bound by changing the MAC of the stations that transmit real time traffic and enable them a higher priority in their access. Therefore, from now on we denote the stations that generate real-time traffic by *High Priority* stations and the other stations which will continue to access the bus by the regular 1-persistent CSMA/CD by *Low Priority* stations.

Throughout the paper, we will use several notations that we now explain. According to the IEEE 802.3 specification [1], a station always transmits a Preamble before its actual data frame, and in the event of collision it begins to transmit a Jam. Also, a station can initiate transmissions only after waiting an Inter Frame Gap from when the last end of carrier was detected. We will denote the times that it takes to transmit a Preamble, a Jam, and the time of the Inter Frame Gap by P, J and IFG respectively. Also, we denote by EOC the event in which an End Of Carrier is detected on the bus, by τ we denote the one way propagation delay in the bus and by δ we denote the maximum time that it takes for a station to detect a change in the transmission pattern on the bus, i.e. to detect carrier, detect collision, detect that a collision is over or to detect that the channel becomes idle. In our model we assume that τ is in the order of the transmission time of 2 bits, i.e. $0.2\mu s$ in a $10Mb/s$ channel, which corresponds to about 40 meters bus. δ is in the order of 10 bit time, i.e. about $1\mu s$.

The mechanism to resolve collisions between Low and High priority stations in our new MAC requires that the inequality $J > 2\tau + 2\delta$ will hold [21]. J is 32 bit times according to [1]. With $\delta = 10$ bit times it must hold that $\tau < 6$ or that the system can be at most 120 meters in length. This seems to be a reasonable restriction for a Home LAN.

As already mentioned, we assume a 10BASE5 or 10BASE2 physical media with no repeaters [1]. Notice that according to the IEEE 802.3 standard, a repeater blocks the forwarding of a signal from a segment if it is longer than a pre-defined number of bit times. We do not encounter such a problem because our system model does not contain repeaters. This problem can happen in other proposals such as [17] and [18].

3 The new MAC

Besides changing the rules by which High Priority stations access the bus, the High Priority stations also use a new frame structure for their transmissions. We begin by describing this new structure and later describe the access rules.

3.1 Frame structure

We distinguish between two possible kinds of frames. In Figure 1(A) we show the standard frame structure of Ethernet [1]. This structure is extended in Figure 1(B) to be the structure

Preamble	SFD	DA	SA	Type/ Length	DATA	FCS

(A)

Preamble	SFD	DA	SA	Type/ Length	Collision Bit	TAG	DATA	FCS	EFD	FILLER

(B)

Figure 1: (A) The standard frame structure in Ethernet (B) The frame structure used by High Priority stations

used by the High Priority stations.

First, the new frame structure will contain a new type value in the Type/Length field. We denote this type by 'H' since frames of this type are only transmitted by High Priority stations. When High Priority stations receive frames of type 'H' and transfer them to the MAC layer client, this new type value is also transfered to the client. This is similar to what is done in the IEEE 802.3 standard [1]. By this new type value a received frame is transferred to the real-time application.

Second, we argue that Low Priority stations that do not run real-time applications cannot be the destination of frames of the new type. It can be that either real-time transmissions will be destined to configured destination addresses, as might be the case with alarm signals, or a connection in the transport layer must first be established, before information transfer. During this connection establishment, the MAC addresses of the recipients are discovered, e.g., through ARP [20] .

It turns out that High Priority frames are transmitted in the network and are received by Low Priority stations. However, they will be immediately rejected by the recognize address test in the MAC layer, which is the first to be performed when a frame is received [1].

In addition to the new Type value, frames of Type 'H' also have four new fields, the TAG, Collision Bit, EFD (Ending Frame Delimiter) and Filler. The TAG field relates to a special number that every High Priority station has in the new MAC protocol, and that we call a TAG. The TAG field contains the TAG number of the High Priority station that transmitted the frame. The need for this field will be explained later. The EFD is used to signal the end of the frame. The Collision Bit field is used by the transmitting High Priority station to signal if its transmission was preceded by a collision with other High Priority station(s) (Collision Bit= 1). The way in which a collision with other High Priority station(s) is detected is explained later. The Filler is a random sequence of bits of length $2\tau + \delta$ time units, i.e. in the order of 14 bits. The Filler prevents Low Priority stations from sensing an end of carrier between successive High Priority frames. At this point, other High Priority stations with frames to transmit will transmit a Jam which has a maximum duration which is proportional to the stations' TAG number (a number which is unique to the individual station). This provides a mechanism which ensures that the station with the Highest TAG number, of those attempting to transmit, will ultimately find that it is the only station transmitting on the bus. At this point, it will begin to transmit the Preamble followed by the rest of the frame. We describe the operation of the protocol in more detail in the following sections.

3.2 Access rules - general description

In this section we give a general description of the access rules in the new MAC protocol. The only event that prevents successful transmissions is stations colliding. Therefore, the main change in the MAC of the High Priority stations, compared to the standard IEEE 802.3 CSMA/CD, is the way in which High Priority stations handle collisions. We distinguish between the three possible collisions that can happen in the network:

1. Collisions in which only Low Priority stations are involved.
2. Collisions in which Low Priority station(s) and one High Priority station are involved.
3. Collisions in which at least two High Priority stations are involved.

The first type of collisions is handled by the usual, standard way of CSMA/CD. Every Low Priority station that is involved in a collision either defers and retransmits the frame again according to its backoff interval, or it drops the frame if the number of collisions that the frame already had exceeds an upper bound of 16 [1].

In the second type of collisions, the Low Priority stations defer as in the first type. However, the High Priority station persists with its transmission until it detects that the channel is clear of all the other transmissions. This will happen because the Low Priority stations defer. When the collision is over, the High Priority station begins to transmit its frame again and this time it is guaranteed that its transmission is successful because the carrier on the bus prevents other stations from attempting transmission.

The above is accomplished by lengthening the Jam (as compared to normal Ethernet transmitters) that a High Priority station transmits, when it collides, to $J + (2\tau + 2\delta)$, i.e. we expand the standard Jam by about 24 bits. We denote this Jam by short-Jam. The High Priority station transmits the short-Jam until it either detects that the collision is over or the short-Jam ends. In [21] we show that if a High Priority station collides with Low Priority stations only, then if it is transmitting the short-Jam it finds it remains the only station that transmits in the system. It then will detect that the collision is over and it can begin to transmit its frame successfully.

In the third type of collisions, Low Priority stations defer as in the first and second types. The only necessary procedure is to distinguish between the transmissions of the High Priority stations that were involved in the collision. The short-Jam is not sufficient for this purpose because it is transmitted identically by all the High Priority stations that are involved in the collision.

First we note that in [21] we show that all the High Priority stations that are involved in the collision will recognize that they collide with other High Priority stations by detecting a collision all through the transmission of the short-Jam. In order to distinguish between the colliding High Priority stations, and to enable one of them to continue and transmit its frame successfully, High Priority stations continue and transmit a long-Jam after their short-Jam is finished, i.e. after transmitting the short-Jam, a High Priority station continues to transmit an additional Jam which is denoted by long-Jam. However, the length of the long-Jam is different at every High Priority station. The length is determined by a special number, denoted by TAG, different at every High Priority station. The long-Jam is set to be TAG$\cdot(2\tau + 2\delta)$ time units which is about TAG$\cdot$24 bits. If we assume that no more than 6 High Priority stations transmit at the same time then the length of the long-Jam is bounded by 144 bits. Thus, a High Priority station with a higher TAG number transmits a longer long-Jam, persists longer, and thus has a higher precedence than a High Priority station with a lower TAG number in acquiring the bus. The minimum value of a TAG number is 1 .

If during the entire transmission of the long-Jam, a High Priority station continues to detect a collision, it defers when the transmission of the long-Jam ends. Otherwise, if during the transmission of the long-Jam, a High Priority station suddenly detects that the collision is over, it immediately begins with its frame transmission again.

At this point, before continuing with the protocol description, we would like to explain two terms that we later use in the paper. The first term is 'Cycle' which is the maximum continuous transmission of frames by High Priority stations such that the order among the transmitting High Priority stations is decreasing in their TAG numbers. During a cycle the

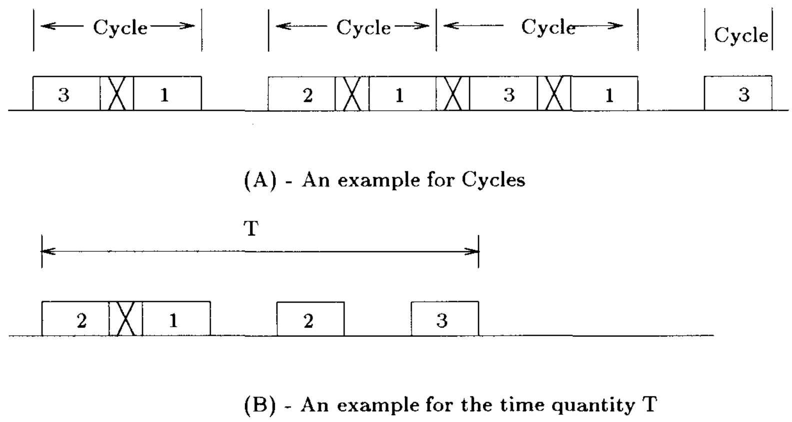

Figure 2: Examples for Cycles and the time quantity T

bus is not idle and a High Priority station begins to transmit its frame during the Filler of the frame transmitted by the previous High Priority station (which has a higher TAG number).

The second definition we introduce is the definition of a time quantity T which is defined as the shortest time interval in which it is guaranteed that every High Priority station transmits a frame. The value of T depends on the pattern of traffic generation of the High Priority stations. For example, consider Figure 2. In this Figure we assume that only High Priority stations with TAG numbers 1, 2 and 3 are transmitting in the system. Figure 2(a) shows four cycles: in the first one only the stations with TAG numbers 1 and 3 are transmitting. In the second cycle only the stations with TAG numbers 1 and 2 are transmitting and in the third cycle only the stations with TAG numbers 1 and 3 are transmitting. Notice that the second and third cycles are continuous. Finally, in the fourth cycle only the station with TAG number 3 is transmitting.

On the other hand, Figure 2(b) shows the time interval T in which all the stations with a TAG number have transmitted at least one frame. In the example the station with TAG number 2 transmitted two frames while the stations with TAG numbers 1 and 3 transmitted one frame.

We now return to the description of the new MAC and demonstrate its operation by using Figure 3. We assume that the High Priority stations I, J and K with the TAG numbers 1, 2 and 3 respectively collide, as depicted in the upper diagram of the figure. All the three High Priority stations begin to transmit when they detect an idle bus, i.e. they do not detect any carrier. They detect a collision of the third type and transmit a long-Jam, the length determined by their TAG numbers. Thus, K persists for the longest time and as we show in [21], K remains the only station transmitting in the system and it transmits its frame

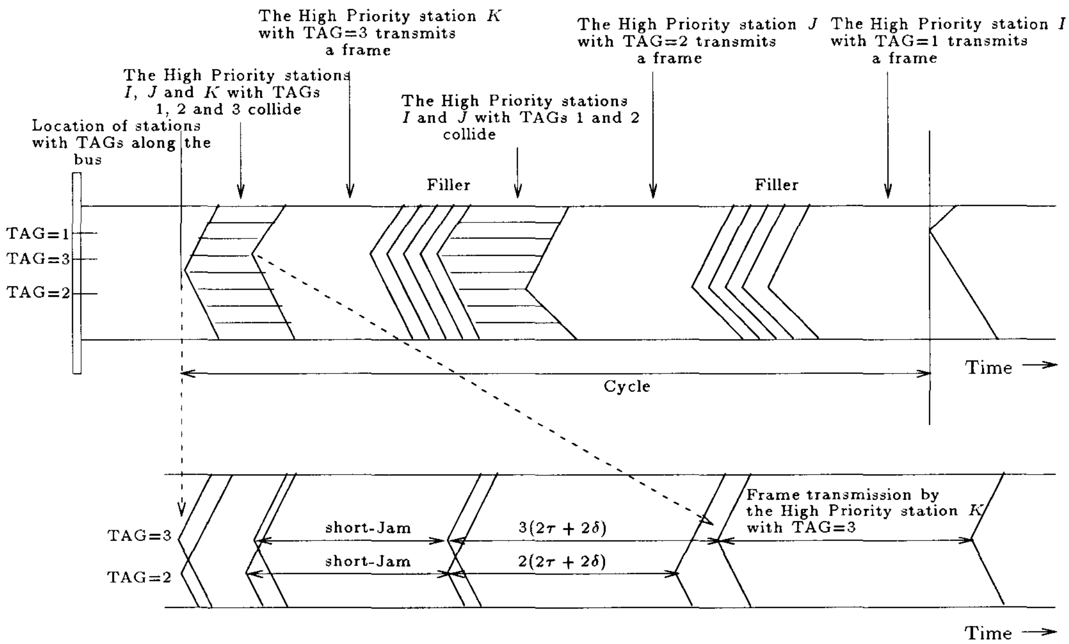

Figure 3: Collision resolution among High Priority stations. The lower part of the figure is an extension to a collision resolution interval in the upper part.

successfully. The collision resolution between K and J is shown in the lower diagram of Figure 3. The same resolution process happens between K and I.

I and J wait until they detect that K completes its transmission and then they both attempt transmission again. In fact the stations wait until K begins to transmit the Filler and then begin to transmit the long-Jam only. The Filler is sent after the useful information and thus the collision with the Filler does not damage the frame data of K which has already been sent at this point. As will become clear later, it is enough to resolve the collision between I and J by the long-Jam only, and therefore they begin their transmission attempt with the long-Jam.

The collision with the Filler keeps the bus occupied in order that Low Priority stations will not attempt to transmit after the transmission of K. I and J collide but now J persists for the longest time period and it will transmit its frame successfully. Finally, I is the only High Priority station transmitting and it will be able to transmit its frame successfully. Notice that I will collide with the Filler of the frame of J but with no other High Priority stations. As mentioned, we denote the time interval from when the three stations begin to transmit and until I finishes to transmit its frame by a *Cycle* (see Figures 2 and 3).

In addition, High Priority stations maintain a Round Robin service. In order to achieve this service, a High Priority station is allowed to transmit in a Cycle only once, and only if stations with higher TAG numbers have transmitted in the Cycle so far. The way by which a High Priority station recognizes the end of a cycle is explained later. In addition, by receiving the frames transmitted in a cycle so far, a High Priority station recognizes if they were all transmitted by High Priority stations with higher TAG numbers. As we show later, these rules guarantee an upper bound on the access delay of High Priority stations of one maximum Cycle length plus the maximum transmission time of a frame.

Also, notice that in principle, a High Priority station can join and transmit in a Cycle in its middle, e.g. assume the same scenario as in Figure 3 but that stations with TAG numbers 1,2 and 4 are involved in the first collision instead of stations with TAG numbers 1,2 and 3 as the figure shows. The station with TAG=4, say L, will win the collision and assume that during L's transmission station K with TAG=3 receives a frame to transmit. K will try to transmit its frame after L terminates its transmission because it notices a new cycle, it is allowed to transmit one frame in a cycle and only stations with a higher TAG number have transmitted in the cycle so far. K will collide in the second collision with the stations that have the TAG numbers 1 and 2 and it will 'win' the collision.

3.3 Access rules - detailed description

We now specify in detail the access rules of High Priority stations. A High Priority station is allowed to begin to transmit in three cases:

1. The bus is idle and more than IFG time units have elapsed since the last EOC on the bus.
2. The bus is idle and less than IFG time units have elapsed since the last EOC on the bus.
3. Immediately after detecting the EFD field of another frame transmitted by a High Priority station.

We now describe the three cases in detail which are shown in Figure 4.

1. The bus is idle and more than IFG time units have elapsed since the last EOC:

   In this case a High Priority station begins to transmit its Preamble, and if it detects a collision, it transmits the short-Jam and long-Jam, as was described in subsection 3.2. In [21] we prove that the High Priority station with the highest TAG number will 'win' the collision and will transmit its frame successfully.

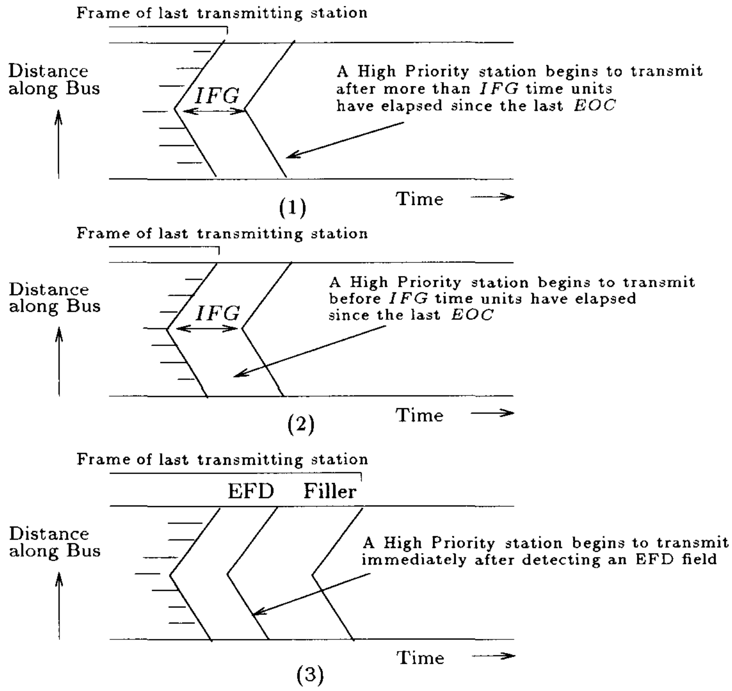

Figure 4: The cases when a High Priority station can begin to transmit

2. The bus is idle and less than IFG time units have elapsed since the last EOC:

   This case is similar to Case 1 except that the High Priority station can also collide with Low Priority stations that begin to transmit within at most $\tau + \delta + IFG$ time units after the High Priority station began its transmission in contrast to at most $\tau + \delta$ time units in Case 1. This is because according to the IEEE 802.3 standard [1], Low Priority stations can begin to transmit after IFG time units elapsed since the last EOC was detected, regardless of whether there is a carrier on the bus or not. This possibility is shown in Figure 5 . These possible collisions with Low Priority stations distinguish between this case and Case 1.

   In order to handle the above mentioned collisions with Low Priority stations, a High Priority station always starts a timer after detecting an EOC, irrespective of whether it has a frame to transmit or not. The timer is set to $IFG + P$ time units. If the High Priority station begins to transmit and it collides before the timer expires, it transmits a Jam until the timer expires. Then it begins to transmit the short-Jam and the long-Jam as in Case 1.

   By using the timer, when High Priority stations collide, they transfer themselves to the case where they started to transmit after at least IFG time units have elapsed since the

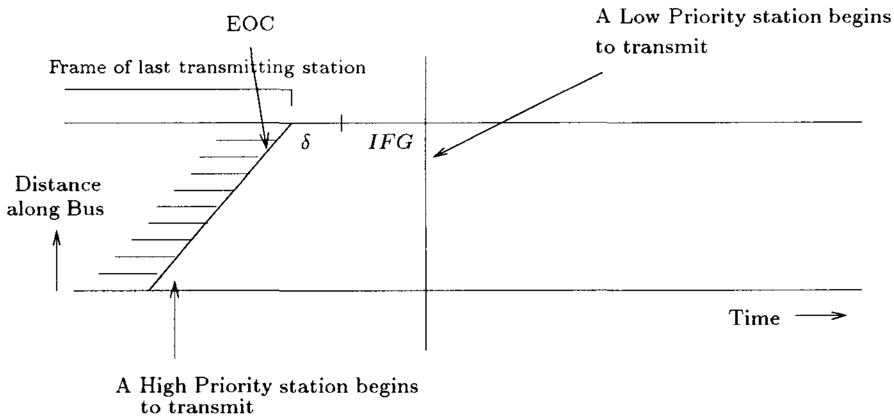

Figure 5: A High Priority station collides with a Low Priority station that begins to transmit IFG time units after an EOC is detected

last EOC and collisions are resolved in the same way as in Case 1.

3. Immediately after detecting an EFD field.

   In this case a High Priority station begins to transmit immediately after detecting the EFD field of a previously transmitted 'H' frame. Since the Filler field is of length $2\tau + \delta$ time units, it can be easily verified that there is actually a 'hand shaking' between the High Priority station that finished transmitting and the High Priority station(s) that begin to transmit after the EFD, and so no station in the system detects an idle bus.

   Keeping the entire bus occupied prevents Low Priority stations from attempting transmission. Therefore, a High Priority station that begins to transmit after an EFD can collide with other High Priority stations only, and it clearly collides with the Filler that follows the EFD field. As we prove in [21], it is sufficient in this case that High Priority stations will transmit the long-Jam only. We show in [21] that the High Priority station with the highest TAG will 'win' the collision in this case also.

   Finally, if the High Priority station that transmitted the EFD is allowed to try to transmit an additional frame immediately, then it can begin to transmit the long-Jam after the EFD. However, it must transmit the first $2\tau + 2\delta$ time units of the long-Jam, even if it does not detect a collision during this time. The first $2\tau + \delta$ time units serve as a Filler.

Notice that in Cases 1 and 2 a High Priority station can detect a collision with other High Priority station(s) if throughout the transmission of the short-Jam it detects a collision. In this case, the High Priority station that 'wins' the collision sets the Collision Bit in its frame. In Case 3 a High Priority station collides with the Filler of the previous 'H' frame. A Filler is of length $2\tau + \delta$ time units. Therefore, if after $2\tau + 2\delta$ time has elapsed since it began the transmission of the long-Jam the High Priority station does not detect a collision anymore, it means that no other High Priority station has begun to transmit. Otherwise, if a collision is detected then it means that at least one another High Priority station has begun to transmit and there is a collision between High Priority stations.

As mentioned, in order to guarantee a Round Robin service to the High Priority stations, and by this to guarantee a bounded access delay to the bus, we defined the term *Cycle*. A Cycle is a continuous transmission of High Priority stations in a decreasing order of their TAG numbers, where a transmission of a High Priority station, except possibly the first one, always begins after the EFD field of the frame transmitted by the previous High Priority station. An example of a Cycle with three frame transmissions is shown in Figure 3.

A High Priority station is allowed to transmit at most one frame in a Cycle. In addition, in order to ensure an upper bound on the access delay (the precise upper bound is computed in Section 5), a High Priority station can also attempt to transmit in a Cycle only if other High Priority stations with higher TAG numbers have transmitted in the Cycle so far. Therefore, the last frame in a Cycle is transmitted with no collision with any other High Priority station and a frame with a Collision bit=0 signals the end of a Cycle. Clearly, an EOC can also signal the end of a Cycle when High Priority station(s) fail or do not have frames to transmit.

Notice that the new protocol is robust in the sense that if there are no transmissions errors, the protocol operates as described. This can be explained as follows: the protocol is based on different length Jam signals, different TAG numbers assigned to High Priority stations and on a Round Robin access policy.

Clearly, the main issue is the assignment of unique TAG numbers. As we explain in Section 4 it is guaranteed that if there are no transmission errors, High Priority stations will be assigned unique TAG numbers. Moreover, if transmission errors occur and lead to a situation where several High Priority stations have the same TAG number, it is guaranteed that after the transmission errors are over the High Priority stations will be re-assigned unique TAG numbers.

Therefore, it is guaranteed that when there are no transmission errors, eventually High Priority stations transmit with unique TAG numbers. They also detect the TAG numbers of High Priority stations transmitting in a cycle, they detect the Collision Bit in frames and with the detection of EOCs, they maintain a Round Robin access policy.

Notice also that the protocol is simple for implementation. The detection of collisions, and their termination, are basic features of the current standard hardware of 802.3 stations. Also, setting the Jam signal to different lengths is an easy task.

Next, the end of a cycle is detected by reading the Collision Bit of received frames or by detecting EOC. Every High Priority station shall process any frame received because it might be destined to itself. Therefore, the task of reading the Collision Bit in received frames is not a complication. Also, the detection of EOC is a basic feature of the current hardware of 802.3 stations. Therefore, maintaining the Round Robin access policy is a simple matter.

It turns out that the only complex component of the new protocol is the procedure for the assignment and update of the TAG numbers. However, this is a simple procedure for implementation since as will be described in Section 4 it involves only comparisons and additions of equal bit length numbers.

4 Assignment and update of TAG numbers

In this section we define a procedure by which High Priority stations are dynamically assigned and update their TAG numbers. The procedure guarantees that High Priority stations always transmit with unique TAG numbers and that they always reduce their TAG numbers, if possible, in order to shorten the collision resolution intervals. These attributes are proved later, in Section 6.

In the following we use the term 'a High Priority station *with* a TAG number'. This term means that the High Priority station has a TAG number and it can use it for transmissions as described in Section 3. This is in contrast to the possibility of a High Priority station not having a TAG number or it has one but it cannot use it for transmissions as described in Section 3.

We emphasize the term '*with* a TAG number' because the procedure is based on the time quantity T that was already mentioned (see Figure 2), such that it is guaranteed that in any time interval of T time units at least one 'H' frame is transmitted by every High Priority station *with* a TAG number. We will show later in Section 6 that such a time quantity T can be defined, and its size is based on how often High Priority stations transmit traffic. The size of T has an impact on the size of the time interval from when a High Priority station wants to begin to transmit High Priority traffic and until it obtains a TAG number and can actually begin to transmit. For this matter, T shall be as small as possible. However, if there is an High Priority application in which the time interval between successive transmissions is large, then this will enforce a large T. In this application, High Priority stations shall transmit dummy packets, without any data, only to inform on their TAG numbers. An actual, appropriate value for T depends on the implemented applications.

The procedure to obtain and update TAG numbers is based on the High Priority stations monitoring the bus in consecutive intervals of T time units. The procedure is as follows:

Step 1: At this step a High Priority station does not have any TAG number yet. In order to get one, it reads all the frames of type 'H' that are transmitted in the system for a time interval of T time units. By the definition of T, it is guaranteed that the High Priority station will read during this interval all the TAG numbers that are currently in use by High Priority stations. At the end of the T interval the High Priority station adopts for itself the maximum TAG number that was received from the bus during the interval, plus 1, i.e. if x is the maximum TAG number that was read then the station adopts number $x + 1$. At this stage the TAG number is considered by the station to be *negotiated*, for a reason specified below. After deciding on a negotiated TAG number, a High Priority station keeps on monitoring the bus in consecutive T intervals.

Step 2: After deciding on a negotiated TAG number, the High Priority station tries to transmit frames of type 'H' with the negotiated TAG written in the TAG field of the frame. However, the transmissions are still with the High Priority station acting as a Low Priority station, carrying out the standard Ethernet CSMA/CD protocol. In order to guarantee unique TAG numbers to High Priority stations, a High Priority station, while trying to transmit a 'H' frame with the negotiated TAG, also listens to the bus. If it detects that another High Priority station has succeeded in transmitting a frame with an equal or a higher TAG number, it defers and does not try to transmit its 'H' frame anymore. The High Priority station then waits up until the end of the next T interval and then goes back to the start of Step 1, adopting a new negotiated TAG number, and then proceeds again to Step 2 and so on.

Notice that if two or more High Priority stations begin to transmit at about the same time, it can happen that two or more of them will adopt the same negotiated TAG. In this event, one of these stations will transmit a frame first and proceed to Step 3, and the other stations will go back to Step 1, adopting a new TAG number.

When a High Priority station finally succeeds in transmitting its first 'H' frame with the negotiated TAG, this number becomes a *permanent* TAG number. A High Priority station with a permanent TAG considers itself to be *with* a TAG number and it transmits as described in Section 3.

Step 3: A High Priority station I with a permanent TAG number x continues to read all the frames of type 'H' that pass on the bus in continuous intervals of T time units. (Notice that a High Priority station does not reset its timer that measures the T intervals after transmitting its first 'H' frame with the negotiated TAG number). At the end of each T interval I updates its TAG number to be the maximum TAG number that was received from the bus and that is smaller than x, plus 1, e.g. if $x = 6$ and the maximum TAG number that is received from the bus and that is smaller than 6 is 3, then I updates its TAG to be $3+1=4$. Clearly, when a frame with TAG equals to $x-1$ is monitored, I's TAG number does not change. As mentioned, the attempt to reduce the TAG numbers is in order to shorten the collision resolution interval among High Priority stations, and thus maximize network efficiency.

As we prove later, this scheme ensures a unique, permanent TAG numbers to the High Priority stations. However, if because of transmission errors or others, while monitoring the High Priority frames, a High Priority station notices the error situation of another station transmitting with the same TAG number as its own, then it returns to Step 1 to obtain another TAG number which is unique.

5 Performance and overhead computation

In this section we prove two upper bounds: one is on the length of a Cycle and the second is on D, the maximum access time, i.e. the time that elapses at a High Priority station from when a frame arrives at the head of its transmission queue until the end of its transmission into the bus. We also make some notes about the scaling of the protocol and its overhead.

We assume in our proofs that High Priority stations always transmit with unique permanent TAG numbers. We prove in Section 6 that the procedure of Section 4 indeed guarantees this property.

We denote the length of a Cycle by CYC and by MFL the largest transmission time of a frame of any type.

Lemma 5.1: Let *1,2,...,M* be the set of possible TAG numbers. Then CYC is bounded by $\frac{M(M-1)}{2} \cdot (2\tau + 2\delta) + M.(MFL + \tau + \delta) + IFG + P + J + 2(2\tau + 2\delta)$ time units.

Proof: A Cycle is composed of continuous transmissions of High Priority stations. A High Priority station attempts to transmit immediately after detecting the EFD field of the previous 'H' frame, except possibly for the first 'H' frame in the Cycle. We now compute three quantities: First, when several High Priority stations begin to transmit after detecting an

EFD field, the length of time it takes, after the transmission of the EFD field is finished, to resolve the collision among them and then for the winning High Priority station to complete transmission of its frame. The second quantity is the time it takes to transmit the last 'H' frame in a cycle, which does not encounter a collision with other High Priority stations. Finally, the third quantity we compute is the time it takes to resolve a collision among Low and High Priority stations and to transmit a frame of type 'H'. This scenario can happen when transmitting the first 'H' frame in a cycle.

First, consider the scenario in which several High Priority stations contend for transmission after detecting the EFD field of a previously transmitted 'H' frame. In [21] we prove that the time that can elapse from when the transmission of the EFD field is finished, until the High Priority station with the highest TAG number completes its frame transmission, is bounded by $\text{TAG}_{max,2}.(2\tau + 2\delta) + (\tau + \delta) + MFL$ where $\text{TAG}_{max,2}$ is the second highest TAG number used by any of the colliding High Priority stations.

In a Cycle, the last transmission of a High Priority station does not collide with other stations but only with the Filler of the previously transmitted 'H' frame. In this case the High Priority station needs to transmit $(2\tau + 2\delta)$ time units of its long–Jam. Then, it is guaranteed that it detects that the collision with the Filler is over and it can transmit its 'H' frame. Notice also that it can take at most $(\tau + \delta)$ time units from a High Priority station finishing transmitting the EFD field until the last transmitting High Priority station in the Cycle begins to transmit its frame.

Finally, the first transmission in a Cycle can encounter collisions with Low and High Priority stations. In this case, from [21], the time interval from any station involved in the collision beginning to transmit until the High Priority station with the highest TAG number finishing transmitting its 'H' frame is bounded by $IFG + P + J + (\text{TAG}_{max,2} + 1)(2\tau + 2\delta) + (\tau + \delta) + MFL$ where again $\text{TAG}_{max,2}$ is the second highest TAG number used by a High Priority station that participates in the collision. This case corresponds to the case where High Priority stations begin to transmit within less than IFG time units after an EOC.

Summing the above, CYC is bounded by $\sum_{i=1}^{y-1}(\text{TAG}_{max,2}(i)(2\tau + 2\delta) + MFL + (\tau + \delta)) + IFG + P + J + (2\tau + 2\delta) + (MFL + (2\tau + 2\delta) + (\tau + \delta))$ where y is the number of 'H' frames transmitted in the cycle and $\text{TAG}_{max,2}(i)$ is the second highest TAG number used by a colliding High Priority station before the i_th 'H' frame is transmitted in the Cycle. The above term equals to $\sum_{i=1}^{y-1}\text{TAG}_{max,2}(i)(2\tau + 2\delta) + y(MFL + \tau + \delta) + IFG + P + J + 2(2\tau + 2\delta)$. In [21] we show that the maximum length of a Cycle is received when $y = M$ and in this case the maximum of $\sum_{i=1}^{y-1}\text{TAG}_{max,2}(i)$ is $1 + 2 + \ldots\ldots + (M - 1) = \frac{M(M-1)}{2}$. Thus, CYC is bounded by $\frac{M(M-1)}{2}.(2\tau + 2\delta) + M.(MFL + \tau + \delta) + IFG + P + J + 2(2\tau + 2\delta)$ time units.

□

Lemma 5.2: D is bounded by $CYC + MFL$ time units.

Proof: Consider a High Priority station with a frame to transmit. We divide the proof into two cases, based on whether or not the High Priority station can attempt transmission of the frame immediately after it arrives at the head of its transmission queue.

1. If the station cannot attempt transmission immediately, then it is only because the frame arrives during a Cycle in which either the station and/or another High Priority station with a lower TAG number has already transmitted.

Assume that the High Priority station under consideration has TAG number n. In the ongoing Cycle only stations with TAG numbers $1,...,n-1$ can still transmit. The next Cycle begins immediately after the transmission of the last 'H' frame in the on going Cycle, i.e. the High Priority station(s) begin to transmit after detecting the end of the EFD field of the last 'H' frame of the ongoing cycle. In the next Cycle the station under consideration will be the $(M - n + 1)\_th$ station to transmit its 'H' frame at the latest. Thus, in [21] we show that $D \leq CYC - (IFG + P + J + (2\tau + 2\delta))$ time units.

2. If the High Priority station can attempt transmission immediately when the frame arrives at the head of its transmission queue, then, if the bus is idle, clearly the station can transmit within CYC time units. If the bus is occupied then it can happen that the station attempts transmission (i) while a Low Priority station transmits, (ii) while a High Priority station transmits, (iii) when only Low Priority stations collide, (iv) when only High Priority stations collide or (v) when Low and High Priority stations collide. From [21], the time intervals that can pass in each of the above cases until the station transmits its frame are $MFL + CYC$, $MFL + CYC - (IFG + P + J + (2\tau + 2\delta))$, $P + J + (3\tau + 2\delta) + CYC$, CYC and $P + J + (2\tau + 2\delta) + MFL + CYC - (IFG + P + J + (2\tau + 2\delta))$ respectively. Among all these possibilities, the longest time interval is $MFL + CYC$ time units.

□

As we already mentioned, our new protocol is suitable for short buses and is efficient for a relatively small number of High Priority stations. This can be shown quantitatively by observing the overhead due to the collision resolutions and the handshake mechanism in a cycle. From Lemma 5.1 this overhead sums to $M^2 \cdot (\tau + \delta)$ time units. If we also add the SFD and the Filler fields of the High Priority packets to the overhead, and we neglect the fact that several long-Jam transmissions overlap with the Filler transmissions, we receive a total overhead of $M^2 \cdot (\tau + \delta) + M \cdot (2\tau + \delta)$ time units, where M is the number of High Priority stations, τ is the one way propagation delay in the system and δ is the maximum time to detect a change in the transmission pattern on the bus (whether the bus is clear, occupied or there is a collision). δ is not dependent on the size of the network. Also notice that the overhead is not linear in M but is a function of M^2.

In order to evaluate the overhead lets assume three possible values for M : $M = 2$, $M = 10$ and $M = 20$. Lets also assume that when $M = 2$ one station transmits a 1.5Mb/s CBR video stream and the other transmits a 64Kb/s PCM voice stream. When $M = 10$ and $M = 20$ two stations transmit CBR video and the rest transmit PCM voice. Packets are generated every 40ms. Lets also assume two system sizes, of 40 meters ($\tau = 2$ bits times) and of 100 meters ($\tau = 5$ bit times). We also assume that $\delta = 10$ bit times and notice that with this figure the system cannot be longer than 120 meters from the correctness point of view. See Section 2.

It can be easily computed now that for $M = 2$ the overhead is about 0.1% for the two system sizes. When $M = 10$ the overheads are 0.9% and 1.1% for the 40 meters and the 100 meters system sizes respectively, and for $M = 20$ the overheads are 2.9% and 3.7% respectively. Notice that $M = 20$ does not seem to be realistic for a Home LAN. Also, for $M = 10$ if we assume that less than 8 stations transmit 64Kb/s PCM voice but several of them transmit real-time streams with higher rates, then their packets are longer and the total overhead is less than 1% .

6 Correctness

Until now we defined a procedure by which High Priority stations are assigned and update their TAG numbers. The procedure is based on the assumption that a time quantity T can be defined such that it is guaranteed that in any interval of T time units every High Priority station *with* a TAG number transmits a frame. Then, in Section 5 we proved that if High Priority stations transmit with unique TAG numbers, then their access delay is bounded.

In this section we prove that given T, the procedure of Section 4 guarantees that High Priority stations transmit *with* unique TAG numbers or by the terms of this procedure, they have unique permanent TAG numbers. (Lemma 6.1). Then, in Lemma 6.2 we prove that given that the access delay of High Priority stations is bounded, we can define a T such that every High priority station with a TAG number, which has a frame to transmit, transmits a frame in every interval of T time units. Finally, in Theorem 6.1 we prove that when the system operates with the procedure of Section 4 with T taken from Lemma 6.2, High Priority stations always transmit with unique TAG numbers and their access delay is bounded by D of Lemma 5.2 .

In lemma 6.1 below we prove that given T, the procedure of Section 4 ensures that High Priority stations always have unique permanent TAG numbers. In order to prove Lemma 6.1 we first prove two claims, Claims 6.1 and 6.2 . In these claims we also assume that a time quantity T exist and that High Priority stations follow the procedure of Section 4.

Claim 6.1: Let High Priority station J adopt its first permanent TAG number j at time t_1 by transmitting its first 'H' frame. Assume that at this time station I has a permanent TAG number i. Consider now the time interval after t_1 when both I and J have a permanent TAG number.
(a) If $j < i$ then J will always have a smaller permanent TAG number than that of I.
(b) If $j > i$ then J will always have a larger permanent TAG number than that of I.
Proof:
(a) After adopting a permanent TAG number a High Priority station monitors the bus in consecutive intervals of T time units. At the end of every T interval the High Priority station decides whether it can reduce its TAG number or not. Therefore, time t_1 is included in a T interval of I and I detects a frame with TAG number j such that $j < i$. Therefore, I cannot adopt a TAG number smaller than $j + 1$ at the end of its considered T interval. By the definition of T, in the following T interval I will detect again a 'H' frame from J and will not be able to adopt a permanent TAG smaller than that of J. Similarly, this holds at the end of every T interval of I.

(b) By the definition of T, during $[t_1, t_1+\mathrm{T}]$ station J detects a frame from I with a TAG number i' such that $i' < j$. Thus, J cannot adopt a TAG number smaller or equal to that of I at the first time after t_1 when a T interval is finished. This condition holds later, at the end of every T interval of J.

$\square$

Claim 6.2: Let High Priority station J adopt its first permanent TAG number j at time t_1 by transmitting its first 'H' frame. Then, at t_1 there is no other High Priority station with the same permanent TAG number j.
Proof: Assume that the claim is not true and that at t_1 it is violated for the first time, i.e. at t_1 there is another High Priority station I with permanent TAG number j. The fact that J adopts a permanent TAG number at t_1 means that there is an earlier T interval, $[t_0-\mathrm{T}, t_0]$, through which J monitored the bus, decided on j as its negotiated TAG number and then,

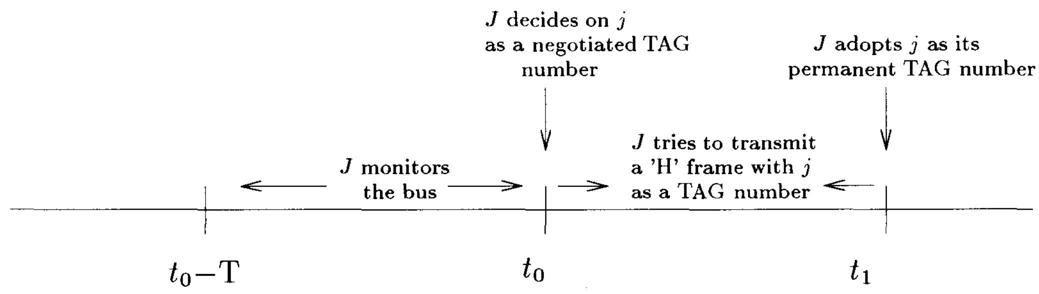

Figure 6: Time diagram for the proof of Claim 6.2

during $[t_0, t_1]$, it tried to transmit its first 'H' frame until success at t_1. See Figure 6.

Consider the interval $[t_1-\mathrm{T}, t_1]$ which is included in $[t_0-\mathrm{T}, t_1]$. It must hold now that I transmitted a 'H' frame in $[t_1-\mathrm{T}, t_1]$. This is because that otherwise, if this is not true, then I has a permanent TAG number all through $[t_1-\mathrm{T}, t_1]$ and it turns out that there is an interval of T time units before t_1 in which I did not transmit a 'H' frame. This contradicts the assumption that a High Priority station with a permanent TAG number transmits at least one packet in every time interval of T time units.

The 'H' frame that I transmitted in $[t_1-\mathrm{T}, t_1]$ must have an equal or a higher TAG number than j but this contradicts the assumption that J still transmits a 'H' frame with a negotiated TAG number j at t_1 (rule 2 in Section 4). Thus, the claim holds.

□

Lemma 6.1: The procedure of Section 4 guarantees that High Priority stations always have unique permanent TAG numbers.

Proof: Assume by contradiction that the lemma does not hold and let t_1 be the first time when there are at least two High Priority stations, I and J, with the same permanent TAG number x. Also assume that station J is the station that adopts x as its permanent TAG number at t_1, i.e. x is already the permanent TAG number of I at t_1.

Notice that Claim 6.2 holds at t_1 and so it cannot happen that J adopts x as its first permanent TAG number at t_1. Therefore, it must hold that J adopts x at the end of a T interval by updating its TAG number. In this case assume without loss of generality that J adopted its first permanent TAG number after I, at time t_0 say. Assume that at t_0 station J adopted permanent TAG number j and by the assumption on t_1, I has a different permanent TAG number i at this time. Let $i < j$. By Claim 6.1 there cannot be a scenario that will lead station J to adopt a TAG number which is equal or smaller than that of I, contradicting the situation at t_1. The same line of proof holds if $i > j$ or if I adopted its first TAG number after J.

$\square$

In Lemma 6.2 below we prove that given that the access delay of High Priority stations is bounded by a quantity D', then a time quantity T can be defined. In the lemma we use X which is defined as the maximum time interval between the generation of two consecutive frames at any High Priority station.

Lemma 6.2: Given a finite D', during an interval of T$= max(X, MFL) + D'$ time units, every High Priority station with a permanent TAG number transmits at least one complete 'H' frame. Moreover, the above is the minimum necessary and sufficient value for T.

Proof: Consider a High Priority station I with a permanent TAG number and an interval of T time units that begins at a time t_0. If I does not have a frame to transmit at t_0 then sometime before time $t_0 + X$ it generates a frame and by the given D' sometime before time $t_0 + X + D' \leq t_0$+T this frame is transmitted successfully.

If station I has a frame to transmit at T_0 but it is not in the middle of transmitting the frame, then by the given D' it is guaranteed that sometime before time $t_0 + D' < t_0$+T this frame is completely transmitted. If station I is transmitting the frame at t_0 then sometime before time $t_0 + X$ it generates another frame. The attempt to transmit this new frame will start at $t_0 + max(MFL, X)$. By the given D', the transmission of the new frame will end at $t_0 + max(MFL, X) + D'$ at the latest, i.e. at t_0+T . This case dictates the minimum necessary and sufficient value for T.

$\square$

We would like to mention at this point that if X above is large then it might happen that T would be too large for a High Priority station to start a session since it would take it a long time until it acquires a TAG number. In this case High Priority stations will need to transmit dummy packets in order to reduce X and T.

Finally, we prove in Theorem 6.1 below that High Priority stations always transmit with unique TAG numbers and their access delay is bounded.

Theorem 6.1: Assume that High Priority stations follow the procedure of Section 4 by using T as defined in Lemma 6.2, and D' equals to D as defined in Lemma 5.2.
Then:

(a) High Priority stations always have unique permanent TAG numbers.

(b) The access delay of High Priority stations is bounded by D as defined in Lemma 5.2.

Proof: When the system begins to operate at time t_0 say, High Priority stations do not have TAG numbers. Thus, it is true that at t_0 High Priority stations transmit with unique TAG numbers because they do not transmit by the scheme of Section 3 at all.

Assume now that (a) does not hold for the first time at time t. During $[t_0, t]$ High Priority stations have unique permanent TAG numbers. Thus, Lemma 5.2 holds during this time interval and so Lemma 6.2 holds. Thus, by Lemma 6.1 it cannot happen that (a) is violated at t. Therefore, (a) always hold and so also (b).

□

7 Transmission methods and simulation results

As we proved in Sections 5 and 6, the new MAC guarantees a bounded access delay to High Priority stations. It also enforces Round Robin transmissions in the case that High Priority stations have a continuous stream of frames to transmit. We suggest two methods by which High Priority stations gather information and decide on the times when they attempt transmissions.

7.1 The 'Stream method'

In this method time is divided into intervals of L time units. Any data generated, e.g. video, is placed by the video module into a buffer as soon as it is generated. At the end of each interval a High Priority station collects all the data in its buffer and generates a frame to transmit, e.g. if a High Priority station is a $4Mb/s$ CBR video source and $L = 4ms$, then the station generates frames of 16000 bits every $L = 4ms$ time units.

Notice that according to the Ethernet and IEEE 802.3 standard [1], the maximum size of frames is about 12000 bits. Here we allow High Priority frames to be longer, e.g. in the example above to be of about 16000 bits.

In this method the data is considered to be simply a stream of bytes and no attention is given to any internal syntax of the stream, e.g. the video framing structure as Transport stream packets when an MPEG video source is considered [22] (hence the term 'Stream'). Clearly, the total transmission time of the frames from High Priority stations in every L-interval shall be at most L time units and usually it shall be less in order to give some residual bandwidth to Low Priority stations. Also, a restriction on the size of L is sometimes imposed e.g. if stations have telephony applications which are transmitting voice. Here, to comply with telephony standards, voice applications require a tight delay limit, from when a voice sample is generated, until when it is received at the destination.

7.1.1 Simulation results

We have simulated the new MAC using a system composed of 10 stations, of which 4 are Low Priority and 6 are High Priority. The Low Priority stations produced fixed length frames of 6000 data bits and 208 overhead bits. The frames were generated according to a Poisson process.

The High Priority stations are of several types: two stations generate a constant bit rate of 64Kb/s and represent telephone sources. One station generates a constant bit rate of 128Kb/s and represents an ISDN termination (or private branch exchange) transmitting voice packets for sessions arriving at the two voice stations from outside sources. Finally, two stations generate 1.5Mb/s and one station generates 4Mb/s respectively, and these represent CBR MPEG sources. The length of the system was set to be 5 bits long, with equal distances between the stations. The detection time was set to 10 bits, the InterFrameGap time was set to 96 bits , the slot time to 512 bits and the Backoff scheme of the Low Priority stations was the Binary Exponential Backoff, all according to the IEEE 802.3 standard [1].

We simulated the new MAC as described in Section 3. However, we also simulated a version of the MAC where High Priority stations can transmit only after detecting an idle bus and waiting at least IFG time units after the last EOC. The first version is denoted by NoGaps and the second one is denoted by WithGaps. The main deference between the two schemes is that in the WithGaps method, cycles of High Priority transmissions are not contiguous. Consequently, the Low Priority stations can initiate transmissions between the transmissions of the High Priority stations; however, they lose the collisions to the High Priority stations and their collisions counter is incremented. Thus, there is a danger that with the WithGaps scheme more packets of Low Priority stations are lost due to an excessive number of collisions. On the other hand, this scheme has a simpler implementation.
In the simulations we measured the following:

1. The access delay of the High Priority stations : the results of these measurements are given in the form of histograms. This is the only performance measurement that was computed for the High Priority stations because the access delay is bounded by the new MAC and it must be lower than the generation rate of the frames. Notice that otherwise, if stations generate frames in a higher rate than that in which they transmit, their packets queues become overloaded. If the simulation results indeed verify the analytic computation of the bound on the access delay then the access delay is also the

queuing delay of the frames and buffers shall contain two frames for the most.

2. Throughput of the Low Priority stations - we compared between the throughputs achieved in the NoGaps and WithGaps schemes.

3. We have simulated Low Priority stations with a limited buffer size of 100 frames. Thus, we also measured the percentage of frames that arrive at full buffers and are discarded.

In the graphs below we use the parameter LOAD which is the ratio between the traffic offered to the network by Low Priority stations and the maximum capacity of the network. In histograms where LOAD does not appear, the results refer to LOAD=1.

In Figure 7 we plot the histograms of the access delay of the High Priority stations in the WithGaps scheme. The variation in the access delay is due to the delay that the first frame in a cycle encounters due to transmissions of Low Priority stations. Each bin in these histograms corresponds to 0.1 ms which equals to the transmission time of 1000 bits. Each histogram contains 7 values of access delays which correspond to each 1000 bits in the Low Priority stations frames. Recall that a frame of a Low Priority station is 6208 bits long. All the access delays are below the upper bound. The differences in the results among the stations are due to the different TAG number that each of them used for its transmissions.

In Figure 8 we plot the results for the NoGaps scheme. As expected, similar results are obtained. However, a prominent difference between the schemes is that in the case of the WithGaps scheme, a high number of frames in each station experience the smallest access delay while a small number experiences the next six access delay regions. On the other hand, in the NoGaps scheme the number of frames in each bin is about the same. The reason for this difference comes from the transmission pattern of the Low Priority stations. In the WithGaps scheme, Low Priority stations collide more times than in the NoGaps scheme and their throughput is lower. This means that in the WithGaps scheme, Low Priority stations transmit fewer complete frames (i.e. successful transmissions) than in the NoGaps scheme and therefore cycles of High Priority stations are rarely delayed due to the transmissions of the Low Priority stations. Therefore, most of the frames concentrated in the bin that represents the lowest access delay which occurs when there is no deferring to transmissions of Low Priority stations. In the NoGaps scheme, the throughput of the Low Priority stations is higher. Therefore, Cycles of High Priority stations are delayed more often due to the transmissions of the Low Priority stations and the division of the frames into the various access delay bins is more balanced.

Figures 9 and 10 show respectively the throughput of the Low Priority stations and the percentage of frames that are discarded due to full buffers. These graphs show that the NoGaps scheme is superior because in this scheme the percentage of the lost frames is lower and the throughput is higher than in the WithGaps scheme. The reason for the superiority of the NoGaps scheme is because in this scheme frames from Low Priority stations collide much less than in the WithGaps scheme.

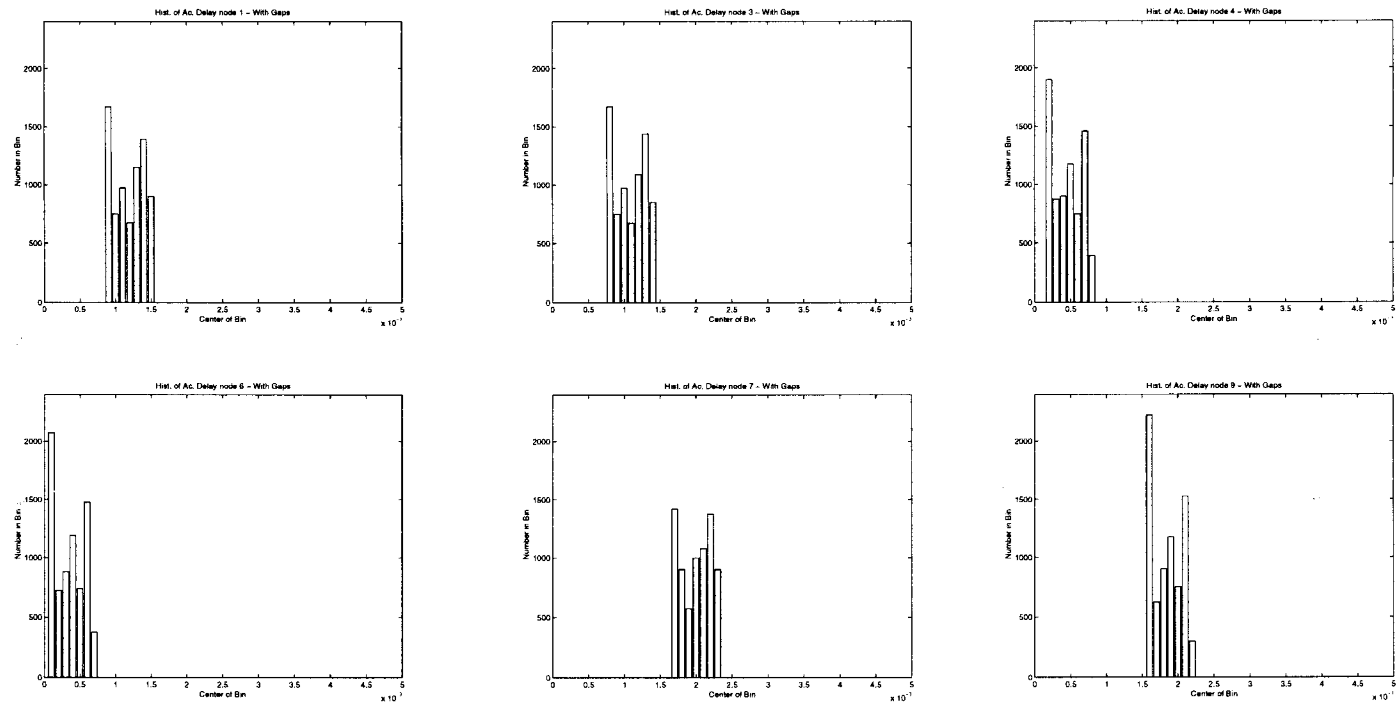

Figure 7: Access delay of High Priority stations - WithGaps scheme. On X axis, 0.001 corresponds to $1ms$

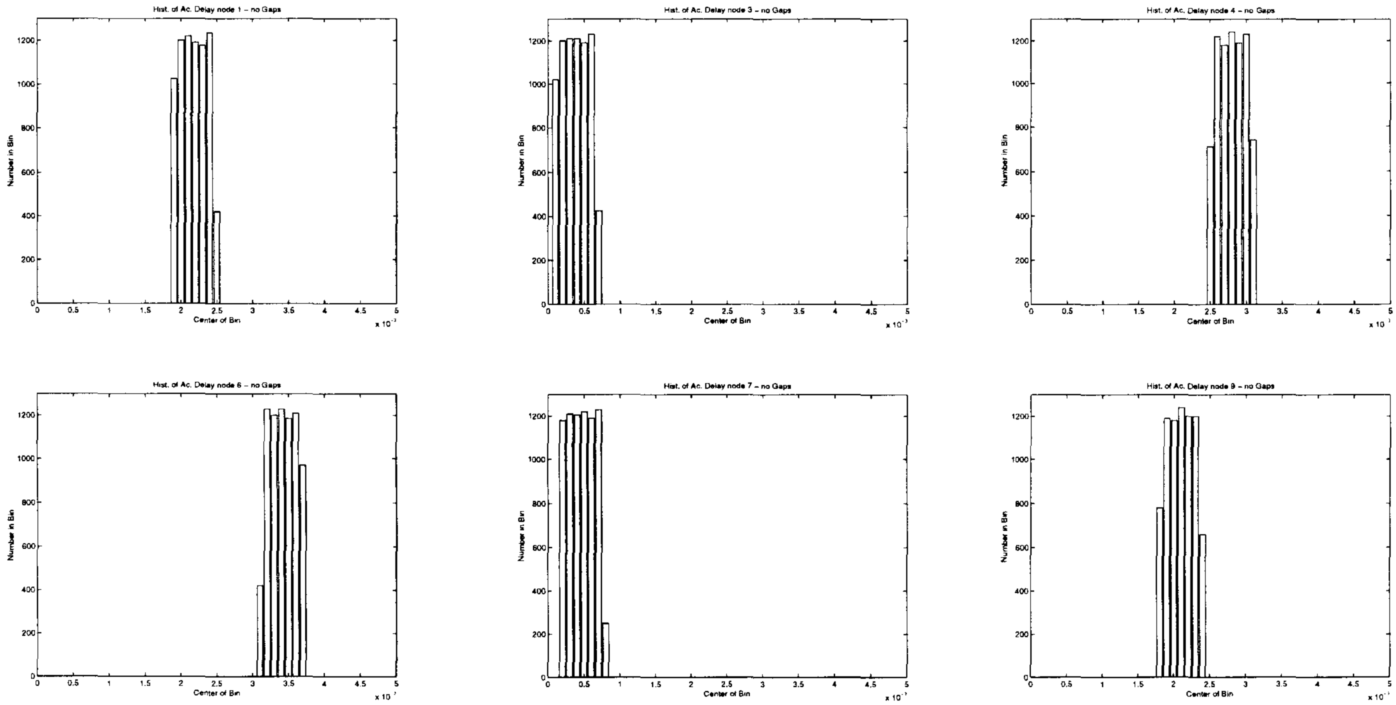

Figure 8: Access delay of High Priority stations - NoGaps scheme. On X axis, 0.001 corresponds to $1ms$

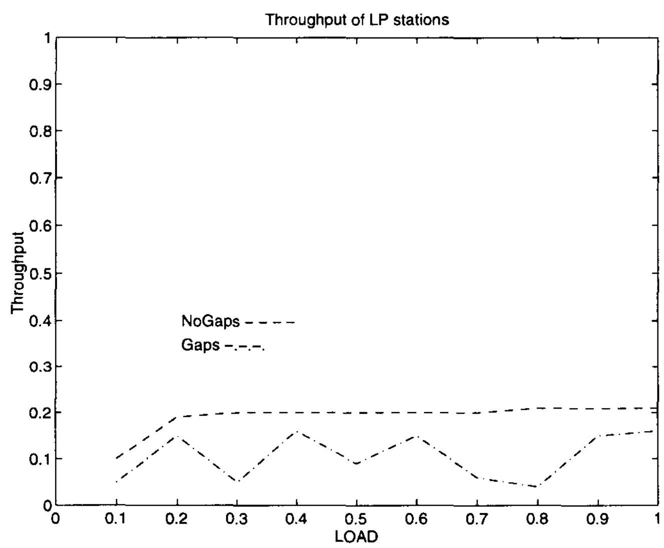

Figure 9: Throughput of Low Priority stations, with respect to their total offered load on network

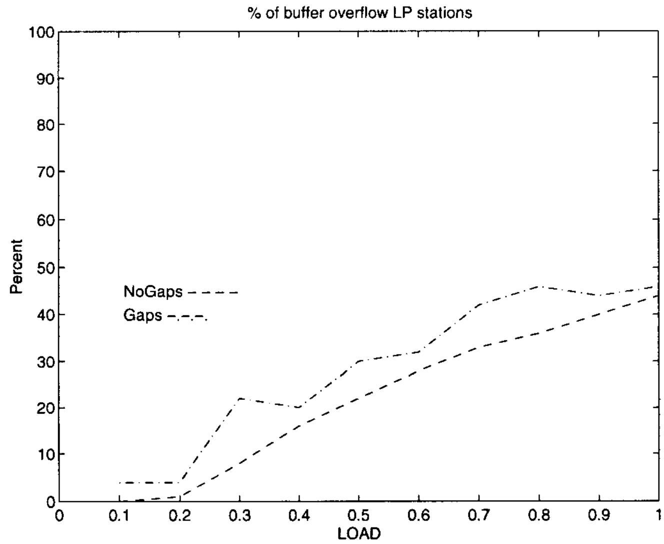

Figure 10: Percentage of lost frames in Low Priority stations, with respect to their total offered load on network

7.2 The 'packet method'

Sometimes it may be better to consider a method which takes into account the internal syntax of the bit stream which is being transmitted. For example, the bit stream of an MPEG video source is divided into Transport packets, each contains 188 bytes= 1504 bits [22], and it may sometimes be desirable to transmit an integral number of MPEG packets in a frame. We use the MPEG example further in this section.

In the packet method frames contain an integral number of MPEG Transport packets. We continue to use time intervals of L time units, as in the Stream method, in order to ensure that the High Priority stations (particularly those transmitting telephony) have a strict upper bound on the access delay. However, the Stream method involves dividing up MPEG packets, e.g. if $L = 4ms$ and we consider a $4Mb/s$ CBR source, then this source transmits in the Stream method $4 \cdot 10^6 \cdot 4 \cdot 10^-3 = 16000\text{bit} = 10.63$ MPEG Transport packets.

In the packet method we suggest that an MPEG video source will generate a frame composed of all the complete Transport packets that it has in its Application buffer when a L time units interval expires, as it is shown in Figure 11. The conceptual model is that video data is continuously fed into the Application buffer and at intervals when MPEG Transport packets become full, they are ready for transmission.

Thus, in the example above, the video source will generate frames which contain either 10 or 11 complete Transport packets (this claim is proved later). Let NFS (Nominal Frame Size) be the number (not necessarily integral) of Transport packets that are generated during an interval of L time units. Also, let MaxFS = $\lceil$ NFS $\rceil$ and MinFS = $\lfloor$ NFS $\rfloor$. Notice that a Cycle that only contains frames of length MaxFS is longer than a Cycle with the same

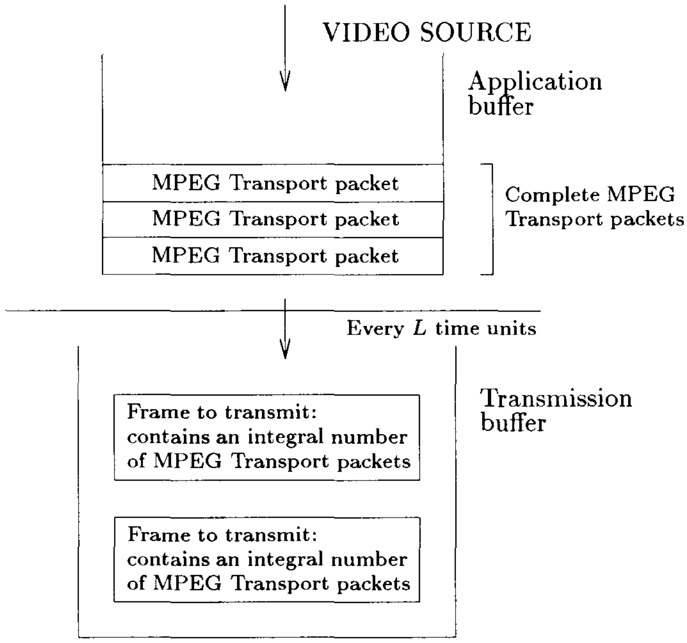

Figure 11: The transmission model in the 'packet method'

number of frames but that also contains frames of length NFS or MinFS .

Notice now that if a station always completes transmitting its frame within less than L time units from when the frame was generated, then no frames are accumulated in the Transmission buffer since it is guaranteed that every frame is transmitted before the next one is generated. In the case that a Cycle is longer than L time units, then it can happen that more than one frame will be accumulated in the Transmission buffer. This possibility requires more space in this buffer. We now prove that the requirement to transmit an integral number of packets (MPEG or other) will cause a High Priority station to complete transmitting a 'H' frame no more than L time units plus the transmission time of K packets after its generation, where K is the number of High Priority stations. This imposes a bound on the size of buffers required.

In the following discussion we denote an interval of L time units by L-interval and we assume the following:

1. The transmission time of one packet is defined to be a 1 time unit.
2. High Priority station I generates $x_i + \frac{\alpha_i}{\beta_i}$ packets during an L-interval where x_i is an integer and GCD(α_i, β_i)=1.

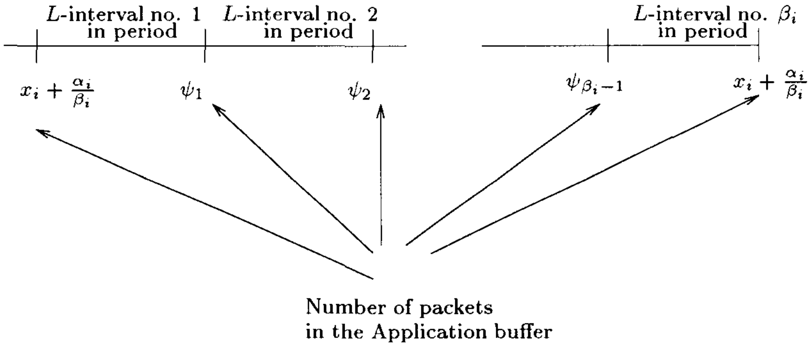

Figure 12: The period of High Priority station I

Claim 7.1: Assume that a High Priority station generates $x + \phi$ packets in an L-interval, where x is an integer and $0 \leq \phi < 1$. Then the station generates frames of size x or $x + 1$ packets only.

Proof: Assume that the station begins to generate bits for transmission at time $t = 0$ and this is the time when its first L-interval begins. We now prove the claim by induction on the times when the L-intervals terminate.

$t = L$: The station has $x + \phi$ packets and it generates a frame of x packets.

$t = 2L$: The station has $x + 2\phi < x + 2$ packets and it can generate a frame of size x or $x + 1$ packets only.

Assume correctness for the first n L-intervals. We now prove for the L-interval that ends at time $t = (n + 1) \cdot L$. At time $n \cdot L$, after generating a frame, the station has ϕ' packets left, where $0 \leq \phi' < 1$. At time $(n + 1) \cdot L$ it has $x + \phi + \phi' < x + 1 + 1 = x + 2$ packets and thus it can only generate a frame of size x or $x + 1$ packets.

□

Consider High Priority station I when it starts to transmit. After the first L-interval it has $x_i + \frac{\alpha_i}{\beta_i}$ packets in its Application buffer. Then, in the next β_i L-intervals it generates $x_i \cdot \beta_i + \alpha_i$ packets, which is an integral number. Therefore, at the end of the $(\beta_i + 1)th$ L-interval station I has again $x_i + \frac{\alpha_i}{\beta_i}$ packets in its Application buffer and so on. We denote the above β_i consecutive L-intervals by a *period.* See Figure 12.

Let ψ_v be the number of packets in the Application buffer of station I at the end of the v_th L-interval in a period but before the actions described in the next sentence. At the end of the v_th L-interval, I takes the integral number of packets from its Application buffer, which can be x_i or $x_i + 1$, and generates a frame to transmit. It then has a remainder of ε_v packets left in the Application buffer, where $0 \leq \varepsilon_v < 1$.

Claim 7.2: Consider the end of the v_th interval in a period:

(a) If I generates a frame to transmit of size x_i packets then $\varepsilon_v \geq \frac{\alpha_i}{\beta_i}$.

(b) If I generates a frame to transmit of size $x_i + 1$ packets then $\varepsilon_v < \frac{\alpha_i}{\beta_i}$.

Proof: By induction on v.
At the beginning of a period, I has $x_i + \frac{\alpha_i}{\beta_i}$ packets in its buffer. Later, at the end of the first L-interval in the period it has $x_i + 2 \cdot \frac{\alpha_i}{\beta_i}$ packets in its buffer. If I generates a frame of size x_i packets then $\varepsilon_1 = 2 \cdot \frac{\alpha_i}{\beta_i}$ and (a) holds. If I generates a frame of size $x_i + 1$ packets then $\varepsilon_1 = x_i + 2 \cdot \frac{\alpha_i}{\beta_i} - x_i - 1 = 2 \cdot \frac{\alpha_i}{\beta_i} - 1$. If $2 \cdot \frac{\alpha_i}{\beta_i} - 1 \geq \frac{\alpha_i}{\beta_i}$ then $\frac{\alpha_i}{\beta_i} \geq 1$ and this cannot hold. Therefore, $\varepsilon_1 < \frac{\alpha_i}{\beta_i}$ and (b) holds.

We now assume correctness for the first n L-intervals in a period. We prove for the end of the $(n+1)\_th$ L-interval. We separate the proof into two cases:

1. At the end of the n_th L-interval station I generates a frame of size x_i packets. Therefore, $\frac{\alpha_i}{\beta_i} \leq \varepsilon_n < 1$. At the end of the $(n+1)\_th$ L-interval station I has $x_i + \frac{\alpha_i}{\beta_i} + \varepsilon_n$ packets in its buffer. If it generates a frame of size x_i packets then (a) holds. If it generates a frame of size $x_i + 1$ packets then $\varepsilon_{n+1} = \frac{\alpha_i}{\beta_i} + \varepsilon_n - 1$. If $\frac{\alpha_i}{\beta_i} + \varepsilon_n - 1 \geq \frac{\alpha_i}{\beta_i}$ then $\varepsilon_n - 1 \geq 0$ or $\varepsilon_n \geq 1$ and this cannot hold. Therefore, $\varepsilon_{n+1} < \frac{\alpha_i}{\beta_i}$ and (b) holds.

2. At the end of the n_th L-interval station I generates a frame of size $x_i + 1$ packets. Thus, $0 \leq \varepsilon_n < \frac{\alpha_i}{\beta_i}$. At the end of the $(n+1)\_th$ L-interval station I has $x_i + \frac{\alpha_i}{\beta_i} + \varepsilon_n$ packets in its buffer. If it generates a frame of size x_i packets then (a) holds. If the frame is of size $x_i + 1$ packets then $\varepsilon_{n+1} = \frac{\alpha_i}{\beta_i} + \varepsilon_n - 1$. If $\frac{\alpha_i}{\beta_i} + \varepsilon_n - 1 \geq \frac{\alpha_i}{\beta_i}$ then $\varepsilon_n \geq 1$ and this cannot hold. Therefore, $\varepsilon_{n+1} < \frac{\alpha_i}{\beta_i}$.

□

Conclusion: By Claims 7.1 and 7.2 it is clear that for any two L-intervals v and u in a period holds $|\psi_v - \psi_u| \leq 1$.

Consider now N consecutive Cycles in which station I transmits frames.
Claim 7.3: In N consecutive Cycles station I transmits at most $N \cdot (x_i + \frac{\alpha_i}{\beta_i}) + 1$ packets.
Proof: Consider arbitrary N L-intervals as depicted in Figure 13. Assume without loss of generality that the first L-interval, out of the considered N, is the $(\beta_i - P_1 + 1)\_th$ L-interval in a period of station I and the last one is the P_2_th L-interval in a period. The number of transmitted packets in the N L-intervals is $\psi_{\beta - P_1} + P_1(x_i + \frac{\alpha_i}{\beta_i}) + z \cdot (\beta_i \cdot x_i + \alpha_i) + P_2(x_i + \frac{\alpha_i}{\beta_i}) - \psi_{P_2} = N \cdot (x_i + \frac{\alpha_i}{\beta_i}) + \psi_{\beta - P_1} - \psi_{P_2} \leq N \cdot (x_i + \frac{\alpha_i}{\beta_i}) + 1$. The last inequality is due to Claim 7.2 .

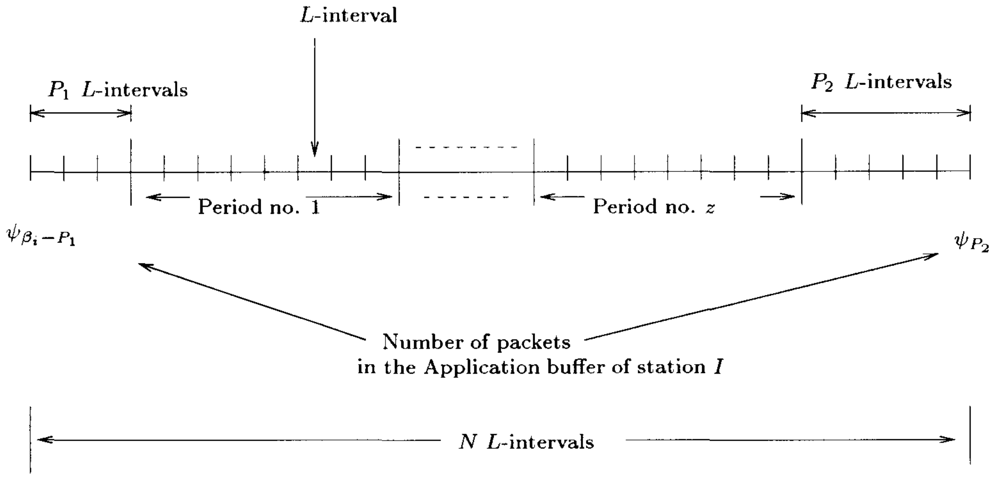

Figure 13: N consecutive L-intervals in which station I generates frames to transmit

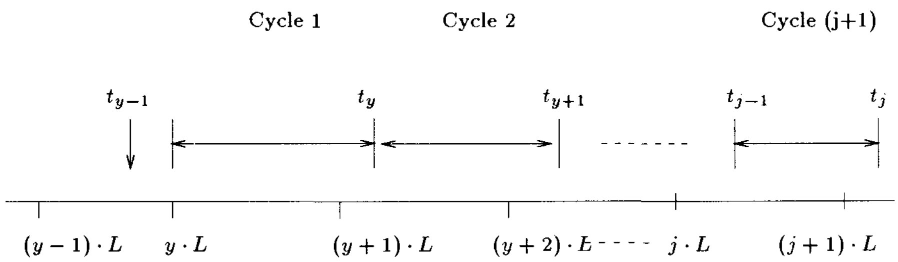

Figure 14: Time diagram in which a frame is delayed for more than $L + K$ time units

□

Consider a High Priority station P that begins to transmit at time 0. At time L it generates its first frame, at time $2L$ it generates its second frame and so on. Let t_j be the time when P completes the transmission of its j_th frame.

Theorem 7.1: For every frame j , $t_j - j \cdot L < L + K$ where K is the number of High Priority stations.

Proof: Assume by contradiction that the Theorem does not hold and let frame j be the first frame such that $t_j - j \cdot L > L + K$. Let y be the lowest index such that for any frame i , $y \leq i \leq j$, $t_i > (i+1) \cdot L$ is true. ($t_i < (i+1) \cdot L$ is not true for all i , $y \leq i \leq j$). See Figure 14.

Notice that in the interval $[y \cdot L, t_j]$ station P always has a frame to transmit. Therefore, all this time interval is devoted to transmissions of High Priority stations. Therefore, $(j + 1 -$

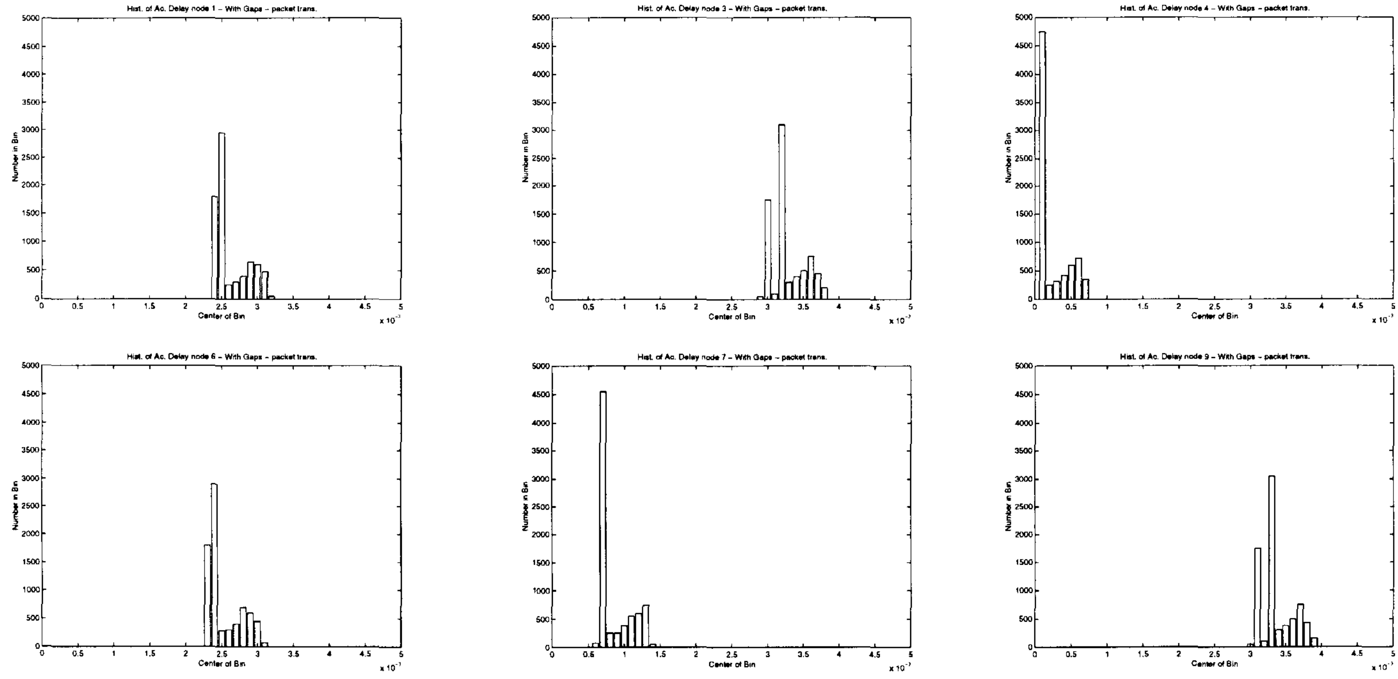

Figure 15: Access delay of High Priority stations - WithGaps scheme. On X axis, 0.001 corresponds to $1ms$

$y) \cdot L + K + \psi$ packets are transmitted during this interval, where $t_j - (j+1) \cdot L = K + \psi$.

In Claim 7.3 we showed that station I can transmit at most $(j+1-y) \cdot (x_i + \frac{\alpha_i}{\beta_i}) + 1$ packets in the considered $(j+1-y)$ cycles. Summing over all the High Priority stations we deduce that the total number of packets that the High Priority stations can transmit during this interval is: $\sum_{i=1}^{k}((j+1-y) \cdot (x_i + \frac{\alpha_i}{\beta_i}) + 1)$.

Now, $(j+1-y) \cdot L + K + \psi - \sum_{i=1}^{k}((j+1-y) \cdot (x_i + \frac{\alpha_i}{\beta_i}) + 1) \geq \sum_{i=1}^{k}(j+1-y) \cdot (x_i + \frac{\alpha_i}{\beta_i}) + K + \psi - \sum_{i=1}^{k}((j+1-y) \cdot (x_i + \frac{\alpha_i}{\beta_i}) + 1) = \sum_{i=1}^{k}(-1) + K + \psi \geq \psi > 0$ and this is not possible since more data is transmitted than the High Priority stations could generate, a contradiction. Therefore, the access delay of a frame of a High Priority station is bounded by $L + K$ time units.

□

7.2.1 Simulation results

We performed the same simulation tests as in the Stream method. In Figures 15 and 16 we plot the histograms of the access delay for the High Priority stations in the WithGaps and NoGaps schemes respectively. The explanations for the results are similar to those for Figures 7 and 8 except that the variation in the access delay is also due to the differences in the Cycle lengths that are caused by the differences in the frame lengths that the High Priority stations generate.

We omit the rest of the simulation results for the Low Priority stations because they show a similar behavior to that in the Stream method.

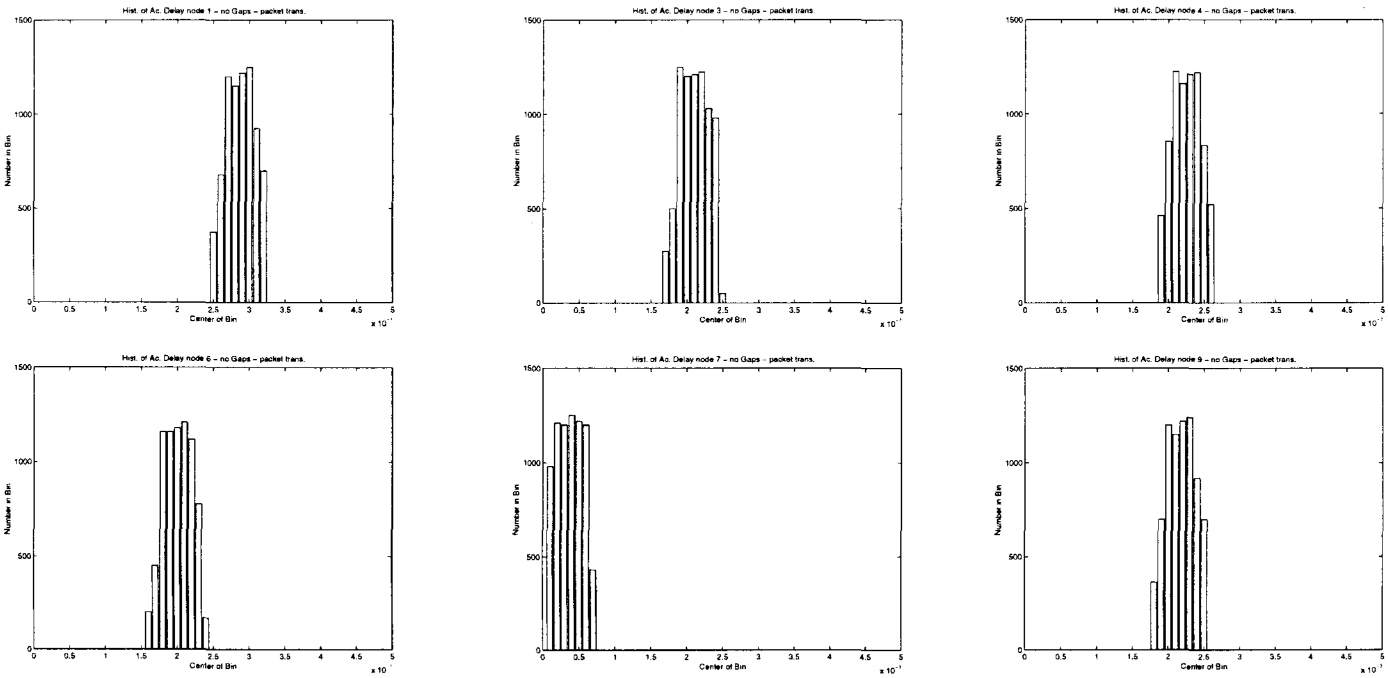

Figure 16: Access delay of High Priority stations - NoGaps scheme. On X axis, 0.001 corresponds to $1ms$

8 Summary

A novel CSMA/CD compatible MAC for real time transmissions is presented. The MAC is suitable for short buses, e.g. in the Home or small office environments and it is efficient with a small number of real time stations, say up to 10. The MAC is based on real time stations transmitting various length Jams by which collisions are resolved. The Jam lengths are determined by special numbers that real time stations obtain through a distributed procedure. The MAC does not assume any periodic nature of the real time traffic transmissions but only requires an upper bound on the time interval between consecutive transmissions.

Two schemes for the transmission of real time traffic are suggested, one ignoring and one taking account of the internal syntax of the transmitted traffic stream respectively. Correctness proofs are given for the operation of the MAC, as well as a proof of an upper bound on the size of the buffers required at the real time stations due to possible jitter in their transmissions. Finally, simulation results are given, demonstrating the performance of the MAC.

References

[1] ANSI/IEEE Std. 802.3 , "Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", Institute of Electrical and Electronics Engineers Inc. 1998 Edition.

[2] 3com Corporation, "PACE Solutions Guide," Networking Solutions Center, April 1997 , URL: *http://www.3com/nsc/100231.html.*

[3] F. Ross, D. Vaman, "IsoEthernet: An Integrated Services LAN," IEEE Communications, Vol. 34, No. 8, pp. 74-84, August 1996.

[4] T. Gonsalves, F. Tobagi, "Comparative Performance of Voice/Data Local Area Networks," IEEE Journal on Selected Areas in Communication, Vol. 7, No. 5, pp. 657-669, June 1989 .

[5] J. Zdepski, K. Joseph, D. Raychaudhuri, "Packet Transport of VBR Interframe DCT Compressed Digital Video on a CSMA/CD LAN," Proceedings of IEEE GLOBECOM, pp. 886-892, 1989 .

[6] F. Edwards, M. Schulz, "Performance of VBR Packet Video Communications on an Ethernet LAN: A Trace Driven Simulation Study," Proceedings of the 13th IEEE International Phoenix Conference on Computers and Communications, pp. 427-433, 1994

[7] I. Dalgic, W. Chien, F. Tobagi, "Evaluation of 10base-T and 100Base-T Ethernets Carrying Video, Audio and Data Traffic," Proceedings of IEEE INFOCOM, pp.1094-1102, 1994 .

[8] F. Tobagi, I. Dalgic, "Performance Evaluation of 10Base-T and 100base-T Ethernets Carrying Multimedia Traffic," IEEE Journal on Selected Areas in Communication, Vol. 14, No. 7, pp. 1436-1454, September 1996 .

[9] S. Deng, "Capture Effect in Residential Ethernet LAN," Proceedings of IEEE GLOBECOM, pp. 1678-1682, 1995 .

[10] M. Bassiouni, M. Georgiopoulos, M. Chiu, R. Guha, " Performance of Standard and Modified Network Protocols in Real-Time Application," Proceedings of the 1997 IEEE International Performance, Computing and Communications Conference, pp. 26-32, 1997

[11] S. Deng, A. Bugos, P. Hill, "Design and Evaluation of an Ethernet-Based Residential Network," IEEE Journal on Selected Areas in Communication, Vol. 14, No. 6, pp. 1138-1150, August 1996 .

[12] M. Molle, "A New Binary Logarithmic Arbitration Method for Ethernet," Technical Report CSRI-298, Computer Systems Research Institute, University of Toronto, 1994 .

[13] Enhanced Media Access Control Algorithm for IEEE 802.3 CSMA/CD, Document #80232/D1.8, Institute of Electrical and Electronic Engineers, New York, 1997 .

[14] M. Molle, K. Christensen, "The Effects of Controlling Capture on Multimedia Traffic for Shared Ethernet Systems," Journal of Telecommunication Systems, Vol. 9, No. 3-4, pp. 287-314, September 1998 .

[15] M. Molle, personal communication .

[16] I. Chlamtac, "An Ethernet Compatible Protocol for Real-Time Voice/Data Integration", Computer Networks ISDN systems 10 (1985) 81-96 .

[17] C. Szabo, "An Ethernet Compatible Protocol to Support Real Time Traffic and Multimedia Applications," Computer Networks and ISDN Systems, Col. 29, No. 3, pp. 335-342, February 1997 .

[18] F. Jia, B. Mukherjee, "The superchannel Scheme for Integrated Services on Multiple Access Broadcast Networks," Computer Networks and ISDN Systems, Vol. 27, No. 11, pp. 1523-1543, October 1995 .

[19] M. F. Maxemchuk, "A variation on CSMA/CD that yields movable TDM slots in integrated voice/data local networks", Bell Syst. Tech. J. vol. 61, pp. 1527–1550, Sept. 1982

[20] D.C. Plummer, " An Ethernet Address Resolution Protocol," RFC–826, November 1982

[21] O. Sharon, M.P. Spratt, " A CSMA/CD compatible MAC for real–time transmissions based on varying collision intervals"; URL: http://www–iri.hpl.hp.com/research/abstracts/oran–html.

[22] British Standard Implemetation of ISO/IEC 13818-1 Information Technology - Generic Coding of Moving Pictures and Associated Audio Information : Systems , document number 94/645296 IST/37: N 2122, July 1994