



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

Oran Sharon*, Michael Spratt
Network Technology Department
HP Laboratories Bristol
HPL-97-161
December, 1997

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

media access control,
MAC, real-time,
local area network,
LAN, Ethernet,
voice over LAN,
multimedia,
voice over Ethernet

This paper suggests a CSMA/CD compatible MAC protocol for real time transmissions in a Home or Small Office Local Area Network. The protocol allows high priority devices transmitting real-time traffic such as digital video and audio to communicate with each other, with guaranteed low delay and jitter, across an existing Ethernet, which is also supporting standard Ethernet devices such as PCs, printers and ISDN routers. A device implementing the new MAC can also communicate across the Ethernet with the standard devices.

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. In a somewhat similar way, in the last 10 years, the MIDI network has revolutionized the production and performance of music, and created a huge market in MIDI-compatible equipment.

Following these examples, 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, Digital VCR, CD player producing digital audio, PCs, printer and ISDN routers. Such a network could also find use in the small office environment.

Currently there are multiple standards bodies examining the home LAN, which we do not list here. The proposals tend to concentrate on carrying just video traffic and are not compatible with the Ethernet interfaces which are already being installed in large numbers in the home to connect PCs to printers and ISDN routers. However, Ethernet is very widespread, and 'is becoming the RS-232 of the 1990s.'

In this paper an Ethernet compatible MAC for real time transmission is proposed, allowing the same network to be used for piping real time traffic around the home, as the transport of computer data. The existing Ethernet cards can be connected directly to the same network as the cards implementing the new MAC.

The main characteristics of the new MAC protocol are:

1. It operates on Ethernet networks, allowing computer data nodes to utilize existing Ethernet cards and real time traffic to be transmitted on the Ethernet with 'higher priority'. The nodes implementing the new MAC can also communicate across the Ethernet to standard Ethernet nodes.
2. It enables very low jitter on real-time transmissions. This is particularly important in interfaces to home appliances, where it is desirable to have minimal buffering at the receiver to overcome network jitter, due to the need to minimize network buffering cost.
3. The protocol is simple and implementable purely in hardware.
4. The protocol is robust.
5. The protocol supports real time transmissions of arbitrary bit rates up to the maximum capacity of the bus in an efficient way. This is in contrast to previous proposals for real time transmissions in Ethernet, such as [2] and [3]. The scheme of [2] assumes that the real time traffic streams have the same rate and characteristics and thus it is based on

chains of transmissions with a fixed order among the transmitting stations. This order is kept for many transmissions since the stations have similar traffic characteristics. In our protocol stations can have arbitrary traffic rates. Also, in [3], frames contain an overhead field with length proportional to the traffic rates. For low rates such as $64kb/s$ voice stations, the overhead is low. However, for higher rates such as video transmissions, the overhead is very high and the scheme in [3] becomes inefficient. Our protocol requires some overhead but it is low and fixed, independent of the traffic rates.

The main assumptions used when designing the protocol are:

1. The home (or small office) LAN will be restricted in length - we initially assume that it will not be greater than 40 meters in length.
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 streams i.e. that need to transmit a frame typically every few 10s of ms , with a strict constraint on the time from when a frame is generated until it is transmitted successfully. 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 streams 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 can take 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. $0.2\tau s$ in a $10Mb/s$ channel, which corresponds to about $40m$ bus. δ is in the order of 10 bit time, i.e. about $1\tau s$.

Finally, although from now on we discuss the new protocol in terms of a bus protocol, it can be potentially viewed as a MAC for wireless transmissions.

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 three 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 that is used by the High Priority stations.

One way the two kinds of frames could be distinguished is by different values in the Type field. Low Priority standard stations send Ethernet frames with the Type field having a currently valid value. The other kind contains frames with a new value in the Type field, which is not contained in the current range of Ethernet Types. We denote this value by ‘H’ since frames of this type are only transmitted by High Priority stations. High Priority stations use frames of Type ‘H’ only for their real-time transmissions. Low Priority stations ignore these frames because they do not recognize the Type ‘H’. Finally, it would also be possible to distinguish the High Priority frames by having High Priority frames utilize a particular range of IEEE 48-bit MAC addresses.


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



(A)



(B)



(C)

Figure 1: (A) The standard frame structure in Ethernet (B) The frame structure used by High Priority stations (C) The IEEE 802 SNAP frame structure

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.

Finally, in Figure 1(C) we show how the frame in Figure 1(B) can be changed into the SNAP format [4]. By this structure applications in Low Priority stations can receive frames from High Priority stations [4]. We do not emphasize on this possibility any further and in the rest of the paper will only assume the structure of Figure 1(B).

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 [5] 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 and it 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 [5] 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 $\text{TAG} \cdot (2\tau + 2\delta)$ time

units which is about TAG-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 bus is not idle and a High Priority station begins to transmit a frame on the Filler of the frame transmitted by the former High Priority station which have 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 [5], K remains the only station transmitting in the system and it transmits its frame 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



(A) - An example for Cycles



(B) - An example for the time quantity T

Figure 2: Examples for Cycles and the time quantity T

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. 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. in Figure 3 assume that the stations with TAGs 1,2 and 4 are involved in the first collision. Later, if K with TAG=3 has a frame to transmit, it will collide in the second collision with the stations that have the TAG numbers 1 and 2. K will 'win' the collision and will not be allowed to transmit in the Cycle again.




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.

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




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

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 [5], it is sufficient in this case that High Priority stations will transmit the long-Jam only. We show in [5] 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 given 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.

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 of $x - 1$ 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

In this section we give 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 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 $1, 2, \dots, M$ be the set of possible TAG numbers. Then CYC is bounded by $\frac{M(M-1)}{2} \cdot (2\tau + 2\delta) + M \cdot (MFL + \tau + \delta) + IFG + P + J + 2(2\tau + 2\delta)$ time units.

Proof: Omitted due to space limits. Can be found in [5].

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

Proof: Omitted due to space limits. Can be found in [5].

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 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 permanent TAG numbers.

- (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: Omitted due to space limits. Can be found in [5].

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: Omitted due to space limits. Can be found in [5].

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 to 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 .

□

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 .

□

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 new High Priority transmitting 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 [6] (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 difference 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 6 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 7 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 8 and 9 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.




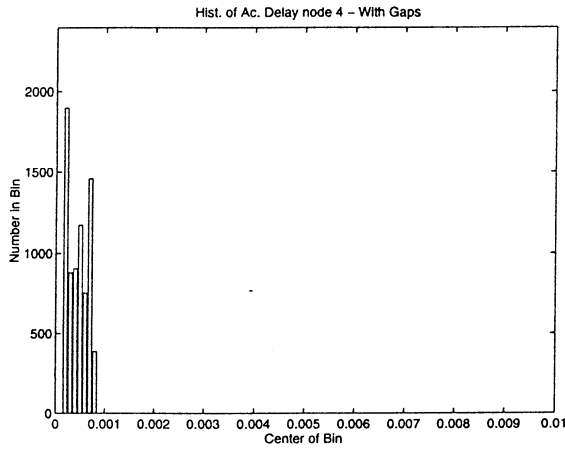
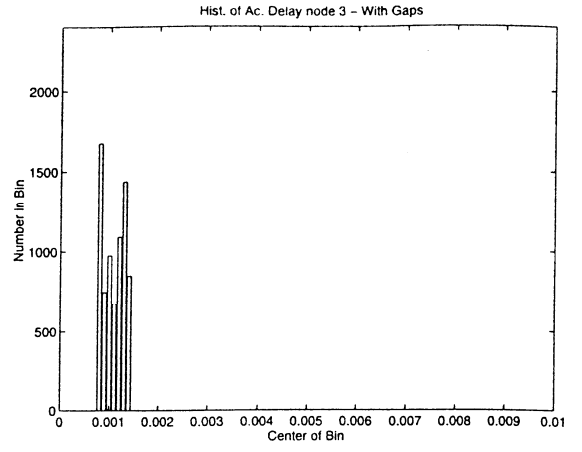
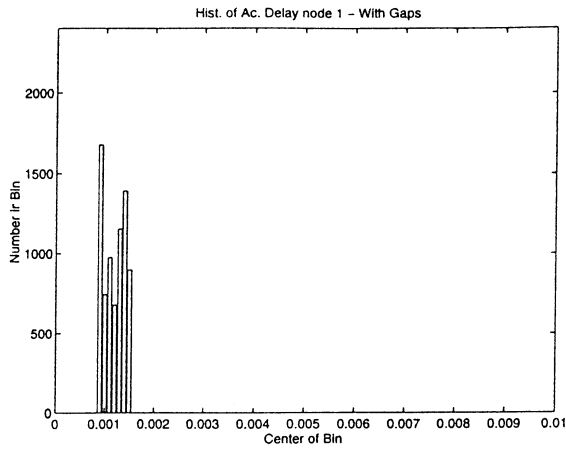


Figure 6: Access delay of High Priority stations - WithGaps scheme. On X axis, 0.001 corresponds to 1ms



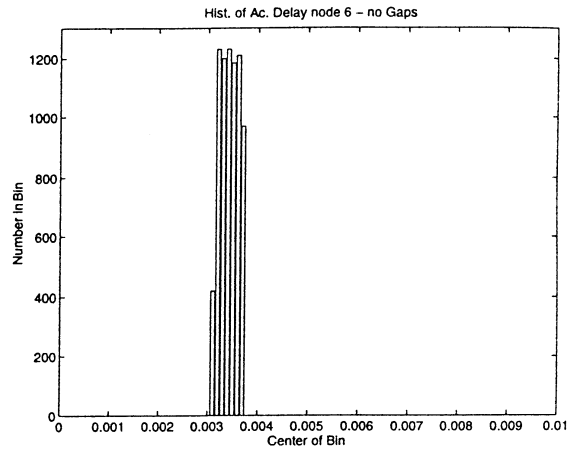
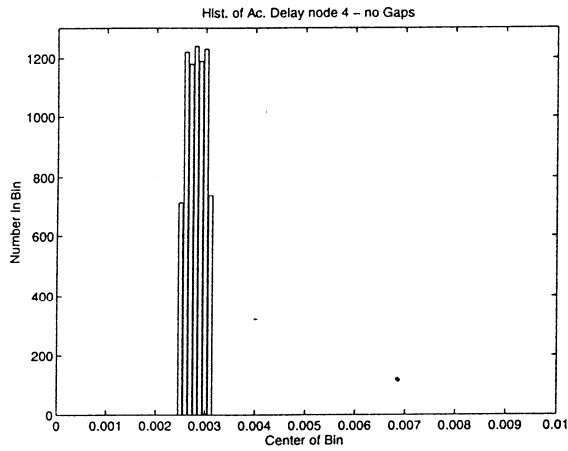
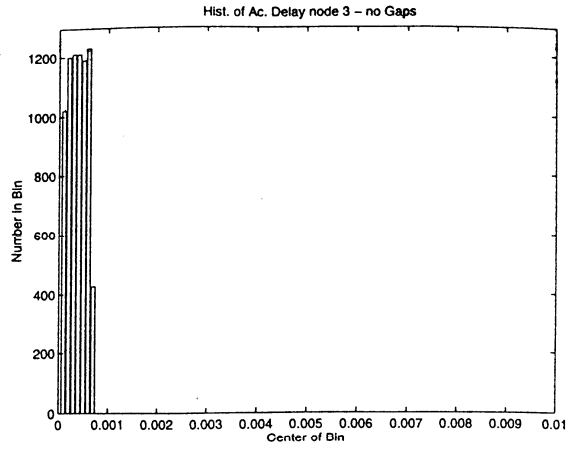



Figure 7: Access delay of High Priority stations - NoGaps scheme. On X axis, 0.001 corresponds to 1ms

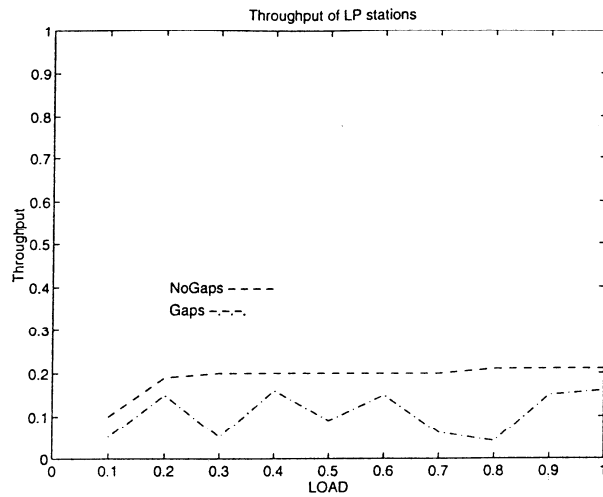


Figure 8: Throughput of Low Priority stations, with respect to their total offered load on network

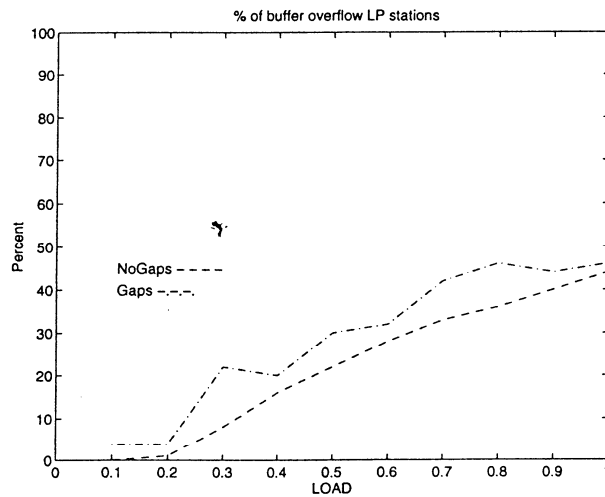


Figure 9: Percentage of lost frames in Low Priority stations, with respect to their total offered load on network

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 [6], 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 $4 \cdot 10^6 \cdot 4 \cdot 10^{-3} = 16000bit = 10.63$ 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 10. 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.

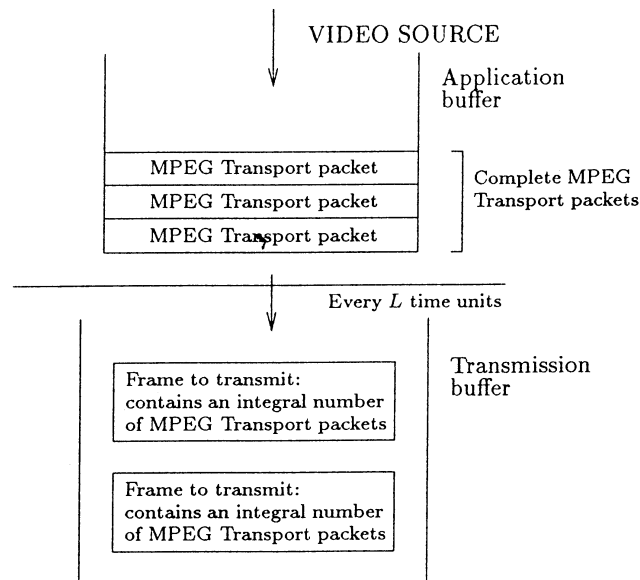


Figure 10: The transmission model in the ‘packet method’

Thus, in the example above, the video source will generate frames which contain either 10 or 11 complete Transport packets (this claim is proved in [5]).

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. In [5] we prove that the requirement to transmit an integral number of packets (MPEG or other) will cause a High Priority station to complete transmitting a 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.

7.2.1 Simulation results

We performed the same simulation tests as in the Stream method. We omit the histograms and graphs due to space limits. They can be found in [5]. Concerning the High Priority stations the results are similar to those of Figures 6 and 7 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.

Concerning the Low Priority stations the results show a similar behavior to that in the Stream method.

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 assumes a small number of real time stations. 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 protocol. 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 also given for the operation of the MAC. 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", Institute of Electrical and Electronics Engineers Inc. 1985

- [2] I. Chlamtac, "An Ethernet Compatible Protocol for Real-Time Voice/Data Integration", Computer Networks ISDN systems 10 (1985) 81-96 .
- [3] 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
- [4] Radia Perlman, "Interconnections - Bridges and Routers", Addison-Wesley, November 1995 .
- [5] O. Sharon, M. P. Spratt, "A CSMA/CD compatible MAC for real-time transmissions based on varying collision intervals"; Extended version, unpublished manuscript. Available upon request.
- [6] British Standard Implementation 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