27

A TDMA Based Access Control Scheme for an ATM Passive Optical Tree Network

Frans J.M. Panken
University of Nijmegen, Computing Science Institute
Toernooiveld 1, 6525 ED Nijmegen, The Netherlands
Tel. +31 24 3652450; E-mail: fransp@cs.kun.nl

Abstract

This paper proposes a MAC protocol for a broadband access facility, using an ATM Passive Optical Network with a tree structure. Access to a shared medium is controlled by means of a request/permit mechanism and allows a flexible way of access for services which are sensitive for (variation in) delay and for Available Bit Rate & Connection Less services. In order to guarantee a fast reaction time on changing traffic characteristics and a flexible use of available data rate, Request Access Blocks are introduced. Quantitative simulation studies show that the traffic distortion caused by the proposed MAC protocol is very small. An additional complex traffic shaper is therefore not necessary. Moreover, the protocol still provides a good performance with respect to transfer delay.

Keywords

Protocol Design and Analysis, Multiple access, Medium Access Control Protocol, TDMA, APON, Performance, Tree Network.

1 INTRODUCTION

So far, research regarding access to broadband networks mainly addresses the needs for big business customers, located in traffic intense metropolitan areas. This is natural, since these users need and can afford a dedicated fiber to have access towards B-ISDN. On the other hand, small business customers and residential users who want access to B-ISDN often neither need nor can afford the full offered data rate. In order to offer these users services from the B-ISDN, a flexible way of sharing resources is required, with the aim to concentrate traffic and to reduce costs. The ATM Passive Optical Network (APON) technology is a good candidate for such an access network, leading to the set up as depicted in Figure 1. This figure shows that the *Optical Line Terminator* (OLT) is located at the root of the tree and the fiber is splitted into several branches which on their turn contain an *Optical Network Unit* (ONU), with connections towards customer(s).

L. Mason et al. (eds.), *Broadband Communications*

calculates the required electronic delay for each ONU in order to know how long an ONU should wait before transmitting a cell. This gives each ONU the same perception of time, but is transparent for the MAC protocol to which the distance from OLT to each ONU appears to be the same. The way ranging can be realized is beyond the scope of this paper. For general issues with respect to ranging, we refer to d'Ascoli (1994).

2 THE PROPOSED MAC PROTOCOL

The main function of the MAC protocol is to avoid collisions between upstream ATM cells, originating from the different users. Moreover, a MAC protocol should aim at a good:

- *Efficiency :* the overhead introduced by the MAC protocol should be low.
- *Performance :* the delay and the delay variation introduced by the MAC protocol should be kept within certain bounds, particularly for services whose delay requirements were guaranteed during connection set up.
- *Priority :* cells originating from non-delay sensitive services (e.g. ABR and CL services) must be handled in such a way that the presence of these services does not have a large impact on the performance of services which are sensitive for (variation in) delay.
- *Fairness :* one customer should not be subject to more delay than another.
- *Robustness :* when an error occurs (especially when it is introduced by the medium), the MAC protocol should be able to recover from this error.

As described in Section 1, a maximum of 6 terminations are grouped into one ONU which saves all cells originating from these terminations in two separate buffers: the sensitive buffer and the non-sensitive buffer. An ONU consequently advertises its transmission requirements through requests (see Section 2.1) which are sent in a RAB to the controller, located at the OLT. These requests contain information about the state of the queues at the ONU. Using these requests, together with parameters agreed on during call set up, the MAC protocol distributes dynamically the available upstream data rate among the ONUs by means of the bandwidth allocation algorithm (BAA, see Section 2.3). The ONUs are then informed about the allocated transmission capacity by means of permits. When an ONU receives a permit, it is allowed to send one cell from the appropriate buffer in the upstream direction.
The request/permit mechanism allows us to assign the protocol to the class of *reservation* based protocols, with centralized control (see Kurose,Schwartz and Yemini (1984)).

2.1 Requests

In order to be able to declare the need of transmission capacity, two *Request Counters* (of 5 bits each) are assumed to be present at each ONU. A request counter counts the number of arrivals per ONU-buffer. Each time a new cell enters one of the ONU-buffers, the corresponding counter is incremented by one. The need for an empty slot for each of the buffers is passed on to the MAC controller by sending *Request Access Blocks* (RABs).
An RAB consists of a collection of requests originating from all 16 ONUs and is sent in the upstream direction every twenty slots. Figure 2 shows the upstream information structure. During two successive RABs a maximum of 5 cells can be generated by each termination, such that a total of $6 \cdot 5 = 30$ cells

may arrive at an ONU during this period. This requires an information field of 5 bits. However, these arrivals may originate from either solely the non-sensitive buffer or solely the sensitive buffer. Therefore, the contribution to an RAB of each ONU needs a field of $2 \cdot 5 = 10$ bits in an RAB in order to inform the controller about all new arrived cells. Every twenty slots the MAC controller issues a permit which forces all ONUs to form an RAB. Each ONU knows its position at the PON and consequently sends after a fixed period of time the value of its request counters (with a maximum of 31), such that an RAB is formed. At the same time, the value of the request counters is decreased with this value.

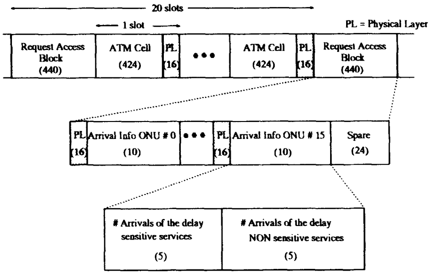

Figure 2 The upstream transmission format. All values given in bits.

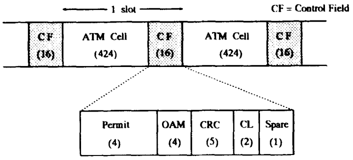

Figure 3 The downstream transmission format. All values given in bits.

Besides the 10 bits field containing information about new arrivals, the contribution of each ONU to one RAB also consists of a preamble of 16 bits. This preamble includes a gap (≈ 6 bits) and bits which are required for fine ranging, bit/byte synchronization and correction for drifts due to temperature changes. It also includes 3 bits which are used for tuning. This is necessary since due to distance equalization, the optical intensity of each ONU at the OLT may differ such that the transceiver in the OLT needs to tune for this. For the same reason such a preamble is attached to each upstream ATM cell. Time is therefore divided into slots of $424 + 16 = 440$ bits. Since there are 16 ONUs, each RAB occupies $16 \cdot (10 + 16) = 416$ bits. To make synchronization easier, an RAB simply takes the position of an ATM cell, enforcing a 24 bit spare field (see Figure 2).

2.2 Permits

The distribution of data rate is handled by the *Bandwidth Allocation Algorithm* (see next subsection), whose output consists of *permits*. A permit is sent downstream and contains the address of an ONU. As soon as an ONU receives a permit, it is allowed to send one ATM cell in upstream direction.

The permit makes part of a *Control Field* which has a length of 2 bytes. We refer to Figure 3 for the downstream transmission format. This control field travels together with a downstream ATM cell from OLT to termination. Because of the broadcast nature of the PON, the destination address of the ATM cell and the permit do not need to be the same. When at the OLT no ATM cell was scheduled for downstream transmission, the control field (containing the permit) can just take the place of an ATM cell.

Three kind of permits can be distinguished:

1. Permits for cells in the sensitive queue.

2. Permits for cells in the non-sensitive queue.
3. Permits to contribute to a Request Access Block.

We decided to let all ONUs wait for a permit to send an RAB (instead of an autonomous decision at the ONUs) to reduce the complexity at the side of an ONUs and for making synchronization easier. A *Permit Distinguish* (PD) field of two bits is added to the downstream control field in order to indicate the functionality of the permit (see Figure 3).

2.3 The Bandwidth Allocation Algorithm (BAA)

This algorithm distributes the available data rate of the shared medium among all terminations and is the *engine* of the MAC protocol. The main characteristics of the BAA of this MAC protocol are:

1. Approximation of a *VP-based UPC device.* By doing this, we concentrate any possibly miscellaneous traffic behaviour to one ONU and therefore supply a protection mechanism for the access network.
2. Use spare transmission capacity of the shared access medium to grant access to cells originating from non-sensitive services.

Upon reception of an RAB, the OLT has a clear overview of the number of cells arrived in both the sensitive buffer and in the non-sensitive buffer of all ONUs during a period of 20 slots. Generating permits for these requests can be realized very easily, since the sequence of ONUs within an RAB is fixed, such that each 10 bit information field within an RAB corresponds with a pointer to an ONU-address. However, if permits are generated according to this fixed sequence, an unfairness between the ONUs is introduced. To accomplish a fair access mechanism, it would be desirable to consider all requests in one RAB to the same collection. Consequently one request can be removed randomly from this collection and a permit for the corresponding ONU can be generated. This process may continue until no requests are left. However, since the number of requests carried within one RAB is not known in advance, this solution is difficult to implement in hardware. Furthermore, it is also desirable to have a mechanism that *protects* the access network.

The solution of fairness and protection is solved at the same time, by making use of 4 sets of k buffers, all located in the OLT. Let us first consider the protection mechanism in case of delay sensitive requests. This mechanism decides whether the permit is assigned to the first set or the second set of k buffers, by verifying if the number of requests corresponds with the parameters agreed on during connection set up. In other words, the MAC controller checks if the terminations, connected to an ONU, are requesting a data rate that agrees with their contracts. By enforcing this, we *expand* the basic functionality of the MAC protocol. Such an enlargement of the MACs functionality was also proposed in Casals, García and Blondia (1993), by making use of a *bundle spacer* (see also García,Blondia,Casals and Panken (1994)). We propose to protect the access network by making use of an algorithm that approximates the well-known *killing window* UPC mechanism, based on the sum of all connections which are shared by one ONU. We shall describe shortly how this can be implemented, by considering the parameter X, defined by

$$X = X_{old} + Nquantum \cdot (Time - Last) \cdot Alloc. \tag{1}$$

In this equation, *Nquantum* is the number of elementary quantum in the link bit rate (measure of *granularity*). *Alloc* is the number of quantum allocated to the policed connections which share an ONU. *Alloc* can easily be determined by $\lceil$ (Nquantum / nett link capacity) $\cdot \sum_i$ peak bit rate$_i\rceil$, where we sum over all VPI/VCI combinations which share an ONU. Finally, there is a parameter *Window*, representing the discarding threshold. Three situations can occur:

1. $X < 0$. When this occurs the BAA creates a permit and assigns it for the first set of k buffers. The value X_{old} is reset to zero and *Last* is set to *Time*.
2. $0 \leq X \leq Window$. When this occurs, the BAA creates a permit and assigns it for the first set of k buffers. X_{old} is updated and *Last* is set to *Time*.
3. $X > Window$. For this situation the request is considered to be *non compliant*. The BAA consequently creates a permit and assigns it for the second set of k buffers. X_{old} is not updated and *Last* is set to *Time*.

This same procedure can be carried out for requests of delay unsensitive cells, by simply substituting the first set of k buffers by the third set and the second set of k buffers by the fourth set.
After all requests are served, we start serving queue i (where i is randomly chosen out the interval [1,k]) and forward the permits from this queue to a queue Q1. When all permits in buffer i are served, we repeat this procedure with buffer $i + 1 \pmod{k}$ until all k buffers of the first set are empty. We repeat this procedure for sets 2,3 and 4, whose permits are forwarded to queues Q2, Q3 and Q4 respectively.
Clearly, this procedure makes the access scheme less unfair. It is not a very complex procedure, with respect to the required amount of additional hardware (except for the sets of buffers, of course). Both the fairness and the costs of the access mechanism are an increasing function in k, such that there is a trade-off between costs and fairness. In Section 3 we show the impact of the number of buffers k on the mean delay, as experienced by cells in each of the ONU buffers.

All permits waiting in queues Q1, Q2, Q3 and Q4 are *served*, i.e. we put them in a control field which is attached to a downstream ATM cell, in the following way. Whenever queue Q1 contains permits, this queue is served first. Only when queue Q1 is empty, queue Q2 is served. When queue Q1 and Q2 are empty, queue Q3 is served. Finally, queue Q4 is served whenever queue Q1, Q2 and queue Q3 are empty. By choosing an appropriate queue length of queues Q2 and Q4, this scheme provides a very efficient protection mechanism for the access network. The robustness scheme of the protocol (see Section 2.4) eliminates the possibility that cells keep waiting in the ONU buffer forever in case a permit is lost due to buffer overflow from queue Qi (i=1..4).

In general, X_{old}, *Last, Nquantum, Alloc* and *Window* can be defined for each ONU. However, in order to reduce the amount of memory to store all these parameters, only *Alloc, Last* and X_{old} are recommended to be defined for each ONU. The parameters *Nquantum* and *Window* are recommended to be chosen in a restricted set of 4 pair, each pair based on one value of *Nquantum* and one value of *Window*.

2.4 Robustness of the MAC protocol

As described in Subsection 2.1, the MAC protocol generates permits according to the information received from RABs. However, the loss of a request or a permit due to bit errors on the medium, leads to the situation that cells remain in the buffer of the ONU for a long time, ruining the low CDV behaviour. In order to recover from this situation, one additional *robustness counter* is necessary for each of the ONU buffers. These counters are used in the following way to recover from errors:
Whenever either a cell is sent in upstream direction or a cell enters an empty buffer, the corresponding robustness counter is set to MAX-VALUE. MAX-VALUE equals the time between two successive RABs (i.e. 20 slots) increased with the round trip delay, expressed in slots of the shared medium. When during a slot no ATM cell from a buffer is sent in the upstream direction, the corresponding robustness counter

is decremented by one. As soon as a robustness counter reaches zero, the maximum delay due to both the round trip delay and the time required to wait until an RAB was sent in upstream direction, is reached. When this occurs, the *critical period* starts. An additional time due to buffering in the OLT queues may be the cause of the fact that no permit has arrived yet. However, when during the critical period an ONU spots either a permit for a cell in a non-sensitive buffer (addressed to any ONU) or an empty permit (i.e. a permit with no ONU address), an error must have occurred. An error for non-sensitive queue is ascertained as soon as the ONU observes an empty permit while its robustness counter is zero and the non-sensitive queue is not empty.

As soon as an error is detected, both the robustness counter of this queue is rest to MAX-VALUE and the appropriate request counter is increased with one. Consequently, after a while, another request will be issued by this ONU.

This scheme allows the MAC protocol to recover from errors in a fast and efficient manner. Since we assumed a 5-bits request counter and a maximum of 30 cells at one ONU can arrive between two RABs, the last bit of this counter guarantees the possibility to correct for errors under each traffic condition.

Since an appropriate distance between ONU and OLT is $\approx$ 10 km, an 8 bits robustness counter for both cell buffers is sufficient. Finally, we remark that the round trip delay is already available, since the ranging procedure requires it.

2.5 Efficiency of the MAC protocol

In the PON we divided time into slots of 440 bits. The assumed overhead of 16 bits per ATM cell which is inherited from the physical layer, corresponds very well with the standardized overhead to transport ATM cells such as SOH and POH in case of SDH and F1, F2 and F3 flows in case of cell based transmission. So, by using an APON slot of 440 bits, the transmission rate of 155 Mbit/s can be used while still preserving 149.76 Mbit/s for ATM cells.

However, the introduction of RABs in the upstream direction requires an additional reservation of 440 bits, every 20 slots. This reduces the actual capacity for ATM cells by $100\%/20 = 5$ %, to about 142.27 Mbit/s on a standardized 149.76 Mbit/s interface. Since it would be very unlikely that a connection with such a large capacity would be established, this is not much of a problem. However, the CAC mechanism has to take this reduced capacity into account (i.e. change the maximum data rate from 149.76 Mbit/s to 142.27 Mbit/s).

3 PERFORMANCE EVALUATION

Since the MAC protocol controls the information flow at the entrance of a B-ISDN network, it has a large impact on the overall performance of the system. In particular, the profile of the traffic as it enters the network is highly influenced by the MAC protocol. Therefore, a performance evaluation of the proposed MAC protocol is carried out in this section.

In order to study the impact of the presence of the proposed MAC protocol, we show the complementary distribution function (i.e. $\Pr\{X > x\}$) of two important performance measures: the transfer delay and the CDV.

The *Transfer Delay* of the access network is defined as the time difference between the sending of a cell at the T_b-Interface and the receiving of this cell at the U_{1b}-Interface.

An accurate performance characterization of CDV is a network performance parameter, known as *1-Point-*

CDV. The 1-Point-CDV describes the variation of the arrival times pattern with respect to the negotiated peak cell rate. It is measured by observing successive upstream cell arrivals at the U_{1B}-Interface and only considers cell clumping, i.e. the effect of cell interarrival distances which are shorter than the reciprocal of the peak cell rate. The characterization of CDV by means of 1-Point-CDV was given in CCITT (1992) and is recommended for CDV assessment by ITU-T.

All presented results in this section are obtained from simulation. In order to make a comparison between certain aspects easier, several related curves are shown in one figure. To prevent that figures become disordered, confidence intervals are omitted. We remark that the 95 % confidence intervals did not show a deviation larger than 10 %, for values larger than 10^{-3}.

To obtain the results presented in Figures 4- 8, we simulated an access network where at each ONU three delay sensitive CBR cell stream of 1, 4 and 10 Mbit/s respectively were offered. In addition, a total of 3 ABR cell streams with the bit rates 1, 4 and 10 Mbit/s were offered at ONU 5,12 and 15 respectively. This accomplishes a load of 0.4. In order to study the impact of the presence of the MAC protocol under a load of 0.8, three delay sensitive CBR cell streams with a bit rate of 5 Mbit/s were added to each ONU. From the two curves in Figures 5 - 12 which represent the performance of connections sharing the same ONU, the one with the lowest quantiles corresponds with a load of 0.4. For each load condition the 1 Mbit/s and the 10 Mbit/s connections were *tagged* during successive simulations and their performance was studied.

Fairness

In order to try to make the delay for each connection of the access network the same, 4 sets of k buffers were proposed in Subsection 2.3. As explained in this subsection, the fairness of the MAC protocol increases with k.

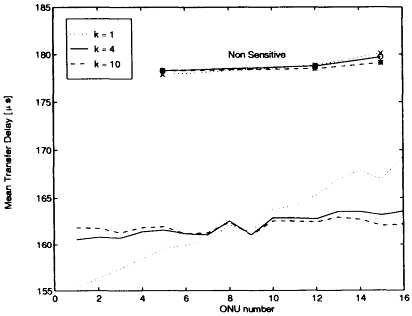

Figure 4 Mean transfer delay as function of the ONU number, for different values of k.

Figure 4 shows the mean delay as function of the ONU number (consider only integer values) under a load of 0.8. We observe that $k = 4$ buffers already shows a rather horizontal curve, corresponding with a fair MAC protocol. For this reason, $k = 4$ buffers were chosen to obtain the rest of the presented results.

Transfer Delay and CDV

From Figures 5 - 8 we conclude that, compared to either centralized MAC protocols or basic equipment in an ATM network such as switches, the proposed access scheme guarantees low CDV and both bounded and acceptable transfer delays for CBR traffic streams at all ONUs. Furthermore, we observe that the performance seen by a connection does not depend very much on its bit rate, as occurs e.g. for some of the MAC protocols analyzed in Chapter 3.6 of Killat (1996). The load of the access network appears to be an important parameter for those connections which are **not** sensitive for (variations in) delay. Figures 9 and 10 show that the above statements are still valid when we increase the bit rate of a tagged connection above the bit rate which corresponds with the situation that interarrival times between successive cells are smaller than the time between two successive RABs.

In Figures 11 and 12 we study the performance in case of VBR input traffic. For VBR traffic we used an *on-off* source, where cells are generated according to an underlying Markov chain which alternates between two states: "on" and "off". When the Markov chain is in the on-state, cells are generated according to a CBR pattern, whereas during the off-state no cells are emitted. Both the number of cells generated during the on-period and the sojourn time in the off-state are chosen to be geometrically distributed. In this configuration, temporary overload situations may occur, causing longer tails.
We observe that the priority scheme for sensitive services becomes effective when the load of the access network is high or when the input traffic is more variable.

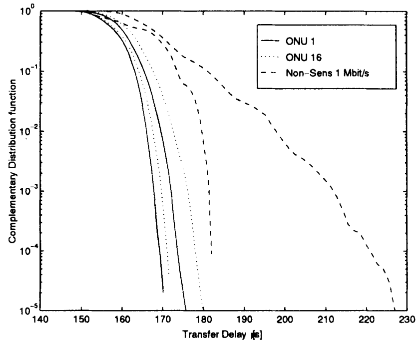

Figure 5 Complementary distribution of the transfer delay of a 1 Mbit/s CBR connection, under the condition that the access network is loaded up to 0.4 and 0.8 respectively.

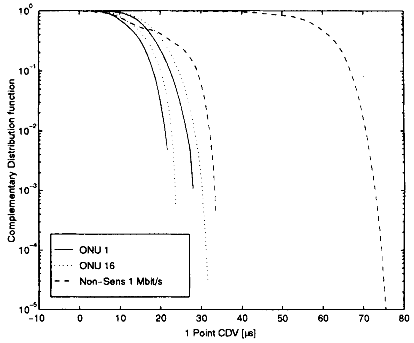

Figure 6 Complementary distribution of the 1 point CDV of a 1 Mbit/s CBR connection under the condition that the access network is loaded up to 0.4 and 0.8 respectively.

4 CONCLUSIONS

In this paper we propose a MAC protocol for an ATM Passive optical tree network. We show how terminations must be grouped into one Optical Network Unit (ONU) and how the ONUs can declare their needs for transmission capacity to a centrally situated controller. Besides this, a robustness scheme and the bandwidth allocation algorithm of the protocol are explained in detail. We showed how the bandwidth allocation algorithm ensures a protection mechanism for the access network, expanding the functionality of the access control scheme.
The MAC protocol can be used for both fiber to the home and fiber to the kerb and allows ABR services in such a way that it can not harm the performance of delay sensitive services. In order to make a make

a statement with respect to the performance of the protocol, simulation results are presented. These results show that the proposed protocol can supply a relative fair access scheme for all users and ensures a bounded delay and very little traffic distortion (CDV), especially for delay sensitive services.

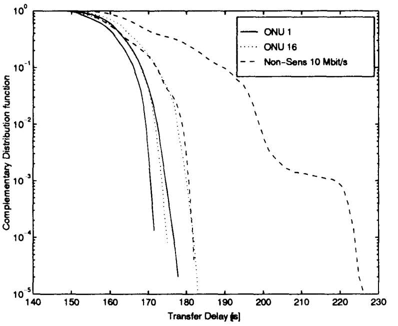

Figure 7 Complementary distribution of the transfer delay of a 10 Mbit/s CBR connection under the condition that the access network is loaded up to 0.4 and 0.8 respectively.

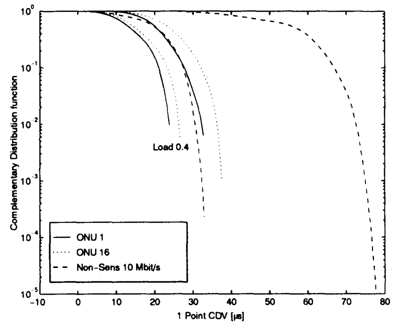

Figure 8 Complementary distribution of the 1 point CDV of a 10 Mbit/s connection under the condition that the access network is loaded up to 0.4 and 0.8 respectively.

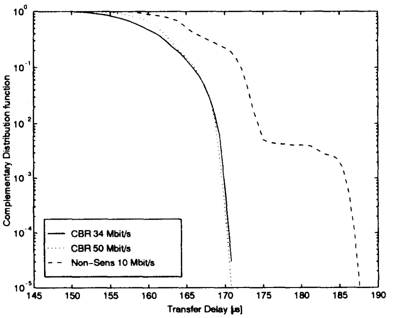

Figure 9 Complementary distribution of the transfer delay when the access network is loaded with CBR traffic of bit rates 34 Mbit/s and 50 Mbit/s. There is one delay sensitive connection and the total load equals 0.8.

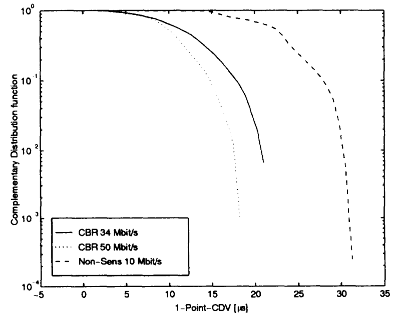

Figure 10 Complementary distribution of the 1 Point CDV when the access network is loaded with CBR traffic of bit rates 34 Mbit/s and 50 Mbit/s. There is one delay sensitive connection and the total load equals 0.8.

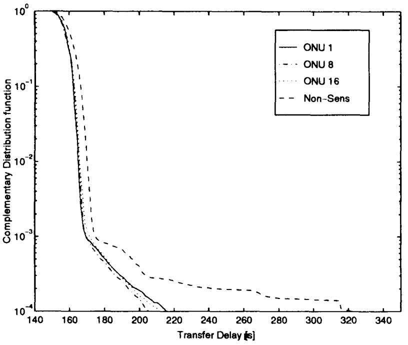

Figure 11 Complementary distribution of the transfer delay when the access network is symmetrically loaded with a total of 45 VBR sources with a peak bit rate of 50 Mbit/s, a mean bit rate of 5 Mbit/s and a burst length of 100 cells.

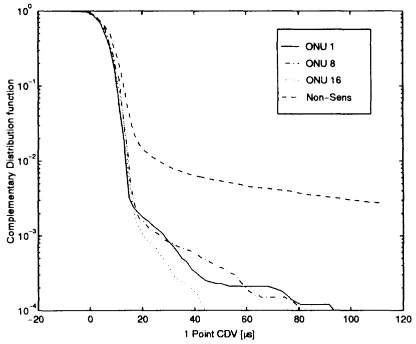

Figure 12 Complementary distribution of the 1 Point CDV when the access network is symmetrically loaded with a total of 45 VBR sources with a peak bit rate of 50 Mbit/s, a mean bit rate of 5 Mbit/s and a burst length of 100 cells.

REFERENCES

Angelopoulos, J.D. and Venieris, I.S. (1994) A Distributed FIFO Spacer/Multiplexer for Access to Tree APONs. *Proceedings of SUPERCOMM/ICC Conference* New Orleans, Louisiana, May 2.

Angelopoulos, J.D. and Venieris, I.S. and Stassinopoulos, G.I. (1993) A TDMA Based Access Control Scheme for APON's. *IEEE Journal Lightwave Technology*, **vol. 11, no. 5/6**, 1095-1103.

Blondia, C. and Verri, L. and Toniatti, T. and Casals, O. and García, J. and Angelopoulos, J.D. and Venieris, I. (1994). Performance of Shared Medium Access Protocols for ATM Traffic Concentration. *European Transactions on Telecommunication*, Special Issue on Teletraffic Research for B-ISDN in the RACE program **Vol. 5, No. 2**, 219-226.

Casals, O and García, J. and Blondia, C. (1993) A Medium Access Control Protocol for an ATM Access Network. *Proceedings of the Fifth International Conference on Data Communication Systems and their Performance, High Speed Networks.*, Eds. H.Perros and Y.Viniotis North Carolina, USA.

CCITT (1992) Draft Recommendation I.35B. B-ISDN ATM Layer Cell transfer Performance. Geneva, June 1992.

d'Ascoli, L. et al. (1994) Out-of-Band Ranging Method for ATM over PON Access Systems. *ECOC, Firenze, Italy.*

García, J. and Blondia, C. and Casals, O and Panken, F. (1994) The Bundle-Spacer: A Cost Effective Alternative for Traffic Shaping in ATM Networks. *Proceedings of IFIP Conference on Local And Metropolitan Communication Systems, Kyoto.*, 177-196.

Killat, U (1996) Access to B-ISDN via PONs. To be published by Wiley.

Kurose, J.F. and Schwartz, M. and Yemini, Y. (1984) Multiple-access protocols and time-constrained communication. *COMPUT. SERV.*, **16**, 1, 43-70.

Van der Plas, G. (1993) APON: An ATM-based FITL system. *EFOC & N'93*, 91-95.