

A MAC PROTOCOL FOR AN ATM PON SUPPORTING EXPLICIT RATE CONGESTION CONTROL FOR ABR TRAFFIC

F. Panken^a, C. Blondia^b, O. Casals^c, J. Garcia^c

^aUniversity of Nijmegen; Computing Science Institute; Nijmegen, The Netherlands;
Tel: +31243652450; Fax: +31243553450; e-mail: fransp@cs.kun.nl

^bUniversity of Antwerp, Dept. Mathematics and Computer Science, Antwerp, Belgium;
Tel: +3238202404; Fax: +3238202421; e-mail: blondia@uia.ua.ac.be

^cPolytechnic University of Catalonia; Computer Architecture Dept., Barcelona, Spain;
Tel; +3434016985; Fax: +3434017055; e-mail: olga@ac.upc.es

Abstract

This paper presents and evaluates the performance of a MAC protocol for a broadband access network using an ATM Passive Optical Network (PON) supporting different service categories. The protocol is an enhancement of an earlier version proposed by the authors and implemented in the demonstrator of a project in the European RACE programme. A request/permit mechanism controls the access to the shared medium. The bandwidth allocation algorithm gives priority to the CBR/VBR traffic and ensures a minimum cell rate for the ABR traffic. The unused bandwidth is allocated to UBR traffic.

Subject Area : ATM Networks

1. Introduction

A major problem when designing an ATM PON (APON) is the definition of the mechanism that controls the access to the shared medium. Such a mechanism, called Medium Access Control (MAC) protocol, should primarily avoid collisions of cells originating from different users. In addition, the MAC protocol should fulfil a number of requirements: it should introduce low overhead, make efficient use of the resources and guarantee the requested QoS for all service categories.

In this paper we propose a MAC protocol which is an enhanced version of a protocol

proposed by the authors and developed and implemented in a demonstrator within the Broadband Access Facilities project of the European RACE program (project R2024), now being used in several national broadband networks. These enhancements allow the customers to set up connections belonging to one of the service categories CBR/VBR, ABR or UBR. The stringent requirements of CBR/VBR traffic impose priority to access the shared medium. The ABR service only needs a guarantee of a Minimum Cell Rate (MCR), together with a low Cell Loss Ratio (CLR), while the remaining bandwidth capacity should be used for UBR traffic. In [PAN97], a MAC protocol supporting these service categories has been proposed, where the congestion control mechanism used for ABR is of binary type. Now we improve this scheme by proposing a congestion control mechanism based on the explicit rate approach. As in [CAS93] and [PAN97], the enhanced MAC protocol has a centralised architecture, it is cell based and it uses a request/permit mechanism to exchange information between OLT and NT1s.. The MAC protocol allocates the bandwidth to the NT1s and among the service classes according to a well-defined algorithm. This bandwidth allocation algorithm gives the highest priority to the CBR/VBR traffic, but ensuring a

minimum cell rate for ABR traffic. The remaining bandwidth is used for UBR traffic.

2. System Configuration

Figure 1 shows the reference configuration of the system implemented in the RACE project "BAF". The system uses a passive optical tree with splitting factor 32 and the maximum fiber length between the OLT and the ONUs is typically 10 km. The maximum bit rate the system is able to support is 622 Mbit/s, both in the upstream direction (from the customer premises network to the core network) and in the downstream direction (from the core network to the customer premises network). The maximum number of Tb interfaces that can be connected to an OLT is 81 and their net bit rate is 149.76 Mbit/s.

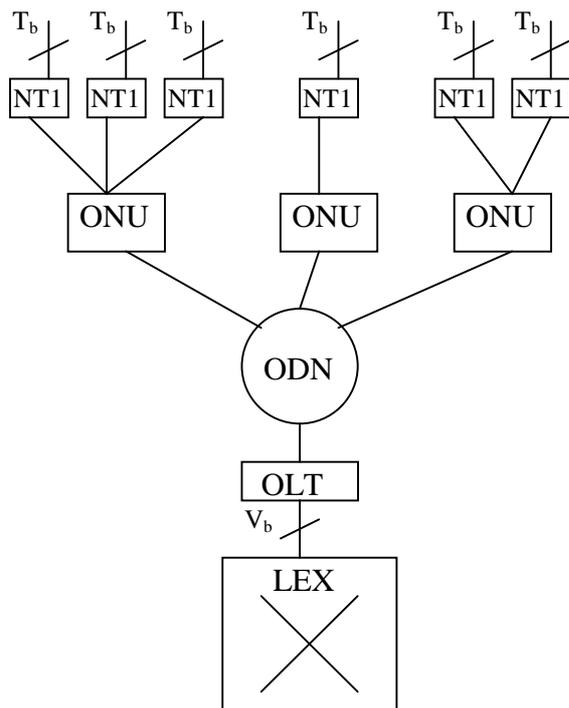


Figure 1 : System Configuration

3. The MAC Protocol

The main function of a MAC protocol is to resolve contention experienced when more than one ONU wants to use the shared medium. In view of the topology of the PON, the MAC

protocol has centralised control located at the OLT. It is defined by:

- the way the control is informed about the bandwidth requirements of the NT1s : *requests*
- the way the NT1s are informed about the bandwidth they obtain : *permits*
- the way the bandwidth is allocated among the NT1s : *bandwidth allocation algorithm*.

3.1. Requests

An NT1 declares its bandwidth needs by means of requests. Figure 2 shows the structure of the upstream information flow. Each NT1 can send *requests* in two different ways:

1. Requests coupled with upstream cells.

When an NT1 is allowed to transmit a cell, a tag is added to this upstream cell. This tag contains the number of ABR (QIABR) and CBR/VBR (QICBR/VBR) cells which are waiting in the buffers (see Figure 2). If only this type of request would be used, the MAC protocol would react too slow to changing traffic situations.

2. Requests Blocks.

A *Request Block* (RB) contains the requests of a number of consecutive NT1s. An RB consists of 9 request fields

Remark that no information about the UBR buffer length is sent as the UBR traffic will only use the spare bandwidth (see next section). In order to have a robust protocol which is able to recover from loss or corruption of requests, the request has to contain the queue length of the NT1. To extract from this information the number of new arrivals in the queue since the previous request, a scheme based on mirror counters has been proposed. For more details we refer to [KIL96].

3.2. Downstream Information.

The OLT issues two kinds of permits (Figure 3).

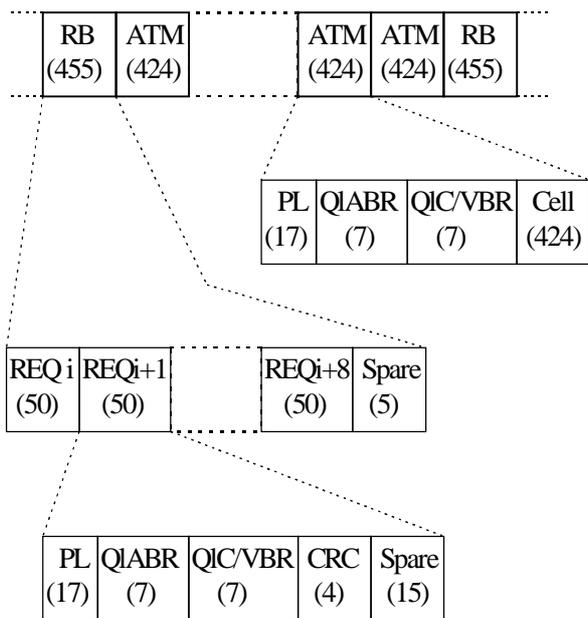


Figure 2 : Requests

1. Permits for ATM cells.

When the controller decides that an NT1 is allowed to send a cell of a certain service category, the OLT generates a *permit for an ATM cell*. This permit consists of an address of the NT1 and the CL-field, which informs the NT1 from which buffer the cell should be selected (CBR/VBR, ABR or UBR). The permit is transmitted in the next downstream time slot together with an ATM cell or an empty cell. Downstream traffic is broadcasted to all NT1s and the NT1 that recognizes its address in the permit is allowed to send a cell in the next upstream slot.

2. Permits for a Request Block.

The OLT generates a *permit for a Request Block* periodically. This permit contains the address of the first NT1 in a series of N=9 consecutive NT1s that can use the RB to send a request. When the NT1s recognise a permit for an RB, they compare their address with the one in the permit. If the address of an NT1 is one of the 9 consecutive NT1 addresses, then this NT1 is allowed to send a request in the next

upstream slot. The 9 requests which are merged together form an RB.

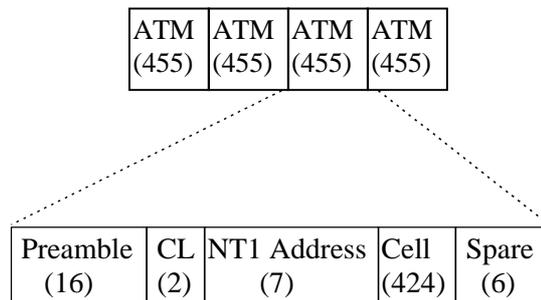


Figure 3 : Permits

3.3. The Bandwidth Allocation Algorithm and its Implementation

The bandwidth has to be allocated on two different levels: among the different NT1s connected to the APON and among the different service categories supported by the APON (CBR/VBR, ABR and UBR).

CBR/VBR Traffic

The OLT computes the number of new CBR/VBR cell arrivals in the NT1s (based on the information in the requests) and generates the necessary number of permits. These permits are put in a global permit FIFO queue, in which permits for cells in the NT1 buffers are stored. Due to the presence of periodic request blocks, no spacing of permits is required (compare with [CAS93]).

ABR Traffic

The Minimum Cell Rate (denoted MCR) of an NT1 is calculated as the sum of the MCRs of each of the ABR connections carried by this NT1. The OLT computes the number of new ABR-cell arrivals (deduced from the requests) and assigns the necessary ABR-permits to the NT1, according to the following rules:

- When the number of slots since the last permit for an ABR cell of that NT1 was put in the global FIFO equals $1/MCR$, and there are ABR cells waiting for transmission, then a permit for an ABR cell is generated and put in the global FIFO.

- When the Global FIFO queue is empty, then an ABR permit can be generated for an NT1 which has ABR cells waiting for transmission (even at a higher rate than the MCR). This NT1 is determined on basis of a cyclic service order.

UBR Traffic

When the NT1s have no CBR/VBR cells nor ABR cells to transmit, the remaining bandwidth can be used for UBR traffic. Permits for UBR cells are distributed in a cyclic order.

Now we show how this distribution algorithm can be implemented in an efficient way. Consider N NT1s, $NT1(1), \dots, NT1(N)$ each provided with three buffers, one for CBR/VBR cells, one for ABR cells and one for UBR cells. Let $m(i)$ be the inverse of the sum of the minimum cell rates of each of the ABR connections carried by $NT1(i)$. The following features are needed.

1. The *global permit FIFO queue* contains permits for cells in the NT1 buffers.
2. The OLT maintains *two global counters C1 and C2*. The counter C1 is used to determine which ABR NT1 queue will receive a permit at the first slot during which the Global FIFO queue is empty. The counter C2 is used to determine which UBR NT1 queue will receive a permit at the first slot during which no CBR/VBR nor ABR traffic is present.
3. Two counters for ABR traffic per NT1 are maintained, a *Countdown Counter* and a *Request Counter*. For connection i , $i=1, \dots, N$, these counters operate as follows:

The countdown counter $CNTD_CNTR(i)$ is given the value $m(i)$, at the moment a permit for $NT1(i)$ is put in the global FIFO queue. Its value is decreased by 1, at each time slot. Counting down stops as soon as a permit for an ABR cell of $NT1(i)$ is put in the global FIFO or if $CNTD_CNTR(i)=0$. The request counter $REQ_CNTR(i)$ is increased each time a request comes in by the number of arrivals since the last request. It is decreased by 1 whenever a permit for an ABR cell for $NT1(i)$ is put in the global permit FIFO queue.

4. ABR Congestion Control

As an integral part of the B-ISDN network, the APON has to advise periodically ABR sources about the rate at which they should be transmitting. The congestion control mechanism proposed in this paper is based on the ERICA switch algorithm [JAI96].

Hardware support

We add three counters in the OLT. A counter $CNTR1$ which keeps track of the time since the last computation of the *overload factor* was made (i.e. the duration of the *observation period* of length T). Two counters $CNTR2$, resp. $CNTR3$, keeping track of the total number of CBR/VBR, resp. ABR, request arrivals during an observation period.

ER calculation algorithm

The counters defined above are used to calculate the new rate which is written in the ER field of the backward RM cells in the following way.

When the first counter reaches 0 (i.e. when $CNTR1=0$), then the overload factor is computed:

ABR Input Rate := $CNTR3/T$;

CBR/VBR Input Rate := $CNTR2/T$

Target ABR Capacity :=

$\text{Max}\{0, \text{Target Util.} * \text{APON Rate} - \text{CBR/VBR Input Rate}\}$;

Overload factor $O := \text{ABR Input Rate} / \text{Target ABR Capacity}$.

Each time the first counter ($CNTR1$) reaches 0 the value of O is updated and stored in a memory. When an RM cell passes through the OLT, then the ER field is updated using the last computed value of O as follows. In order to achieve fairness, we need to compute the fair share for each NT1:

FairShare := $\text{ABR Capacity} / \text{Number of active NT1s}$;

ABR Capacity := $\text{Max}\{0, \text{APON Rate} - \text{CBR/VBR Input Rate}\}$;

Number of active NT1s := Number of NT1s whose ABR request counter is > 0 .

We assume here that we have one ABR VP per NT1. In case we want to have several ABR

VPs per NT1, the FairShare should be calculated as:

FairShare := ABR Capacity/Total Number of VPs of active NT1s;

The NT1Share is obtained using the Current Cell Rate (CCR) field of the RM cell and the last computed overload factor O: NT1Share := CCR/O. The ER value to put in the RM cell is then calculated as follows:

ER Calculated := Min(ABR Capacity, Max(FairShare, NT1Share))

ER in RM Cell := Min(ER in RM cell, ER Calculated).

	S1	S2	S3
CBR			
-Peak (Mbit/s)	10	34	50
-Mean (Mbit/s)	10	34	5
-Number	23	6	20
ABR			
-Peak (Mbit/s)	20	68	34
-Mean (Mbit/s)	10	34	0
-off/on period	0	0	10
-Number	23	7	20
Line Type	_____	- - - - -

Table 1 : Scenarios

5. Performance Evaluation

5.1. The System Model and Traffic Scenarios

The impact of the MAC protocol on the performance is evaluated quantitatively by studying both the transfer delay and the 1-Point CDV experienced by the CBR/VBR traffic class and by studying the distribution of the queue length of the ABR buffers by simulation. The scenarios chosen to study the performance of the MAC protocol are shown in Table 1. All results presented assume a 90% target utilization and an observation period T of 180 slots.

5.2. Delay and 1-point CDV of CBR/VBR Traffic

Figures 4 and 5 depict the complementary distribution of respectively the transfer delay and the 1-Point CDV experienced by connections of the CBR/VBR service class. From Figure 4 we conclude that constant high bit rate connections, as in scenario S2, clearly benefit from the fact that requests can be coupled to upstream cells. The high tail corresponding to scenario S3 is due to overload situations. From Figure 5 we conclude that the 1-Point CDV introduced by the MAC protocol is relatively low and acceptable to support real time applications without the need to re-space cells at the site of the OLT.

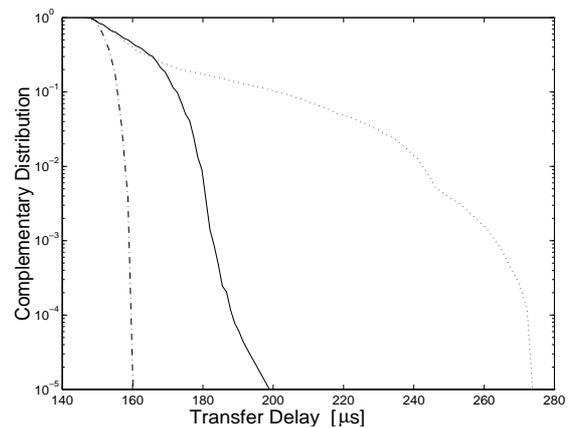


Figure 4: CPDF transfer delay CBR/VBR sources

5.3. Queue Length of ABR Traffic

Figure 6 shows the probability distribution of the queue length of an ABR buffer for all scenarios mentioned in Table 1. The algorithm to calculate the ER in the access network is effective and prevents cell loss of ABR connections, even in situations with a (temporarily) extremely high load. Recall that in the case of scenario S3 the production of cells belonging to the CBR/VBR service class can already cause an overload situation,

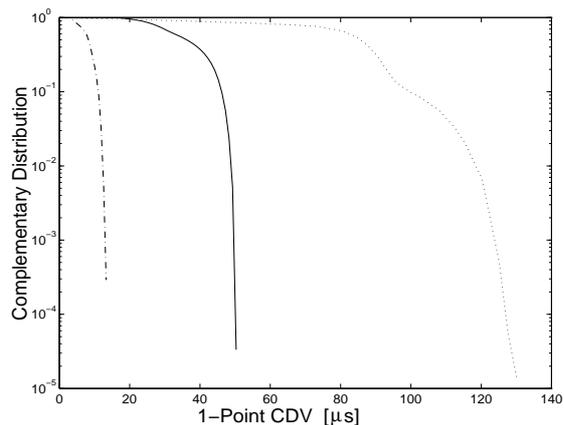


Figure 5: CPDF 1-point CDV CBR/VBR sources

and there is no way the algorithm to calculate the ER can change this. Cell loss in this extreme case could be avoided completely by allowing a zero MCR for scenario S3. For the scenarios defined in table 1, ABR buffers of 127 places are sufficient to support the ABR service category in the access network. The observation period of 180 slots provides good results.

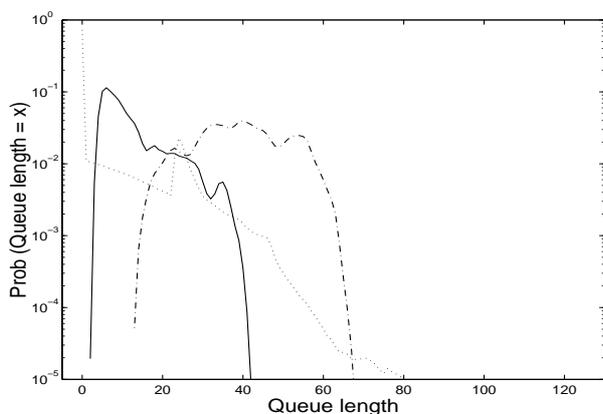


Figure 6: ABR queue length distribution

6. Conclusions

This paper proposes a MAC protocol for an ATM PON supporting different service categories. It allows users to set up connections belonging to one of the service categories CBR/VBR, ABR or UBR, guaranteeing the required QoS of the corresponding category, while making efficiently use of the network resources. In view of its stringent requirements,

CBR/VBR traffic receive priority to access the shared medium, but with a guarantee of a minimum cell rate for the ABR traffic. The remaining capacity is used for UBR traffic. As integral part of the B-ISDN, the access network has to implement a congestion control mechanism for the ABR traffic. The protocol implements a scheme based on the Explicit Rate Congestion Control mechanism similar to the ERICA algorithm. The performance of the MAC protocol is evaluated by simulations, revealing that both the MAC protocol and the adapted and integrated algorithm to prevent congestion control can be used to fully support the use of the ABR service category in the access network in a cost effective manner.

Bibliography

- [CAS93] O.Casals, J. Garcia, and C. Blondia, *A Medium Access Control Protocol for an ATM Access Network*, Proc. of the 5th Int. Conf. on Data Comm. Systems and their Performance, High Speed Networks, Eds. H.Perros and Y.Viniotis North Carolina, USA, October 1993.
- [ATM96] ATM Forum: *Traffic Management Specification*, Version 4.0, AF-TM 0056, April 1996
- [ITU95] ITU-T Recommendation I.371, *Traffic Control and Congestion Control in B-ISDN*, 1995
- [JAI96] R. Jain, S. Fahmi, S. Kalyanaraman, R. Goyal, *ABR Switch Algorithm Testing : A Case Study with ERICA ; ATM Forum contribution 96-1267*, October 1996.
- [KIL96] U. Killat (editor), *Access to B-ISDN via PONs; ATM Communication in Practice*, John Wiley & Sons Ltd B.G. Teubner, 1996.
- [PAN97] F. Panken, C. Blondia, O. Casals, J. Garcia, *A MAC Protocol for APONs Supporting Different Service Categories*, Proc.15th ITC, Washington, USA, June 1997, Eds. V. Ramaswami and P.Wirth, Elsevier, Vol.2, pp.825-834.