

Impact of time division duplexing on delay variation in slotted access systems

Sašo STOJANOVSKI and Maurice GAGNAIRE

Ecole Nationale Supérieure des Télécommunications

46 rue Barrault - 75634 - Paris cedex 13 - FRANCE

Email: {sassos, gagnaire} @res.enst.fr

Abstract

Slotted access systems and Time Division Duplexing (TDD) are two techniques well suited to various local loop configurations. Access networks are built on a shared physical medium with point-to-multipoint architecture. A centralised approach with one master entity and several slave entities is adopted. A polling-based MAC protocol enables the slave entities to dynamically share the medium. In this paper we assume that the master entity grants permits to the slaves using either the Virtual Clock or the RCSP service discipline. Both constant- and variable-length TDD frames are considered. The impact of time duplexing on delay variation in constant bitrate connections with variable bitrate background traffic is particularly investigated. For that purpose, several traffic parameters are considered, such as offered load and number of connections.

Keywords

Asynchronous Transfer Mode (ATM), Time Division Duplexing (TDD), access network, Cell Delay Variation (CDV), Rate-Controlled Static Priority (RCSP), Virtual Clock, MAC protocol.

1 INTRODUCTION

Asynchronous Transfer Mode (ATM) has been chosen in the mid '80 as the switching and multiplexing technique for the future B-ISDN networks. Its standardisation reaches the final stage. At the present time the network configuration to be adopted for the last mile of broadband networks remains an open problem. The only solution initially proposed within ITU-T was the Fibre-To-The-Home (FTTH) paradigm, which turned out to be a very expensive solution, since it requires a dedicated fibre from the local broadband exchange to each customer's premises. Recently a multitude of cost-effective solutions for the access network have been proposed, such as: several Digi-

tal Subscriber Line techniques (xDSL), Hybrid Fibre-Coax networks (HFC), ATM-based Passive Optical Networks (APON, SuperPON) or Wireless ATM (WATM). Some of these solutions tend to re-use the existing copper and coaxial infrastructures (xDSL, HFC), others require a new infrastructure, at least in the feeder and the distribution part of the access (APON, SuperPON), while WATM only requires a new spectrum band. Another way for classification of these techniques is their architecture. Namely, all the xDSL techniques are point-to-point, while the HFC, APON and WATM techniques are point-to-multipoint. Therefore, the latter techniques require a specific MAC protocol in the upstream direction. This MAC protocol must have a centralised architecture, since that is the only way to offer QoS support to the ATM layer. This is the single common point for all the MAC protocols considered so far in the literature (Acampora. 1996), (Angelopoulos *et al.* 1996), (Angelopoulos SuperPON *et al.* 1996), (Ayanoglu *et al.* 1996), (Charzinski. 1996), (Correia *et al.* 1997), (Miah *et al.* 1997), (Raychaudhuri. 1996), (Walke *et al.* 1996). On the other hand, each solution may differ by the duplexing mechanism (TDD, FDD), the access mechanism (reservation, contention, polling), the introduced MAC overhead, the service discipline for the bandwidth requests etc. In this paper, we consider a slotted access system that uses TDD scheme. Although we do not insist on identifying this access system with one of the existing ATM proposals, we may just note that it is mostly related to WATM for two reasons. First, the majority of recent WATM proposals consider TDD scheme (Raychaudhuri. 1996), (Acampora. 1996). Second, TDD schemes represent an efficient solution only when the propagation delay is small, which again is the case of WATM (cell radius of 100 metres). However, if channel utilisation is not an issue, then other environments may also be considered. In section 2, we describe the architecture of the access system, the details of the polling-based MAC protocol and the service disciplines considered. In section 3 we present the results obtained by simulation and give an intuitive explanation. The conclusions are given in section 4.

2 ACCESS SYSTEM

The ATM technology, as initially devised, is a purely switched technology with no notion for shared medium. However, recent trends in the access portion of broadband networks introduced the notion of shared medium for which a medium access control protocol has to be defined.

Only a MAC protocol with a centralised architecture can offer QoS support to the ATM layer. Therefore, the access system considered here consists of a single master entity and several slave entities, as shown in Figure 1. The master entity is the Optical Line Termination (OLT) in APON systems, the headend in HFC systems or the base station in WATM systems. Each slave entity corresponds to a single virtual connection. Note that this approach is rather different from the majority of proposals available in the literature where




Figure 1 The access system.

the slave entity corresponds to a device which represents an *aggregation* of virtual connections. Such device is the Optical Network Unit (ONU) in APON systems, the cable modem in HFC networks and the mobile stations in WATM. In the following the terms *slave* and *connection* will be used interchangeably.




Figure 2 The TDD frame structure.

The TDD frame structure is shown in Figure 2. Note that the system considered here is the one already presented in (Stojanovski *et al.* 1997). As one can see from the figure, the TDD frame structure consists of two parts: a downstream half-frame (master to slaves) and an upstream half-frame (slaves to master). The downstream half-frame consists of j downstream slots (*down 1* to *down j*), whereas the upstream half-frame consists of k upstream slots (*up 1* to *up k*). The variables j and k may take any value up to a predefined maximum value. The downstream slots in general have a different size from the upstream slots. The latter are usually bigger, since they contain a physical layer preamble. The MAC protocol has a mechanism for the slave entities to request bandwidth and a mechanism for the master entity to distribute the transmission permits. The transmission permits are announced at the beginning of each downstream half-frame (in Figure 2 this is the OverHead (OH)

slot). Also, polling of a number of slave entities is performed at the beginning of the downstream half-frame within the OH slot. The bandwidth requests of the slave entities may be sent to the master in two ways. First, as a response to the polling in the OH slot, they can send their requests in specific minislots at the beginning of the upstream half-frame. Second, they can also piggyback their requests along with the upstream data slots. The frame size is variable and each of its downstream and upstream parts can vary dynamically and independently from one another. The range of fluctuation of this size may go from zero data slots up to a predefined maximum value. These two parts are separated by a guard interval, whose length is at least equal to the propagation time between the master and the most distant slave. From the MAC layer viewpoint, we consider that all the slave entities experience the same propagation delay from the master. This is either almost true in WATM context (negligible distance disparity), or can be achieved by means of ranging in all other contexts. The master keeps track of a separate counter for each of the connections. These counters reflect more or less accurately the buffer state at the slaves. A service discipline is being applied to these counters. We consider two service disciplines: Virtual Clock and Rate-Controlled Static Priority (RCSP).

Virtual Clock is a *work-conserving* discipline that approximates the very popular WFQ (Weighted Fair Queueing) discipline. A state variable $AuxVC$ is associated to each connection. This variable is updated upon arrival of every new request n according to the formula:

$$AuxVC(n) = \max\{TNow, AuxVC(n-1)\} + 1/Peak. \quad (1)$$

where $Peak$ is the negotiated peak rate for the connection and $TNow$ is the instant of the request arrival. The newly calculated value for $AuxVC$ is used for tagging the corresponding request. When the server wants to send a permit to some of the slaves, it chooses the request with the smallest tag value. Virtual Clock has the interesting property that conforming connections are properly served even in presence of misbehaving connections.

The RCSP is a *non-work conserving* service discipline, which means that the scheduler may not serve an outstanding request even if there is available bandwidth. As a consequence, RCSP increases the average delay when compared to Virtual Clock or other work-conserving disciplines. Unlike Virtual Clock, RCSP has the very interesting property of controlling the delay variation. As shown below, this property is preserved even in the TDD context. From Figure 3 one can see that the RCSP server has two components: a rate-controller and a static priority scheduler. The rate-controller consists of a set of traffic regulators, a separate traffic regulator being associated to each connection. The parameters for the regulators are chosen according to the traffic contract of the underlying connection. Each newly arriving request

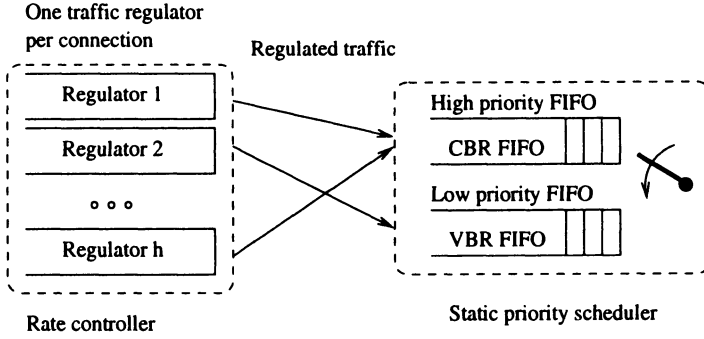


Figure 3 RCSP server.

is being tagged with an *eligibility time* ET which is calculated according to the formula:

$$ET(n) = \max\{TNow, ET(n-1)\} + 1/Peak. \quad (2)$$

This formula is identical to the one for Virtual Clock. The difference is in the way the tags are interpreted. Whereas in Virtual Clock the *AuxVC* tag is a *virtual* retransmission instant and is used just for sort-out purposes, in RCSP the ET tag corresponds to the *real* retransmission instant in case there are no other competing requests. Only after this instant is reached, is the queued request transferred from the traffic regulator to one of the static priority queues in the scheduler. In other words, the traffic regulators *re-shape* the traffic according to the underlying traffic contract. Another very interesting property of the RCSP algorithm is that it decouples the bandwidth allocation from the delay requirements. Namely, the bandwidth allocation is handled by the traffic regulators, whereas the delay requirements can be met by placing the eligible requests into the appropriate scheduler queue. There are as many scheduler queues as there are different delay classes. The service discipline in the scheduler itself is a non-preemptive one with static priorities: it always serves the request which is at the head of the highest priority non-empty queue. More details about these disciplines can be found in (Zhang. 1995).

3 PERFORMANCE EVALUATION

3.1 Simulation context

Our simulations focus on the constant bitrate traffic. We investigate the impact of the TDD scheme on access delay and access delay variation. The results are expressed as a cumulative probability density function for both

the delay and its variation. In all the simulations we consider a constant bitrate traffic in presence of a variable bitrate background traffic. The constant bitrate traffic at the sources is as ideal as can be i.e. the generated packets are equally spaced at $1/Peak$. However, since the system is slotted, we round the generation instant to the beginning of the next slot. After a number of simulations we came to the conclusion that the choice for the background variable bitrate traffic is not very important. Many types of background traffic lead to equal results. However, for the sake of preciseness, we define the traffic type that we use in our simulations. It is a single variable bitrate traffic source characterised by its peak rate $Peak$, sustainable rate $Mean$, maximum and minimum burst size ($MaxBS$, $MinBS$). The duration of the bursts is uniformly distributed between $MinBS$ and $MaxBS$. Once the burst is generated, the subsequent period of inactivity is determined in such a way that the overall mean of the generated traffic can reach the declared sustainable rate. This traffic model is conforming to the ATM Forum's VBR traffic model (ATM Forum TM. 1996). A value of 5 is chosen for the $Peak/Mean$ ratio (or *burstiness*) and the simple GCAC formulas (ATM Forum PNNI. 1996) are used to calculate the equivalent bandwidth for the VBR connection. The traffic regulator for VBR traffic is implemented in the following way: for each newly arriving request we calculate two tags, one for the peak cell rate ET_p and another for the sustainable cell rate ET_s , according to the formulas:

$$ET_p(n) = \max\{TNow, ET_p(n-1)\} + 1/Peak. \quad (3)$$

$$ET_s(n) = \max\{TNow, ET_s(n-1)\} + 1/Mean. \quad (4)$$

The request becomes eligible only when both of the following conditions are met:

$$ET_p(n) \leq TNow. \quad (5)$$

$$ET_s(n) \leq TNow + BT. \quad (6)$$

where BT is the *Burst Tolerance* and is calculated according to the formula in (ITU-T I.371. 1996), which we reproduce here as:

$$BT = (MaxBS - 1)(1/Mean - 1/Peak). \quad (7)$$

Note that an implicit policing mechanism is integrated in the VBR traffic regulator, since non-conforming VBR bursts will experience intolerable delays. The static priority scheduler distinguishes two traffic delay classes. A high priority and a low priority are associated to the CBR traffic and to the VBR traffic, respectively.

Normally, the slot size would be different for the upstream and downstream directions, as already mentioned above. However, we assume throughout the simulations that the upstream slots have the same size as the downstream ones. The number of connections which are polled at the beginning of the downstream half-frame is fixed to 4. In the Virtual Clock case, the connections are polled in a simple round-robin fashion. On the contrary, in the RCSP case only those connections which have no outstanding requests at the master are polled. We note that these different choices are well suited to each of the considered disciplines.

3.2 Variable vs constant frames

We consider a scenario with 32 constant bitrate connections with identical peak rate and a single variable bitrate connection. All the connections are symmetrical. The constant bitrate connections are dephased randomly, the phase being uniformly distributed within one period $1/Peak$. All the CBR connections occupy 50% of the link capacity, the other half being allocated to the VBR connection. The latter has a burstiness of 5, and according to the simple GCAC formulas, one can calculate that the sustainable rate (*Mean*) is equal to 34.2% of the allocated capacity. We assume a guard time of one slot between the upstream and downstream frame. Figure 4 and Figure 5 consider the cases of variable frames, for the RCSP and Virtual Clock disciplines, respectively. In these diagrams, as well as in the remainder of this paper, the bold curves represent the cumulative probability density function (pdf) for the access delay, whereas the thin curves represent the 1-point CDV pdf. The abscissa values stand for the number of time slots. On each figure there are three curves corresponding to three different values for the maximum half-frame size: 12, 18 and 24 data slots. Note that these values do not take into account the constant overhead (OH slot, minislots and guard times).

By comparing Figure 4 and Figure 5, one can draw several conclusions. First, in both cases the access delay increases with the maximum half-frame size. Second, the access delay distribution strongly depends on the service discipline: in case of Virtual Clock it is concentrated around much lower values than for RCSP. This is another confirmation of the well known result that work-conserving disciplines imply lower *average* delay. In the RCSP case, it is the traffic regulators that increase the delay by re-shaping the traffic. Nevertheless, the values for the *maximum* access delay are comparable in both cases. Third, the delay variation (measured with the 1-point CDV) increases with the maximum half-frame size, whatever the service discipline. However, while the maximum values for the delay variation are of the order of the access delay for Virtual Clock, these maximum values are *bounded* for RCSP. This is the most interesting result and it will be confirmed on the following figures, as well. Note that the delay variation bound depends on the maximum half-




Figure 4 Variable frames, RCSP, max half-frame size of 12, 18 and 24.




Figure 5 Variable frames, VClock, max half-frame size of 12, 18 and 24.

frame size, and its value roughly corresponds to three times the maximum half-frame size. Namely, the maximum values for the delay variation are: 39, 57 and 75 for the maximum half-frame size of 12, 18 and 24, respectively.




Figure 6 Constant frames, RCSP, max half-frame size of 12, 18 and 24.

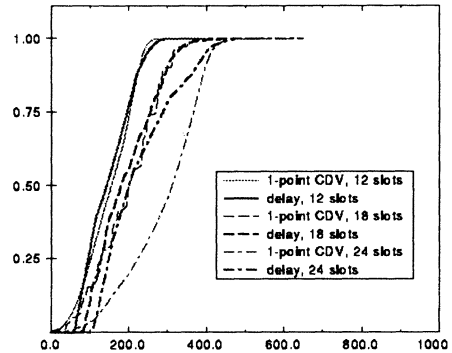


Figure 7 Constant frames, VClock, max half-frame size of 12, 18 and 24.

On Figure 6 and Figure 7 we compare the service disciplines in the case of constant frames, with half-frame size equal to: 12, 18 and 24. One can note that access delays increase and the delay variation decreases in the case of fixed-length frames, compared to the case of variable frames. Again, one can note that the delay variation is of the order of the access delay for Virtual Clock, while in case of RCSP the delay variation is bounded. Furthermore, one can see that this bound is even smaller than in the case of variable frames. In spite of the fact that with constant frames one can reduce the delay variation (at the expense of slightly increasing the access delay), this case will not be considered further. The fact that statistical gain in case of asymmetrical traffic is achievable only with variable-sized frames justifies this choice.

3.3 Impact of traffic parameters

Figure 8 and Figure 9 present the variable-frame case in which the parameter is the number of CBR connections: 1, 16 and 64. In all the cases, the generated CBR traffic represents 50% of the link capacity i.e. the cell interarrival time for the CBR connections is equal to 5, 80 and 320 slot times.




Figure 8 Variable frames, RCSP, parameter is number of CBR connections: 1, 16 and 64.




Figure 9 Variable frames, VClock, parameter is number of CBR connections: 1, 16 and 64.

One can see that delay variation is insensitive to the number of connections with RCSP. In fact, the three curves representing the 1-point CDV cumulative pdf are indistinguishable from one another and appear as a single curve in Figure 8 (the first curve from the left). On the contrary, in the Virtual Clock case the delay variation increases with the number of connections and has maximum values very close to the maximum access delays.

Figure 10 and Figure 11 consider variable-frame cases with 32 CBR connections, the changing parameter being the CBR traffic load. One can see that by increasing the CBR traffic load (25%, 50% and 75% of the link capacity), the delay variation increases with Virtual Clock, whereas it remains unaffected with RCSP.

3.4 Comments about the RCSP discipline

In this section we try to justify the values obtained for the delay variation bound. Let us consider the case with RCSP service discipline and variable frames with maximum half-frame size of 12. In a "steady state" the CBR bandwidth requests arrive more or less periodically into the traffic regulators. After being held for some time, they become eligible and arrive as credits into the scheduler's CBR FIFO. At the beginning of each downstream half-frame, the server announces the permits for the following upstream half-frame. It




Figure 10 Variable frames, RCSP, parameter is the CBR load: 25%, 50% and 75% of the link capacity.




Figure 11 Variable frames, VClock, parameter is the CBR load: 25%, 50% and 75% of the link capacity.

collects all the available CBR credits from the CBR FIFO, and then looks for the credits in the VBR FIFO. The server collects VBR credits up to the moment when either the VBR FIFO gets empty or the total number of credits reaches the value of 12 (the maximum half-frame size).

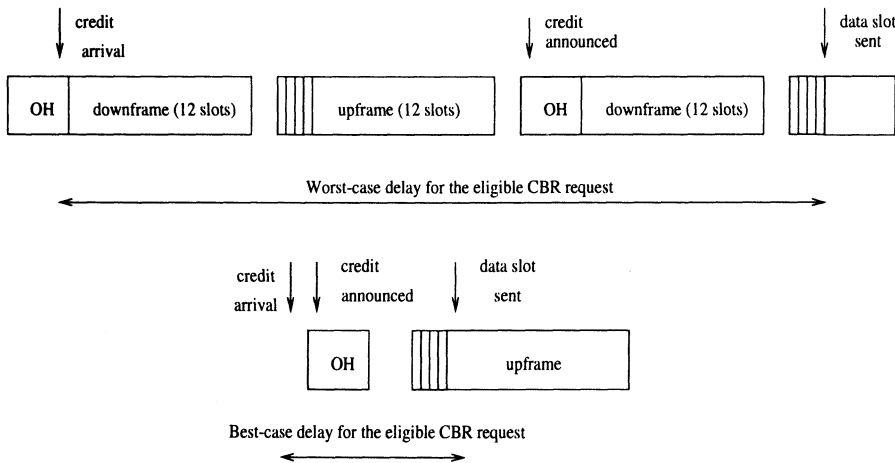


Figure 12 Worst-case and best-case waiting time.

Whenever the VBR connection has simultaneous bursts in both directions, the frame size reaches the value of $12 + 12$ data slots. In the worst case, as shown in Figure 12, a new CBR request (or a "credit", as marked in the figure) becomes eligible and arrives in the CBR FIFO just after the OH slot i.e. the instant when the server already announced the 12 credits. Consequently, this request will wait for the duration of a full-sized frame ($12 + 12$) before being announced. After being announced in the next OH slot, there is again a whole downstream half-frame (12 slots) of waiting time before the upstream half-

frame starts. Therefore, the worst-case waiting delay T_w for the credit is equal to:

$$T_w = 3 * 12DataSlots + 3 * GuardTime + 1 * OHslot + 2 * MiniSlots. \quad (8)$$

On the opposite side, when the VBR connection has no burst for transmission in any of the two directions, the frame size is much shorter. It is even possible to have frames with zero data slots, given that the CBR connections are dephased and the two traffic directions are uncorrelated. In the best case, as shown in Figure 12, a new CBR request becomes eligible and arrives in the CBR FIFO just before the beginning of a new downstream frame. With the assumption that there are no downstream data for the current downframe, the best-case delay T_b experienced by the credit is the delay induced by the constant frame overhead:

$$T_b = 1 * OHslot + 1 * GuardTime + 1 * MiniSlots \quad (9)$$

Hence, the maximum delay variation experienced by the credits is roughly equal to the difference ($T_w - T_b$), which in this specific case is equal to 39 slot durations. With similar reasoning one can conclude that in the constant-frame case the maximum delay variation is equal to two half-frames (instead of three).

4 CONCLUSION

In this paper, we have investigated a generic slotted system for access to broadband distribution networks. It uses time-division duplexing and a centralised polling-based MAC protocol. Although we have tried to remain as general as possible throughout the paper, it is true that the applicability of the TDD approach is limited to small-range systems, such as WATM. Nonetheless, this approach could be applied to other contexts, too, if bandwidth efficiency is not an issue. We have considered the advantages of using a non-work-conserving discipline for request scheduling over using a work-conserving one. It has been shown through simulations that the interesting properties of the non-work-conserving disciplines are preserved even in the specific TDD context. The comparison has been performed by changing several parameters, such as the traffic load or the number of connections. Simulations with both constant and dynamically-varying frames have been performed. From all the simulations we have drawn several conclusions. First, delay variation is very influenced by traffic parameters in case of Virtual Clock and is almost insensitive to traffic parameters in case of RCSP. Second, in the RCSP case the maximum value for the delay variation is *predictable* and is approximately equal to the duration of two maximum-sized downframes plus the duration of one maximum-sized up-frame. Clearly, our simulation results show that in case of temporary overloads

due to the background traffic, it is the maximum frame size that determines the upper bound for the delay variation of constant bitrate connections. This predictability for the maximum delay variation is a nice feature, since it corresponds actually to the CDV Tolerance to be used with regard to the first upstream ATM switch. The values expressed above in slot durations should be divided by two since the slot duration in the context of TDD is half of a non-TDD slot duration. Third, by using frames of constant size one can decrease the delay variation, but increase the average access delay. Fourth, one can see that the average access delay is always bigger with RCSP than it is with Virtual Clock. However, they have comparable values for the maximum access delay.

REFERENCES

- Acampora, A. (1996) Wireless ATM: A Perspective on Issues and Prospects *IEEE Personal Comm.*, **August 1996**, 8–17.
- Angelopoulos, J.D., Fragoulopoulos E.K. and Protonotarios E.N. (1996) Efficient Support of Best effort Traffic in Passive Tree Local Loops *SPIE Journal*, **Vol. 2609**, 182–191.
- Angelopoulos, J.D., Koulouris J. and Fragoulopoulos E. (1996) A MAC Protocol for an ATM-based SuperPON *SPIE Journal*, **Vol. 2919**, 279–287.
- ATM Forum (1996) Traffic Management Specification *Version 4.0*, **1996**.
- ATM Forum (1996) ATM Forum PNNI Specification *Version 1.0*, **1996**.
- Ayanoglu, E., Kai Y.E. and Karol M.J. (1996) Wireless ATM: Limits, Challenges and Proposals *IEEE Personal Comm.*, **Vol.3 No.4**, 18–34.
- Charzinski, J. (1996) A New Approach to ATM Access Networks *SPIE Journal*, **Vol.2917**, 108–119.
- Correia, L.M. and Prasad, R. (1997) An Overview of Wireless Broadband Communications *IEEE CommMag.*, **Vol.35 No.1**, 28–33.
- ITU-T (1996) I.371 Traffic Control and Congestion Control in B-ISDN *Geneva*, **June 1996**.
- Miah, B. and Cutbert, L. (1997) An Economic ATM Passive Optical Network *IEEE CommMag.*, **Vol.35 No.3**, 62–68.
- Raychaudhuri, D. (1996) Wireless ATM Networks : Architecture, System Design and Prototyping *IEEE Personal Comm.*, **Vol.3 No.4**, 42–49.
- Stojanovski, S. and Gagnaire, M. (1997) A New Wireless ATM Access Protocol for the Local Loop *First IFIP/IEEE International Workshop on Mobile and Wireless Communications Networks*, **May 1997**, Paris.
- Walke, B., Petras, D. and Plassmann, D. (1996) Wireless ATM: Air Interface and Network Protocols of the Mobile Broadband System *IEEE Personal Comm.*, **Vol.3 No.4**, 50–56.
- Zhang, H. (1995) Service Disciplines for Guaranteed Performance in Packet-Switching Networks *Proceedings of the IEEE*, **Vol.83 No.10**.