



HAL
open science

Traffic conformance issues in passive optical networks: challenges and solutions

Minh Thanh Ngo, Annie Gravey

► **To cite this version:**

Minh Thanh Ngo, Annie Gravey. Traffic conformance issues in passive optical networks: challenges and solutions. International journal of communication networks and distributed systems, 2009, 2 (2), pp.260-280. 10.1504/IJCNDS.2009.022335 . hal-02161302

HAL Id: hal-02161302

<https://hal.science/hal-02161302>

Submitted on 14 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Traffic conformance issues in passive optical networks: challenges and solutions

M. Thanh Ngo and Annie Gravey

Department of Computer Science,
Institut TELECOM/TELECOM Bretagne,
CS 83818, Brest, France
Fax: ++33 2 29 00 12 82
E-mail: mt.ngo@telecom-bretagne.eu
E-mail: annie.gravey@telecom-bretagne.eu

Abstract: Broadband passive optical networks (PONs) share a single optical fiber transmission system between several customers. A dynamic bandwidth allocation (DBA) is used to avoid collisions for upstream traffic. An ideal DBA should support both committed bandwidth and best effort services while efficiently utilising the available bandwidth and providing isolation between customers in order to offer fair access to resources. In the present paper, we assume that each customer negotiates traffic profiles for its upstream committed bandwidth traffic. We show that a DBA that does not take into account conformance to the negotiated traffic profiles cannot isolate compliant customers from non-compliant customers, leading to QoS degradations for compliant customers. We then propose a framework for taking into account the traffic conformance within the DBA. This study is illustrated by an original, priority based DBA designed for Ethernet optical networks (EPONs).

Keywords: fiber to the home; full service access network; passive optical network; dynamic bandwidth assignment; multiservice network; class of service; CoS; quality of service; QoS; traffic profiles; service level agreement; traffic conformance; conformance checking.

Reference to this paper should be made as follows: Ngo, M.T. and Gravey, A. (2009) 'Traffic conformance issues in passive optical networks: challenges and solutions', *Int. J. Communication Networks and Distributed Systems*, Vol. 2, Nos. 2/3, pp.260–280.

Biographical notes: Minh Thanh Ngo received his BE in Computer Science from INSA, Toulouse, France, in 2005, MEng in Networks and Telecommunications at INSA, Toulouse, France in the same year. He is currently working towards his PhD in Computer Science and Telecommunications at the ENST-Bretagne, France. His research interests are in the area of quality-of-service in broadband access networks, medium access control protocol design and network traffic modeling.

Annie Gravey is currently Professor and Head of the Computer Science Department at Institut TELECOM/TELECOM Bretagne, in Brest (France). Before that, Dr. Gravey spent 20 years with France Telecom R&D where she designed mechanisms for specifying and controlling broadband traffic and conducted extensive operational activities in standards and network deployment. Her current research interests include the design and evaluation of traffic engineering methods for broadband access (optical, wired and wireless) and spontaneous networking in sensor or P2P networks.

1 Introduction and problem statement

An access network is the set of links and nodes that allows business and residential subscribers to reach the interconnection network. It is known as either the last mile or the first mile bottleneck. The demand for high capacity services such as interactive multimedia communications, real-time video streaming, high-definition television (HDTV), interactive games and video-on-demand (VoD) is continuously increasing. Green (2004) has shown that although existing access technologies based on copper (xDSL) and cable modems (CATV) are rapidly evolving to partially fulfill these needs, these technologies may well not satisfy future users' requests. Access networks relying on optical technologies, commonly called FTTx access networks, represent an alternative to copper and cable modem based access. The capacity that an optical fiber can support is conceivably large enough to provide high broadband services to subscribers. FTTx has thus become the preferred alternative to other technologies for full service access network (FSAN), especially in areas when deploying new access networks.

The broadband passive optical network (PON) is one of the most promising approaches to support FTTx. Due to cost issues, a PON deploys passive optical components (for example, splitters and couplers) with no active element in the path of signals between the end points and the interconnection network. The tree, ring and bus are several topologies suitable for PONs, with the tree topology being the most common architecture. Typically, a PON is a point-to-multipoint network with a tree topology. It consists of an optical line terminal (OLT), a 1: N passive star splitter, and multiple optical network units (ONUs). The OLT, located at the provider's central office, acts as the distributor, arbitrator and aggregator of traffic. In the downstream direction (from OLT to ONU), data frames are broadcasted by the OLT to every ONU. ONUs filter and discard frames that are not addressed to them. In the upstream direction (from ONU to OLT), data frames from different ONUs transmitted simultaneously may collide. Thus, PONs need arbitration mechanisms to avoid collisions and to share the fiber capacity between ONUs for upstream traffic. This paper addresses the upstream traffic arbitration issue.

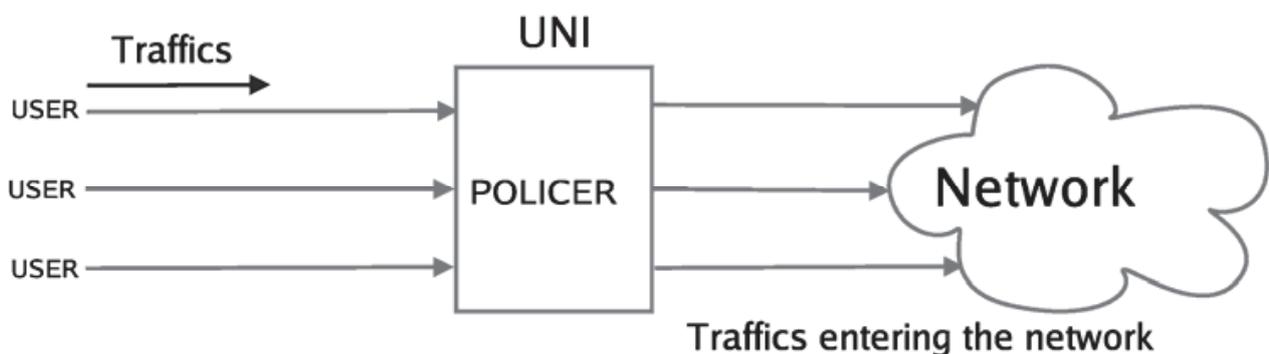
The original PON standards (Green, 2004), known as APON (ATM-PON) and BPON (broadband PON), focused on ATM as MAC layer. APON standards were promoted by the FSAN consortium and were published by the ITU-T as the G.983 x specifications. The original APON technology suffers from some shortcomings such as large overhead, complexity and cost. Recently, an alternative PON standard, namely gigabit PON (GPON), has been specified in the G.984.x ITU-T Recommendations and is promoted by the FSAN. A GPON uses a faster interface (up to 2.5 Gbps in the current standards) and is able to transmit variable-length packets, including Ethernet frames and ATM cells. Another alternative is the Ethernet PON (EPON), which is being discussed within the (IEEE 802.3ah, 2004) by IEEE Ethernet in the first mile (EFM) task force. An EPON is a PON that carries data encapsulated in Ethernet frames.

A simple mechanism for arbitrating between upstream traffic flows in EPONS is fixed time division multiplexing (TDM) (Kramer et al., 2002) in which every ONU gets a fixed timeslot to transmit upstream traffic. While TDMA (TDM Access) is very simple and well suited to support balanced and permanent traffic offers, its drawback is an obvious lack of efficiency for highly bursty sources, or unbalanced demands. Indeed, in these situations, time slots allocated to idle ONUs would be unused while active ONUs would be congested due to insufficient transmission opportunities. Another drawback of TDMA is the lack of support for dynamic topology changes (for example node failure, node activation etc.). Many mechanisms supporting a dynamic allocation of transmission opportunities have been proposed in the literature (Zheng and Mouftah, 2005; McGarry et al., 2004). They are all based on two complementary types of mechanisms:

- A dynamic bandwidth allocation (DBA) mechanism centralised in the OLT that distributes transmission opportunities to the ONUs. The DBA relies on a signaling mechanism that allows an ONU to describe its requirements and the OLT to allocate transmission opportunities. It implements a method to optimise the allocation of transmission opportunities based on the ONUs' demands, the negotiated Service level agreements (SLAs) and specific traffic management policies that are proprietary to the access operator.
- An intra-ONU traffic management policy that, within an ONU, allocates transmission opportunities to the upstream traffic queued for transmission, according to the global allocation forwarded by the OLT.

It is important to note here that the DBA should not rely only on the dynamic demands transmitted by the ONUs. Indeed, it should check that these requirements are consistent with the negotiated SLAs, since the SLAs globally specify what type of service should be delivered to each ONU. If an SLA corresponding to pure TDMA is negotiated, this is easily done. It is much harder when a dynamic TDM policy is chosen to operate the PON. Indeed, the usual approach in an active access network (for example point-to-point architecture) for enforcing the negotiated SLAs is to filter incoming traffic according to the supported traffic classes or connections and then to check its conformance with the traffic profiles mandated in the SLAs. The policer which implements traffic enforcement is usually implemented in the line cards of the provider edge equipment (Figure 1), which, in a PON scenario, would then be the OLT.

Figure 1 Location of policing function in the network



However, checking conformance of upstream traffic at the OLT is already too late if the user has submitted non-conforming traffic for which transmission opportunities have been allocated by the OLT. As we show in Section 4.2, non conforming traffic competing with conforming traffic may indeed lead to QoS degradation for the conforming traffic unless conformance is checked before transmission opportunities are allocated. On the other hand, checking for conformance within the ONU, prior to requiring transmission opportunities, is also sub-optimal, similarly to a TDMA policy, since the OLT cannot take advantage of unused resources to reallocate them to active ONUs. To the best of our knowledge, this mandatory relationship between traffic conformance and DBA policies has not been completely addressed by published EPON DBAs. The present paper addresses this issue by proposing a framework for merging conformance checking with the DBA controlling the access to transmission resources of upstream traffic.

This paper is organised as follows. Section 2 presents related work. Section 3 identifies the main traffic classes which are to be supported and discusses their characteristics and requirements. Section 4 presents a multi-class DBA and demonstrates the QoS degradation due to non-conforming traffic when traffic conformance is not taken into account in the DBA. The next section presents our proposal for merging the multi-class DBA with a policy strategy enforcing traffic profiles. This section illustrates the behaviour of our DBA and shows its efficiency in protecting conforming traffic. The last section concludes the paper by outlining the perspectives for future development.

2 Background knowledge and related research

The present section first describes the mechanisms used to control upstream traffic transmission. It then reviews the state of the art regarding these mechanisms and conformance checking issues.

2.1 MPCP and polling schemes

2.1.1 Multi point control protocol (MPCP)

MPCP is a signaling protocol developed by the IEEE 802.3ah (2004) task force to facilitate the dynamic bandwidth allocation of transmission opportunities. MPCP in itself does not specify any particular algorithm for bandwidth allocation. This protocol simply provides a tool for exchanging requests and grants. MPCP uses two control messages (gate and report) for coordinating upstream data transmission and three other signaling messages (register request, register, and register ACK) for automatic ONU discovery. Each message is a standard 64-byte MAC control frame.

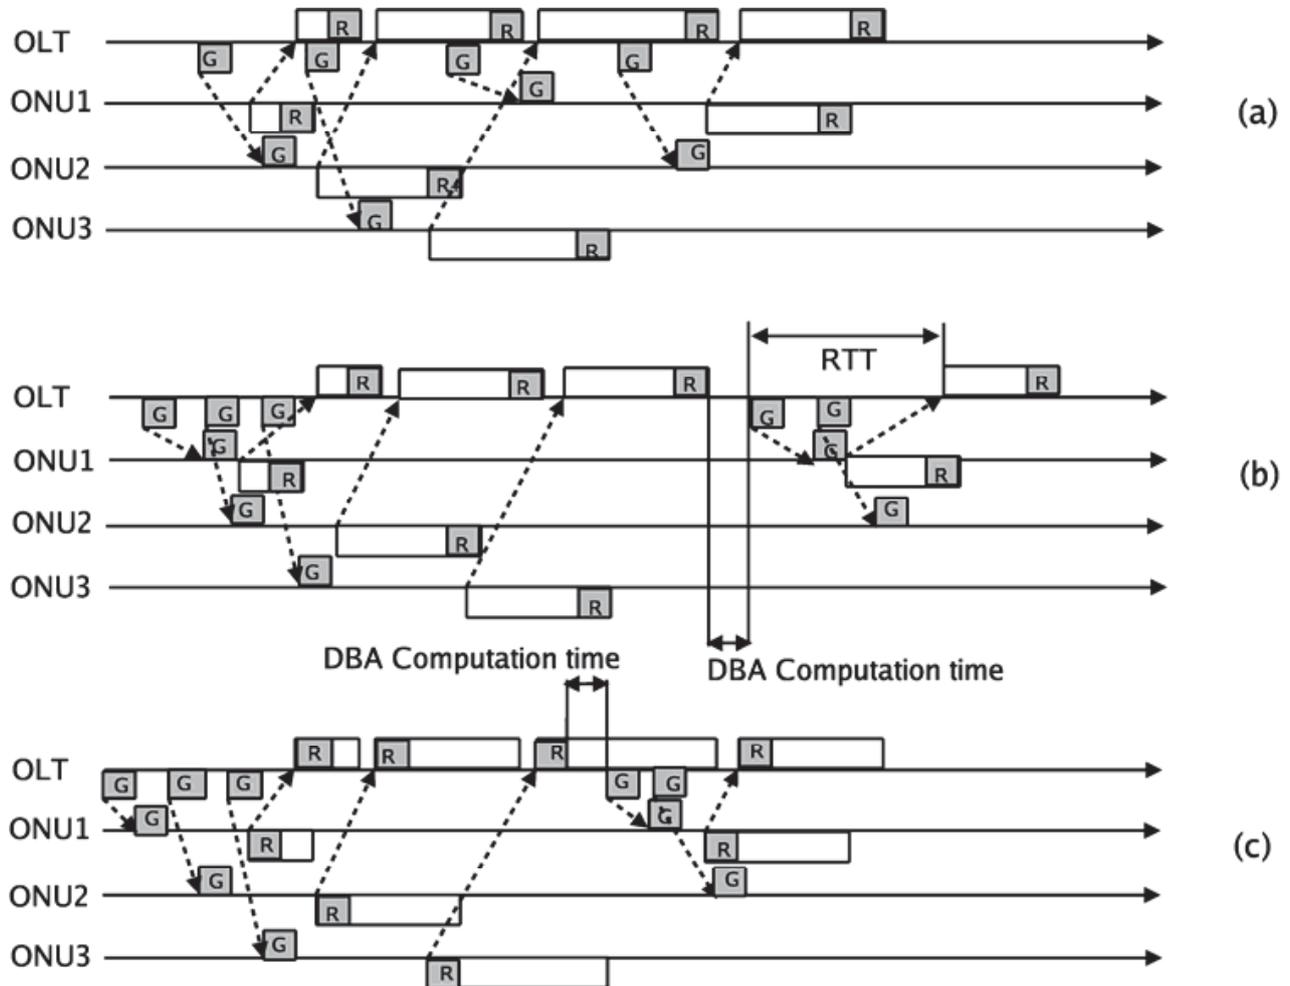
The gate message is used to grant transmission windows (TW) to ONUs. Each gate message contains the transmission start time and transmission length of the time period allocated to the corresponding ONU. In the report messages, ONUs indicates their current requirements. These requirements can be the exact number of bytes waiting in each queue or estimation by each ONU of its own requirements depending on some local policy. A report message is sent by an ONU in the allocated TW with data packets.

2.1.2 Polling schemes

A polling protocol can poll multiple ONUs for transmission based on different policies.

A simple policy, called poll-and-stop polling, in which, the ONUs are polled and allowed to transmit one after the other, with a complete round-trip message walking overhead time required for each ONU. Obviously, this protocol wastes a lot of bandwidth on the upstream channel, especially when the OLT is serving a large number of ONUs, which would largely reduce global channel utilisation and thus increase packet delay.

Figure 2 Polling schemes



A more efficient way is to interleave polling messages, allowing the OLT to send a gate message to the next ONU before the data and report message(s) from the previous polled ONU(s) arrive, as shown in Figure 2a. Since this protocol allows the OLT to allocate bandwidth based only on already received bandwidth demands, the OLT is therefore unable to take into account the bandwidth demands of all ONUs and make an intelligent decision based on global demand.

Figure 2b illustrates an effective way to overcome this drawback, called cyclic polling policy. Unlike interleaved polling, in cyclic polling policy, the OLT does not start the next polling cycle before all report messages are received. This allows the OLT to perform bandwidth allocation based on the bandwidth demands of all ONUs at the end of each polling cycle and thus make a more intelligent decision. However, as shown in Figure 2b, the upstream channel is not utilised from the instant the transmission of the last polled ONU in the previous cycle is completed to the instant the transmission of the

first polled ONU in the next cycle starts, which can result in underutilisation of the upstream channel. This issue is addressed by Choi (2004) by proposing a more efficient cyclic polling method in which ONUs send report messages at the beginning of the transmission window. This polling scheme is illustrated in Figure 2c and is used in the remainder of the present paper.

2.2 *Related research work*

Interleaved polling with adaptive cycle time (IPACT), which uses interleaved polling, has been proposed for EPON in Kramer et al. (2002). We note that in IPACT, the OLT issues global grants to ONUs, which means that the OLT does not dictate how many bytes exactly from a particular queue an ONU must transmit. Instead, the ONU uses a strict priority policy to determine the order in which the queues are processed. It has been shown that the global grant approach proposed by IPACT cannot provide different QoS guarantees to different end users.

Ma and Zhu (2003) have proposed a bandwidth guaranteed polling protocol, which allows the upstream bandwidth to be shared based on the service level agreement (SLA) between each ONU and the operator. This protocol is able to provide a bandwidth guarantee for premium ONUs based on SLAs while providing best-effort service to other ONUs. However, in this model, an ONU is unable to request different QoS services for different traffic streams. An and Hsueh (2003) proposed a similar approach, named hybrid slot-size/rate (HSSR), in which traffic is classified into two priority classes. It allocates a fixed amount of bandwidth to the high priority class in order to minimise the delay and jitter of packets. The high utilisation and hybrid granting (HUHG) proposed by Kim (2006) focused on improving upstream bandwidth utilisation and the QoS offered to the expedited forwarding (EF) service class. This is achieved by dividing the cycle into two sub-cycles, the EF sub-cycle for the EF class and the AF sub-cycle for the assured forwarding (AF) and best effort (BE) classes and by granting the next EF sub-cycle in advance, using the deterministic feature of EF service class. ONUs with unsatisfied AF demands will evenly share the remaining resource. Another DBA with support for DiffServ is proposed in Choi (2004). In this approach, high priority (EF) traffic would be allocated a guaranteed amount of bandwidth regardless of the number of packets waiting in the queue. Medium priority (AF) traffic would be allocated bandwidth as per their proportion in the total demand. Typically, these DBAs (Ma and Zhu, 2003; An and Hsueh, 2003; Kim, 2006; Choi, 2004) do not yet explicitly guarantee a minimum bandwidth for low priority streams, especially when one or more ONUs send a large amount of low priority data into the network. Furthermore, the fixed bandwidth for high priority traffic cannot always satisfy instantaneous demands.

Assi et al. (2003) proposed a DBA scheme in which ONUs were partitioned into two groups, under-loaded and overloaded, according to their minimum guaranteed transmission window sizes. Here, the total bandwidth saved from the under-loaded group is proportionally re-allocated to over loaded ONUs. A contribution of this scheme is to fairly distribute the excessive bandwidth among highly loaded ONUs. However, with this scheme, the OLT still only focuses on how to satisfy bandwidth requests from different ONUs and does not provide a better service to high priority traffic. Furthermore, this mechanism does not provide any mechanism to prevent a user from consuming more available transmission opportunities by sending excessive requests.

Our approach supports both committed bandwidth (CB) and best effort (BE) traffics and also allows discrimination between different classes of CB traffic. It does not rely on a static allocation of transmission opportunities for CB traffic, but dynamically allocates transmission opportunities based on actual demands. On the other hand, it relies on a connection access control (CAC) policy that limits the amount of CB traffic in order to support committed QoS.

Before describing this approach in Section 4, the following section addresses the selection of a minimal set of QoS classes and describes QoS indicators used in our performance study.

3 Traffic classification and QoS requirements

Multimedia applications, such as voice over IP, video conferencing, and VoD generate traffic with characteristics that significantly differ from traffic generated by data applications (file upload and download), and they usually have more stringent delay and loss requirements.

3.1 Traffic classes

In this section, we identify a minimal set of traffic classes that can be used to differentiate network services. In order to efficiently support diverse application requirements, the usual approach in Ethernet networking is to classify packets into a limited number of traffic classes. The ANSI/IEEE Standard 802.1D (1998) identifies eight traffic classes but we have chosen to take into account only a limited subset of three classes as shown in Table 1. This does not preclude the definition of several subclasses in each global class.

Table 1 Traffic types

<i>Classes</i>	<i>Characteristics</i>	<i>Requirements</i>
T_0	Real-time CB	Limited loss rate, low delay and jitter, guaranteed bandwidth
T_1	Data CB	Limited loss rate, limited delay and jitter, guaranteed bandwidth
T_2	Best effort	None

Basically, we assume that there are two CB classes and a single best BE class. A major difference between CB and BE traffic is that CB traffic is characterised by a traffic profile and expects its QoS requirements to be fulfilled as long as the offered traffic complies with the negotiated traffic profiles. On the other hand, there are no BE traffic profiles and BE traffic cannot request QoS commitments. We have chosen here to assume that traffic profiles are specified as classical token buckets (or equivalently leaky buckets) since they are the traffic profiles used in ATM, IP or MPLS networks. The difference between the two CB classes is that the first one is intended to support real time applications (interactive multimedia communications such as VoIP) whereas the second one should support less time sensitive applications that require either guaranteed bandwidth (for example IPTV, VoD) or some level of interactivity (for example e-banking, database synchronisation, etc).

3.2 QoS metrics

QoS requirements should be fulfilled end-to-end. Achieving the end-to-end QoS requirement (delay and jitter in particular) of an individual flow traversing a network path is usually addressed by partitioning the end-to-end QoS requirement into the local QoS requirements in each individual network element across that network path and then achieving them in each network element. In our work, we focus on the values taken by QoS metrics in an EPON architecture.

3.2.1 Packet delay (latency)

Latency is defined as the time interval between the departures of data from the source to the arrival at the OLT. The required performance is expressed in terms of the maximum delay bound. According to the ITU-T Recommendation G.114, a maximum of 1.5 ms is recommended for the T_0 packet delay in an access network. The delay should be limited for both T_0 and T_1 traffic. However, it does not make any sense to put an absolute upper bound on the delay since this does not make a performance parameter that can be computed in a statistically reliable way (an absolute maximum delay increases as connection time increases). Therefore, we define the delay constraint by upper bounding a p -percentile of its distribution. Let D_{max} be the upper bound for packet delay, relevant for T_0, T_1 . Given the delay value D of the p -percentile, the probability that the delay is larger than D is at most $1-p$.

$$D_{max} = D, \text{ and } P(\text{Delay} \geq D) \leq 1 - p$$

3.2.2 Delay variation (jitter)

Jitter is an important QoS parameter for real-time CB traffic (T_0). Since jitter in packet networks can never be completely eliminated, information streams that have very low jitter requirements must use a playout buffer to smooth out the delay variability. We also define the jitter constraint using p -percentile values (D_{max}, D_{min}).

$$\text{Jitter } J = D_{max} - D_{min}$$

4 Performance of a multi-class DBA

In this section, we describe a multi-class DBA that implements a priority based allocation mechanism to share transmission opportunities between ONUs. As a first step, we show how the allocation is performed under the assumption that all demands are compliant with the negotiated traffic profiles. We then investigate the performance of this DBA, and demonstrate the impact of non conformance on delivered QoS.

4.1 Definition of a multi-class DBA

In order to support differentiated QoS, two separated, but inter-related scheduling mechanisms are required.

- A high-level scheduling mechanism (inter-ONU scheduling): Implemented at the OLT, the DBA arbitrates between the upstream transmissions by allocating an appropriate TW to each ONU.
- A low-level scheduling mechanism (intra-ONU scheduling): Implemented at each ONU, this scheduling mechanism distributes the bandwidth allocated by the OLT among its local queues.

As discussed in Section 2.1.1, the ONUs and the OLT use MPCP to exchange control messages. We have chosen a cyclic allocation process in order to have good knowledge of all pending requests before actually computing the allocations. The OLT receives the information about per-class demands relative to each ONU. The DBA uses this information to determine the amount of bandwidth allocated to each queue of each ONU. Gate messages carry this information to each ONU, informing the ONU when to transmit and how much it is allowed to transmit (start time and stop time of next TW). The ONU sends a new report at the beginning of each TW. The report carries per-class demands that will not be accommodated in the remainder of the present TW. The TW is then used to actually transmit the data. There are many possible ways of expressing per-class demands. We have chosen in the simple DBA scheme described below to report actual queue lengths, although nothing precludes its extension to more complex indicators. As usually assumed in all papers proposing DBA policies, each ONU is expected to comply with its allocated TW. Furthermore, we also assume that the OLT describes per-class transmission opportunities and that the ONU complies with these per-class allocations. There are several arguments in favour of this conservative policy:

- usually, traffic is marked to indicate its class, and the OLT could implement a filtering process that discards traffic sent outside its allocated TW
- supporting QoS relies on enforcing a scheduling policy that is global to the PON and not local to each ONU; by favouring one class over others, an ONU can degrade the QoS of lower classes with little advantage to the favoured traffic class.

We have chosen to implement a simple DBA policy which:

- is strictly priority based at the inter-ONU level
- bases its computations on the actual transmission requested carried in REPORT messages, and does not anticipate requests.

The variables used below are now defined. We consider an EPON with N ONUs and assume that the link capacity is $CMbps$. DBA schemes frequently define a so called ‘polling cycle’. A polling cycle is defined as the time that elapses between two successive TW assigned to the same ONU. The polling cycle can have a fixed or a variable length. Clearly, the polling cycle is an important parameter of the system; it impacts on the QoS that can be supported (both the latency and the jitter increase with longer polling cycles), on the ONU buffer size (longer polling cycles imply larger buffers) and on the efficiency of the allocation process (short polling cycles increase both protocol overhead and total guard times while costing more CPU time for the OLT). As many before us, we have chosen to allow variable length polling cycles while enforcing an upper bound on polling cycle length.

We denote the maximum polling cycle in bytes and in seconds by B^{max} and T^{max} respectively. We also denote the maximum amount of bandwidth (in bytes) B_i^{lim} offered to a given ONU_{*i*} in the case where all ONUs are heavily loaded. Based on ONU_{*i*}'s negotiated traffic profile, B_i^{lim} can be computed as follows:

$$B_i^{lim} = \frac{(T^{max} - Nt_g)C}{8} w_i, \text{ with } \sum_{i=1}^N w_i = 1 \quad (1)$$

where t_g is the inter-ONU guard time and w_i is the weight of the *i*th ONU according to its SLA (for example the ratio of the bandwidth committed to this ONU to the total amount of committed bandwidth), w_i takes into account reservations for both T0 and T1 traffic. With the above equation, we note that B_i^{lim} varies depending on the negotiated SLAs. B_i^{lim} is NOT an upper limit on the amount of resources allocated per cycle time to ONU_{*i*}. In some cases, if some ONUs are not very active, ONU_{*i*} can get larger transmission opportunities than B_i^{lim} . However, if all ONUs are sufficiently active, B_i^{lim} is the effective limit which enforces the upper bound to the cycle time. We denote $R_i^{T_0}, R_i^{T_1}, R_i^{T_2}$ as the high, medium and low priority request (in bytes) of the *i*th ONU respectively. The DBA algorithm serves each class successively and attempts to satisfy all the demands from one class before allocating transmission opportunities to a lower class. This translates as follows, denoting $G_i^{T_k}$ as the grant allocation computed for traffic T_k of the *i*th ONU:

Step 1 The OLT begins by allocating transmission opportunities for class T_0 :

$$G_i^{T_0} = \begin{cases} R_i^{T_0} & \left\{ \begin{array}{l} \text{if } \sum_{i=1}^N R_i^{T_0} < B^{avail} \\ \text{or } R_i^{T_0} < B_i^{lim} \end{array} \right. \\ B_i^{lim} + B_i^{aT_0} & \text{otherwise} \end{cases} \quad (2)$$

$$\text{with } B_i^{aT_0} = \frac{B_{exc}^{T_0} R_i^{T_0}}{\sum_{O^{T_0}} R_i^{T_0}}, B_{exc}^{T_0} = \sum_{U^{T_0}} (B_i^{lim} - R_i^{T_0}), \text{ and } B^{avail} = B^{max}$$

$$O^{T_0} = \{ONU_i \mid R_i^{T_0} > B_i^{lim}\}$$

$$U^{T_0} = \{ONU_i \mid R_i^{T_0} < B_i^{lim}\}$$

Step 2 The OLT then computes $G_i^{T_1}$ similarly with

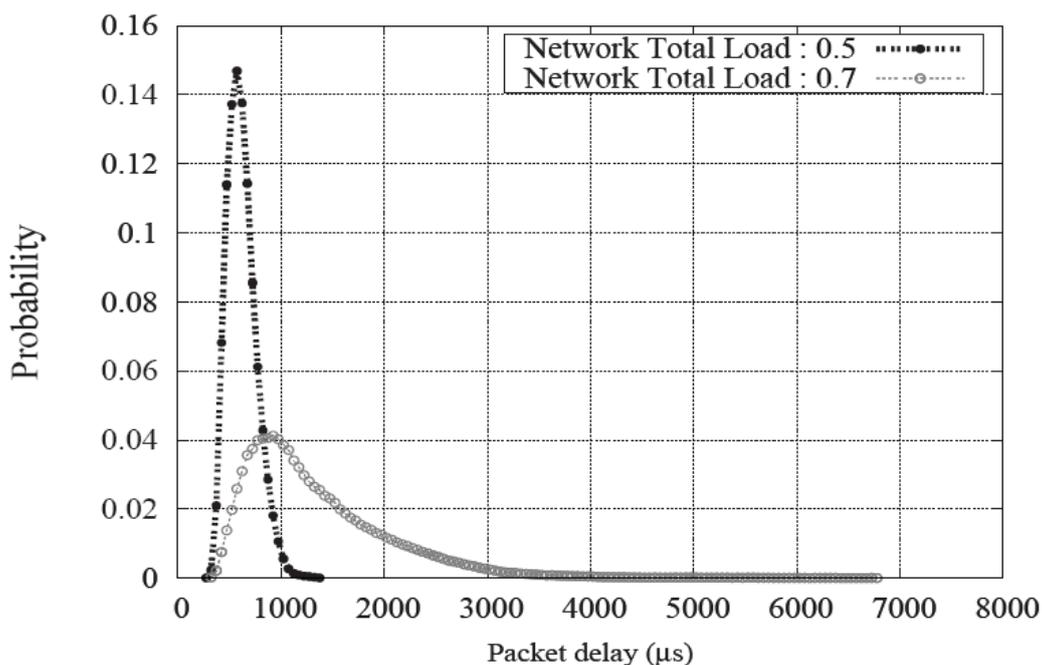
$$B^{avail} = (B^{max} - \sum_{i=1}^N G_i^{T_0}). \quad (3)$$

- Step 3 If there are remaining grants, the low-priority T_2 grant bandwidth $G_i^{T_2}$ is calculated in a similar manner.
- Step 4 The grants computed for each class are sent using the MPCP protocol. Once a ONU receives the gate message, it performs the intra-ONU policy. As argued previously, we assume that each ONU applies exactly the allocated per class grants. In particular, frames that arrived after the previous request cannot be served within the current cycle. As noted before, it can be seen that the total number of grants given to ONU_{*i*} may exceed B_i^{lim} . We show in the next section how this can impact on the QoS of committed bandwidth traffic supported in the system.

4.2 Impact of non-conforming traffic on the performance of the multi-class DBA

We have analysed the performance of the above DBA using a simulation platform developed in (NS-2). The simulation framework we have developed allows us to simulate both the signaling process of the DBA and the actual transfer of data packets in the upstream traffic. In the following, we assess the performance of the DBA algorithm in the case where in the system; there are not only compliant users but also users who violate their negotiated traffic profiles by sending for example more traffic than specified in the traffic profiles. We have chosen to demonstrate the performance degradation by computing the upper quantile of the packet delay (called maximum packet delay in the remainder of the paper). We consider an EPON system with 16 ONUs, and a total uplink bandwidth of 1 Gbps. The maximum cycle time T^{max} is set to 1.5 ms and the guard time t_g between two consecutive ONUs transmissions is 4 μs . The distance between OLT and ONUs is 5–20 kms.

Figure 3 Histogram of the packet delay for the T_0 traffic (all traffic is conforming)



As discussed above, three traffic classes are simulated, namely T_0 , T_1 and T_2 . In our experiments, class T_0 is simulated as a real-time VBR stream which is represented as an ON-OFF model. The duration lengths D_{on}, D_{off} of an ON and OFF period are generated according to an exponential distribution with a same mean of 10 ms. Class T_1 is also represented as an ON-OFF exponential model, ON and OFF average duration periods are both being set to 200ms. In the ON-OFF model, the average rate R_a can be obtained by

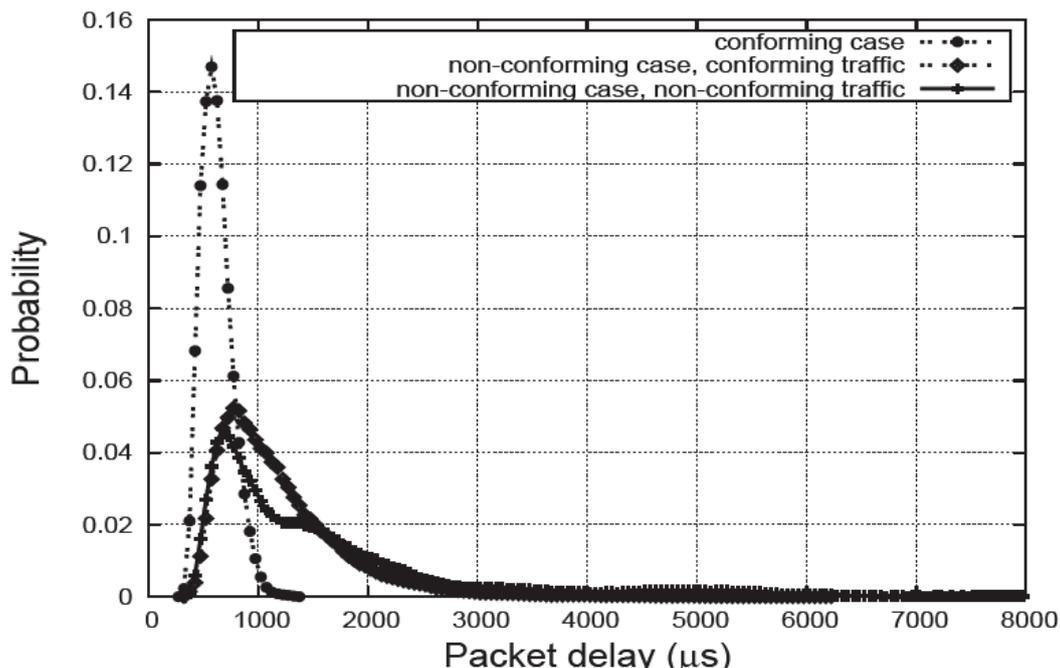
$$R_a = \frac{D_{on}}{(D_{on} + D_{off})} R_p \quad (4)$$

where R_p is the rate at which a source, during the ON state, transmits. Class T_2 is modelled as a BE traffic which is generated following a Pareto distribution or simply as a FTP source. Each ONU may have different traffic profiles specifying T_0 and T_1 traffic flows.

Figure 3 shows the histogram of the packet delay for T_0 traffic in the system in two cases where only T_0 traffic is supported. In the first case, all ONUs send T_0 traffic at the same average rate of 30 Mbps, which results in a total network load of 0.5 while in the second case, all ONUs send only T_0 traffic at the same average rate of 45 Mbps which results in a total network load of 0.7.

Results such as shown in Figure 3 have to be considered when designing a multi-class connection acceptance control (CAC) for the EPON. Indeed, we see that if the offered load is 0.7, the maximum packet delay for T_0 traffic may exceed a 1.5 ms limit (ITU-T Recommendation G.114), whereas at 0.5, most of the delays are well below 1.5 ms. Discussing a multi-class CAC for the above DBA is outside the scope of the present paper. However, for the relevant discussion here, we assume in the following that the CAC mandates to limit the accepted load of T_0 traffic to 0.5.

Figure 4 Comparison of T_0 packet delay in conforming and non-conforming case (non conformance control)



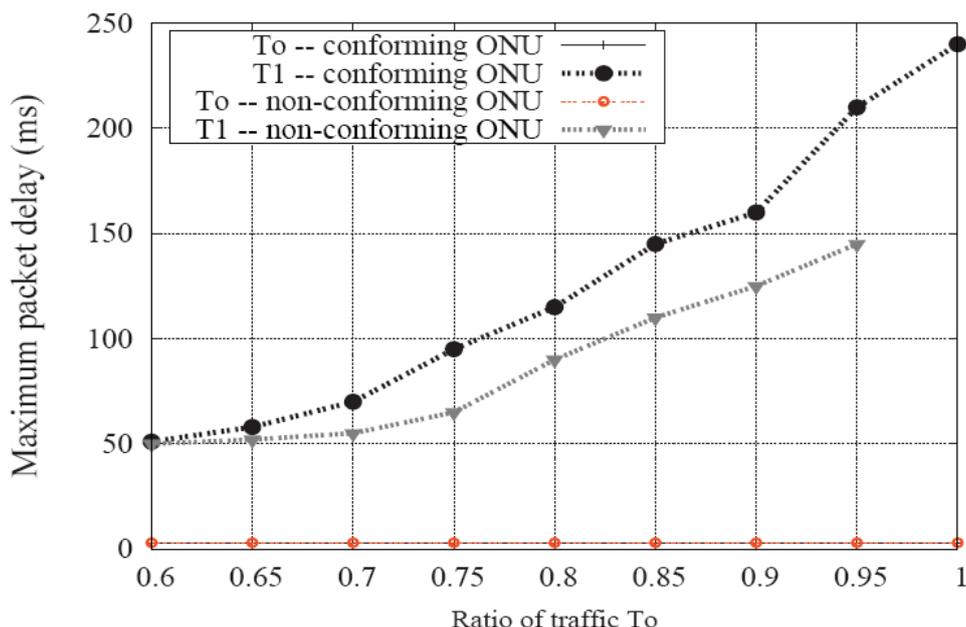
We now introduce some non compliant ONUs. We define here a non compliant ONU as an ONU whose R_a of T_0 traffic is larger than 30 Mbps (due to a larger R_p). Figure 4 shows

the histogram of the packet delay for conforming T_0 traffic in the system in two cases. In the first case, all ONUs send only T_0 traffic at the same average rate of 30 Mbps. In the second case, four ill-behaved ONUs send an average rate of 80 Mbps of T_0 traffic while the 12 remaining ONUs are compliant.

Figure 4 shows that in the first case (compliant ONUs only), all delays are below 1.5 ms whereas the delay distribution presents a long tail up to 6 ms in the second case even for compliant ONUs. This clearly shows the impact of non compliant ONUs on the performance delivered to compliant ONUs: both the latency and the packet jitter for conforming T_0 are seriously degraded when several ONUs overload the system with non-conforming high priority traffic.

We now investigate another scenario where both T_0 and T_1 traffic is supported. In our study, each ONU sends T_0 and T_1 traffic and the average rate of each ONU is kept constant and equal to 50 Mbps. In the nominal case (compliant ONUs) 60% of this load is offered by T_0 traffic (i.e. 30 Mbps) and this ratio is kept constant during the simulation. Non-compliant ONUs vary the ratio of T_0 traffic between 60% (i.e. 30 Mbps) and 100% (i.e. 50 Mbps). We assume that there are eight compliant and eight non-compliant ONUs. Figure 5 shows the maximum packet delay performance of conforming and non-conforming traffic versus the ratio of T_0 traffic sent by malicious ONUs. We see in Figure 5 that the performance offered to T_0 traffic, either for conforming or non-conforming ONUs is not degraded since T_0 traffic has a high priority in the DBA. On the other hand, non-conforming T_0 traffic has a negative impact on T_1 traffic, for both conforming and non-conforming ONUs since it uses more resources than expected. A side effect is that the T_1 performance is actually less degraded for non-conforming ONUs due to the fact that less T_1 traffic is offered by these ONUs. The simple scenarios in this section indubitably show that a DBA algorithm that does not take into account traffic conformance when allocating transmission opportunities cannot prevent non-compliant ONUs from affecting the performance of compliant ONUs. This observation led us to integrate conformance control aspects within the DBA in order to allow the OLT to isolate compliant ONUs from non-compliant ONUs and thus fulfill SLA requirements.

Figure 5 Impact of non-conforming T_0 on QoS for T_1 (see online version for colours)



5 Integrating conformance control within a DBA

In this section, we propose a strategy that merges our priority based DBA with a mechanism that checks conformance of the upstream traffic. Our primary goal is to prevent non-compliant traffic from affecting the QoS of compliant traffic. In order to achieve this goal, we designed a method that checks traffic (i.e. committed rate and burst size) before allocating grants.

5.1 Design principles

The basic idea of our mechanism is fairly simple: instead of implementing a policing mechanism (such as a token bucket filter) for upstream traffic at the OLT, the OLT filters the demands carried in report messages and takes account of a virtual policing scheme when computing the number of grants to be sent to the ONU. Checking conformance within the DBA does not preclude implementing an active policing mechanism within the OLT e.g. to check that packet marking is performed according to the allocated grants when upstream traffic reaches the OLT. The advantage of merging conformance control with the DBA instead of implementing it in the ONU is the following: if the ONU filtered its demands, and requested grants for example only for conforming traffic, the OLT could not take advantage of unused transmission opportunities in order to satisfy excess requests. On the other hand, a centralised conformance control plus DBA algorithm can optimise grant allocation on CB traffic, excess (non-conforming) traffic and BE traffic according to a policy specified by the network operator.

Figure 6 Proposed architecture for the conformance checking DBA

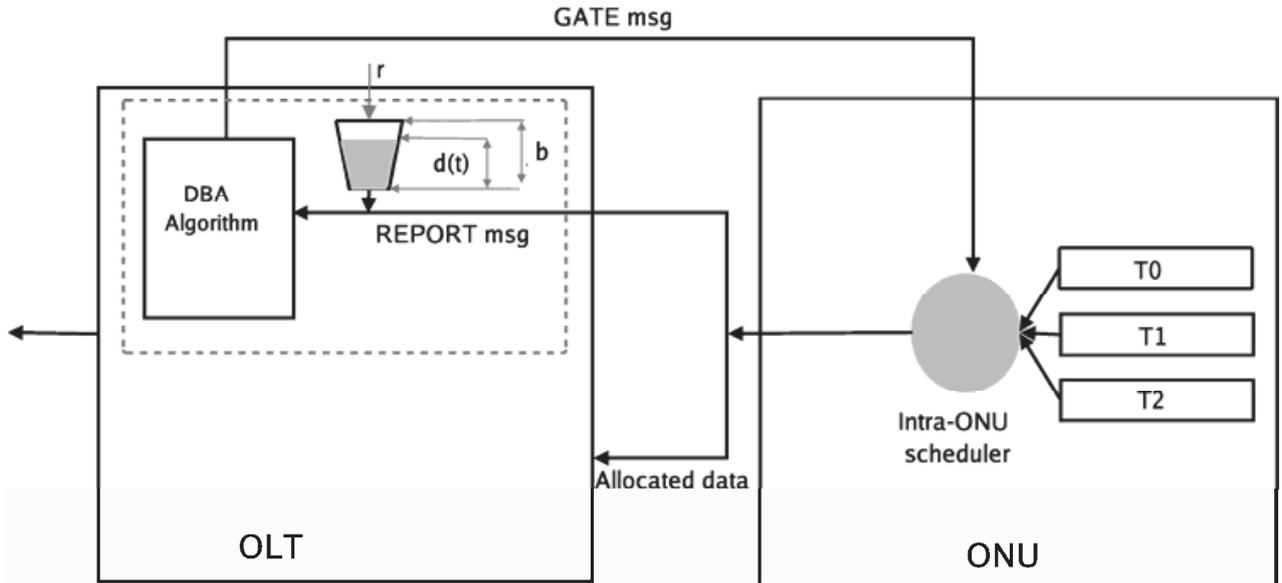


Figure 6 shows our proposed architecture for a conformance checking DBA that includes a traffic policing module (TPM) that checks the conformance of requests carried in the report messages. For ONU_i , the TPM controls the T_0 traffic flow by means of tokens which are generated at the rate of r_i tokens per second. The tokens are stored in a bucket of finite size b_i . The goal is to regulate the traffic T_0 in such a way that over long time periods the average permitted rate approaches the desired token rate r_i asymptotically and

over short time periods the amount by which the data rate can exceed r_i is upper bounded by the bucket size b_i . Each pair of values b_i and r_i represents the T_0 traffic profile.

Before reaching the DBA algorithm module in Figure 6, a request $R_i^{T_0}$ has to find an equivalent number of tokens (in bytes) in the bucket in order to be considered as conforming.

$$R_i^{T_0} = \begin{cases} R_i^{T_0} & \text{if } R_i^{T_0} < d_i(t) \\ d_i(t) & \text{otherwise} \end{cases} \quad (5)$$

The number of available tokens $d_i(t)$ in bucket i is then decremented by $G_i^{T_0}$, which corresponds to the grant bandwidth allocated for the i th ONU after transmission opportunities are computed by the OLT. A similar conformance process is also implemented for T_1 traffic. When some demands are in excess of the negotiated traffic profile, the OLT could take one or more actions depending on the traffic management strategy desired by the network operator:

- 1 Allocate: although the traffic is out of profile, grants are allocated to excess traffic
- 2 Suggest-discard: in some cases (e.g. EF) traffic profiles have to be strictly respected and enforced
- 3 Suggest-mark: Non-conforming packets are allowed into the network, but are marked. How these marked packets are later treated depends on traffic and congestion control strategies in the OLT and other parts of the network.
- 4 Suggest-buffer: all non-conforming packets are to be buffered at the ONU until enough tokens are accumulated.

In order to illustrate the efficiency of our proposal, we now show the behaviour of a system where excess traffic is buffered at the ONU.

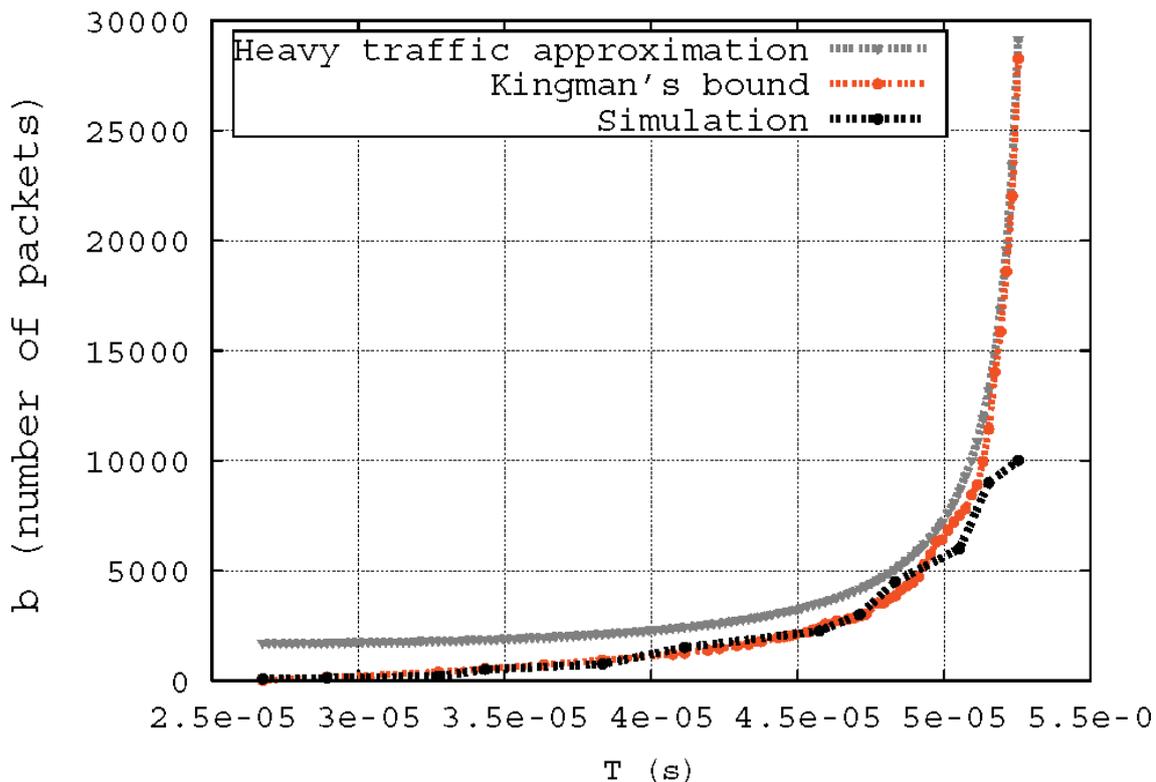
5.2 How to determine token bucket parameters

There are many methods allowing token bucket parameters to be computed according for example to the statistics of application generated traffic. The parameters can also be mandated by the network operator as part of its global traffic management strategy. In this paper, since we assess the performance of a DBA on synthetic ON-OFF traffic, we need a method to compute the values of r and b , parameters of the token bucket.

We have used two approaches described in Gravey et al. (1996). The first one, namely Kingman's bound, is based on estimates for the tail behaviour of the virtual waiting time in the GI/G/I queue obtained by Kingman (1970) and extended by Ross (1974). The second one, namely heavy traffic approximation, is based on a standard heavy traffic approximation for the virtual waiting time in the GI/D/I queue. Figure 7 shows the accuracy of these approaches. Note that T is calculated here as follows:

$$T = \frac{8pkts}{r}, \text{ where } pkts \text{ is the packet size} \quad (6)$$

Figure 7 b vs T in the case where $D_{on} = D_{off} = 0.01s$ and $R_p = 60$ Mbps (see online version for colours)



5.3 QoS delivered to CB traffic in case of non conformance

To verify the effectiveness of our proposed scheme, we repeat the experiments chosen in Section 4.2. We first assume that for each ONU, a pair of values of $R_a=30$ Mbps and $R_p=60$ Mbps are negotiated for T_0 traffic. A set ($r = 32$ Mbps; $b = 7.2$ Mbits) is determined using the above method and is used in our experiments.

Figure 8 first shows the results of the scenario already considered in Figure 4 in which we replace the simple DBA by the conformance checking DBA. As we can see, the QoS offered to compliant T_0 traffic is now identical to the QoS offered to all ONUs when all are compliant. On the other hand, non-compliant ONUs now receives a degraded QoS for their T_0 traffic. This is because conformance is checked before transmission opportunities are allocated. Therefore, the excess T_0 traffic is not considered by the DBA which can then safely allocate transmission opportunities to conforming T_0 traffic only.

Figure 9 shows the results of the scenario considered in Figure 5 in which we replace the simple DBA by the conformance checking DBA. Once again, we can observe that the committed QoS is indeed delivered to a compliant ONU, thanks to the isolation from non-compliant ONUs offered by our conformance checking DBA policy.

Figure 8 Comparison of T_0 packet delay in conforming and non-conforming case (conformance control applied) (see online version for colours)

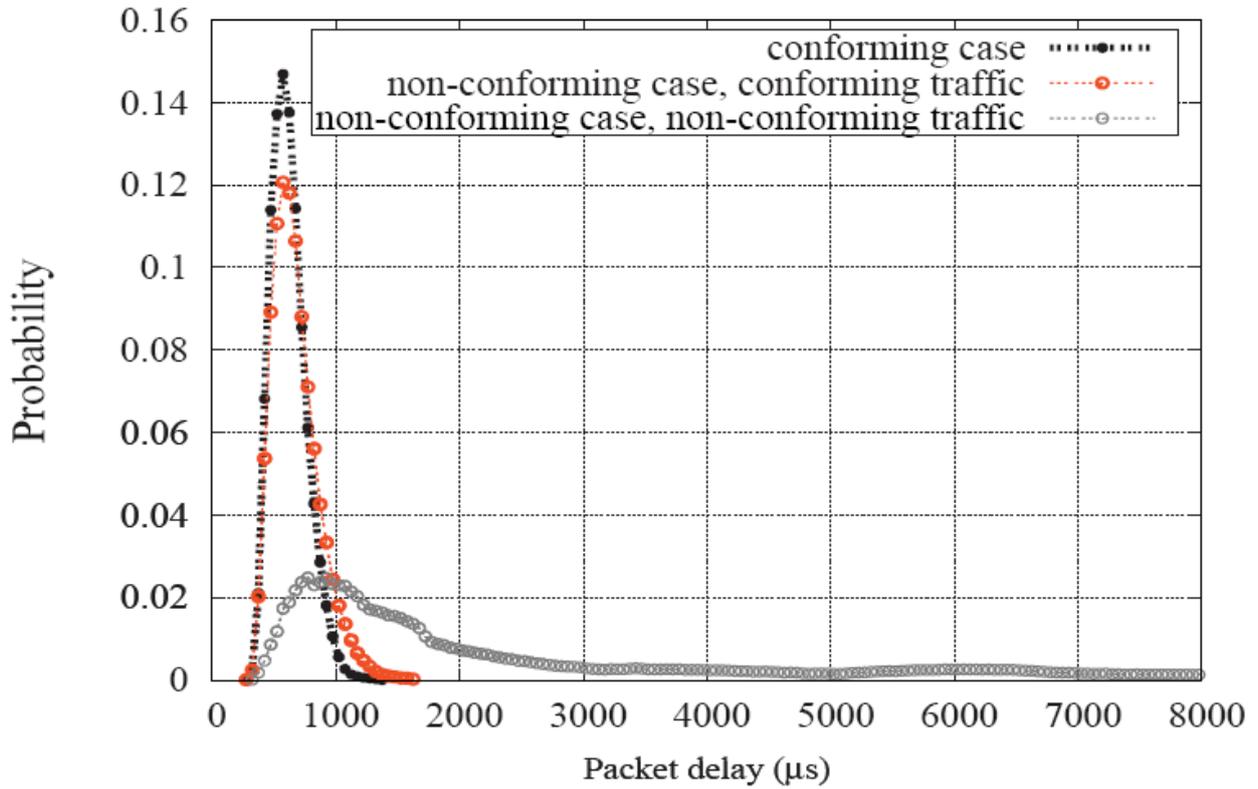
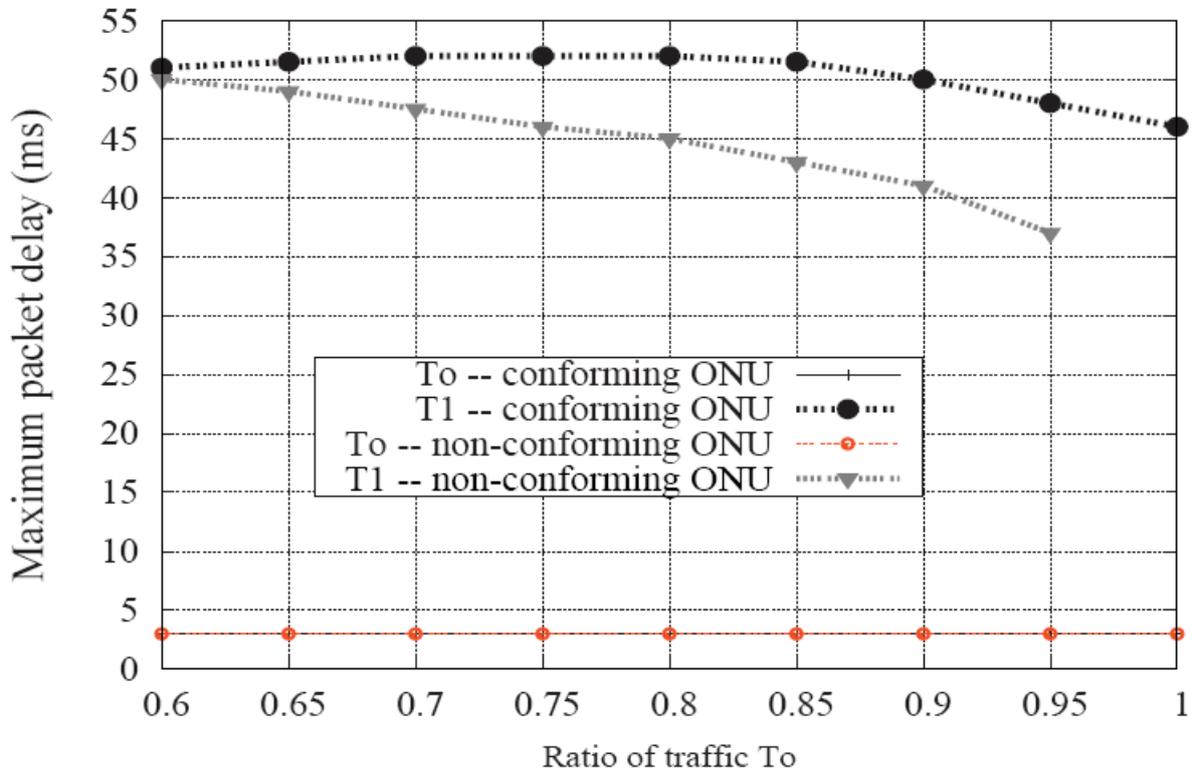


Figure 9 Maximum packet delay performance for conforming and non-conforming traffic (see online version for colours)



5.4 Upstream channel utilisation

As shown in Section 5.3, the performance delivered to the non-compliant ONUs is very degraded when the system only accepts the negotiated traffic rate for each ONU. However, since the OLT knows all requests and is able to check conformance, it could allocate more transmission opportunities than specified by the strict application of the conformance test if there are enough available resources (due for example to idle ONUs), and if the operator wishes to increase upstream channel utilisation. In order to archive this goal, the OLT could for example take the following actions:

- the OLT allocates transmission opportunities first to conforming T_0 and T_1 traffic
- T_2 traffic gets grants if there are remaining grants after serving all conforming T_0 and T_1 traffic demands
- if there are remaining grants, the OLT then handles non-conforming T_0 and T_1 traffic.

The above policy is only one simple example of a policy applied by the access network operator.

Figure 10 Upstream channel utilisation with both BE and conforming traffic

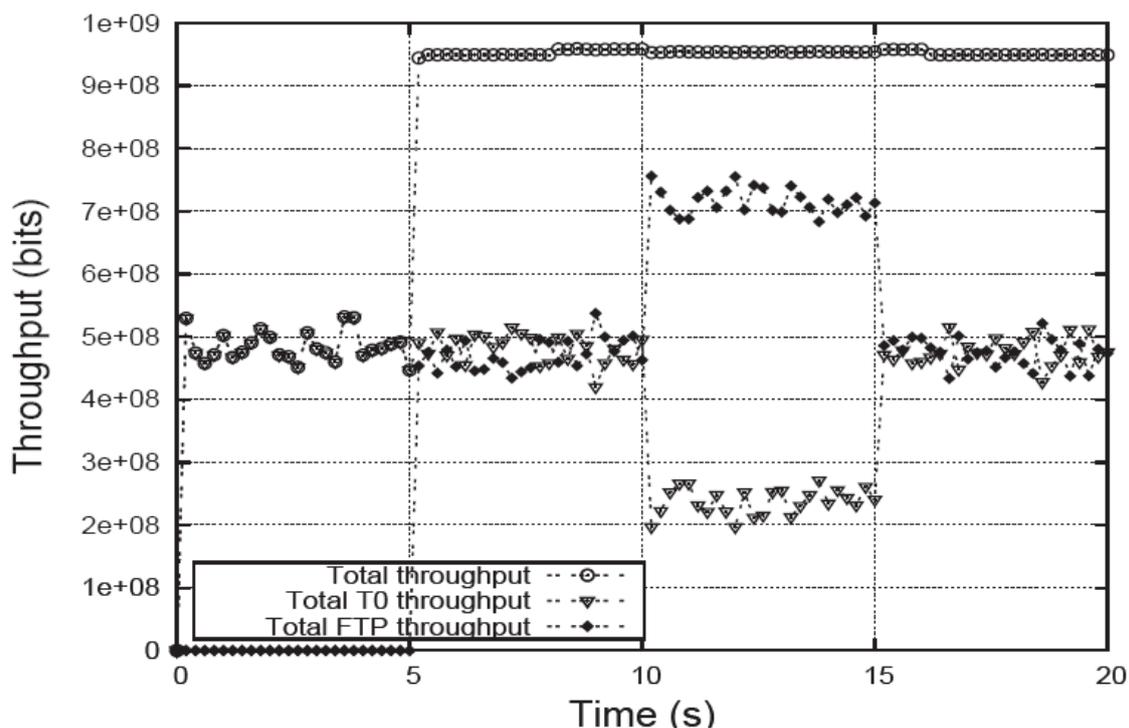
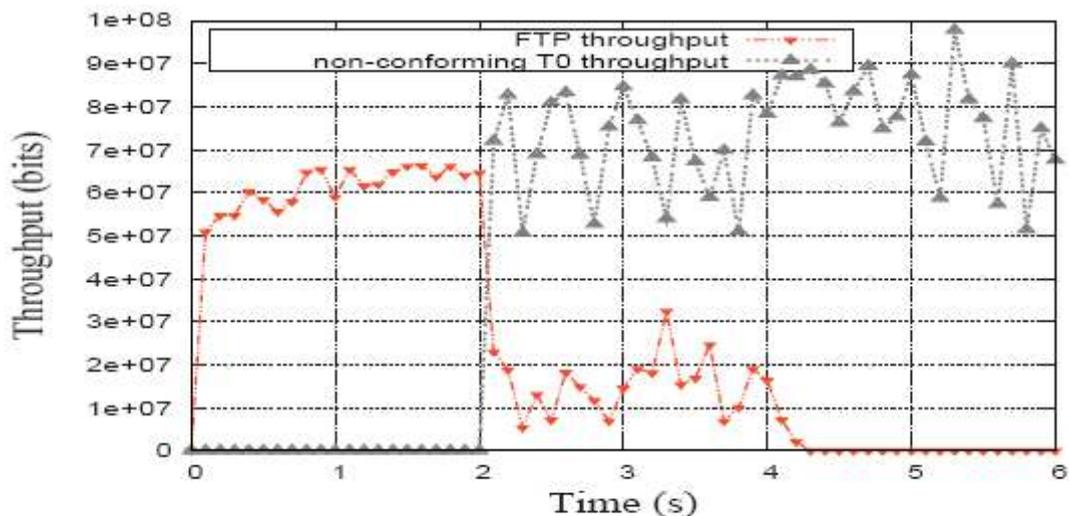


Figure 10 first shows the performance in terms of throughput in the upstream direction when there is only compliant T_0 traffic and BE traffic. In this particular scenario, 8 ONUs always send T_0 traffic during the whole simulation time while the T_0 sources in eight other ONUs are active only from 0.0s to 10.0s and from 15.0s to 20.0s. Each ONU also starts an FTP session at 5.0s. Figure 10 shows that, as soon as the FTP sources are activated, the upstream channel is almost full. Indeed, FTP traffic grabs all available bandwidth (see the throughput achieved between 10.0s and 15.0s). However, as long as T_0 traffic is sent, it receives the expected throughput, whether or not FTP sources are activated. These results clearly show the performance of our approach in terms of

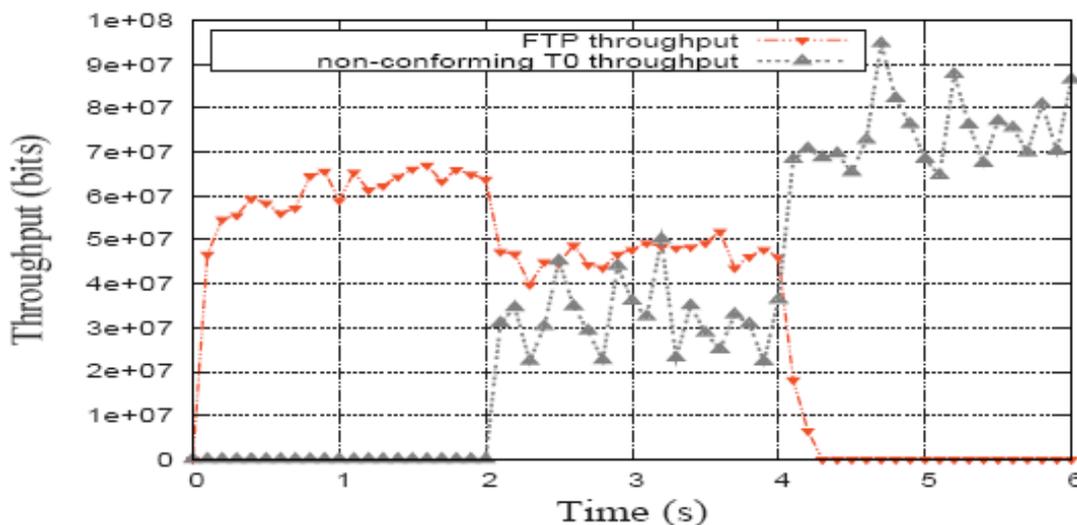
delivering high upstream channel utilisation while ensuring the QoS requirements of CB traffic.

We now investigate scenarios where conforming and non-conforming T_0 traffic are sent, together with T_2 traffic. In these experiments, each of 12 conforming ONUs sends an FTP source and conforming T_0 traffic (i.e. 30Mbps on average) while each of four other non-conforming ONUs sends an FTP source and T_0 at a rate of 70 Mbps; therefore, part of the latter traffic is non-conforming. For each non-conforming ONU, T_0 traffic is active between 2.0s and 6.0s. For every ONU, the FTP source is active between 0.0s and 4.0s. Figure 11b shows the throughput performance for T_0 traffic versus simulation time. Here, we can see that, between 2.0s and 4.0s, T_0 traffic from non-conforming ONUs is enforced to a conformance rate of roughly 30 Mbps while between 4.0s and 6.0s; the T_0 traffic is allowed to be served at a higher rate (70 Mbps in average) due to the fact that the FTP sources are idle. In Figure 11a, non-conforming T_0 traffic is served before T_2 traffic, which is then partially starved. This shows the impact of the selection of the policy implemented in the DBA relative to the support by the network operator of non-conforming traffic.

Figure 11 Throughput of BE and non-conforming traffics with different policies (see online version for colours)



(a)



(b)

6 Conclusions

In this work, we have indicated the need to merge a multi-class DBA with a policy strategy enforcing traffic profile. We have first shown that it is not realistic to provide uncontrolled access to the system if applications with stringent delay and loss requirements are to be supported. Moreover, we have also shown that even within the set of committed bandwidth classes, it is important to check conformance in order to allow the lower class to get its negotiated QoS in the presence of a higher than negotiated ratio of high priority traffic.

We have thus proposed to integrate conformance checking into the DBA process. This strategy is shown to prevent non-conforming traffic from affecting the performance of conforming traffic. This has been illustrated via a simple multi-class, priority based DBA that has been designed for EPONs. Finally, we have shown that our modified DBA fully protects conforming traffic and thus allows a compliant ONU to be isolated from the impact of non-compliant ONUs requesting more transmission opportunities than specified in their negotiated SLAs. Centralising the conformance checking and DBA computation in the OLT allows the allocation of transmission opportunities to ONUs to be optimised. No policy local to the ONU (intra-ONU scheduling) would be able to fulfil this goal as efficiently.

This approach exemplifies the relationship between the management plane (responsible for the SLA and the network operator strategy) the command plane (responsible for the DBA) and the transfer plane (actual use of transmission opportunities). It is well known that these planes are inter-related and have to be synchronised in order to implement a QoS and performance management strategy. However, in the PON case, this relationship is more intricate than in point-to-point architectures due to the polling method implemented for supporting upstream traffic.

In this paper, we also show the advantage of centralising both conformance checking and DBA. Indeed, since the OLT knows all requests and is able to check conformance, it can allocate more transmission opportunities than specified by the conformance if there are enough available resources (due for example to idle ONUs). This would not be possible if conformance was checked within the ONU, and only conforming requests were transmitted in the REPORTS (a typical customer edge shaping scenario). The polling mechanism allows us to adapt the allocation process to the number of effectively available transmission opportunities, while protecting conforming traffic from non compliant behaviours and thus supporting contractual (SLA based) QoS guarantees.

Acknowledgements

This work was partially funded by France Telecom in Research Contract 46134298 on Optical Access Networks.

References

- An, F. and Hsueh, Y. (2003) 'A new dynamic bandwidth allocation protocol with quality of service in Ethernet-based passive optical networks', *International Conference on Wireless and Optical Communication (WOC)*, Banff, Canada.
- ANSI/IEEE Standard 802.1D (1998) *IEEE Standard for Information Technology, Telecommunications and Information Exchange between Systems, Local and Metropolitan Area Networks, Common Specifications, Part 3: Media Access Control (MAC) Bridges*, <http://standards.ieee.org/getieee802/download/802.1D-1998.pdf>.
- Assi, C.M., Ye, Y. Dixit, S. and Ali, M.A. (2003) 'Dynamic bandwidth allocation for quality-of-service over Ethernet PONs', *IEEE J. Sel. Areas Commun*, Vol. 21, No. 9, pp.1467–1477.
- Choi, S. (2004) 'Cyclic polling based dynamic bandwidth allocation for differentiated class of service in Ethernet PONs', *Photonic Network Communications*, Vol. 7, No. 1, pp.87–96.
- FSAN Full Service Access Network group, <http://www.fsanweb.org/>
- Gravey, A., Romoeuf, L., Sevilla, K. and Blaabjerg, S. (1996) 'Models for setting ATM parameter values', *Telecommunication Systems*, Vol. 5, No. 1, pp.85–107.
- Green, P.E. (2004) 'Fiber to the home: the next big broadband thing', *IEEE Communications Magazine*, Vol. 42, No. 9, pp.100–106.
- IEEE 802.3ah (2004) *Ethernet in the First Mile Task Force IEEE Std 802.3ah*, <http://www.ieee082.org/3/efm/>.
- ITU-T Recommendation G.114 *One-way Transmission Time*, <http://www.itu.int/itudoc/itu-t/aap/sg12aap/history/g.114/>.
- ITU-T Recommendations, <http://www.itu.int/ITU-T/publications/recs.html/>.
- Kim, J.S. (2006) 'HUHG: high utilization and hybrid granting algorithm for EPON', in *Proc. AICT-ICIW'06*, Guadeloupe, France.
- Kingman, J.F.C. (1970) 'Inequalities in the theory of queues', *Journal of the Royal Statistical Society, Series B*, Vol. 32, pp.102–110.
- Kramer, G., Mukherjee, B. and Pesavento, P. (2002) 'Interleaved polling with adaptive cycle time (IPACT): a dynamic bandwidth distribution scheme in an optical access network', *Photonic Network Communications*, Vol. 4, No. 1, pp.89–107.
- Ma, M. and Zhu, Y. (2003) 'A bandwidth guaranteed polling MAC protocol for Ethernet passive optical networks', in *Proc. IEEE INFOCOM*, San Francisco, CA, USA, Vol. 1, pp.22–31.
- McGarry, M.P., Maier, M. and Reisslein, M. (2004) 'Ethernet PONs: a survey of dynamic bandwidth allocation (DBA) algorithms', *IEEE Communications Magazine*, Vol. 42, No. 8, pp.S8–15.
- NS-2, *The NS Manual*, <http://www.isi.edu/nsnam/ns/nsdocumentation.html>.
- Ross, S.M. (1974) 'Bounds on the delay distribution in GI/G/I queues', *Journal of Applied Probabilities*, Vol. 11, pp.417–421
- Zheng, J. and Mouftah, H.T. (2005) 'Media access control for Ethernet passive optical network: an overview', *IEEE Communications Magazine*, Vol. 43, No. 2, pp.145–150.