

Controlling QoS in EPON-based FTTX access networks

M. Thanh Ngo · Annie Gravey · Deepak Bhadauria

Published online: 26 May 2010

© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract The Ethernet PON (EPON) is viewed by many as an attractive solution to deliver very high-speed broadband access and is widely deployed in some geographical areas. While downstream traffic is broadcast to all customers, the access of upstream traffic to the fiber has to be arbitrated in order to avoid collisions. This arbitration mechanism and more generally, bandwidth distribution and QoS Provisioning, have been left to the implementer. One solution is to enforce static Time Division Multiplexing Access (TDMA) between end-users. This however precludes an efficient usage of resources. Interleaved Polling with Adaptive Cycle Time (IPACT) is one of the earliest proposed schemes for Dynamic Bandwidth Assignment (DBA) in EPON and has been extensively used as a benchmark by many subsequent allocation schemes. In this paper, we first propose an analytical model which yields approximate values for mean queue length and mean packet delay in an EPON using IPACT with Gated Service (GS) scheme under the assumption of heterogeneous Poisson arrivals. We use the model to demonstrate that all users experiment performance degradation in

case of local overload, thus showing the necessity of correcting somehow IPACT-GS in order to avoid this phenomenon. This is achieved by designing a control plane for EPON, which includes a priority based DBA together with a framework for enforcing Service Level Agreements (SLAs), and fairly sharing available resources. The proposed framework is easily configured (all the control being centralized at the OLT or in the backbone) while allowing the support of large varieties of services. It is shown to behave more efficiently than other QoS sensitive DBAs in the literature.

Keywords Optical access · Ethernet PON (EPON) · Dynamic bandwidth allocation (DBA) · Quality of service (QoS) · Traffic conformance · Mean value analysis (MVA) · Interleaved polling with adaptive cycle time (IPACT)

1 Introduction

Broadband Passive Optical Network (PON) is one of the most deployed technologies for delivering fiber based (FTTX) access networks [9]. The Ethernet PON (EPON), standardized by the IEEE Ethernet First Mile Alliance [8], is viewed by many operators as an attractive solution to deliver very high-speed broadband access.

An EPON architecture consists in an Optical Line Termination (OLT), located at the central office, and multiple remote Optical Network Units (ONUs) that deliver broadband voice, data, and video services to subscribers. Downstream (i.e., from the network to the customers) and Upstream (i.e., from the customers to the network) traffics use two different wavelengths on a single fiber. Due to the topology of PON, the transmission modes for downstream and upstream traffics are different. While downstream traffic is broadcast to all customers, the access of upstream traffic to the fiber has to be arbitrated in order to avoid collisions. One simple

This work has been partially funded by the French National Research Agency in the ECOFRAME project and by France Telecom in Research Contract 46134298 on Optical Access Networks. It was initiated while Deepak Bhadauria was doing a research internship at Telecom Bretagne.

M.T. Ngo (✉) · A. Gravey · D. Bhadauria
Institut TELECOM, TELECOM Bretagne,
Department of Computer Science, Technopole Brest Iroise
CS 83818, 29238 Brest, France
e-mail: mt.ngo@telecom-bretagne.eu

A. Gravey
e-mail: annie.gravey@telecom-bretagne.eu

M.T. Ngo · A. Gravey · D. Bhadauria
Université Européenne de Bretagne, Rennes, France

mechanism is Time Division Multiplexing Access (TDMA) in which every ONU gets a fixed time slot [11] to transmit upstream traffic. While TDMA is very simple and well suited for supporting balanced and permanent traffic offers, its drawback is an obvious lack of support for highly bursty sources, or unbalanced demands. In fact, in this type of situation, 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 (e.g. a node failure, a new node enters the network, etc.).

Many mechanisms that allow a dynamic allocation of transmission opportunities have been proposed in the literature [15]. This is achieved by designing both a Dynamic Bandwidth Assignment (DBA) that shares bandwidth opportunities between customers, and an intra-ONU scheduling policy, which is used by customers to take advantage of the allocated bandwidth opportunities. The DBA and the settings of the intra-ONU scheduling policy are part of the control plane of the PON, and thus characterize which types of services can be supported by the operator. Several DBAs have been proposed for EPONs, which differ in the way grants are computed. They also differ in the performance delivered to traffic streams sent by ONUs.

Due to the high complexity of network behavior introduced by the DBA, EPON systems are usually analytically intractable. We propose here an analytical approach applicable to one of the most popular DBAs (IPACT, with Gated Service). We use this model to assess the QoS delivered to the users in terms of mean transmission delay. We then demonstrate that all users experience performance degradation in case of local overload, which indicates that IPACT is only suitable to deliver a Best Effort service.

However, the offer for broadband access is not identical over all network operators: some operators offer only a best effort access; others operators already support streaming services such as Video on Demand, and interactive services such as VoIP. In the near future, it is fairly likely that some operators would like to support Committed Bandwidth services (especially for SME or SOHO customers) similar to the well known Frame Relay services that specified a Committed Information Rate (CIR) and an Excess Information Rate (EIR). Other could wish to support several Classes of Services, implementing either relative priorities between classes or enforcing some kind of access control on some Classes. An operator could also wish to support different types of services for different markets (e.g. best effort access for the residential market, but committed bandwidth services for SME/SOHO market). While IPACT efficiently supports a best effort service, a more complex strategy is required to support all these types of services.

In this paper we present a versatile control plane for EPONs that can be configured to support either a class-

less Best Effort network, or several relative Classes of Services, or a mix of committed bandwidth and Best Effort services. This control plane includes a framework for enforcing Service Level Agreements, and fairly sharing available resources. A central element for this control plane is a simple priority based upstream traffic management policy which is shown to be able to support several types of broadband access service offers. The proposed policy can easily be re-configured if the service offer evolves over time. For several typical scenarios, we identify which features should be configured and show how our control plane delivers the expected behavior.

The rest of the article is organized as follows. Section 2 provides some background knowledge on the analytical analysis of IPACT performance. Section 3 describes our analytical model for computing mean transmission delay in an EPON operated with IPACT and gated service. Section 4 addresses the impact of heterogeneous users on the PON and also validates by simulation the approximate analytical model. Our contribution regarding a control plane for EPON is given in Sect. 5. Section 6 assesses how the control plane can be configured to support various access network services. The last section concludes the paper.

2 Background knowledge

The IPACT scheme has been proposed by Kramer et al. in [11]. IPACT requires the OLT to poll ONUs individually in a round-robin fashion and to dynamically assign transmission opportunities. These opportunities are allocated according to the buffer occupancy status reported by each ONU. IPACT is based on interleaving polling messages in order to minimize switchover times.

Several service disciplines, i.e., ways for the OLT to determine the transmission window (TW) size $TW_i(n)$ (bytes) for ONU_i in cycle n , have been proposed. This transmission window $TW_i(n)$ (bytes) is computed according to the window $Q_i(n-1)$ (in bytes) requested in the cycle $(n-1)$:

- **Fixed Service scheme:** This scheduling discipline ignores the requested window size and always grants the maximum transmission window.
- **Gated Service scheme:** In this discipline, the OLT always grants each ONU the amount of bytes that this ONU requested in the previous cycle:

$$TW_i(n) = Q_i(n-1)$$

- **Limited Service scheme:** In this discipline, the OLT grants the requested number of bytes, but no more than a maximum threshold TW^{\max} :

$$TW_i(n) = \min(Q_i(n-1), TW^{\max})$$

In the present paper, we focus on the Gated Service discipline, which is more tractable than others due to its similarity with polling systems.

2.1 Mean value analysis for polling systems

Polling systems have been researched for a number of years. The interested reader can find a comprehensive survey in [22].

A recent paper [23] proposes a novel approach based on Mean Value Analysis (MVA) to compute the mean delays for Exhaustive or Gated polling systems. More specifically, the authors have derived a set of linear equations for these delay figures in the case of Exhaustive and Gated Service, respectively, with the help of the following two basic queueing results: the PASTA property, i.e., Poisson Arrivals See Time Averages [24] and Little's Law [13].

The transmission of upstream traffic in EPONs can be modeled by a polling system: each ONU is cyclically "served" (i.e. uses the transmission grants) to empty its queue. However, the results in [23] cannot be directly used. Indeed, the gated polling service analyzed in [23] differs from the IPACT gated service. In [23], all packets waiting at the start of the server's visit to this queue are served in the current cycle. In the IPACT EPON system, grants cannot be used for packets that were not already present in the system at the end of the previous serving period. This means that the set of equations used here differs from that in [23].

2.2 Analytical models for EPONs

We now review some recent papers dealing with the Gated Service policy for EPON. A good performance analysis study of EPON network is proposed in [19]. The authors obtain results regarding the performance of a gated interleaved polling system with a two-stage queue.

S. Bhatia et al. in [5] developed a recursive model for the Gated Service IPACT scheme suggested in [11]. Based on this recursive model the authors derived closed form expressions for the average queue size at the ONU, for both single-ONU and multiple-ONU network setups. The authors in [4] developed an analytical framework to explore the performance characteristic of EPONs with a quasi-leaved polling operation and gated resource allocation policy. The derived analytical expression by this approach is useful to evaluate the network performance for memoryless traffic inputs, as well as to closely estimate the performance of lightly loaded networks for bursty input traffic profiles.

However, the above references deal only with the symmetric case. Recently, the analytical model developed by B. Lannoo et al. in [12] allows an approximation for the mean packet delay performance in IPACT Fixed and Gated Service to be computed in both symmetric and asymmetric cases.

In this paper, we present a more exact method for computing the mean queue length and the mean queueing delay in an IPACT system with memoryless, symmetric and asymmetric traffic. We use the MVA approach shown in [23] for polling systems, by adapting the method to the queueing process specific to a Gated IPACT system, which differs from the queueing process in a polling system.

3 Analytical model for IPACT

3.1 The queueing model

Consider a polling system with $N \geq 1$ queues (which correspond to $N \geq 1$ ONUs in PON terminology), indexed by $i = 1, 2, \dots, N$, which are served by a single server in a cyclic manner. Each queue has an infinite buffer. Packets arrive at all queues according to independent Poisson processes with rates λ_i , $i = 1, 2, \dots, N$. Packets in a queue are served in the order of their arrival, and packet service times at queue i are independent, with mean $E[B_i]$ and second moment $E[B_i^2]$.

The *setup time* of queue i is the time taken by the server to move from queue i to queue $i + 1$ (i.e. the time interval between the moment when transmission at queue i stops and the moment when transmission at queue $i + 1$ starts). This setup time models the EPON guard time. Its characteristics are the mean $E[S_i]$, and the second moment $E[S_i^2]$. The mean total setup time $E[S]$ in a cycle is given by

$$E[S] = \sum_{i=1}^N E[S_i] \quad (1)$$

The derivation also uses the mean residual service time and the mean residual setup time for queue i which can be expressed as follows

$$E[R_{B_i}] = \frac{E[B_i^2]}{2E[B_i]} \quad (2)$$

$$E[R_{S_i}] = \frac{E[S_i^2]}{2E[S_i]}, \quad i = 1, 2, \dots, N$$

The occupation rate of the system ρ_i (excluding setup times) at queue i is defined by $\rho_i = \lambda_i E[B_i]$ and the total occupation rate ρ is given by $\rho = \sum_{i=1}^N \rho_i$. A necessary condition for stability of this system is $\rho < 1$.

For a stationary system, the cycle length C , which is the time between two successive arrivals of a server at a queue does not depend on the queue. Its mean value is obviously the solution to the following equation

$$E[C] = \sum_{i=1}^N (\rho_i E[C] + E[S_i]) = \rho E[C] + E[S] \quad (3)$$

which yields $E[C] = \frac{E[S]}{1-\rho}$. However, in EPON, the cycle time cannot become lower than a minimum value which is defined as the maximum Round Trip Time (RTT) [11] between the OLT and each ONU. So, we choose to take account of this constraint by approximating:

$$E[C] = \begin{cases} \frac{E[S]}{1-\rho}, & \text{if } \frac{E[S]}{1-\rho} > RTT \\ RTT, & \text{otherwise} \end{cases} \quad (4)$$

Let the visit time θ_i of queue $i, i = 1, 2, \dots, N$ be defined as the service period of queue i , plus the setup time of queue $i - 1$. The mean visit period of queue i is:

$$E[\theta_i] = \rho_i E[C] + E[S_{i-1}], \quad i = 1, 2, \dots, N \quad (5)$$

Define an (i, j) -period $\theta_{i,j}$ as the sum of j consecutive visit times starting in queue $i, j = 1, 2, \dots, N$. The corresponding mean is given by:

$$E[\theta_{i,j}] = \sum_{k=i}^{i+j-1} E[\theta_k], \quad i = 1, 2, \dots, N, \quad j = 1, 2, \dots, N \quad (6)$$

Notice that in the case $j = 1$ and $j = N, E[\theta_{i,j}]$ is equal to the mean visit period $E[\theta_i]$ of queue i and the mean cycle length $E[C]$, respectively. The fraction of time $q_{i,j}$ the system is in (i, j) -period equals:

$$q_{i,j} = \frac{E[\theta_{i,j}]}{E[C]}, \quad i = 1, 2, \dots, N, \quad j = 1, 2, \dots, N \quad (7)$$

where $q_{i,N}$ equals 1 since $E[\theta_{i,N}]$ is equal to the mean cycle length $E[C]$.

The mean residual (i, j) -period is given by:

$$E[R_{\theta_{i,j}}] = \frac{E[\theta_{i,j}^2]}{2E[\theta_{i,j}]}, \quad i = 1, 2, \dots, N, \quad j = 1, 2, \dots, N \quad (8)$$

3.2 Mean value analysis for IPACT

An EPON with gated IPACT DBA behaves as follows: at the end of a serving period of ONU_i , this ONU requests grants for all queued packets that have not been served. In the next serving period of ONU_i , these packets will be served (Gated Service).

At any given time during a cycle, in the queue of ONU_i , there are packets that are to be served in the cycle (the packets “behind” the gate), and packets that will be served in the next cycle (the packets “before” the gate).

Let $E[\tilde{L}_{i,j}]$ be the mean queue length before the gate at queue i at an arbitrary epoch within the visit time of queue $j, i, j = 1, 2, \dots, N$ and $E[\bar{L}_i]$ be the mean queue length behind the gate during the service time of queue i (keep in

mind that the visit time defined previously differs from the service time).

When the server is not serving queue i , the queue length behind the gate equals the number of arrivals during the previous cycle. Its mean value equals $\lambda_i E[C]$.

For all $i = 1, 2, \dots, N$, the mean queue length before the gate $E[\tilde{L}_i]$, and the mean unconditional queue length $E(L_i)$ are given by:

$$E[\tilde{L}_i] = \sum_{k=1}^N q_{k,1} E[\tilde{L}_{i,k}] \quad (9)$$

$$E(L_i) = \sum_{k=1}^N q_{k,1} E[\tilde{L}_{i,k}] + \rho_i E[\bar{L}_i] + (1 - \rho_i) \lambda_i E[C] \quad (10)$$

Let us consider a tagged packet at its arrival at queue $i, i = 1, 2, \dots, N$. Based on the PASTA property [24] we know that this packet sees the system in equilibrium. The components of the delay for the tagged packet are the following

- the time that elapses between the tagged packet’s arrival and the time at which the server next leaves queue i ; this is a residual $(i + 1, N)$ -period.
- the time between the server’s departure from queue i and its next arrival at the same queue; this is a cycle minus the service time for queue i .
- the time taken to serve all the packets waiting in queue i before the gate at the tagged packet’s arrival.

This yields the following formula for $E[W_i]$

$$E[W_i] = E[R_{\theta_{i+1,N}}] + (1 - \rho_i) E[C] + E[\tilde{L}_i] E[B_i] \quad (11)$$

In combination with Little’s Law [13] $E(L_i) = \lambda_i E[W_i]$, we get the following relation

$$E(L_i) = \rho_i E[\tilde{L}_i] + \lambda_i (E[R_{\theta_{i+1,N}}] + (1 - \rho_i) E[C]) \quad (12)$$

Substituting $E(L_i)$ from (10) in (12), we get

$$(1 - \rho_i) \sum_{k=1}^N q_{k,1} E[\tilde{L}_{i,k}] + \rho_i E[\bar{L}_i] = \lambda_i E[R_{\theta_{i+1,N}}] \quad (13)$$

When the server finishes serving queue i and moves toward queue $i + 1$, the number of packets before the gate at queue i is 0 as all packets are moved behind the gate. So, the mean length of queue i before the gate at an arbitrary epoch of the visit time of queue $i + 1$ equals the number of arrival during the visit time of queue $i + 1$: $E[\tilde{L}_{i,i+1}] = \lambda_i E[R_{\theta_{i+1,1}}]$. We can also compute the number of packets “before” the gate at queue i for an arbitrary

moment within an $(i + 1, j)$ -period, which is equal to the number of arrivals during an $(i + 1, j)$ -period.

$$\sum_{k=i+1}^{i+j} \frac{q_{k,1}}{q_{i+1,j}} E[\tilde{L}_{i,k}] = \lambda_i E[R_{\theta_{i+1,j}}], \quad i, j = 1, 2, \dots, N \quad (14)$$

Now, we see that (13) and (14) give us $N(N + 1)$ equations for $E[\tilde{L}_{i,j}]$, $E[\bar{L}_i]$, and $E[R_{\theta_{i,j}}]$. In order to eliminate the unknown mean residual (i, j) -period $E[R_{\theta_{i,j}}]$, we will express $E[R_{\theta_{i,j}}]$ in terms of $E[\bar{L}_i]$.

$$\begin{aligned} E[R_{\theta_{i,1}}] &= \frac{E[S_{i-1}]}{E[\theta_i]} (E[R_{S_{i-1}}] + \rho_i E[C]) \\ &\quad + \frac{\rho_i E[C]}{E[\theta_i]} (E[\bar{L}_i] E[B_i] + E[R_{B_i}]) \\ i &= 1, 2, \dots, N \end{aligned} \quad (15)$$

$$\begin{aligned} E[R_{\theta_{i,2}}] &= \frac{q_{i,1}}{q_{i,2}} (E[R_{\theta_{i,1}}] + q_{i+1,1} E[C]) \\ &\quad + \left(1 - \frac{q_{i,1}}{q_{i,2}}\right) E[R_{\theta_{i+1,1}}], \quad i = 1, 2, \dots, N \end{aligned} \quad (16)$$

$$\begin{aligned} E[R_{\theta_{i,j}}] &= \frac{q_{i,1}}{q_{i,j}} \left(E[R_{\theta_{i,1}}] + \sum_{k=i+1}^{i+j-1} q_{k,1} E[C] \right) \\ &\quad + \left(1 - \frac{q_{i,1}}{q_{i,j}}\right) E[R_{\theta_{i+1,j-1}}] \end{aligned} \quad (17)$$

We focus first on computing the value of $E[R_{\theta_{i,1}}]$. Here, $E[R_{\theta_{i,1}}]$ equals the sum of a residual setup time and a service time of queue i with probability $\frac{E[S_{i-1}]}{E[\theta_i]}$ (the probability that the server is in the setup period), while with probability $\frac{\rho_i E[C]}{E[\theta_i]}$ the server is serving a packet at queue i , so $E[R_{\theta_{i,1}}]$ equals, with probability $\frac{\rho_i E[C]}{E[\theta_i]}$, the sum of the service time of $E[\bar{L}_i]$ packets and the residual service time of the currently serviced packet. Consequently, the value of $E[R_{\theta_{i,1}}]$ is computed as shown in (15).

In the case of an $(i, 2)$ -period, $E[R_{\theta_{i,2}}]$ equals $E[R_{\theta_{i+1,1}}]$ with probability $(1 - \frac{q_{i,1}}{q_{i,2}})$ (i.e. the probability that the server is in the visit time of queue $i + 1$), while with probability $\frac{q_{i,1}}{q_{i,2}}$ this residual period $E[R_{\theta_{i,2}}]$ equals the sum of $E[R_{\theta_{i,1}}]$ and a visit time of queue $i + 1$. Consequently, the value of $E[R_{\theta_{i,2}}]$ is computed as shown in (16).

Finally, for general, $E[R_{\theta_{i,j}}]$ is computed similarly as shown in (17) for $i = 1, 2, \dots, N$ & $j = 2, 3, \dots, N$. The elimination of $E[R_{\theta_{i,j}}]$ in (13) and (14) with the help of (17) yields a set of $N(N + 1)$ linear equations for equally many unknown $E[\tilde{L}_{i,j}]$ and $E[\bar{L}_i]$. Consequently, we can obtain the mean queue length $E[L_i]$ from $E[\tilde{L}_{i,j}]$ and $E[\bar{L}_i]$, and we can also obtain the mean packet delay $E[W_i]$ with the help of Little's Law [13].

3.3 Discussion of the results

Note that the method described previously is exact in the sense that it does not assume light load or heavy load situations like the method presented in [12] which has to rely on this type of assumption. This allows us to claim that our method is more exact than that in [12] since it applies to any load situations. However, we point out that our method is still approximate since (4) is used to compute the mean cycle time duration which cannot be smaller than an RTT in an EPON [11].

Moreover, our method applies only to Poisson Arrivals, which is not a realistic assumption for access networks, where traffic is probably very bursty. However, as we show in the next section, it allows to point out a major defect in IPACT with gated service scheme, i.e. its lack of support in isolating one customer from the others.

4 Performance of IPACT with gated service

In the following, we investigate the general case of a non homogeneous case, in which the ONUs offer different amounts of traffic. We consider an EPON with gated IPACT with $N = 16$ ONUs. The rate of the upstream channel is 1 Gbps, the maximum RTT between OLT and ONU is 160 μ s and the setup time (guard time) is 2 μ s.

We first consider the case where several ONUs send much more traffic than the others; this could represent e.g. misbehaving ONUs, sending more traffic than expected. We thus assume that regular ONUs have the same arrival rate 50 Mbps and 4 overloaded ONUs have a higher arrival rate, varying from 50 to 90 Mbps. Thus, the global load for the EPON varies from 0.8 to 0.96. Figure 1 shows the mean number of waiting packets in both overloaded ONUs and regular ONUs versus arrival rate of overloaded ONUs. Figure 2 illustrates the mean packet delays for overloaded ONUs and regular ONUs in this case. Those figures compare the results obtained by simulation and theoretical analysis. The simulation experiments were performed using Network Simulator NS-2 [18].

These figures show the accuracy of the proposed analytical derivation for the mean queue size and the mean waiting time. In these figures, and in all cases analyzed so far, the approximate values obtained by our MVA approach slightly undervalues the exact values.

One could think, looking at Fig. 1, that the performance in terms of mean waiting time for overloaded ONUs is worse than the one for regular ONUs. However, Fig. 2 shows that this is not the case: overloaded ONUs and regular ONUs experience the same mean packet delay performance which increases with the overloaded ONUs offered load. This is because, as noted in [12], the mean waiting time is tightly

Fig. 1 Impact of overload over mean queue lengths

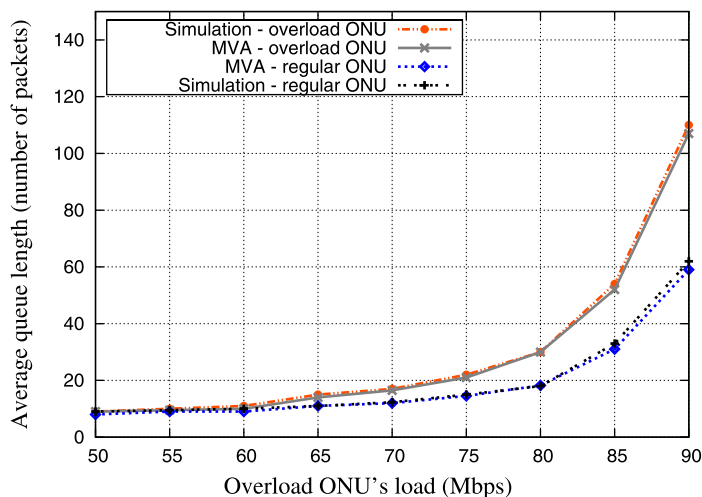
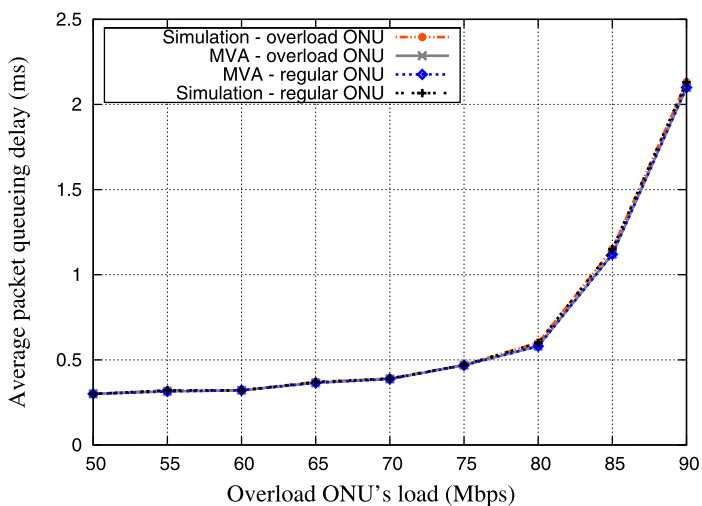


Fig. 2 Impact of overload over mean packet delays



related to the mean cycle time, which is obviously the same for both regular and overloaded ONUs.

In order to further illustrate the impact of traffic heterogeneity on system's performance, we now consider the case where the total network load is kept constant while the proportion of overloaded ONUs vary. We split the 16 ONUs into two sets; in the first set of n_1 ONUs, the arrival rate is λ_1 while in the second set of n_2 ONUs, the arrival rate is λ_2 ($\lambda_1 > \lambda_2$). The total offered load $n_1\lambda_1 + n_2\lambda_2$ is constant and equals 800 Mbps. For each pair (n_1, λ_1) , Figs. 3 and 4 respectively show the mean queue length of one ONU in the first set (overloaded ONU) and the mean packet delay for this ONU (actually, the mean packet delay does not depend on whether the ONU is overloaded or underloaded). Figure 4 shows clearly that the mean queuing delay (which takes the constant value 0.34 ms in this experiment) depends only on the total offered load, and not on the load offered by the tagged ONU.

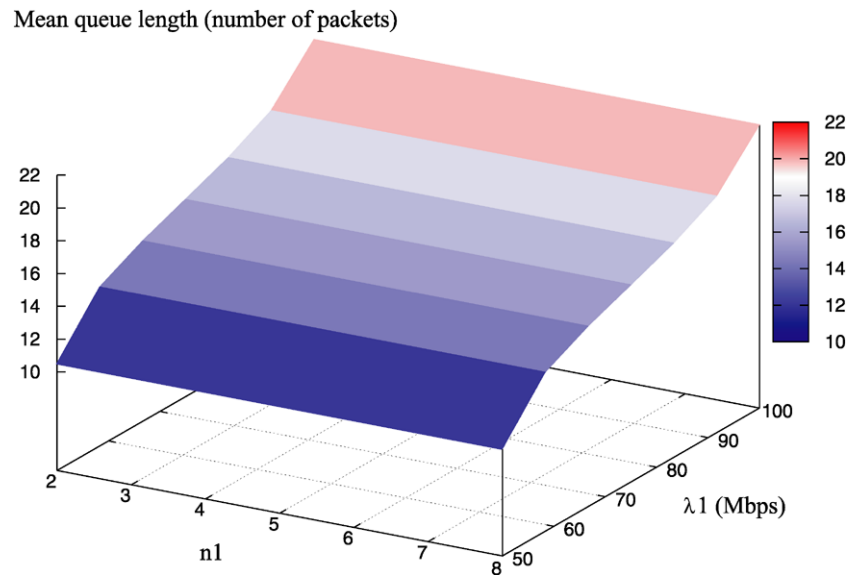

This raises a serious issue for the access network operator: unless all ONUs send the traffic that is expected by the

operator and that is e.g. specified in an appropriate Service Level Agreement (SLA), the network operator cannot ensure even as simple a performance metric as the mean waiting time. Indeed, a single, misbehaving ONU can monopolize the uplink capacity and degrade the QoS of the whole system. In other words, the network operator cannot isolate the customers from one another, which is a major problem in any access network.

This is why mechanisms that can efficiently serve ONUs, such as IPACT, but can also differentiate between committed traffic and best effort traffic, are needed. We address this requirement in the next section.

5 Control plane architecture for EPONs

The objective of this section is to propose a generic traffic management policy for EPONs that can isolate customers from one another, enforce traffic profiles dictated by SLAs

Fig. 3 Mean queue length in an overloaded ONU versus (n_1, λ_1) **Fig. 4** Mean packet queuing delay versus (n_1, λ_1) 

while efficiently distribute excess resources to Best Effort Traffic.

The control plane implements many functions. Some of them are related to Operation Administration and Maintenance (OAM) procedures that are used to initialize the PON, to insert new ONUs, to provide synchronization in the system, etc. The present paper does not address these issues but concentrates on the traffic management functions addressed by the control plane. These include:

- Connection Acceptance Control (CAC)** that accepts or rejects the establishment of a new connection;
- Conformance Checking** that consists in enforcing the traffic description negotiated in SLAs;
- QoS Control** that dynamically allocates resources in such a way that the QoS levels negotiated in the SLAs are delivered and the non committed resources are fairly shared.

The configuration of the EPON scheduling policy (i.e. DBA and intra-ONU scheduling) is thus obviously part of the control plane of EPONs.

In the following, we first present a brief review of the various proposals for supporting multiservice in EPONs. We then propose a comprehensive traffic management policy that shall be shown, in the remained or the paper, to support versatile implementations of the above traffic management functions.

5.1 State of the art in resource management in EPON

This section presents a brief review of upstream resource management in EPONs.

The Multi-Point Control Protocol (MPCP) is developed to facilitate the dynamic bandwidth allocation of transmission opportunities [8]. This protocol provides a tool for

exchanging requests and grants between the OLT and the ONUs. MPCP uses two control messages (GATE and REPORT) for coordinating upstream data transmission.

The GATE message is used to grant Transmission Windows (TW) to ONUs. We assume here that, in a REPORT message, an ONU indicates the exact number of bytes waiting in each queue.

The EPON standards do not specify any specific algorithm for bandwidth allocation; in the following, we review the major approaches proposed in the recent past.

The most well known DBA is called Interleaved Polling with Adaptive Cycle Time (IPACT) [11]. IPACT, in its usual Limited Service (LS) implementation, answers directly to each received REPORT by assigning the required transmission slots in the limit of a given predetermined maximum. This discipline has been shown to provide good performance in terms of packet delay and throughput for a single traffic class.

An ideal DBA should not only efficiently allocate resources between users, but should also deliver the QoS levels requested by different services. A large variety of QoS-enabled DBA mechanisms have been proposed in this field.

In [3], the authors propose a DBA scheme in which ONU nodes are partitioned into two groups, “underloaded” and “overloaded”, according to their minimum guaranteed transmission window sizes. In that scheme, the bandwidth saved from the underloaded group is proportionally re-allocated to overloaded ONUs. The authors here use priority queueing and intra-ONU queue scheduling to provide differentiated service for each class. This scheme is shown to fairly distributes the excessive bandwidth among highly loaded ONUs. However, the OLT still only supports a single traffic class.

Reference [14] proposes 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.

In [10], the authors propose IPACT-LS-QoS DBA in which every ONU maintains a separate queue for each class of service in its buffer but addresses only “global” requests to the OLT. The OLT then issues “global” grants to each ONU and each ONU uses priority-based intra-ONU scheduling to determine the order in which its queues are processed. The OLT can then neither favour demands for high priority class traffic, nor control how the ONU shares its grants between classes. Therefore, the mechanisms proposed in [10] cannot provide bandwidth guarantees for traffic flows specified by SLAs.

Authors in [2, 21] proposed “per class subframing” approach to allocate the bandwidth. This approach consists of

separating the transmission of high priority and low priority packets into two sub-cycles. By protecting high priority traffic in a separate and fixed sub-cycle, its delay jitter performance is improved. However, “per class subframing” schemes consume upstream bandwidth since each ONU transmits multiple times during each polling cycle, degrading the efficiency by increasing the amount of “guard times”.

Authors in [6] also propose a DBA using cyclic polling scheme in order to support DiffServ. Expedited Forwarding (EF) traffic is allocated a guaranteed amount of bandwidth regardless of the number of packets waiting in the queue, while Assured Forwarding (AF) traffic is allocated bandwidth as per their proportion in the total demand. This DBA precludes however AF traffic to take advantage of inactive EF sources. Moreover [6] does not relate EF and AF classes to typical traffic classes.

None of the reviewed traffic management policies are versatile enough to support either an efficient Best Effort service, or a service taking into account strict SLAs or QoS sensitive services.

5.2 Connection Acceptance Control in EPONs

A CAC is applied in order to enable the operator to accept or refuse a new traffic flow. This is the case for many communication services (Voice or Conference) and video services (IPTV or VoD).

A recent paper [7] proposes a CAC framework for EPONs, that is distributed between a local CAC, performed by the ONU and a global CAC that is under the control of the OLT. However, CAC is a global issue for a network operator and it cannot rely on local, uncoordinated, decisions taken in every portion of the network. We thus prefer to delegate the complete CAC procedure to the OLT. More precisely, we assume that the EPON does not implement the CAC, but enforces CAC decisions taken by the network operator. This is because the EPON standard does not support a connection oriented transfer mode: no MAC primitives have been defined for implementing connection SETUP requests, connection CONNECTs, nor to terminate connections in EPONs. Apparently, this point has been overlooked in [7].

In our framework, the characteristics of the EPON access should be taken into account by the operator because it may constitute a traffic bottleneck, and also because the end-to-end QoS delivered to the flow depends on the QoS delivered by the EPON. Moreover, the CAC also has a direct impact on other traffic control functions: Conformance Control should be aware of the amount of resources committed by the network operator, and QoS control should deliver the committed QoS to any accepted flow.

Therefore, we assume that a CAC function directly interacts with the OLT, so that the latter can enforce dynamic traffic conformance decisions and control the QoS offered to various traffic flows.

5.3 Definition of service classes

The usual multiservice approach in Ethernet networking is to classify packets into a limited number of traffic classes, and to apply a strict priority scheduling on packets from different classes. Our approach relies on the packet marking method provided by Ethernet [1], but for the sake of simplicity, takes account of only 3 classes:

Real-time Committed Bandwidth T_0 is designed to support interactive real-time applications that require a guaranteed bandwidth, low delay and low jitter.

Data Committed Bandwidth T_1 is dedicated to non-interactive guaranteed bandwidth applications that are less delay and jitter-sensitive.

Best Effort T_2 refers to data that can be sent at the leisure of the system. A major difference between CB and BE traffic is that CB traffic is characterized by a traffic profile specified in a SLA 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.

For illustration purposes, we consider in the present paper the following target values for maximum packet delay: 3 ms for T_0 traffic and 30 ms for T_1 traffic. However, note that both the proposed DBA and/or the Connection Admission Control (CAC) allow to consider different values.

In the following, we evaluate the performance delivered to QoS classes by computing the upper quintile of the packet delay (called maximum packet delay in the remainder of the paper). Indeed, both delay and jitter performances have to characterize the “maximum values”, and not the average values taken by delay and jitter. However, absolute maxima are not good performance indicators since they directly depend on the duration of the measurement period, which is why we introduce quintiles. 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 $(1 - p)$.

5.4 Implementing conformance control and QoS control

We now describe an EPON control plane framework, based on a *DBA with Traffic Conformance Mechanism* (DBA-TCM), which includes a DBA for implementing QoS differentiation and a conformance checking mechanism for enforcing negotiated SLAs.

The DBA-TCM has the following characteristics:

- it forces ONUs to comply with their negotiated SLAs by checking traffic conformance: the OLT filters the demands and takes account of a virtual policing scheme

when computing the amounts of grants to be sent to the ONUs.

- it allows variable length polling cycles while enforcing an upper bound on the maximum polling cycle length T_{\max} .
- it is priority based at the inter-ONU level, i.e. DBA-TCM serves each class successively and attempts to satisfy all the demands from one class before allocating transmission opportunities to a lower class.
- it directly controls ONU scheduling through the OLT, i.e. each ONU applies exactly the allocated per-class grants as specified in the GATE message.

A short description of DBA-TCM is given below in order to clarify its role in the proposed control plane for EPON; more details are available in [17].

5.4.1 Traffic Conformance Checking Mechanism

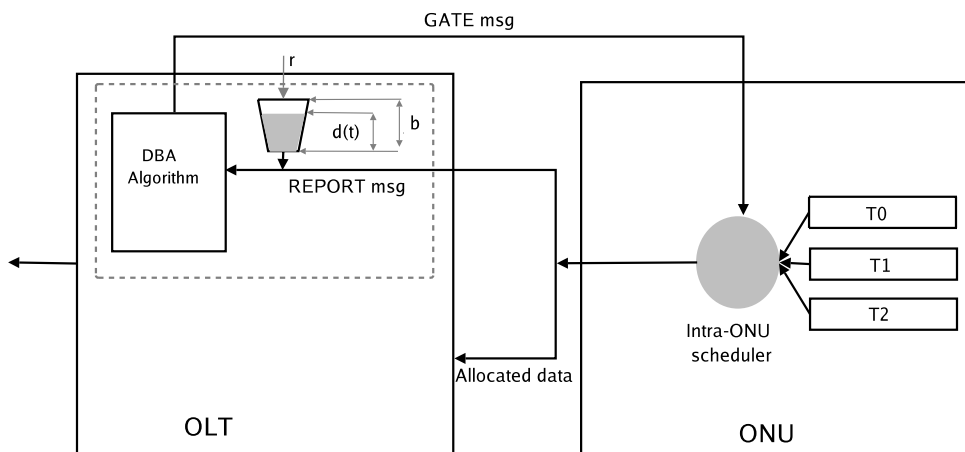
The basic idea of our mechanism is fairly simple: instead of implementing a policing mechanism (such as a Token Bucket filter) for upstream traffic at the OLT, the OLT filters the demands carried in REPORT messages and takes account of a virtual policing scheme when computing the number of grants to 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 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 centralized conformance control plus DBA algorithm can optimize grant allocation for CB traffic, excess (non-conforming) traffic and BE traffic according to a policy specified by the network operator.

Figure 5 shows, for illustration purposes, a simplified version of our proposed architecture for a conformance checking DBA that includes a Traffic Conformance Mechanism (TCM) that checks the conformance of requests carried in the REPORT messages.

For ONU_i , the TCM 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 . A pair of values (r_i, b_i) represents a profile for T_0 traffic. Before reaching the DBA algorithm module in Fig. 5, a high priority request $R_i^{T_0}$ (in bytes) has to find an equivalent number of tokens (in bytes) in the bucket in order to be considered as conforming.

$$R_i^{T_0} = \begin{cases} R_i^{T_0} & \text{if } R_i^{T_0} < d_i(t) \\ d_i(t) & \text{otherwise} \end{cases} \quad (18)$$

Fig. 5 Proposed architecture for the conformance checking DBA



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 (marking, buffering or discarding) depending on the traffic management strategy desired by the network operator.

5.4.2 Dynamic Bandwidth Assignment

Conforming T_0, T_1 demands are now used for DBA computation. We consider an EPON with N ONUs and assume that the link capacity is C Mbps. 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:

$$B_i^{\lim} = \frac{(T^{\max} - Nt_g)C}{8}w_i, \quad \text{with } \sum_{i=1}^N w_i = 1 \quad (19)$$

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). Note that w_i takes into account reservations for both T_0 and T_1 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.

DBA-TCM algorithm serves each class successively and attempts to satisfy all the conforming demands from one

Table 1 Mean delay performance for different policies

Overload ONU load	50 Mbps	80 Mbps	85 Mbps	90 Mbps
Overload ONU-LS	0.34 ms	0.58 ms	1.49 ms	6.98 ms
Regular ONU-LS	0.34 ms	0.58 ms	0.96 ms	0.98 ms
Overload ONU-GS	0.34 ms	0.58 ms	1.2 ms	2.13 ms
Regular ONU-GS	0.34 ms	0.58 ms	1.2 ms	2.13 ms
Regular ONU-TCM	0.34 ms	0.34 ms	0.36 ms	0.35 ms

Regular ONU arrival rate = 50 Mbps

class before allocating transmission opportunities to a lower class.

Once a ONU receives the GATE message, it performs the intra-ONU policy. As argued previously, we assume here 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.

In order to illustrate the behavior of DBA-TCM, we compare the mean delay performance results of the scenario already considered in Fig. 2 using IPACT with Limited Service scheme (LS), IPACT with Gated Service scheme (GS), and DBA-TCM. In LS, the maximum cycle time can be upper bounded thanks to the definition of TW^{\max} (see Sect. 2). This partly prevents a malicious users from monopolizing the uplink capacity and from degrading the QoS of all traffic. DBA-TCM also ensures that an ONU is not allowed more grants than specified in its SLA [17].

In Table 1, the results for LS and DBA-TCM cases are obtained by simulation, while the results for GS are obtained by the model developed in Sect. 3. We assume that $TW^{\max} = 7500$ bytes (which results in a maximum cycle time value of 1.0 ms) [11] and that an arrival rate of 50 Mbps is specified in the SLA for each regular ONU (called conforming ONU in [17]).

As shown in Table 1, both regular and overloaded ONUs experience a large waiting time for GS, whereas LS limits

the queueing time for regular ONUs when the overload increases: the mean waiting times for regular and overloaded ONUs differ when the overloaded ONUs send 85 Mbps or more, but not when they send 80 Mbps or less. DBA-TCM policy offers a better protection to regular ONUs, as the mean waiting time for a regular ONU does not depend on the overloaded ONUs behavior. On the other hand, the performance delivered to the overloaded ONUs is very degraded since the system only accepts 50 Mbps for each ONU. This is why no value is presented for the overloaded ONUs, since their queue is in a transient state (increase).

The simple experiment described in Table 1 justifies a more thorough investigation of the capabilities offered by DBA-TCM. This is done in the next section.

6 Supporting Various Traffic Offers

In this section, we assess how the control plane can be configured to support various offers or access network services, and the performance of this support.

6.1 Simulating EPON traffic management

The simulated network architecture and the various traffic types that are considered in the following analysis are now specified. The simulation platform used for this work has been developed using Network Simulator (ns2) and is described in [18].

We consider an EPON system with 16 ONUs, and a total uplink bandwidth C of 1 Gbps. The maximum cycle time T_{\max} is set to 1.5 ms and the guard time t_g between 2 consecutive ONU transmissions is 4 μ s.

In our experiments, T_0 and T_1 traffic are simulated as VBR, “exponential ON-OFF models” (using the ns2 terminology). The mean rate (respectively the peak rate) is denoted R_a (respectively R_p). Class T_2 is modeled as best-effort traffic which is generated following a Pareto distribution or simply as a greedy (i.e. always active) FTP source. Each ONU may have different traffic profiles specifying T_0 and T_1 traffic flows. Depending on the scenario, T_2 traffic is modeled either as a Pareto source, or as a greedy FTP source.

For each simulated point, we simulated more than 10^6 packets and estimated the upper quintile of the packet delay with $(1 - p)$ equals 10^{-3} .

6.2 A single class, best effort service offer

We first investigate the case where the operator wishes to support only a classless, Best Effort service. This offer is the most common one in the residential market. For the operator, the issue is twofold: first ensure that users are fairly

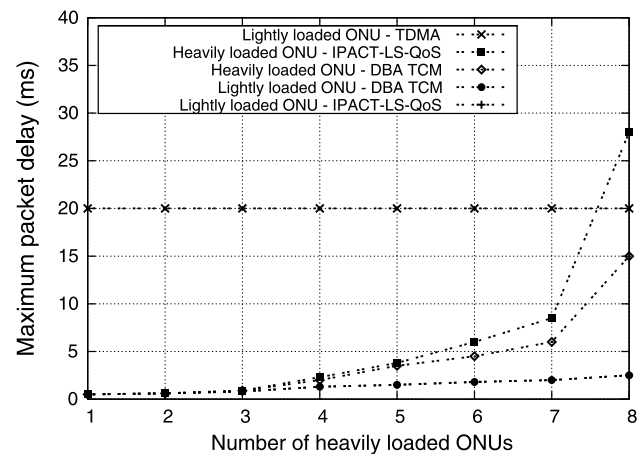


Fig. 6 Asymmetric case: maximum packet delay for Best Effort traffic

served (e.g. avoid that one group of users starves another group), and secondly ensure that resources are efficiently used (e.g. serve as much users demands as possible).

The configuration of DBA-TCM is simple in this case: each ONU is allocated an equal share of available capacity that is slightly smaller than C/O (O is the number of ONUs) in order to account for transmission overhead; we assume here that the OLT allocates 60.0 Mbps to each of the 16 ONU. All traffic is handled as T_1 traffic with a traffic profile described as a Leaky Bucket with the configured rate and a 0 Burst tolerance. Moreover, Conformance Control is activated; the grant allocation process takes account of conformance, but may allocate grants to non-conforming packets. Neither conforming nor non-conforming packets are marked.

We compare below the performance delivered by an EPON using static Time Division Multiplexing (TDMA), IPACT-LS-QoS (which behaves here as IPACT-LS) and our proposed scheme DBA-TCM.

In the scenario illustrated in Fig. 6, we consider unbalanced traffic offers, and wish to assess whether heavily loaded ONUs can use unallocated resources without degrading the performance offered to lightly loaded ONUs. We consider 8 lightly loaded ONUs ($R_a = 25$ Mbps, $R_p = 70$ Mbps) and a varying number of heavily loaded ONUs ($R_a = 60$ Mbps, $R_p = 120$ Mbps). Figure 2 shows the maximum packet delay performance for the two types of sources, versus the number of active heavily loaded ONUs.

The performance delivered by TDMA for heavily loaded sources is not shown on Fig. 6 because it is very bad (500 ms). TDMA behaves badly compared to the 2 dynamic DBAs for both heavily and lightly loaded sources. IPACT-LS-QoS and DBA-TCM behave similarly for lightly and heavily loaded sources; both protect lightly loaded sources, and efficiently share resources amongst heavily loaded sources. Furthermore, Fig. 2 also shows that DBA-TCM behaves significantly better towards heavily loaded

sources than IPACT-LS-QoS when the number of heavily loaded sources is large. This is because IPACT-LS-QoS applies an Inter-Leaved Polling scheme in which the OLT does not have a global view of all traffic demands when computing transmission opportunities; the OLT has then to blindly limit the number of bytes allocated to one ONU in each cycle, irrespective of the activity of the other ONUs. On the other hand, DBA-TCM allows the OLT to perform bandwidth allocation based on the global traffic bandwidth and thus take more intelligent decisions, allowing the heavily loaded sources to take advantage of unused resources.

This example has demonstrated the usefulness of implementing a dynamic DBA; moreover, DBA-TCM is shown to behave always better than the well known IPACT, with very little configuration effort. Although we have chosen here to identically configure all ONUs, it is possible with DBA-TCM to allocate different level of resources to different types of stations, while allowing heavily loaded sources to take advantage of available resources, whatever is their configured resource level. This is a distinct advantage over IPACT-based mechanisms.

6.3 A multiclass service offer with committed bandwidth

We have previously shown the limits of a classless service offer, which does not support stringent delay performance when the EPON is heavily loaded. This is why some operators choose to offer multiservice support. For illustration purposes, we assume that all 3 classes described in Sect. 3.2 are supported.

If multiple traffic classes are supported, traffic profiles should be negotiated for each class with committed QoS, since only relative QoS can be supported otherwise. For relative QoS offers, a high priority class receives better QoS than a low priority class, but only global QoS levels can be guaranteed, since each ONU is de facto allowed to send only high priority traffic!

For a Multiclass Service Offer, the operator should be able to ensure that each class receives the negotiated QoS levels (delay performance for T_0 traffic are more stringent than for T_1 traffic).

The configuration of DBA-TCM in case of multiservice offer is as follows: each ONU is allocated a share of available capacity computed according to its SLA. In each ONU, each traffic class has its own traffic profile (and possibly dynamically modified depending on CAC decisions). We assume that each traffic profile is known by the OLT. Traffic Conformance Control is activated.

We wish to assess and compare the QoS supports offered by DBA-TCM and IPACT-LS-QoS. We investigate a scenario where some ONUs require Committed Bandwidth services while others require only a “Best Effort” access. This scenario could represent an EPON supporting both residential

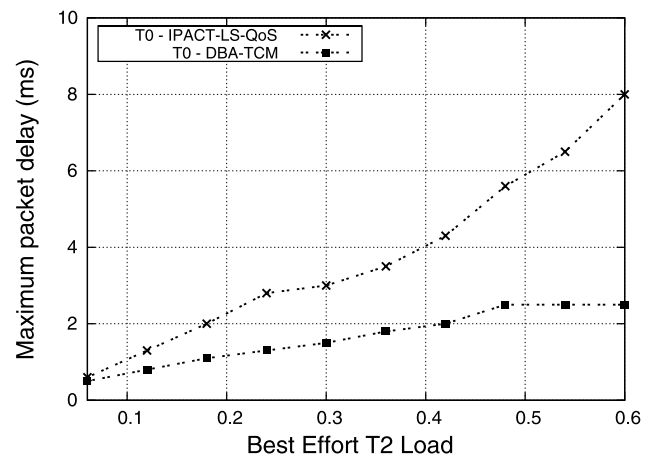


Fig. 7 Maximum packet delay for T_0 traffic in a multiservice scenario

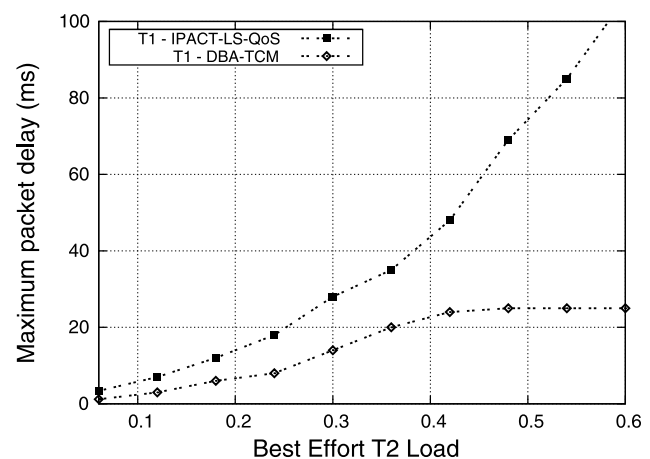


Fig. 8 Maximum packet delay for T_1 traffic in a multiservice scenario

and SOHO/SME customers; it also corresponds to a scenario where only part of the ONUs request committed bandwidth services at a given time. In a first set of 6 ONUs, each ONU send T_0 traffic at an average rate of 20 Mbps and T_1 traffic at an average rate of 40 Mbps. The 10 ONUs in the second set send only T_2 (BE) traffic. Figures 7 and 8 show the maximum packet delay performance for T_0 and T_1 traffic respectively versus the load offered by T_2 traffic. The QoS delivered to T_2 is not shown since it is actually very bad (but there is no actual commitment on it).

A first comment on Figs. 7 and 8 is that both IPACT-LS-QoS and DBA-TCM offer different levels of QoS to T_0 and T_1 traffics. However, their respective reaction to the increase in T_2 traffic load significantly differ. For DBA-TCM, the QoS delivered to T_0 and T_1 traffics is not impacted by a large T_2 traffic load. This is because in DBA-TCM, the OLT, knowing all demands, first attempts to satisfy T_0 traffic demands, then T_1 traffic demands and lastly T_2 traffic if there are some remaining resources. The curves for DBA-TCM indeed become flat (i.e. the delay performance is not

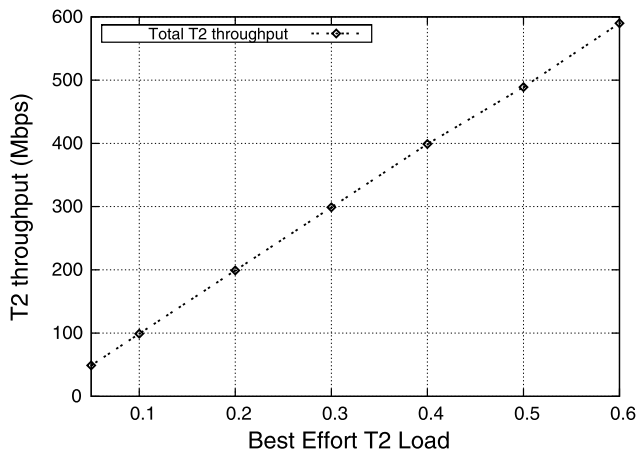


Fig. 9 Total T_2 throughput in a multiservice scenario

degraded) for T_2 traffic load of (roughly) 0.4 because then, the frame size is systematically limited to 1.5 ms. This is not the case for IPACT-LS-QoS which cannot efficiently protect CB traffic from T_2 traffic: the delay performance delivered to T_0 traffic is as bad as 8 ms, which is much too large for real time service.

We observe on Fig. 7 (respectively Fig. 8) that the target delay value for T_0 (respectively T_1) traffic is enforced by DBA-TCM. Moreover, Fig. 9 shows that DBA-TCM allows BE traffic to use the available bandwidth with a good efficiency.

A detailed analysis of the behavior of DBA-TCM, and its multiservice support is presented in [16].

As shown by the above examples, DBA-TCM can be easily configured to support different types of service offers. The complete control of upstream traffic management is assured by the OLT which allows designing simple ONUs; in particular, there is no configuration to carry out in ONUs, since conformance is enforced in the OLT only. If a customer wishes to change its current SLA, no operation has to be performed on the ONU, but only on the conformance parameters in the OLT.

6.4 Admission control for EPON systems

The goal of this section is to discuss the relationship between CAC and traffic management in the EPON. A huge literature addresses the design of CAC in broadband networks using Point-to-Point links. A good overview of the major approaches is given in [20]. Although a discussion of the pros and cons of CAC approaches is outside the scope of the present paper, we shall consider here a CAC designed to deliver statistical and not absolute QoS guarantees, because delivering absolute guarantees leads to a severe underutilization of resources.

For illustration purposes, we consider homogeneous, independent ON-OFF sources with exponential ON and OFF

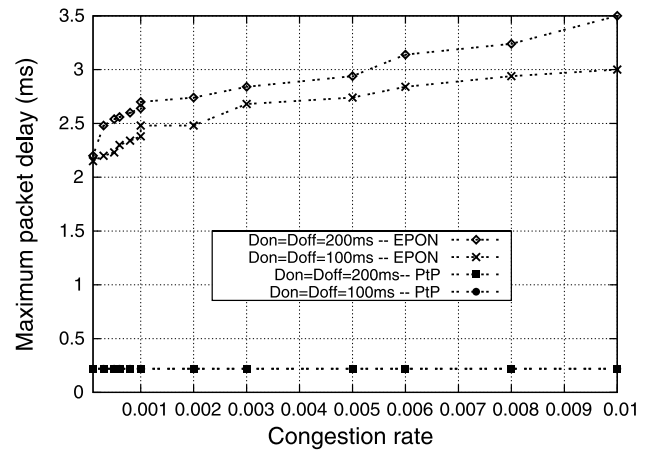


Fig. 10 Maximum packet delay for Real-time Committed Bandwidth; Point-to-Point versus PON

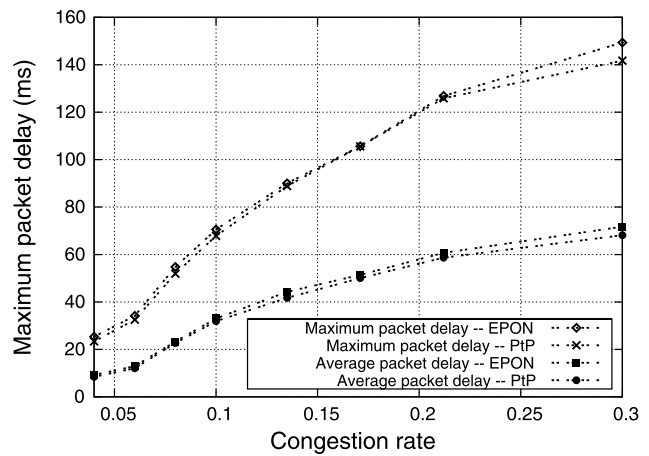


Fig. 11 Packet delay for Data Committed Bandwidth; Point-to-Point versus PON

distribution and with Peak Rate $R_p = 10$ Mbps. The number of active sources at time t is obviously a continuous time Markov Chain.

We consider that the system is “congested” when M (here, $M = 100$) or more sources are in the ON state. Let ϵ be a target upper value for the congestion probability. A small target value for ϵ (i.e. smaller than 0.01) is applied for T_0 traffic, and a larger value for ϵ (i.e. between 0.05 and 0.3) is accepted for T_1 traffic. This is admittedly rather naive, but it conveniently models the fact that for T_0 traffic, there should be minimal queuing whereas T_1 traffic is allowed to queue in the system.

Let us then consider the following CAC for homogeneous sources: the number of admissible sources is the maximum N for which the congestion probability is smaller than ϵ .

In the following, we consider scenarios with a single class of traffic (either T_0 or T_1). Figures 10 and 11 compare the delay performance versus target congestion rate achieved respectively by Point-to-Point and EPON systems.

Figure 10 considers T_0 traffic (ϵ is smaller than 0.01). We consider 2 cases, corresponding to different ON-OFF period durations. We see that the delay performance achieved by the Point-to-Point system is excellent and not sensitive to ϵ , regardless of the ON-OFF period duration. This is however NOT the case for the EPON system: the performance is significantly worse and degrades when ϵ increases, and it also depends on the ON-OFF period duration. This is due to the impact of polling: although the cycle time is limited, all traffic has to wait at least a cycle time, and more if there is significant buffering in the ONUs. This is why the delay performance can be as bad as 3.5 ms (a bit more than twice the maximum cycle time) for $\epsilon = 0.01$. We see that, in order to limit the packet delay of T_0 traffic to 3 ms in the present case, the congestion rate should be limited to less than 0.005. This implies that designing a CAC for T_0 traffic in an EPON is not straightforward, since it is necessary to explicitly take account of the polling behavior.

On the other hand, what we observe in Fig. 11 shows that the delay performances for both Point-to-Point and EPON systems are very similar. Indeed, the impact of the polling delay is negligible compared to the queueing time in the ONU. The practical consequence is that any classical CAC designed for a Point-to-Point system is likely to directly apply to an equivalent PON access.

Figure 10 also shows both the average and the quintile for the delay distribution. In this particular case, the upper quintile is roughly twice as large as the average delay. This highlights the fact that the delay performance should be analyzed in terms of quintile and not in terms of average delay as in [7].

7 Conclusion

This paper addresses the issue of delivering QoS in EPON based access networks. In order to assess the QoS delivered by the well known benchmark IPACT with Gated Service, we have derived a MVA based method for computing the mean waiting time and the mean queue length in an EPON with gated IPACT, under the assumption of Poisson arrivals. As shown by comparing analytical results with simulation results, this method is indeed very accurate. We have used the MVA method to investigate the impact of offered traffic heterogeneity on the delivered performance.

We have shown that, although the mean queue lengths depend on individual offered load, the mean waiting time is only sensitive to the global offered load. This has a significant impact since it means that a single misbehaving ONU may induce global performance degradations in the EPON. This argues in favour of modifying the Gated policy, either by simply limiting the maximum amount of traffic granted to an ONU in a cycle (Limited Service policy), or by using a

DBA that computes grants according to a traffic control policy that protects conforming ONUs from overloaded ONUs.

This paper further shows how a simple priority based DBA with a Traffic Conformance checking Module (DBA-TCM) is a major building block for an EPON control plane, that can be configured to support both a simple best effort access service offer and a multiclass service offered with tight QoS guarantees.

In the recent past, several authors advocate in favour of de-centralizing traffic management procedures by giving more autonomy to ONUs. The present approach does not follow this trend, in order to keep ONU design as simple as possible and to delegate all major decisions to the OLT. The OLT is indeed clearly under the control of the network operator, whereas ONUs, as DSL or cable modems, are distant and less easily managed.

Although DBA-TCM is fairly simple in its principles, this paper shows that it easily supports different service offers, and is capable of delivering Committed Bandwidth services with stringent QoS guarantees whereas the well known IPACT-LS-QoS (IPACT with limited Service and QoS support) is slightly less efficient and cannot support stringent QoS guarantees. A previous paper [17] has shown that DBA-TCM can protect users who comply with a negotiated traffic profile from non conforming traffic, or from Best Effort Traffic that is not controlled by a traffic profile. Two of the main functions of a control plane (conformance checking and QoS control) are thus fulfilled by DBA-TCM.

The paper also addresses the important problem of the integration of CAC into an EPON control plane. DBA-TCM in itself does not implement a CAC procedure because the MAC for EPON does not support the necessary mechanisms. However, any CAC implemented in the network node where the OLT is located can directly interact with DBA-TCM through the OLT (e.g. by dynamically modifying the traffic profiles associated to an ONU). The present paper argues that while any CAC designed for Committed Bandwidth services with no stringent delay requirements can be directly used for EPONs, this is not the case for services with stringent delay requirements. In this particular case, the polling procedures adds a delay which can be as large as several ms. It is thus necessary to be more stringent with EPONs than with a classical Point-to-Point access architecture.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. ANSI/IEEE Standard 802.1D available at <http://www.standards.ieee.org/>.

2. An, F. T., & Hsueh, Y. (2003). A new dynamic bandwidth allocation protocol with quality of service in Ethernet-based passive optical networks. In *Proceeding of WOC 2003*.
3. Assi, C. M., Ye, Y., & Dixit, S. (2003). Dynamic bandwidth allocation for quality-of-service over Ethernet PONs. *IEEE Journal Selected Areas in Communications*, 21(9), 1467–1477.
4. Bai, X., Shami, A., & Ye, Y. (2007). Delay analysis of Ethernet passive optical networks with quasi-leaved polling and gated service scheme. In *Proc. IEEE Second International Conference on Access Networks (AccessNets2007)*, Ottawa, Canada, 2007.
5. Bhatia, S., Garbuzov, D., & Bartos, R. (2006). Analysis of the gated ipact scheme for EPONs. In *Proc. IEEE International Conference on Communications (ICC06)*, Turkey, June 2006 (Vol. 6, pp. 2693–2698).
6. Choi, S.-i. (2004). Cyclic polling based dynamic bandwidth allocation for differentiated class of service in Ethernet PONs. *Photonic Network Communications*, 7(1), 87–96.
7. Dhaini, A., Assi, C., Maier, M., & Shami, A. (2007). Per-stream QoS and admission control in Ethernet passive optical networks (EPONs). *IEEE/OSA Journal of Lightwave Technology*, 25(7), 1659–1669.
8. Ethernet in the First Mile Task Force IEEE Std 802.3ah. Available at <http://www.ieee082.org/3/efm/>, 2004.
9. Green, P. E. (2004). Fiber to the home: the next big broadband thing. *IEEE Communications Magazine*, 42(9), 100–106.
10. Kramer, G., Mukherjee, B., & Pesavento, P. (2002). On supporting differentiated classes of service in EPON-based access network. *Journal of Optical Networks*, 280–298.
11. Kramer, G., Mukherjee, B., & Pesavento, P. (2002). Interleaved polling with adaptive cycle time (IPACT): a dynamic bandwidth distribution scheme in an optical access network. *Photonic Network Communications*, 4(1), 89–107.
12. Lannoo, B., et al. (2007). Analytical model for the IPACT dynamic bandwidth allocation algorithm for EPONs. *Journal of Optical Networking*, 6(6), 677–688.
13. Little, J. D. C. (1961). A proof of the queueing formula $L = \lambda W$. *Operations Research*, 9, 383–387.
14. Ma, M., & Zhu, Y. A bandwidth guaranteed polling MAC protocol for Ethernet passive optical networks. In *Proc. IEEE INFOCOM*, San Francisco, CA, USA (Vol. 1, pp. 22–31).
15. McGarry, M. P., Maier, M., & Reisslein, M., (2004). Ethernet PONs: a survey of dynamic bandwidth allocation (DBA) algorithms. *IEEE Communications Magazine*, 42(8), 8–15.
16. Ngo, M. T., & Gravey, A. (2009). In *Lecture notes in computer science: Vol. 5464. A priority-based multiservice dynamic bandwidth allocation for Ethernet passive optical networks* (pp. 185–199). Berlin: Springer.
17. Ngo, M. T., & Gravey, A. (2009). Traffic conformance issues in passive optical networks: challenges and solutions. *International Journal of Communication Networks and Distributed Systems*, 2(2), 260–280.
18. Ngo, M. T., Gakhar, K., & Gravey, A. (2008). Upstream traffic management in EPONs: a simulation based analysis. In *Proc. ICST Simutools 2008*, Marseille, France, 2008.
19. Park, C. G., et al. (2005). Packet delay analysis of symmetric gated polling system for DBA scheme in an EPON. *Telecommunication Systems*, 30, 13–34.
20. Roberts, J., Mocchi, U., & Virtamo, J. (1996). *Broadband network teletraffic*. Final Report of Action COST 242.
21. Shami, A., Bai, X., Assi, C. M., & Ghani, N. (2005). Jitter performance in Ethernet passive optical networks. *Journal of Lightwave Technology*, 23(4), 1745–1753.
22. Takagi, H. (2000). Analysis and application of polling models. *Performance Evaluation*, 423–442.
23. Winands, E., Adan, I., & van Houtum, G. (2006). Mean value analysis for polling systems. *Queueing Systems*, 54(1), 35–44.
24. Wolff, R. W. (1982). Poisson arrivals see time averages. *Operations Research*, 30, 223–231.



M. Thanh Ngo received his BE in Computer Science from INSA, Toulouse, France, in 2005, MESC in Networks and Telecommunications at INSA, Toulouse, France in the same year. He prepared his PhD in Computer Science and Telecommunications at Telecom Bretagne, France, and received his degree from Evry University in 2009. His research interests are 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 Engineering Department at Telecom Bretagne, in Brest (France). Before year 2000, 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 and metro networks (optical, wired and wireless).



Deepak Bhaduria received the BTech Degree in Computer Science from Indian Institute of Technology, Guwahati, India, in 2007. He is currently a PhD degree candidate at University of Minnesota, Twin Cities, USA. His research interests include motion planning, graph theory and approximation algorithms.