Toward Quality of Service Protection in Ethernet Passive Optical Networks: Challenges and Solutions

Chadi Assi, Martin Maier, and Abdallah Shami

Abstract

This article surveys various solutions proposed to date to support quality of service in EPONs. Namely, we overview the inter-and intra-ONU scheduling solutions and argue, and later show, that although these solutions can help in supporting QoS, they fall short of protecting the requirements of the admitted traffic, especially as the network becomes saturated. Hence, we present an admission control framework that is capable of supporting and protecting the QoS of real-time traffic while guaranteeing a minimum bandwidth for best effort traffic, and present an appropriate dynamic bandwidth allocation algorithm. Finally, we present numerical results to show the effectiveness of the proposed solution and compare the performance with that of intra-ONU scheduling solutions.

Ethernet passive optical network (EPON) [1] represents the convergence of inexpensive and ubiquitous Ethernet equipment with low-cost fiber infrastructure; it is viewed as an attractive solution for the broadband access network bottleneck and has been under intensive research activities recently [1–3].

The EPON is a point-to-multipoint access network; it has been standardized by the IEEE 802.3ah Task Force, and comprises one optical line terminal (OLT) and a number of optical network units (ONUs). The OLT resides at the provider central office and connects the access network to the metropolitan or wide area network. The ONU is usually located at either the curb (i.e., fiber to the curb [FTTC]) or the end-user location (i.e., fiber to the building [FTTB] and fiber to the home [FTTH]), and provides broadband video, data, and voice services to end customers. EPON systems currently deploy only one channel for downstream traffic and another for upstream traffic. In the downstream direction, Ethernet frames are broadcast by the OLT and selectively received by each ONU. However, it is worthwhile noting that roughly 95 percent of deployed lines are covered with a video overlay service, and additional channels are reserved for future applications, such as the optical time domain reflectometer (OTDR) at 1600 nm. In the upstream direction multiple ONUs share the same channel to transmit control and data packets to the OLT. Since ONUs are unable to detect collisions and due to the difficulty of implementing carrier sense multiple access with collision detection (CSMA/CD), it is necessary to design a mechanism that arbitrates the access of ONUs to the shared medium; this is achieved by designing an appropriate medium access control (MAC) protocol (e.g., MPCP).

Currently, broadband access providers view quality of service (QoS) and multimedia-capable networks as essential to offer residential customers video on demand, audio on demand, voice over IP (VoIP), and high-speed Internet access. Furthermore broadband access networks, and EPONs in particular, are especially appropriate for peer-to-peer (P2P) applications. The authors of [4] have shown that P2P applications represent a high fraction of the upstream traffic in hybrid fiber-coax cable access networks. Unlike early file sharing applications (e.g., Napster and Gnutella), many recent P2P applications include live media broadcasting, high-bandwidth content distribution, and real-time audio conferencing, and require high-performance access networks in order to deliver satisfying QoS to users. Hence, in order to provide QoS in the access network, bandwidth management on the upstream channel is essential.

Various inter-ONU and intra-ONU scheduling approaches have recently been proposed in order to enable the support of QoS. However, in order to support and "protect" the QoS of real-time traffic, one needs, in addition to bandwidth allocation and service differentiation, an admission control (AC) algorithm that makes decisions on whether or not to admit a new stream at a particular ONU based on the stream's requirements and upstream channel usage conditions. This problem of QoS protection is significant in access networks [5] and particularly in EPON. Furthermore, appropriately controlling the admission of real-time traffic will prevent malicious users from manipulating the upstream channel by sending more traffic into the network than their service level agreement (SLA) allows.

This article surveys various research work related to supporting QoS in EPONs through either inter-ONU or intra-ONU scheduling. We briefly review the MAC in EPONs. We provide an up-to-date overview of dynamic bandwidth allocation (DBA) algorithms for EPONs. We present a framework for admission control in EPONs and present plausible solutions. We provide numerical results and conclude the article.

Medium Access Control in EPONs

In the upstream direction, all ONUs share the same transmission medium; therefore, EPON systems must employ a MAC mechanism to arbitrate access and avoid collisions [3]. Conventional contention-based multiple access (e.g., CSMA/CD) is difficult to implement because ONUs are unable to easily detect a collision that may occur at the OLT. Although the OLT is able to detect a collision and inform the ONUs by sending a collision message, transmission efficiency would be greatly affected. (For the sake of completeness, we mention that research on decentralized MAC has begun recently. To enable distributed medium access, however, the original EPON architecture has to be modified such that each ONU's upstream transmission is echoed at the splitter to all ONUs. In doing so, all ONUs are able to monitor the transmission of every ONU and arbitrate upstream channel access in a distributed manner, similar to CSMA/CD as used in Ethernet LANs. For further information on decentralized access control in EPON networks we refer the interested reader to [6, 7].) One possible solution is to use wavelength-division multiplexing (WDM) technology and allow each ONU to operate at a different wavelength, thus avoiding interference from transmissions of other ONUs. This solution is simple to implement but requires a tunable receiver or an array of receivers at the OLT to receive data transmitted on multiple channels. Additionally, it requires each ONU to be either equipped with a tunable transceiver or a transceiver that operates at a fixed and network unique wavelength. The former solution is costly, while the latter may result in an inventory problem. Alternatively, time-division multiple access (TDMA) on a single wavelength appears to be more attractive. Here, each ONU is allocated a time slot or transmission window (TW, fixed or dynamic) for data transmission by the OLT and the ONUs are scheduled periodically for transmission either using a simple round robin or a more elaborate scheduling scheme. Each TW carries multiple Ethernet packets; packets received from one or more users are buffered at each ONU until the time window for that ONU arrives. Upon the arrival of its turn, the ONU will send out its buffered packets at full transmission rate. Accordingly, the TDMA solution avoids data collisions from different ONUs while it requires only a single wavelength and a single transceiver at the OLT, which makes it highly cost effective. The OLT uses the services offered by the multipoint control protocol (MPCP) to communicate the assigned TWs to their appropriate ONUs [3]. MPCP is a signaling protocol used for periodic bandwidth scheduling. The OLT gathers information from different ONUs and allocates them bandwidth through the use of REPORT and GATE messages of the MPCP protocol. Within each cycle, ONUs use REPORT messages to report their bandwidth requirements (e.g., buffer occupancy) to the OLT. Upon receiving REPORT messages, the OLT performs the appropriate bandwidth allocation computation and sends a GATE message to each ONU, containing the appropriate transmission grants; this is referred to as *dynamic bandwidth allocation* (DBA). Unlike dynamic bandwidth allocation, in static bandwidth allocation, each ONU is allocated a fixed portion of the upstream bandwidth regardless of its bandwidth requirements; static allocation however is not useful especially when the network traffic is quite bursty.

Recently, research on WDM upgrades of single-channel TDM EPONs has been gaining momentum. Several DBA algorithms for WDM enhanced EPONs have been proposed and investigated. Among others, a multichannel extension of the above mentioned IPACT, termed WDM IPACT with a single polling table (WDM IPACT-ST), and the so-called Byte Size Clock (BSC) dynamic wavelength allocation (DWA) algorithms have received attention. For an overview of recently proposed WDM EPONs and DWA algorithms as well as WDM extensions to MPCP that provide backward compatibility with TDM EPONs and future-proofness against arbitrary WDM ONU structures, we refer the interested reader to [2].

Dynamic Bandwidth Allocation Algorithms

Various DBA algorithms have been proposed; they can be categorized into algorithms with statistical multiplexing, such as the Interleaved Polling with Adaptive Cycle Time (IPACT) [1] and its various extensions and algorithms with QoS assurances. In the former (i.e., IPACT), the OLT polls the ONUs individually and issues transmission grants to them in a round robin fashion. In the latter, the algorithms are further subdivided into algorithms with absolute QoS assurances, such as bandwidth guaranteed polling (BGP) [8] and deterministic effective bandwidth [9], and relative QoS assurances [10]. Limited sharing with traffic prediction (LSTP) for dynamic bandwidth allocation is presented in [11] wherein prediction techniques for traffic arrival during waiting times is used in order to limit packet delays. For a detailed review of the various DBA algorithms, see [2].

In addition to bandwidth allocation at the OLT (inter-ONU scheduling), each ONU may also deploy a local scheduling discipline to transmit packets, for example, according to their priorities. Initially, each ONU upon receiving packets from different traffic streams of end users performs three main operations [12]. First, it classifies every newly arriving packet using a "packet-based" classifier. Next, and before placing packets in the corresponding priority queues, the ONU may decide whether a packet should be admitted depending on the adopted traffic policing mechanism (e.g., leaky bucket). Finally, the ONU will schedule packets from its queues for transmission in the assigned transmission window.

There are two types of intra-ONU scheduling: strict and non-strict scheduling [10, 13]. In strict priority scheduling, a lower-priority queue is scheduled only if all queues with higher priority are empty. However, under light load conditions, and when bandwidth is abundant, the strict priority scheduling may result in low-priority queue starvation, a phenomenon also known as *light-load penalty* [13]. The authors of [13] proposed two methods to deal with this problem: a two-stage queuing method and a traffic prediction method. Non-strict priority scheduling, on the other hand, addresses this problem by allowing reported packets (regardless of their priority) to be transmitted first [10]. Furthermore, the transmission order of different queues is based on their priorities; as a result, all traffic classes have access to the upstream channel, which enables fairness in scheduling. The authors of [14] proposed a new intra-ONU scheduling scheme named Modified Start-Time Fair Queuing (M-SFQ). Here, the scheduler selects for transmission the queue with the minimal start time, derived from the head-of-line (HOL) packet in each queue, and synchronized with a global virtual time. Kramer *et al.* [15] proposed a new hierarchical scheduler that fairly divides the excessive bandwidth resulting from lightly loaded ONUs among priority queues (PQs) from different ONUs. An intra-ONU scheduling approach that is based on deficit weighted round-robin (DWRR) is presented in [12] to achieve adaptive fairness. Here, each ONU can adaptively set the weights of its queues; accordingly each class is guaranteed to receive a fair access to the upstream bandwidth. Most recently [16], an OLT-centric DBA that employs a credit pooling technique combined with a weighted-share policy of the upstream channel was proposed. The scheme provides superior fairness among various CoS of different ONUs; the authors argue that

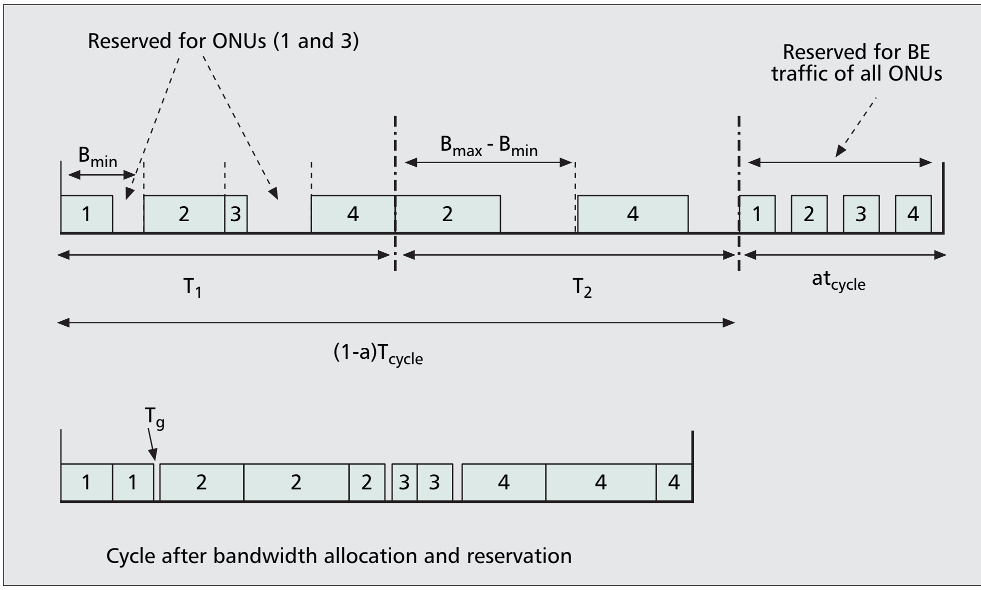

■ Figure 1. *Proposed cycle framework.*

since the OLT has global control and knowledge about the traffic requirements of each ONU, it is capable of allocating bandwidth in an efficient manner, and hence the ONU is only left to perform buffer management and request/receive grants from the OLT. In order words, both inter-and intra-scheduling are performed now at the OLT.

Finally, the authors of [9] noted that in order to satisfy real time services with heterogeneous QoS characteristics, it is very important to provide QoS guaranteed network access while utilizing the bandwidth efficiently. They therefore proposed a dual Deterministic Effective Bandwidth-Generalized Processor Sharing (DEB-GPS) scheduler to provide delay-constraint and lossless QoS guaranteed services and maximize the bandwidth for best effort. The work we present in this article is similar to that of [9] in determining an effective bandwidth for real-time services to meet their QoS requirements. However, unlike the work of [9], we propose a method for protecting the QoS of admitted real-time streams by committing the effective bandwidth once a flow is admitted.

Admission Control in EPON

In order to support and protect the QoS of real-time streams, one needs, in addition to bandwidth allocation (inter-ONU) and service differentiation and scheduling (intra-ONU), an admission control algorithm which makes a decision on whether or not to admit a new real-time traffic stream based on its requirements and the upstream channel usage condition [19]. Admission control further helps in protecting the QoS of existing traffic and admits new flows only if their QoS requirements can be guaranteed.

We note that the problem of QoS protection is not trivial because the bandwidth allocated to every ONU can only be guaranteed for one small cycle (T_{cycle}, $T_{cycle} \leq 2$ms [1]) and may vary from one cycle to another depending on the traffic demand at other ONUs.

Bandwidth reservation may be required to cope with this issue and will provide guaranteed bandwidth to the ONUs; accordingly, each ONU is required to reserve bandwidth for its streams in order to satisfy their QoS requirements. Once this bandwidth is reserved, the OLT can no longer allocate it to other ONUs. Every ONU may be guaranteed a minimum bandwidth (B_{min}) and could be allocated up to a maximum bandwidth (B_{max}) in order to allow other ONUs to also receive their share of the channel [18, 19]. Best effort (BE) traffic shares as well a fraction of the total cycle (e.g., $\alpha \times T_{cycle}$ where $\alpha < 1$). The remainder of the cycle ($(1 - \alpha) \times T_{cycle}$) is used to provide services for bandwidth guaranteed traffic. This new cycle in turn is divided into two subcycles (T_1, T_2); the OLT computes the minimum bandwidth guaranteed for each ONU using T_1:

$$B_{\min} = \frac{(T_1 - N \times T_g) \times \xi}{8 \times N},$$

where ξ is the transmission speed of the PON in megabits per second, N is the number of ONUs, and T_g is a guard time between the transmission of ONU_n and ONU_{n+1}). The ONU has total control over this bandwidth, while the bandwidth of the second subcycle T_2 is under the control of the OLT (please refer to Fig. 1 for a graphical illustration, with $N = 4$). This new system will enable us to implement a two-step admission control; the first is local at the ONU, and the second is global at the OLT (as explained later). Note that, although the minimum guaranteed bandwidth is under the control of the ONU, the scheduling of various ONUs is still done centrally at the OLT in order to achieve a collision free access to the channel. The two sub-cycles are selected of equal length; however, if $T_1 < T_2$, then the OLT will have more control over the bandwidth with less bandwidth guaranteed per ONU. Conversely, the ONU is guaranteed more bandwidth, which may be under-utilized if the load at a particular ONU is not high. Under our assumption of equal lengths for the sub-cycles, we set the maximum bandwidth that a highly loaded ONU can be allocated, $B_{max} = \delta \times B_{min}$ ($\delta = 3$ in Fig. 1). We note that the guaranteed bandwidth per ONU is predetermined and communicated to the ONU or hard coded at the deployment time. Intuitively, one may think that this method may result in bandwidth inefficiency, should one or more ONUs get disconnected for maintenance or some other reason. However, this potential bandwidth efficiency problem can easily be resolved by allowing the service provider (e.g., the OLT) to utilize the bandwidth (committed to the disconnected ONU) to schedule the transmissions of flows from other ONUs or allocate for best effort traffic. However, this bandwidth cannot be used by the ONUs to make admission decisions on new real-time requests.

For real-time applications, QoS metrics can be predefined in a policy control unit (PCU), and various thresholds could be specified/defined. For example, if the expected drop rate or delay requirement for a certain flow/application cannot be respected, the flow should not be admitted. Not only will admitting such a stream make it experience a degraded level of service, but it will also degrade the QoS of existing streams. Alternatively, best effort traffic is never rejected and is always guaranteed a minimal bandwidth (B_{min} BE). In every cycle, the ONU reports (using the MPCP protocol) to the OLT the BE buffer occupancy for bandwidth allocation in the next cycle; for real-time streams that the ONU has already admitted, the OLT will schedule only their transmission since the bandwidth of each stream can already be predetermined and reserved, and it is guaranteed per cycle for the rest of the lifetime of the stream.

With respect to admission, upon the arrival of a new flow a decision should be made according to both admission policies and QoS requirements, often supplied by the application layer. The set of parameters that characterize the traffic stream vary from one traffic class to another. For example, while CBR traffic is non-bursty and characterized by its mean data rate (μ), which makes it quite predictable, it requires stringent packet delays and delay variations. Alternatively, VBR traffic is quite bursty and may be characterized by: mean data rate (μ in bits per second), peak arrival data rate (δ in bits per second), maximum burst size (ρ in bits), and delay bound (θ). Finally, BE traffic is bursty and requires neither delay requirements nor guaranteed bandwidth (note that network operators may set a certain minimum bandwidth that should be guaranteed for BE traffic; e.g., by appropriately adjusting α).

When these parameters are specified by the end user [17], the problem left for the admission control unit (ACU, which is either at the ONU or OLT) is simply to determine whether a new stream should be admitted and whether its QoS requirements can be guaranteed while the QoS requirements of the already admitted streams can be protected. For constant bit rate (CBR) traffic, the admission decision is straightforward: if the mean data rate can be supported, the stream is admitted. Hence, enough bandwidth per cycle should be reserved to guarantee the stream data rate. Here, the average delay of CBR traffic is guaranteed to be bounded by the length of cycle. For variable bit rate (VBR) traffic, the ACU may decide to admit a stream only if its peak rate can be supported (for the best QoS) or may admit the stream as long as the mean data rate is available [5]. The former approach ends up admitting few streams, and the latter approach barely supports QoS for bursty streams. Therefore, a guaranteed bandwidth (g_i) based on the traffic parameters for a flow i could be derived upon regulating the stream (e.g., using a dual token bucket filter; note that for CBR traffic, $g_i = \mu_i$). Consequently, conventional rate-based admission control can be used to determine whether a new stream can be admitted or not. For example, if S_j^{TW} is the bandwidth (bits per second) allocated and reserved for ONU j, a new flow $i + 1$ could be admitted if

$$g_{i+1}^{j} + \sum_{i=1}^{h_j} g_i^{j} \le S_j^{TW}, \quad (1)$$

where h_j is the number of real-time streams (CBR or VBR) at ONU j.

The challenge, however, comes from the fact that in EPON the bandwidth assigned per ONU is not guaranteed. Hence, we propose a two-step admission control scheme that will provide bandwidth guaranteed for each real-time stream. First, since each ONU is guaranteed a minimum bandwidth per cycle, B_{min}, the ONU can locally perform rate-based AC according to the bandwidth requirement of the new arriving flow and the bandwidth availability. For example, if g_f^j is the guaranteed rate for a flow f arriving at ONU j, the bandwidth requirement (in bytes) per cycle for the new flow is $R_f^j = g_f^j \times T_{cycle}$. Therefore, this new flow will be admitted if $R_j + \Sigma_{fi}^{hj} \le R_j \le B_{min}$, where h_j is the total number of flows already admitted by the ONU. If the condition is satisfied, the ONU will conditionally admit the flow and monitor its QoS for a predefined number of cycles (e.g., for 20 ms) [19]. If the requirements of the newly admitted flow are satisfied and the QoS of existing flows remains intact, then the flow is ultimately admitted; otherwise, it is dropped.

When a flow f cannot be admitted locally at the ONU (due to bandwidth insufficiency), the ONU reports the arrival of a new flow to the OLT. This reporting is assumed to be done over the MPCP protocol. Although the MPCP protocol currently does not provide mechanisms for such information reporting, it is assumed that the protocol can be extended by exploiting the optional fields reserved for future applications in the REPORT message to enable for such signaling.

The OLT may admit this new flow only if there is bandwidth available in the second subcycle (T_2) and the ONU sending the request has not been allocated more than B_{max}. Hence, the OLT maintains a variable for every ONU designating the bandwidth allocated so far, $B_{alloc}^j = \Sigma_{i=1}^{hj} R_i^j$, where R_i^j denotes the bandwidth guaranteed for already admitted h_j flows. The OLT maintains as well another variable that indicates the bandwidth still available, B_{avail}, (i.e., not committed yet) in T_2. The new flow may be admitted if the following two conditions hold simultaneously:

- $R_f^j + \Sigma_{i=1}^{hj} R_i^j$
- $R_f^j \le B_{avail}$

If both conditions are satisfied, the OLT will conditionally admit the new flow and monitor its QoS parameters for the subsequent n cycles. A flow will be rejected if at least one of the above two conditions is not satisfied. Upon admitting a flow, the OLT will reserve additional bandwidth for the ONU and update the total available bandwidth: $B_{avail} = B_{avail} - R_f^j$. Similarly, when a flow terminates, the ONU reports to the OLT, and the latter will update the available bandwidth accordingly: $B_{avail} = B_{avail} + R_f^j$.

Admission-Control-Enabled DBA Scheme (AC-DBA)

As presented earlier, every stream, if admitted, is guaranteed a bandwidth per cycle that is computed based on the derived guaranteed or effective rate of the flow. The OLT then allocates a transmission window for each ONU computed according to the bandwidth requirements of its admitted flows (CBR and VBR) and *not* according to its actual buffer occupancy. This, however, may either result in a waste of bandwidth (if the ONU does not have enough traffic to transmit in the assigned window), or, alternatively, yield increased delays if the ONU has more traffic to send than can be accommodated in the assigned window. This issue is attributed to the burstiness of the real-time traffic (i.e., VBR) and to the fact that the bandwidth allocation is no longer *on demand*. Furthermore, estimating the bandwidth requirement of a flow based on its guaranteed rate, although sound statistically, may not reflect the real need of a flow especially that the bandwidth needs for a flow is estimated for a short period of time (i.e., the cycle) and hence the inefficiency in the bandwidth prediction and reservation.

To resolve the above problems, we propose a two-branch solution [18]. First, the OLT selects a super-cycle ($T_{sc} = \lambda \times T_{cycle}$, λ is a constant), and every admitted flow is guaranteed a bandwidth per T_{sc} instead, while the allocation by the OLT is done every T_{cycle}. The purpose of this proposal is to mitigate the inefficiency of the bandwidth reservation caused by the short-time prediction. The period $(1 - \alpha) \times T_{sc}$ is divided into two periods, T_1 and T_2; each ONU is now guaranteed a bandwidth of B_{min}^{new} which is computed based on T_1. The OLT controls the remaining bandwidth of the super-cycle. Upon the arrival of a new flow f at ONU j with bandwidth guaranteed B_g^f, the flow is either admitted or rejected as described earlier. Second, we ensure that the reservation does not waste any bandwidth. Here, we apply a *crediting system*; every time a flow is admitted, the OLT will be informed, and it will estimate a total credit (number of bytes available per T_{sc} for this flow) $C_{fi}^j = B_g^{fi} \times T_{sc}^r$ where T_{sc}^r is the period between the arrival of the flow and the end of the current super-cycle. The OLT maintains as well a total credit per type of traffic (C_{CBR}^j for CBR and C_{VBR}^j for VBR) per ONU; for example, $C_{CBR}^j = \Sigma_{i=1}^{Nj} C_{fi}^j$ where N_j is the number of CBR flows at ONU $_j$. In every cycle, the OLT deducts the requested/allocated bandwidth of this flow from its reserved credits until the next super-cycle; at this point, the credits are reset to the estimated ones.

Now with respect to bandwidth allocation, as with conventional DBAs, the ONU reports (in cycle $n - 1$) to the OLT its buffer occupancy ($Q_{CBR}(n - 1)$, $Q_{VBR}(n - 1)$, and $Q_{BE}(n - 1)$) and requests transmission bandwidth accordingly. Let $A_{CBR}^j(n)$, $A_{VBR}^j(n)$, $A_{BE}^j(n)$ be the bandwidth allocated for ONU j (whose values will be shortly determined); then the two following conditions should be met:

$$\sum_{j=1}^{N} (A_{CBR}^j(n) + A_{VBR}^j(n)) \le B_{cycle} - T_{gt}^t - (N \times B_{BE}^{\min}) \quad (2)$$

$$\sum_{j=1}^{N} A_{BE}^j(n) \le N \times B_{BE}^{\min} \quad (3)$$

where B_{cycle} is the total bandwidth available in T_{cycle} and T_{gt}^t is the total guard time (in bytes) between ONUs transmissions and B_{BE}^{min} is the minimum bandwidth guaranteed for BE traffic. Every time the OLT allocates bandwidth to one ONU, it will adjust its remaining credits; if the credits for a particular traffic class at one ONU are depleted, the OLT can no longer allocate any bandwidth for this class and its transmission will be deferred to the next supercycle. As for the computation of the available bandwidth for each class, the OLT waits until all requests are received from all ONUs. If $\Sigma_{j=1}^{N} (Q_{CBR}^j(n-1) + Q_{VBR}^j(n-1)) \le B_{cycle} - T_{gt}^t - N \times B_{BE}^{min}$, then $A_{CBR}^j(n) = min(Q_{CBR}^j(n-1), C_{CBR}^j(n-1))$and the credits are updated accordingly; similarly for VBR traffic. Otherwise, the OLT will compute a total guaranteed bandwidth for each ONU, for example, according to the requests from all ONUs:

$$B_j(n-1) = \frac{R_j(n-1) \times (B_{cycle} - T_{gt}^t - (N \times B_{BE}^{\min}))}{\sum_{j=1}^{N} R_j(n-1)},$$

$$R_j(n-1) = Q_{CBR}^j(n-1) + Q_{VBR}^j(n-1),$$

and then allocate bandwidth for each traffic class.

Next, the OLT will allocate bandwidth to BE traffic; the total BE bandwidth per cycle is $B_{BE} = N \times B_{BE}^{min}$, which is shared among all ONUs, and the bandwidth assigned to each ONU could be determined according to all bandwidth requests. Once the bandwidth for each traffic class is determined, the OLT will send one GATE message to inform the ONU of this bandwidth allocation and schedule its transmission accordingly. The GATE message will contain three grants, one for each class of service.

Numerical Results

In this section we study the performance of the proposed AC framework using simulations; we compare its performance with two intra-ONU scheduling schemes: strict priority (SP) and modified deficit weighted round robin (M-DWRR) [12]. We consider a realistic traffic profile where real-time streams (voice and video) and BE traffic arrive dynamically at the ONUs (Fig. 3a). A voice (video) flow is generated at a mean rate of 64 kb/s (*guaranteed rate* of 4 Mb/s) with a delay bound θ_{CBR} = 2~4 ms (θ_{VBR} = 25~30 ms) and each BE flow at a mean rate of 5 Mb/s. Voice traffic is modeled by a CBR source, and video traffic is modeled using a VBR source. VBR and BE traffic are highly bursty, and we use self-similar traffic for modeling these classes; packet sizes are uniformly distributed between 64 and 1518 bytes. Alternatively, a Poisson distribution can approximately model CBR traffic, and the packet size is fixed to 70 bytes. Here, the load increases incrementally as more flows are admitted into the network. We choose B_{BE}^{min} = 4100 bytes (in each cycle), which means that each ONU is guaranteed a BE throughput of 15 Mb/s. The total number of ONUs N =16, and the PON speed = 1 Gb/s. The guard time between different transmission windows is equal to 1 ms, the cycle time T_{cycle} = 2 ms, the ONU buffering queue size to 10 Mbytes, and we select T_{sc} = 500 ms for our simulations.

Figure 2 shows the instantaneous average packet delays. Clearly, using M-DWRR and SP schedulers, CBR traffic shows the optimal performance where its average packet delay remains under 2 ms even when the load continuously increases (i.e., as the simulation time continues to increase). This shows the advantage of M-DWRR; that is, although it divides the cycle among the various traffic classes based on their preassigned weights [12], it also provides optimal performance for CBR traffic. This is due to the fact that the assigned weights are adaptively set based on the QoS requirements. Under strict priority, the scheduler always selects packets from a higher-priority queue until satisfied (i.e., until it is empty); therefore, CBR traffic will exhibit the best performance. Alternatively, AC-DBA (Fig. 2a) makes sure to satisfy the QoS requirements defined previously (in terms of delay bound and throughput) by *crediting* every real-time traffic the appropriate bandwidth and reserving it in every supercycle/cycle, since a CBR flow is admitted only if its guaranteed bandwidth is assured in every cycle. Hence, AC-DBA maintains a CBR average packet delay of 2~4 ms with a noticeable slight decrease pattern that repeats every supercycle. As for VBR traffic, AC-DBA maintains a respectable delay performance (Fig. 2d) that meets the specified target QoS requirements of the streams (i.e., 25~30 ms), while this delay experiences an exponential increase under intra-ONU scheduling schemes that do not deploy any AC (Figs. 2e and 2f). This behavior actually emphasizes the need for admission control in EPONs, because when the system reaches a saturation point and the network keeps admitting all arriving streams, the QoS performance can no longer be maintained (not only for new applications but for existing applications as well).

We further investigate the performance with respect to satisfying a minimal bandwidth requirement for BE traffic. We measure total throughput received by BE traffic; as shown in Fig. 3b, BE throughput increases to reach a total of ≈ 400 Mb/s under all schemes (i.e., with and without AC) when the load is low and decreases when more flows are admitted into the network. However, when the system reaches saturation, AC-DBA

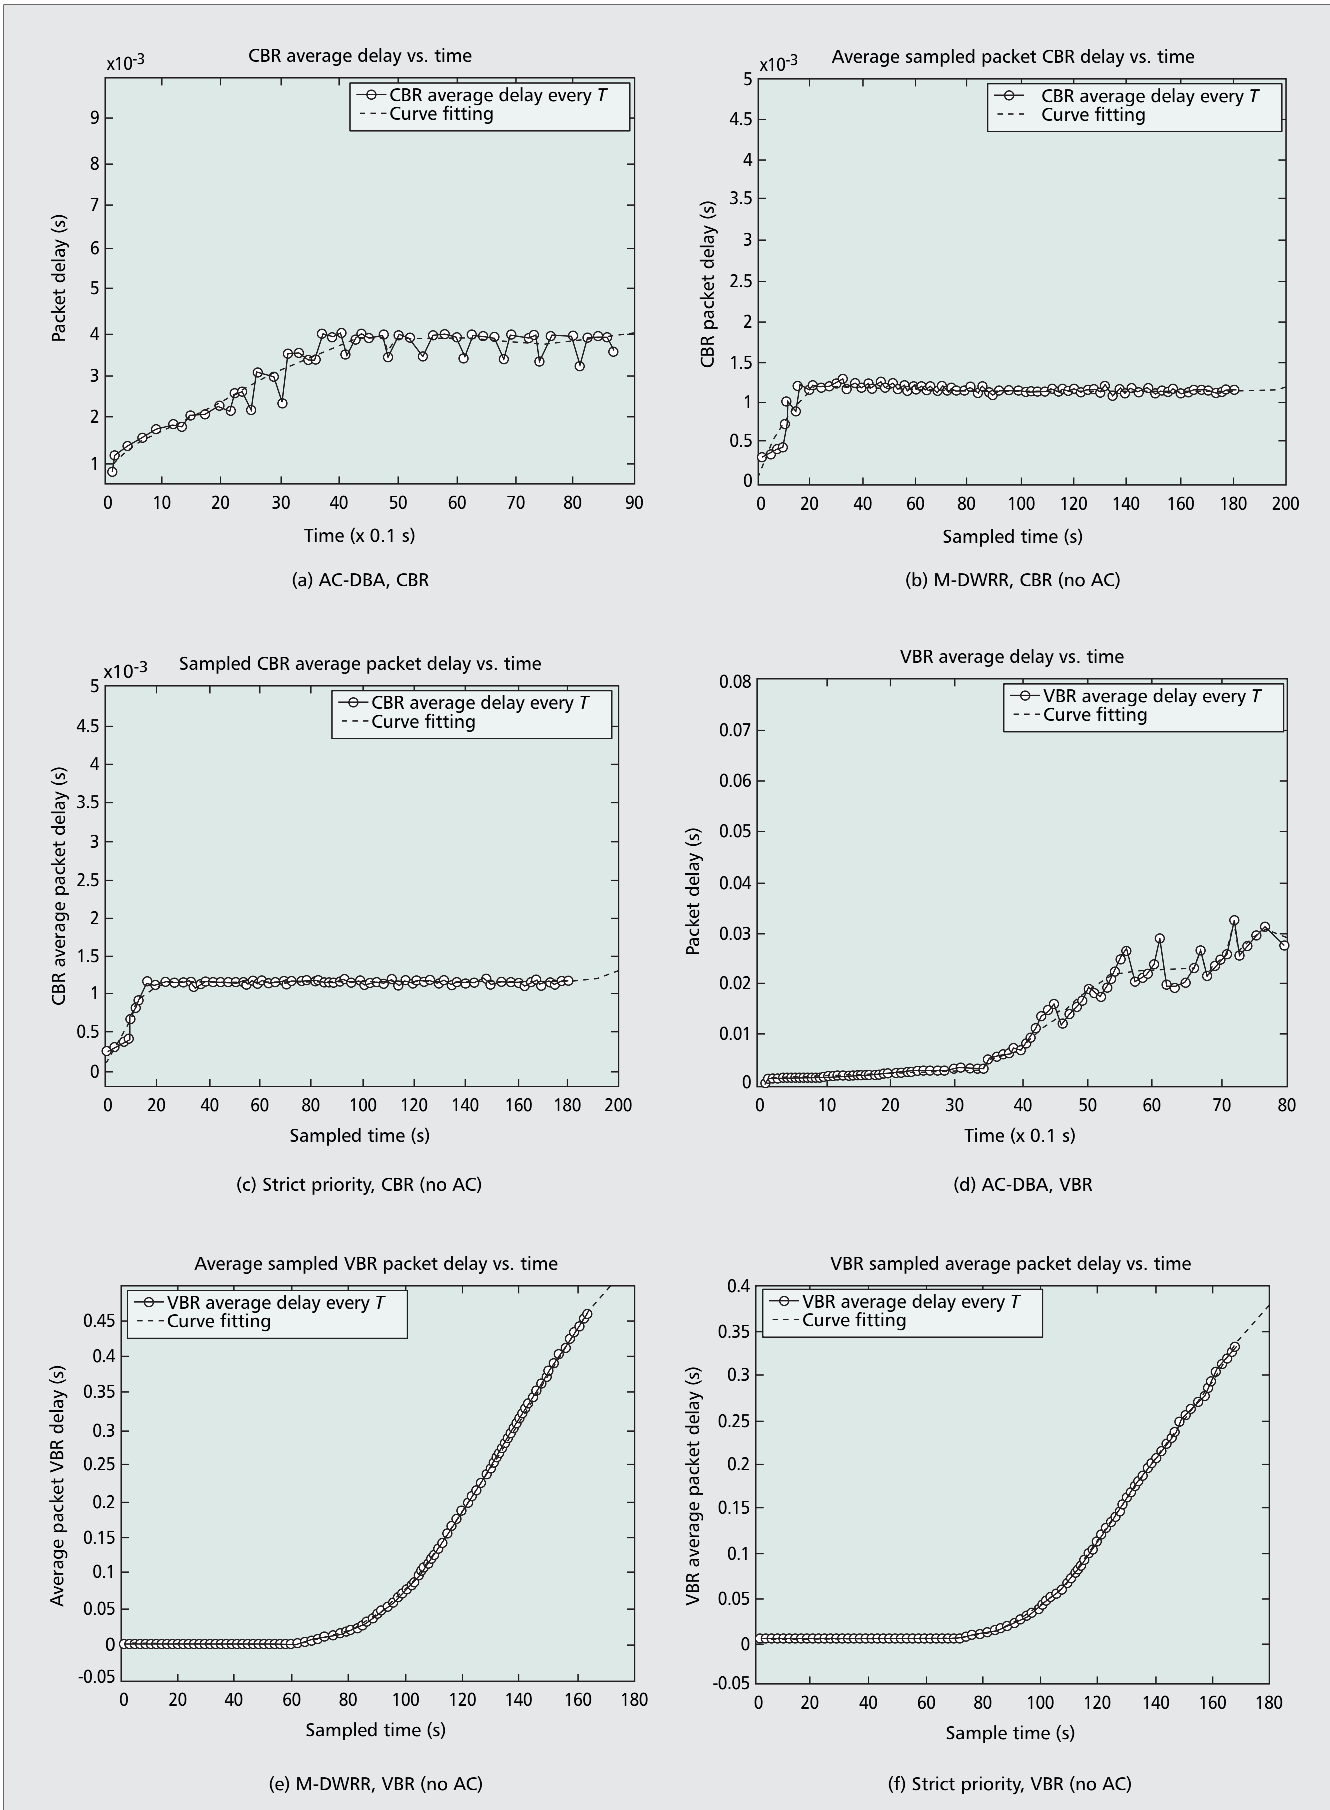

■ Figure 2. *Average packet delay [18].*

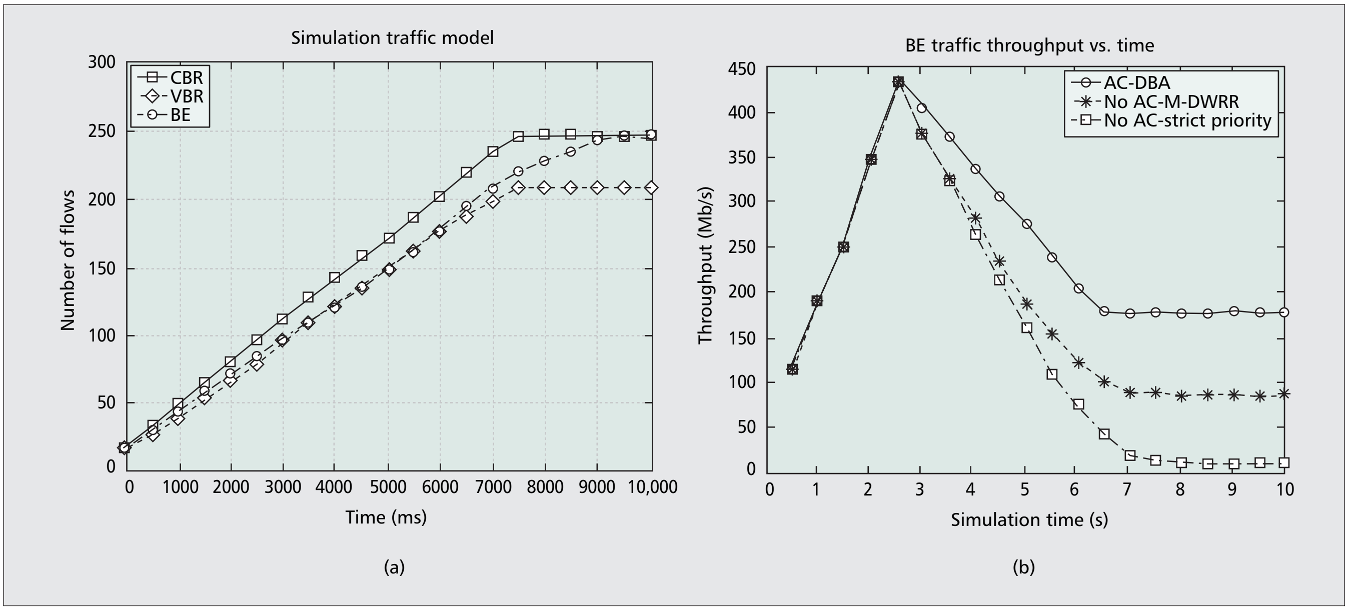

■ Figure 3. *a) Traffic model ; b) BE throughput [18].*

Number of generated CBR flows	252
Number of admitted CBR flows	234
Number of rejected CBR flows	18
CBR admission rate	≈ 92%
Number of generated VBR flows	209
Number of admitted VBR flows	173
Number of rejected VBR flows	36
VBR admission rate	≈ 83%
Number of generated BE flows	247
Number of admitted BE flows	247
Number of rejected BE flows	0
BE admission rate	100%

■ Table 1. *Traffic control statistics [18].*

makes sure to preserve the minimum predefined throughput; while with M-DWRR and strict priority schedulers, the throughput is not guaranteed. Nevertheless, M-DWRR still provides a minimum throughput (which is one of the advantages of M-DWRR) by forcing the weight policy, while it reaches a very low one (≈1 Mb/s) with SP, a phenomenon known as BE traffic starvation. Finally, Table 1 shows some interesting statistics collected from our simulations. These results show that ≈ 92 percent of the generated CBR traffic are admitted into the network while their overall QoS and bandwidth requirements are guaranteed; ≈ 83 percent of VBR flows are admitted as well; and finally, all BE flows arriving are admitted. Note that under M-DWRR and SP scheduling, all traffic is admitted; however, as shown, their QoS requirements are not guaranteed (except for CBR traffic).

Conclusion

In this article we present a brief survey of the various solutions to support QoS in EPON networks; namely, the scheduling (inter- and intra-ONU) solutions that have been adopted thus far. However, we show that these solutions, although they support respectable QoS for real-time traffic, are unable to protect the QoS of existing traffic when the network gets saturated. Admission control, in addition to scheduling and bandwidth allocation, is a solution presented in this article; we show that a network deploying this solution can indeed protect the QoS of admitted real-time traffic, as opposed to solutions that rely only on traffic scheduling.

Acknowledgments

The authors would like to thank the reviewers for their constructive and helpful comments. The authors would also like to thank Ahmad Dhaini for his work on simulating the AC-DBA.

References

[1] G. Kramer, B. Mukherjee, and G. Pesavento, "IPACT: A Dynamic Protocol for an Ethernet PON (EPON)," *IEEE Commun. Mag.*, vol. 40, no.2, Feb. 2002, pp. 74–80.
[2] M. P. McGarry, M. Maier, and M. Reisslein, "WDM Ethernet Passive Optical Networks," *IEEE Commun. Mag.*, vol. 42, no. 8, Feb. 2006, pp. S18–S25.
[3] M. P. McGarry, M. Maier, and M. Reisslein, "Ethernet PONs: A Survey of Dynamic Bandwidth Allocation (DBA) Algorithms," *IEEE Commun. Mag.*, vol. 44, no. 2, Aug. 2004, pp. S8-S15.
[4] J. Zheng and H. T. Mouftah, "Media Access Control for Ethernet Passive Optical Networks: An Overview," *IEEE Commun. Mag.*, vol. 43, no. 2, Feb. 2005, pp. 145–50.
[5] M. Garcia *et al.*, "Analysis and Modeling of Traffic on a Hybrid Fiber-Coax Network," *IEEE JSAC*, vol. 22, no. 9, Nov. 2004, pp. 1718–30.
[6] C.-T. Chou, S. N. Shankar, and K. G. Shin, "Achieving Per-Stream QoS with Distributed Airtime Allocation and Admission Control in IEEE 802.11e Wireless LANs," *Proc. IEEE INFOCOM*, Miami, FL, Mar. 2005, vol. 3, pp. 1584–95.
[7] C. H. Foh et al., "FULL-RCMA: A High Utilization EPON," *IEEE JSAC*, vol. 22, no. 8, Oct. 2004, pp. 1514–24.
[8] S. R. Sherif *et al.*, "A Novel Decentralized Ethernet-Based PON Access Architecture for Provisioning Differentiated QoS," *IEEE/OSA J. Lightwave Tech.*, vol. 22, no. 11, Nov. 2004, pp. 2483–97.
[9] M. Ma, Y. Zhu, and T. H. Cheng, "A Bandwidth Guaranteed Polling MAC Protocol for Ethernet Passive Optical Networks," *Proc. IEEE INFOCOM*, San Francisco, CA, Mar./Apr. 2003, vol. 1, pp. 22–31.
[10] L. Zhang *et al.*, "Dual DEB-GPS Scheduler for Delay-Constraint Applications in Ethernet Passive Optical Networks," *IEICE Trans. Commun.*, vol. E86-B, no. 5, May 2003, pp. 1575–84.

[11] C. M. Assi *et al.*, "Dynamic Bandwidth Allocation for Quality-of-Service over Ethernet PONs," *IEEE JSAC*, vol. 21, no. 9, Sept. 2003, pp. 1467–77.
[12] Y. Luo and N. Ansari, "Bandwidth Allocation for Multiservice Access on EPONs," *IEEE Commun. Mag.*, vol. 43, no. 2, Feb. 2005, pp. S16–S21.
[13] A. R. Dhaini *et al.*, "Adaptive Fairness through Intra-ONU Scheduling for Ethernet Passive Optical Networks," *Proc. IEEE ICC*, Istanbul, Turkey, June 2006, vol. 6, pp. 2687–92.
[14] G. Kramer *et al.*, "Supporting Differentiated Classes of Service in Ethernet Passive Optical Networks," *OSA J. Opt. Net.*, vol. 1, no. 9, Aug. 2002, pp. 280–98.
[15] N. Ghani *et al.*, "Intra-ONU Bandwidth Scheduling in Ethernet Passive Optical Networks," *IEEE Commun. Lett.*, vol. 8, no. 11, Nov. 2004, pp. 683–85.
[16] G. Kramer *et al.*, "Fair Queuing with Service Envelopes (FQSE): A Cousin-Fair Hierarchical Scheduler for Subscriber Access Networks," *IEEE JSAC*, vol. 22, no. 8, Oct. 2004, pp. 1497–1513.
[17] H. Naser and H. T. Mouftah, "A Joint-ONU Interval-Based Dynamic Scheduling Algorithm for Ethernet Passive Optical Networks," *IEEE/ACM Trans. Net.*, vol. 14, no. 4, Aug. 2006, pp. 889–99.
[18] A. R. Dhaini *et al.*, "Per-Stream QoS and Admission Control in Ethernet Passive Optical Networks (EPONs)," *IEEE/OSA J. Lightwave Tech.*, vol. 25, no. 7, July 2007, pp. 1659–69.
[19] A. R. Dhaini, "Design and Analysis of Next Generation Ethernet-Based Passive Optical Access Networks," Master's thesis, Concordia Univ., Sept. 2006.
[20] ITU-T Rec. G.1010, "End-User Multimedia QoS Categories," 2001.

Biographies

CHADI M. ASSI (assi@ece.concordia.ca) is an associate professor with the Concordia Institute for Information Systems Engineering at Concordia University, Montreal, Canada. He received his B.Eng. degree from Lebanese University, Beirut, in 1997 and his Ph.D. degree from the Graduate Center, City University of New York, in April 2003. Before joining the Institute in August 2003 as an assistant professor, he was a visiting scientist for one year at Nokia Research Center, Boston, Massachusetts, working on QoS in optical access networks. He received the prestigious Mina Rees Dissertation Award from the City University of New York in August 2002 for his research on WDM optical networks. His current research interests are in the areas of optical networks, wireless and ad hoc networks, and security.

MARTIN MAIER (maier@ieee.org) is an associate professor at Institut National de la Recherche Scientifique (INRS), Montreal, Canada. He was educated at the Technical University of Berlin, and received M.Sc. and Ph.D. degrees (both with distinction) in 1998 and 2003, respectively. In the summer of 2003 he was a postdoctoral fellow at Massachusetts Institute of Technology, Cambridge. He was a visiting professor at Stanford University, California, from October 2006 through March 2007. His current research activities focus on service-oriented access and metro networks.

ABDALLAH SHAMI received his B.E. degree in electrical and computer engineering from Lebanese University, Beirut, in 1997, and a Ph.D. degree in electrical engineering from the Graduate School and University Center, City University of New York, in September 2002. At that time he joined the Department of Electrical Engineering, Lakehead University, Ontario, Canada, as an assistant professor. Since July 2004 he has been with the University off Western Ontario, London, where he is currently an assistant professor in the Department of Electrical and Computer Engineering. His current research interests are in the area of data/optical networking, EPON, WiMAX, WLANs, and software tools. He held the Irving Hochbert Dissertation Fellowship Award at City University of New York and a GTF Teaching Fellowship.