Per-Stream QoS and Admission Control in Ethernet Passive Optical Networks (EPONs)

Ahmad R. Dhaini, *Student Member, IEEE*, Chadi M. Assi, *Member, IEEE*, Martin Maier, *Member, IEEE*, and Abdallah Shami, *Member, IEEE*

***Abstract*—Ethernet passive optical networks (EPONs) are designed to deliver services for numerous applications such as voice over Internet protocol, standard and high-definition video, video conferencing (interactive video), and data traffic. Various dynamic bandwidth allocation and intra–optical network unit (ONU) scheduling algorithms have been proposed to enable EPONs to deliver differentiated services for traffic with different quality of service (QoS) requirements. However, none of these protocols and schedulers can guarantee bandwidth for each class of service nor can they protect the QoS level required by admitted real-time traffic streams. In this paper, we propose the first framework for per-stream QoS protection in EPONs using a two-stage admission control (AC) system. The first stage enables the ONU to perform flow admission locally according to the bandwidth availability, and the second stage allows for global AC at the optical line terminal. Appropriate bandwidth allocation algorithms are presented as well. An event-driven simulation model is implemented to study the effectiveness of the proposed schemes in providing and protecting QoS.**

***Index Terms*—Access networks, admission control, ethernet passive optical network (EPON), quality of service (QoS).**

I. Introduction

THE 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; EPON is a point-to-multipoint access network with no active elements in the signal's path from source to destination. It has been standardized by the IEEE 802.3ah Working Group [4], and it comprises one optical line terminal (OLT, at the Central Office) and a number of optical network units (ONUs) and provides broadband video, data, and voice services to end customers. EPON systems currently deploy only one channel for downstream traffic and another channel for upstream traffic. In the downstream, Ethernet frames are broadcast by the OLT and are selectively received by each ONU. In the upstream, multiple ONUs share the same transmission channel to transmit data and control packets to the OLT. Since ONUs are unable to detect collision and due to the difficulty to implement a *carrier sense multiple access with collision detection*, it is necessary to design a mechanism that arbitrates the access of ONUs to the shared medium [13]. This is achieved by designing an appropriate medium access control (MAC) protocol. Current MAC supports time-division multiple access (TDMA), where each ONU is allocated a fixed or dynamic time slot to transmit data to the OLT, and each ONU buffers packets received from different subscribers until they are transmitted in the assigned window. Currently, broadband access providers view quality of service (QoS) and multimedia-capable networks as essential ingredients to offer residential customers video-on-demand, audio-on-demand, voice over Internet protocol, and high-speed Internet access. Furthermore, broadband access networks, particularly EPON, are especially appropriate for peer-to-peer (P2P) applications. Garcia *et al.* [3] have shown that P2P applications represent a high fraction of the upstream traffic in a hybrid fiber–coax cable access network. Unlike early file-sharing applications (such as 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 the 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 been recently proposed in order to enable the support of QoS. However, in order to support and "protect" the QoS of real-time traffic streams, one needs, in addition to bandwidth allocation and service differentiation, an admission control (AC) algorithm that makes decision on whether or not to admit a new real-time flow based on its requirements and the upstream channel usage condition. We note that the problem of QoS protection is significant in EPONs because the bandwidth allocated by the OLT to one ONU can only be guaranteed for a significantly short time (e.g., one cycle). Furthermore, appropriately controlling the admission of real-time traffic streams will prevent malicious users from manipulating the upstream channel by sending traffic into the network more than their service level agreement (SLA). Accordingly, in this paper, we will present a suitable AC scheme that may be deployed in EPONs in order to support QoS and protect it and enable the transmission of emerging real-time traffic with guaranteed performance. The rest of the paper is organized as follows. In Section II, we present the recent work related to supporting QoS in EPONs. A solution for intra-ONU

Manuscript received October 6, 2006; revised March 13, 2007.

A. R. Dhaini and C. M. Assi are with the Faculty of Engineering and Computer Science, Concordia Institute for Information Systems Engineering, Concordia University, West Montreal, QC H3G 1M8, Canada (e-mail: a_dhaini@ciise.concordia.ca; assi@ciise.concordia.ca).

M. Maier is with the Institut National de la Recherche Scientifique (INRS)-Énergie, Matériaux et Télécommunications, Montréal, QC H5A 1K6, Canada (e-mail: Martin.Maier@emt.inrs.ca).

A. Shami is with the Department of Electrical and Computer Engineering, University of Western Ontario, London, ON N6A 5B9, Canada (e-mail: ashami@eng.uwo.ca).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2007.899160

scheduling based on the deficient weighted round robin (DWRR) [6] is presented in Section III; this scheme ensures that every class of traffic gets a fair share of the assigned bandwidth at the ONU. Section IV presents our AC scheme with a detailed analysis. Section V presents the performance evaluation, and finally, we conclude in Section VI.

II. Background and Related Work

Dynamic bandwidth allocation (DBA) is deployed at the OLT to assign bandwidth for different ONUs; DBA uses the services offered by the multipoint control protocol (MPCP) to communicate the assigned transmission windows (TWs) to their appropriate ONUs. MPCP is a signaling access protocol that is developed and standardized by the IEEE 802.3ah Task Force [4] and used for inter-ONU bandwidth scheduling. The OLT gathers information from different ONUs and dynamically allocates bandwidth to ONUs through the use of REPORT and GATE messages of MPCP. Within each cycle, ONUs use REPORT messages to report its bandwidth requirements (e.g., buffer occupancy) to the OLT. Upon receiving REPORT messages from the ONUs, the OLT performs the appropriate bandwidth allocation computation and broadcasts a GATE message to each ONU, containing the appropriate transmission grants (transmission start $T_{\rm start}$ and transmission end $T_{\rm end}$). Note that MPCP does not specify any particular bandwidth allocation algorithm. Instead, it is designed to facilitate the implementation of DBA algorithms. Various DBAs have been proposed so far; they can be categorized into algorithms with statistical multiplexing and its various extensions [1] and algorithms with QoS support [2], [8], [10]. For a detailed review about DBA, we refer the reader to [1]. Locally, at the ONU, upon receiving traffic "flows" from end users, the ONU performs three main operations before transmitting its packets on the upstream channel. First, it classifies every newly arriving packet using a "packet-based" classifier. Next, and before placing packets in the corresponding priority queues (PQs), the ONU may decide whether a packet should be admitted, depending on the adopted traffic policing (e.g., AC) mechanism. Finally, the ONU will schedule packets from its queues (also known as intra-ONU scheduling) for transmissions in the assigned TW allocated by the OLT (depending on the inter-ONU scheduling algorithm). There are two types of intra-ONU scheduling: strict priority (SP) and non-SP scheduling algorithms. In SP scheduling, a queue with a lower priority is scheduled only if all queues with a higher priority are empty. However, this may result in a starvation for low-priority traffic, which results in the so-called "light-load penalty." Non-SP scheduling, on the other hand, addresses this problem by allowing reported packets (regardless of their priority) to be transmitted first as long as they are transmitted in the allocated TW [2]. Further, the transmission order of different PQs is based on their priorities. As a result, all traffic classes have access to the upstream channel while maintaining their priorities; this enables fairness in scheduling. On the other hand, Ghani *et al.* [7] 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*, which is derived from the head-of-line (HOL) packet in each queue and synchronized with a *global virtual time*. Kramer *et al.* [5] recently proposed a new hierarchical scheduler that fairly divides the excessive bandwidth that results from lightly loaded ONUs among PQs from different ONUs. Most recently [15], a n OLT-centric DBA, which employs a credit-pooling technique, combined with a weighted-share policy of the upstream channel, was proposed. The scheme provides superior fairness among various classes of service (CoSs) of different ONUs.

III. Decentralized Intra-ONU Scheduling

To date, a wide range of scheduling schemes have been studied [e.g., weighted fair queuing (WFQ), self-clocked fair queuing (SCFQ), start-time fair queuing (SFQ), weighted round robin (WRR), and stratified round robin (SRR)]. One distinguished scheme for achieving fairness with low complexity is the DWRR [6]. In this paper, we propose a modified algorithm (M-DWRR) to enforce fairness among the various CoSs.

A. DWRR Scheduling Discipline

DWRR, as proposed, defines the following three main parameters for each CoS or PQ i:

1) a "weight" α_i that defines the percentage of the output port bandwidth that is allocated to the queue;
2) a "deficit counter" $DC(i)$ that specifies the total number of bytes that the queue is permitted to transmit in each scheduler's visit; the DC saves "credits" remaining from previous scheduling visit and adds them to the DC of the next visit until the queue is empty, and hence, $DC(i)=0$;
3) a "quantum" $Q(i)$ that is proportional to α_i and is expressed in bytes.

First, a round robin (RR) scheduler initializes the deficient counters, $DC(i)=0_{,i=0\ldots x}$, where x is the number of PQs, and then visits each nonempty queue and determines the size (in bytes) of the HOL packet. $Q(i)$ is computed from the available port bandwidth as follows:

$$Q(i)=\lceil \alpha_i \times B_{\rm port}\rceil \tag{1}$$

where $B_{\rm port}$ is the bandwidth that is available on the transmission port (in bytes). Next, the scheduler computes

$$DC(i)=DC(i)+Q(i). \tag{2}$$

At this time, it checks if the HOL packet is greater than $DC(i)$; if yes, it moves to the next queue and "saves" the remaining *credits* in $DC(i)$; otherwise, it will select the packet for transmission and updates its deficient counter as follows:

$$DC(i)=DC(i)-S_i^{\rm HOL} \tag{3}$$

where $S_i^{\rm HOL}$ is the size of the HOL packet in queue i. When queue i is empty, $DC(i)$ is reset to 0, and the pointer of the

RR scheduler moves to the queue with a lower priority. The advantages of DWRR over other schemes are listed as follows.

1) DWRR accurately supports weighted fair bandwidth distribution for CoS queues of variable-length packets.
2) DWRR combines both the class-based queuing approach along with the WRR scheduling scheme.
3) DWRR has a lower complexity than WFQ and can be implemented in hardware.

B. Integrating DWRR With EPON

In EPON, every ONU maintains a number of PQs where incoming packets are classified and queued based on their priorities. Unlike the system that was discussed in Section III-A, in EPON, the ONU accesses the channel during the assigned TW that is specified by T_{start} and T_{length}. Hence, ONU j will compute the quantum for each queue i based on the weight assigned to the queue and the TW allocated by the OLT. Therefore, DWRR will have to set its three defined parameters, namely, $\alpha_{i,j}$, $DC(i,j)$, and $Q(i,j)$, for each queue i. Suppose that the allocated TW is of size S_j (in bytes) and is computed as follows:

$$S_j = \min\left(B_{\min} + B^j_{\text{excess}}, \sum_{i=1}^{x} R_{i,j}\right) \quad (4)$$

where $R_{i,j}$, $i = 1 \ldots x$ is defined as the requested size of each queue i, B^j_{excess} is the excess bandwidth that is allocated to ONU j, and $B_{\min}$ is the minimum guaranteed bandwidth [2]. Then the quantum is computed as follows:

$$Q(i,j) = \lceil \alpha_{i,j} \times S_j \rceil. \quad (5)$$

The update of the deficient counter is computed as in (2). Note that $Q(i,j)$ can be set by the OLT and incorporated in the GRANT message.

C. Modified DWRR (M-DWRR)

As mentioned earlier in this paper, the DWRR scheduling discipline visits each PQ in an RR fashion. Moreover, after each visit made by the scheduler to all PQs, the deficient counter is updated according to the rules explained in Section III-A On the other hand, in M-DWRR, once the scheduler has finished visiting all the queues, the remaining bandwidth from the assigned TW of the current cycle is distributed to all the PQs based on the corresponding weights, i.e.,

$$DC(i,j) = DC(i,j) + \left\lceil \alpha_{i,j} \times B^j_{\text{remain}} \right\rceil \quad (6)$$

where B^j_{remain} is the remaining bandwidth (in bytes) from the assigned TW of the same cycle. This remaining bandwidth is found from the unutilized bandwidth after the first scheduling visit to all PQs. In other words, since the TW is divided among PQs, depending on their weights (and not their needs), some queues might not utilize all their corresponding assigned bandwidth. Thus, in order to eliminate the waste of bandwidth, we reallocate this portion to the PQs based on the same weight assignment. Alternatively, the ONU might follow a different "update scheme" and, hence, revalidates the deficient counters based on a different weight assignment scheme, which may be derived from the different traffic requirements and queues occupancies rather than the original weight agreement. Furthermore, another "update discipline" might be implemented, where $DC(i,j)$ is computed as in (6), but if the allocated bandwidth of a higher priority is not needed (i.e., queue is empty), it will be distributed to the queues with a lower priority. However, since high-priority traffic is delay sensitive and since incoming packets might arrive after the described distribution, the scheduler *must* permit transmission of these packets by setting a *flag* that triggers its pointer, upon the arrival of these packets, to the appropriate queue. In this way, a high-priority traffic delay is preserved, and its jitter is protected. On the other hand, the scheduler might allocate the remaining bandwidth in a traditional RR fashion while assigning bandwidth for each nonempty queue such that this allocated bandwidth is "*just*" equal to the HOL packet of each queue.

The advantage of such a scheme and of DWRR in particular is that each ONU can *adaptively* set (depending on the traffic demand and the SLA) its own weights in both phases (i.e., initially and/or after computing B^j_{remain}). Hence, every CoS is guaranteed to receive locally at the ONU a fair share or a fair access to the bandwidth allocated by the OLT. However, the drawback of this scheme and of other schemes proposed so far is that there is no guarantee that each ONU will get the bandwidth that is required to service its admitted streams while satisfying their QoS requirements. A bandwidth-guaranteed polling (BGP) scheme was proposed in [8] to provide guaranteed QoS; here, the ONUs are divided into bandwidth-guaranteed (e.g., premium subscribers) and best effort (BE) ONUs. However, BGP does not consider the case of multiservice ONUs, where both bandwidth- and QoS-guaranteed and BE users coexist. Further, BGP does not provide any QoS protection for existing streams in a more dynamic environment.

IV. AC IN EPON

A. Preliminaries

In order to provide sustainable QoS in the access network, bandwidth management on the upstream channel is essential. In order to support and protect the QoS of real-time traffic streams, one needs, in addition to bandwidth allocation and service differentiation, an AC algorithm that makes decision on whether or not to admit a real-time traffic stream based on its requirements and the upstream channel usage condition. As we mentioned earlier, the problem of QoS protection is significant because the bandwidth allocated by the OLT to one ONU can only be guaranteed for one cycle. Furthermore, appropriately controlling the admission of real-time traffic will prevent malicious users from manipulating the upstream channel by sending traffic into or requesting bandwidth from the network more than their SLA. Accordingly, AC helps in protecting the QoS of existing traffic and admit new flows only if their QoS requirements can be guaranteed.

In current EPONs, the bandwidth of the upstream channel is shared among different ONUs using a TDMA scheme. The OLT allocates a transmission bandwidth for every ONU that

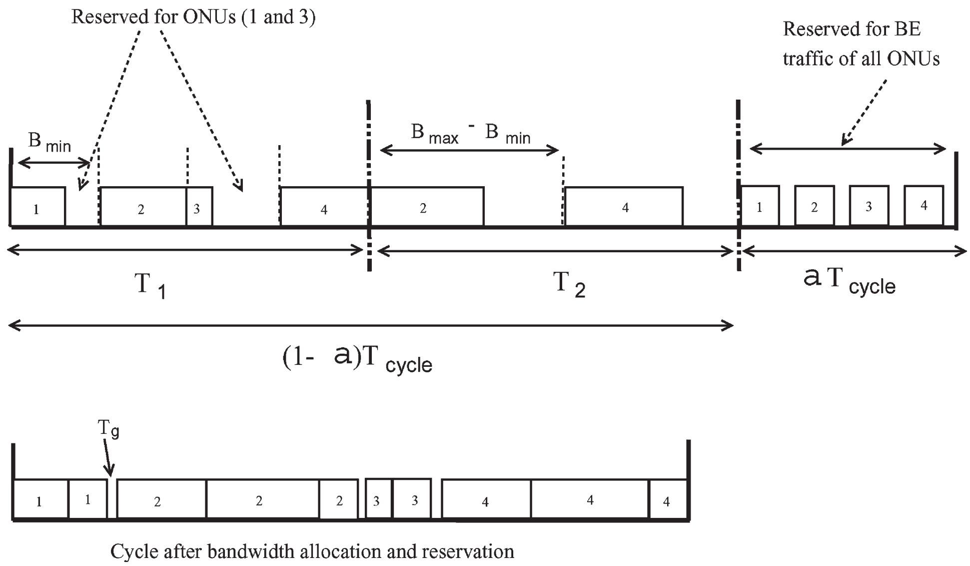

Fig. 1. Proposed cycle framework.

is either equal to its bandwidth request from the previous cycle, equal to the minimum guaranteed bandwidth ($B_{\min}$), or equal to the minimum guaranteed bandwidth plus a surplus bandwidth that may remain unused in the cycle. Clearly, the bandwidth of one ONU cannot be guaranteed and may vary from one cycle to another according to the load at other ONUs.

Bandwidth reservation resolves the uncertainty in allocating enough bandwidth that results from the load variations at different ONUs. Hence, each ONU is required to reserve bandwidth for its real-time 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 is guaranteed a new minimum bandwidth ($B_{\min}$) and could be allocated up to a maximum bandwidth ($B_{\max}$) in order to allow other ONUs to receive their share of the channel. BE traffic shares a fraction of the total cycle T_{cycle} ($T_{\text{cycle}} \leq 2$ ms in EPONs [2]), e.g., $\alpha \times T_{\text{cycle}}$, where $\alpha < 1$. When $\alpha = 0$, all the bandwidth of the upstream channel is used to transmit bandwidth-guaranteed traffic.

The new cycle ($(1-\alpha) \times T_{\text{cycle}}$) is used to provide services for bandwidth-guaranteed traffic. This new cycle, in turn, is divided into two subcycles T_1 and T_2; the OLT computes the minimum guaranteed bandwidth ($B_{\min}$) for each ONU using T_1, i.e., $B_{\min} = (T_{\text{cycle}} - N \times T_g) \times \xi / 8 \times N$, where ξ is the transmission speed of the passive optical network (PON: in megabits per second), N is the number of ONUs, and T_g is the guard time that separates the TW for every ONU_n and ONU_{n+1}. The ONU has total control over this bandwidth, whereas the bandwidth of the second subcycle is under the control of the OLT (please refer to Fig. 1 for a graphical elaboration, with $N = 4$). This new system enables us to implement a two-step AC: the first is a local AC (LAC) at the ONU, and the second is a global AC (GAC) at the OLT (as will be explained later in this paper). Note that, although the minimum guaranteed bandwidth is under the control of the ONU, the scheduling of various ONUs is still centrally done at the OLT in order to achieve a collision-free access to the upstream channel. The two subcycles are selected to be of equal length; however, if $T_1 < T_2$, then the OLT will have more control over the bandwidth with less guaranteed bandwidth per ONU. Conversely, the ONU is guaranteed more bandwidth, which may be unutilized if the load at a particular ONU is not high. Under our assumption of equal lengths for the subcycles, we set the maximum bandwidth that a highly loaded ONU can be allocated, i.e., $B_{\max} = \delta \times B_{\min}$. For example, when $\delta = 3$, a highly loaded ONU may or could be assigned a maximum of $2 \times B_{\min}$ from the second half cycle and, hence, a total of $3 \times B_{\min}$ per cycle. 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 the delay requirement for a certain flow/application cannot be respected, the flow should not be admitted. Such a stream if admitted will experience a degraded level of service and will also degrade the QoS of existing streams. Alternatively, BE traffic is never rejected and is always guaranteed a minimal bandwidth ($B_{\text{BE}}^{\min}$). Hence, to achieve these goals, the following two rules should not be violated before and after admitting a new real-time flow.

1) The QoS of each real-time stream (existing or new) should be guaranteed.
2) The BE traffic throughput $\text{BE}_{\text{throughput}} \geq B_{\text{BE}}^{\min}$.

In every cycle, the ONU reports 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 has already been predetermined and reserved, and it is guaranteed per cycle for the rest of the lifetime of each stream.

B. Traffic Characteristics and QOS Requirements

The admission decision for a new flow should be made according to both admission policies and QoS requirements often supplied by the application layer at the end users. The set of parameters that characterize the traffic stream varies from one traffic class to another. For example, constant bit rate (CBR) traffic is nonbursty and characterized by its mean data rate μ, which makes it quite predictable. With respect to QoS, CBR traffic requires stringent packet delays and delay variations (jitter). Alternatively, variable bit rate (VBR) traffic is quite bursty and may be characterized by the following parameters [9]: mean data rate μ (in bits per second), peak arrival data rate σ (in bits per second), and maximum burst size ρ (in bits). The delay bound θ, which is the maximum amount of time in units of microseconds allowed to transport a traffic stream (flow) measured between the arrival of the flow to the MAC layer and the start of transmission in the network. BE traffic, on the other hand, 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, the problem left for the AC unit (ACU, which is either at the ONU or OLT) is simply to determine whether a new stream i should be admitted and whether its QoS requirements can be guaranteed while the QoS requirements for the already admitted streams can be protected. For CBR traffic, the admission decision is straightforward: if the mean data rate μ_i ($g_i = \mu_i$ is the guaranteed rate) can be supported, then 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 a cycle. For 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 [9]. The former approach ends up admitting few streams, and the latter approach barely supports QoS for bursty streams. Therefore, a guaranteed bandwidth based on the traffic parameters could be derived, and we use a dual-token bucket for traffic regulation; this dual-token bucket is situated at the entrance of the MAC buffer and is associated with each stream. The guarantee rate for a VBR flow i can be easily derived [9] as follows:

$$g_i = \frac{\rho_i}{\theta_i + \frac{\rho_i}{\sigma_i}}. \tag{7}$$

Consequently, a conventional rate-based AC [11] can be used to determine whether a new stream can be admitted or not. For example, if S_j^{TW} is the bandwidth (in bits per second) that is allocated and reserved for ONU j, then a new flow $i+1$ could be admitted if

$$g_{i+1}^j + \sum_{i=1}^{h_j} g_i^j \le S_j^{\mathrm{TW}} \tag{8}$$

where h_j is the number of real-time streams (CBR or VBR) at ONU j. Note, however, that in EPON, the bandwidth that is assigned per ONU is not guaranteed, as mentioned earlier. Hence, we next propose a two-step AC scheme that will provide guaranteed bandwidth for each stream.

C. LAC

Each ONU is guaranteed a minimum bandwidth per cycle $B_{\min}$, and accordingly, the ONU can locally perform rate-based AC that is based on the bandwidth requirement of the new arriving flow and the bandwidth availability. For example, if g_f^j is the guarantee rate for the new flow f arriving at ONU j, then the bandwidth requirement (in bytes) per cycle for the new flow is: $R_f^j = g_f^j \times T_{\mathrm{cycle}}$. Therefore, this new flow will be admitted if

$$R_f^j + \sum_{i=1}^{h_j} R_{f_i}^j \le B_{\min} \tag{9}$$

where h_j is the total number of flows already admitted by the ONU, $R_{f_i}^j$ is the bandwidth requirement for a flow f_i, $R_{f_i}^j = g_{f_i}^j \times T_{\mathrm{cycle}}$, and $g_{f_i}^j$ is the guarantee rate (in bits per second). The LAC classifies the arriving flow into BE traffic or real-time traffic. BE traffic is always admitted. For real-time traffic, the ONU will derive the guarantee rate and check (9). If (9) holds, then the ONU will conditionally admit the flow and monitor its QoS for a predefined number of cycles (e.g., for 20 ms). If the QoS requirements of the newly admitted flow are satisfied and the QoS of existing flows remain intact, then the flow is admitted. Otherwise, the flow is dropped.

D. GAC

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.[1] The OLT may admit this new flow only if there is bandwidth that is available in the second subcycle T_2 and if the ONU that is 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 to this ONU, i.e., $B_{\mathrm{alloc}}^j = \sum_{i=1}^{h_j} R_i^j$, where R_i^j denotes the guaranteed bandwidth for already admitted h_j flows for ONU j. The OLT also maintains another variable that indicates the bandwidth that is still available (i.e., not committed yet), namely, B_{avail}, in T_2. The new flow may be admitted if the following two conditions hold simultaneously:

$$R_f^j + \sum_{i=1}^{h_j} R_i^j \le B_{\max} \tag{10a}$$

$$R_f^j \le B_{\mathrm{avail}}. \tag{10b}$$

Upon admitting a new flow, the OLT will reserve additional bandwidth for ONU j and update accordingly the total available bandwidth: $B_{\mathrm{avail}} = B_{\mathrm{avail}} - R_f^j$. Similarly, the OLT performs the algorithm above for every admission request of a new flow at any ONU. A flow will be rejected if at least one of the two conditions above is not satisfied. If both conditions are satisfied, then the OLT will conditionally admit the new flow and monitor its QoS parameters for the subsequent n cycles in order to determine whether it finally should admit the flow. When a flow leaves the network, the ONU reports to the OLT, and the latter will update the available bandwidth accordingly: $B_{\mathrm{avail}} = B_{\mathrm{avail}} + R_f^j$.

[1]This reporting is assumed to be done over the MPCP protocol. However, if another control plane is assumed for EPONs, then this information could be carried by the signaling protocol, e.g., the reservation protocol, of this plane.

E. Issues and Solutions

In the proposed AC scheme, every real-time stream is provided a guaranteed bandwidth that is computed based on the guarantee rate of the flow and is reserved and fixed per cycle. The OLT then allocates a TW that encompasses all the guaranteed bandwidth for every ONU per cycle. A subtle issue may arise, however; i.e., guaranteeing bandwidth per flow per cycle could ultimately waste the bandwidth. This issue is attributed to the burstiness of real-time traffic. In other words, if one ONU is being reserved bandwidth for a particular flow per cycle and has no traffic from this flow to transmit, then this bandwidth is not utilized and wasted. This issue arises because the allocation is done now statically (i.e., reservation) and not dynamic as in traditional EPON systems, where the bandwidth is allocated on demand. Moreover, if a flow had more bytes to be sent than the reserved ones (i.e., guaranteed), then the purpose of providing *guaranteed bandwidth* in every cycle will be defeated. This is because estimating the bandwidth requirement for a flow based on its guarantee rate does not accurately reflect the real nature of the traffic, especially with respect to the arrival of its packets in a short period (i.e., the short length of the cycle) and, hence, the inefficiency of the bandwidth prediction and reservation.

To resolve the above problem, we propose a two-branch solution. First, the OLT selects a supercycle ($T_{\rm sc} = \lambda \times T_{\rm cycle}$, where λ is a constant), and every admitted real-time flow is now guaranteed a bandwidth per $T_{\rm sc}$. The purpose of this proposal is to mitigate the inefficiency of the bandwidth reservation caused by the short-time prediction, and thus, a more accurate bandwidth estimation will take place. Here, as before, the period $(1-\alpha) \times T_{\rm sc}$ is divided into two periods, namely, T_1 and T_2. Each ONU is now guaranteed a bandwidth of $B_{\rm min}^{\rm new}$, which is computed based on T_1. The OLT controls the remaining bandwidth of the supercycle. Upon the arrival of a new flow f at ONU j with guaranteed bandwidth B_g^f, the flow is either admitted/rejected locally at the ONU or globally by the OLT, as described earlier in this paper. Second, we ensure that the reservation does not waste any bandwidth. Here, we apply a *crediting system* where each flow's estimated bandwidth is saved as credits at the OLT. In other words, every time a flow is admitted, the OLT will be informed, and it will compute/estimate a total credit (number of bytes available per $T_{\rm sc}$ for this flow) $C_{f_i}^j = B_g^{f_i} \times T_{\rm sc}^r$, where $T_{\rm sc}^r$ is the period between the arrival of the flow and the end of the current supercycle. The OLT also maintains a total credit per type of traffic (i.e., $C_{\rm CBR}^j$ for CBR and $C_{\rm VBR}^j$ for VBR) per ONU; e.g., $C_{\rm CBR}^j = \sum_{i=1}^{N_j} C_{f_i}^j$, where N_j is the number of CBR flows at ONU j. Now, in every cycle, the OLT deducts the requested/allocated bandwidth of this flow from its reserved credits until the time of a new supercycle. At this point, the credits are reset to the estimated ones. Next, we will explain how this solution will help in designing a DBA with effective reservation scheme.

F. AC-Enabled DBA Scheme (AC-DBA)

To apply the solutions proposed in the pervious section, we propose a new hybrid DBA that will perform both bandwidth allocation and reservation at the same time. The ONU reports to the OLT, in every cycle, its buffer occupancy $Q_{\rm CBR}(n-1)$, $Q_{\rm VBR}(n-1)$, and $Q_{\rm BE}(n-1)$, where n is the cycle number) and requests transmission bandwidth accordingly. However, here, the OLT will allocate bandwidth to each CoS at each ONU according to its remaining available credit in the current supercycle, as well as based on the requests received from other ONUs. Let $A_{\rm CBR}^j(n)$, $A_{\rm VBR}^j(n)$, and $A_{\rm BE}^j(n)$ be the bandwidth allocated for ONU j; then, we have

$$\sum_{j=1}^{N}\left(A_{\rm CBR}^j(n) + A_{\rm VBR}^j(n)\right) \le B_{\rm cycle} - T_{\rm gt}^t - \left(N \times B_{\rm BE}^{\min}\right) \tag{11}$$

$$\sum_{j=1}^{N} A_{\rm BE}^j(n) \le N \times B_{\rm BE}^{\min} \tag{12}$$

where $B_{\rm cycle}$ is the total bandwidth available in $T_{\rm cycle}$, $T_{\rm gt}^t$ is the total guard time (in bytes) between ONU transmissions, and $B_{\rm BE}^{\min}$ is the minimum guaranteed bandwidth (in bytes) for BE traffic, which is computed as follows:

$$B_{\rm BE}^{\min} = \frac{T_{\rm cycle} \times \frac{\alpha \times T_{\rm sc}}{N}}{8 \times T_{\rm sc}} \times \xi = \frac{T_{\rm cycle} \times \alpha}{8 \times N} \times \xi \tag{13}$$

where ξ is the PON speed (1 Gb/s). Every time the OLT allocates bandwidth to one ONU, it will adjust the available credit for every CoS accordingly: $C_{\rm CBR}^j(n) = C_{\rm CBR}^j(n-1) - A_{\rm CBR}^j(n)$. The credit for VBR traffic is similarly updated. If the ONU has run out of credits, then the OLT does not allocate any bandwidth for this CoS at this ONU during this supercycle. As for the computation of the available bandwidth for each CoS, the OLT waits until all requests [i.e., $R(Q_{\rm CBR}^j(n-1) + Q_{\rm VBR}^j(n-1) + Q_{\rm BE}^j(n-1)]$ are received from all ONUs. If $\sum_{j=1}^{N}(Q_{\rm CBR}^j(n-1) + Q_{\rm VBR}^j(n-1)) \le B_{\rm cycle} - T_{\rm gt}^t - N \times B_{\rm BE}^{\min}$, then $A_{\rm CBR}^j(n) = \min(Q_{\rm CBR}^j(n-1), C_{\rm CBR}^j(n-1))$; similarly, for VBR traffic, $A_{\rm VBR}^j(n) = \min(Q_{\rm VBR}^j(n-1), C_{\rm VBR}^j(n-1))$ and their credits for both CBR and VBR are updated accordingly. Otherwise, the OLT will compute the total guaranteed bandwidth B_j for each ONU j as follows:

$$B_j(n-1) = \frac{R_j(n-1) \times \left(B_{\rm cycle} - T_{\rm gt}^t - \left(N \times B_{\rm BE}^{\min}\right)\right)}{\sum_{j=1}^{N} Rj(n-1)} \tag{14}$$

where $R_j(n-1) = Q_{\rm CBR}^j(n-1) + Q_{\rm VBR}^j(n-1)$. Then, the OLT allocates bandwidth as follows:

$$A_{\rm CBR}^j(n) = \min\left(Q_{\rm CBR}^j(n-1), C_{\rm CBR}^j(n-1)\right) \tag{15a}$$

$$A_{\rm VBR}^j(n) = \min\Big(B_j(n-1) - Q_{\rm CBR}^j(n-1), C_{\rm VBR}^j(n-1)\Big). \tag{15b}$$

Next, the OLT will allocate bandwidth to BE traffic based on the requests received from the ONUs. The total BE bandwidth per cycle is $B_{\rm BE} = N \times B_{\rm BE}^{\min}$, which is shared by all ONUs.

If $\sum_{j=1}^{N}(A^j_{\mathrm{CBR}}(n) + A^j_{\mathrm{VBR}}(n)) \leq B_{\mathrm{cycle}} - T^t_{\mathrm{gt}} - N \times B^{\min}_{\mathrm{BE}}$, then the total bandwidth available for BE becomes

$$B_{\mathrm{BE}} = N \times B^{\min}_{\mathrm{BE}} + \left(B_{\mathrm{cycle}} - T^t_{\mathrm{gt}} - \sum_{j=1}^{N}\left(A^j_{\mathrm{CBR}}(n) + A^j_{\mathrm{VBR}}(n)\right)\right). \quad (16)$$

If $Q^j_{\mathrm{BE}}(n-1) \leq B^{\min}_{\mathrm{BE}}$, then $A^j_{\mathrm{BE}}(n) = Q^j_{\mathrm{BE}}(n-1)$. Otherwise, the OLT allocates to the ONU requesting less than $B^{\min}_{\mathrm{BE}}$ and computes the excess bandwidth to distribute them to other ONUs requesting more BE traffic. Accordingly, if $Q^j_{\mathrm{BE}}(n-1) > B^{\min}_{\mathrm{BE}}$, then $A^j_{\mathrm{BE}}(n) = B^{\min}_{\mathrm{BE}} + \chi_j$, where χ_j is the excess bandwidth allocated for ONU j and is expressed as

$$\chi_j = \frac{\alpha_j \times B^{\mathrm{rem}}_{\mathrm{BE}}(n)}{\alpha_t} \quad (17)$$

where $\alpha_j = Q^j_{\mathrm{BE}}(n-1) - B^{\min}_{\mathrm{BE}}$, $\alpha_t = \sum_{j=1}^{N}\alpha_j$, and $B^{\mathrm{rem}}_{\mathrm{BE}}(n)$ is the remaining bandwidth in the cycle n after allocating all ONUs bandwidth for their BE traffic such that

$$B^{\mathrm{rem}}_{\mathrm{BE}} = B_{\mathrm{BE}} - (N-L) \times B^{\min}_{\mathrm{BE}} - \sum_{j=1}^{L} Q^j_{\mathrm{BE}}(n-1) \quad (18)$$

where L is the number of ONUs requesting bandwidth for BE less than the minimum guaranteed bandwidth. In order to prevent the waste of bandwidth and control the allocation of surplus to various ONUs, the excess bandwidth allocated for the BE traffic at a highly loaded ONU χ_j is computed as follows:

$$\chi_j = \min(\chi_j, \alpha_j). \quad (19)$$

V. Performance Evaluation

We will study the performance of both the proposed intra-ONU scheduler (M-DWRR) and AC schemes for their QoS support and protection via simulations. The performance is measured with respect to maintaining satisfiable QoS requirements for real-time streams while guaranteeing a minimum required service for BE traffic. The total number of ONUs $N = 16$, and the PON speed is 1 Gb/s. The guard time is 1 μs, the cycle time $T_{\mathrm{cycle}} = 2$ ms, and the ONU buffering queue size is 10 MB.

We consider a more realistic traffic profile where real-time bandwidth-guaranteed streams (voice traffic is modeled by a CBR source, and video traffic is modeled using a VBR source) and BE traffic arrive dynamically at the ONUs. 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 B. Alternatively, a Poisson distribution can approximately model CBR traffic, and the packet size is fixed to 70 B. This scenario is appropriate to study the performance of the AC scheme presented. Each CBR flow is generated at a mean rate of 64 kb/s with a delay bound $\theta_{\mathrm{CBR}} = 2 \sim 4$ ms as QoS requirement [12], each VBR flow at a *guarantee rate*

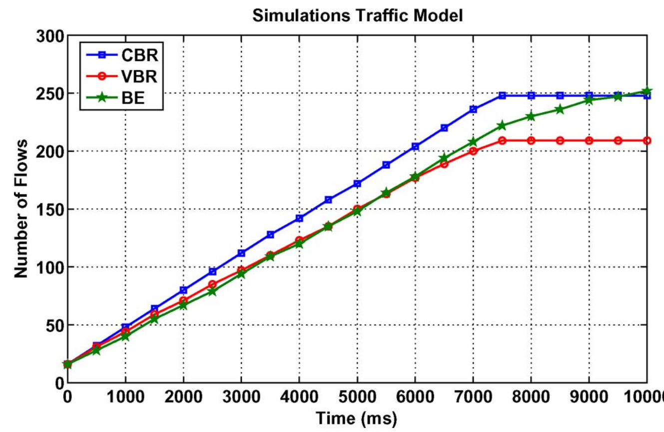

Fig. 2. Traffic model used for the AC framework.

[based on (7)] of 4 Mb/s with a delay bound $\theta_{\mathrm{VBR}} = 25 \sim 30$ ms as QoS requirement [12], and each BE flow at a mean rate of 5 Mb/s. Here, the load increases incrementally as more flows are admitted in the network. Fig. 2 depicts the chronological arrival of flows sorted by CoS and injected in the network. As shown, we stop generating real-time flows (i.e., CBR and VBR) at time 7500 ms, whereas BE flows continue to arrive until the end of the simulation. As for AC rule (2), we choose $B^{\min}_{\mathrm{BE}} = 4100$ B (in each cycle), which means that each ONU is guaranteed a BE throughput of 15 Mb/s (i.e., $\mathrm{BE}_{\mathrm{throughput}} > \simeq 20$ Mb/s, if available) if $T_{\mathrm{cycle}} \leq 2$ ms. Consequently, 16 ONUs will equally share a maximum of 20%–24% of PON's available bandwidth. For selecting the value of T_{sc}, we note that a large value (i.e., $\lambda \gg 1$, where $T_{\mathrm{sc}} = \lambda \times T_c$) should be selected especially when real-time streams are highly bursty (e.g., VBR traffic). Recall that the bandwidth requirement for a real-time flow is estimated according to the derived guarantee rate of the stream. The larger the period during which we estimate the bandwidth requirement is, the more accurate the bandwidth estimate becomes; this is mainly due to the bursty nature of the streams. For this reason, we select $T_{\mathrm{sc}} = 500$ ms for our simulations. The metrics of comparisons are average packet delays for CBR and VBR streams, throughput for CBR, VBR and BE flows, and the flow rejection rate. Note that, in this paper, we do not consider the flow QoS monitoring described earlier. We evaluated the performance of the proposed scheduling algorithm, i.e., M-DWRR, and we compared its performance with other scheduling disciplines such as DWRR and M-SFQ. Results (average packet delay and packet loss) have shown that M-DWRR outperforms the other two intra-ONU scheduling disciplines [14]. Accordingly, we only use M-DWRR for quantitative comparison with the proposed AC scheme. In addition, we also use an SP scheduler as a base scheme for our comparisons.

We begin by testing the behavior of our AC by showing in Fig. 3 the number of admitted real-time traffic streams. As shown, our system reaches saturation (i.e., no more real-time flows can be admitted in the network) at 7000 ms. As we continue generating real-time flows until 7500 ms, all these flows arriving afterward are rejected. However, this does not imply that no flows were rejected earlier since condition (9) or

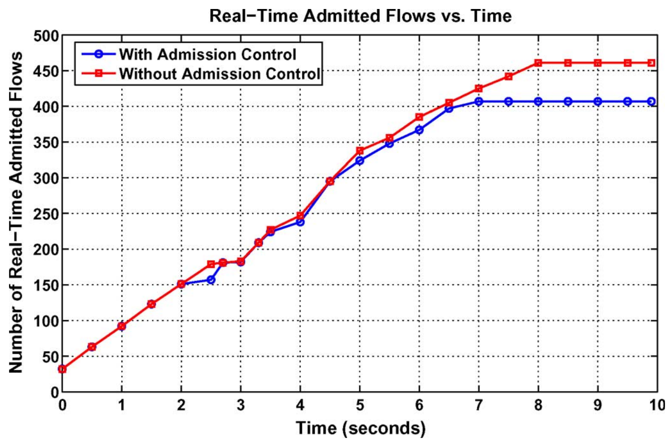

Fig. 3. Admission control behavior.

(10a) and (10b) need to be respected to admit a new arriving real-time flow; otherwise, a flow is rejected. The figure shows that starting at 450 ms, the network starts rejecting some newly arriving flows.

Next, we study the performance of real-time traffic by measuring the instantaneous average packet delays. To reduce the measurement complexity, we choose the sampling period $T = T_{\rm sc} = 500$ ms. Figs. 4 and 5 show these measurements, with AC (i.e., AC-DBA) and without AC [using M-DWRR ($\alpha_1 = 20\%$, $\alpha_2 = 70\%$, and $\alpha_3 = 10\%$) and SP schedulers]. Note that with AC-DBA, there is no intra-ONU scheduling required since the OLT allocates bandwidth for each CoS, per each ONU, every cycle. 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. This shows the advantage of M-DWRR; that is, although it divides the cycle among the CoS queues based on their assigned weights, it also provides an optimal performance for CBR traffic. This is due to the fact that the assigned weights are adaptively set based on the QoS requirements. On the other hand, under the SP scheduler that always selects packets from a queue with a higher priority until satisfied (i.e., until it is empty), CBR traffic will exhibit the best performance. As for AC-DBA, it makes sure to satisfy the QoS requirements defined previously (in terms of delay 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. As for VBR traffic (Fig. 5), AC-DBA maintains its delay performance to meet the specified target QoS requirements of the stream (i.e., 25–30 ms) while the delay witnesses an exponential increase under both adopted schedulers [Fig. 5(b) and 5(c)], i.e., a system that does not deploy any AC. This behavior highlights the need for the application of AC in EPON, because when the system reaches saturation (as described earlier in this paper) and all the arriving streams are admitted, the performance is no longer maintained. More specifically, no bandwidth is guaranteed for all types of traffic, and the QoS requirements are no longer met (not only for new application but for existing applications as well). On the other hand, the deployment of AC in EPON allows for a bandwidth-guaranteed service with guaranteed protected QoS.

We further investigate our AC framework by measuring/monitoring the throughput of one flow from each CoS (i.e., CBR and VBR) with AC (i.e., AC-DBA) and with no AC (i.e., M-DWRR and SP) in Fig. 6. As shown and expected, the selected CBR flow exhibits the same performance with and without AC, whereas the selected VBR flow shows a different behavior. Here, the VBR flow with AC maintains its derived 4-Mb/s throughput throughout the simulation, even after the system reaches saturation. On the other hand, when no AC is applied, the VBR flow does not show a stable throughput behavior. Moreover, when the system reaches saturation, the throughput of the VBR flow starts decreasing. This is due to the fact that when more real-time flows are admitted and no AC is applied, the bandwidth that was guaranteed for the already admitted flows (before saturation) is now shared by more flows. Hence, the bandwidth is no longer guaranteed for the already admitted flows and for the newly admitted ones. This, again, shows the need for AC in EPON to stabilize and guarantee the throughput for all admitted flows and reject the flows that will break this theme. This, in real and practical settings, will deny all malicious users from monopolizing the bandwidth provided; and at the same time, it will allow for *protection* to the bandwidth assigned for other well-behaved users. As for BE traffic, our concern is to guarantee a minimum total throughput that meets rule (2) in the AC scheme. For that reason, we measure its total throughput rather than the per-flow throughput as we did for CBR and VBR traffic. Here, the BE throughput increases to reach a total of ≈400 Mb/s under all schemes (i.e., with AC and with no AC) when the load is low and decreases when more flows are admitted. However, when the system reaches saturation, AC-DBA makes sure to preserve the minimum predefined throughput; whereas with M-DWRR and SP schedulers, the throughput is not guaranteed, and hence, the predefined throughput is no longer respected. Nevertheless, M-DWRR still provides a minimum throughput (which is one of its advantages) by forcing the weight policy while it reaches a very low one (≈0 Mb/s) with SP, a phenomenon known as BE traffic starvation.

We now study the impact of the size of the supercycle on the algorithm performance. A small supercycle will result in an inaccurate estimation of the bandwidth requirement for a stream due to its highly bursty nature. For example, if $T_{\rm sc} = 2$ ms (i.e., $\lambda = 1$), according to the derived guarantee rate, a flow will be guaranteed a certain bandwidth per $T_{\rm sc}$. If the flow has more traffic to transmit (since the peak rate is usually higher than the guarantee rate), then the OLT can only allocate the guaranteed bandwidth. If the flow has less traffic to send, the OLT allocates enough bandwidth for the requested traffic. Unlike longer supercycles, where for a given flow we assign total credits and the credits will be saved (carried on) for subsequent cycles if the flow does not have enough traffic to send in the current cycle, for the case where $T_{\rm sc} = 2$ ms, if a flow requests less than the guaranteed, then the flow will be allocated only the requested, and this results in wastage of the bandwidth of the supercycle for real-time streams. The problem even persists for small values of λ (this

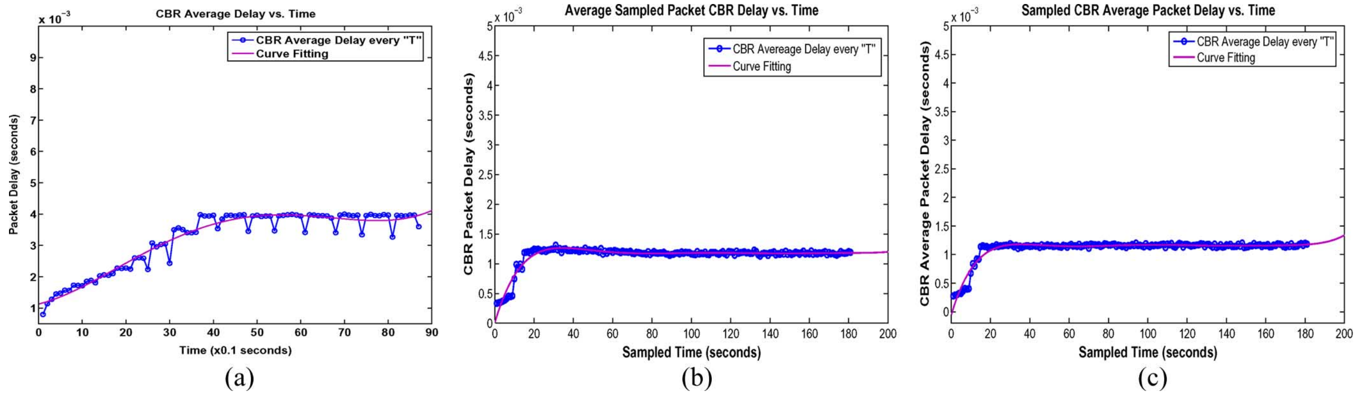

Fig. 4. CBR packet delay. (a) AC-DBA. (b) No-AC using M-DWRR. (c) No-AC using SP.

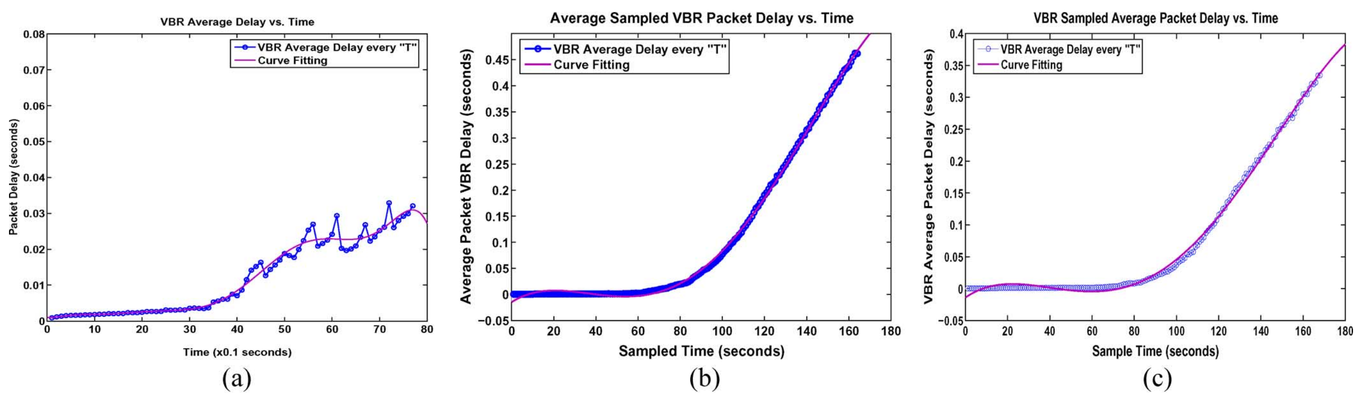

Fig. 5. VBR packet delay. (a) AC-DBA. (b) No-AC using M-DWRR. (c) No-AC using SP.

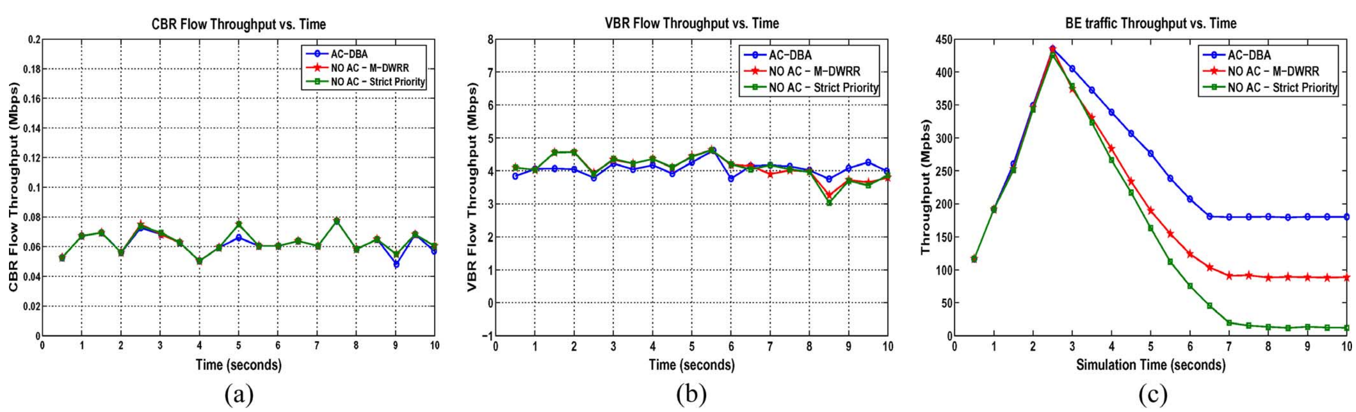

Fig. 6. CoS traffic throughput. (a) CBR flow throughput. (b) VBR flow throughput. (c) BE total throughput.

is because the period of the bursts is large in comparison to the length of the cycle, i.e., 2 ms). We run simulations for $T_{\mathrm{sc}} = 2$ ms, and the results show that there is almost 1 Mb/s degradation in the per-flow throughout for VBR streams. A lower degradation (around 5–10 kb/s) is seen for CBR traffic. We have, however, witnessed some increase in the total BE throughput (total increase of 50 Mb/s at higher loads). That is because when VBR requests less than the guaranteed bandwidth, there is some remaining bandwidth in the supercycle that can be used by the OLT to schedule the transmission of BE traffic; accordingly, a BE throughput increase is observed. Finally, Table I shows some interesting statistics collected from our simulations. These results show that $\approx$92% of the generated CBR traffic is admitted while their overall QoS and bandwidth requirements are guaranteed; $\approx$83% of VBR flows are admitted as well; and finally, all BE traffic arriving is admitted. Note that under M-DWRR and SP, all traffic is admitted; however, their QoS requirements are not guaranteed (except for CBR traffic). Note also that these collected results are traffic model dependent. In other words, more flows can be admitted or rejected,

TABLE I
TRAFFIC CONTROL STATISTICS

Number of Generated CBR Flows	252
Number of Admitted CBR Flows	234
Number of Rejected CBR Flows	18
CBR Admission Rate	$\approx 92\%$
Number of Generated VBR Flows	209
Number of Admitted VBR Flows	173
Number of Rejected VBR Flows	36
VBR Admission Rate	$\approx 83\%$
Number of Generated BE Flows	247
Number of Admitted BE Flows	247
Number of Rejected BE Flows	0
BE Admission Rate	100%

depending on all of the required guaranteed throughput for real-time and BE traffic, the generated flows' mean rates, and the number of flows generated.

VI. CONCLUSION

Providing bandwidth-guaranteed service in EPON is a challenging subject that has not been addressed in the literature. In this paper, we presented the first complete EPON framework that supports the application of AC in EPON. This framework implements a two-stage AC (i.e., at the ONU and at the OLT) with all its rules and functionalities, along with a new hybrid AC-enabled DBA that performs both bandwidth allocation and reservation simultaneously. We have also presented a new simulation model that is designed to test this framework. We showed that although some of the scheduling (intra-ONU) mechanisms can provide QoS for various types of traffic in the network, none of these schedulers could protect these QoS requirements. Our AC system has shown a good performance in terms of maintaining the QoS level for already existing traffic while providing an overall acceptable minimal throughput for BE traffic even under network saturation.

REFERENCES

[1] M. P. McGarry, M. Maier, and M. Reisslein, "Ethernet PONs: A survey of dynamic bandwidth allocation (DBA) algorithms," *IEEE Commun. Mag.*, vol. 42, no. 8, pp. S8–S15, Aug. 2004.
[2] C. M. Assi, Y. Ye, S. Dixit, and M. Ali, "Dynamic bandwidth allocation for quality-of-service over Ethernet PONs," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 9, pp. 1467–1477, Nov. 2003.
[3] M. Garcia, D. F. Garcia, V. G. Garcia, and R. Bonis, "Analysis and modeling of traffic on a hybrid fiber–coax network," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 9, pp. 1718–1730, Nov. 2004.
[4] *IEEE 802.3ah, Ethernet in the First Mile Task Force*. [Online]. Available: http://www.ieee802.org/3/efm/index.html
[5] G. Kramer *et al.*, "Fair queuing with service envelopes (FQSE): A cousin-fair hierarchical scheduler for subscriber access networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 8, pp. 1497–1513, Oct. 2004.
[6] C. Semeria, *Supporting Differentiated Service Classes: Queue Scheduling Disciplines*. Sunnyvale, CA: Juniper Netw., Jan. 2002. white paper.
[7] N. Ghani, A. Shami, C. Assi, and Y. Raja, "Intra-ONU bandwidth scheduling in Ethernet passive optical networks," *IEEE Commun. Lett.*, vol. 8, no. 11, pp. 683–685, Nov. 2004.
[8] M. Ma, Y. Zhu, and T. H. Cheng, "A bandwidth guaranteed polling MAC protocol for Ethernet passive optical networks," in *Proc. IEEE INFOCOM*, San Francisco, CA, Apr. 2003, pp. 22–31.
[9] 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," in *Proc. IEEE INFOCOM*, Apr. 2005, pp. 1584–1595.
[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, pp. 1575–1584, May 2003.
[11] S. Jamin, P. Danzig, S. Shenker, and L. Zhang, "A measurement-based admission control algorithm for integrated services packet networks," *IEEE/ACM Trans. Netw.*, vol. 5, no. 1, pp. 56–70, Feb. 1997.
[12] *End-user multimedia QoS categories*, 2001. ITU-T Recommendation G.1010.
[13] J. Zheng and H. T. Mouftah, "Media access control for Ethernet passive optical networks: An overview," *IEEE Commun. Mag.*, vol. 43, no. 2, pp. 145–150, Feb. 2005.
[14] A. R. Dhaini, C. Assi, A. Shami, and N. Ghani, "Adaptive fairness through intra-ONU scheduling in Ethernet passive optical networks," in *Proc. IEEE ICC*, Istanbul, Turkey, Jun. 2006, pp. 2687–2692.
[15] H. Naser and H. T. Mouftah, "A joint-ONU interval-based dynamic scheduling algorithm for Ethernet passive optical networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 4, pp. 889–899, Aug. 2006.

Ahmad R. Dhaini (S'04) received the B.S. degree in computer science from the American University of Beirut, Beirut, Lebanon, in 2004 and the M.S. degree in electrical and computer engineering, with a best thesis award nomination, from Concordia University, Montreal, QC, Canada, where he is currently working toward the Ph.D. degree.

His current research interests include wired/wireless access networks, quality of service, and scheduling algorithms. More specifically, he is currently working in the area of passive optical networks.

Chadi M. Assi (M'03) received the B.S. degree in engineering from the Lebanese University, Beirut, Lebanon, in 1997 and the Ph.D. degree from the City University of New York, in April 2003.

From September 2002 to August 2003, he was a Visiting Researcher at the Nokia Research Center, Boston, MA, working on quality of service in optical access networks. He is currently with the Faculty of Engineering and Computer Science, Concordia Institute for Information Systems Engineering, Concordia University, Montreal, QC, Canada, which he joined in August 2003 as an Assistant Professor. His current research interests include provisioning and restoration of optical networks, wireless and *ad hoc* networks, and security.

Dr. Assi received the Mina Rees Dissertation Award from the City University of New York in August 2002 for his research on wavelength-division-multiplexing optical networks.

Martin Maier (M'03) received the Dipl.-Ing. and Dr.-Ing. degrees (both with distinctions) in electrical engineering from the Technical University Berlin, Berlin, Germany, in 1998 and 2003, respectively.

He is currently an Associate Professor with the Institut National de la Recherche Scientifique (INRS)—Énergie, Matériaux et Télécommunications, Montréal, QC, Canada. He was a Visiting Researcher with the University of Southern California, Los Angeles, and Arizona State University, Tempe. In summer 2003, he was a Postdoctoral Fellow with the Massachusetts Institute of Technology, Cambridge. He is the author of the book *Metropolitan Area WDM Networks—An AWG Based Approach* (Springer, 2003). His research interests include network and node architectures, routing and switching paradigms, protection, restoration, multicasting, and the design, performance evaluation, and optimization of MAC protocols for optical wavelength-division-multiplexing (WDM) networks, automatically switched optical networks (ASONs), and Generalized Multi-Protocol Label Switching (GMPLS), with a particular focus on metro and access networks. Recently, his research interests have concentrated on evolutionary WDM upgrades of optical metro ring networks and access networks.

Dr. Maier was the recipient of the two-year Deutsche Telekom doctoral scholarship from June 1999 to May 2001. He is also a corecipient of the Best Paper Award presented at the International Society of Optical Engineers (SPIE) Photonics East 2000 Terabit Optical Networking Conference. He served on the Technical Program Committees of IEEE INFOCOM, the IEEE Global Communications Conference, and the IEEE International Performance, Computing, and Communications Conference and is an Editorial Board Member of the *IEEE Communications Surveys and Tutorials*.

Abdallah Shami (M'02) received the B.E. degree in electrical and computer engineering from the Lebanese University, Beirut, Lebanon, in 1997 and the Ph.D. degree in electrical engineering from the City University of New York, in September 2002.

In September 2002, he joined the Department of Electrical Engineering, Lakehead University, Thunder Bay, ON, Canada, as an Assistant Professor. Since July 2004, he has been with the University of Western Ontario, London, ON, where he is currently an Assistant Professor with the Department of Electrical and Computer Engineering. His current research interests include data/optical networking, EPON, WIMAX, WLANs, and software tools.

Dr. Shami is a recipient the Irving Hochberg Dissertation Fellowship Award from the City University of New York and a GTF Teaching Fellowship.