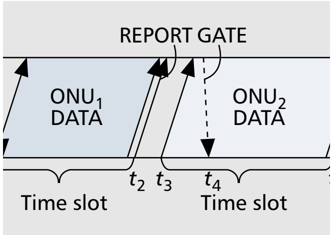

Bandwidth Allocation for Multiservice Access on EPONs

Yuanqiu Luo and Nirwan Ansari, New Jersey Institute of Technology

Abstract

Ethernet passive optical networks are a low-cost high-speed solution to the bottleneck problem of the broadband access network. A major characteristic of EPONs is the shared upstream channel among end users, mandating efficient medium access control to facilitate statistical multiplexing and provision multiple services for different types of traffic. This article addresses and provides an overview of the upstream bandwidth allocation issue for multiservice access provisioning over EPONs, and proposes an algorithm for dynamic bandwidth allocation with service differentiation. Based on the multipoint control protocol (MPCP) and bursty traffic prediction, our algorithm enhances QoS metrics such as average frame delay, average queue length, and frame loss probability over other existing protocols.

Introduction

Ethernet passive optical networks (EPONs) address the first mile of the communication infrastructure between service provider central offices and customer sites. With the expansion of services offered over the Internet, a dramatic increase of bandwidth has been facilitated in the backbone network through the use of wavelength-division multiplexing (WDM), providing tens of gigabits per second per wavelength. At the same time, LANs have been scaled up from 10 to 100 Mb/s and are being upgraded to Gigabit Ethernets. Such a growing gap between the capacity of the backbone network and end users' needs results in a serious bottleneck of the access network between them [1]. An access network technology is desired that can provide low-cost efficient equipment to facilitate multiservice access to end users. EPONs are considered an attractive and promising solution for the broadband subscriber access network. As an inexpensive, simple, and scalable technology, and capable of delivering integrated services, EPONs are being deliberated in the standardization process of the IEEE 802.3ah Ethernet in the First Mile (EFM) Task Force [2], which aims to significantly increase broadband service performance while minimizing equipment, operation, and maintenance costs.

Typically, an EPON consists of an optical line terminal (OLT) located at the provider central office and a set of associated optical network units (ONUs) that deliver broadband voice, data, and video services to end users. The optical distribution network (ODN) comprises fibers with a passive splitter lying between each OLT and its associated ONUs. As shown in Fig. 1, a single fiber extends from an OLT to a 1:*N* passive optical splitter. The splitter fans out to multiple single fiber drops, which are connected to different ONUs. EPONs eliminate the active electronic components such as regenerators and amplifiers in ODN, and replace them with less expensive passive optical splitters, which are simpler and easier to maintain than active components. With data encapsulated in IEEE 802.3 Ethernet frames, EPONs rely on the multipoint control protocol (MPCP) mechanism for the operations, administration, and maintenance (OAM) to control the point-to-multipoint (P2MP) connection between each OLT and its associated ONUs. Compared to the point-to-point network and curb-to-switched network, this architecture has the advantages of minimizing the number of optical transceivers and eliminating intermediate powering [3].

The process of transporting data downstream to end users in EPONs is different from that of transporting data upstream to the OLT. In downstream transmission, data are broadcast from the OLT to each ONU using the entire bandwidth of the downstream channel, and all the downstream data are carried in one wavelength. ONUs selectively receive frames destined to themselves by matching the addresses in the Ethernet frames. The broadcasting nature of Ethernet perfectly matches the EPON downstream transmission, and the "broadcast and select" architecture allows downstream multimedia services like video broadcasting. In the upstream direction, multiple ONUs share the common upstream channel, and another wavelength is employed for the upstream traffic. Only a single ONU may transmit during a time slot in order to avoid data collisions. Because of the directional nature of the passive optical splitter, each ONU transmits directly to the OLT, but not to other ONUs. An ONU buffers the frames from end users until its time slot arrives. The buffered frames would be "burst" out to the OLT in the exclusively assigned time slot at full channel speed. In order to provide diverse quality of service (QoS), bandwidth management of the upstream channel is essential for successful implementation of EPONs. In this article we focus on the issue of upstream bandwidth allocation and propose an algorithm to dynamically allocate bandwidth among end users with service differentiations.

The rest of the article is organized as follows. The next sec-

 0163-6804/05/$20.00

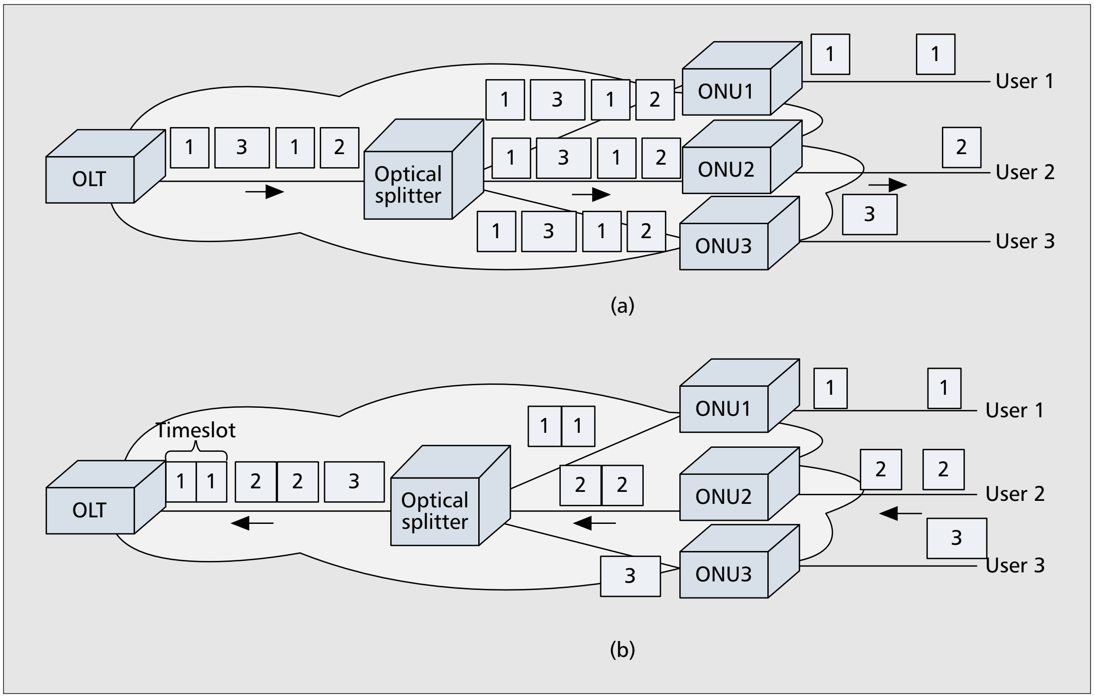

FIGURE 1. EPON transmission: a) downstream; b) upstream.

tion addresses the provisioning of multiple services among ONUs (i.e., schemes for upstream bandwidth allocation among ONUs). We then discuss the issue of providing multiple services among end users of a single ONU, which can be facilitated by a combination of queuing, scheduling, and class-based bandwidth allocation at the ONU. We then propose a dynamic bandwidth allocation algorithm that employs bursty traffic prediction, and its performance is demonstrated by simulations. Conclusions are drawn in the last section.

MULTIPLE SERVICES AMONG ONUS

The major challenges of EPONs include medium access control (MAC) protocol design and multiservice provisioning. Because of the directional property of the passive optical splitter, it is difficult for an ONU to detect frame collisions by the conventional carrier sense multiple access with collision detection (CSMA/CD) MAC protocol designed for Ethernet. Therefore, an efficient MAC protocol is crucial to ensure high bandwidth utilization. Since Ethernet does not support QoS directly, while the access network is required to accommodate various kinds of traffic, multiservice access is a distinguished feature EPONs are expected to provide. Due to differences in subscribers' service level agreements (SLAs), different ONUs may have different bandwidth requirements. A pragmatic approach is to employ time-slot-based bandwidth allocation by providing various lengths of time slots to different ONUs. Below we describe the MPCP mechanism for bandwidth allocation and a few approaches the OLT employs to allocate bandwidth among ONUs.

MULTIPOINT CONTROL PROTOCOL

MPCP [2] is being developed by the IEEE 802.3ah Task Force to specify the mechanism between an OLT and the associated ONUs to facilitate efficient upstream data transmission. MPCP is a frame-based protocol that introduces 64-byte MAC control messages to provide optimal transmission of Ethernet frames. REGISTER_REQUEST, REGISTER, and REGISTER_ACK messages are utilized in the auto-discovery process to harmonize a new ONU, register the ONU, and negotiate parameters.

Without specifying any particular bandwidth allocation algorithm, MPCP provides the REPORT/GATE mechanism to manage upstream bandwidth. The control messages for bandwidth assignment illustrated in Fig. 2 are REPORT and GATE, respectively. An ONU sends a REPORT message to the OLT, containing the timestamp and queue status of the ONU. The OLT calculates the round-trip time (RTT) from the reported timestamp. The bandwidth allocation algorithm uses the queue status to make the allocation decision. The OLT sends a GATE message downstream, containing the information of the timestamp, grant start time, and grant length. The destined ONU updates its local clock by the received timestamp, and transmits frames from the grant start time in the grant length. No packet fragmentation is allowed within a time slot, and the "unfit" Ethernet frame will be deferred to the next time slot.

FIXED BANDWIDTH ALLOCATION

Fixed bandwidth allocation (FBA) grants each ONU a fixed time slot length in every service cycle. A service cycle is defined as the time each ONU transmits its data once to the OLT. FBA works exactly like time-division multiple access (TDMA), in which the time slot of each ONU is fixed beforehand and is ignorant of the actual traffic arrival rate. Without the overhead of the queue status report and transmission grant, FBA is simple to implement. On the other hand, an ONU will occupy the upstream channel for its assigned time slot even if there is no frame to transmit, thus resulting in the increased delay for all the Ethernet frames buffered in other ONUs. Many frames could be backlogged in the buffers while the upstream channel is lightly loaded or even idle, hence leading to underutilization of the upstream channel.

LIMITED BANDWIDTH ALLOCATION

Limited bandwidth allocation (LBA) [4] monitors the incoming traffic by the REPORT/GATE mechanism. The time slot length of an ONU is upper bounded by the maximum time slot length B_{max}, which could be specified by SLA or other system parameters. When the reported queue size is less than the limit, the OLT grants the bandwidth request; otherwise, B_{max} is granted. LBA tracks the traffic load by means of the

REPORT message	GATE message
Destination address (6 octets)	Destination address (6 octets)
Source address (6 octets)	Source address (6 octets)
Type (2 octets)	Type (2 octets)
Opcode (2 octets)	Opcode (2 octets)
Timestamp (4 octets)	Timestamp (4 octets)
Report bitmap (1 octet)	Number of grants (1 octet)
Queue #0 (0/4 octet)	Grant level (1 octet)
Queue #1 (0/4 octet)	Grant start time (4 octets)
Queue #2 (0/4 octet)	Grant length (2 octets)
⋮	⋮
Pad	Pad
FCS (4 octets)	FCS (4 octets)

FIGURE 2. MPCP control messages for bandwidth management.

queue status reports, the granted time slot length varies according to the dynamic traffic, and the service cycle varies because ONUs are assigned with different time slot lengths in different service cycles. The conservative feature of LBA confines each ONU by its own limit, thus restricting aggressive competition for the upstream bandwidth.

CREDIT-BASED BANDWIDTH ALLOCATION

Under the REPORT/GATE mechanism, each ONU experiences a waiting time from sending the REPORT message to sending the buffered frames, as illustrated in Fig. 3. When sending a REPORT message at time t_1, an ONU only reports the already buffered frames to the OLT; therefore, frames that arrive during the waiting time (i.e., from t_3 to t_1) have to be deferred to the next time slot even if the upstream channel is lightly loaded. Credit-based bandwidth allocation (CBA) [4] takes such frames into consideration, and when the OLT allocates the upstream bandwidth, it adds a credit into each ONU's requirement. $\boldsymbol{B}_{grant} = \boldsymbol{B}_{queue} + \boldsymbol{C}$, where $\boldsymbol{B}_{grant}$ is the granted bandwidth to an ONU, $\boldsymbol{B}_{queue}$ is the frames (in terms of bandwidth) queued up in the buffer, and $\boldsymbol{C}$ is the credit. $\boldsymbol{C}$ could be a constant or linear credit. The incoming frames during the waiting time are expected to be transmitted (or partially transmitted) within the current time slot.

EXCESSIVE BANDWIDTH REALLOCATION

In LBA there might be some lightly loaded ONUs with bandwidth requirements less than the limits. The sum of the underexploited bandwidth of lightly loaded ONUs is called *excessive bandwidth*, $\boldsymbol{B}_{excess}$. As an extension of LBA, excessive bandwidth reallocation (EBR) [5] exploits $\boldsymbol{B}_{excess}$ by redistributing it among the heavily loaded ONUs. The heavily loaded ONU$_k$ obtains an additional bandwidth, $\boldsymbol{B}_{add,k}$, where $\boldsymbol{B}_{add,k} = (\boldsymbol{B}_{excess} * \boldsymbol{B}_{max,k})/(\Sigma_i \boldsymbol{B}_{max,i})$, and $\boldsymbol{B}_{max,i}$ is the bandwidth limit of ONU$_i$ specified in LBA.

MULTIPLE SERVICES AT AN ONU

Other than upstream channel bandwidth allocation among different ONUs, it is necessary for a single ONU to provide multiple services to its different end users. Residing at customer premises, an ONU must be capable of supporting data, voice, and video services to end users. This can be approached by means of a combination of queuing, scheduling, and class-based bandwidth allocation.

As shown in Fig. 2, one 64-byte GATE message carries up to six grants to a particular ONU. The *number of grants* field specifies how many grants are in the message, and the *grant level* field indicates the order of the queues to which grants are generated. Each grant contains a grant start time and grant length. One 64-byte REPORT message from an ONU reports up to eight queues' status. The *report bitmap* field identifies the order of the reported queues. The OLT processes the queue status report, and sends back the GATE message including at least one grant, depending on the bandwidth allocation algorithm. Therefore, reporting multiple queues of an ONU and granting multiple requirements to an ONU are possible, thus making multiservice provisioning to the end users of an ONU feasible.

PRIORITY QUEUING

Categorizing the traffic of an ONU into different classes is a practical approach to multiservice provisioning [6]. The high-priority class is expedited forwarding (EF), which is delay-sensitive and requires bandwidth guarantees. The medium-priority class is assured forwarding (AF), which is not delay-sensitive but requires bandwidth guarantees. The low-priority class is best effort (BE), which is neither delay-sensitive nor bandwidth guaranteed. Frames belonging to different classes are enqueued into their corresponding priority queues. All queues share the same buffer. When the buffer is full, incoming frames with higher priority replace lower-priority frames while incoming low-priority frames are dropped immediately. The multiple queue status at an ONU is reported to the OLT by means of the REPORT message. MPCP allows one ONU to transmit eight queue status reports at once.

PRIORITY SCHEDULING

Buffered frames are transmitted according to a specific scheduling scheme. As defined in IEEE 802.1D [7], *strict priority scheduling* serves buffered higher-priority frames first. BE frames can only be transmitted when the other two queues are empty. Adhering to the priority order, strict priority scheduling serves higher-priority frames that arrive during the waiting time ahead of lower-priority frames that may already be queued up in the buffer. In Fig. 3 the EF frames that arrive during the waiting time (i.e., from t_5 to t_7) will be served ahead of the AF and BE frames arriving earlier (i.e., before t_5). Therefore, the lower-priority frames suffer uncontrolled increasing delay (if the buffer is not full) or unfair drop (if the buffer is full).

Priority-based scheduling [5] tackles unfairness by employing strict priority scheduling within a specific time interval. After an ONU transmits all buffered frames in an interval, the frames arriving after this interval will be served if the current time slot can still transmit more frames. By configuring the interval as the time between sending the REPORT messages (i.e., from t_1 to t_5 in Fig. 3), higher-priority frames arriving in the waiting time (i.e., from t_5 to t_7) will be served after all classes of frames of the previous interval (i.e., from t_1 to t_5) have been served. This scheme provides a bounded delay for low-priority frames.

CLASS-BASED BANDWIDTH ALLOCATION

Reference [8] handles class-based bandwidth allocation by collecting REPORT messages from all ONUs before making decisions. This algorithm is referred to as D1. The OLT assigns fixed bandwidth to the EF traffic of all ONUs regardless of its dynamics. The AF requests are granted as follows: if the sum of the AF requests of all ONUs is less than or equal to the leftover bandwidth after serving the EF services, all AF requests are granted; otherwise, the leftover bandwidth is equally distributed among all AF requests. The leftover bandwidth after serving EF and AF traffic is distributed among all BE requests. The major drawbacks of D1 include the fixed bandwidth allocation for EF traffic, which penalizes AF and BE traffic by increasing the frame delay; and the long report collection time, which does not end until reports are received from all ONUs.

The algorithm proposed in [5] estimates the incoming EF traffic in the waiting time by the amount of such frames in the previous cycle, and is referred to as D2. The reported EF traffic is the sum of the buffered EF frames plus the estimation, while the reported AF and BE traffic is the actually buffered amount. Bandwidth requests are granted by EBR, with lightly loaded queues receiving instantaneous grants, and grants for heavily loaded queues being deferred until all reports have been received. D2 alleviates the delay of collecting all reports by granting lightly loaded queues immediately. This algorithm estimates incoming EF frames arriving during the waiting time, and gives priority to EF traffic by allocating the estimated bandwidth. The drawback is that the service order of ONUs changes in every service cycle, with heavily loaded ONUs always being served after lightly loaded ones; therefore, estimation of the incoming EF frames is severely impaired because the waiting time of each ONU may change drastically.

Dynamic Bandwidth Allocation with Multiple Services

In this section we present our proposed dynamic bandwidth allocation with multiple services (DBAM) algorithm to accommodate various types of traffic in EPONs. Instead of providing multiple services among ONUs and among end users separately, the approach of DBAM is to incorporate both of them into the REPORT/GATE mechanism with class-based bandwidth allocation.

To overcome the drawbacks of D1, DBAM applies *priority queuing* to enqueue the EF, AF, and BE frames, and gives preference to higher-priority traffic. Priority-based scheduling is exploited to schedule the buffered frames, and the schedule interval is the time between sending REPORT messages. DBAM adopts LBA to arbitrate bandwidth allocation among ONUs, thus prohibiting aggressive bandwidth scrambling. DBAM employs class-based traffic prediction to take the frames arriving during the waiting time into account , and such prediction is expected to reduce frame delay and queue length. The OLT serves all ONUs in a fixed round robin order to avail traffic prediction.

The maximum bandwidth parameter of a specific class of traffic is determined by the SLA between the end user and service provider. Let S_i^{EF}, S_i^{AF}, and S_i^{BE} be the maximum bandwidth parameters for the EF, AF, and BE traffic at ONU_i, respectively, and thus the maximum bandwidth parameter of ONU_i is S_i, where $S_i = S_i^{EF} + S_i^{AF} + S_i^{BE}$.

Figure 4 illustrates the upstream transmission in an EPON with two ONUs. The OLT serves the two ONUs alternately. As mentioned earlier, the interval of an ONU is the time between sending REPORT messages. In terms of ONU_1, interval n ranges from time t_1 to time t_6, (i.e., $T_{1,n} = t_6 - t_1$). ONU_1 bursts out the buffered frames in its time slot from time t_1 to time t_2, and piggybacks a REPORT message at the end of the time slot. Time t_2 to time t_4 is the RTT between ONU_1 and the OLT plus the REPORT processing time. Time t_2 to time t_6 is the waiting time for ONU_1 in interval n (i.e., $T_{1,n}^W = t_6 - t_2$), during which ONU_1 is idle, and more frames from end users are enqueued. Interval (n + 1) of ONU_1 begins at time t_6, and the granted time slot from time t_6 to time t_8 is based on the REPORT message sent at time t_2. In terms of ONU_2, interval n begins at time t_3 and ends at time t_9 (i.e., $T_{2,n} = t_9 - t_3$). Time t_3 to time t_5 is the exclusive time slot for ONU_2, and a report of its three-queue status is sent at time t_5. Time t_5 to time t_9 is the waiting time in interval n (i.e., $T_{2,n}^W = t_9 - t_5$).

After transmitting data in its time slot of interval n, ONU_i requests its bandwidth of interval (n + 1) by piggybacking a REPORT message. The bandwidth request for interval (n + 1) for each service class, c, is

$$R_{i,n+1}^c = (1+\alpha)B_{i,n}^c, c \in \{EF, AF, BE\}. \tag{1}$$

$R_{i,n+1}^c$ is the requested bandwidth for class c traffic at ONU_i for interval (n + 1). $B_{i,n}^c$ is the amount of class c traffic already queued up in the buffer when ONU_i is making a request during interval n. α is the *estimation credit*, which is the ratio of the waiting time of ONU_i in interval n vs. the length of interval n,

$$\alpha = T_{i,n}^w \big/ T_{i,n}, \tag{2}$$

where $T_{i,n}^W$ is the waiting time of ONU_i in interval n, and $T_{i,n}$ is the length of interval n. In other words, the frames arriving in the waiting time of interval (n + 1) are estimated from the information of interval n. The idea behind this is network traffic self-similarity, which implies that network traffic exhibits long-rang dependence [9]. $T_{i,n}^W$ and $T_{i,n}$ are calculated based on the timestamps in the REPORT and GATE messages.

The OLT instantaneously makes the bandwidth allocation decision after having received a REPORT. The granted bandwidth of all three classes of traffic at ONU_i for interval (n + 1) is

$$B_{i,n+1} = \min\{\textstyle\sum_c R_{i,n+1}^c, S_i\}, c \in \{EF, AF, BE\}. \tag{3}$$

The multiple services among the ONUs are incorporated with

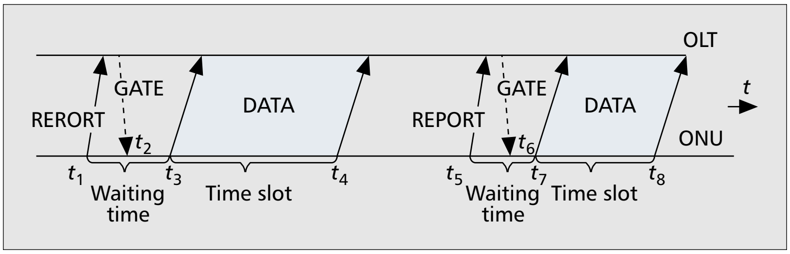

Figure 3. REPORT/GATE mechanism and waiting time.

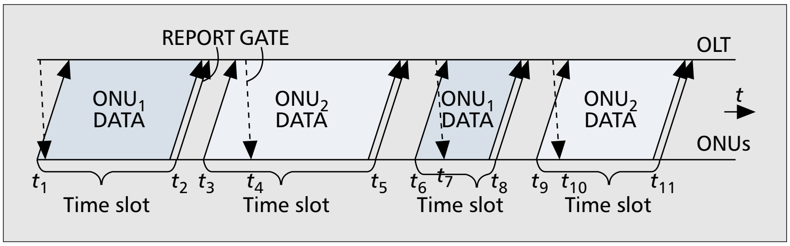

Figure 4. Upstream transmission in DBAM.

their bandwidth requirements and SLA limits. The bandwidth of ONU_i is upper bounded by the smaller value of its request and its maximum bandwidth parameter. The assigned bandwidth changes with incoming traffic. The OLT allocates bandwidth to the EF and AF traffic first, both of which require bandwidth guarantees. The amount of bandwidth granted to different classes is thus

$$B^{c}_{i,n+1} = \min\{R^{c}_{i,n+1}, S^{c}_{i}\}, c \in \{EF, AF\}, \tag{4}$$

$$B^{BE}_{i,n+1} = B_{i,n+1} - B^{BE}_{i,n+1} - B^{AF}_{i,n+1}. \tag{5}$$

Multiservice access at an ONU is based on the traffic classification. The ONU gives priority to EF and AF traffic. BE traf-

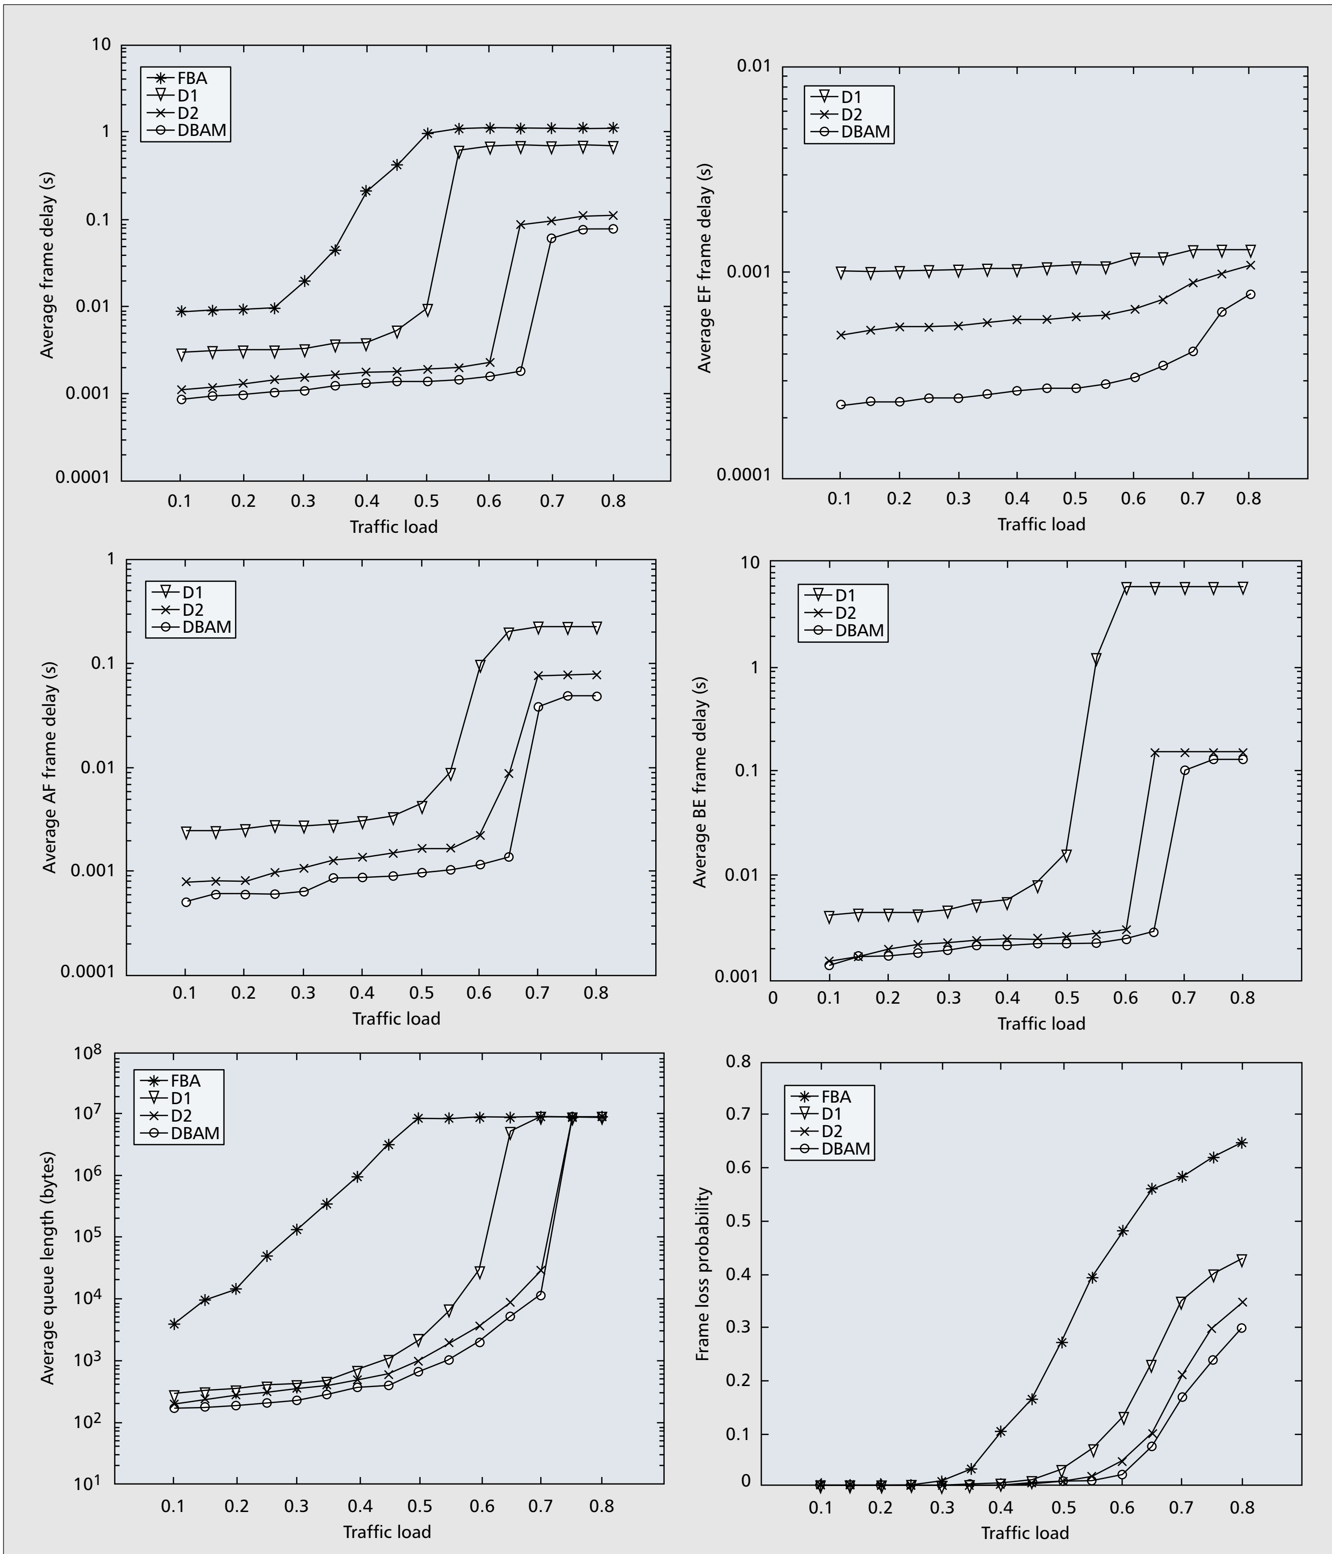

FIGURE 5. Simulation results: a) Average frame delay of all frames; b) average fame delay of EF frames; c) average frame delay of AF frames; d) average fame delay of BE frames; e) average queue length; and f) frame loss probability.

fic makes use of the unused bandwidth. The granted bandwidth of EF and AF traffic is upper bounded by their maximum bandwidth parameters, which restrict them from aggressively consuming the upstream channel. On receiving the GATE message, the ONU updates its local clock, and programs the local registers with the grant start time and grant length values. When its dedicated time slot comes, the ONU bursts out the scheduled frames to the OLT without contention from other ONUs.

PERFORMANCE EVALUATION

The system model is set up as in Fig. 1 with one OLT and 32 ONUs. The downstream and upstream channels are both 1 Gb/s. Each ONU has a finite buffer of 10 Mbytes, and the distance from an ONU to the OLT is assumed to be from 10 to 20 km. The length of Ethernet frames randomly varies from 64 to 1518 bytes. Self-similar traffic is generated with the Hurst parameter of 0.7. The total traffic load of the entire network is changing from 0.1 to 0.8; 20 percent, 30 percent, and 50 percent of the traffic are EF, AF, and BE traffic, respectively. FBA is simulated as the benchmark in which the upstream bandwidth is evenly distributed among 32 ONUs and there is no classification among incoming traffic. The service policy is first-in first-out (FIFO). Other studied algorithms include D1 and D2. D1 adopts strict priority scheduling with FBA for EF traffic, and the OLT serves the ONUs after having processed all of the REPORT messages in a service cycle. D2 employs EBR with priority-based scheduling, estimates the EF frames, and serves lightly loaded ONUs ahead of heavily loaded ones. The figures of merits include *average frame delay*, *average queue length*, and *frame loss probability*.

Figure 5a–d illustrates the relationship between average frame delay and network traffic load. The average frame delay is defined as the average time between enqueuing a frame in the buffer and sending out the last bit of the frame. FBA experiences the longest delay even in light network load. This is attributed to the fact that FBA disregards the dynamics of the traffic, and thus lightly loaded ONUs backlog frames at other ONUs. Since D1 pre-reserves a fixed bandwidth for EF traffic, delay of EF frames in D1 is relatively stable, and AF and BE frames in D1 both suffer the longest delay. DBAM outperforms D2 by reporting the bandwidth request with the estimation of all classes. Other factors contributing to the shortest frame delay of DBAM include the instantaneous response of each REPORT message and the unchanged service order among ONUs. All of the above make traffic estimation in DBAM more accurate by not deferring the time slots of heavily loaded ONUs.

The performance of FBA, D1, D2, and DBAM in terms of average queue length exhibits a similar trend to that of average frame delay. Again, as shown in Fig. 5e, FBA has the longest queue, and DBAM has the shortest. A shorter average frame delay means that the ONUs transmit the frames faster, and therefore, less frames are held in the buffer. In FBA, a frame is dropped if the buffer is full. In D1, D2, and DBAM, a frame of BE traffic is dropped if the buffer is full, while an EF or AF frame is dropped if the buffer is full and there are no lower-priority frames in the buffer that can be replaced. The *frame loss probability* is defined as the ratio of the number of dropped frames vs. the total number of frames. Once again, as shown in Fig. 5f, the traffic estimation and immediate bandwidth allocation provisioned in DBAM alleviate the frame loss by requesting the predicted bandwidth, thus reducing the number of backlogged frames in the buffer.

CONCLUSIONS

As an inexpensive, simple, and scalable solution to the access network, EPONs have the capability to deliver integrated broadband services by employing efficient medium access control in the upstream channel. In this study we have investigated the bandwidth allocation issue of the upstream channel in EPONs. Our proposed DBAM algorithm enhances multiple services among ONUs with dynamic and diverse bandwidth requests. Multiservice access for different end users is realized by means of class-based traffic estimation and SLA-limited bandwidth allocation. The simulation performance results show that the fixed service order among ONUs in DBAM enhances the accuracy of traffic estimation, and the improved traffic estimation thus contributes to the reduction of frame delay, queue length, and frame loss.

REFERENCES

[1] 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, pp. 22–31.

[2] IEEE 802.3ah task force, http:// www.ieee802.org/3/efm

[3] G. Kramer and G. Pesavento, "Ethernet Passive Optical Network (EPON): Building a Next-Generation Optical Access Network," *IEEE Commun. Mag.*, vol. 40, no. 2, Feb. 2002, pp. 66–73.

[4] 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.

[5] C. M. Assi *et al.*, "Dynamic Bandwidth Allocation for Quality-of-Service over Ethernet PONs," *IEEE JSAC*, vol. 21, no. 9, Nov. 2003, pp. 1467–77.

[6] IETF RFC 2475, "An Architecture for Differentiated Services," Dec. 1998.

[7] ANSI/IEEE 802.1D, pt. 3, "Media Access Control (MAC) Bridges," 1998.

[8] S. Choi and J. Huh, "Dynamic Bandwidth Allocation Algorithm for Multimedia Services over Ethernet PONs," *ETRI J.*, vol. 24, no. 6, Dec. 2002, pp. 465–68.

[9] W. Willinger *et al.*, "Self-Similarity through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *IEEE/ACM Trans. Net.*, vol. 5, no. 1, Feb. 1997, pp. 71–86.

BIOGRAPHIES

YUANQIU LUO [S'02] (yl6@njit.edu) received her B.S. (honors) degree in electronics and information systems and her M.E. degree in electrical engineering from Shandong University, China, in 1997 and 2000, respectively. She is currently working toward a Ph.D. degree in electrical engineering at New Jersey Institute of Technology (NJIT), Newark. Her research interests include broadband access networks, network survivability, resource allocation, and network modeling. She received first place in the IEEE North Jersey Section Student Presentation Contest.

NIRWAN ANSARI [S'78, M'83, SM'94] (Nirwan.Ansari@njit.edu) received his B.S.E.E. (summa cum laude), M.S.E.E., and Ph.D. from NJIT, University of Michigan, and Purdue University in 1982,1983, and 1988, respectively. He joined the Department of Electrical and Computer Engineering, NJIT, as an assistant professor in 1988, and was promoted to a full professor in 1997. He authored with E.S.H. Hou *Computational Intelligence for Optimization* (1997, translated into Chinese in 2000), and edited with B. Yuhas *Neural Networks in Telecommunications* (1994), both published by Kluwer. He is a technical editor of *IEEE Communications Magazine*, *Computer Communications*, *ETRI Journal*, and *Journal of Computing and Information Technology*. His current research focuses on various aspects of high-speed networks and multimedia communications. He initiated and organized (as General Chair) the First IEEE International Conference on Information Technology: Research and Education (ITRE 2003); was instrumental, while serving as its Chapter Chair, in rejuvenating the North Jersey Chapter of the IEEE Communications Society, which received the 1996 Chapter of the Year Award and a 2003 Chapter Achievement Award; served as Chair of the IEEE North Jersey Section and on the IEEE Region 1 Board of Governors during 2001–2002, and currently serves on various IEEE committees.