Media Access Control for Ethernet Passive Optical Networks: An Overview

Jun Zheng and Hussein T. Mouftah, University of Ottawa

ABSTRACT

Medium access control is one of the crucial issues in the design of Ethernet passive optical networks. To ensure efficient transmission, an EPON system must employ a MAC mechanism to arbitrate access to the shared medium in order to avoid data collisions in the upstream direction and at the same time efficiently share the upstream transmission bandwidth among all ONUs. The purpose of this article is to provide a good understanding of the MAC issue, discuss the major problems involved (e.g., multiple access, bandwidth allocation, transmission scheduling, and quality of service support), and present an overview of the state-of-the-art solutions proposed thus far to the problems. It is also our purpose to motivate further studies on the problems described in this article.

INTRODUCTION

The explosive growth of the Internet and emerging broadband applications, such as Internet telephony, high-definition television (HDTV), interactive games, and video on demand, have imposed a huge demand for network bandwidth on the underlying telecommunications infrastructure. In the past decade, the backbone networks have experienced tremendous growth in bandwidth capacity to meet the ever-increasing bandwidth demand of network users. However, the access networks, which cover the "last mile" area and serve numerous residential and small business users, have not scaled up commensurately and have therefore become a bandwidth bottleneck between end users and backbone networks [1]. To alleviate this bottleneck, an effective solution must not only provide more bandwidth to end users, but also meet the low-cost requirement of access networks. In seeking a cost-effective solution to the problem, passive optical networks (PONs) have received a lot of attention from both industry and academia. A great amount of effort has already been made and is currently going on in developing and standardizing various PON technologies, including IEEE Ethernet PON (EPON) [2] and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Gigabit PON (GPON) [3].

As one promising solution, EPON has attracted tremendous interest in recent years. EPON combines low-cost Ethernet equipment and low-cost passive optical components, and thus has a number of advantages over traditional access networks, such as larger bandwidth capacity, longer operating distance, lower equipment and maintenance cost, and easier updating to higher bit rates [2]. An EPON is a point-to-multipoint fiber optical network with no active elements in the transmission path from source to destination. It can be deployed in different multipoint topologies, such as bus, ring, and tree [2]. The most popular EPON architecture based on a tree topology consists of an optical line terminal (OLT), a 1:N passive star coupler (or splitter/combiner), and multiple optical network units (ONUs), as shown in Fig. 1. The OLT resides in a central office (CO) that connects the access network to a metropolitan area network (MAN) or wide area network (WAN), and is connected to the passive star coupler through a single optical fiber. The passive coupler is located a long distance away from the CO but close to the subscriber premises. Each ONU is located at either the curb or subscriber premises, and is connected to the passive coupler through a dedicated short optical fiber. The distance between the OLT and each ONU typically ranges from 10 to 20 km. In an EPON system all data are encapsulated in Ethernet packets for transmission, which are compatible with IEEE 802.3 Ethernet standards. All transmissions are performed between the OLT and the ONUs. In the downstream direction, an EPON is a point-to-multipoint network in which the OLT broadcasts data to each ONU through the 1:N splitter, where N is typically between 4 and 64. Each ONU extracts the data destined for it based on its medium access control (MAC) address. In the upstream direction, an EPON is a multipoint-to-point network in which multiple ONUs transmit data to the OLT through the 1:N passive combiner. The line data rate from an ONU to the OLT and the user access rate from a user to an ONU do not nec-

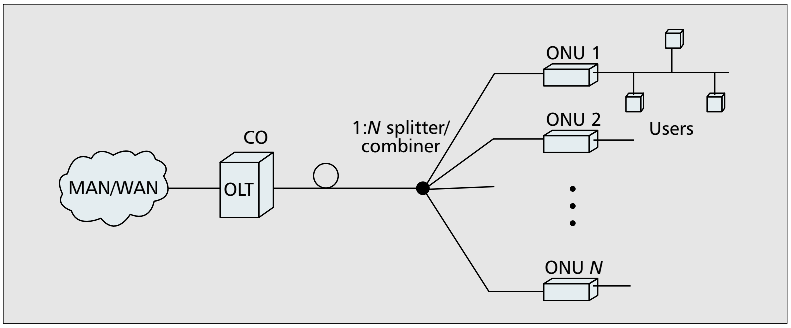

■ **Figure 1.** *EPON architecture.*

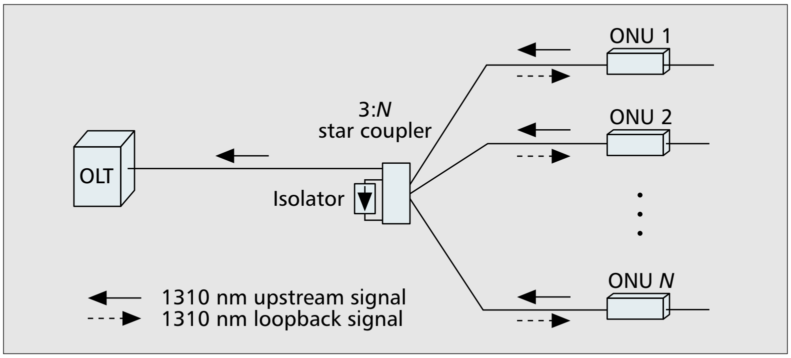

■ **Figure 2.** *An EPON using a loopback star coupler.*

essarily have to be equal, and the line data rate is usually much higher than the user access rate. Since all ONUs share the same upstream transmission medium, an EPON system must employ a MAC mechanism to arbitrate access to the shared medium in order to avoid data collisions in the upstream direction and at the same time efficiently share the upstream transmission bandwidth among all ONUs.

The purpose of this article is to provide a good understanding of the MAC issue, discuss the major problems involved, and present an overview of the state-of-the-art solutions to these problems.

Medium Access Control

Channel separation for upstream and downstream transmissions, and multiple access for upstream transmission are essential to the MAC problem.

Channel Separation

To increase the transmission efficiency of an EPON system, the upstream and downstream transmission channels should be separated appropriately. A simple solution is to use space-division multiplexing, where two separate optical fibers and passive couplers are used, one for upstream and the other for downstream transmission. To reduce network cost, a more attractive solution is to use a single coupler and a single fiber for both directions with one wavelength for upstream transmission and another for downstream transmission. Currently, the most popular solution to channel separation is to use a 1550 nm wavelength for downstream transmission and another 1310 nm wavelength for upstream transmission [2].

Multiple Access

In the upstream direction, multiple ONUs transmit data packets to the OLT through the common passive combiner and share the same optical fiber from the combiner to the OLT. Due to the directional property of a passive combiner, data packets from an ONU can only reach the OLT but not the other ONUs. For this reason, conventional contention-based multiple access (e.g., carrier sense multiple access with collision detection, CSMA/CD), is difficult to implement because the 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 largely reduced because of considerable propagation delay between the OLT and the ONUs. To address this problem, an optical loopback technique was proposed in [4] to achieve high channel efficiency with CSMA/CD. With this loopback technique, a portion of the upstream signal power transmitted by each ONU is looped back to the other ONUs at the star coupler by using a $3 \times N$ coupler and connecting two ports of the coupler together through an isolator, as shown in Fig. 2. If two or more ONUs transmit data simultaneously, collisions will be detected at each ONU and all data transmissions will be stopped immediately. The optical CSMA/CD protocol is applied to all upstream transmissions [5]. The OLT will receive the data packets transmitted by each ONU and discard those packets with collisions. However, to implement the optical CSMA/CD protocol, each ONU has to use an additional receiver operating at the upstream wavelength and a carrier sensing circuit, which would largely increase the network cost. On the other hand, contention-based multiple access is unable to provide guaranteed bandwidth to each ONU. So it is difficult to support any form of quality of service (QoS). For these reasons, contention-based multiple access is currently not the preferred solution for upstream multiple access.

Another 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 either a tunable receiver or a receiver array at the OLT to receive the data transmitted in multiple channels. In particular, it also requires each ONU to use a fixed transmitter operating at a different wavelength, which could result in an inventory problem. Although the inventory problem can be solved by using tunable transmitters, such devices are costly at the current stage. Due to these facts, the WDM solution is cost prohibitive and thus not feasible in the near term.

Compared to contention-based multiple access and WDM, time-division multiple access (TDMA) on a single wavelength is more attractive for upstream transmission. In this solution, each ONU is allocated a time slot or transmission window for data transmission by the OLT. Each time slot is capable of carrying several Ethernet packets. Packets received from one or more users are buffered in an ONU until the time slot for that ONU arrives. Upon the arrival of its time slot, the ONU will send out its buffered packets at the full transmission rate of the upstream channel. Accordingly, TDMA avoids data collisions from different ONUs. Moreover, it requires only a single wavelength for all ONU transmissions and a single transceiver at the OLT, which is highly cost effective.

In TDMA the time slot size allocated to each ONU can be either fixed or variable. Fixed time slot allocation, also called static bandwidth allocation, is simple to implement. However, due to the bursty nature of network traffic, it may result in a situation in which some time slots overflow even under very light load, causing packets being delayed for several time slots, while other time slots are not fully used even under very heavy traffic, leading to the upstream bandwidth being underutilized. For this reason, static allocation is not preferred. To increase bandwidth utilization, it is desirable that the OLT dynamically allocate a variable time slot to each ONU based on the instantaneous bandwidth demand of the ONUs. To this end, a polling mechanism has been widely used [6–8]. With polling, the OLT can dynamically allocate bandwidth for each ONU and flexibly arbitrate the transmissions of multiple ONUs, which can significantly increase bandwidth utilization and improve network performance.

Polling Protocol Design

In this section we first introduce the multipoint control protocol (MPCP) being standardized by the IEEE 802.3ah Ethernet in the First Mile Task Force [9] and then discuss several problems in designing a polling protocol based on MPCP.

Multipoint Control Protocol

MPCP is a signaling protocol for facilitating dynamic bandwidth allocation and arbitrating the transmissions of multiple ONUs. It resides at the MAC control layer and has two operation modes: normal and auto-discovery. In normal mode, MPCP relies on two Ethernet control messages, GATE and REPORT, to allocate bandwidth to each ONU. The GATE message is used by the OLT to allocate a transmission window to an ONU. The REPORT message is used by an ONU to report its local conditions to the OLT. In auto-discovery mode, the protocol relies on three control messages, REGISTER, REGISTER_REQUEST, and REGISTER_ACK, which are used to discover and register a newly connected ONU, and collect relevant information about that ONU such as round-trip delay and MAC address.

In normal operation, MPCP in the OLT gets a request from the higher MAC client layer to transmit a GATE message to a particular ONU. Upon getting such a request, MPCP will timestamp the GATE message with its local time and then send the message to the ONU. The GATE message typically contains a granted start time, a granted transmission window, and a 4-byte timestamp, which is used to calculate the round-trip time between the OLT and the ONU. Once the ONU receives the GATE message, it programs its local register with the values contained in the GATE message. Meanwhile, it also updates its

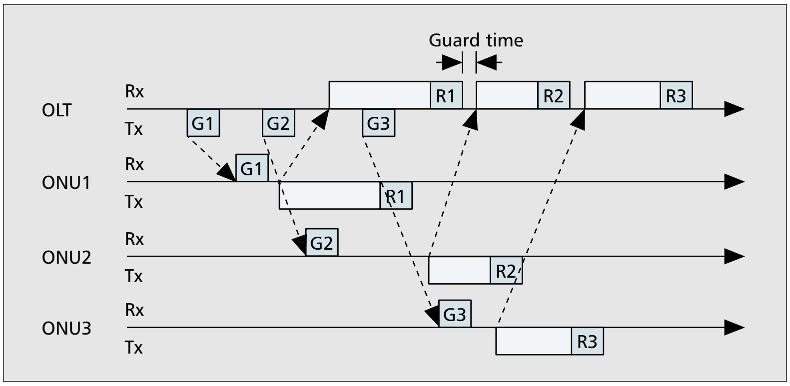

Figure 3. *A flow of GATE and REPORT messages.*

local clock to that of the timestamp extracted from the received GATE message in order to maintain synchronization with the OLT. At the granted start time, the ONU will start to transmit data for up to the window size. The transmission may include multiple data packets, depending on the window size and the queue length in the ONU. No packet fragmentation is allowed during transmission. If the next packet cannot be transmitted in the current window, it will be deferred to the next window.

A REPORT message is sent by an ONU in the allocated transmission window together with a data packet. It can be transmitted automatically or on demand at either the start or the end of a window. A REPORT is generated at the MAC client layer and is timestamped at the MAC layer. It typically contains the bandwidth demand of an ONU based on the instantaneous queue length of that ONU. The ONU should also account for additional overhead in its request, including a 64-bit frame preamble and a 96-bit interframe gap associated with each Ethernet packet. Once a REPORT message is received by the OLT, it is passed to the MAC client layer, which is responsible for bandwidth allocation and recalculation of the round-trip delay to the source ONU. Figure 3 illustrates a flow of GATE (G) messages and REPORT (R) messages for upstream transmission of three ONUs.

It should be pointed out that MPCP is not concerned with any particular bandwidth allocation scheme and transmission scheduling algorithm, and allows them to be vendor-specific. To design an efficient polling protocol based on MPCP, several problems must be considered, including maximum bandwidth limit, channel utilization, and transmission scheduling.

Maximum Bandwidth Limit

A polling protocol is typically cycle-based. In each polling cycle, each ONU is polled once and is allocated a transmission window based on its bandwidth demand. If the OLT allows each ONU to send all its buffered packets in one transmission, ONUs with high traffic load may monopolize the entire bandwidth of the upstream channel. This is unfair to those ONUs with low traffic load. To address this problem, the OLT should limit the maximum transmission bandwidth of each ONU. The maximum window size can be either fixed based on some criterion,

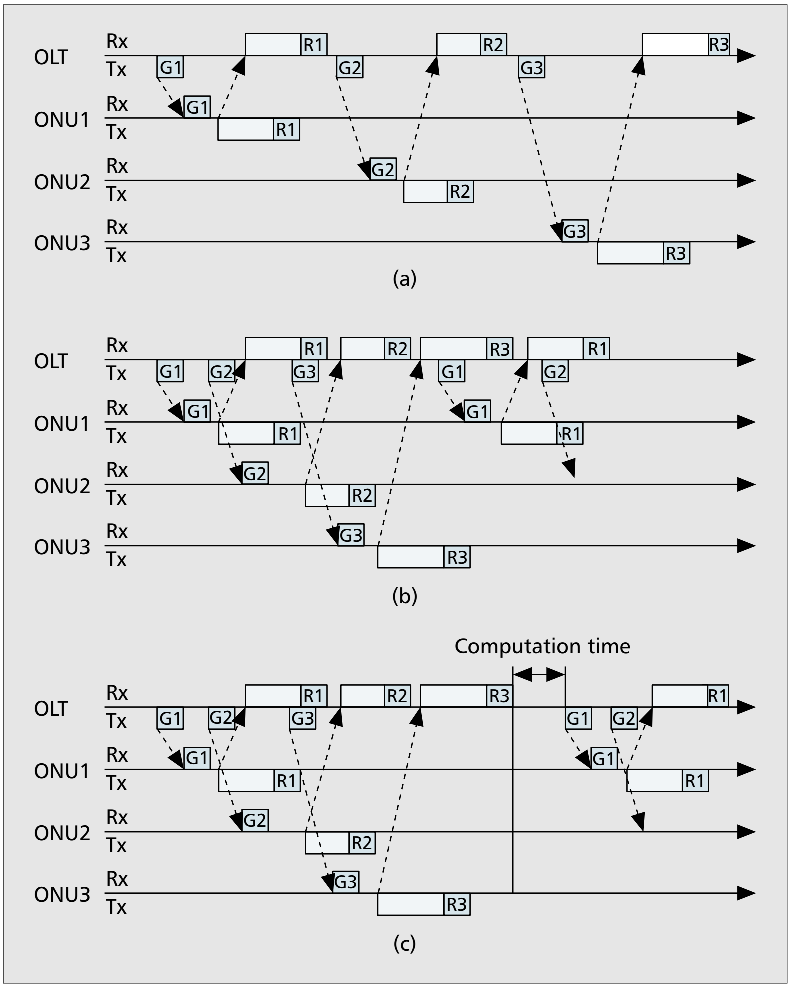

■ **Figure 4.** *Polling policies: a) poll-and-stop polling; b) interleaved polling; c) interleaved polling with stop.*

such as a service level agreement (SLA), or variable based on instantaneous network conditions. Under high traffic load, the maximum window size determines the maximum polling cycle. In general, making the maximum polling cycle too long will result in larger delay for all packets under high traffic load, including high-priority packets. On the other hand, making the maximum cycle too short will result in more bandwidth being wasted by interframe gaps (or guard times). Accordingly, the maximum window size has a great impact on network performance.

While the maximum window size imposes a limit on the maximum bandwidth that can be allocated to each ONU in each polling cycle, it is also the guaranteed bandwidth available to each ONU. In fact, only when all other ONUs use all their available bandwidth will an ONU be limited to its guaranteed bandwidth. If any ONU requests less bandwidth, it will be allocated a smaller window size, making the polling cycle shorter and thus increasing the actual bandwidth available to all other ONUs. This is the benefit of dynamic bandwidth allocation.

Channel Utilization

Since the upstream channel is shared by multiple ONUs, channel utilization is of great concern. A polling protocol can poll multiple ONUs for transmission based on different policies. A simple policy, called poll-and-stop polling, is to send a GATE message to an ONU and then stop for the data and REPORT message to come back from that ONU before the OLT sends a GATE message to the next ONU, as shown in Fig. 4a. Obviously, this protocol wastes a lot of bandwidth on the upstream channel, which would largely reduce channel utilization and increase packet delay.

A more efficient way is to use interleaved polling [5], which allows the OLT to send a GATE message to the next ONU before the data and REPORT message(s) from the previous polled ONU(s) arrive, as shown in Fig. 4b. This is feasible because the upstream and downstream channels are separated, and the OLT maintains relevant information about each ONU in a polling table, including the bandwidth demand of and round-trip time to each ONU. The results obtained in [5] indicate that the interleaved polling protocol can significantly improve network performance in terms of channel utilization and average packet delay. However, this protocol allows the OLT to allocate bandwidth only based on already received bandwidth demands. The OLT is unable to take into account the bandwidth demands of all ONUs and make a more intelligent decision on bandwidth allocation.

An effective way to overcome this drawback is to use a variation of interleaved polling called interleaved polling with stop. Like interleaved polling, this protocol allows the OLT to send a GATE message to the next ONU before the transmission and REPORT message(s) from the previous polled ONU(s) arrive. Unlike interleaved polling, the OLT does not start the next polling cycle before the transmissions and REPORT messages from all ONUs are received. This allows the OLT to perform bandwidth allocation based on the bandwidth demands of all ONUs at the end of each polling cycle and thus make a more intelligent decision. However, such intelligence is obtained at the cost of upstream channel utilization because the upstream channel is not utilized from the instant the transmission of the last polled ONU in the previous cycle is completed to the instant the transmission of the first polled ONU in the next cycle starts. Figure 4c illustrates an example of control message flows with interleaved-polling-with-stop.

Transmission Scheduling

To ensure efficient transmission, a polling protocol must schedule the transmissions of multiple ONUs in a manner that avoids data collisions from different ONUs. This is not difficult to implement because such scheduling is based on the granted window size and the round-trip time to each ONU. Since the OLT knows the granted window size and the round-trip time to the last polled ONU, it can calculate the transmission start time and window size for the next ONU. Note that to allow the receiver in the OLT to prepare for receiving the transmissions, a minimum gap or guard time is usually required between the transmissions of different ONUs.

On the other hand, the OLT must also be responsible for scheduling the transmission order of different ONUs, which may have a great impact on network performance. This is not difficult to implement because the order of the transmissions is usually determined one cycle ahead by performing a scheduling algorithm. The most widely used scheduling algorithm is round-robin (RR), which has been adopted by many polling protocols [6-8]. RR schedules the transmissions of different ONUs in the order of their indexes in the polling table and is simple to implement. However, it does not take into account the instantaneous traffic conditions at each ONU and thus may not be able to provide the best performance in terms of packet delay and data loss. To improve network performance, it is desirable to use an adaptive scheduling algorithm that can dynamically schedule the order of different ONU transmissions based on the instantaneous traffic conditions at each ONU. For example, an adaptive scheduling algorithm can schedule ONU transmissions in descending order of the instantaneous queue length of each ONU (i.e., longest queue first, LQF) or ascending order of the arrival time of the first packet queuing in each ONU (i.e., earliest packet first, EPF) [10].

Dynamic Bandwidth Allocation

Due to limited upstream bandwidth, an EPON system may not always be able to provide sufficient bandwidth to meet the bandwidth demand of all end users. To better serve end users, it is desirable to dynamically allocate bandwidth to each ONU based on the instantaneous traffic conditions of the ONUs. In this context several bandwidth allocation schemes have already been proposed in [6], including limited allocation, constant credit, linear credit, and elastic allocation.

In the limited allocation scheme, the OLT simply grants an ONU the number of bytes the ONU requested, not exceeding a maximum window size. This is the most conservative scheme because it assumes that no more packets arrived after the ONU sent its request. In practice, however, because of the round-trip time between the OLT and each ONU, there might be more packets arriving between the instant an ONU sends a REPORT message and the instant the ONU receives a GATE message. In this case, those newly-arriving packets may not be able to be transmitted in the current cycle, resulting in increased average packet delay. To address this problem, the constant credit and linear credit schemes were proposed.

In the constant credit scheme, a credit is added to the requested window size and is considered in the granted window size. The size of the credit is constant no matter how large the requested window size. Once an ONU receives a GATE message, it can send packets for up to the requested window size plus the constant credit. The choice of credit size may have an impact on network performance. A too small size will not be able to improve packet delay a lot. A too large size will reduce the bandwidth utilization of the upstream channel. The choice should be based on traffic characteristics or some empirical data.

In the linear credit scheme, a similar credit is added to the requested window size. However, the size of the linear credit is proportional to the requested window size. The basis of this scheme is that network traffic usually has a certain degree of predictability. This means that if a long burst of data is observed, this burst is very likely to continue for a longer time.

In the elastic allocation scheme, there is no limit imposed on the maximum window size. The only limit is the maximum cycle time. The maximum window size W_{max} is granted in such a way that the accumulated size of the last N grants (including the one being granted) does not exceed $N \times W_{max}$, where N is the number of ONUs. Thus, if only one ONU has data to send, it may get a granted window size up to $N \times W_{max}$.

Among all these schemes, the limited scheme exhibits the best performance [2].

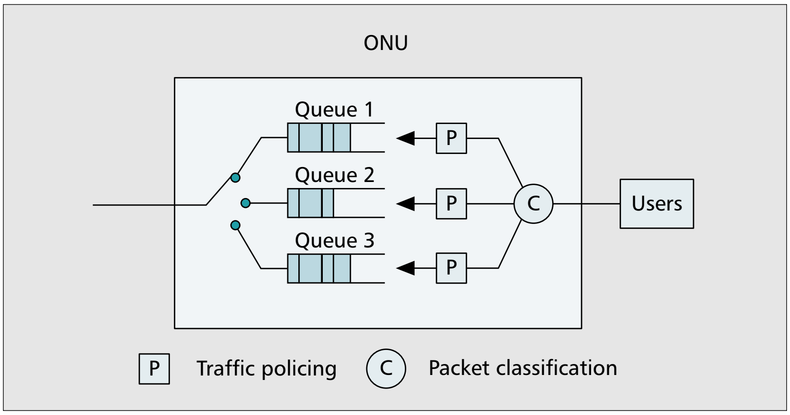

Figure 5. *Priority queuing and intra-ONU scheduling.*

Quality of Service Support

To support a variety of network services with diverse QoS requirements, an EPON system must consider differentiated QoS in its MAC design.

Priority Queuing

An effective way to support differentiated QoS is to use priority queuing. With priority queuing, network traffic is classified into a set of classes with diverse QoS requirements, and for each traffic class a priority queue is maintained at each ONU. Figure 5 illustrates an example of priority queuing in which an ONU maintains three priority queues that share the same memory buffer of fixed size. Data packets from end users are first classified by checking the type of service (ToS) field of the IP packets encapsulated in the Ethernet packets and then buffered in corresponding priority queues. If a higher-priority packet finds the buffer full at the time of its arrival, it can preempt a lower-priority packet. If a lower-priority packet arrives and finds the buffer full, it will be dropped. As a result, lower-priority traffic may experience very high packet loss and even resource starvation. To address this problem, an ONU should perform some kind of traffic policing to control the amount of higher-priority traffic from each end user.

With future advances in enabling technologies, optical devices that are currently costly may become affordable. That would open up new research opportunities for new EPON architectures and related MAC problems.

INTER-ONU SCHEDULING AND INTRA-ONU SCHEDULING

In supporting differentiated QoS, there are two types of scheduling paradigms: inter-ONU and intra-ONU scheduling. Inter-ONU scheduling is responsible for arbitrating the transmissions of different ONUs, and intra-ONU scheduling is responsible for arbitrating the transmissions of different priority queues in each ONU. There are two strategies to implement these two scheduling paradigms. One is to allow the OLT to perform both inter-ONU and intra-ONU scheduling. In this case, the OLT is the only device that arbitrates upstream transmissions. Each ONU can request the OLT to allocate bandwidth for each traffic class. For this purpose, an ONU must report the status of its individual priority queues to the OLT through REPORT messages. MPCP specifies that each ONU can report the status of up to eight priority queues [8]. The OLT can then generate multiple grants, each for a specific traffic class, to be sent to the ONU using a single GATE message. The format of the 64-byte MPCP GATE message can be found in [9].

The other strategy is to allow the OLT to perform inter-ONU scheduling and allow each ONU to perform intra-ONU scheduling. In this case, each ONU requests the OLT to allocate bandwidth for it based on its buffer occupancy status. The OLT only allocates the requested bandwidth to each ONU. Each ONU will divide the allocated bandwidth among different classes of services based on their QoS requirements and schedule the transmissions of different priority queues within the allocated bandwidth. For intra-ONU scheduling, there are two types of scheduling algorithms: strict and non-strict priority scheduling. In strict priority scheduling, a lower-priority queue is scheduled only if all queues with higher priority are empty. Obviously, this may potentially result in infinite packet delay and high packet loss for low-priority traffic. To address the problem, a non-strict priority scheduling algorithm was proposed in [7]. In non-strict priority scheduling, only those packets that were reported are transmitted first as long as they can be transmitted within the allocated time slot. The transmission order of different priority queues is based on their priorities. If the packets that were reported are all scheduled and the current time slot can still accommodate more packets, newly arriving packets that were not reported are also transmitted based on their priorities. As a result, all traffic classes can have access to the upstream channel within the allocated time slot as reported to the OLT while their priorities are maintained, which ensures fairness in scheduling.

CONCLUSIONS

The performance and QoS guarantees of EPON systems largely depend on efficient MAC protocols. We discuss the major MAC problems in EPONs and present an overview of various solutions proposed thus far to those problems. Since MPCP is being standardized by the IEEE 802.3ah Task Force, current research efforts are primarily focused on bandwidth allocation and transmission scheduling, seeking more efficient solutions to QoS provisioning. With future advances in enabling technologies, optical devices that are currently costly may become affordable. That would open up new research opportunities for new EPON architectures and related MAC problems.

REFERENCES

[1] J. Zheng and H. T. Mouftah, *Optical WDM Networks: Concepts and Design Principles*, Wiley-IEEE Press, 2004.

[2] 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.

[3] J. D. Angelopoulos *et al.*, "Efficient Transport of Packets with QoS in an FSAN-Aligned GPON," *IEEE Commun. Mag.*, vol. 42, no. 2, Feb. 2004, pp. 92–98.

[4] B. N. Nesai *et al.*, "An Optical Implementation of a Packet-Based (Ethernet) MAC in a WDM Passive Optical Network Overlay," *Proc. OFC '01*, vol. 3, 2001, pp. WN5-1–3.

[5] C.-J. Chae *et al.*, "Optical CSMA/CD Media Access Scheme for Ethernet over Passive Optical Network," *IEEE Phot. Tech. Lett.*, vol. 14, no. 5, May 2002, pp. 711–13.

[6] 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.

[7] M. Ma *et al.*, "A Bandwidth Guaranteed Polling MAC Protocol for Ethernet Passive Optical Networks," *Proc. IEEE INFOCOM '03*, vol. 1, Mar./Apr. 2003, pp. 22–31.

[8] 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.

[9] IEEE 802.3ah, Ethernet in the First Mile Task Force, http://www.ieee802.org/3/efm/index.html

[10] J. Zheng and H. T. Mouftah, "An Adaptive MAC Polling Protocol for Ethernet Passive Optical Networks (EPONs)," *Proc. IEEE ICC '05*, May 2005.

BIOGRAPHY

JUN ZHENG [M] (junzheng@site.uottawa.ca) is a research scientist with the University of Ottawa, Canada. He received his Ph.D. in electrical and electronic engineering from the University of Hong Kong in 2000. He has conducted extensive research in the field of telecommunications and computer networks, ranging from IP networks to ATM networks, wireless networks, and optical networks. His current research interests include MAC protocol design for shared medium access networks. He is the first author of *Optical WDM Networks: Concepts and Design Principles*, and has published over 40 technical papers in international journals and conference proceedings. He is lead guest editor for a feature issue on optical access networks of *OSA Journal of Optical Networking* and a special issue on wireless sensor networking of *IEEE Network*. He is serving as technical program chair for the 2005 International Conference on Sensor Networks and has served on the technical program committees of many international conferences, including IEEE ICC and GLOBECOM.

HUSSEIN T. MOUFTAH [F] (mouftah@site.uottawa.ca) joined the School of Information Technology and Engineering (SITE) of the University of Ottawa in September 2002 as a Canada Research Chair (Tier 1) Professor in Optical Networks. He was with the ECE Department at Queen's University (1979–2002), where he was prior to his departure a full professor and Department associate head. He has three years of industrial experience mainly at BNR of Ottawa, now Nortel Networks (1977–1979). He served as Editor-in-Chief of *IEEE Communications Magazine* (1995–1997) and IEEE ComSoc Director of Magazines (1998–1999), and Chair of the Awards Committee (2002–2003). He has been a Distinguished Speaker of IEEE ComSoc since 2000. He is the author or coauthor of five books, 22 book chapters, and more than 700 technical papers and 8 patents. He was the recipient of the 1989 Engineering Medal for Research and Development of the Association of Professional Engineers of Ontario (PEO), and the Ontario Distinguished Researcher Award of the Ontario Innovation Trust (2002). He is the joint holder of the Best Paper Award for a paper presented at SPECTS 2002, and the Outstanding Paper Awards for IEEE HPSR 2002 and the IEEE ISMVL '85.