Toward Energy-Efficient 1G-EPON and 10G-EPON with Sleep-Aware MAC Control and Scheduling

Jingjing Zhang and Nirwan Ansari, New Jersey Institute of Technology

ABSTRACT

To rapidly meet increasing traffic demands from subscribers, the IEEE 802.3av task force was charged to study the 10G-EPON system to increase the data rate to 10 times the line rate of 1G-EPON. With the increase of the line rate, the energy consumption of 10G-EPON increases accordingly. Achieving low energy consumption with 10G-EPON has attracted broad research attention from both academia and industry. In this article we first briefly discuss the key features of 10G-EPON. Then, from the perspective of MAC-layer control and scheduling, we discuss challenges and possible solutions to put optical network units into low-power mode for energy saving. More specifically, we detail sleep-mode control and sleep-aware traffic scheduling schemes in two scenarios: sleep for over one DBA cycle, and sleep within one DBA cycle.

INTRODUCTION

As one of the major fiber to the home/curb/cabinet/and so on (FTTx) technologies, Ethernet passive optical network (EPON) is developed based on Ethernet technology, and enables seamless integration with IP and Ethernet technologies [1]. Due to the advantages of fine scalability, simplicity, and multicast convenience, as well as the capability of providing full-service access, EPON has been rapidly adopted in Japan and is also gaining momentum with carriers in China, Korea, and Taiwan since the IEEE ratified EPON as the IEEE 802.3ah standard in June 2004.

On the other hand, video-centric applications and services such as HDTV are growing and emerging in the network [2]. As compared to traditional voice and data traffic, these multimedia applications are more bandwidth-hungry. For example, one HDTV channel requires as much as 10 Mb/s bandwidth. Motivated by satisfying these emerging high bandwidth demands, the IEEE 802.3av 10G-EPON task force was charged to increase the downstream bandwidth to 10 Gb/s, and to support two upstream data rates: 10 Gb/s and 1 Gb/s.

While the line rate is significantly increased to satisfy subscribers' demands, the power consumption of 10G-EPON may be increased as well [3]. Power consumption of 10G-EPON has become a big concern of network service providers as it contributes to part of their operational expenditure (OPEX). Moreover, energy consumption is becoming an environmental and therefore social and economic issue because one big reason for climate change is the burning of fossil fuels and the direct impact of greenhouse gases on the Earth's environment [4]. Previously, Baliga *et al*. [5] estimated that the Internet currently consumes around one percent of the total electricity consumption in broadband enabled countries. It is also shown that currently and in the medium term future, access networks consume the majority of the energy in the Internet. The analysis in [5, 6] showed that, among various access technologies, such as WiMAX, fiber to the node (FTTN), and point-to-point optical access networks, PON is the most power-efficient solution in terms of energy consumption per transmission bit attributed to the nearest approach of optical fibers to users.

Although PON consumes the smallest power among all access network technologies, it is desirable to further reduce its power consumption, especially when the line rate is increased to 10 Gb/s. With the increase in line rate, the optical dispersion increases as well. Compensating for higher dispersion exerts higher requirements on optical lasers, which may incur an increase of power consumption of lasers. In addition, the electronic circuit should be sufficiently powered such that it can process 10 times faster than that of 1G-EPON. Thus, 10G-EPON will consume more energy than 1G-EPON [4, 5].

Reducing power consumption of 10G-EPON requires efforts across both the physical and medium access control (MAC) layers. Efforts are being made to develop optical transceivers and electronic circuits with low power consumption. Besides, multi-power-mode devices with the ability of disabling certain functions can also help reduce the energy consumption of the network. However, low-power-mode devices with some functions disabled may result in degradation of network performances. To avoid service degradation, it is important to properly design

0163-6804/11/$25.00 © 2011 IEEE

	Data rate (Gb/s)		FEC	Line coding	Tx type and launch power (dbm)		
	Upstream	Downstream					
1G-EPON	1.25	1.25	Optional RS(255, 239)	8b/10b	PX10: OLT: [–3,2] ONU: [–1,4]	PX20: OLT: [2,7] ONU: [–1,4]	
10G-EPON	10.3125	10.3125	Enabled RS(255, 223)	64b/66b	PR10: OLT: EML [1,4] ONU: DML [–1,4]	PR20: OLT: EML+AMP [5,9] ONU: DML [–1,4]	PR30: OLT: EML [2,5] ONU: HP DML [4,9]
	1.25	10.3125	Enabled RS(255,223)	64b/66b	PRX10: OLT: EML [1,4] ONU: DML [–1,4]	PRX20: OLT: EML+AMP [5,9] ONU: DML [–1,4]	PRX30: OLT: EML [2,5] ONU: DML [.6,5.6]

Table 1. *Comparison between 1G-EPON and 10G-EPON.*

MAC-layer control and scheduling schemes that are aware of the disabled functions. This is the focus of this article.

We first discuss the evolution from 1G-EPON to 10G-EPON. Then, we discuss the challenges in empowering optical network units (ONUs) in low-power mode. Finally, we detail our proposed control and scheduling schemes to achieve energy saving without degrading user services.

EVOLUTION FROM 1G-EPON TO 10G-EPON

10G-EPON supports both symmetric 10 Gb/s downstream and upstream, and asymmetric 10 Gb/s downstream and 1 Gb/s upstream data rates, while 1G-EPON provides only the 1 Gb/s symmetric data rate. With a focus on the physical layer, the IEEE 802.3av Task Force specifies the reconciliation sublayer (RS), symmetric and asymmetric physical coding sublayers (PCSs), physical media attachments (PMAs), and physical media-dependent (PMD) sublayers. Table 1 lists several key physical layer features of 10G-EPON [7]. Instead of using the 8B/10B line coding adopted in 1G-EPON, 10G-EPON employs 64B/66B line coding, with which the bit-to-baud overhead is reduced to as small as 3 percent. To relax the requirements for optical transceivers, Reed-Solomon code (255, 223) is chosen as the mandatory forward error correction (FEC) code in 10G-EPON to enhance the FEC gain, while Reed-Solomon code (255, 239) is specified as optional for 1G-EPON. 10G-EPON defines the PRX power budget for asymmetric-rate PHY of 10 Gb/s downstream and 1 Gb/s upstream, and the PR power budget for symmetric-rate PHY of 10 Gb/s both upstream and downstream. Each power budget further contains three power budget classes: low power budget (PR(X)10), medium power budget (PR(X)20), and high power budget (PR(X)30). PR(X)10 and PR(X)20 power budget classes are defined in 1G-EPON as well, while PR(X)30, which can support 32-split with a distance of at least 20 km, is an additional one defined in 10G-EPON. Due to limited space, we only list the transmitter (Tx) type along with its launch power of 10G-EPON in Table 1. As compared to 1G-EPON, advanced transmitters and higher launch power are employed in 10G-EPON to guarantee a sufficient signal-to-noise ratio (SNR) at the receiver side for accurate recovery of data with a rate of 10 Gb/s. Because of the increased launch power, the power consumption of the optical transmitter should be increased accordingly. Also, due to the mandatory FEC mechanism and increased line rate, the electronic circuit has to enable more functions and process faster than that in 1G-EPON, thus consequently incurring high power consumption and possibly larger heat dissipation. Therefore, to accommodate 10 Gb/s in the physical layer, the power consumption of the optical line terminal (OLT) and ONU may increase significantly.

For the MAC layer and layers above, in order to achieve backward compatibility such that network operators are encouraged to upgrade their services, 10G-EPON keeps the EPON frame format, MAC layer, MAC control layer, and all the layers above almost unchanged from 1G-EPON. This further implies that similar network management system (NMS), PON-layer operations, administrations, and maintenance (OAM) system, and dynamic bandwidth allocation (DBA) used in 1G-EPON can be applied to 10G-EPON as well.

Next, with a focus on MAC layer control and DBA, we propose a scheme to reduce energy consumption of 1G-EPON and 10G-EPON. In this article, we focus on reducing the energy consumption of ONUs.

CHALLENGES IN SAVING ENERGY OF 1G-EPON AND 10G-EPON

Formerly, the *sleep* mode was proposed to be introduced into ONUs to save energy when ONUs are idle [8–10]. ITU-T Recommendation G. sup 45 [11] specified two energy saving modes for ONUs in GPON. One is *doze* mode, in which only the transmitter can be turned off when possible. Another one is *cyclic sleep* mode, in which both transmitter and receiver can be turned off. Since the access network traffic is rather bursty [12], ONUs may be idle for quite long periods, implying that putting idle ONUs into *sleep* mode is an effective way to reduce energy consumption. However, it is challenging to wake up *sleep* ONUs in time to avoid service disruption when downstream or upstream traffic arrives in 1G-EPON and 10G-EPON.

The major challenge lies in downstream transmission. In EPON, the downstream data traffic of all ONUs is time-division multiplexed (TDM) into a single wavelength, and is then broadcasted to all ONUs. An ONU receives all downstream packets and checks whether the packets are destined to itself. An ONU does not know when the downstream traffic arrives at the OLT and the exact time the OLT schedules its downstream traffic. Therefore, without proper sleep-aware MAC control, receivers at ONUs need to be awake all the time to avoid missing their downstream packets.

To address this problem, Mandin [9] proposed to implement a three-way handshake process between OLT and ONUs before putting ONUs to sleep. Since an OLT is aware of the *sleep* status of ONUs, it can queue the downstream arrival traffic until the *sleep* ONU wakes up. However, to implement the three-way handshake process, extended multipoint control protocol (MPCP) is required to introduce new MPCP packet data units (PDUs). In addition, the negotiation process takes at least several round-trip times that further incurs large delay. Lee *et al*. [13] proposed to implement fixed bandwidth allocation (FBA) when the network is lightly loaded. By using FBA, the time slots allocated to each ONU in each cycle are fixed and known to the ONU, and thus ONUs can go to sleep in the time slots allocated to other ONUs. However, since traffic of an ONU changes dynamically and from cycle to cycle, FBA may result in bandwidth under- or overallocation, and consequently degrade services of ONUs to some degree.

Besides the downstream scenario, an efficient sleep control mechanism should also consider upstream traffic and MPCP control message transmission. For upstream transmission, the wake-up of a *sleep* ONU can be triggered by the arrival of upstream traffic. However, this arrival traffic cannot be transmitted until the ONU is notified of the allocated time from the OLT. Before the OLT allocates bandwidth to an ONU, the newly *awake* ONU first needs to request upstream bandwidth. To realize this, some periodic time slots may need to be allocated to ONUs to enable them to access the upstream channel in time even when they are asleep.

Regarding the MPCP control message transmission, to keep a watchdog timer in the OLT from expiring and deregistering the ONU, both IEEE 802.3ah and IEEE 802.3av specify that ONUs should send MPCP REPORT messages to the OLT periodically to signal bandwidth needs as well as to arm the OLT watchdog timer even when no request for bandwidth is being made. The longest interval between two reports as specified by `report_timeout` is set as 50 ms in both 1G-EPON and 10G-EPON. Besides, the OLT also periodically sends GATE messages to an ONU even when the ONU does not have data traffic. The longest interval between two GATE messages, specified by `gate_timeout`, is set as 50 ms. Therefore, to comply with MPCP, *sleeping* ONUs need to wake up every 50 ms to send the MPCP REPORT messages and receive the GATE messages.

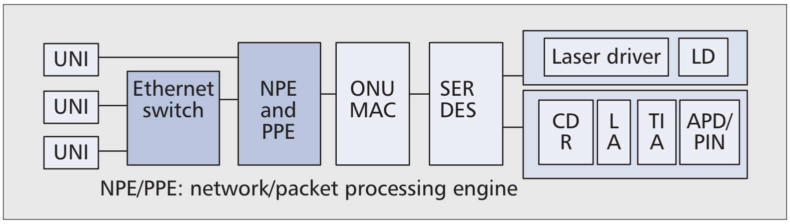

Figure 1. *The constitution of an ONU.*

SLEEP-AWARE MAC CONTROL AND SCHEDULING

In this section, we discuss our proposals on saving energy in ONUs. Our basic idea is still putting ONUs into sleep mode whenever possible. Different from existing proposals, which put the whole ONU to sleep, we investigate the constitution of an ONU and put different components of an ONU to sleep under different conditions.

SLEEP STATUS OF ONUS

Figure 1 illustrates the typical constitution of an ONU. The optical module consists of an optical transmitter (Tx) and an optical receiver (Rx). The electrical module mainly contains serializer/deserializer (SERDES), ONU MAC, network/packet processing engine (NPE/PPE), Ethernet switch, and user-network interfaces (UNIs). When neither upstream nor downstream traffic exists, every component in the ONU can be put to *sleep*. When only downstream traffic exists, the functions related to upstream transmission can be disabled. Similarly, the functions related to receiving downstream traffic can be disabled when only upstream traffic exists. Even when upstream traffic exists, the laser driver and laser diode (LD) do not need to function all the time, but only during the time slots allocated to this ONU. Thus, each component in the ONU can likely *sleep*, and potentially higher power savings can be achieved.

By putting each component of an ONU to sleep, an ONU ends up with multiple power levels. Figure 2 shows the power levels of an ONU and the transition between different power levels. The *wakeup* of UNI, NPE/PPE, and switch can be triggered by the arrival of upstream traffic and the forwarding of downstream traffic from the ONU MAC [14]. They are relatively easily controlled as compared to the other components. Thus, we only focus on the ONU MAC, SERDES, Tx, and Rx. As shown in Fig. 2a, two power levels, *all:awake* and *all:sleep*, result from putting the whole ONU to sleep. In our proposal two more sleep statuses, *Rx:sleep* and *Tx:sleep*, are introduced; thus, four power levels are generated. When the ONU is in the *all:awake* status, if Tx does not need to work, it enters into *Tx:sleep* status, and further enters into *all:sleep* status if Rx does not need to function either. In *all:sleep* status, besides Rx and Tx, the ONU MAC and SERDES *sleep* as well. Similarly, transitions happen between the *all:awake*, *Tx:sleep*, and *Rx:sleep* statuses.

Transitioning among these statuses should be properly designed so as to maximize energy sav-

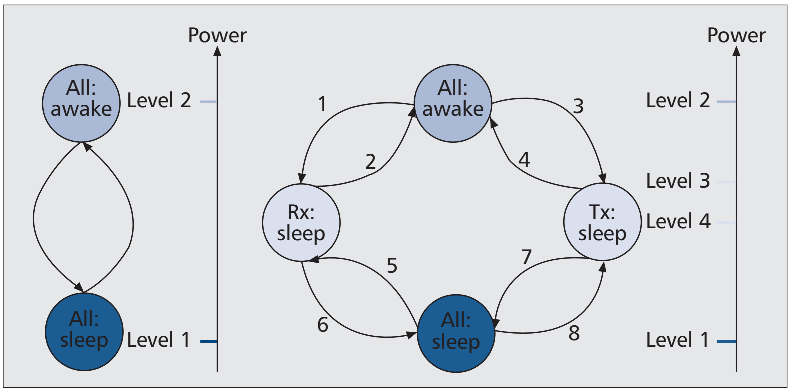

Figure 2. *Multi-power-level ONUs.*

```
a: If the Tx has not transmitted traffic for the time duration of idle_threshold
      s = 1;
      b: Tx enters into sleep status;
      sleep_time = $2^{s-1}$*short_active + ($2^{s-1}$ – 1) * idle_threshold;
      If sleep_time > 50 ms
         sleep_time = 50 ms
      Endif
      Tx wakes up after sleep time duration
      The ONU checks the queue length and reports the queue status
      If there is queued traffic
         Keep Tx awake
         s = 0;
         go to line a;
      Else
         s = s + 1;
         go to line b;
         Endif
   Endif
```

Algorithm 1. *Decide the transition between* all:awake *and* Tx:sleep.

ing without degrading services. We present solutions in determining the transitions under two respective scenarios: sleep for more than one DBA cycle, and sleep within one DBA cycle.

SCENARIO 1: SLEEP FOR MORE THAN ONE DBA CYCLE

In this scenario the transition is decided by the incoming traffic status. Tx/Rx is put to sleep if no upstream/downstream traffic exists for some time.

Whether or not downstream/upstream traffic exists can be inferred based on the information of the time allocated to ONUs and queue lengths reported from ONUs, which is known to both OLT and ONUs. If no upstream traffic arrives at an ONU, the ONU requests zero bandwidth in the MPCP REPORT message. Then, the OLT can assume that this ONU does not have upstream traffic. If no downstream traffic for an ONU arrives at the OLT, the OLT will not allocate downstream bandwidth to the ONU. Assume that, out of fairness concerns, an OLT allocates some time slots in a DBA cycle to every ONU with downstream traffic. Then, considering the uncertainty of the exact time allocated to an ONU in a DBA cycle, the ONU can infer that no downstream traffic exists if it does not receive any downstream traffic within two DBA cycles.

The next question is to decide the transition between different statuses. In this scenario, the status *Rx:sleep* actually does not exist since there will still be some downstream MPCP control packets for an ONU to facilitate the upstream transmission even when no downstream data traffic exists. Hence, we next discuss the transition between *all:awake* and *Tx:sleep*, and the transition between *Tx:sleep* and *all:sleep*.

Formerly, Kudo *et al*. [10] proposed periodic wakeup with sleep time adaptive to the arrival traffic status. We also decide the sleeping time based on traffic status. More specifically, we set the sleep time as the time duration in which traffic stops arriving. When putting Tx to sleep, for example, Algorithm 1 describes the transition between *all:awake* and *Tx:sleep*. We assume that Algorithm 1 is known to the OLT as well. Then, the OLT can accurately infer the time that Tx is asleep or awake.

Let `idle_threshold` be the maximum time duration a transmitter stays idle before being put to sleep, `short_active` be the time taken for an ONU to check its queue status and send out the report, and `sleep_time` be the time duration each time an ONU sleeps. If the transmitter is idle for `idle_threshold`, Tx will be put into sleep status, and the `sleep_time` for the first sleep equals `idle_threshold`. Then, Tx wakes up to check its queue status and sends a report to the OLT, which takes `short_active` time duration. If there is no upstream traffic being queued, Tx will enter sleep status again. Until now, the elapsed time since the last time Tx transmitted data packets equals `idle_threshold` + time duration of the first sleep + `short_active`. So for the second sleep, the sleep time duration `sleep_time` is set as `idle_threshold` + time duration of the first sleep + `short_active`. According to MPCP, ONUs send MPCP REPORT messages to the OLT every 50 ms when there is no traffic. Thus, we set the upper bound of `sleep_time` as 50 ms to be compatible with MPCP and also to avoid introducing too much delay of the traffic which arrives during sleep mode. This process repeats until upstream traffic arrives. For the *s*th sleep, the `sleep_time` equals `idle_threshold` + the total time durations of the former *s* – 1 sleep + (*s* – 1)* `short_active`, which also equals 2 s – 1 * `short_active` + (2^{s-1} – 1) * `idle_threshold`.

For the transition between *Tx:sleep* and *all:sleep*, the transitioning algorithm is similar to Algorithm 1 with the exception that line **a** should be changed to: "If the Rx has not received downstream traffic destined to the ONU for the time duration of `idle_threshold`." The remaining codes are similar.

Figure 3 shows an example of the sleep time control process with `short_active` = 2.5 ms and `idle_threshold` = 10 ms. Then, `sleep_times` of the first, second, third, and fourth sleeps are as follows:

- First sleep: 10 ms
- Second sleep: `idle_threshold` + 10 ms + `short_active` = 22.5 ms
- Third sleep: `idle_threshold` + 32.5 ms + 2 * `short_active` = 47.5 ms
- Fourth sleep: min{50 ms, `idle_thresh-`

old + 80 ms + 3 * `short_active`} = 50 ms

In deciding the sleep time, `idle_threshold` and `short_active` are two key parameters which are set as follows.

idle_threshold — Setting `idle_threshold` needs to consider the time taken to transit between *sleep* and *awake*. Considering the transition time, the net sleep time will be reduced by the sum of the transit time from awake to sleep and the transit time from sleep to awake. Hence, `idle_threshold` should be set longer than the sum of two transit times in order to save energy in the first sleep. Currently, the time taken to power up the whole ONU is around 2–5 ms [8]. Hence, `idle_threshold` should be greater than 4 ms in this case. In addition, we assert that the upstream/downstream traffic queue is empty if no bandwidth is allocated to upstream/downstream traffic for `idle_threshold`. To ensure this assertion is correct, `idle_threshold` should be at least one DBA cycle duration, which typically extends less than 3 ms to guarantee delay performance for some delay-sensitive service. So Tx/Rx must sleep for over one DBA cycle with this scheme.

short_active — During the short awake time of Tx, an ONU checks its upstream queue status and reports to the OLT. Thus, `short_active` should be long enough for an ONU to complete these tasks. In addition, using some upstream bandwidth for an ONU to send a report affects the upstream traffic transmission of other ONUs. In order to avoid interruption of the traffic transmission of other ONUs, we set `short_active` to be at least one DBA cycle duration such that an OLT can have freedom in deciding the allocated time for an ONU to send its report. For Rx, during the short awake time, the OLT begins sending the queued downstream traffic if there is any. Similar to the Tx case, `short_active` is set to be at least one DBA cycle to avoid interrupting services of ONUs in the Rx case.

Scenario 2: Sleep within One DBA Cycle

In the former scenario, the sleep and awake durations of Tx and Rx are greater than one DBA cycle. In this section, we discuss the scheme of putting Tx and Rx to sleep within one DBA cycle.

Consider a PON with 16 ONUs. During a DBA cycle, on average, only 1/16 of time duration is allocated to an ONU. This means that even if upstream/downstream traffic exists, Tx/Rx need only be awake for 1/16 of the time and can go to sleep for the other 15/16 of the time. Therefore, significant energy savings can be achieved.

To enable an ONU to sleep and wake up within a DBA cycle, the transit time between awake and sleep should be less than half the DBA cycle duration such that the net sleep time can be greater than zero, and thus energy can be saved. Formerly, Wong *et al*. [15] reduced the transition time to as small as 1–10 ns by keeping some of the back-end circuits awake. Thus, with advances in speeding up transition time, it is physically possible to put an ONU to sleep within one cycle to save energy.

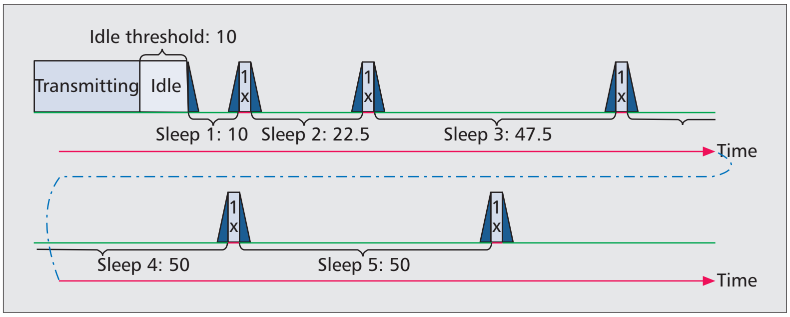

Figure 3. *An example of sleep time control of the transmitter.*

For the upstream case, waking up Tx can be triggered by ONU MAC when the allocated time comes. Tx can go to sleep after data transmission. For the downstream case, however, it is difficult to achieve since Rx does not know the time when the downstream traffic is sent and has to check every downstream packet. To address this problem, we propose the following sleep-aware downstream scheduling scheme.

For downstream transmission, an OLT schedules the downstream traffic of ONUs one by one, and the interval between two transmissions of an ONU is determined by the sum of the downstream traffic of all other ONUs. Again, due to the bursty nature of ONU traffic, the ONU traffic in the next cycle does not vary much from that in the current cycle. Accordingly, we can make an estimation of traffic of other ONUs and put this particular ONU to sleep for some time.

More specifically, for a given ONU, denote Δ as the difference between the ending time of its last scheduling and the beginning time of its current scheduling. Then we set the rule that the OLT will not schedule this ONU's traffic until $f(\Delta)$ time after the ending time of the current scheduling. As long as the ONU is aware of this rule, it can go to sleep for $f(\Delta)$ time durations.

Figure 4 illustrates one example of putting an ONU to sleep within one DBA cycle. In this example, one OLT is connected to four ONUs, and $f(\Delta)$ is set as $0.8 * \Delta$. The interval between the first two schedulings of ONU 4 is 9. Hence, the OLT will not schedule the traffic of ONU 4 until 7.2 time duration later; ONU 4 can sleep for 7.2 time duration and then wake up. However, this wakeup is an early wakeup since the actual transmission of the other ONUs takes 9.5 time durations, which is 2.3 longer than the estimation. Similarly, the duration of the second sleep is set as 7.6. However, this wakeup is a late wakeup since the actual time taken to transmit the other ONUs' traffic is 6.5. The late wakeup incurs 1.1 idle time duration on the downstream channel.

As can be seen from the example, early wakeup and late wakeup are two common phenomena of this scheme. Early wakeup implies that energy can be further saved, while late wakeup results in idle time durations, and thus possibly service degradation. From the network service providers' perspective, avoiding late wakeup and the subsequent service degradation is more

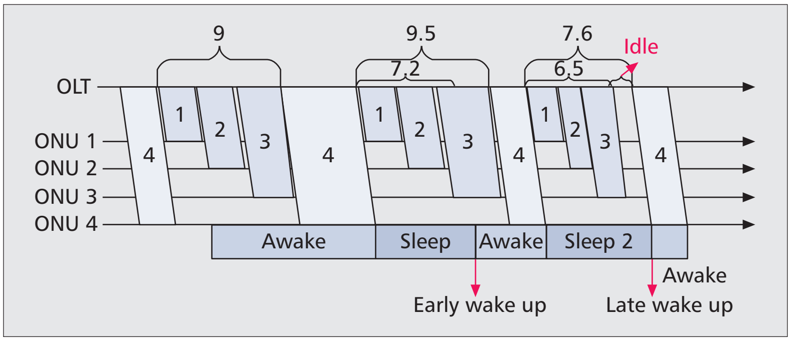

Figure 4. *One example of putting ONUs into sleep within one DBA cycle.*

desirable than avoiding early wakeup. Hence, $0 < f(\Delta) < \Delta$ is suggested to be set. If $f(\Delta)$ is set as small as 0.5Δ, on average an ONU can still sleep 15/32 of the time when a PON supports 16 ONUs and traffic of ONUs is uniformly distributed. Therefore, significant power savings can be achieved with this scheme.

CONCLUSION

With the increase of line rate, energy consumption of 10G-EPON may be significantly increased as compared to that of 1G-EPON. It is important to reduce power consumption of 10G-EPON for low OPEX and environmental friendliness. In this article, with the focus on saving energy of ONUs, we have proposed to disable some functions of ONUs whenever possible, and introduced four power levels for ONUs. Then, from the perspective of MAC control and scheduling, we have investigated schemes to properly transit between power levels of ONUs. Two schemes have been proposed for two respective scenarios: sleep for over one DBA cycle, and sleep within one DBA cycle. For the former scenario, we have described the scheme of implicitly conveying sleep information between ONU and OLT, and proposed an algorithm to adapt the sleeping time to user traffic. For the latter scenario, to address the challenging issue in the downstream case, we have designed a sleep-aware downstream scheduling scheme to achieve energy saving without degrading user service.

REFERENCES

[1] Y. Luo and N. Ansari, "Bandwidth Allocation for Multiservice Access on EPONs," *IEEE Commun. Mag.*, vol. 43, no. 2, 2005, pp. S16–S21.
[2] K. Cho *et al.*, "The Impact and Implications of the Growth in Residential User-to-User Traffic," *Proc. Conf. Apps., Tech., Architectures, Protocols for Comp. Commun.*, 2006.
[3] R. Kubo *et al.*, "Sleep and Adaptive Link Rate Control for Power Saving in 10G-EPON Systems," *Proc. 2009 IEEE GLOBECOM*, 2009, pp. 1573–78.
[4] J. Baliga *et al.*, "Energy Consumption in Optical IP Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 27, no. 13, 2009, pp. 2391–2403.
[5] J. Baliga *et al.*, "Energy Consumption in Access Networks," *Proc. Optical Fiber Commun. Conf.*, 2008.
[6] C. Lange and A. Gladisch, "On the Energy Consumption of FTTH Access Networks," *Proc. Optical Fiber Commun. Conf.*, 2009
[7] K. Tanaka, A. Agata, and Y. Horiuchi, "IEEE 802.3av 10G-EPON Standardization and Its Research and Development Status," *IEEE/OSA J. Lightwave Tech.*, vol. 28, no. 4, Feb.15, 2010, pp. 651–61.
[8] S. Wong *et al.*, "Sleep Mode for Energy Saving PONs: Advantages and Drawbacks," *IEEE GreenCom*, 2009.
[9] J. Mandin, "EPON Power Saving via Sleep Mode," IEEE P802. 3av 10GEPON Task Force Meeting, 2008.
[10] R. Kubo *et al.*, "Adaptive Power Saving Mechanism for 10 Gigabit Class PON Systems," *IEICE Trans. Commun.*, vol. E93-B, no.2, 2010, pp. 280–88.
[11] ITU-T Rec. G. sup. 45.
[12] G. Kramer, B. Mukherjee, and G. Pesavento, "Ethernet PON (ePON): Design and Analysis of an Optical Access Network," *Photonic Net. Commun.*, vol. 3, no. 3, 2001, pp. 307–19.
[13] S. Lee and A. Chen, "Design and Analysis of a Novel Energy Efficient Ethernet Passive Optical Network," *Proc. 9th Int'l. Conf. Net.*, 2010.
[14] E. Trojer and P. Eriksson, "Power Saving Modes for GPON and VDSL," *Proc. 13th Euro. Conf. Netw. & Optical Commun.*, Austria, June 30–July 3, 2008.
[15] S. Wong *et al.*, "Demonstration of Energy Conserving TDM-PON with Sleep Mode ONU Using Fast Clock Recovery Circuit," *Proc. Optical Fiber Commun. Conf.*, 2010.

BIOGRAPHIES

JINGJING ZHANG [S'09] received a B.E. degree from Xi'an Institute of Posts and Telecommunications, China, in 2003, and an M.E. degree from Shanghai Jiao Tong University, China, in 2006, both in electrical engineering. She is working toward her Ph.D. degree in electrical engineering at the New Jersey Institute of Technology (NJIT), Newark. Her research interests include planning, capacity analysis, and resource allocation of broadband access networks, QoE provisioning in next-generation networks, and energy-efficient networking. She received the New Jersey Inventors Hall of Fame Graduate Student Award in 2010.

NIRWAN ANSARI [S'78, M'83, SM'94, F'09] (Nirwan.Ansari@NJIT.EDU) received his B.S.E.E. (summa cum laude, with a perfect GPA) from the New Jersey Institute of Technology (NJIT) in 1982, his M.S.E.E. degree from the University of Michigan, Ann Arbor, in 1983, and his Ph.D. degree from Purdue University, West Lafayette, Indiana, in 1988. He joined NJIT's Department of Electrical and Computer Engineering as an assistant professor in 1988, tenured associate professor in 1993, and full professor since 1997. He has also assumed various administrative positions at NJIT. He authored *Computational Intelligence for Optimization* (Springer, 1997, translated into Chinese in 2000) with E.S.H. Hou, and edited *Neural Networks in Telecommunications* (Springer, 1994) with B. Yuhas. His current research focuses on various aspects of broadband networks and multimedia communications. He has also contributed over 350 technical papers, over one third of which were published in widely cited refereed journals/magazines. For example, one of his published works was the sixth most cited article published in *IEEE Transactions on Parallel and Distributed Systems*, according to the journal EIC report in February 2010. He has also guest edited a number of special issues, covering various emerging topics in communications and networking. He was visiting (chair) professor at several universities. He has served on the Editorial and Advisory Boards of eight journals, including as a Senior Technical Editor of *IEEE Communications Magazine* (2006–2009). He has served the IEEE in various capacities such as Chair of the IEEE North Jersey Communications Society Chapter, Chair of the IEEE North Jersey Section, member of the IEEE Region 1 Board of Governors, Chair of the IEEE Communications Society Networking Technical Committee Cluster, Chair of the IEEE Communications Society Technical Committee on Ad Hoc and Sensor Networks, and Chair/Technical Program Committee Chair of several conferences/symposia. Some of his recent recognitions include a 2007 IEEE Leadership Award from the Central Jersey/Princeton Section, the NJIT Excellence in Teaching Award in Outstanding Professional Development in 2008, a 2008 IEEE MGA Leadership Award, the 2009 NCE Excellence in Teaching Award, several best paper awards (ICNIDC 2009 and IEEE GLOBECOM 2010), a 2010 Thomas Alva Edison Patent Award, and designation as an IEEE Communications Society Distinguished Lecturer (2006–2009, two terms).