OThB4.pdf

An SLA-Based Energy-Efficient Scheduling Scheme for EPON with Sleep-Mode ONU

Lei Shi [I], Sang-Soo Lee [II], and Biswanath Mukherjee [I]

I. Networks Lab, Department of Computer Science, University of California, Davis, CA 95616

II. Electronics and Telecommunications Research Institute, Daejeon, Korea

E-mail: {leishi, bmukherjee}@ucdavis.edu; soolee@etri.re.kr

Abstract: We propose an SLA-based scheduling scheme for EPON in which OLT can adjust sleep time and ONU can quit sleep mode for sending expedited frames. Considerable energy can be saved under practical power consumption settings.

OCIS: (060.4250) Networks

1. Introduction

As the Information and Communication Technologies (ICT) have experienced unprecedented growth for almost thirty years, only recently a previously-neglected aspect of ICT, viz. energy cost, has been recognized as a serious issue and caught the world's attention. This has led to intensive energy conservation research in many aspects of ICT. In the access network area, IEEE P802.3az Energy Efficient Ethernet (EEE) Task Force is working on improving energy efficiency of Ethernet, the dominant technology in access, and has proposed two different methods—shutting down links and slowing down link rate—with the former being more applicable to Ethernet on fiber physical layer devices (PHYs) as transitioning to lower-speed PHYs is generally not feasible for optical components.

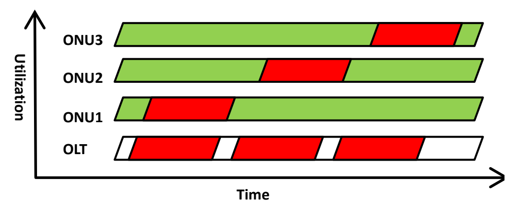

Fig. 1. In EPON, an ONU can switch to sleep mode and thus save considerable energy in empty slots.

Enabling the optical network units (ONU) in Ethernet passive optical network (EPON) to enter sleep mode for energy saving [1] is based on the fact that the average utilization of Ethernet linked to end users is only about 1% to 5% [2]. Even if the network is heavily used, the broadcast nature of the EPON downstream, together with the fact that the usage behavior follows the Pareto principle where 80% of the traffic is generated by 20% heavy users [3], still left a lot of empty slots for a lightly-used ONU to switch to sleep mode, as shown in Figure 1. However, unlike Ethernet with copper PHY, a predetermined sleep time must be specified by either the optical line terminal (OLT) or the ONU as the ONU is unable to sense the channel in sleep mode.

How to implement the ONU sleep mode in an efficient but practical way is still an open question. Ref. [4] proposed downstream scheduling schemes that are tightly coupled with upstream bandwidth allocation, so that an ONU could download and upload simutaneously during its active period. Ref. [5] proposed a scheme where the system enters power-saving mode, and each ONU wakes up periodically within its designated slot when the load is lower than a pre-determined threshold. Other research works that demonstrate their energy saving capability usually use static bandwidth allocation on a fixed time-division multiplexing (TDM) setting when conducting simulations.

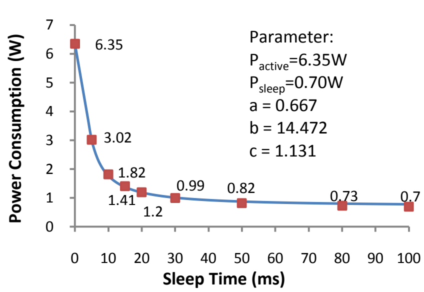

Fig. 2. Sleep mode power consumption of an ONU, with respect to sleep time.

2. SLA-Based Energy-Efficient Scheduling Scheme

All the proposals so far prefer frequent mode switching with periods of a few miliseconds. This would harm the efficiency as energy is wasted in frequent waking up processes which introduce extra time overhead estimated at around 2ms with aggressive assumptions [1]. Besides, these studies are optimistic in predicting the energy saving rate. According to Figure 2 which shows measured data [1], it takes around 20ms for the ONU to reduce its power by 80%, thus the ONU could never reach the claimed 90% energy saving rate with frequent mode switching. Comparing these values to 12 s, the time for sending one 1500-byte Ethernet frame in 1G EPON, it is quite inefficient to transmit just a few frames during an active period, as this will result in a frame efficiency of less than 1%.

As problems arise from frequent mode switching, long sleep time could lead to bad user experience, especially for delay-sensitive applications. To solve the problem, we propose an SLA-based energy-efficient scheduling scheme which adjusts the sleep time according to experiened traffic on both upstream and downstream. Its unique arrangement allows an ONU to quit sleep mode before the predetermined sleep time is up in order to send expedited frames with strict service level agreement (SLA) and an extra guard time is introduced at the end of the active period to keep the ONU from frequent mode switching. Our proposal is within the IEEE 802.3ah control scheme and requires minor changes in the control signal flow. The scheme follows the procedure in [1] for initiating and quiting sleep mode.

Our scheduling scheme categorizes Ethernet frames into three classes with different SLAs. The high-priority class is expedited forwarding (EF), which is delay-sensitive and requires bandwidthguarantees. The medium-priority class is assured forwarding (AF), which is not delay-sensitive but requires bandwidth guarantees. Finally, the low-priority class is best effort (BE), which is neither delay-sensitive nor bandwidth guaranteed.

When an ONU quits sleep and enters active mode, all frames in the buffers of the OLT and ONU will be sent under normal protocols regardless of its class for the purpose of maximizing frame efficiency. Thus, in the worst scenario, the sleep mode will cause a BE frame to experience an extra delay of a whole sleep time. EF frames on the upstream direction will be guaranteed a much smaller delay using the methods introduced below. Preemptive queues are used in both the OLT and ONU to ensure that higher class frames are always sent first. Credit-based bandwidth allocation [6] is used by the OLT to take into consideration frames that arrived after REPORT message is sent.

After the end of the transmission on both directions and before the OLT and ONU agree on entering sleep mode for the next time, they wait for some extra guard time to prevent frequent mode switching. This time is set to twice the largest inter-frame time of two EF frames arrived in the past one minute. If no EF frames are received at both ends during this period, the OLT will initiate sleep mode for the ONU.

25ms standard sleep time is used as a tradeoff between energy saving rate and user experience. If no frames are sent in either direction of the last active period, the sleep time is added by 25ms with a probability of 50% until it reaches maximum sleep time, which is set to 250ms. Sleep time is reset to 25ms as long as the OLT sends or receives frames during the last active period.

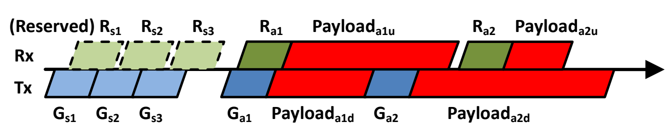

Fig. 3. OLT sends GATE to sleeping ONUs, reserving slots for possible REPORT.

On the upstream direction, upon receiving an EF frame from end users, the ONU will wake itself up immediately, regardless of whether the predetermined sleep time is over. In order for the unexpected active ONU to send the EF frames, OLT will send a special GATE message G_s periodically to all ONUs in sleep mode, reserving each ONU a time slot on the upstream channel for sending the REPORT message R_s, as shown in Figure 3. The frequency of sending these GATE messages are related to the SLA, and for protocol simplicity we could just send them at the beginning of every polling cycle. The upstream reserved slots are not expected to introduce a large overhead as the time for transmiting one REPORT message plus the guard band needed is below 2 s, which only counts for 0.2% of the typical 1ms cycle time.

As shown in Figure 4, after quiting the sleep mode, it takes the ONU 2ms for time recovery and syncronization. At most in one cycle time (1ms), the ONU will receive the GATE message and start sending the REPORT message using the reserved slot. After another cycle (1ms), the ONU will receive a GATE message that allocates upstream bandwidth and actual data transmission will usually start in less than one cycle since EF frames are of high priority.

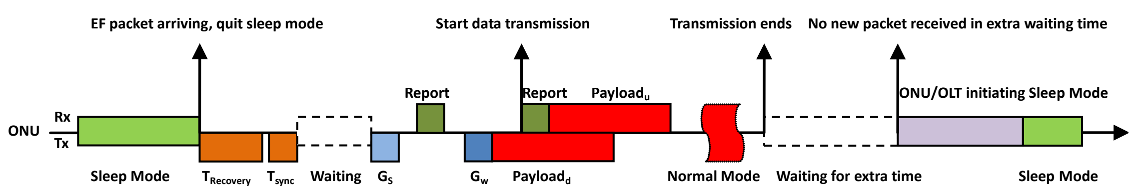

Fig. 4. Timeline for the ONU to end sleep mode before the predetermined sleep time is over.

OThB4.pdf

Based on the above delay analysis, we expect no more than 5ms extra delay for the first incoming EF frame. Considering an 1Mbps EF stream, a small 1000-bit frame size will only have 5 frames in 5ms. The extra delay caused by sleep mode should only affect the first few EF frames reception. The delay for these frames could be further reduced if the OLT finds the upstream utilization low and thus is able to allocate a larger reserved slot in which not only the REPORT message, but also several data frames could be sent. The scheme could be applied in other situations also. For example, when AF frames fill up the buffer and thus a waking up is inevitable for preventing the violation of SLA.

3. Simulation Results

The energy saving rate and delay characteristic of an EPON with 32 ONUs are evaluated. Among the 32 ONUs, 5 are heavy users and others are light users. Ethernet frames are generated using a Poisson distribution [7] and the frame size distribution is set according to [8].

To compute the actual energy cost, the dynamic sleep mode power consumption of an ONU is approximated by function (1), where and are the static power consumptions of the two modes. With proper parameter values (shown in Fig. 2), we found the function fits the known data very well, as depicted by the blue line in Fig. 2. Using indefinite integral, we can then get the approximate power consumption in sleep mode, as shown in function (2). Figures 5 compares them with the theoretical power/energy saving rate ( ).

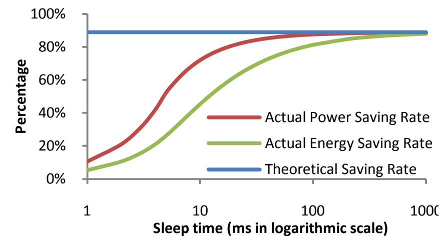

Fig. 5. Actual/theoretical energy/power saving rate comparison.

(1) (2)

Figure 6 shows ONU energy saving rate with different downstream transmission rates. Under extremely light use (5kbps), the ONU could save 84% of the energy. Under moderate use (1Mbps), the energy saving rate is still 60%. As the sleep time is almost always set to 25ms and the transmission time is small under moderate use, the energy saving rate is flat until the transmission time becomes significant which causes a decrease in energy saving rate when the ONU is heavily loaded.

The queueing delay for the upstream EF frames is 1.5ms on average. For BE frames, the average queueing delay is about 14ms as long as EPON is under its throughput limits.

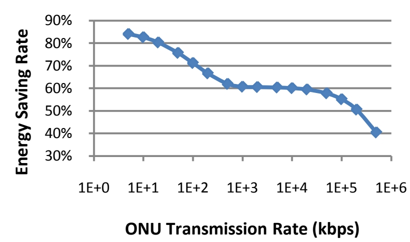

Fig. 6. ONU energy saving rate under different downstream transmission rates.

4. Conclusion

Our proposed SLA-based energy-efficient scheduling scheme solves the problem of frequent mode switching. It adjusts the sleep time according to traffic, waits for extra guard time before entering new sleep period, and allows ONUs to quit sleep mode for sending expedited frames with strict SLA. Simulation results show that the ONU could save 60% of the energy under moderate use and practical power consumption settings.

References

[1] J. Mandin, "EPON powersaving via sleep mode," IEEE P802.3av 10GEPON Task Force Meeting, Sept. 2008.

[2] K. J. Christensen, C. Gunaratne, B. Nordman, and A. D. George, "The next frontier for communications networks: power management," *Computer Comm.*, vol. 27, pp. 1758–1770, 2004.

[3] R. N. Clarke, "Costs of Neutral/Unmanaged IP Networks," *Review of Network Economics*, vol. 8, no. 1, 2009.

[4] Y. Yan, S.-W. Wong, L. Valcarenghi, S.-H. Yen, D. R. Campelo, S. Yamashita, L. Kazovsky and L. Dittmann, "Energy Management Mechanism for Ethernet Passive Optical Networks (EPONs)," *IEEE ICC 2010*, May 2010.

[5] S. S. W. Lee and A. Chen, "Design and Analysis of a Novel Energy Efficient Ethernet Passive Optical Network," *ICN 2010*, Apr. 2010.

[6] G. Kramer, B. Mukherjee, and G. Pesavento, "IPACT: a dynamic protocol for an Ethernet PON (EPON)," *IEEE Comm. Mag.*, vol. 40, no. 2, pp. 74-80, Feb. 2002

[7] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido, "A nonstationary Poisson view of Internet traffic," *IEEE INFOCOM'04*, Mar. 2004

[8] D. Sala and A. Gummalla, "PON Functional Requirements: Services and Performance," IEEE 802.3ah meeting in Portland, July 2001.