Next Generation Ethernet Access Networks: GPON vs. EPON

T. ORPHANOUDAKIS[1], H.-C. LELIGOU[2], J. D. ANGELOPOULOS[2]
1: University of Peloponnese, Karaiskaki str., 22100, Tripolis, Greece
2: National Technical University of Athens, Heroon Poloytechneiou 9, Zographou Campus, Athens, Greece

Abstract: Recently both ITU and IEEE have standardized solutions for Passive Optical Networks operating at gigabit per second line rates and optimized for the transport of packet-based traffic to improve the efficiency of previously standardized broadband PONs, which were based on ATM. The efficiency and performance of PON systems, depends mainly on the implemented medium access protocol. The latter is not part of the standards and left to the implementer, however the standards describe a set of control fields that constitute the tool-set for the MAC operation. In this paper we compare the efficiency and performance of the two systems under as close as possible MAC.

Key-Words: - Passive Optical Networks, Broadband Access Networks, MAC protocol, Last mile

1. Introduction

In response to the steadily increasing demand for bandwidth and networking services for residential users as well as enterprise customers, Passive Optical Networks (PONs) have emerged as a promising access technology that offers flexibility, broad area coverage, and cost-effective sharing of the expensive optical links compared to the conventional point-to-point transport solutions. In addition, they inherently concentrate traffic and greatly reduce the number of input ports in the access multiplexer, both important for the cost-sensitive residential access market. Due to these advantages, PONs have generated during the last decade substantial commercial activity also reflected in the work of several standardization bodies. Since the initial standardization of ATM-based PONs (APONs or alternatively named in ITU-T G.983.1 standard [1] Broadband PONs - BPONs) newer standards support multi-gigabit rates and better adapt to the packet-based Internet applications. In January 2003, the GPON (Gigabit PON) standards were ratified by ITU-T and were included in the G.984.x series of ITU-T Recommendations ([2], [3]). Driven by a closed group of worldwide system vendors and national telecom operators, they are designed to support a mix of TDM, ATM and packet based services, reaching symmetrical transmission rates of up to 1.244Gbps or 2.488Gbps

At the same time IEEE, through the activities of Ethernet in the First Mile (EFM) 802.3ah group, has standardized a Gigabit Ethernet-friendly technology ([4]) called Ethernet PON (EPON), with the objective to leverage the great success of Ethernet as a LAN technology and exploit the economies of scale that the dominance of Ethernet has generated.

Although PONs can achieve economical deployment and operation, which are a major concern to operators and service providers, high and fair resource utilization is equally important. Due to the multiple access nature of PONs in the upstream direction, the performance of a PON in terms of delay, delay variation and throughput strongly depends on the upstream bandwidth allocation function of the Medium Access Controller (MAC) residing at the Optical Line Termination (OLT).

In the following section we discuss the details of the TDMA based operation of the MAC protocol in each case, identify the critical parameters of each technology that affects performance and derive relevant measures used in our simulations for comparison and collection of quantitative performance results that we present in the last section.

2. TDMA based upstream burst allocation

In the case of EPON, DBA is supported by means of the REPORT messages of MPCP, while in GPON by the fields of the DBRu headers, which in both cases announce the requests for additional bandwidth in order to serve the packets waiting in the ONU queues. Since both standards support transmission of variable length packets the allocations are expressed in number of bytes. The temporal data transmission bandwidth allocated to each ONU by the OLT is inferred from the number of allocated bytes (indicated in the GATE messages of MPCP and the Upstream Bandwidth Map of the GPON downstream frames) and the time elapsed between successive allocation intervals, determined by the scheduling algorithm implemented in the OLT. The scheduling of grants by the OLT (not defined in the standards since it is considered vendor specific) results in the proportional distribution of upstream bandwidth among ONUs as well as among different queues of the same ONU taking into account specific parameters that actually comprise a Service Level Agreement (SLA). The OLT scheduling also results in the spacing of grants and corresponding byte allocations in time and should be performed in a cyclic manner among all active ONUs forming a kind of scheduling windows. Scheduling windows can be periodic with a certain scheduling period and during each period a number of accumulated requests from previous reporting periods are served while new requests are collected.

In GPON, each queue is associated with a T-CONT type and is typically allocated by the MAC at activation time (on the basis of SLA information) two parameters: SDI (Successive Data Interval, i.e the time distance between two successive grants) and TB (Transmit Bytes, i.e how many bytes are to be sent with each allocation). Upper and lower bounds of these parameters are defined in the service agreement and are related to the specified peak and average rates as well as the allowed variation. This provides the tool to specify a guaranteed part (based on minTB, MaxSDI) allowing the surplus bandwidth to be assigned dynamically up to the peak rate (defined by MaxTB, MinSDI).

In IEEE 802.3ah there is no specific reference to traffic profile specifications. The specified MPCP protocol procedures though favor the periodic scheduling of grants calculating allocations for all ONUs in cumulative way within a scheduling period D_m. This is common among most proposed EPON MAC algorithms proposed in the literature ([5]-[11]), although initial proposals allowed for variable D_m, which is not recommended when service guarantees must be enforced.

As a basis for our comparison we take the EPON compliant MAC protocol proposed by the authors in [11], which extends the approach investigated in [10] to collectively handle four allocation strategies with enhanced bandwidth efficiency and optimized scheduling of granted upstream transmission windows. The MAC protocol described in [11] is totally aligned with the GPON MAC protocol described in [3] with respect to the service class definition and only differs in the scheduling of upstream allocations in time due to the different MAC mechanisms implemented in each standard. The different scheduling implementation in these two cases as discussed in [11] was actually selected due to two fundamental differences between the two technologies. The first reason is the time consumed for the transmission of the preamble, the delimiter plus the required guard time (denoted as T_{pre} in [11]), which introduces larger overhead in EPONs when interleaved burst transmissions are attempted by multiple ONUs. The second reason is the segmentation functionality provided by the GEM encapsulation adopted in GPON, which allows for complete utilization of each upstream allocation independently of its length, whereas in EPONs small upstream slot allocations increase the probability for waste of bandwidth.

The overhead introduced following different scheduling alternatives is graphically explained in Figure 1 where upstream burst allocations and actual data transmissions (covering a time window D_m) from different ONUs and CoS queues are shown. In Figure 1a, a possible scheduling of grants accommodating the requests collected in earlier polling cycles is shown. The order of the allocations (which target individual queues) affects the achieved efficiency. In Figure 1a, higher priority allocations precede lower priority allocations. Taking into account that the duration of the CoS 1 allocations is fixed, the position of the two higher priority allocations in every scheduling cycle D_m are fixed, forming sub-frames per CoS, (a scheduling discipline also selected in [9] and [10]). In EPONs, part of the allocations may be wasted since the exact ONU queue occupancy and packet delineation is not known by the grant scheduler at the OLT. It is then very likely that the left-over time at the end of the allocated window does not match the length of the next packet in the FIFO (First In First Out) queue. This phenomenon called Unused Slot Remainder (USR), is also shown in Figure 1 only for the 2nd ONU as an example. Figure 1b shows the alternative of serving all CoS queues from the same ONU before allocating slots to other ONUs creating sub-frames per ONU.

Obviously, this schedule introduces less physical layer overheads and the efficiency improvement depends on the scheduling period D_m as well as on the number of supported queues and ONUs. The longer the D_m, the higher is the efficiency achieved (minimizing physical layer overheads). However, assuming that this will also be the service period for all services (including delay sensitive ones), the scheduling period also directly affects (and is close to) the delay observed by CBR-like services. To achieve a delay guarantee of 1.5ms for voice services ([4], [5]), a scheduling period of about equal duration should be selected.

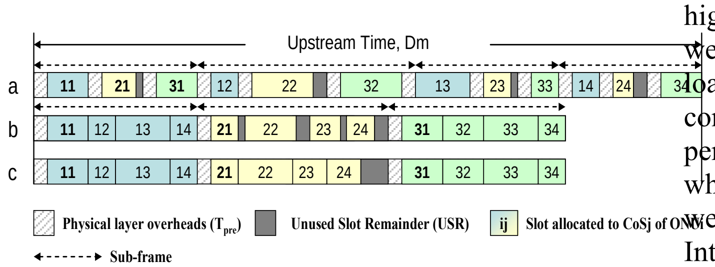

Figure 1: Example of OLT burst allocation leading to the Unused Slot Remainder effect

The GPON protocol described in [3] allows scheduling of upstream transmissions per ONU queue (Alloc-ID) at arbitrarily small scheduling periods as needed in order to respect their maxSDI parameter. It is worth noting that GPON specifications support even TDM services, explaining the choice of 125μs frames. To allow a fair comparison of the GPON MAC presented in [3] with the EPON MAC presented in [11] the D_m scheduling period will be chosen small enough to support similar source granularity as the GPON. These conditions provide the grounds for a meaningful comparative performance evaluation which is given in the following section.

3. Performance evaluation

The main motivation for this work has been to compare the efficiency of EPONs and GPONs as multi-service broadband access systems when operating under the full extent of their dynamic grant scheduling capabilities available in the EPON MPCP protocol and the dba mechanisms of the GPON MAC. A simulation model was developed using the OPNET simulator including 16 ONUs, each equipped with four CoS queues for both systems.

The offered load is shared uniformly among all ONUs. The traffic mix included on average 10% high priority traffic, while 2nd, 3rd and 4th priority were injecting at 15%, 20% and 55% of the total load respectively. High priority sources were of constant bit rate generating short fixed-size packets periodically (a model suitable for voice traffic) while the sources for the other 3 types of traffic were of the ON-OFF type (modeling self-similar Internet traffic), with different burstiness factors, namely, 2, 5 and 5 for 2nd, 3rd and 4th priority respectively. The packet size followed the tri-modal distribution characterizing traffic generated from IP-based applications (packet sizes of 64, 500, 1500 bytes appear with probability 0.6, 0.2, and 0.2 respectively according to [12]).

The guardband and physical layer overhead (i.e. T_{pre}) were assumed equivalent to 1μs for EPON, while for GPON this is an order of magnitude lower (GPON operates with only 15Bytes overhead for burst ONU transmission). As discussed in the previous section a final important parameter is the scheduling period D_m. The performance of MAC was measured while varying the scheduling period D_m from sub-millisecond values (750μsec) up to 2msec, (1msec and 2 msec are the values most frequently assumed in the literature).

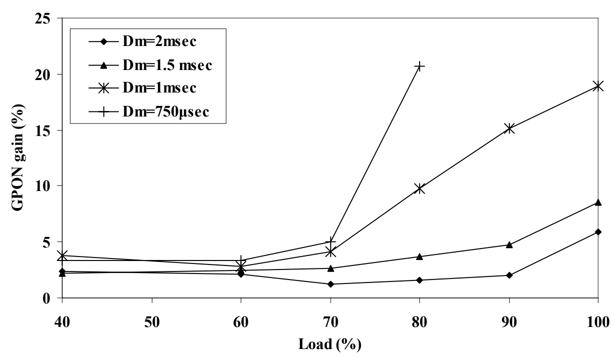

Figure 2: GPON gain vs. load as a function of D_m

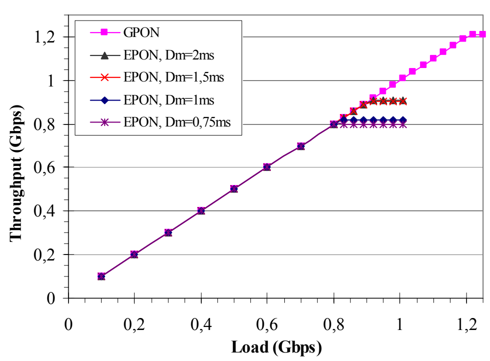

Figure 3: GPON vs. EPON bandwidth efficiency as a function of D_m

In the performance evaluation we focus on bandwidth efficiency as the other performance parameters, (delay and delay variation), are not very different once a similar D_m value is used. In this respect, GPON allows better utilisation of the upstream link.. Apart from the inefficiency introduced by the 8/10bit coding adopted in EPONs which limits the available upstream rate to 1Gb/s instead of 1.24Gb/s in GPON, their efficiency when transporting Ethernet traffic is also better in terms of allocated slot utilization due to their GEM which combats the Unused Slot Remainder effect as explained earlier. Although GEM encapsulation introduces a header of 5 bytes on each upstream frame or segment of frame, it allows segmentation and reassembly of Ethernet frames achieving a nearly perfect fit. Hence, while in GPON the inefficiency stems from the GEM overhead, in EPON inefficiency is introduced by the USR, the longer physical layer overhead and the longer reporting message (3bytes per queue in GPON vs. 64bytes for all 8 queues in EPON).

The gain of GPON is quantified in Figure 2, which shows the extra payload bytes that a GPON ONU would transmit compared to an EPON ONU for the same sequence of upstream time allocations as a percentage of the total transmitted payload. This gain is due to the complete suppression of the USR by GEM, which turns out to save more bandwidth than lost to the 5-byte overhead per packet, as well as the lower protocol overheads of GPON. This gain slightly decreases as system load rises from 35% towards 65% because longer upstream windows are granted to the ONUs hence the USR in EPON (which is limited by the maximum packet size) decreases as a percentage of the allocated time. The total GPON gain due to the adoption of segmentation and re-assembly ranges from 2 to 20%, depending on the selected D_m, and must be added on top of the 25% higher transport efficiency of GPON due to the 8/10bit coding in EPON. The overall bandwidth efficiency (maximum throughput achieved in each case) taking into account all the above effects, is summarized in Figure 3. Obviously not as good a trade-off between bandwidth efficiency and delay variation as can be reached in GPON is possible in EPON, since the reduction of D_m in order to make polling more frequent (and bring down average access delay and delay variation for delay sensitive applications) results in high losses due to the lack of segmentation mechanisms and the USR effect .

4. Conclusions

EPON design aimed at exploiting the widespread and mature Ethernet technology for reducing component development effort, design cycles and overall cost. GPONs on the other hand aimed at higher line rates accepting higher receiver circuit costs while targeting a set of mechanisms for flexible traffic multiplexing, detailed traffic management specifications and Quality of Service guarantees with better control of network resource allocation as well as operation and maintenance. In this paper we compare the two technologies mainly in terms of MAC efficiency. It is shown that GPON enjoys improved performance by the more elaborate MAC features introduced into the standard and specifically: segmentation and reassembly, reduced physical layer (including burst preambles and line coding) and reduced MAC protocol overheads. Of course the overall merit of the two competing solutions cannot be judged on the performance alone.

Acknowledgement: This work was funded by the research project NGEAN.

References

1. ITU-T, Rec. G.983.1. Study group 15: Broadband optical access systems based on passive optical networks (PON), October 1998
2. ITU-T, Rec. G.984.1, Study Goup 15, “Gigabit-capable Passive Optical Networks (GPON): General characteristics”, March 2003.
3. J. Angelopoulos, H. Leligou, T. Argyriou, S. Zontos, E. Ringoot, T. Van Caenegem, “Efficient transport of packets with QoS in an FSAN-aligned GPON”, IEEE Communications mag. **42**, 92-98, (2004).
4. “802.3 Media Access Control (MAC) parameters and minimal augmentation of the MAC operation”, IEEE Std 802.3ah-2004 (IEEE, 2004).
5. A. Shami, X. Bai, N. Ghani, C. M. Assi, and H. T. Mouftah, “QoS Control Schemes for Two-Stage Ethernet Passive Optical Access Networks” IEEE J. Selected Areas In Comm. **23**, 1467-1478 (2005).
6. M. P. McGarry, M. Maier, and M. Reisslein, “Ethernet PONs: A Survey of Dynamic Bandwidth Allocation (DBA) Algorithms,” IEEE Commun. Mag. **42**, S8–S15 (2004).
7. H N. Ghani, A. Shami, C. Assi, M. Raja, “Intra-ONU bandwidth scheduling in Ethernet Passive Optical Networks,” IEEE communication let. **8**, 683-685 (2004).
8. M. Ma, Y. Zhu, T. H. Cheng, “A Bandwidth Guaranteed Polling MAC Protocol for Ethernet Passive Optical Networks,” in Proc. Infocom 2003, April 1 - 3, San Francisco, CA, USA, 2003
9. F. An et al., “A new dynamic bandwidth allocation protocol with quality of service in ethernet-based passive optical networks,” in Proc. Int. Conf. on Wireless & Optical Comm. (WOC03), Banff, Canada, Jul. 2003
10. A. Shami, X. Bai, C. M. Assi, and N. Ghani, “Jitter Performance in Ethernet Passive Optical Networks,” IEEE J. of Lightwave Techn. **23**, 1745-1753 (2005).
11. T. Orphanoudakis, H.-C. Leligou, E. Kosmatos, J. D. Angelopoulos, K. Kanonakis, G. Prezerakos, I. Venieris, “Efficient Resource Allocation with Service Guarantees in Passive Optical Networks”, OSA Journal of Optical Networking, Vol. 6, No. 6, June 2007.
12. K. Claffy, G. Miller and K. Thompson, “The nature of the beast: recent traffic measurements from an Internet

backbone," in Proc. of ISOC INET'98, Washington, DC, 1998.
13. J.D. Angelopoulos, I.S. Venieris, G.I. Stassinopoulos, "A TDMA based Access Control Scheme for APON's," IEEE/OSA Journal of Lightwave Technology, Special Issue: Broadband Optical Networks, Vol. 11, No. 5/6, May/June 1993, pp. 1095-1103.