Introduction to IEEE 802.3av 10Gbit/s Ethernet Passive Optical Networks (10G EPON)

White Paper

by Steve Gorshe, Principal Engineer &
Jeff Mandin, Technical Advisor

Issue No. 1: January 2009

PMC-Sierra, Inc.

Legal Information

Copyright

PMC-2081212, Issue 1

Disclaimer

None of the information contained in this document constitutes an express or implied warranty by PMC-Sierra, Inc. as to the sufficiency, fitness or suitability for a particular purpose of any such information or the fitness, or suitability for a particular purpose, merchantability, performance, compatibility with other parts or systems, of any of the products of PMC-Sierra, Inc., or any portion thereof, referred to in this document. PMC-Sierra, Inc. expressly disclaims all representations and warranties of any kind regarding the contents or use of the information, including, but not limited to, express and implied warranties of accuracy, completeness, merchantability, fitness for a particular use, or non-infringement.

In no event will PMC-Sierra, Inc. be liable for any direct, indirect, special, incidental or consequential damages, including, but not limited to, lost profits, lost business or lost data resulting from any use of or reliance upon the information, whether or not PMC-Sierra, Inc. has been advised of the possibility of such damage.

Trademarks

For a complete list of PMC-Sierra's trademarks and registered trademarks, visit: http://www.pmc-sierra.com/legal/

Patents

The technology discussed in this document may be protected by one or more patent grants.

Abstract

Providing the “triple play” services of voice, video, and high-speed data access is an important way for carriers to increase their revenue and compete with other access providers such as the CATV operators. For both telephone network providers and CATV providers, the most flexible and future-proof medium for providing triple play services is fiber, with its virtually unlimited bandwidth availability. Since providing a direct optical connection between the CO and each subscriber is cost prohibitive, most optical access systems share a passive optical network (PON) among multiple subscribers. The existing 1 Gbit/s IEEE 802.3ah Ethernet PON (EPON) and ITU-T G.984 2.5 Gbit/s PON (GPON) are currently being deployed. While these protocols provide high speed data, IPTV and VoIP access, more bandwidth is required for high definition switched digital video over Internet Protocol television (IPTV). This white paper provides a tutorial overview of the IEEE 802.3av 10Gbit/s Ethernet PON (10G EPON) standards, including the ways in which it differs from EPON.

About PMC

PMC-Sierra is a leading provider of high-speed broadband communications and storage semiconductors and MIPS-Powered processors for Enterprise, Access, Metro Optical Transport, Storage Area Networking and Wireless network equipment. The company offers worldwide technical and sales support, including a network of offices throughout North America, Europe and Asia. The company is publicly traded on the NASDAQ Stock Market under the PMCS symbol.

About the Authors

Steve Gorshe, Ph.D. is a Principal Engineer in PMC-Sierra’s Chief Technology Officer’s organization and oversees technology for SONET/SDH, optical transmission and access systems.

Currently Steve is a Fellow of the IEEE, Director of Magazines for the IEEE Communications Society, Associate Editor-in-Chief and former Broadband Access Series co-editor for IEEE Communications magazine. He has worked on the development and research of access and transport systems since 1983 and has been with PMC-Sierra since 2000. He has also authored/co-authored over 250 contributions to six standards bodies, including the ATIS OPTXS (T1X1) Committee responsible for SONET and optical network interface standards where he serves as Chief Editor and technical editor for multiple standards. He received ATIS’s two highest awards for his technical and editorial contributions. He has also serves as technical editor for multiple ITU-T standards including G.7041/Y.1303 (Generic Framing Procedure – GFP) and G.8011.1 (Ethernet Private Line Service) recommendations. He has 30 patents issued or pending, over 24 published papers, and is co-author of a telecommunications textbook.

Jeff Mandin is a Technical Advisor in PMC-Sierra's Fiber-to-the-Home Business unit, and leads PMC's activities in the IEEE 802.3 10G EPON Task Force.

Jeff has worked in communication systems development and research since 1990 and has been with PMC-Sierra since 2006. Since 2003, Jeff has held positions and contributed extensively to communication standards in IEEE 802, IETF, and CableLabs. In the IEEE 802.16e (Mobile WiMAX link layer) Task Force he chaired security activities and also acted as the group's liaison to the IETF. He is author of several issued and pending patents.

Revision History

Issue No.	Issue Date	Details of Change
1	January 2009	Document created

Table of Contents

Legal Information 2
Copyright 2
Disclaimer 2
Trademarks 2
Patents 2
Abstract 3
About PMC 3
About the Authors 3
Revision History 4
Table of Contents 5
List of Figures 6
List of Tables 6
1 Introduction 7
2 Introduction to PON 8
3 IEEE 802.3av 10Gbit/s Ethernet-based PON (10G EPON) 10
3.1 10G EPON Physical Layer 10
3.2 Signal Formats and MAC Protocol 15
3.2.1 Signal Formats 15
3.2.2 MAC-layer Control Protocol 18
3.3 FEC (Forward Error Correction) 20
3.4 ONU Discovery and Activation 23
3.5 Ranging Mechanism 25
3.6 EPON OAM 26
3.7 Dynamic Bandwidth Allocation 26
4 Summary of EPON and 10G EPON Features 28
5 Conclusions 29
Appendix A – 64B/66B Line Code 30
6 References 35
7 Glossary of Acronyms 36

List of Figures

Figure 1 PON Network Example 8
Figure 2 10G EPON optical spectrum allocation 11
Figure 3 Illustration of EPON and 10G EPON ONUs sharing the same PON 12
Figure 4 Dual rate receiver option illustration 14
Figure 5 10 Gbit/s upstream burst transmission illustration 16
Figure 6 MPCPDU illustration, including the GATE and REPORT MPCPDU information fields 17
Figure 7 Ethernet PON MAC Operation Example 18
Figure 8 FEC overhead locations (downstream) 21
Figure 9 10G EPON FEC code block formatting and transmission 22
Figure 10 ONU Discovery handshake 24
Figure 11 64B/66B block code structure 31
Figure 12 64B/66B block code self-synchronous scrambler 34

List of Tables

Table 1 Summary of EPON and 10G EPON Protocols and Features 28
Table 2 Encoding for non-ordered set control codes 33
Table 3 Encoding for ordered set control codes 33

1 Introduction

Telephone companies and community access television (CATV) providers, also called "cable" providers or multiple system operators (MSOs), are competing to offer subscribers the triple play services of voice, video, and high-speed data access. Both have historically relied on copper cables to connect through the last mile to their subscribers. The coaxial cable of the CATV companies has superior bandwidth capabilities relative to the twisted pair wiring from telephone companies. However, the coaxial cable must be shared by many subscribers in order to be economical. Also, most of the coaxial cable's bandwidth is divided into channels for downstream broadcast video channels, leaving limited bandwidth available for broadband upstream traffic and services such as video on-demand (VOD). Clearly, the most flexible and future-proof medium is fiber, with its virtually unlimited bandwidth availability, including the ability to provide broadband data services and broadcast video services over different wavelengths. For telephone network providers, fiber connections have been seen as the key to leapfrog the capabilities of the MSOs. As the telephone companies deploy their all-fiber access networks, MSOs are also exploring all-fiber networks in order to remain competitive.

Since providing a direct optical connection between the CO and each subscriber is cost prohibitive, most optical access systems share a passive optical network (PON) among multiple subscribers. The PON systems currently being deployed are those of the 1 Gbit/s EPON [3] and 2.5 Gbit/s GPON [8] standards. At least two broad factors are already at work to push a migration to the emerging 10G EPON [1]. The first is the increased bandwidth of home networks. Both the IEEE 802.11 wireless networks and the wireline home networks are increasing in capacity beyond 100 Mbit/s. Part of this factor is the decreasing cost and increasing availability of 802.11n and 1 Gbit/s Ethernet interfaces on new personal computers. The other factor is the expectation of increased customer desire for more on-demand digital video delivery. While the current generation of PON systems can satisfy some of this demand, the migration to HD video will require higher speed PONs.

In order to better understand IEEE 802.3av 10 Gbit/s Ethernet PON (10G EPON), the white paper begins with a brief introduction to PON. For additional information on PON technology and legacy PON protocols, see PMC white paper PMC-2061015. The 10G EPON tutorial that constitutes the main body of the white paper includes highlights of how it differs from the legacy IEEE 802.3ah Ethernet PON (EPON).

2 Introduction to PON

In order to better understand 10G EPON, the white paper begins with a brief introduction to PON systems and technologies. For a more extensive discussion of PON and fiber access system evolution, see [2].

Figure 1 PON Network Example

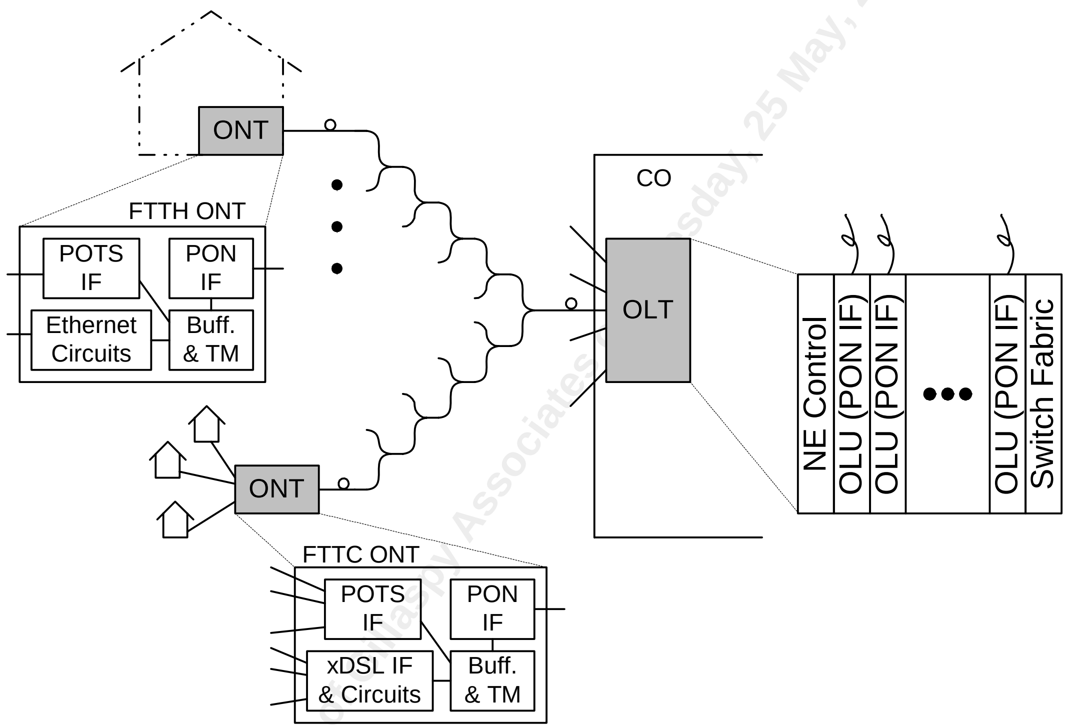

Passive optical networks (PONs) have become popular as a way to reduce the number of optical transceivers and fibers by moving the splitters closer to the subscriber. As illustrated in Figure 1, a PON system uses a single optical transceiver at the optical line terminal (OLT) to serve multiple subscribers over a fiber tree/bus network constructed with passive optical signal splitters. For fiber to the home (FTTH), the optical network unit (ONU) is at the customer's premises, either mounted on an outside wall by the network interface or inside the house. The FTTH ONU optionally provides the POTS interface to the subscriber and an Ethernet interface for data services. For non-FTTH ONUs such as for fiber to the curb (FTTC), the ONU provides the final copper drop to the subscriber. This interface will typically use either a CAT5 Ethernet connection or a DSL technology for the data service with analog POTS riding along in its native frequency range. With FTTC, the loop lengths are typically less than 1000 feet, making very high rate digital subscriber lines (VDSL) very practical for video delivery.

As shown in Figure 1, the OLT consists of a number of PON interface units, a switch fabric, and a NE controller. The ONUs are ultimately also managed by the NE controller, which is responsible for all ONU provisioning and OAM&P reporting. The OLT and ONUs together form the PON system, making it logically function as a single NE. In this very limited sense, the fiber interconnection can be conceptually thought of as an extended backplane.

In the downstream direction, the OLT broadcasts the data for all ONUs. This downstream signal comprises the downstream data for all the ONUs, the overhead for OAM&P, and the synchronization information for the upstream transmissions. The ONUs extract their downstream data based on cell/packet frame address information.

In the upstream direction, a time domain multiple access (TDMA) is the most common approach for sharing bandwidth among the ONUs. With TDMA, the OLT assigns (grants) each ONU a time slot in which to transmit their upstream data. A guardband time is required between the upstream burst transmissions of different ONUs so that their transmissions don't overlap at the OLU receiver. In order to minimize this guardband time, the OLT uses a protocol to determine the round trip delay time between itself and each ONU and takes it into account when assigning the ONU upstream transmission times.

In basic PON systems, the OLT makes the upstream bandwidth assignments such that each ONU has a fixed average portion of the upstream bandwidth. In order to achieve greater efficiency, PON systems commonly allow dynamic bandwidth allocation (DBA) among the ONUs. With DBA, each ONU communicates its bandwidth requirements to the OLT. For example, this information could be based on its input queue fill level, including the levels for data in different classes of service. The OLT evaluates the requests from the ONUs and assigns the bandwidth for the next upstream transmission frame. The information used by the OLT in determining the appropriate bandwidth allocations can include the service level agreements (SLAs) associated with the data flows associated with the ONUs. These bandwidth assignments are sent in a downstream transmission frame, and are typically communicated as a transmission start time and either a stop time or transmission duration. In some systems, the ONU is responsible for determining how to accommodate the relative priorities of its transmit data within the granted upstream transmission slot.

3 IEEE 802.3av 10Gbit/s Ethernet-based PON (10G EPON)

The IEEE 802.3av PON standard was developed to increase the data rate of EPON systems from 1 Gbit/s to 10 Gbit/s, in keeping with the 10 Gbit/s Ethernet interface. 10G EPON shares much of its protocol with EPON. A combination of coarse wave division multiplexing (CWDM) and time division multiplexing (TDM) is used in order to allow EPON and 10G EPON systems to co-exist on the same PON. As with EPON, 10G EPON relies on VoIP for carrying voice traffic and circuit emulation service (CES) for carrying other TDM clients.

3.1 10G EPON Physical Layer

The downstream data rate of 10G EPON is 10 Gbit/s. Both 1 Gbit/s and 10 Gbit/s rates are supported in the upstream direction. The 64B/66B block line code, described in Appendix I, is used for all of the 10Gbit/s signals with a resulting signal line rate of 10.3125 Gbit/s. The 1 Gbit/s upstream uses the same 8B/10B block line code as EPON, giving a line rate of 1.25 Gbit/s.

The downstream and upstream data is transmitted over a single PON fiber, using WDM to separate the upstream and downstream signals. The wavelengths used by the different upstream and downstream signals are shown in Figure 2. Since there are many ONUs on the PON and only a single OLT, the wavelength bands were chosen to allow using less expensive lasers at the ONUs[1].

For 1 Gbit/s upstream operation, 10G EPON uses the same 1310nm wavelength as the EPON upstream signal. This allows the OLT to use the same receiver for all 1 Gbit/s signals. The dynamic bandwidth allocation algorithm (see [2]) allocates the bandwidth of the 1 Gbit/s upstream signal between the EPON and 10G EPON ONUs.

The 10 Gbit/s upstream signal use a separate wavelength band, however it overlaps with the 1 Gbit/s upstream wavelength band. When an OLT supports both 1 Gbit/s and 10 Gbit/s operation on the same PON it is referred to as a dual rate mode. The dual rate OLT can either separate the 10 Gbit/s and 1 Gbit/s upstream signals by dividing the signal in the optical domain or in the electrical domain. With optical domain separation, one approach is to use a 1:2 optical splitter to feed the signal to both the OLT's 1 Gbit/s and 10 Gbit/s receivers. Since the splitter introduces ≥3dB additional loss, a low gain optical amplifier may be required to achieve the desired performance for the longer reach interfaces. Multiple techniques exist for splitting the signal in the electrical domain after the optical receiver. There are tradeoffs between these implementations relative to their performance versus complexity.

[1] Lasers that operate in the 1270nm and 1310nm regions are less expensive than those operating in the 1500-1600nm range due both the to the technology required for their fabrication and the relative market volumes for the lasers.

Figure 2 10G EPON optical spectrum allocation

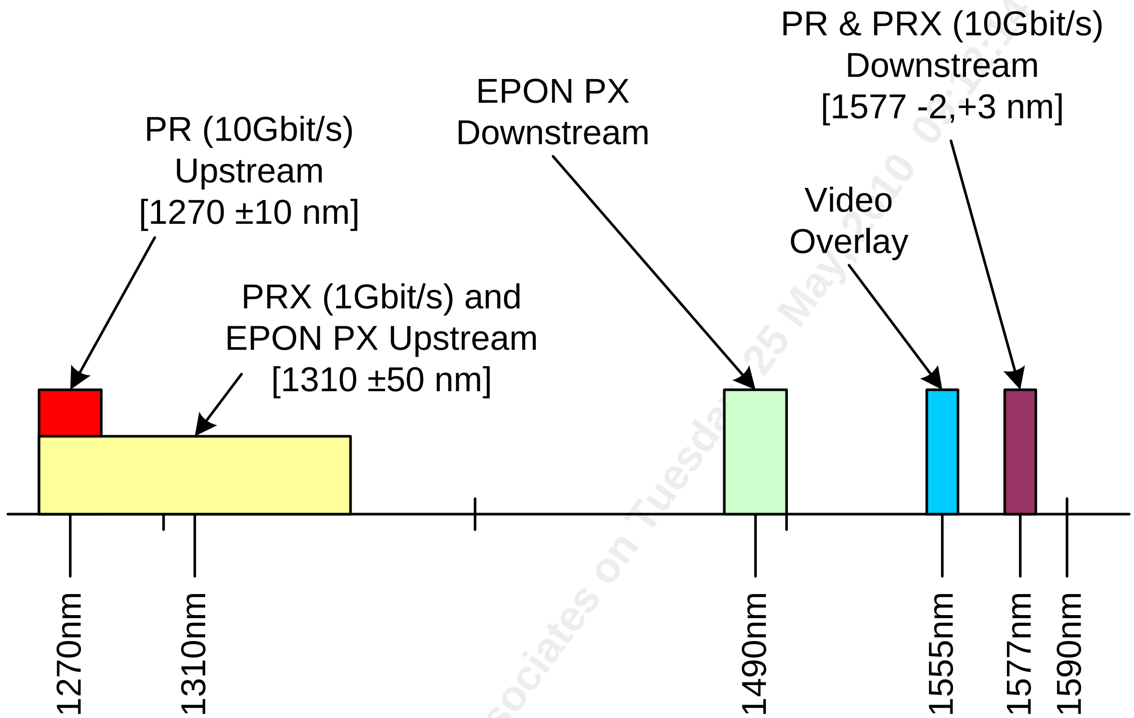

One of the key features of the 10G EPON specification is that it allows 10G EPON to operate over the same PON optical distribution network that is already being used for EPON. The advantages to this feature include:

- Allowing customers to use the most cost-effective ONU for the desired service
- Allowing a network to migrate from EPON to 10G EPON by upgrading the OLT then migrating the ONUs as needed
- Continued operation of the existing network and services during the upgrade of the network

Figure 3 illustrates a network where an OLT supports a mix of EPON ONUs, ONUs with 10Gbit/s downstream and 1Gbit/s upstream, and ONUs 10Gbit/s upstream and downstream. For convenience, the wavelength color key in Figure 3 is consistent with the key for Figure 2. Note that WDM is used to separate the 1Gbit/s and 10Gbit/s traffic in the downstream direction, and combination of WDM and TDM is used in the upstream direction. The discovery and other protocol extensions to support the co-existence of EPON and 10G EPON ONUs are discussed in the appropriate sections below.

Figure 3 Illustration of EPON and 10G EPON ONUs sharing the same PON

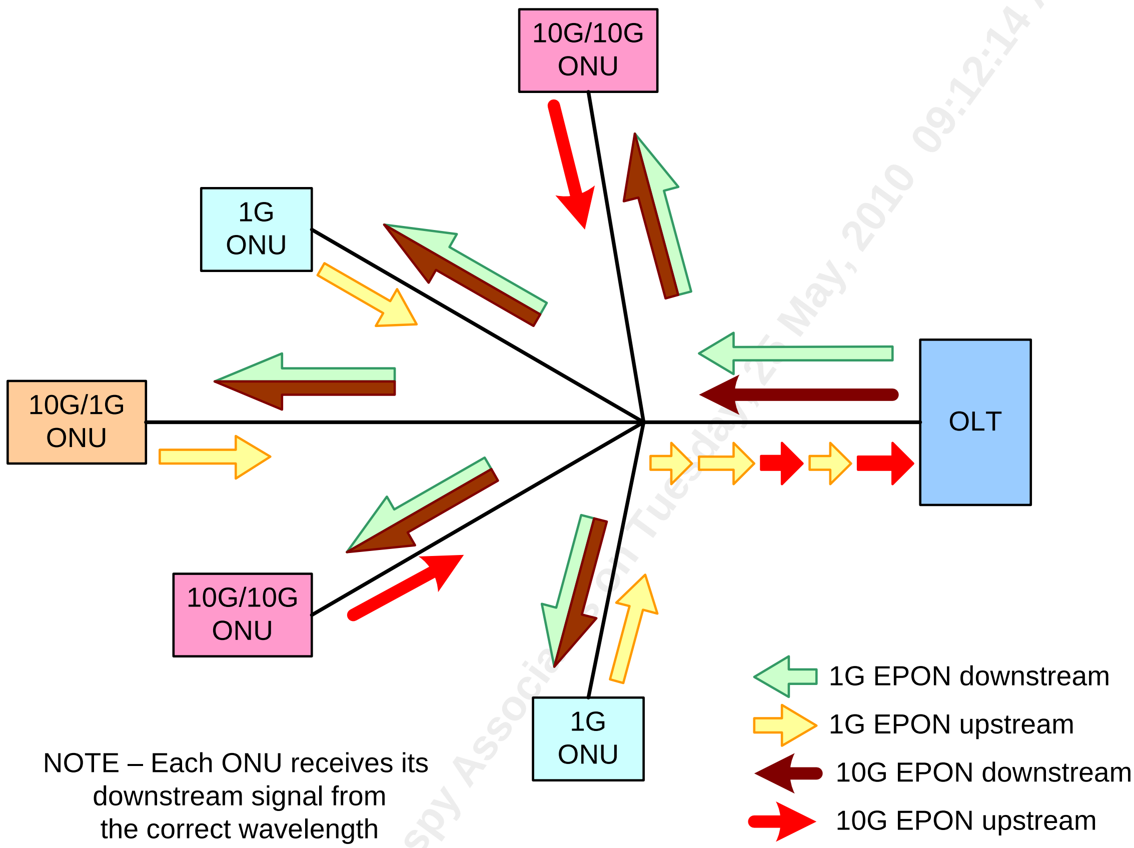

As its reference for the optical link loss budgets, the 802.3av specification uses a split ratio of either 1:16 (i.e., 16 ONUs on a single PON connecting to one OLT interface) or 1:32. In practice, larger split ratios such as 1:64 or 1:128 can be used if the other optical losses (e.g., the length of the fiber) are constrained to offset the additional 3dB loss that is incurred when the split ratio is doubled.[2] All of the interfaces are specified to operate at an uncorrected bit error rate no worse than (BER) 10^{-3}. After FEC correction, the bit error rate will be no worse than 10^{-12}. The nomenclature adopted to identify the different optical interface options may be summarized as follows:

- PRX interfaces use 10 Gbit/s downstream and 1 Gbit/s upstream transmission
- PR interfaces use 10 Gbit/s for both downstream and upstream transmission

---

[2] The limits on the split ratio are a combination of the functions of the optical parameters (e.g., loss budget) and the desired per-ONU bandwidth.

- PR-Dn and PRX-Dn (n = 10, 20, 30) refer to the OLT optical interface specification
- PR-Un and PRX-Un (n = 10, 20, 30) refer to the ONU optical interface specification
- PR10 and PRX10 specifies an optical channel insertion loss of ≤20 dB for ≥10 km reach with 1:16 split ratio
- PR20 and PRX20 specifies an optical channel insertion loss of ≤24 dB for ≥20 km reach with a 1:16 split ratio or ≥10 km reach with a 1:32 split ratio
- PR30 and PRX30 specifies an optical channel insertion loss of ≤29 dB for ≥20 km reach with a 1:32 split ratio

As with EPON, the 1550-1560nm-wavelength band is reserved for downstream video transmission.

Following the same approach as EPON, the upstream burst timing is relaxed for 10G EPON in order to allow using existing off-the-shelf components. The standard has mechanisms to allow for future tighter timing to be implemented with better components for increased bandwidth efficiency.

Dual-rate operation

Dual-rate operation refers to an OLT that simultaneously receives upstream signals from ONUs using 1 Gbit/s and 10 Gbit/s rates. As illustrated in Figure 4, the received 1 Gbit/s and 10 Gbit/s streams can either be separated in the optical domain or electrical domain. Since both signals time-share the same upstream wavelength, it is not possible to use WDM filters to separate them in the optical domain.

Separating the signals in the optical domain involves using a 1:2 optical splitter. Each of the two splitter outputs goes to its own photodetector followed by a receiver with a filter optimized for its bandwidth in order to maximize the receiver's sensitivity. The drawback with this approach is the roughly 3dB additional optical loss introduced by the 1:2 optical splitter. If this additional loss cannot be tolerated, a low-gain optical amplifier must be used in the receiver.

Splitting in the electrical domain allows using a single photodetector and introduces no additional optical signal loss. In the electrical domain, one approach is to design the receiver filter a compromise that allows reception of both the 1 Gbit/s and 10 Gbit/s signals. This means that the receiver sensitivity is not optimal for either signal, lowering it by about 1dB for each. Alternatively, the OLT can adjust (switch) the transimpedance of its transimpedance amplifier (TIA) filter for that burst's rate. The APD bias can either be set to a compromise value or switched along with the transimpedance.[3] While the performance of an adaptable receiver is optimum, it's additional complexity impacts the receiver cost. Detecting the rate of the current incoming burst must be performed fast enough to switch the receiver. The burst rate could be detected by looking for spectral energy that would only be present for a 10 Gbit/s burst. Alternatively, the OLT could exploit its knowledge of which upstream burst is scheduled to arrive. However, since this knowledge is in the MAC layer and not the PMD, it would be a layer stack restriction violation to require it.

Figure 4 Dual rate receiver option illustration

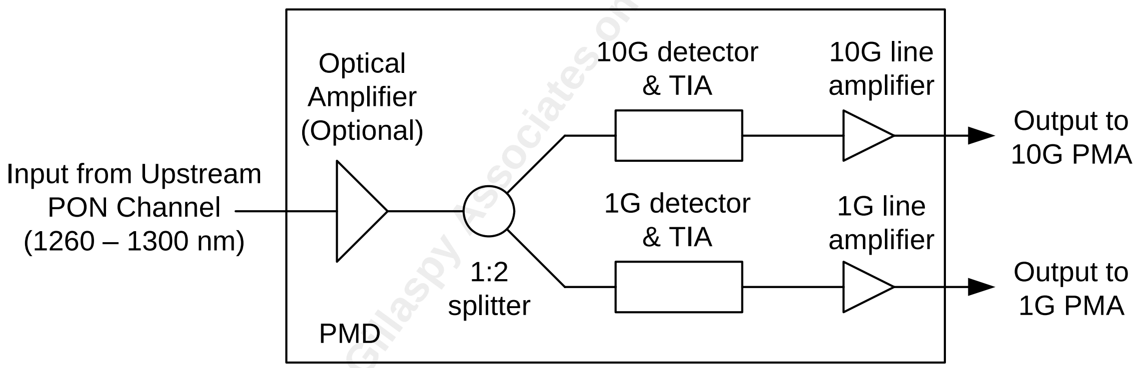

a) Optical domain rate splitting

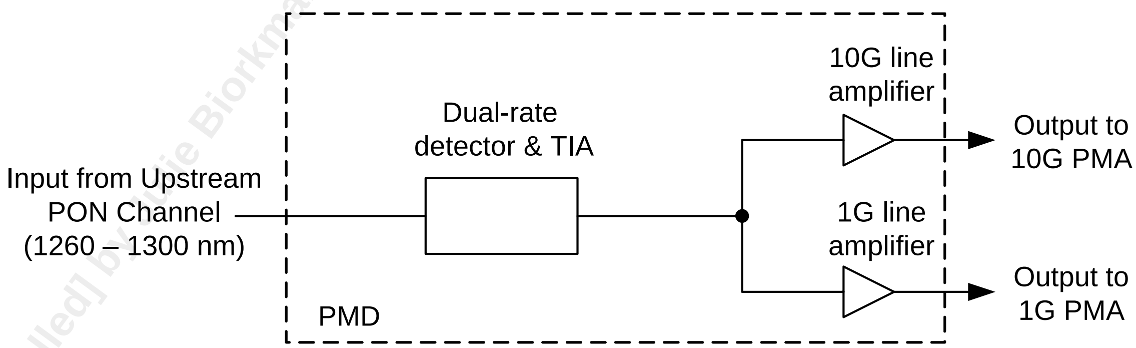

b) Electrical domain rate splitting

[3] Using a compromise APD bias results in a loss of around 1dB receiver sensitivity, which is 1dB better than using the compromise transimpedance value.

3.2 Signal Formats and MAC Protocol

3.2.1 Signal Formats

The downstream signal is simply a stream of Ethernet frames and Idle characters, as with a point-to-point 10 Gbit/s Ethernet signal.[4] The upstream signal uses a TDMA format, in which the upstream bandwidth is shared among the ONUs on a time division basis with the ONUs transmitting their upstream bandwidth in bursts. When the OLT assigns the transmission start and stop times each ONU, it must leave an adequate guard time between transmissions of the ONUs to insure that it is only receiving the signal from one ONU at time. As discussed in [2], this guard time is minimized by the OLT taking the relative distance to each ONU into account.

The guard time between the upstream transmissions of different ONUs must take into account the time needed for an ONU laser to turn on and off. If the ONUs do not turn their lasers off when they are not transmitting, spontaneous emission noise from ONUs closer to the OLT would interfere with data transmissions from ONUs further from the OLT.

The beginning of an upstream burst is illustrated in Figure 5. The synchronization patterns at the beginning of an upstream transmission burst allow the OLT to synchronize its receiver to new burst from an ONU. The Burst Delimiter pattern is used by the OLT to determine the start of 66B block transmission and the FEC codeword alignment. The 66-bit value of the Burst Delimiter is 0x 4 97 BA C4 69 F0 4C 88 FD (which results in a transmission bit sequence of 10 11101001 01011101 00100011 10010110 00001111 00110010 00010001 10111111). The FEC codeword alignment can be achieved in the presence of transmission errors. This burst delimiter is followed by two 66-bit blocks containing Idle characters. These Idle characters allow the OLT to synchronize its descrambler and delineate the start of the actual data frame. The first two blocks of Idle characters are included in the initial FEC codeword.

[4] As explained in section 3.3, both the upstream and downstream streams are encoded into FEC blocks in a manner that preserves the 64B/66B block stream format.

Figure 5 10 Gbit/s upstream burst transmission illustration

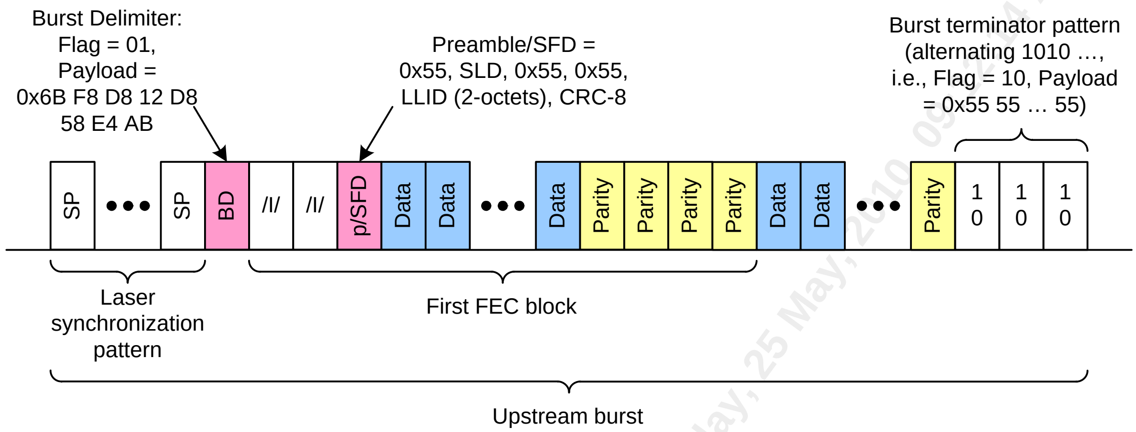

The preamble and start of frame delimiter (SFD) are modified for EPON and 10G EPON from their normal values for Ethernet. Specifically, the preamble bytes are replaced by the transmitting MAC's MODE and LLID variables. While the Ethernet 8-character preamble/SFD consists of/S/, 0x55, 0x55, 0x55, 0x55, 0x55, 0x55, and 0xd5, the EPON and 10G EPON preamble/SFD consists of 0x55, 0x55, SLD, 0x55, 0x55, 2-octet LLID, and CRC-8. The SLD is the Start of LLID Delimiter, and has the value 0xd5. The LLID is the two-octet logical_link_id field that uniquely identifies the ONU MAC. The MSB of the two octets that contain the LLID is the MODE indication bit. As discussed in section 3.4, the LLID is assigned to the ONU by the OLT during the registration phase of the discovery process. The CRC-8 covers the SLD through the LLID octets, and uses the generator polynomial $x^8 + x^2 + x + 1$.

The upstream transmission ends with a burst terminator pattern comprised of three 66-bit blocks of alternating zeros and ones (1010 … 10) after the last FEC codeword of the burst. The ONU turns off its laser at the beginning of the burst terminator pattern, which insures that it will be completely off by the end of the burst.

Each ONU has one unique Logical Link Identifier (LLID) that the OLT associates to the ONU for uni-cast traffic. In other words, these MAC instances are used emulate a point-to-point connection between and ONU and the OLT over the PON. In addition, the OLT has two Single Copy Broadcast (SCB) MAC instances that are used as an efficient mechanism to broadcast downstream traffic to the ONUs. Such a broadcast is used for broadcast data or for when the OLT must communicate with unregistered ONUs. In the upstream direction, an SCB MAC is only used for client registration. The LLID value of 7F-FF is associated with the SCB MAC for 1 Gbit/s downstream operation and the LLID value of 7F-FE is associated with the SCB MAC for 10 Gbit/s downstream operation. An ONU can use higher layer networking processing, such as VLAN filtering and IGMP snooping, to narrow the amount of received multicast traffic that is passed to applications. It is possible that these higher layers may require addition multicast MAC instances at the OLT, in which case an OLT can have more MACS than two plus the number of ONUs.

Multi-Point Control Protocol PDUs (MPCPDUs) are control frames used by the ONUs to make their requests for bandwidth, and by the OLT to assign it. As illustrated in Figure 6, the MPCPDU[5] frame is a basic 802.3 MAC control frame containing a 4-byte timestamp and a 40-byte field filled with data and padding as needed. MPCPDU messages are also used for the discovery and ranging processes, as discussed in sections 4.3 and 4.4. MPCPDUs are layered below the data interface, and have higher priority than any data packet. This ensures that the bandwidth requests and grants are sent in a timely manner.

Figure 6 MPCPDU illustration, including the GATE and REPORT MPCPDU information fields

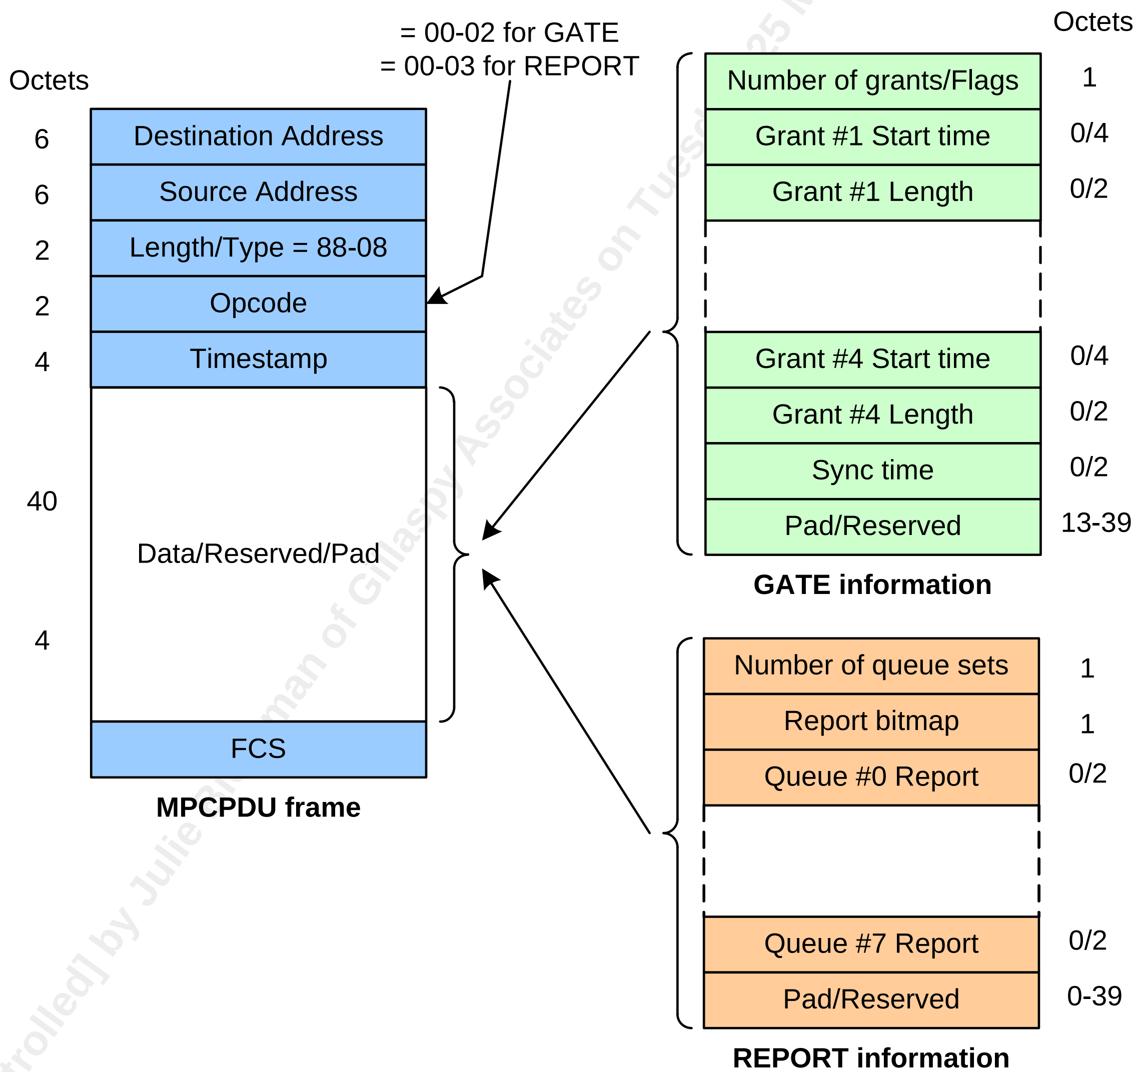

[5] The MPCPDU Ethertype is 0x88-08. The opcodes are assigned between 00-02 and 00-06.

3.2.2 MAC-layer Control Protocol

The 10G EPON MAC-layer control protocol is based on the protocol for EPON and includes enhancements for management of 10G FEC and inter-burst overhead. This MAC protocol operates on the basis of the ONUs informing the OLT of their upstream bandwidth requirements, and the OLT scheduling and granting bandwidth to the ONUs to transmit their upstream data. The details of the MAC protocol are described in this section and are illustrated in Figure 7 with the downstream and upstream data flows.

Figure 7 Ethernet PON MAC Operation Example

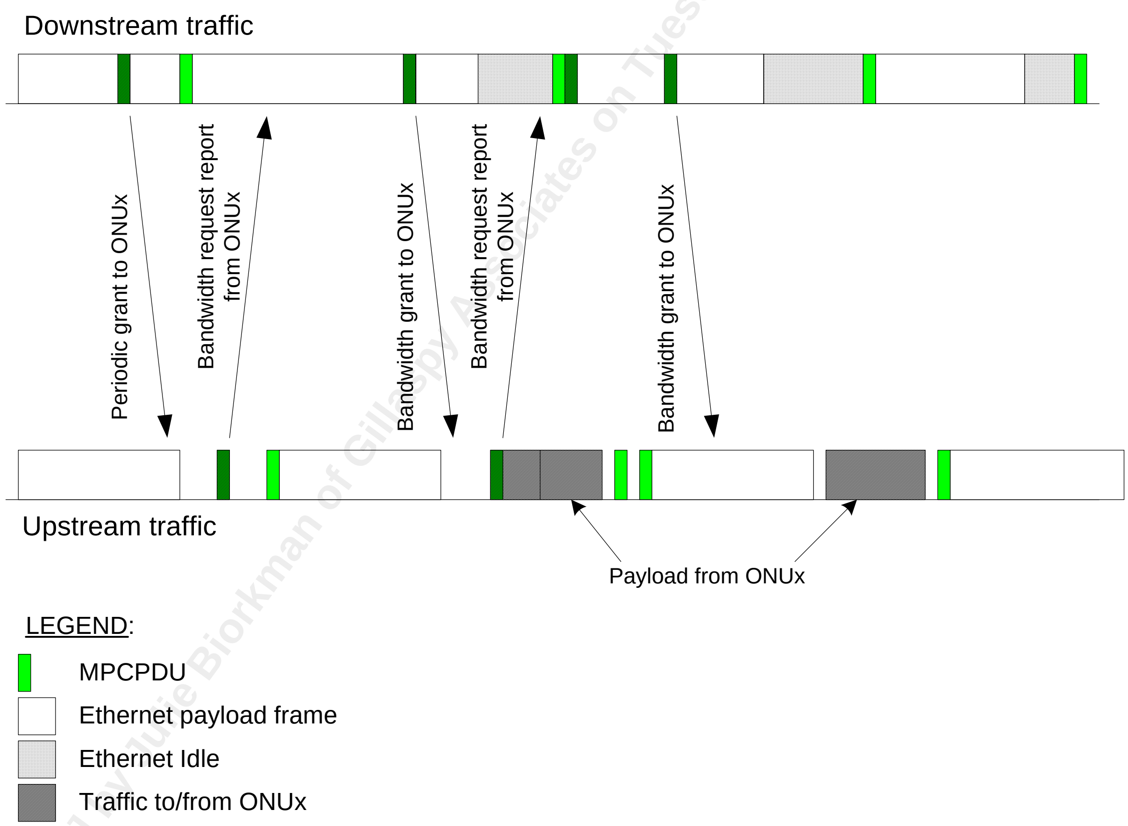

GATE Messages for Upstream Bandwidth Grants

The OLT grants bandwidth to an ONU in the GATE MPCPDU message that is transmitted downstream. Gating is the function that controls when the ONUs are allowed to transmit upstream data. The Gating function relies on a local timer that is synchronized to the OLT timer. (See the ranging protocol description in section 3.5 for a further discussion of the timers.) The GATE message specifies the ONU upstream start time and transmission length. The bandwidth grants from OLT are always made for at least 1024 time quanta into the future so that the ONU has time to process the GATE message and be ready to transmit. An ONU turns its laser on when the local time equals the start time specified in the GATE message. The length field gives the length of time that the ONU is allowed to transmit its burst[6]. The ONU stops it transmission such that the laser can be completely turned off before the end of the granted transmission length. The start time is a 32-bit number (the same length as the local timer counter), and the length is a 16-bit number. The OLT includes the time required to turn the ONU laser on and off and the time to send the upstream synchronization patterns when it assigns the grant time.

The OLT sends GATE messages to each ONU periodically so that they can report their upstream bandwidth needs. The first (left-most) grant in Figure 7 is an example of such a periodic grant. The GATE MPCPDU is illustrated in Figure 6. The ONUs have watchdog timers that are reset whenever a GATE message is received[7]. Note that a flag in the GATE message can be used to force the ONU to respond with the REPORT MPCPDU described below.

The number of upstream bandwidth grants in a GATE message can range from 0-4. As discussed in section 4.3, the ONU advertises the number of outstanding grants it can accept during the discovery process. A GATE message with no grants is used as a keep-alive for the ONU timeout counters and to communicate the timestamp information. Multiple grants may be made in a single GATE message, as illustrated with the first payload transmission from ONUx in Figure 7. For example, the different grants could correspond to different priority queues at the ONU. In practice, however, this adds considerable complexity to the OLT bandwidth assignment process. Sending a single grant in the GATE message gives much finer resolution and faster response for the upstream bandwidth assignments, and has only a small impact on the downstream overhead bandwidth. The preferred approach, especially for per-flow DBA, is for the OLT to send each ONU a single grant in each GATE message to service its bandwidth requests and let the ONU decide which data should be sent in that upstream grant.

---

[6] The time window is specified with respect to the number of periods of the ONU's local clock.

[7] For the purposes of keeping the ONU's watchdog timer alive, an OLT can also periodically send empty GATE messages when it has no pending bandwidth requests for that ONU.

The "Sync Time" field in the GATE MPCPDU is a report from the OLT to communicate the amount of time it needs at the beginning of the upstream transmission burst to synchronize its receiver to the new burst. As illustrated in Figure 5, each burst begins with a synchronization pattern, followed by a Burst Delimiter pattern, followed by two blocks of Idle characters. The ONU transmits the 66-bit synchronization pattern repeatedly and then transmits the Burst Delimiter so that the duration of the entire sequence is the same as the Sync Time requested by the OLT.

REPORT Messages for Upstream Bandwidth Requests

The ONUs communicate their upstream bandwidth requirements by sending REPORT MPCPDU messages. The OLT grants the upstream bandwidth for these REPORT messages in its GATE messages. In addition to the timestamp, the REPORT message consists of a summary of its requests for upstream bandwidth, and the specific amount of bandwidth it needs. Like EPON, 10G EPON supports the eight queue priority levels defined in IEEE 802.1Q. The summary field of the REPORT message indicates how many and which, if any, of these queues have data to send. The summary is followed by binary numbers to indicate the specific number of bits to be transmitted from each queue. The bit count is a 16-bit number, and includes Inter-Packet Gap (IPG) characters. Unlike EPON, the bandwidth value carried by the 10G EPON REPORT does not include burst overhead or FEC overhead

Each ONU sends REPORT messages periodically, even if it has no data waiting for transmission, in order to reset a watchdog timer at the OLT. If the watchdog timer expires, the OLT deregisters that ONU from the network.

3.3 FEC (Forward Error Correction)

EPON and 10G EPON use different FEC approaches. FEC allows a link to operate with a higher bit error rate at the receiver. Consequently, FEC effectively increases the optical link budget, which in turn allows increased distance or split ratios. EPON specified an optional RS(255, 239) FEC that is applied to Ethernet frames. As illustrated in Figure 8a, the Ethernet frame is divided into blocks for purpose of computing the FEC block codes, and the resulting FEC check symbols for all the blocks are gathered and placed at the end of the Ethernet frame.

When the RS(255,223) FEC parity is taken into account, the effective data rate of a 10G EPON link is approximately 8.7 Gbit/s.

FEC becomes increasingly important as bit rate increases. For this reason, FEC is mandatory in 10G EPON. Additionally, the 10G EPON FEC differs in two ways from EPON. First, a more powerful RS(255, 223) code is used for error correction of 16 symbols rather than the 8 symbols that can be corrected with the RS(255,239) code.[8] Second, FEC is applied to fixed-length sequences of streaming data rather than Ethernet frames as illustrated in Figure 8b. Figure 8 illustrates the downstream transmission direction, which is a continuous stream of FEC codewords that include the Ethernet frames and all inter-packet information such as IPG and Ordered Set data. The upstream transmission is similar except, as illustrated in Figure 5, the first FEC codeword of an upstream burst is aligned with the beginning of the burst in order to allow the OLT FEC decoder immediate codeword synchronization for each burst.

Figure 8 FEC overhead locations (downstream)

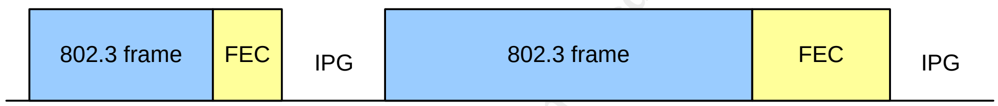

a) FEC overhead illustration for 1G EPON

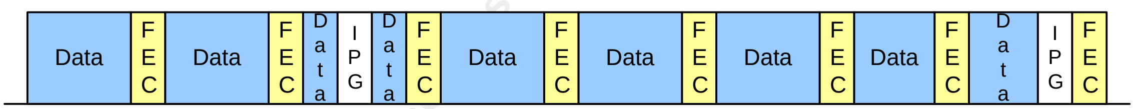

b) FEC overhead illustration for 10G EPON

[8] The generator polynomial $G(x) = x^8 + x^4 + x^3 + x^2 + 1$ is specified for the RS(255,223) here.

One of the challenges in adding FEC to the 10GE PON stream is extending the 64B/66B block code format so that a 10GbE receiver can receive and synchronize to the stream that now includes FEC parity data. The method used is illustrated in Figure 9. Each FEC codeword covers a group of 27 64B/66B blocks. As illustrated in Figure 9, the first step of FEC encoding is removing the first flag bit of the 64B/66B block.[9] The resulting 27x65 = 1755-bit block is padded with 29 leading zeros to get a total of 1784 bits (223 bytes). The RS(255,223) encoding produces 32 parity bytes. In the final stage, the zero pad bits are removed, the original 27 64B/66B blocks are restored, and the parity bytes are converted into a sequence of 64B/66B blocks for transmission. Specifically, the 32 FEC parity bytes are treated a four groups of 64 bits. Each of these 64-bit parity groups is then given a pair of leading header bits in order to create 64B/66B blocks. In order to create a recognizable header pattern, the header bits for the four parity blocks are 00, 11, 11, and 00, respectively. The string of 31 64B/66B characters is then transmitted.

The receiver can then synchronize to the 64B/66B character stream and extract the original data through the reverse process, performing error correction as it decodes the FEC blocks.

Figure 9 10G EPON FEC code block formatting and transmission

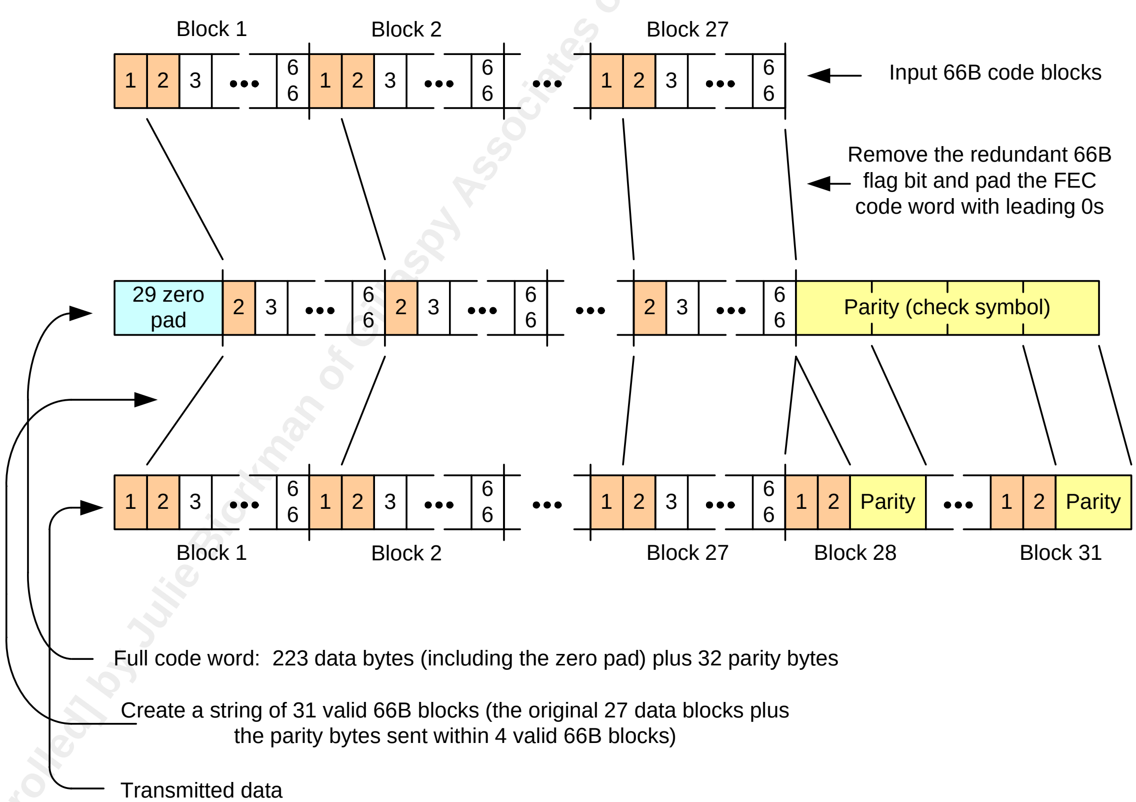

[9] Since the two leading flag bits of the 64B/66B block are intentionally redundant, only one of them needs to be protected by the FEC.

3.4 ONU Discovery and Activation

The message handshake is illustrated in Figure 10. The OLT periodically opens a Discovery Window in order to allow new ONUs to announce themselves[10]. The OLT opens the Window by transmitting a discovery GATE message, which includes the length of the Window and its start time. A discovery GATE MPCPDU includes a Discovery information field that communicates to the ONUs whether the OLT is capable of receiving 1Gbit/s upstream signals, capable of receiving 10Gbit/s upstream signals, and whether the discovery window is being opened is for 1Gbit/s or 10Gbit/s upstream signals from the ONUs.

[10] The period is implementation dependent.

Figure 10 ONU Discovery handshake

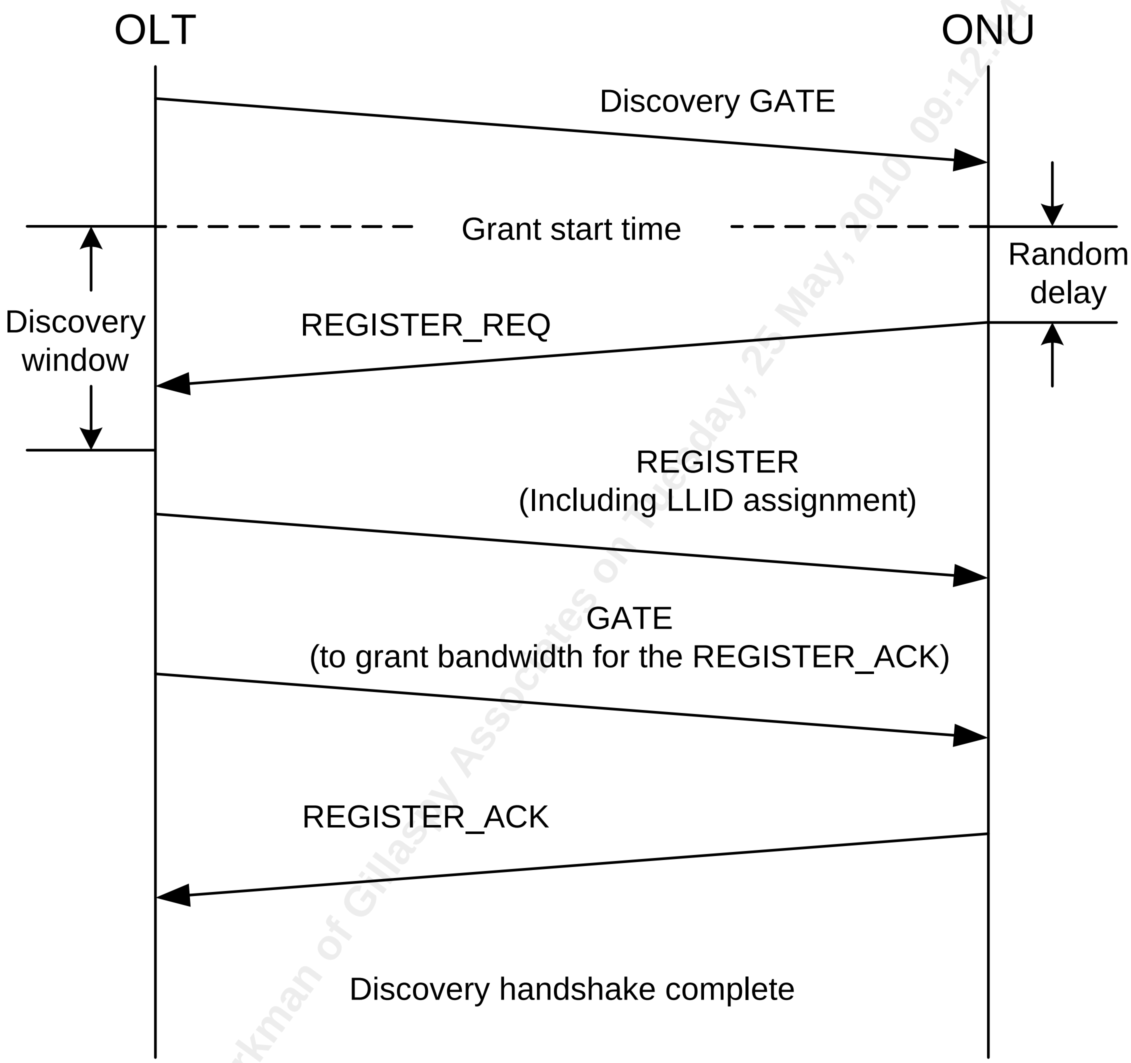

The unregistered ONUs respond to the discovery GATE message by transmitting a REGISTER_REQ message. The REGISTER_REQ message includes the ONU's MAC address and the number of outstanding grants that it can accept. (See section 3.2.2 regarding multiple grants.) The REGISTER_REQ has been expanded to include fields for Laser On time and Laser Off time to specify the time required by the ONU to turn its laser on and off. A contention algorithm is used in order to minimize the chance of collision when multiple ONUs are attempting to register during the same Discover Time Window. The contention algorithm operates by having each ONU delay its transmission by a random time relative to the beginning of the Discover Window[11].

When the OLT receives the REQUEST_REQ, it assigns a LLID to the ONU, bonding the LLID to the ONU's MAC address. The OLT then sends a REGISTER message to the ONU in order to communicate the ONU's LLID and the required OLT synchronization time, to echo the maximum number of pending grants that the ONU can accept, and to echo the ONU LaserOn and LaserOff time fields. The synchronization time is the amount of time the OLT will require in order to reliably synchronize to the ONU's upstream transmission burst. The synchronization time is specified in multiples of 66-bit data patterns that the ONU sends at the beginning of the burst.

After the ONU has processed the Register message, it sends a Register_ACK message to the OLT in response to a standard GATE message from the OLT.

Note that the Discovery GATE, REGISTER_REQ, and REGISTER messages are sent on the broadcast channel since ONU doesn't know its LLID until it receives the REGISTER message. After the ONU receives its LLID, the remaining GATE and REGISTER_ACK messages are sent on the unicast channel.

Mechanisms exist in the protocol to deregister an ONU (e.g., if a watchdog timer expires) and re-register.

3.5 Ranging Mechanism

The ranging mechanism for EPON makes use of local clocks that are maintained by the OLT and each ONU. A counter has 32 bits, and is incremented once every 16 ns.

[11] The upper bound requirement on this hold-off time is that it must be short enough that the ONU can transmit its entire REGISTER_REQ before the end of the Discovery Window.

The OLT counter is the PON master. When the OLT transmits a MPCPDU message, it loads its current counter value into the message's 32-bit timestamp field. When the ONU receives a MPCPDU, it resets its own local counter to the value contained in the MPCPDU timestamp field. When the ONU sends a MPCPDU to the OLT, the ONU loads its updated counter value into the timestamp field. The OLT then compares the offset between its current count and the value it receives in the MPCPDU timestamp field, with the difference being the round trip time (RTT). The RTT is then used to establish the ONU's range, which is taken into account when the OLT assigns the start times for upstream bandwidth grants.

Some drift may occur in the RTT over time. When the drift exceeds a specified threshold, a timestamp drift error condition is declared. Either the ONU or OLT can detect this condition as an offset between the expected value received in the MPCPDU and the one actually received.

3.6 EPON OAM

Ethernet link OAM was developed as part of the same IEEE 802.3ah project that developed EPON. Several Ethernet frames were defined to communicate the link OAM information.

3.7 Dynamic Bandwidth Allocation

With Dynamic Bandwidth Allocation (DBA), the OLT assigns the bandwidth to the different ONUs based on the data they have to send rather than a static allocation per ONU. EPON uses the REPORT messages from the ONUs to inform the OLT of their current bandwidth needs. Their bandwidth requests are reported in terms of the number of characters they have in the different priority queues awaiting upstream transmission. The OLT can also take into account the service level agreements (SLAs) that have been specified for the service flows associated with an ONU. For example, an ONU with an active VoIP service will need a fixed amount of bandwidth on a regular basis. Consequently, the ONU does not need to waste upstream bandwidth reporting bandwidth requests for this service flow. As another example, if the OLT receives upstream bandwidth requests from multiple ONUs, it can grant more bandwidth to ONUs that have been recently consistently requesting more bandwidth than to ONUs that have had few recent requests. In this example, the DBA algorithm needs to insure the nodes with fewer bandwidth requests don't become starved or encounter high latency while the ONUs with more bandwidth requests are services.

EPON DBA has the flexibility to customize EPON network behavior to meet various carrier needs. The flexible nature of the EPON DBA is defined in both EPON standards (1G & 10G) and allows quick adaptation to possible carrier challenges, making the EPON infrastructure compliant with the ever growing, ever changing carrier's requirements. It is possible to map both user and service flows in to specific containers that are managed by the DBA and provide the QoS that is needed for every customer and service. Two straightforward adjustable parameters related to EPON DBA are latency and total system performance (upstream BW utilization). An example of this is the Service DBA algorithm used in PMC's EPON OLT devices.

One benefit of the 10GEPON system is the ability to overcome system bottlenecks via adjustments in the EPON DBA algorithm. The DBA cycle length and BW allocation per ONU can be adjusted so that the total OLT upstream transmission going into the switch will be "smoother", less bursty in nature, allowing carriers to overcome blocking elements in their network topology (e.g. assigning more BW to the OLT ports then the uplink ports in the switch connected to the OLT to save CAPEX).

In order to flexibly support a variety of DBA algorithms, PMC-Sierra OLT devices include powerful embedded processors. PMC-Sierra also provides the software for multiple DBA algorithms that customers can use directly or modify according to their needs.

4 Summary of EPON and 10G EPON Features

Table 1 gives a summary of the key features of EPON and 10G EPON.

Table 1 Summary of EPON and 10G EPON Protocols and Features

Feature	EPON	10G EPON	Comment
Responsible standards body	IEEE 802.3 (IEEE 802.3ah)	IEEE 802.3 (IEEE 802.3av)	
Data rate	1 Gbit/s upstream & downstream	10 Gbit/s downstream and either 1 or 10 Gbit/s upstream	
Split Ratio (ONUs/PON)	1:64	1:64	While the optical parameters were designed around 1:16 for EPON, and 1:16 or 1:32 for 10G EPON, in practice both support 1:64.
Line code	8B/10B	Down: 64B/66B Up: 8B/10B (1 Gbit/s) or 64B/66B (10 Gbit/s)	The 8B/10B line code results in a 1.25 Gbit/s line rate for the 1 Gbit/s MAC rate, and the 64B/66B line code results in a line rate of 10.3125 Gbit/s for the 10 Gbit/s MAC rate.
Number of fibers	1	1	
Wavelengths	1490 nm down & 1310 nm up	1577 nm down & 1310 nm up (1 Gbit/s) or 1270 nm up (10 Gbit/s)	Both support 1550 nm for downstream video overlay.
Maximum OLT to ONU distance	10 and 20 km (1:16 split)	10 km with 1:16 split 20km with 1:16 split 20km with 1:32 split	
Optics			
Protection switching	None	None	Defined in regional specifications or other bodies
Data format (encapsulation)	None (uses Ethernet frames directly)	None (uses Ethernet frames directly)	
TDM Support	CES	CES	
Voice Support	VoIP	VoIP	
Multiple QoS levels	Yes (802.1Q priority levels)	Yes (802.1Q priority levels)	All traffic types are handled using a single LLID.
FEC	RS(255, 239) – frame oriented	RS(255, 223) – stream oriented	
Encryption	Defined by regional standards	Not in the scope of 802.3av. Defined by regional standards	A multi-company protocol using 128-bit key AES has become the de facto standard for EPON. IEEE 802.1ae is anticipated as the encryption standard to be commonly used with 10G EPON.
OAM	802.3ah Ethernet OAM frames	802.3ah Ethernet OAM frames	

5 Conclusions

Due to its very high bandwidth capability, fiber is the most flexible medium for broadband service delivery to the home. After years of being a promising "next generation" technology, FTTH has finally become an economically viable option for providing residential triple-play services. The various technical and operational hurdles that have slowed large-scale deployment of FTTH have largely been resolved.

PON is the most cost-effective approach to providing FTTH. By providing a highly flexible platform for different services, and by eliminating the active electronics from the access plant, PON provides carriers with substantial ongoing OAM savings over copper-based technologies such as DSL or coaxial cable with cable modems. The 10G EPON standard provides a powerful extension to the existing EPON standard to deliver either more bandwidth per subscriber, or potentially serve more subscribers from same PON. A key aspect of 10G EPON is that it can co-exist on the same PON as EPON. This co-existence allows carriers a cost-effective mechanism for upgrading the bandwidth of an existing PON network. It also allows using 10G EPON ONUs to serve enterprise customers with high bandwidth demands on the same PON as lower cost EPON ONUs to serve subscribers with lower bandwidth demands.

PMC-Sierra has developed a full line of ICs for both ONUs and OLTs. The EPON ICs are Chinese Standard Compliant in addition to being IEEE 802.3ah compliant, and the GPON ICs are FSAN compliant. These devices offer high integration and advanced capabilities that enhance the system performance and simplify the design and operation of the PON system. The device features include high throughput (full wire-speed packet processing), packet processing capability for Layers 2-4, a powerful programmable engine for DBA, a "virtual scope" for in-service diagnostics, and FEC. PMC-Sierra played a leading role in the development of the EPON standard and continues to be actively involved in FSAN and IEEE 802.3 work to drive next generation PON requirements.

Appendix A – 64B/66B Line Code

There are two primary requirements for any line code. The first is to ensure an adequate number of line code level transitions within a given time window so that the receiver can align the frequency and phase of its clock and data recovery circuit to the incoming signal. This alignment is critical for making the decision about the value of the received bit/symbol at the point with the least impact from noise, distortion, and inter-symbol interference. The second requirement is that there must be a method of determining the boundaries of the line code characters.

Gigabit/s Ethernet typically uses the 8B/10B line code in which eight data bits are mapped into ten bits for transmission. The 256 8-bit data values are mapped into pairs of 2048 10-bit characters in a manner that maintains a running balance between the transmitted ones and zeros (the running disparity). Some of the 10-bit codes not used for data encoding are used to communicate control information. Special Idle control characters are transmitted between packets that allow the receiver to synchronize to the alignment boundaries the 8B/10B characters. The 25% bandwidth overhead of the 8B/10B code becomes a problem with high-speed signals. Consequently, the more bandwidth efficient 64B/66B line code was developed for 10GE.

The 64B/66B block code consists of eight octets (64 bits) of data and/or control characters preceded by two flag bits that indicate whether the block contains control characters or only data characters. The mapping is performed in a manner that allows mapping all the 8B/10B special control characters into the 64B/66B code. The two leading code bits also provide the means for the receiver to synchronize to the code block boundaries. The code maintains an average balance between transmitted ones and zeros through a scrambler, although not as strictly in the short term as the 8B/10B code.

Figure 11 illustrates the 64B/66B code construction. The following bit and byte label conventions are used. Hexadecimal values are used to represent the values of the block type fields, control characters, and data octets. In Figure 11, the information, including any binary values, is transmitted from left to right with the LSB (bit 0) of the hexadecimal values transmitted first. For example, a block type value of 0x87 = 10001110 is transmitted from left to right as 01110001. The transmitted (received) block bits are designated as TxB<65:0> (RxB<65:0>).

The 64B/66B code can either be a data block or a control block. A synchronization header (sync header) value of 01 indicates a data block where the payload of the 64B/66B code contains only Ethernet frame data characters. A sync header of 10 indicates a control block in which the 64B/66B block contains control information such as Ethernet frame Start or Termination characters, or control characters. A sync header of 10 is immediately followed by an 8-bit field (the Block Type Field or BTF) that specifies the structure of the 64B/66B code payload. The receiver can achieve block code alignment by searching for an adjacent pair of bits that always have complementary values (i.e., a transition between their values). The payload of the 64B/66B code is scrambled, which makes it highly unlikely that any two payload adjacent bit positions will retain complementary values for a significant period of time.

Data blocks contain eight data characters, and control blocks contain eight characters that are either all control characters or a combination of control and data characters, as specified by the BTF. The BTF occupies the first eight bits of the 64B/66B block payload area. Two different techniques are used to fit eight characters into the 56 block payload bits after the BTF. First, control characters are encoded as 7-bit values or 4-bit O-codes. Second, the locations of the Ethernet frame Start and Termination characters are implied by the BTF, and hence don't need to be carried explicitly within the block.

Figure 11 64B/66B block code structure

Block Type	Input Data (Block Format)	Sync	Block Payload								
Data	$D_0D_1D_2D_3D_4D_5D_6D_7$	01	D_0	D_1	D_2	D_3	D_4	D_5	D_6	D_7	
Control			B.T.F.								
	$S_0D_1D_2D_3D_4D_5D_6D_7$	10	0x78	D_1	D_2	D_3	D_4	D_5	D_6	D_7	
	$C_0C_1C_2C_3S_4D_5D_6D_7$	10	0x33	C_0	C_1	C_2	C_3		D_5	D_6	D_7
	$O_0D_1D_2D_3S_4D_5D_6D_7$	10	0x66	D_1	D_2	D_3	O_0		D_5	D_6	D_7
	$O_0D_1D_2D_3O_4D_5D_6D_7$	10	0x55	D_1	D_2	D_3	O_0	O_4	D_5	D_6	D_7
	$C_0C_1C_2C_3O_4D_5D_6D_7$	10	0x2D	C_0	C_1	C_2	C_3	O_4	D_5	D_6	D_7
	$O_0D_1D_2D_3C_4C_5C_6C_7$	10	0x4B	D_1	D_2	D_3	O_0	C_4	C_5	C_6	C_7
	$C_0C_1C_2C_3C_4C_5C_6C_7$	10	0x1E	C_0	C_1	C_2	C_3	C_4	C_5	C_6	C_7
	$T_0C_1C_2C_3C_4C_5C_6C_7$	10	0x87		C_1	C_2	C_3	C_4	C_5	C_6	C_7
	$D_0T_1C_2C_3C_4C_5C_6C_7$	10	0x99	D_0		C_2	C_3	C_4	C_5	C_6	C_7
	$D_0D_1T_2C_3C_4C_5C_6C_7$	10	0xAA	D_0	D_1		C_3	C_4	C_5	C_6	C_7
	$D_0D_1D_2T_3C_4C_5C_6C_7$	10	0xB4	D_0	D_1	D_2		C_4	C_5	C_6	C_7
	$D_0D_1D_2D_3T_4C_5C_6C_7$	10	0xCC	D_0	D_1	D_2	D_3		C_5	C_6	C_7
	$D_0D_1D_2D_3D_4T_5C_6C_7$	10	0xD2	D_0	D_1	D_2	D_3	D_4		C_6	C_7
	$D_0D_1D_2D_3D_4D_5T_6C_7$	10	0xE1	D_0	D_1	D_2	D_3	D_4	D_5		C_7
	$D_0D_1D_2D_3D_4D_5D_6T_7$	10	0xFF	D_0	D_1	D_2	D_3	D_4	D_5	D_6	
Bit Position		0 1	2								65

B.T.F. = Block Type Field

The mapping of data and control characters into the 64-bit code payload is designed around the 32-bit XGMII interface. Specifically, data blocks filled with two four-octet transfers from the XGMII. The data octets and control characters are designated according to their position within the code block, which also corresponds to their position within the eight characters of the XGMII bus transfer. For example, if the block contains eight data octets, these octets are designated D0 to D7, with D0 located immediately after the leading flag (synchronization) bits. The control octets are encoded as 8-bit values on the XGMII interface, but as discussed above are encoded into 7-bit or 4-bit values when encoded into a 64B/66B control block. The set of legal control code encodings are shown in Table 2 and Table 3. Any other control codes are regarded as errors if they are received. The control codes for start of packet (/S/) are designated S0 or S4, and start of an ordered set control character (/O/) is designated O0 or O4 since they are only valid on the first octet or the fourth-octet XGMII position. (Note that while the start and end of packet information is communicated across the XGMII, the ordered set information is typically not communicated over the XGMII.) The end of the packet control code (/T/) can occur anywhere within the code block, and hence can be designated T0 to T7. Similarly, control codes other than /S/, /O/, and /T/ can be located in any position within the block and are hence designated C0 to C7.

When the O0 O-codes are encoded as 4-bit values and placed in the first nibble of the fifth octet of the 64B/66B block payload, and the 4-bit value representing the O4 O-codes are placed in the last nibble of the fifth octet of the 64B/66B block payload. O-codes represent the start of a 4-octet ordered set.

The narrow rectangles in Figure 11 represent single pad bits within the control block payload area. These pad bits contain a value of 0 and are used achieve the desired alignment of the control and data characters when their encoding leads to less than 56 total bits following the BTF.

Table 2 Encoding for non-ordered set control codes

Control Character	Notation	XGMII Control Code	10G-BASE-R 7-bit Control Code	8B/10B Code
Idle	/I/	0x07	0x00	K28.0, K28.3 or K28.5
Start	/S/	0xFB	Implied by BTF	K27.7
Terminate	/T/	0xFD	Implied by BTF	K29.7
Error	/E/	0xFE	0x1E	K30.7
reserved0	/R/	0x1C	0x2D	K28.0
reserved1			0x33	K28.1
reserved2	/A/	0x7C	0x4B	K23.3
reserved3	/K/	0xBC	0x55	K28.5
reserved4		0xDC	0x66	K28.6
reserved5		0xF7	0x78	K28.7
Note – The /A/, /R/, and /K/ codes are XAUI interface codes that can only appear on the XGMII due to bit errors.				

Table 3 Encoding for ordered set control codes

Control Character	Notation	XGMII Control Code	10G-BASE-R 4-bit O-code	8B/10B Code
Sequence ordered_set	/Q/	0x9C	0x0	K28.4
Signal ordered_set	/Fsig/	0x5C	0xF	K28.2

A self-synchronous scrambler is used on the block payload in order to help randomize the payload bits. The sync header bits bypass the scrambler. The resulting effect of the scrambler is to, improve DC balance, reduce the likelihood of payload bits mimicking the sync header bits, and increase the average number of data value transitions, which aids the receiver data recovery circuit. Self-synchronous scramblers avoid the need for initialization or periodic explicit synchronization by feeding the data back on itself through exclusive-OR gates, as illustrated in Figure 12. The descrambler uses a circuit similar to the scrambler in order to reverse the process. The scrambler polynomial is: $G(x) = 1 + x^{39} + x^{58}$.

Figure 12 64B/66B block code self-synchronous scrambler

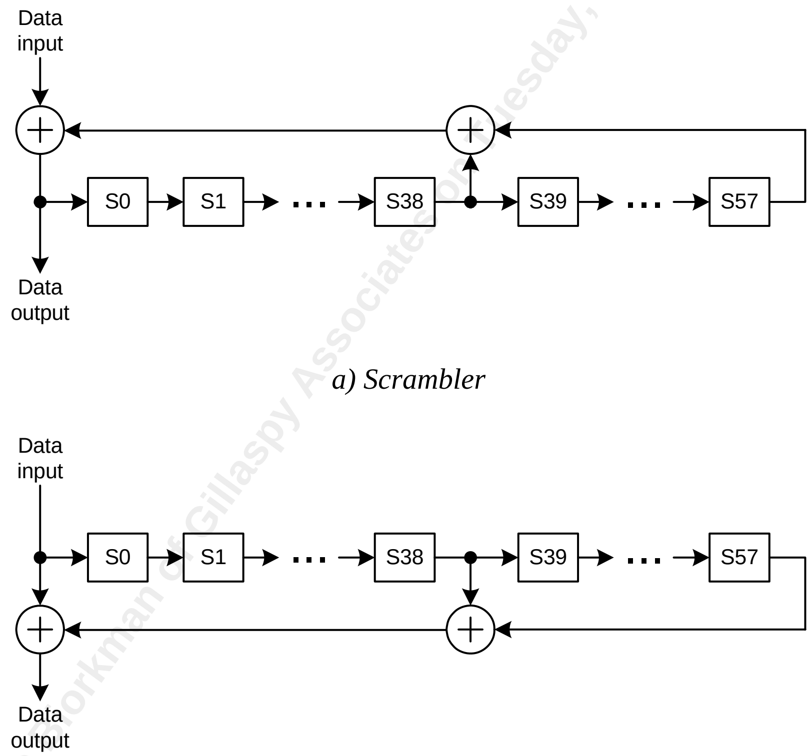

a) Scrambler

b) Descrambler

6 References

[1] IEEE 802.3av (draft) Amendment: Physical Layer Specifications and Management Parameters for 10 Gb/s Passive Optical Networks

[2] PMC-2061015 "Introduction to Broadband Passive Optical Networks (PON)," PMC-Sierra white paper by Steve Gorshe

[3] IEEE 802.3 (2005) Information technology -Telecommunications and information exchange between systems - Local and metropolitan area networks – Specific requirements - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications

[4] Federal Information Processing Standard 197, Advanced Encryption Standard, National Institute of Standards and Technology, U.S. Department of Commerce, November 26, 2001.

[5] O. Haran and O. Hiironen, "CTR Mode for Encryption," http://grouper.ieee.org/groups/802/3/efm/public/jul02/p2mp/haran_p2mp_2_0702.pdf

[6] ITU-T G..983.4 (2001) A broadband optical access system with increased service capability using dynamic bandwidth assignment (DBA)

[7] ITU-T G.983.5 (2002) A broadband optical access system with enhanced survivability

[8] ITU-T G.984 Recommendation series on Gigabit/s PON

7 Glossary of Acronyms

10GE	10 Gbit/s Ethernet
10G EPON	PON based on 10GE (IEEE 802.3av)
APD	Avalanche Photo Diode
AES	Advanced Encryption System
CATV	Community Access TV
CBR	Constant Bit Rate
CES	Circuit Emulation Service
CO	Telephone company Central Office
CTS	Common Technical Specification for GPON
DBA	Dynamic Bandwidth Assignment
DSL	Digital Subscriber Line
EFM	Ethernet in the First Mile
EPON	Ethernet-based PON (IEEE 802.3ah)
FEC	Forward Error Correction
FCS	Frame Check Sequence
FITL	Fiber in the Loop
FSAN	Full Service Access Network consortium
FTTx	x = B (business), C (curb), Cab (cabinet), H (home), P (premises), N (node)
GPON	Gigabit PON (ITU-T G.984 series)
IPTV	Internet Protocol Television
LLID	Logical Link layer ID
MAC	Media Access Control
MPCPDU	Multi-Point Control PDU

NE	Network Element
OAM	Operations, Administration, and Maintenance
OAM&P	Operations, Administration, Maintenance, and Provisioning
OLT	Optical Line Terminal
ONT	Optical Network Terminal
ONU	Optical Network Unit
PDU	Protocol Data Unit
PON	Passive Optical Network
POTS	Plain Old Telephone Service
RT	Remote Terminal
RTD	Round Trip Delay
NSR	Non Status Reporting
NRZ	Non-Return to Zero line code
RS	Reed-Solomon (FEC code)
SFD	Start of Frame Delimiter
SR	Status Reporting
TC	Transmission Convergence layer
TDM	Time Division Multiplexed
TDMA	Time Division Multiple Access
TIA	Transimpedance Amplifier
US BW Map	Upstream Bandwidth Map
VDSL	Very high-speed DSL
VoD	Video on Demand
VoIP	Voice over Internet Protocol
WDM	Wave Division Multiplexing

XGMII	10 Gigabit Ethernet Media Independent Interface

Contacting PMC-Sierra

PMC-Sierra
100-2700 Production Way
Burnaby, BC
Canada V5A 4X1

Tel: +1 (604) 415-6000
Fax: +1 (604) 415-6200

Document Information: document@pmc-sierra.com
Corporate Information: info@pmc-sierra.com
Technical Support: apps@pmc-sierra.com
Web Site: http://www.pmc-sierra.com