The XG-PON System: Cost Effective 10 Gb/s Access

Frank J. Effenberger, *Senior Member, IEEE*

***Abstract*—The ten gigabit passive optical network (XG-PON) system is the newest member of the ITU-T family of passive optical network standards. XG-PON is the result of a 3 year project involving the full service access network (FSAN) group and ITU-T study group 15 (SG15) question 2. This paper reviews the deliberations that led to the selection of the XG-PON system, and then explains the three primary layers of the system: physical, protocol, and management. The paper concludes with information on standards and implementations of the system, and on future work in this area.**

***Index Terms*—Broadband communication, optical fiber communication, optical subscriber loops, passive optical network.**

I. INTRODUCTION

BEGINNING in late 2006, the FSAN group began to contemplate the system that would follow after gigabit PON (G-PON). Initially, the focus of this work was to develop any additional specifications for the G-PON system that would enable a smoother migration to whatever system came later. This work resulted in the G.984.5 recommendation, which refined the spectrum plan for G-PON and defined the blocking filters in the G-PON optical network units (ONUs), which prevent crosstalk from non-GPON wavelengths (such as those used by XG-PON) [1]. With this preliminary task completed, the way was clear for consideration of next generation PONs.

A. *Scope of the Next Generation PON (NG-PON) Systems*

In late 2007, the focus moved towards defining the new system itself. At first, a very wide range of architectures were raised as possible candidates, including TDM-PONs, WDM-PONs, CDMA-PONs, and others. This posed a problem in that many of these systems are quite different in architecture and service profile, so it was difficult to compare them in a reasonable and objective way. The solution to this was to divide the system proposals into two groups, as shown in Fig. 1 [2]. The first group (NG-PON1) included systems that could coexist with G-PON on the very same optical distribution network (ODN). The second group (NG-PON2) included all other systems that either required a different ODN, or that required technologies that were not available in the expected time horizon. This key decision on the scope of NG-PON1 enabled the comparison of a reasonable set of alternative systems.

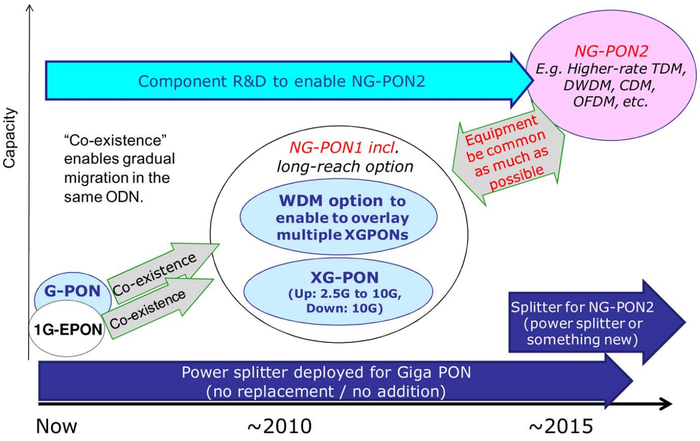

Fig. 1. Definition of scope of NGA1 and NGA2 (circa 2007).

B. *Candidate Systems and Selection*

The FSAN group gathered options for NG-PON1 through early 2008. The candidates are explained in detail in [3]. In brief, they included

- Physical split reduction: This is simply the application of the G-PON with a smaller split ratio, to increase the bandwidth per customer.
- WDM bidirectional split reduction: This system uses four wavelengths in downstream and upstream, creating 10 Gb/s downstream and 5 Gb/s upstream.
- WDM downstream-only split reduction: This system uses four wavelengths in the downstream, and only one in the upstream, for a bandwidth of 10/1.25 Gb/s.
- XG-PON1: 10 G down, 2.5G up: This system uses one wavelength in each direction, with the bandwidth as the name suggests.
- XG-PON2: 10 G symmetric: This system is a version of the XG-PON with symmetrical 10 Gb/s bandwidth.
- Reach enhanced versions of the XG-PONs: This system aimed at higher optical capabilities though the use of more tightly controlled ONU optics.
- Hybrid DWDM/XG-PON: This system aimed at multiplexing many XG-PONs on a single feeder fiber through the use of wavelength seeding.

All of these systems were considered on their merits, with the objective of selecting one system that could meet the requirements of most at the lowest cost and lowest risk.

Some systems could be eliminated on the basis of not maintaining compatibility with the existing ODN, such as the physical split reduction and the WDM downstream-only split reduction system. Some systems were eliminated because they were too forward looking and had too much technical risk, such as the reach enhanced XG-PONs and the hybrid DWDM/XG-PON. Some systems were eliminated because they were not forward

---

Manuscript received June 15, 2010; revised September 19, 2010; accepted September 29, 2010. Date of publication October 07, 2010; date of current version February 08, 2011.

The author is with Huawei Technologies, Santa Clara, CA 95050 USA (e-mail: feffenberger@huawei.com).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2010.2084989

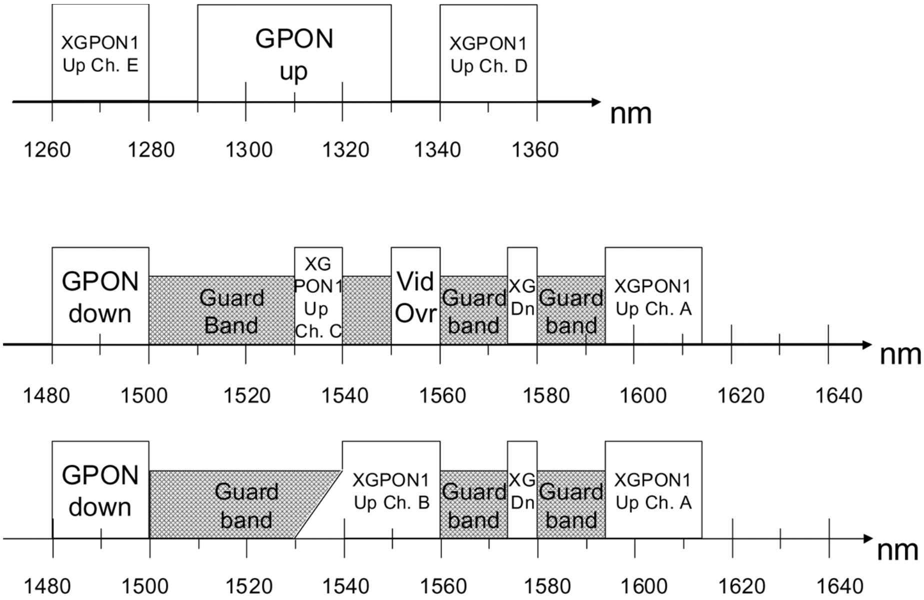

Fig. 2. Wavelength plan options.

looking enough, such as the bidirectional WDM split reduction system. Lastly, after some time, some systems were eliminated because of concerns of high costs, such as the XG-PON2 system. Ultimately, the FSAN group decided that XG-PON1 looked like the best candidate for standardization development.

C. Requirements, Old and New

Alongside the consideration of the candidates, the FSAN operators considered what the system requirements for XG-PON would be. The basic position is that XG-PON would inherit all the requirements of G-PON, with a few additions.

Of course, the primary new requirement for XG-PON that is not found in the G-PON system is that XG-PON must coexist with the G-PON system. This key requirement has far-reaching implications on the physical and protocol layers. These will be explained in the later sections of this paper.

One major new feature is the inclusion of more security. In the original G-PON, the threat model assumed that the upstream channel was physically secure, and this motivated a relatively weak security arrangement. (This was later strengthened in optional amendments to G-PON.) In XG-PON, the PON system is required to support the option of strong mutual authentication, and to use the authentication to protect the integrity of the PON management messages and the PON encryption keys. These enhancements make it quite difficult for an attacker to masquerade as either an ONU or an optical line terminal (OLT), even if he has access to the PON fibers, and even if he can precisely interleave his transmissions with the victim ONU.

Another new feature is the support of equipment power saving. The primary goal of power savings is to reduce the load during power failures such that a given size battery will last longer. The second goal is to reduce power at all times so that consumption of electricity is reduced to the greatest extent possible. The first method of power saving is to turn off those user network interfaces (UNI) that are not actively used. This is quite effective; however, there are sometimes difficulties in determining if a UNI is truly unused. The second mode of power savings is achieved by deactivating the transmitter for routine PON transmissions, when the user has no real data to send (this has been given the name “Dozing”). The third level of power saving is when the ONU deactivates both its transmitter and receiver when the user has no activity (this is called “sleeping”). This last form, while promising to have the lowest power consumption, has the issue that new network-side activity cannot be signaled to the ONU immediately.

II. PHYSICAL LAYER

The physical layer, also commonly referred to as the physical media dependent (PMD) layer, was a subject of considerable debate during the development of the XG-PON specification. While it may seem at first blush to be a low-level engineering issue, it is in fact a design feature that has tremendous impacts on operator-visible features, such as compatibility with other systems, and with fiber plants. This was the driver for its intense consideration.

A. Wavelength Plan

The first topic of major debate was the wavelength plan. The key driver was the coexistence with existing systems (G-PON and the video overlay). First of all, there was the selection of coexistence method. The use of TDM or WDM was considered. After evaluating the deployment scenarios, it was decided that WDM coexistence would be used for both the downstream and upstream paths of the system. This meant that XG-PON would require two wavelength bands that were sufficiently isolated from all other wavelengths. Fig. 2 illustrates the wavelength bands considered. For the downstream, there was relatively little question that a 6 nm band around 1578 nm would be used. This agreed with the wavelength choice already made in the IEEE P802.3av project (10GEPON), and would thereby benefit from economies of scale.

The upstream; however, was not as obvious. Five channel options were considered (labeled A through E in the figure) [4]. Once again, a selection process was undertaken to decide which channel would be used. Channel A (1595 to 1615 nm) was rejected due to fears that the fibers and passive components were not sufficiently specified at those wavelengths. Channel B (1540 to 1560 nm) was dropped because it is incompatible with the video overlay, and many significant PON deployments in the world use this overlay. Channel C (1530 to 1540 nm) was eliminated because such ONUs would be costly, and because existing G-PON ONUs may not have sufficient isolation against this wavelength. Channel D (1340 to 1360 nm) was not selected because the coexistence filter it would require would be quite difficult to make with a low loss. Thus, at the end, channel E (1260 to 1280 nm) was selected as the upstream wavelength for the XG-PON1 system.

B. Line Rates, Codes, and Coexistence

With the wavelengths in hand, the next consideration was made of the exact line rate and code. For the downstream, there were two solutions considered, primarily because they had already been standardized for other applications and were commercialized already. The first was the 9.95328 Gb/s synchronous digital hierarchy (SDH) rate, presumed to employ non-return to zero (NRZ) coding. The second was the 10.3125 Gb/s Ethernet rate, presumed to use 64b66b block coding. The choice of these two systems raised an important issue regarding to coexistence of XG-PON. As covered above, XG-PON had to be compatible with G-PON and video overlay systems. It happens to also coexist with most 1G EPON systems (due to the use of WDM). The question then became: should XG-PON also try to coexist with 10GEPON?

If this complete coexistence situation was desired, then the choice of line code was directed quite simply to use the Ethernet line rate and code. However, the operators made a fateful decision that this kind of intra-generational coexistence would not be needed. It was considered unlikely that an operator would deploy both 10GEPON and XG-PON in the same network. Therefore, the somewhat more complicated option of using the Ethernet line coding system to carry the XG-PON downstream was dropped, and the SDH-based rate and code was selected.

In the upstream, three options were considered. The first was to operate the line at 3.125 Gb/s, which is 25% faster than the desired 2.5 Gb/s payload. This over-rating would be used for line conditioning coding (9b10b), and for forward error correction (FEC). The second was to use a line coding of 2.577 Gb/s (which is 1/4 of the 10.3125 Gb/s downstream). The third would run the line at 2.488 Gb/s (which is 1/4 of the 9.985 Gb/s downstream). Of these choices, the first was eliminated because most of the vendors thought that it would be easier to obtain optics that ran at the standard 2.5 Gb/s rate, albeit that they would need special burst mode design. Then, given the selection of the downstream rate of 9.985 Gb/s, the 2.488 Gb/s upstream rate was the natural choice.

C. Power Budgets

The next item to tackle was the power budget. This topic should have been quite direct to solve, because the XG-PON was supposed to share the same optical distribution network as G-PON, and the G-PON standards quite clearly specify what the ODN's characteristics are (28 dB of loss in the windows from 1260 to 1360, and from 1480 to 1580 nm). However, there were two effects that made the direct reuse of the same specifications impossible.

The first is the introduction of the so-called WDM1r filter. This filter, which is specified in G.984.5, is the interconnection point of the G-PON OLT, the XG-PON OLT, and the ODN. As with any practical WDM filter, it has some loss, and this loss needs to be accounted for. The GPON systems are already deployed, and so every effort must be made to build the WDM1r so that its loss in the G-PON path is minimized. In some respects, this is a zero sum game, so that if loss is optimized for G-PON, then the loss for XG-PON will be increased. It was estimated that the loss differential would be about 0.5 more for XG-PON for this reason.

The second is that some real deployments in the world have tended to deviate from the standard. One reason is that commercially attractive optics were developed that have a bit more loss margin than the standard required. It was determined that the common value actually in use was 29.5 dB. Another factor was that operators were designing their ODNs using the 1310 nm value of the fiber loss, and not the 1260 nm value of the fiber loss (as the standard would dictate). There is approximately an 0.05 dB/km loss differential between these two wavelengths. Thus, a 20 km PON might see a 1 dB higher loss at 1260 nm than at 1310 nm. Operators were keen to avoid re-engineering their ODNs, and so additional loss margin was added to the XG-PON budgets.

Given these factors, the FSAN operators identified two loss budget that were considered "nominal" PON budgets. The so-called nominal 1 budget is 29 dB, and allows for XG-PON coexisting with standardized G-PON (and EPON) systems. The nominal 2 budget is 31 dB, and allows for coexistence with the super-standard 29.5 dB G-PON systems.

On top of this link budget question, there was the issue of detector type at the ONU. The ONU is very cost sensitive, and every possibility to reduce its cost must be used. In general, PIN type photodetectors are less expensive than APD types, and this makes them attractive. On the other hand, APDs are far more sensitive than PINs, and this requires a less powerful OLT transmitter (a benefit for several practical reasons). As it developed, this choice had proponents on either side, and it was impossible to make a definite conclusion. The standard therefore specifies both the APD and PIN solution for the nominal 2 budget. The prospective is that the industry will decide this issue in the commercial market over the next few years. When a winner becomes clear, then the standard can be revised to reflect the implementation reality.

Lastly, there is the issue of extended loss budgets. In the G-PON system, an extended loss budget was developed that had two major features: 4 dB more loss than the nominal budget, and ONU specifications that were unchanged from the nominal budget. After due consideration of the technical issues, it was determined that these same design features could be reused in XG-PON. Since there are two nominal budgets (29 and 31 dB), there are also two extended budgets (33 and 35 dB), and both

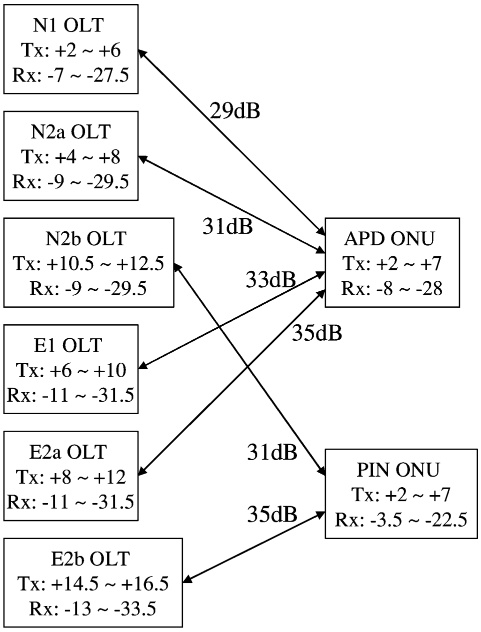

Fig. 3. The power budgets for XG-PON (all values in dBm).

PIN and APD variants are supported. The result is a set of four loss budgets that are implemented with two ONU variants, and 6 OLT variants, as outlined in Fig. 3.

III. Protocol Layer

Given the decision on the line rate, code, and coexistence, it was fairly clear that the basis of the XG-PON protocol layer (properly termed the transmission convergence layer) would be based on G-PON. However, unlike G-PON, which was developed in a more holistic way, the XG-PON protocol was more heavily structured with three distinct sublayers being defined. The first was the physical (PHY) adaptation layer, which handles the unique issues of the XG-PON physical layer. The second was the framing layer, which does the main work of transmission convergence (that is, the control of the PON TDMA system). The third was the client adaptation layer, which works to carry user signals over the XG-PON system. These sublayers are illustrated in Figs. 4 and 5.

A. PHY Adaptation Sublayer (PAS)

This layer takes care of the low level coding of the TC frame over the physical channel. One of the most important design features of the XG-PON PHY is the use of forward error correction (FEC). This is a required feature in both directions (although it can be deactivated in the upstream, if the link quality is good enough.) So, much of the work of the PHY adaptation layer concerns FEC. The FEC used in the downstream is the RS(248,216) code. In each 125 microsecond frame time, 24 bytes of physical synchronization block downstream (PSBd) is set aside for the functions of the PAS. In the downstream, these functions include:

- Framing: The first 64 bits of the PSBd are set to a fixed framing pattern that the receiver can use to find the 125 microsecond frame.
- Super frame counter: The second 64 bits hold a super frame counter, that provides a much longer scale time reference, and also a scrambler preload.
- PON identification: The third 64 bits hold a value settable by the OLT that can serve to identify this particular PON signal (useful in field operations).

The remainder of the downstream frame is exactly 627 codewords of the FEC code.

The upstream of the XG-PON is burst-transmission oriented, and so it has a few differences from the downstream. The first is that the physical synchronization block upstream (PSBu) contains the preamble and delimiter patterns. The second is that the payload is not a fixed size, and for this reason, the number of FEC codewords in each burst is variable. In addition, to reduce the loss of bandwidth for odd-sized bursts, the last codeword in the burst is shortened to fit the available time in the burst.

In both directions, the FEC encoded payload is scrambled using a frame-synchronous cyclic shift register based scrambler. Unlike other systems, this scrambler uses the super-frame counter as the basis for its preload. This, coupled with the scramblers greatly increased size (58 bits), makes it very difficult for an attacker to guess the scrambler pattern and to knock the PON out of service.

B. Framing Sublayer (FS)

This layer takes care of the TDMA aspect of the PON, including activation and normal operation phases, as well as house-keeping functions. The XG-PON transmission convergence (XGTC) downstream header contains three parts. The first is the fixed size part of the header, and carries the lengths of the next two parts of the header (protected with a header error correction (HEC) code). The second part is the bandwidth map, which carries several bandwidth allocations to the ONUs on the PON. The third part carries physical layer operations administration and management (PLOAM) messages to the ONUs on the PON. Following this header, the remainder of the downstream carries the payload.

The concept of the bandwidth map is largely modeled after that found in G-PON, but there are some improvements. Just as in G-PON, each bandwidth allocation is an instruction for a particular ONU to transmit in the upstream, and consecutive allocations to the same ONU can be concatenated together for added efficiency. The XG-PON allocations are now described with a start-time and payload-length concept, rather than the start-time and stop-time in G-PON. The important difference is that the payload-length is given before FEC overheads are added. This makes the calculation of concatenated allocations much easier, and it removes many invalid allocation possibilities that were possible in the G-PON method. The XG-PON bandwidth allocation-identification (ID) address space has been expanded by a factor of 4, providing for wider split PONs. Lastly, each allocation specifies a burst profile to use in this burst. The burst profile includes the length and pattern of the preamble, the delimiter, and whether FEC is active. In this way, each transmission on the PON can be customized to fit its particular situation.

The PLOAM messages are also modeled after those in G-PON, with improvements. One change is that more than one PLOAM message can be sent per downstream frame. This makes the channel much more responsive, which is useful in power saving applications. Another improvement is that the

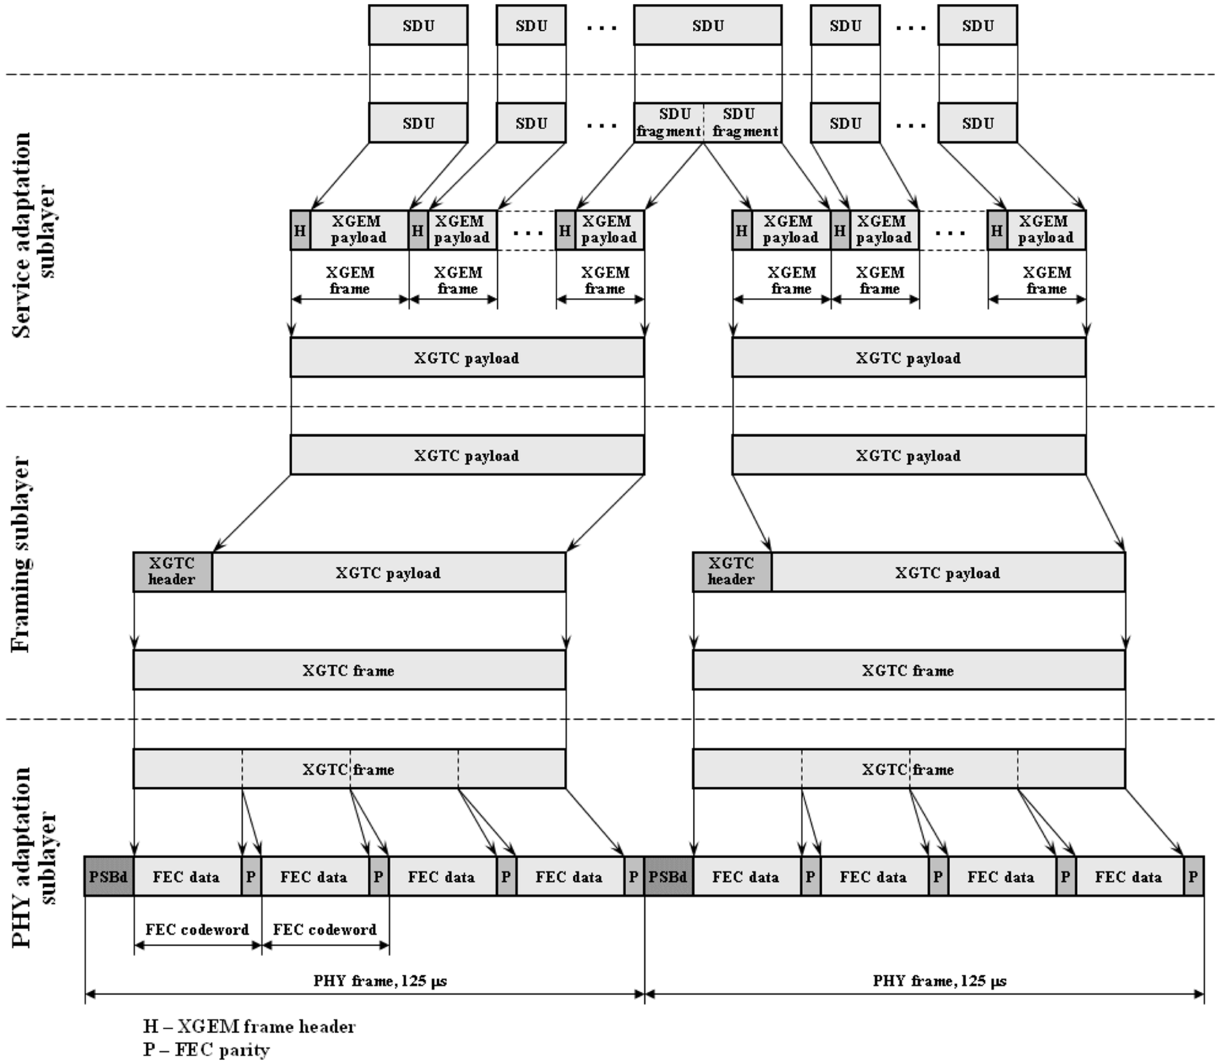

Fig. 4. Downstream XG-PON TC layer organization.

PLOAM messages have been made larger so that they accommodate the known message sets without fragmentation. Of course, this flexibility has raised the possibility of overrunning the ONUs with messages that are too fast. So, some limits on the maximum permissible rate to each ONU have also been established.

In the upstream, each burst also has headers of multiple types. At the beginning of each burst, there is a fixed burst header and a variable burst header. The fixed burst header contains the ONU-ID number, as well as the echo of the control information from the allocation. The ONU-ID number has also been expanded by a factor of 4 over G-PON, to support a wider split PON (1023 ONUs are supported). The variable burst header carries the upstream PLOAM message (if any). At the beginning of each allocation, there is also an optional allocation header, and this carries the dynamic bandwidth report upstream (DBRu).

C. Client Adaptation Layer (CAS)

This sublayer takes the user's payload (data packets) and formats them for transmission over the PON. The generic name for this is the XG-PON encapsulation method (XGEM). There are three major aspects to be taken care of in XGEM. The first is that individual flows of traffic (termed "ports" in XG-PON) must be marked so that they can be accepted by the appropriate client on the other side of the PON. This is done by using a 16-bit Port-ID. This is an expansion of the address space of 16 times over G-PON, and again can support wider split PONs.

The second XGEM function is that of fragmentation. The framing header must occur at its regular periodic time, and so a user packet might straddle this boundary. In the upstream, the burst may end before the current packet can be completed. The XGEM system allows for packets to be fragmented so that the first part is transmitted in the current PON frame or burst, and the second part is transmitted at the next opportunity. The rules regarding the generation of fragments were enhanced over G-PON, so that very small fragments are avoided. This makes implementations easier.

The third XGEM function is data privacy. Each XGEM fragment has a key index associated with it, and this index selects a key that has been previously negotiated between the ONU and OLT. The key index allows for a very well defined key switch-over, and keys can be changed in XG-PON with no loss of data at all. This, coupled with the strong mutual authentica-

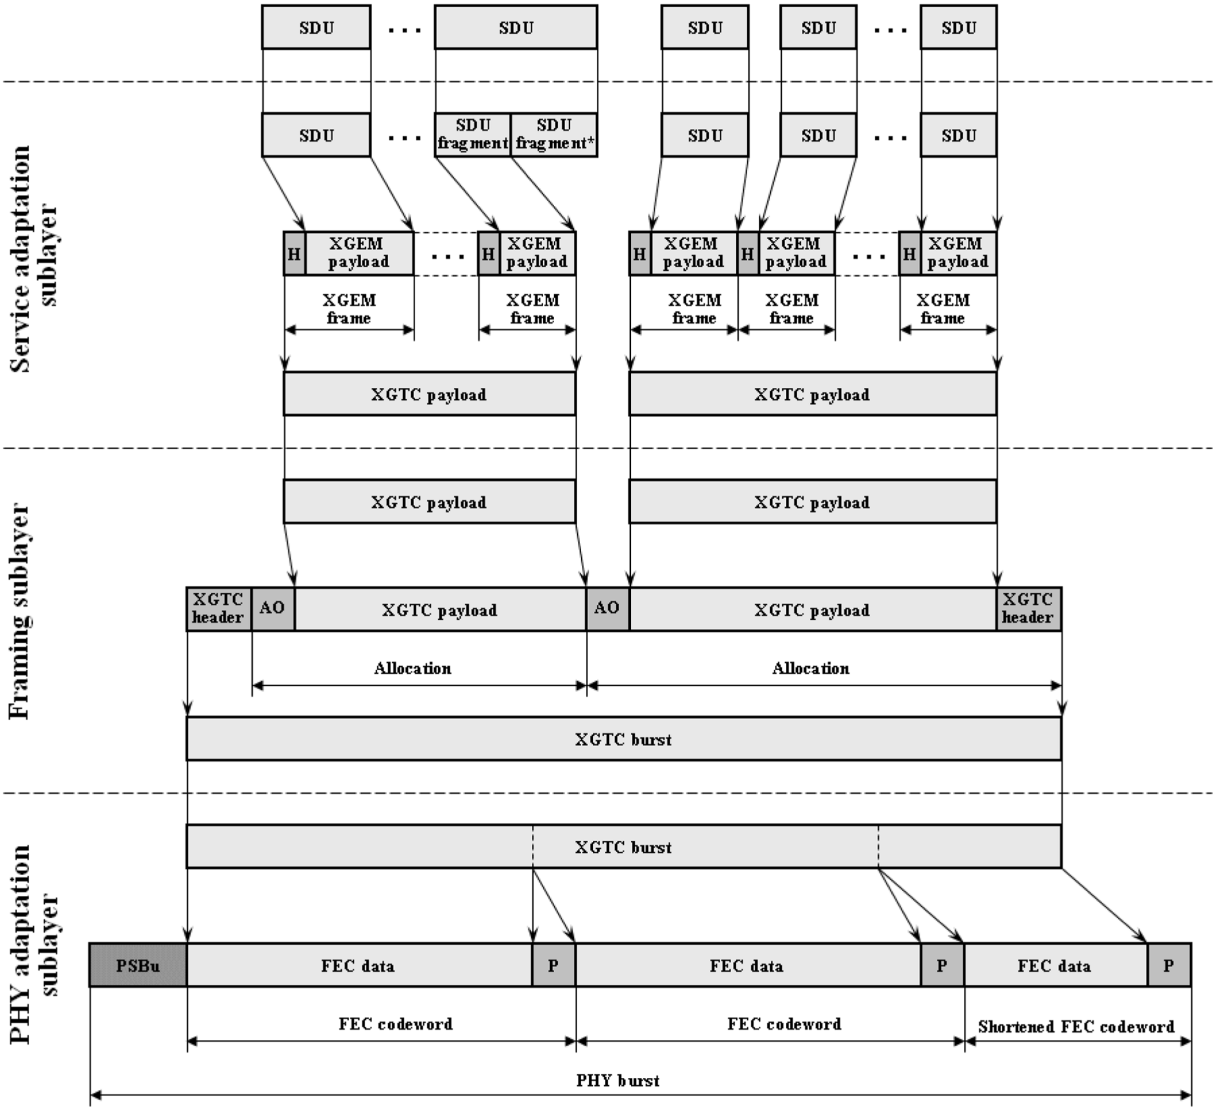

(*) The remaining fragment of the SDU is transmitted in the subsequent allocation with the same Alloc-ID.
H – XGEM frame header
AO – Allocation overhead
P – FEC parity

Fig. 5. Upstream XG-PON TC layer organization.

tion, makes the XG-PON system as secure as any of the other major broadband communications system.

IV. MANAGEMENT AND SERVICE LAYERS

One of the unique capabilities of the ITU-T family of PON systems is the integrated management system, and the definition of service layers. This is what makes the ITU-T optical access standards most suitable for telecom deployment. The following sections describe the development of these features in the ITU and the broadband forum. XG-PON has been lined up in such a way that all the progress that has been achieved with G-PON can be directly inherited, with few modifications.

A. *Generic ONU Management and Control Interface (OMCI)*

If we consider the final goal of standardization to be interoperability, then every interface to the ONU must be described in detail. The previous sections discussed the physical interface and the protocol interface, but this leaves the most complex interface unmentioned. It is the OMCI that contains by far the most complexity (measured in total function points, in variability, and in evolution over time). So, in this sense, the definition of the OMCI is the most difficult and time consuming work in all the PON standardization.

This work was begun in G.983.2, which described the OMCI for B-PON. This work grew over time into a sizable collection of management features for nearly every feature that the industry could build into an ONU. When G-PON began, its OMCI used G.983.2 as a base, and only added G.984.4, which described the small changes needed to adapt G.983.2 to G-PON systems. The situation where both G.983.2 and G.984.4 were developed in parallel continued for a time, but eventually G.984.4 was revised to include all the relevant parts of G.983.2, and the G-PON

OMCI became an independent recommendation. The G-PON OMCI continued to grow and adapt to all the new services and features that PONs were gaining.

When the time came to describe the XG-PON OMCI, the decision was made to avoid repeating the migration process from B-PON to G-PON. Rather, the OMCI standard would be made into a generic OMCI recommendation, G.988. In this way, both G-PON and XG-PON, and any other technology that wanted to use it, could refer to this common document directly. The generic OMCI recommendation would never need to be revised due to a technology shift or other physical change, since it is a general document.

Since this organizational shift was made, interest has been found in using OMCI for other systems. The point to point Gigabit Ethernet systems described in G.986 use OMCI for basic ONU management functions. There is even some interest in using OMCI for some parts of IEEE EPON management, although that has not reached a consensus yet. Nevertheless, the OMCI recommendation continues to grow, and is the most compact and complete reference for ONU management in the world.

B. Broadband Forum TR-156/167

The material developed so far deals with getting the OLT to work with the ONU. It is equally important to get the PON system to work with the network at large. The group that has done the most work in this area has been the Broadband Forum (BBF), and in particular their technical report 101 (TR-101) has been the guiding force behind a large part of the deployment of digital subscriber line (DSL) deployment.

Starting in 2008, work was begun to extend TR-101 to cover G-PON. The result was TR-156 (and TR-167, later). These documents give very detailed specifications for what a practical network application of G-PON technology would be, including such things as quality of service, the arrangement of VLANs, and even the nomenclature of equipment slots and ports for management purposes. These documents have been used to accelerate the interoperability testing of G-PON, and to great success.

Since XG-PON uses the same internal connection models, and especially the same management system, it allows the direct reuse of all the same agreements from TR-156 and TR-167. In fact, the implementing language is already being included in a small revision to these documents.

V. CONCLUSION

The standardization of the core XG-PON system is nearly complete. The terminology and references (G.987), service requirements (G.987.1) and the physical layer (G.987.2) were completed late in 2009 [5]–[7]. The transmission convergence (G.987.3) and the generic OMCI are being consented in June 2010, with final agreement occurring over the next month or two [8], [9]. With these five documents in place, it is possible for implementers to construct fully standards compliant equipment.

Of course, several equipment vendors have already moved ahead with prototypes that anticipate the standard. While these are not fully compliant, they do serve the purpose to test some part of the system, and allow the users to evaluate the suitability of XG-PON. The first field trial for the XG-PON system was completed in Dec. 2009 [10], and this trial confirmed that the system delivered 10 Gb/s performance, and coexisted with both G-PON and video overlay systems.

As with any major access system, there are always further improvements to be considered. One improvement is the description of midspan reach extenders, which can be used to increase the distance or the split ratio of the PON system. This work has begun in ITU, and a draft based on the reach extenders for G-PON has been prepared.

Another improvement is the consideration of higher upstream speeds. One interesting possibility is the use of non NRZ coding to perhaps get more performance out of the existing XG-PON optics. This could deliver 5 Gb/s for not much more cost than the current 2.5G optics, and it would of course coexist with the 2.5G system (and G-PON). In this and other ways, XG-PON promises to be the solid future for ITU PON systems well into the future.

REFERENCES

[1] ITU-T G.984.5, Gigabit-Capable Passive Optical Networks (G PON): Enhancement Band 2007.

[2] J.-I. Kani, F. Bourgart, A. Cui, A. Rafel, M. Campbell, R. Davey, and S. Rodrigues, "Next-generation PON part I—Technology roadmap and general requirements," *IEEE Commun. Mag.*, vol. 47, no. 11, pp. 43–49, Nov. 2009.

[3] F. J. Effenberger, H. Mukai, S. Park, and T. Pfeiffer, "Next-generation PON part II—Candidate systems for next generation PON," *IEEE Commun. Mag.*, vol. 47, no. 11, pp. 50–57, Nov. 2009.

[4] F. J. Effenberger, H. Mukai, J.-I. Kani, and M. Rasztovits-Wiech, "Next-generation PON part III—System specifications for XG-PON," *IEEE Commun. Mag.*, vol. 47, no. 11, pp. 58–64, Nov. 2009.

[5] *10-Gigabit-Capable Passive Optical Network (XG-PON) Systems: Definitions, Abbreviations, and Acronyms*, ITU-T G.987, 2009.

[6] *10 Gigabit-Capable Passive Optical Network (XG-PON): General Requirements*, ITU-T G.987.1, 2009.

[7] *10-Gigabit-Capable Passive Optical Networks (XG-PON): Physical Media Dependent (PMD) Layer Specification*, ITU-T G.987.2, 2009.

[8] *10-Gigabit-Capable Passive Optical Network (XG-PON) Systems: Transmission Convergence Layer Specification*, ITU-T G.987.3.

[9] *ONU Management and Control Interface Specification (OMCI)*, ITU-T G.988.

[10] S. Jain, F. Effenberger, A. Szabo, Z. Feng, A. Forcucci, W. Guo, Y. Luo, R. Mapes, Y. Zhang, and V. O'Byrne, "World's First XG-PON Field Trial," in *Proc. OFC/NFOEC*, May 2010, pp. 1–3.

Frank J. Effenberger (M'94–SM'01) received the B.E. degree from Stevens Institute of Technology, Hoboken, NJ, in 1988, the M.S. degree from the University of Rochester, Rochester, NY, in 1989, and the Ph.D. degree from the University of Central Florida, Orlando, in 1995.

He was a staff scientist at Bellcore where he analyzed all types of access network technologies. He witnessed the early development of the FSAN initiative and the development of the APON standard. In 2000, he moved to Quantum Bridge, where he led the system engineering group. This work supported the development and standardization of advanced optical access systems based on B-PON and G-PON technologies. In 2006, he became Director of FTTx in the advanced technology department of Huawei Technologies, Santa Clara, CA. He remains heavily involved in the standards work, and has been a leading contributor and editor of the major PON standards in the ITU. He is working on forward-looking fiber access technologies, including the 802.3av 10 G EPON and ITU NGA topics. In 2008, he became the chairman of ITU-T Q2/15.