Bandwidth and Latency Analysis of 25GS PON

An Overview of 25GS Passive Optical Networks

A Technical Paper prepared for SCTE by

Edward Boyd
CTO
Tibit Communications, Inc
459 N Gilbert Rd, Suite A255, Gilbert, AZ 85234, USA
ed.boyd@tibitcom.com

Edward Walter
Director Member Technical Staff
AT&T, Inc.
123 Timber Mountain Dr., Boerne, TX 78006, USA
Ew8532@att.com

Fernando Villarruel
Chief Architect, MSO Practice
Ciena Corp
1120 Sanctuary Pkwy, Alpharetta, GA, 30004, USA
fvillarr@ciena.com

Table of Contents

List of Figures

List of Tables

1. Introduction

The fiber access market is largely dominated by passive optical networks (PON). A PON is a point-to-multipoint network supporting multi-service traffic across a single fiber that can be split up to 1:128 in the outside plant and run end-to-end passively (no amplifiers required). The passive splitting feature conserves trunk fiber and minimizes equipment in the central office, node, and headend.

The IEEE and ITU-T define standards for PON at multiple speeds. Ethernet PON (EPON) is specified by the IEEE 802.3 at 1 Gigabit per second (Gbps) in the upstream and downstream direction. The IEEE has also specified a 10 Gbps version called 10G-EPON deployed in Japan and with North American MSOs. The ITU-T defined the Gigabit PON (GPON) that provides 1 Gbps upstream and 2 Gbps downstream. Many operators in China deployed with GPON and selectively upgraded to a higher speed version called NGPON1 with 10 Gbps downstream and 2.5 Gbps upstream. The ITU-T later defined a 10 Gbps downstream and 10 Gbps upstream version of GPON called XGS-PON. Operators in Europe, North America, and other parts of the world started with GPON and upgraded to XGS. XGS has begun to dominate PON deployment and will be the favored solution for residential PON into the future.

Since user demand for bandwidth will continue to increase, it is inevitable that PON speeds beyond 10 Gbps will be needed. The IEEE most recently added a 25 Gbps version with the option of 50 Gbps by optically combining two 25 Gbps signals on the same fiber. Using the IEEE's 25 Gbps optical layer definition, a group of companies formed a Multi-Source Agreement (MSA) to create a 25 Gbps version of ITU-T's GPON. The ITU-T standard has been actively defining multiple versions of 50 Gbps PON. 50 Gbps downstream with 12.5 Gbps upstream or 25 Gbps upstream has been created. A 50 Gbps upstream is currently being defined by the ITU-T. Unlike the IEEE 50 Gbps, the ITU-T uses a single wavelength. Cable Labs has also started a 100 Gbps Coherent PON standard as well.

With multiple operators announcing field trials and devices becoming available, we wanted to write a paper on 25 Gbps PON to explore the bandwidth, latency, upgrade, and cost impacts. The three authors of this paper represent a component supplier, a system provider, and a major operator. From the component level, we want to understand the cost and technology differences between the 10 Gbps PON technologies, 25 Gbps PON technologies, and the future 50 Gbps/100 Gbps PON standards. From a system level, we will explain the coexistence and upgrade paths from lower PON speeds to 25 Gbps. From an operator perspective, we want to explore the use cases for a 25 Gbps symmetric PON solution. Finally, this paper explains the expected bandwidth and latency possibilities with a 25 Gbps PON system. PON has very significant overhead from forward error correction, framing, and burst overheads that will drop the 25 Gbps line rate. In the upstream direction, tradeoffs must be made between low latency and high bandwidth. This paper will present some simple models to explore those tradeoffs.

Since analyses of both the ITU-T and IEEE standards would be prohibitive for a single document, this paper focuses on the 25GS MSA standard that uses the ITU-T's XGS framework. (We believe that a significant amount of the analysis and conclusions would be the same for the IEEE's 25G EPON.) The 25GS MSA defines a downstream rate of 25 Gbps with upstream rates of either 10 Gbps or 25 Gbps. This paper will only consider the 25 Gbps symmetric system where both upstream and downstream are 25 Gbps.

The results in this paper are largely based on modeling of the XGS and 25GS PON systems with some spot checking of the model in XGS mode from lab tests. Since 25GS technology is still in development, large system testing results are not available.

2. Adding 25GS PON to the ODN

Migrating between PON technology can be implemented via two methods:

- New fiber and splitters for each technology
- Wavelength coexistence via a single feeder fiber

For areas where an operator is fiber-rich and has capacity for splitters at the distribution enclosure, they may elect to use additional feeder fiber and additional splitters. In cases where fiber is constrained, an operator may elect to deploy a wavelength coexistence strategy. There are two methods for supporting coexistence. One method, which is internal to the OLT optic, is via a multi-PON module, in which multiple PON technologies (GPON and XGS, or XGS and 25GS PON) are integrated into a single module. The other option is an external coexistence element that combines multiple PON modules.

Table 1 – Wavelength Definitions

Legend	Wavelength Name	Center Wavelength (nm)	Wavelength Range (nm)	Specification
	ITU GPON US	1310	± 20	GPON "Reduced", ITU-T G.984.5
	ITU GPON DS	1490	± 10	ITU-T G.984.2
	ITU XGS PON US	1270	± 10	ITU-T G.9807.1
	ITU XGS PON DS	1577	+3/-2	ITU-T G.9807.1
	25G (UW0) US (MSA & 25G EPON)	1270	± 10	IEEE Std 802.3ca
	25G (UW1) US (MSA & 25G EPON)	1300	± 10	IEEE Std 802.3ca
	25G PON (UW3) US 25GS MSA Only	1286	± 2	25GS-PON MSA
	25G (DW0) DS (MSA & 25G EPON)	1358	± 2	IEEE Std 802.3ca
	25G (UW2) US (50G EPON - 2nd 25G WL)	1320	± 2	IEEE Std 802.3ca
	25G (DW1) DS (50G EPON - 2nd 25G WL)	1342	± 2	IEEE Std 802.3ca
	ITU 50G PON (12.5G/25G/50G) (Option 1) US	1270	± 10	ITU-T G.9804.3
	ITU 50G PON (12.5G/25G/50G) (Option 2) US	1300	± 10	ITU-T G.9804.3
	ITU 50G PON DS	1342	± 2	ITU-T G.9804.3

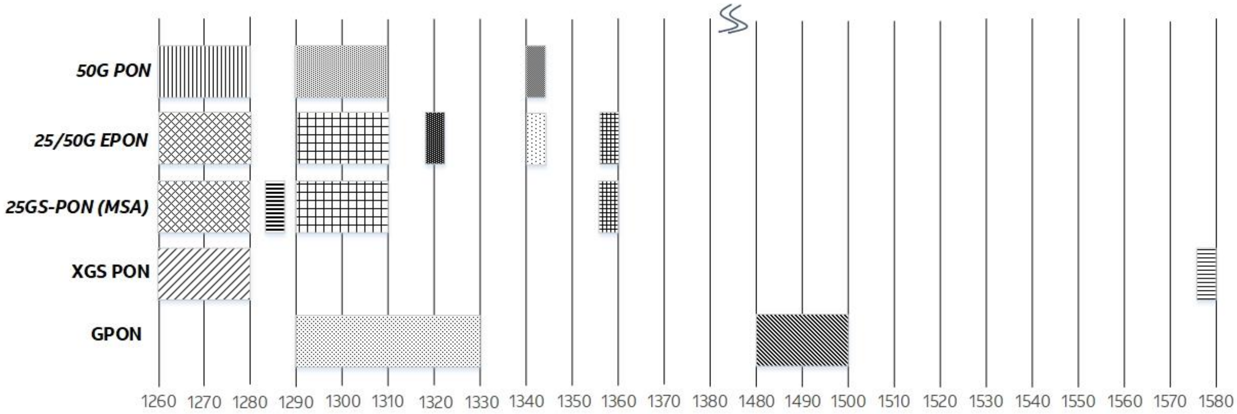

Figure 1 – PON Wavelength Plan

As illustratedin Figure 1, GPON and XGS PON upstream(US) and downstream(DS) wavelengths are separated which allows for coexistence across a single feeder fiber. The simple figure below illustrates how a passive coexistence element (CE) can bring XGS-PON US/DS Wavelengths together with GPON US/DS wavelengths. This would not be possible if the wavelengths overlapped.

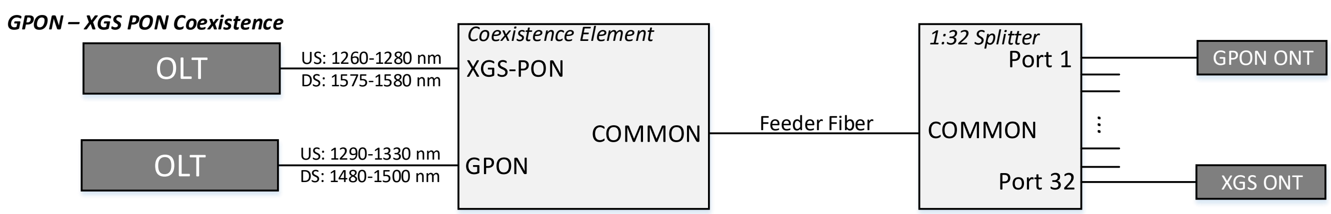

Figure 2 – GPON and XGS PON Coexistence

In conjunction with the use of a CE, two other key factors must be considered. First, the ONTs must be able to filter out or block the downstream wavelengths for the coexisting OLT. In the GPON-XGS coexistence example above, this means that the GPON ONT must block the downstream XGS wavelength and the XGS ONT must block the downstream GPON wavelength. Second, since both PON variants share the same ODN, the OLT and ONT optics for both PON variants should be compatible with the attenuation range of the shared ODN.

Table 2 below from ITU-T G.9807.1 gives the attenuation range classes (B+, N1, etc…) for the ODN. May want to pull in the attenuation range info as a reference of ODN Classes.

Table 2 – ITU-T Optical Class Definitions

PON TYPE	ODN Class	Attenuation Range (dB)
GPON	B+ class	13-28
GPON	C+ class	17-32
XGS/25GS	N1 class	14-29
XGS/25GS	N2 class	16-31

PON TYPE	ODN Class	Attenuation Range (dB)
XGS	E1 class	18-33
XGS	E2 class	20-35

As identified in theTable 2, B+ and C+ are ODN classes that supports a maximum fiber plant loss of 28db and 32db respectively. In addition, N1 (29dB), N2 (31dB), E1 (33dB), and E2 (35dB) are defined in ITU-T, while N1 and N2 are defined in the 25GS PON MSA specification.

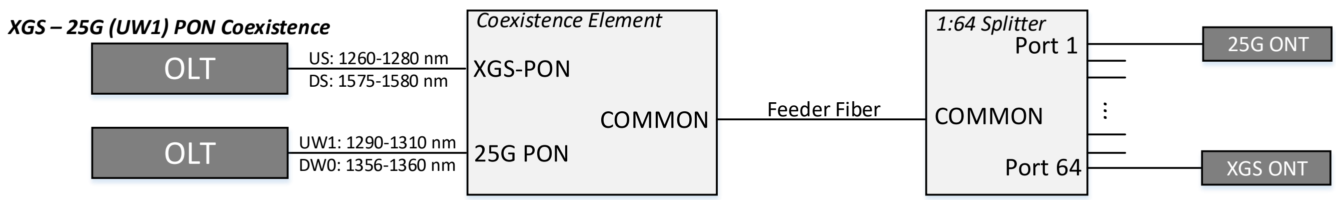

Figure 3 – XGS and 25GS (UW1) PON Coexistence

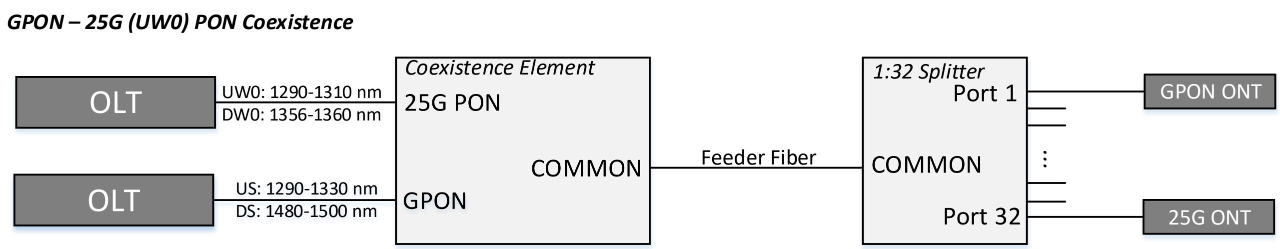

Figure 4 – GPON and 25GS (UW0) PON Coexistence

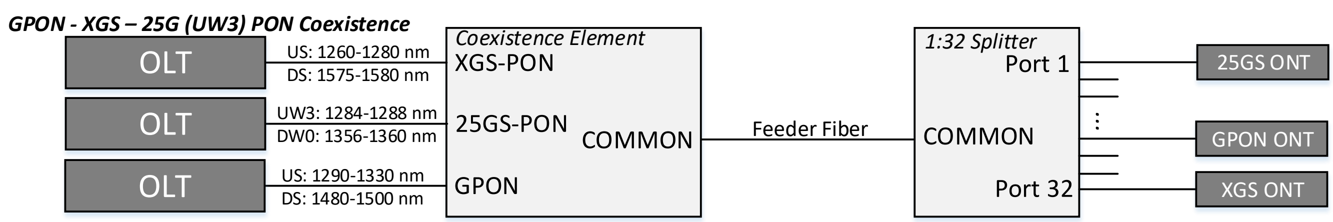

Figure 5 – GPON, XGS, and 25GS (UW3) PON Coexistence

As seen in Figure 3, Figure 4, and Figure 5, tere are three paths to facilitating coexistence between 25GS with GPON, 25GS with XGS, or 25GS with both prior PON technologies. The IEEE 802.3 standard developed the options for UW0 (1290-1310 nm) or UW1 (1260-1280 nm) with DW0 (1356-1360 nm).

The 25GS-PON MSA incorporated these into their specification and added UW3 (1284-1288 nm) with DW0 (1346-1350 nm).

UW0 and UW1 are 20nm wide, supporting the use of a DML for an N1 class ODN. UW3, on the other hand, is 4nm wide and will require an EML for an N1 class ODN. UW3 allows for triple coexistence between GPON, XGS, and 25GS but there are additional cost considerations that need to be weighed by the operator. With UW0 or UW1, an operator achieves double coexistence (GPON and 25GS or XGS and 25GS) with less expensive optics but loses the ability to coexist across three PON technologies. The 3rd PON technology must be placed on a separate fiber and splitter.

While this paper focuses on 25GS PON it should be noted that 50G PON (EPON or ITU-T) can coexist across a single feeder fiber with a prior PON technology.

3. System Upgrade to 25GS PON

Whether an operator is introducing PON for the first time or evolving a PON network to services supported by a 25G solution, a general consideration is how the new deployment implementation will support introduction to or coexistence with 25GS. Historically, PON technology evolutions involve evaluating customer behavior, preferred network topology, platform MAC scale and implementation granularity. These evaluations typically revolve around how and where an OLT is implemented.

There are three OLT model types: the large chassis, the pizza box, and the MSA pluggable OLT. In the following sections, we review the implications of these options with regards to introduction and evolution to 25GS.

3.1. Large Chassis OLT

Figure 6 – Example of Chassis with XGS Cards and 25GS Cards

The large OLT platforms, as seen in Figure 6, offer from the twenties to hundreds of PON ports. They are typically made up of modular, dedicated functional cards integrated into mainframes that distribute signal back and forth via a high-speed backplane. The platforms are physically large in the range of 10 rack units or more, and require locations with dedicated HVAC systems. OLT platforms of this type are quite efficient when dedicated in scale to the particular function or technology they were created for very large contiguous deployments of a particular PON technology, servicing a large number of customers within 20 km of a central office or hub. Their drawbacks are the sunk investment of a large platform when not used in scale and their otherwise inflexibility to work beyond the technology they were designed for.

Figure 7 – Example of Chassis with Dual Speed XGS/25GS Line Cards

The evolution of large OLT platforms to 25GS is implemented in the context of line card scale capability, their modularity, i.e., processing capacity and backplane throughput, with the understanding that if a box was created for, as an example, 100 PON instances of GPON, transitioning through XGS then to 25GS would likely result in a continual reduction of port counts to the point where it is no longer an efficient large scale box. Similarly, although to a lesser extent, evolving from XGS will also result in the PON instance capability being reduced. In one example, 25GS is supported by a new line card with half the number of PON ports. In the other example, the dual speed line card can select the speed by using a different transceiver. Large platforms can also be equipped to function as broadband network gateways (BNG) and collocated for large scale deployments.

3.2. Pizza box OLT

The "pizza box" refers to integrated platforms of a reduced size, typically one rack unit and generally less than four. The attractive quality for pizza box OLTs is flexibility. They can be used to service less dense population areas, and because of their size (and if temperature hardened) they can be placed in remote cabinets to service customers at distances much greater than 20km. They can also act in unison as a large OLT when racked together, interconnected by a top or rack switch and facilitated by a centralized controller.

With regards to evolution to 25GS PON, the flexibility of pizza box platforms is lost. Boxes that are not preconditioned to support a 25GS solution must be replaced with a new box. This is certainly a service-affecting change. If a 25GS option is available, it is most likely at a lower density than for prior PON technologies.

3.3. Micro-OLT

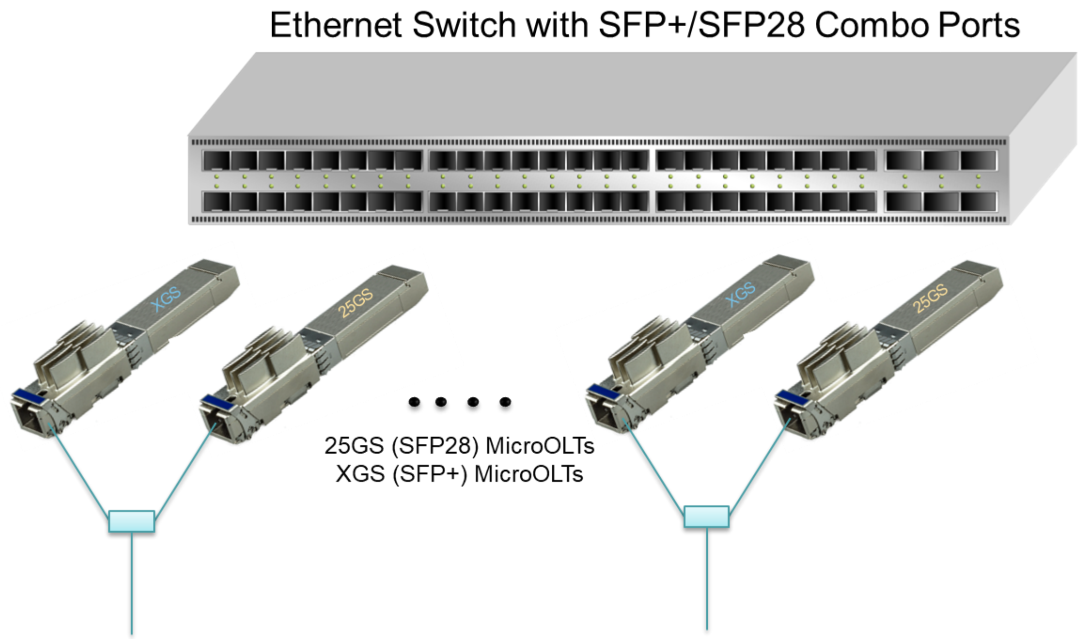

Figure 8 – XGS/25GS MicroOLTs in a multi-rate Ethernet Switch

As seen in Figure 8, the MSA pluggable OLT solution is also known as a micro-OLT (µOLT). (MSA pluggable refers to SFP+, SFP28, CFP2, etc.) These µOLTs facilitate one PON instance at a time through a single port on an Ethernet switch. In contrast to the large chassis and pizza box solutions which combine the MAC domain chipsets with the switching fabric but isolate the optics, the µOLT technology combines the MAC domain chipsets with the optical interface, leaving the switching fabric separate.

The scalability of these platforms is determined by the host switch, which also determines the geographic deployment. The µOLT can service sparsely deployed regions, including being temperature hardened and deployed in remote cabinets, in the most granular manner—one PON instance at a time. The flexibility is nearly boundless when creating PON deployments with µOLTs.

A previously unattainable dimension of flexibility is also available when using µOLTs. Because the PON service is limited to a dedicated switch port per the pluggable, combining other services on the same Ethernet fabric is now possible. Studies in recent years have looked at the convergence of services over the same access network, and the µOLT facilitates service integration at the layer 2 and layer 3 levels. The considerations for deterministic bandwidth assignments are beyond the scope of this paper, but there is a whole field of study that expects to leverage µOLTs as an easy method to enable service convergence.

With regards to the evolution or introduction to 25GS, the flexibility follows the granularity of the µOLT. Transitions to 25GS would be a straightforward port change; µOLTs have been generally available since 10G PON implementations. Note that the technical challenge is for pluggable solutions to maintain the development necessary to include additional capabilities, such as thermal growth into continually small form factors. One particularly interesting dimension of pluggable solutions is their ability to support the coexistence of different PON technologies on the same bridging domain. For example, a single switch can be populated with a mix of both 10G and 25GS services. Such flexibility could potentially support large usage end-line customers, gradually transitioning their services to 25GS. This collection of system design qualities makes the µOLT a very attractive option for 25GS solutions.

4. Use cases for 25GS PON

There are several factors that will determine how 25GS is used:

- Coexistence can be used as a mechanism to move from one PON technology to the next (e.g., GPON to XGS or XGS to 25GS) without having to place additional feeder fiber and splitters.
- XGS is now in scale deployment with many operators either offering or preparing to offer a multi-gig broadband service.
- Most of the cost of fiber-to-the-premises (FTTP) is fiber placement. If fiber is abundant, then it may be more cost effective to use a feeder fiber and place a new splitter. If fiber is constrained, then a coexistence strategy may prove to be more cost effective.
- ONTs are normally dedicated per unit/household and thus component costs are important to constrain (i.e., DML vs EML).
- OLTs are a significant cost but can be shared among multiple subscribers:
    - Large chassis: 128-256 PONs x 1:64 split (8K to 16K subscribers)
    - Pizza box: 24-48 PONs x 1x1:64 split (1.5K to 3K subscribers)
    - µOLT: Varies from 1-48 PONs based on the switch x 1:64 split (64 to 3K subscribers)

So the question is, what do you do with 2.5 times more capacity than XGS?

As shown in Figure 9 the architecture proposal will support consumer, business, and mobility applications across a single feeder fiber and splitter. This design will allow a mix of best effort multi-gig broadband services with a maximum latency of 4ms through the OLT, as well as business and mobility applications with 10G+ multi-gig broadband services with a "protected" (guaranteed) amount of capacity and sub-ms latency across the OLT.

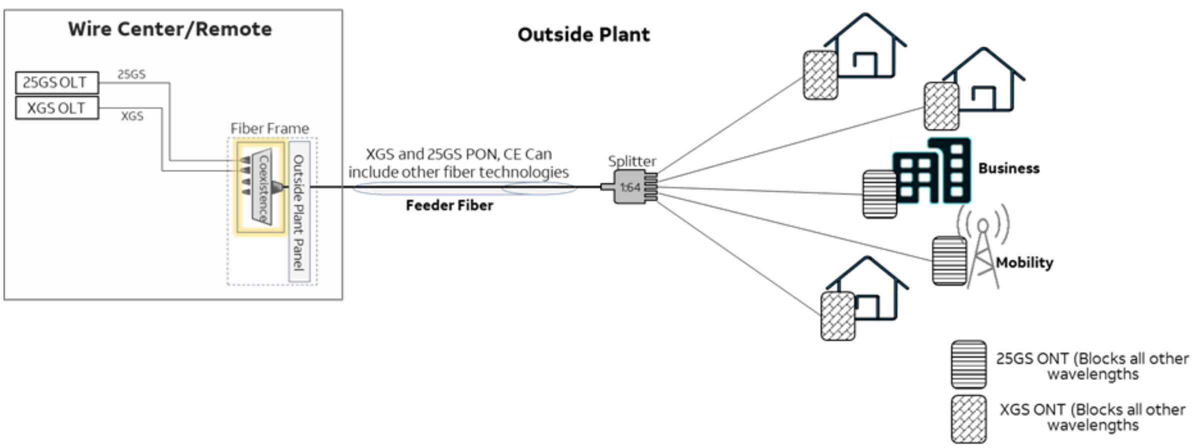

Figure 9 – XGS and 25GS Coexistence Architecture

XGS will continue to be a valuable technology, and currently supports consumer services through multi-gig broadband. Many users will continue to be supported over XGS for years to come. 25GS offers not only the ability to offer 10G+ broadband speed tiers, but also much more capacity at a small incremental increase in cost. 25GS could initially target business, work from home subscribers, mobility, and capacity failover for bandwidth intensive subscribers on XGS. Since consumer scale deployments on

25GS are assumed to be a future phase, a smaller percentage of the total is assigned to the splitter. For example, a splitter of 1:64 is usually capacity-managed to the total. In this design, however, two PON technologies are managed separately. XGS would have access to capacity of all 64 ports while 25GS would be set to capacity trigger at 25% (16 ports of a 1:64 splitter in this example). Placement of the XGS or 25GS ONT and assignment to the appropriate OLT PON port determine what is serviced on the splitter. An operator does not need to lock in specific ports of a passive splitter. Any port on the splitter can support XGS or 25GS as long as the capacity triggers are not exceeded. In this example, XGS can be deployed across more than 48 ports, and the capacity algorithm for 25GS would be decremented to maintain the total of 64.

The four initial use cases for 25GS mentioned above are described below:

- Business: 25GS has the capacity to support large numbers of customers behind a switch with an SFP28 25GS ONT or a purpose-built gateway. In this case, capacity, "protected" bandwidth, and sub-ms latency would benefit this market without driving significant cost.
- Work from home: an operator could use the additional capacity to offer consumer-grade broadband service-mapped interfaces (such as wireless SSID or physical LAN interfaces) while also supporting business services to other interfaces.
- Mobility: 25GS allows for more than 20Gbps x 20Gbps of usable bandwidth. "PON as a transport" for backhaul and mid-haul mobility services is possible with a sub-ms PON.
- Capacity protection: As new applications emerge that drive additional capacity requirements on XGS, 25GS offers a capacity failover from XGS and allows the operator to offer more tailored services to these high-end users. In addition, the XGS is managed for a better customer experience for the remaining subscribers. Note that with a coexistence strategy, the transition for an existing XGS subscriber to 25GS requires the shipment and customer installation of a new 25GS PON ONT or gateway.

Refer to the following Table 3 and Table 4 for details about PON capacity by splitter.

Table 3 - 1:16 Splitter, 25% 25GS PON, Decrementing 25GS

Splitter Port	XGS	25GS (Set at 25%)	
1	1		
2		1	
3	2		
4	3		
5	4		
6		2	
7	5		
8		3	
9	6		
10	7		
11	8		
12	9		
13	10		
14	11		
15	12		
16	13		Assigned to XGS and decrements 25GS

Table 4 - 1:16 Splitter, 25GS PON Capacity Met

Splitter Port	XGS	25GS (Set at 25%)	
1	1		
2		1	
3	2		
4	3		
5	4		
6		2	
7	5		
8		3	
9		4	Capacity of 25% on 25GS reached, new splitter
10	6		
11	7		
12	8		
13	9		
14	10		
15	11		
16	12		

5. Cost Analysis of 25GS PON

5.1. OLT/ONU

The highest equipment cost in a PON deployment is the customer-side device, the Optical Network Unit (ONU), also known as the Optical Network Terminal (ONT). These devices connect to an Optical Line Terminal (OLT) at the operator side. Since the OLT is designed to split a single fiber to 64 ONUs at 20 kilometers, adding only $1 to the cost of the ONU adds $64 to the cost of operating the OLT port. For 25GS to succeed as the next evolution from XGS, it must support a cost-effective ONU. This goal is currently attainable in most cases.

In 2020, the *Journal of Optical Communications* presented a relative cost comparison between components of the various PON speeds beyond XGS. The article noted that 25GS requires the same components as XGS, with minor upgrades for the higher speed. As a result, there is only a small cost increase for the 2.5x data rate. At speeds higher than 25 Gbps, additional components and major technology changes are required, which increase both the expense and the power requirements of the corresponding ONUs, as shown in Figure 10.

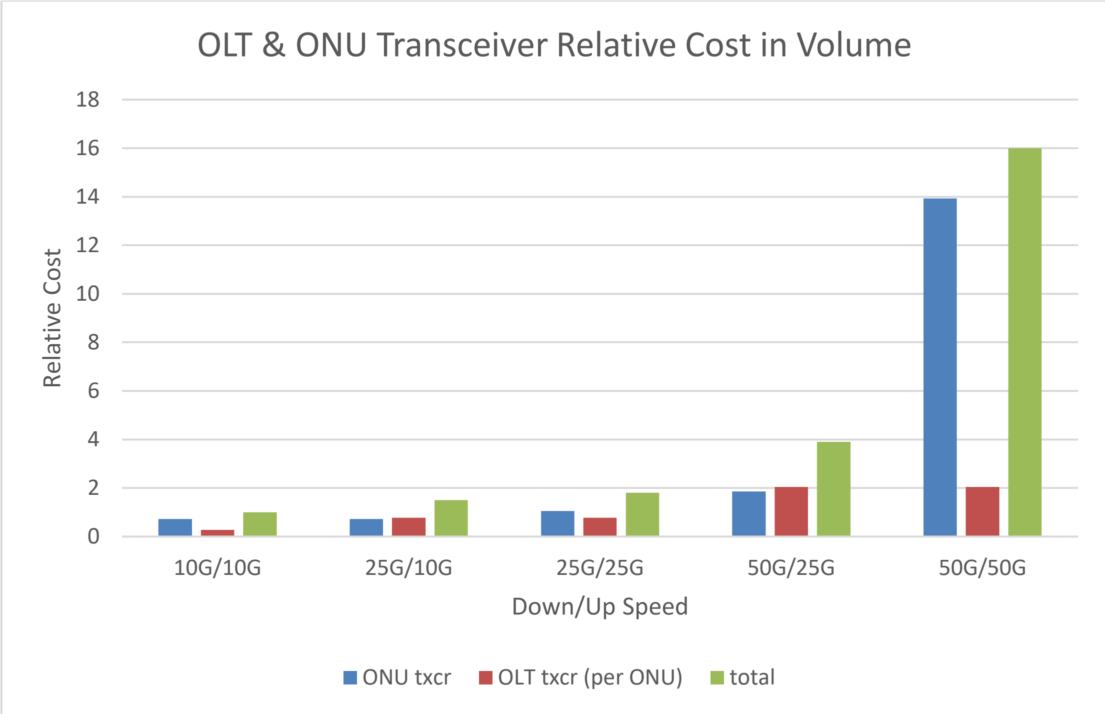

Figure 10 – Relative PON Transceiver Cost in Volume

Note that the specification for 50G/50G was not finalized at the time of this article. While we believe that costs will be significantly higher than for lower speeds, they will likely be less than the 14x shown above.

5.2. Optical Budget

PON speed increases require more optical budget. Going to a 25 Gbps line rate from 10 Gbps causes a 4 dB penalty. Forward error correction, or FEC (which sends redundant data to assist the receiving device with assessing errors) provides a portion of the needed gain. A larger block size and a low-density parity check (LDPC) FEC in 25GS provides additional gain over the Reed Solomon (RS) FEC used in XGS. 25GS FEC requires significantly more logic area in the PON OLT or ONU ASIC. The LDPC makes it difficult or cost prohibitive to use FPGAs for the OLT PON MAC. In ASIC form, the cost difference is not very significant. The 25GS FEC provides roughly 1.5 dB of optical gain over the XGS FEC.

In addition to the gain from FEC, the transmit optical power and receiver sensitivity must recover 2.5dB of the 4dB penalty. The proposed values for transmit optical and receiver sensitivity are near the limits of unamplified optics. For speeds above 25 Gbps in the upstream or downstream, a silicon optical amplifier (SOA) will be required. Adding the SOA can double the cost of the optics and consume an extra one watt of power. While the N1 optical budget doesn't require the SOA at 25 Gbps, the N2 optical budget is still undecided. It is easy to achieve N2 level with the SOA but it also seems very possible to achieve it with proper process control.

5.2.1. Transmission Lasers

PON systems use one of two laser types to transmit data across the fiber line:

- Directly Modulated Laser (DML) creates wider wavelengths with more frequency distortion. They are used for lower speeds and shorter distances.
- Externally Modulated Laser (EML) are temperature controlled with less distortion and narrower wavelengths. They are used for higher speeds and longer distances.

The EML is often 2.5x the cost of the DML. EML also requires up to one watt of additional power for temperature control. For the OLT transmitter, the EML is normally used for all current versions of PON and will likely be the solution for 25GS. As a result, migrating from XGS to 25GS will not result in significant OLT transmitter cost increase if the SOA is not added.

For the ONU transmitter, the DML can be used for GPON and XGS. For 25GS, the requirement is determined by the wavelength. DML can generate the wider UW0 and UW1 wavelengths, but EML is required for the narrower UW3 wavelength. For 50Gbps, EML will be required for all wavelengths.

5.2.2. DSP, High Speed ADC, Coherent

XGS and 25GS don’t require a high-speed analog to digital converter and DSP to process the receive signal. Going to 50 Gbps and beyond will require these functions. These functions add significant cost, complexity, and power to the ONU receiver. In some estimates, 2 to 5 watts of additional power will be required to support these functions. Adding Coherent optics allows for higher split ratio, longer reach, and speeds beyond 100 Gbps. Unfortunately, Coherent optics will also increase the cost and power significantly over 25GS.

5.2.3. Cost Summary

With all things considered, 25GS has been estimated to be 1.5 times the cost of XGS in volume. Since the XGS ONU and 25GS ONU don’t require an EML transmitter, SOA, or DSP, the increment in cost is reasonable for a 2.5 times speed increase. The 25GS ONU optics could be significantly higher cost if the SOA is required to reach the N2 optical budget and/or the EML transmitter is required to use the UW3 wavelength.

6. Bandwidth Analysis of 25GS

A 25GS delivers less than a full 25 Gbps worth of Ethernet bandwidth across the PON. Framing, forward error correction (FEC), and physical layer management eat into the total bandwidth, see Figure 11

6.1. Downstream Bandwidth Calculation

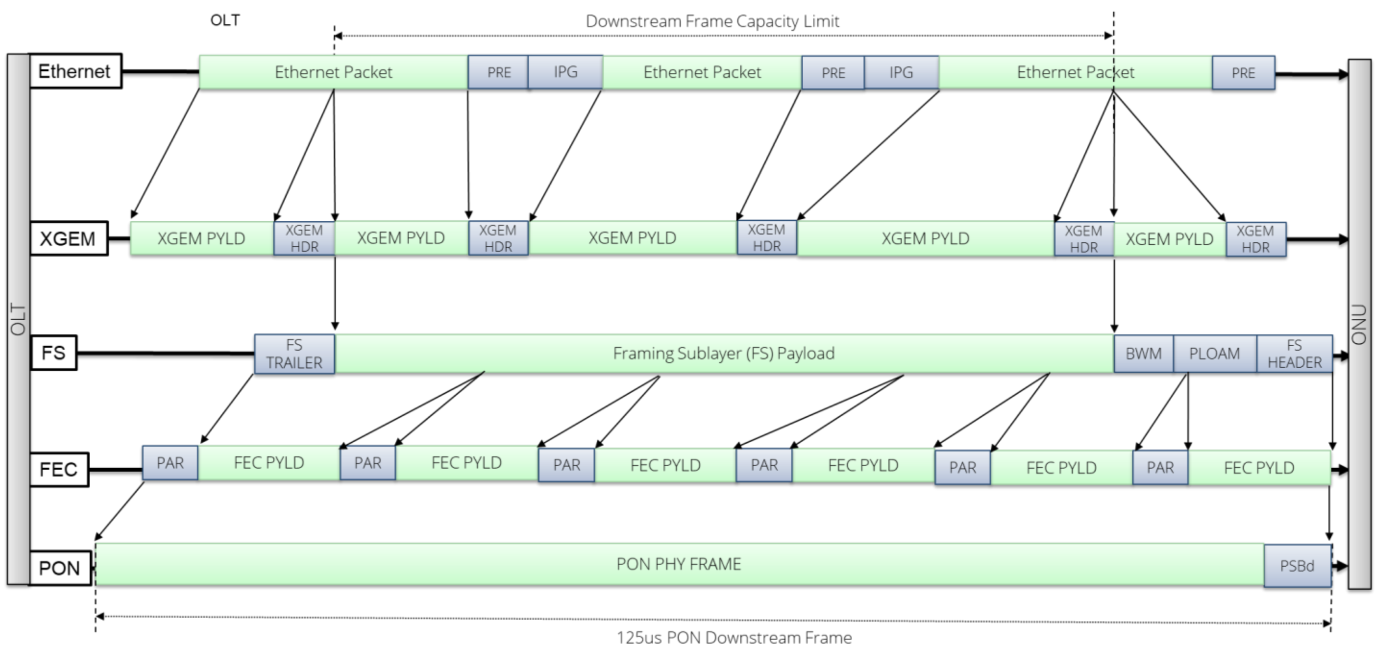

Figure 11 – XGS and 25GS Downstream Framing

To determine the actual downstream bandwidth, overhead costs will be removed and variable overhead costs will be estimated. These overhead costs are described below and listed in Table 5:

- Like XGS, 25GS uses non-return-to-zero (NRZ) data without additional bits for line encoding. At a line rate of 24.8832Gbps, 25GS PON is exactly 2.5x the XGS line rate of 9.95328Gbps; no additional bandwidth loss is incurred by the higher rate alone. These starting rates are shown in the first row of the table.
- In both XGS and 25GS, the downstream contains a 24-byte physical synchronization block (PSBd) every 125µs. This block allows for downstream byte alignment, FEC block alignment, and frame alignment (these bytes are not covered by FEC). The PSBd impact is shown in the second row in the table.
- The most significant difference between XGS and 25GS is the new FEC (mentioned in the previous section). In XGS, a FEC block of 248 bytes contains 216 bytes of payload and 32 bytes of parity, reducing the data rate to 8.667648Gbps (efficiency ≈ 87%). In 25GS, a FEC block of 2144 bytes contains 1824 bytes of payload and 320 bytes of parity, reducing the data rate to 21.154304Gbps (efficiency ≈ 85%). The resulting data rates are shown in the third row in the table.
- After the PSBd, the 125µs downstream frame contains the framing sublayer (FS) header and FS trailer. For both XGS and 25GS, the FS trailer is a 32-bit interleaved parity over the FS data excluding FEC parity. The FS header has a minimum size of 4 bytes and a variable size for the physical layer OAM (PLOAM) and the bandwidth map (BWmap) carrying the upstream grants. The resulting data rates are shown in the fourth row in the table.
    - Each PLOAM message is 48 bytes long. For the bandwidth calculations, we assume one PLOAM every 125µs for both XGS and 25GS.

  - Each upstream allocation in a BWmap is 8 bytes long. The number of allocations required by the upstream is determined by the polling and maximum burst size configuration.

- ITU-T-based ONUs are managed by ONU management control interface (OMCI) frames. These frames allow for setting up the ONU data path and monitoring the performance. The amount of OMCI required for XGS and 25GS should be the same, averaging two frames per second for each ONU; for a 64-ONU system, this is 128 frames per second. OMCI includes a baseline frame of 48 bytes and a variable extended frame size. Since many ONUs support only the baseline size, we will use this value in the calculations below. Finally, every OMCI frame includes an 8-byte XGEM header, bringing the total size to 56 bytes. The resulting data rates are shown in the 5th row inTable 5.

Table 5 – Downstream Overhead

	XGS (10G/10G)		25GS (25G/25G)	
	Deduction (bps)	Available BW (Gbps)	Deduction (bps)	Available BW (Gbps)
Initial Line Rate	-	9.953280000	-	24.883200000
minus: PSBd, FS Header/Trailer	2,560,000	9.950720000	2,560,000	24.880640000
minus: FEC Parity	1,284,096,000	8.666624000	3,727,360,000	21.153280000
minus: OMCI	57,344	8.666566656	57,344	21.153222656
minus: PLOAM	768	8.666565888	768	21.153221888
minus: BW Map	1,677,517	8.664888371	227,253,073	20.925968815
% User XGEM Data		87%		84%

The ITU-T-based PON system uses XGEM framing to carry Ethernet. The XGEM frame is an 8-byte XGEM header that encapsulates each Ethernet frame from the destination address to the CRC-32. XGEM frames do not include the 12-byte interpacket gap (IPG) or the 8-byte preamble found in the Ethernet frame. XGEM framing allows for segmentation of the Ethernet frames, but only at the 125µs frame boundary in the downstream direction. When comparing the XGEM bandwidth (BW) with the Ethernet Layer 1 BW, the per-packet overhead is 20 bytes (preamble + IPG) for Ethernet and 8 bytes for XGEM. This difference allows for a higher BW to be carried on the PON than on the Ethernet side of the network. The difference can be significant with small frames and less significant with large frames.

Finally, our calculations must include the application layer bandwidth available to customers running speed tests, FTP, etc. To calculate these values, we assume a layer 2 (DA/SA/TYPE/CRC-32) of 18 bytes, an IPv4 header of 20 bytes, and a TCP header of 20 bytes. Despite a drop in data rates from the table above, applications on XGS can still operate at download speeds above 8Gbps while 25GS can still operate at download speeds above 20 Gbps, as shown inTable 6.

Table 6 – Downstream Bandwidth Available

	XGS (10G/10G)	25GS (25G/25G)
XGEM Frame BW	8.659910 Gbps	21.146566 Gbps
L1 Ethernet BW (64B)	10.103228 Gbps	24.670994 Gbps
L1 Ethernet BW (1500B)	8.728822 Gbps	21.314841 Gbps
L7 Application BW (1500B)	8.280895 Gbps	20.381963 Gbps

6.2. Upstream Bandwidth Calculation

The upstream for 25GS uses the same 24.8832Gbps line rate (2.5 times the XGS line rate) as the downstream. The calculation of the possible upstream bandwidth is more complicated than for the downstream because the amount of bandwidth lost is determined by the variable size of the upstream burst. We will start by assuming 64 ONUs on the PON and a simple static granting to see the maximum possible performance. We will then add the overhead required to make the upstream a dynamic operation, so it can be shared on demand. Finally, although the PON can support multiple services in multiple upstream traffic containers or prioritized traffic in a single container, we will assume a single service and the same priority level for each ONU for both our bandwidth and efficiency analyses.

For these analyses, we will consider the three primary impacts on upstream bandwidth: burst overhead, static (fixed) BW allocation, and dynamic BW allocation (DBA).

6.2.1. Burst Overhead

Figure 12 – Upstream Burst Overhead

The upstream burst can be split into 3 parts, (as seen in Figure 12 and Figure 13): dead time between bursts, a preamble at the start of the burst, and the block of data. For ITU-T PON, the dead time and preamble are determined by the configurable parameters of guard time and preamble time. The guard time includes the ONU Laser ON time, the ONU Laser OFF time, and the dead time for any jitter in the upstream slot time. The preamble time includes the time required for the OLTs to perform gain control and clock recovery. The preamble is shown as the physical layer synchronization block for upstream (PSBu). Optical components for 25GS are still in development and they will certainly improve over time. For this analysis, we will assume the same time duration values as XGS. For XGS, a value of 256 bytes will be used for guard time and preamble time combined. For 25GS PON, we will multiply the XGS's 256 bytes by 2.5 to get 640 bytes for both values combined.

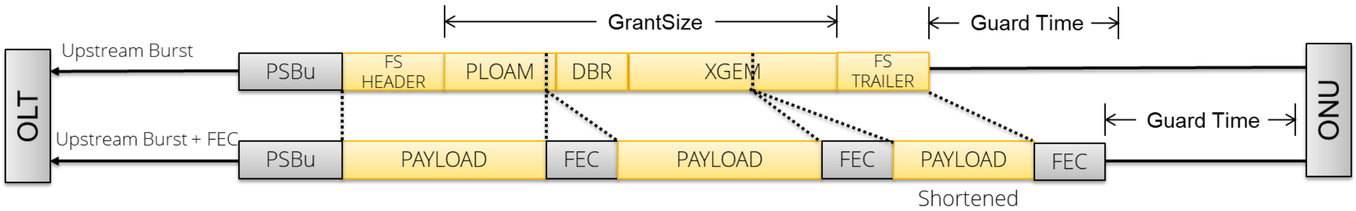

Figure 13 – Upstream Framing Overhead

After the PSBu, the FS burst is the FEC protected data. The FS burst starts with the FS header that is 4 bytes long without a physical Layer OAM (PLOAM) or 52 bytes with a PLOAM. The grant in the BWmap from the OLT selects the inclusion of PLOAM, inclusion of bandwidth request (BWR), and the size of XGEM blocks within the FS burst. Bursts that only carry PLOAM and/or BWR can have a length of 0. The FS burst can contain grants to one or more allocation IDs on the ONU. For our analysis, we will assume a single allocation per burst. The standard allows for a maximum grant size of 9719 upstream slot times. In the case of XGS, upstream slots are defined as 128 bits or 16 bytes. In 25GS, the upstream slots are 2.5 times greater, so they are 320 bits or 40 bytes. In that case, the maximum grant of 9719 upstream slots equals 155,504 bytes in XGS and 388,760 bytes in 25GS. The FS burst ends with the 4-byte FS trailer that contains a parity for checking the data.

6.2.2. Static (Fixed) BW Allocation

With static allocation, the OLT sends a fixed-size grant to each ONU in a simple round robin. This is often called unsolicited granting in DOCSIS. In this analysis, we will assume that all ONUs receive the same fixed allocation.

Physical layer OAM (PLOAM) provides the initial configuration of the PON physical layer and key exchange for encryption. PLOAM in the upstream is infrequent after registration, so required bandwidth consideration is negligible. ONU management control interface (OMCI) is the management traffic from the OLT to the ONU. Data path configuration, firmware download, and statistic gathering are the primary tasks using OMCI. For OMCI frame bandwidth in the upstream, a fixed payload of 56 bytes is granted every 8ms. Discovery slots are 250µs of deadtime for ONUs to register (assuming a 20km PON). In this analysis, a discovery slot is granted every 3 seconds. With the maximum allocation of 9719 upstream slots and the overhead described above, we can calculate the maximum upstream bandwidth for both XGS and 25GS:

- For XGS, a grant of 9719 upstream slots would carry 155,504 bytes of payload data, 8 bytes of FS header/trailer, and 23,040 bytes of FEC parity.
- For 25GS, a grant of 9719 upstream slots would carry 388,760 bytes of payload data, 8 bytes of FS header/trailer, and 68,480 bytes of FEC parity.

It should be noted that while these grants are possible, they are not practical for a 64-ONU system, since the delay between grants would be very long. As an alternative, a static allocation of 20,000 bytes is included in the following table as an example of a viable lower latency configuration.

Table 7 – Static Allocation Analysis for 20km, 64 ONU PON

	XGS (10G/10G)		25GS (25G/25G)	
Line Rate	9.953280 Gbps		24.883200 Gbps	
Static Payload Size	20,000 Bytes	155504 Bytes	20,000 Bytes	388760 Bytes
Minus Preamble, Guard time, FS H/T	(-222,852,025 bps) 9,730,427,975 bps	(-29,233,025 bps) 9,924,046,975 bps	(-1,291,341,739 bps) 23,591,858,264 bps	(-69,785,265 bps) 24,813,414,735 bps
Minus FEC	(-1,256,075,050 bps) 8,474,352,925 bps	(-1,275,622,869 bps) 8,648,424,106 bps	(-3,507,347,934 bps) 20,084,510,330 bps	(-3,687,418,908 bps) 21,125,995,827 bps
Minus OMCI	(-38,912,000 bps) 8,435,440,925 bps	(-38,912,000 bps) 8,609,512,106	(-106,496,000 bps) 19,978,014,330 bps	(-106,496,000 bps) 21,019,499,827 bps
Minus Discovery	(-829,440 bps) 8,434,611,485 bps	(-829,440 bps) 8,608,682,666	(-2,073,600 bps) 19,975,940,730 bps	(-2,073,600 bps) 21,017,426,227 bps
L1 Ethernet BW (1518B packet)	8.507744902 Gbps	8.677271545 Gbps	20.08482188 Gbps	21.09803954 Gbps
L7 Application BW (1518B packet)	8.076272794 Gbps	8.237201857 Gbps	19.066215828 Gbps	20.028047937 Gbps
Latency	1.46 ms	9.46 ms	.762 ms	9.72 ms

The static allocation in Table 7, shows both a very high possible bandwidth (the second column for each rate) and the ability to achieve a lower latency with a smaller burst size consuming lower bandwidth (the first column for each rate). Note that using smaller burst sizes increases the amount of bandwidth required for the combined per-burst overhead (preamble, guard time, FS header and trailer), since smaller bursts require more bursts to be sent, and each burst contains its own overhead.

While static granting is helpful for a simplified overhead analysis, it is impractical for most PONs. For a 64-ONU system with the maximum burst size, the upstream Ethernet bandwidth per ONU is only 135Mbps in XGS and 329Mbps in 25GS. Since operators expect PON networks to meet their requirements for statistical gain, fewer ONUs and different allocations should be used in real-world scenarios to allow for grant sizes to be adjusted accordingly. The static BW analysis above shows the maximum possible upstream bandwidth and cost of the PON overheads.

6.2.3. Dynamic BW Allocation (DBA)

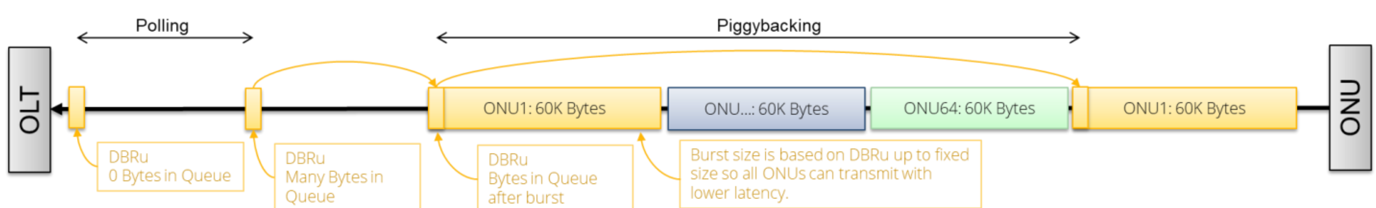

Figure 14 – DBA Upstream Bursts

Adjusting the upstream slot according to subscriber needs is facilitated with DBA in a request/grant methodology. The OLT sends a dynamic bandwidth resource unit (DBRu) in the upstream slot. The returning DBRu consumes 4 bytes in the upstream burst and indicates the amount of packet data remaining in the queue. If the ONU returns the DBRu with a value of 0 (indicating no data in the queue), no grant will be sent from the OLT. Instead, the OLT will send another DBRu to the ONU after a fixed time interval. This process is known as polling.

If the ONU returns the DBRu with a non-zero value, the OLT then sends a variable-size grant to the ONU based on the reported data payload need. The data payload sent upstream from the ONU includes a DBRu reporting additional payload waiting. The granting of a follow-up DBRu with the data payload burst is known as piggyback granting. When a grant is piggybacked, polling is not needed.

Calculating the available upstream bandwidth for a DBA system is more complicated than for a static allocation system since there are endless possible burst sizes from different traffic scenarios and different SLAs. For this analysis, we will establish some basic parameters to model the most complex scenario, as follows:

- We will consider request/grant or status-reporting DBAs only, and not non-status reporting or predictive grants.
- All ONUs will be assigned the same SLA, with the goals of low latency and higher bandwidth for the entire PON.
- The status reporting methodologies used will be polling and piggybacking.

In a DBA PON system, the amount of overhead is determined by the number and size of the bursts. We will examine the extreme conditions to calculate the bandwidth and latency of the PON upstream in a PON of 64 ONUs, considering two scenarios:

- When a single ONU is requesting the maximum rate, it will use piggybacking, while the other 63 ONUs will use polling. The bandwidth for polling the 63 ONUs will decrease the amount of bandwidth available for the single piggybacking ONU. In this scenario, the key factor is the polling interval: a shorter polling interval will have lower latency for the 63 polling ONUs but consume more upstream bandwidth from the single piggybacking ONU, while a longer polling interval will increase latency for the 63 polling ONUs but save upstream bandwidth for the single piggybacking ONU.
- When all 64 ONUs are actively sending payload data and requesting the maximum rate, they will all use piggybacking exclusively. Since all ONUs are treated equally, the DBA will schedule the ONUs in a round robin based on the last time data was granted. The size of the round robin is the number of active ONUs times the maximum size of the burst. In this scenario, the key factor is the maximum size of the burst: a larger burst will be more efficient for the ONU that is sending its payload, but it will create a longer delay for the other ONUs.

Table 8 – Dynamic Bandwidth Allocation Analysis for 20km, 64 ONU PON

	XGS (10G/10G)		25GS (25G/25G)	
Line Rate	9.953280Gbps		24.883200Gbps	
Polling Interval	3.5 ms		3.5 ms	
Max PON Burst Size	60,000 Bytes		140,000 Bytes	
ONU Activity	1 active and 63 idle	64 active	1 active and 63 idle	64 active
Payload Burst Overhead	75,319,749 bps	75,920,301 bps	191,775,180 bps	193,589,029 bps
Payload FEC Overhead	1,259,482,114 bps	1,269,524,418 bps	3,634,877,255 bps	3,669,256,667 bps
Polling Overhead	80,064,000 bps	0 bps	232,128,000 bps	0 bps
L1 Ethernet BW (1500B packet)	8,539,654,048 bps	8,607,743,778 bps	20,761,985,282 bps	20,958,356,379 bps
L7 Application BW (1500B packet)	8,122,924,176 bps	8,187,691,169 bps	19,748,813,154 bps	19,935,601,462 bps
Latency	4.125 ms	3.95 ms	4.125 ms	3.80 ms

Table 8 above shows a practical configuration for XGS applied to 25GS. This configuration allows an XGS system to have a ~4ms maximum upstream latency. In both XGS and 25GS, polling inactive ONUs becomes significant when fewer ONUs need to use the upstream. The polling waste is reduced to 0 when all ONUs are active. XGS shows 8.5 Gbps of Ethernet bandwidth 8.1 Gbps of application bandwidth while 25GS achieves 20 Gbps of Ethernet bandwidth and just under 20 Gbps of application bandwidth. While the goal of 2.5 times the speed of XGS is not reached, the results are very close.

Main Port Traffic Statistics

Name	TX L1 (%)	TX L1 (bit/s)	TX L2 (bit/s)	TX (pps)	TX (bytes)	TX (packets)	RX L1 (%)	RX L1 (bit/s)	RX L2 (bit/s)	RX (pps)
P-0-0-0	89.600	8,959,986,780	8,842,092,220	736,841	39,254,169,000	26,169,446	86.120	8,612,033,640	8,498,717,480	708,226

Figure 15 – DBA Lab Results for 64 active ONUs

Name	TX L1 (%)	TX L1 (bit/s)	TX L2 (bit/s)	TX (pps)	TX (bytes)	TX (packets)	RX L1 (%)	RX L1 (bit/s)	RX L2 (bit/s)
P-0-0-0	90.000	8,999,997,660	8,881,576,700	740,131	23,786,878,500	15,857,919	85.059	8,505,864,680	8,393,945,480
P-0-0-1	90.000	8,999,990,280	8,881,569,320	740,131	23,848,053,000	15,898,702	87.329	8,732,881,510	8,617,672,070

Figure 16 – DBA Lab Results for 1 active ONU on 64 ONU system

A lab test of a 64-ONU system in XGS with the sample configuration shown in the table able was performed to validate the model. With 64 ONUs transmitting 9 Gbps upstream, the "RX L1" of 8.612 Gbps received is very close to model's prediction of 8.607 Gbps, as seen in Figure 15,. With a single ONU transmitting, the 8.505 Gbps "RX L1" on P-0-0-0 closely matches the model's prediction of 8.539 Gbps, as seen in Figure 16. The downstream performance from the lab is also available on P-0-0-1 "RX L1". The lab shows 8.733 Gbps which is very close to the downstream Ethernet L1 model's predication of 8.729 Gbps.

7. Latency Analysis of 25GS

Latency has become a hot topic in the industry. Residential subscribers are looking for low latency for gaming and interactive experiences such as the metaverse. In this paper, we will look at the one-way packet latency between the OLT NNI port and ONU Ethernet UNI port. These delays assume layer 2 switching at the OLT and ONU. Routing devices, Wi-Fi interfaces, etc. will add more latency. This analysis focuses on worst case scenarios with a fully loaded PON. In most cases, the customers will see much better latencies due to a lower take rate or activity.

7.1. Downstream Latency

The downstream latency for PON is very low, and largely based on the functional characteristics of the network switches and the length of the PON fiber. On a 20km fiber, the flight delay is 100µs for both XGS and 25GS. The FEC block for 25GS is 2144 bytes versus 248 bytes for XGS, so the 25GS decoder will require more time for processing. However, the FEC decoder time is still minimal for 25GS, at approximately 1µs. This is a small value compared to the store and forward delays in the OLT and ONU switching. Finally, equipment delays are normally less than 50µs, so the worst-case delay on a 20km fiber is 150µs. This calculation also applies to both XGS and 25GS.

7.2. Upstream Latency

If operators don't have demand for the 20 Gbps of upstream traffic possible in 25GS, they can trade off the upstream bandwidth for a lower latency upstream. Since TCP/IP traffic downstream requires an upstream acknowledge, a lower latency upstream can also improve downstream throughput. By decreasing the interval for polling and reducing the maximum burst size, the latency for all ONUs in a 64-ONU PON could be reduced significantly. The inaccuracy of predictive granting, inflexibility of fixed granting, and other less predictable techniques can be avoided. The ability to lower the latency provides an opportunity to increase the downstream performance and throughput. Since the large downstream TCP/IP bursts require acknowledging in the upstream direction, minimizing the latency for these frames increases the downstream throughput and overall latency. Upstream latency on a PON can be broken down into 3 areas: start latency, continuation latency, and queue delay.

7.2.1. Start Latency

Whenever the upstream packet stream gives a DBRu of 0 to the DBA, it is considered idle and won't be granted until the next polling cycle. Since data arrives randomly compared to the polling cycle, the maximum wait time to be sampled is the polling interval. After the polling interval, the DBRu must be sent to the OLT/DBA to be granted. This transmit time can be 0µs for a 0 km distance ONU and 100µs for an ONU at the end of the fiber. An idle ONU should be at the front of the round robin so it will be granted quickly. The DBA will have a delay to issue the grant. This delay can be 125µs waiting for the start of the downstream framing or additional delay for software processing. In a software DBA, the cycle time is often used to define this time interval. In the example below, a hardware DBA is assumed that only waits for the 125µs downstream frame boundary. After the DBA issues the grant, the grant must travel to the ONU and back up the PON. This time is often referred to as the PON round trip time. On a 20km PON fiber system, it is 250µs. Because ranging an ONU delays the transmitter for closer ONUs, this delay is constant regardless of the ONUs position. The PON round trip time also sets the size of a discovery window to the same 250µs. If the polling or data grant to the ONU is needed after a discovery window has been requested, an additional 250µs of delay/jitter is possible. A small delay for the ONU and OLT hardware should be included as well. While not always the case, it will be considered a fixed delay in this analysis. When testing in the lab, the fixed delay will show up as the min delay on a long test and variable delay can be determined by subtracting the max delay from the min delay.

Start_Up_Fixed_Delay = Upstream_Flight_Time + PON_Round_Trip_Time + ONU_HW_Delay + OLT_HW_Delay

Start_Up_Variable_Delay = Polling_Interval + DBA_Delay + Discovery_Window

Start_Up_Max_Delay = Start_Up_Fixed_Delay + Start_Up_Variable_Delay

The start latency is a big factor in downstream TCP/IP performance. The upstream acknowledge frames are often spread further apart than the polling interval so they will see the start latency as the dominant factor. Large bursts upstream will hit the start latency for the first packets but the delay for the tail of the burst determines the true latency of the transaction, so the start latency might not be the key factor in large upstream bursts.

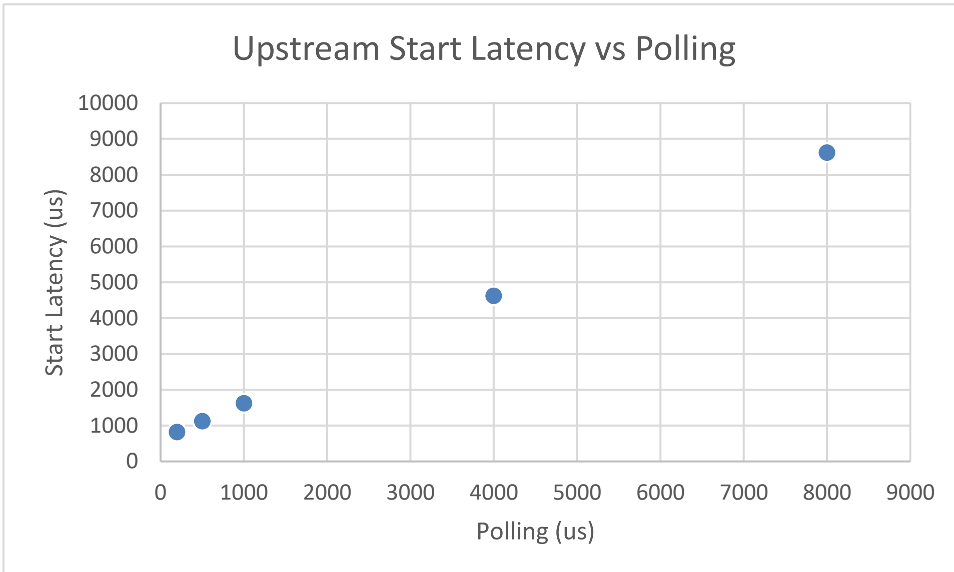

Figure 17 – Upstream Start Latency vs Polling

The polling interval is the only factor in the start latency that can be easily modified. The fiber length, number of ONUs, and hardware delays are fixed inputs. Based on 20km fiber and estimates for the hardware delay, the start delay is 625µs plus the polling interval. Figure 17 shows the direct relationship between the maximum start latency and the polling interval. Both XGS and 25GS have the same start latency equation since the data rate is not a factor.

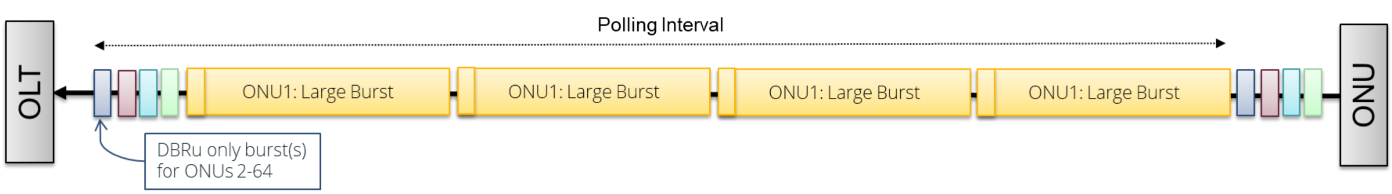

Figure 18 – Polling and the Single Transmitting ONU

Decreasing the polling interval can significantly reduce the bandwidth. Figure 18 shows the scenario with the largest polling penalty. 63 ONUs are idle and 1 ONU is requiring the full bandwidth. In this case, the polling grants are non-traffic carrying blocks of time that limit the bandwidth to the single ONU transmitting. A shorter interval limits the bandwidth. When more ONUs are transmitting, the bandwidth lost to polling idle ONUs decreases but it is minimal until a large percentage of ONUs are active.

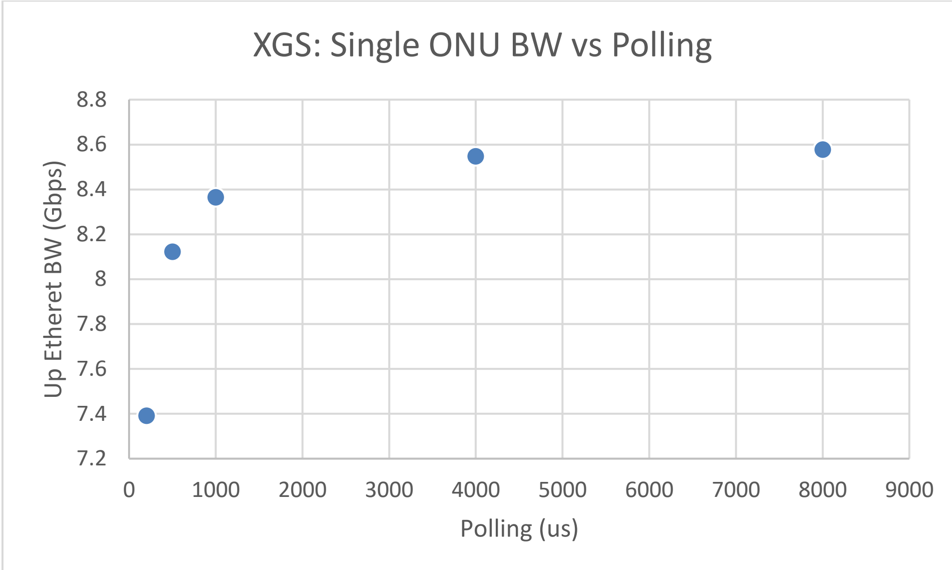

Figure 19 – XGS Upstream Bandwidth versus Polling Interval

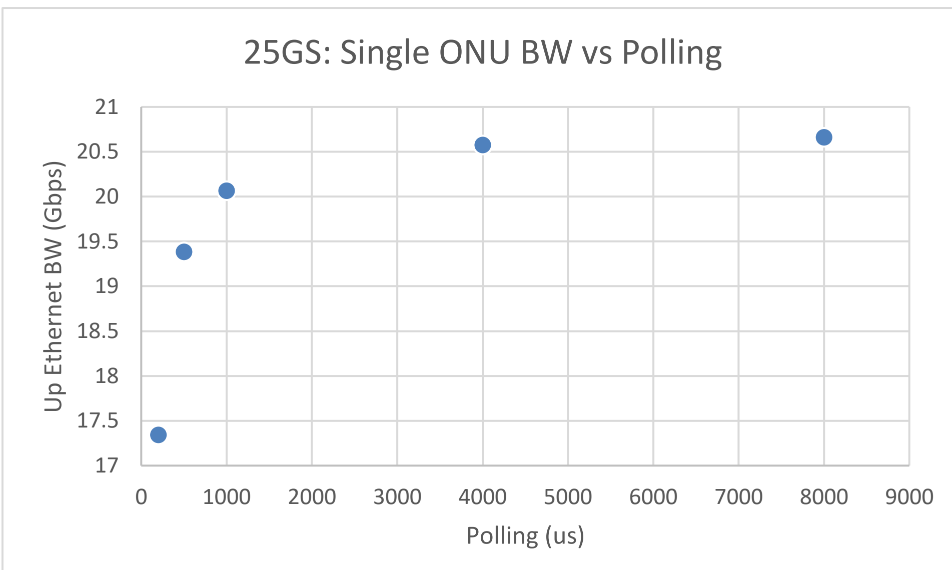

Figure 20 – 25GS Upstream Bandwidth versus Polling Interval

7.2.2. Continuation Latency

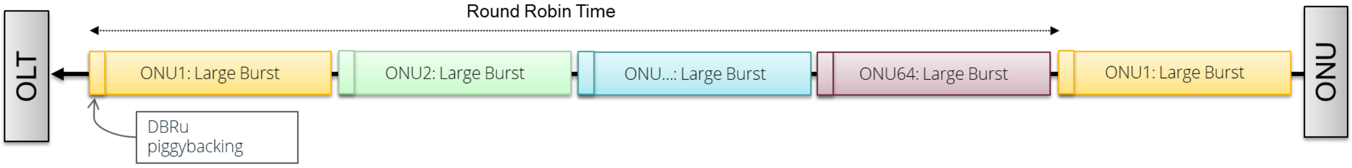

Figure 21 – Piggybacking and 64 ONUs transmitting

After an ONU has received the first grant of a burst, piggybacking will allow the OLT to accurately grant the remaining packets, see Figure 21. In this case, the latency is determined by the number of ONUs actively in the round robin at that time. In the best case, it is a single ONU. In the worst case, it is all ONUs requesting maximum burst sizes at the same time. In this case, the maximum latency is dominated by the number of ONUs and the maximum burst size allowed. Small bursts have a higher percentage of overhead to data (less efficient) with lower latency while large bursts have a lower percentage of overhead to data (more efficient) with greater latency. The continuation latency has similar equations as the start latency. The DBRu must travel upstream and the grant must traverse the entire PON. The big difference is the Round_Robin_Time that replaces the Polling_Interval and DBA_Delay. The DBA_Delay is often absorbed since the grant is known well before the Round_Robin_Time is available.

Cont_Fixed_Delay = Upstream_Flight_Time + PON_Round_Trip_Time + ONU_HW_Delay + OLT_HW_Delay

Cont_Variable_Delay = Round_Robin_Time + Discovery_Window

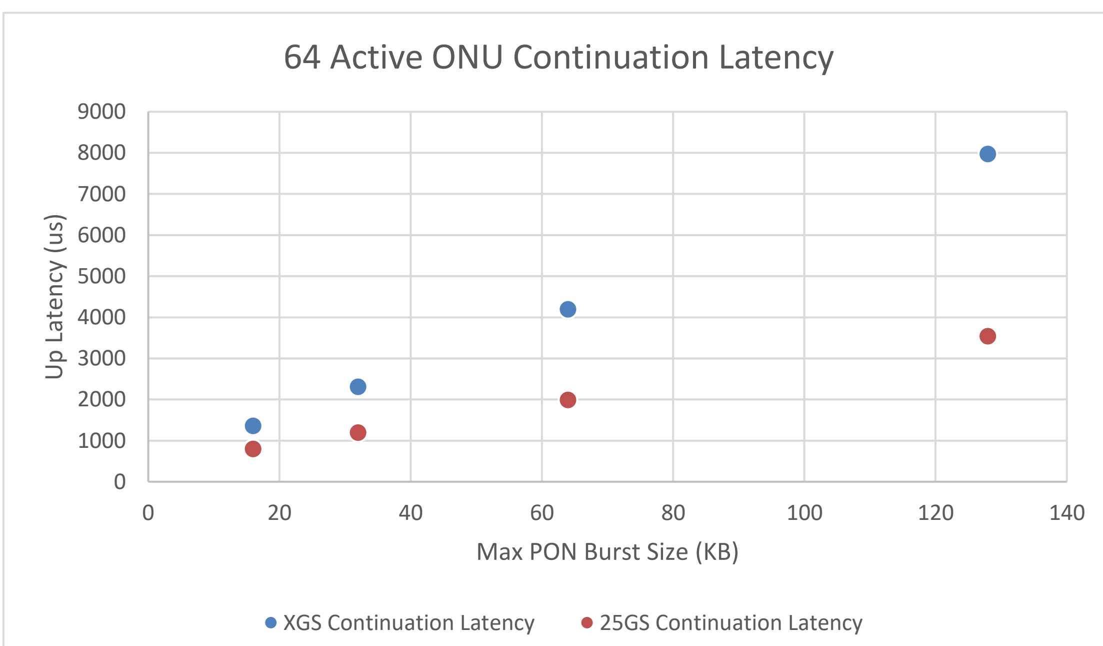

Figure 22 – 64 Active ONU Latency

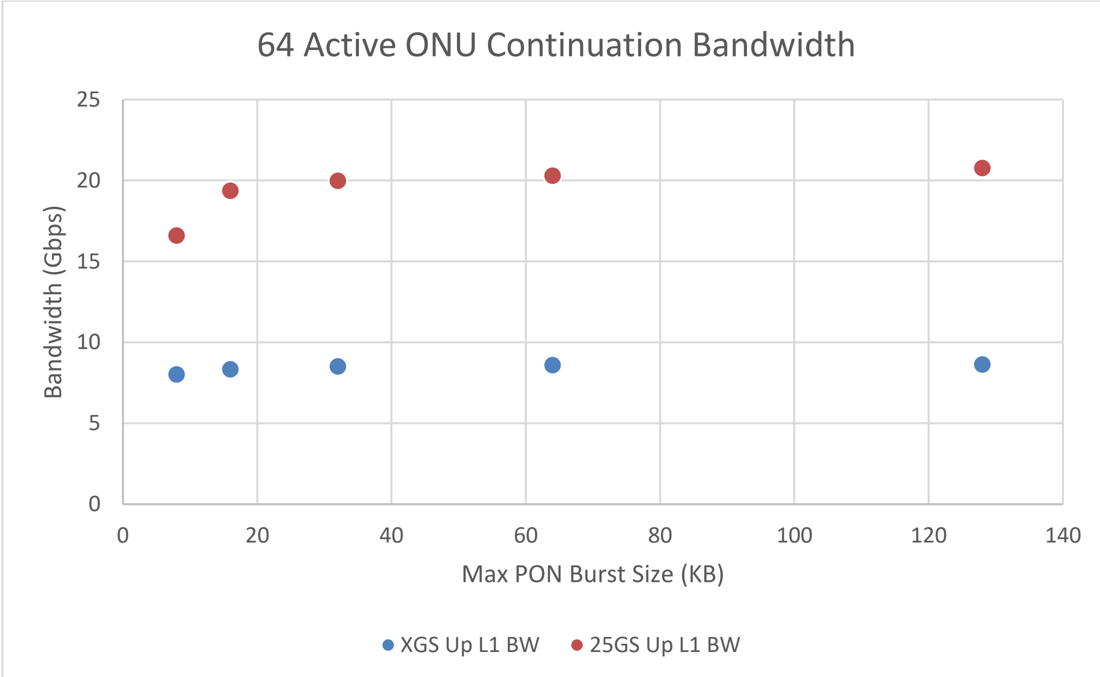

Figure 23 – 64 Active ONU Bandwidth

7.2.3. Queuing Delay

Queuing delay occurs when the input bandwidth exceeds the upstream bandwidth available. In the PON upstream case, it can occur when the subscriber exceeds their SLA or congestion limits the bandwidth available to a smaller number. Queuing delay performs an important function in TCP/IP networks. The delay on the acknowledge frame causes the next block of data to be delayed and thus slowed to a lower rate. Alternatively, a queue that overflows will cause a timeout and retransmit of a data block. Dropping frames has a much more significant impact to the user than the queuing delay. The queues in the ONUs should be sized to absorb large bursts of data and avoid drops. Priority queues can be used to allow higher priority traffic to avoid congestion. It is important to downstream performance to have a stable and low latency upstream. Variable upstream delays can cause the downstream bursts to be delayed and thus lower the throughput. By configuring a priority queue in the upstream for small acknowledge frames, it is possible to minimize the latency jitter to the downstream bursts. Queuing delay can be calculated by determining the difference in data rate between the input and output. If the difference in data rate is multiplied by the queue size, the maximum queuing delay before discard can be determined. The latency in this paper focuses on flows that have adjusted to the data rate and thus are below the maximum capacity of the SLA or total capacity. Since the queuing delay is a short-term factor required for bandwidth adjustment, it isn't something that could or should be removed from the system. A low max delay and low jitter delay when the subscriber is below the maximum capacity is the focal point of this analysis.

7.2.4. Balanced Delay

If the maximum start latency and the continuation latency are roughly the same, a single maximum latency covers both traffic conditions. Balancing the two delays to a common value allows for a consistent performance without wasting efficiency. For upstream traffic, the balancing of start latency and continuation latency is a good goal. To achieve a balanced maximum latency under 1ms in XGS or 25GS, the polling rate and maximum burst must be adjusted to smaller values. The amount of upstream bandwidth available is variable based on the number of ONUs bursting upstream. The table shows the lower of the two values for comparison.

Table 9 – Sub Millisecond Delay for 20km, 64 ONU PON

	XGS	25GS
Up Max Burst Size	9500 Bytes	22000 Bytes
Polling Interval	300 us	300 us
Start Latency	0.925 ms	0.925 ms
Continuation Latency	0.975 ms	0.975 ms
L1 Ethernet BW (1518B packet)	7.346 Gbps	16.739 Gbps
L7 Application BW (1518B packet)	6.992 Gbps	15.932 Gbps

Table 9shows that reaching sub millisecond max latency for 64 ONUs is possible in both XGS and 25GS. In both cases, a significant amount of upstream data is used to guarantee the lower latency. In XGS, almost 7 Gbps of application bandwidth remains while 25GS has almost 16 Gbps of upstream application data. The 25GS can still provide a 10Gbps symmetric commercial service along with the more asymmetric residential service.

7.2.5. XGS and 25GS Sharing the Fiber

In the earlier scenarios, a PON was either 64 XGS ONUs or 64 25GS ONUs. In many deployment scenarios, XGS and 25GS will exist side by side on the same fiber. In this case, it is possible for an operator to selectively move subscribers to a higher speed PON. If subscribers purchasing 5 Gbps, 10 Gbps, or higher SLAs were moved from XGS to 25GS, the number of subscribers on 25GS would be limited and the need for very high upstream bandwidth on XGS would be lessened. For example, the number of high bandwidth SLAs on the PON could be limited to 8 or 16 subscribers. Therefore, the maximum number of 25GS ONUs would be 8 or 16. In this case, the penalty for achieving the sub 1 millisecond latency on 25GS is much less. Table 10 shows an example of up to 8 25GS ONUs on a PON. In this case, the burst size can be significantly increased since the round robin will be only 8 ONUs. Additionally, the number of idle ONUs polling at 300µs is limited to 8 on the 25GS PON. With these two factors, the 25GS can achieve a 20 Gbps upstream Ethernet bandwidth with a sub millisecond worst case delay. With a total upstream bandwidth of 28 Gbps, the combination of XGS and 25GS could have a long-term future in the market. If the number of 25GS ONUs is increased to 16, the upstream drops by 400 Mbps but still stays above 20 Gbps of Ethernet BW, as shown in Table 11.

Table 10 – Shared PON with XGS (up to 64 ONUs) and 25GS (up to 8 ONUs)

	XGS	25GS	Total
Up Max PON Burst Size	9500 Bytes	180000 Bytes	
Polling Interval	300 µs	300 µs	
Start Latency	0.925 ms	0.925 ms	
Continuation Latency	0.975 ms	0.924 ms	

	XGS	25GS	Total
L1 Ethernet BW (1518B packet)	7.347 Gbps	20.818 Gbps	28.165 Gbps
L7 Application BW (1518B packet)	6.992 Gbps	19.802 Gbps	26.794 Gbps

Table 11 – Shared PON with XGS (up to 64 ONUs) and 25GS (up to 16 ONUs)

	XGS	25GS	Total
Up Max PON Burst Size	9500 Bytes	100000 Bytes	
Polling Interval	300 µs	300 µs	
Start Latency	0.925 ms	0.925 ms	
Continuation Latency	0.975 ms	0.988 ms	
L1 Ethernet BW (1518B packet)	7.347 Gbps	20.417 Gbps	27.764 Gbps
L7 Application BW (1518B packet)	6.992 Gbps	19.421 Gbps	26.413 Gbps

8. Conclusion

The IEEE 802.3 and a 25GS MSA group of 50+ companies standardized a 25 Gbps symmetric speed for PON access. 25 Gbps is the last PON speed that doesn't require a DSP, SOA, or EML at the ONU so it can be cost effective and low power. 25 Gbps is a useful speed for that reason. 50 Gbps and 100 Gbps will be available in the future at a higher cost and power. The 25GS standard is essentially 2.5 times the speed of the ITU-T XGS standard with the LDPC FEC defined by the IEEE 802.3. 25GS can be mixed with GPON or XGS on the same fiber plant allowing for a simple upgrade path. Equipment vendors offer simple upgrade paths for 25GS and a way to co-exist in the same box. 25GS allows operators to offer 5 Gbps and 10 Gbps symmetric services to customers. With all overhead considered, operators can expect to get approximately 20 Gbps in the downstream or upstream application bandwidth with 25GS. In addition to higher speed tiers, operators may choose to use the additional upstream bandwidth to lower the upstream latency. It is possible to achieve sub millisecond worst-case latency for 64 subscribers on the PON. By mixing XGS and 25GS on the PON, it is possible to achieve sub millisecond latency and 28 Gbps of upstream bandwidth. Based on the cost, ease of upgrade, simplicity, and additional bandwidth, 25 Gbps PON will provide value to operators looking for higher bandwidth and lower latency.

Abbreviations

25GS	25 Gbps symmetric PON defined by 25GS MSA
bps	bits per second
FEC	forward error correction
Gbps	1,000,000,000 bits per second
GPON	ITU-T Gigabit Passive Optical Network
IEEE	Institute of Electrical and Electronics Engineers
ITU-T	International Telecommunication Union Telecommunication
LDPC	Low-density parity-check
MSA	Multi-source agreement

N1	XGS/25GS 29 dB loss budget
N2	XGS/25GS 31 dB loss budget
OAM	Operation Administration Maintenance
OLT	Optical Line Terminal. Carrier side PON device
OMCI	ONT Management and Control Interface
ONT/ONU	Optical Network Terminal/Unit. Subscriber side PON device
PLOAM	Physical Layer OAM
PON	Passive Optical Network
RS	Reed Solomon
SCTE	Society of Cable Telecommunications Engineers
XGS	ITU-T 10 Gbps symmetric PON

Bibliography & References

25GS-PON Specification: 25 Gigabit Symmetric Passive Optical Network, 10 August 2021, Version 2.0

Institute of Electrical and Electronics Engineers (IEEE) standards:

- 802.3ca-2020: IEEE Standard for Ethernet Amendment 9: Physical Layer Specifications and Management Parameters for 25Gb/s and 50Gb/s Passive Optical Networks

International Telecommunications Union – Telecommunications Sector (ITU-T) Standards:

- G.652: Characteristics of a single-mode optical fibre and cable,
- G.984.2: Gigabit-capable Passive Optical Networks (G-PON): Physical Media Dependent (PMD) layer specification
- G.984.5: Gigabit-capable passive optical networks (G-PON): Enhancement band
- G.9701: Fast access to subscriber terminals (G.fast) - Physical layer specification
- G.9804.3: 50-Gigabit-capable passive optical networks (50G-PON): Physical media dependent (PMD) layer specification
- G.9807.1: 10-Gigabit-capable symmetric passive optical network (XGS-PON)

Journal of Optical Communications and Networking, Sept 2020; IEEE/Optica Publishing Group

