What Could You Do With 100 Gbps Coherent PON?

A technical paper prepared for presentation at SCTE TechExpo24

Edward W. Boyd
Vice President of PON R&D
Ciena Corporation
eboyd@ciena.com

John Bender
Wireline Architect
GCI Communication Corp
jbender@gci.com

Kevin Noll
Principal Architect
CableLabs
k.noll@cablelabs.com

James Harley
Sr Principal Architect
Ciena Corporation
jharley@gci.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

The broadband access industry has seen remarkable growth in its short 30-year life beginning in the late 1990's with technologies delivering 1.5Mbps. Today's broadband networks, delivering access at speeds over 10 Gigabits per second (Gbps), enrich the lives of billions of people around the world by enabling individuals to interactively share their experience and by removing the geographic and economic barriers to educational resources, health care, jobs, and so much more.

Anticipating a continuing need to increase capacity and speeds in the access network, the industry has begun work to create a 100 Gbps Passive Optical Network (PON). Some proposals for 100 Gbps PON focus on using tried-and-true Intensity Modulation-Direct Detection (IM-DD) technology that has enabled PON since the 1990s. Coherent optical transmission has matured significantly over the last decade creating an ecosystem that is primed to apply this technology in the cost-sensitive access network. Multiple standards organizations are hearing proposals for a 100 Gbps Coherent PON, but no specification is complete.

Solutions must have a low Optical Network Unit (ONU) cost, must operate on existing/legacy Optical Distribution Network (ODN), and must reuse existing PON standards where possible. It is also crucial to maximize compatibility with back-office systems, ensure smooth integration into network operations, and to enable reuse of existing system software and PON standards.

This paper will analyze the viability of a specific coherent PON physical layer (PHY) based on coherent transmission that operates seamlessly with the PON channel management protocols defined in ITU-T G.9804.2 and ONU management defined in G.988. The authors have direct PON experience as an operator, standards organization members, and technology vendors. The analysis will include relative cost, optical budget, system throughput, latency, jitter, and the applicability of coherent PON as a replacement for point-to-point fiber solutions used in backhaul, enterprise connectivity, and aggregation.

2. Coherent PON Standard

Currently, no standard exists for Coherent PON. Work, however, is occurring at CableLabs to produce a specification for 100 Gbps PON using coherent reception. ITU-T Q2/SG15 and Full-Service Access Network (FSAN) are also studying the prospect of 100 Gbps+ PON and coherent optics is being considered as an enabling technology.

The effort occurring within CableLabs has produced an architecture specification for Coherent PON. This specification does not detail the implementation of CPON but describes the use cases and requirements for subsequent specifications that follow the architectural framework. We can expect the CPON specifications to follow the typical CableLabs pattern to include a PHY specification and a MAC and Upper Layer Protocols Interface Specification (MULPI). The PHY specification is expected to describe the physical layer parameters and operation, including modulation formats, signaling rates, forward error correction, wavelength allocations, and other PHY-related requirements. The MULPI specification will describe CPON's MAC layer operation and parameters. These would include framing formats, adaptation of user Protocol Data Units (PDUs) to the PON, scheduling/granting protocol, PON maintenance, etc.

One of the primary goals would be to enable manufacturers to build Optical Line Terminals (OLTs) and ONUs that are interoperable under the CPON specifications. Another goal would be to create a specification that can be used across the entire broadband industry which will create cost advantages for all operators. This will likely translate into a specification that reuses components found in other PON standards like the MAC layer from IEEE 802.3 and the Transmission convergence (TC) Layer and management protocols from ITU-T G.9804.2.

3. PON Network Dimensions

Applying coherent optics to point-to-multipoint PON networks opens new options and flexibility. Coherent transmission offers access to higher spectral efficiency through advanced modulation schemes such as Quadrature Phase Shift Keying (QPSK) and Quadrature Amplitude Modulation (QAM) and by adding polarization as a new dimension to the modulation schemes. This means more bits per symbol relative to IM-DD and advances the system capacity from small fractions of a bps/Hz to over 1bps/Hz. At the same time, coherent detection improves receiver sensitivity, enabling the system to operate with higher losses in the ODN. This means longer distances and higher split ratios are achievable in a point-to-multipoint ODN.

Coherent systems use digital signal processing (DSP) to gain additional advantages like chromatic mode dispersion (CMD) and polarization mode dispersion (PMD) compensation. They also implement advanced error correction, equalization and signal recovery.

These improvements potentially change the assumptions under which network designers operate. For example, an increased power loss budget will increase the number of potential splits in the ODN. Increased power loss budget combined with improved PMD and CMD compensation enables longer fiber distances between the OLT and ONUs.

It is inevitable that the access network will have multiple technologies operating on it, each at its own wavelength(s). This means that CPON will need to coexist with those other optical users, like earlier generation Time Division Multiplexing (TDM) PON or Wave Division Multiplexing (WDM) PON. Coexistence is accomplished with a passive coexistence element (CEx) which is a specialized passive optical multiplexer.

Table 1: CPON Split Ratio and Distance

Split Ratio 1:N	Max Distance (km)	Event and CEx Losses (dB)	Splitter Loss (dB)
16	86.5	4	13.7
32	73	3.5	16.9
64	57.5	2.5	21
128	41.5	2.5	24.2
256	28	2	27.4
512	12	2	30.6

There are two primary coexistence scenarios anticipated for CPON. The first, depicted in Figure 2, is an overlay of CPON on existing Point to Multipoint (P2MP) ODNs that conform to legacy architectures reaching 20km or less and split ratios of no more than 1:128. The second scenario, depicted in Figure 3, is new architectures that take advantage of the longer reach, extending to 80km, that CPON will enable.

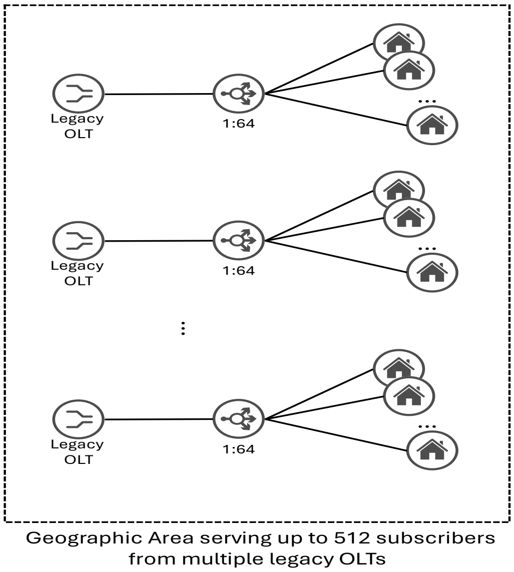

Figure 1: Residential PON Access Networks

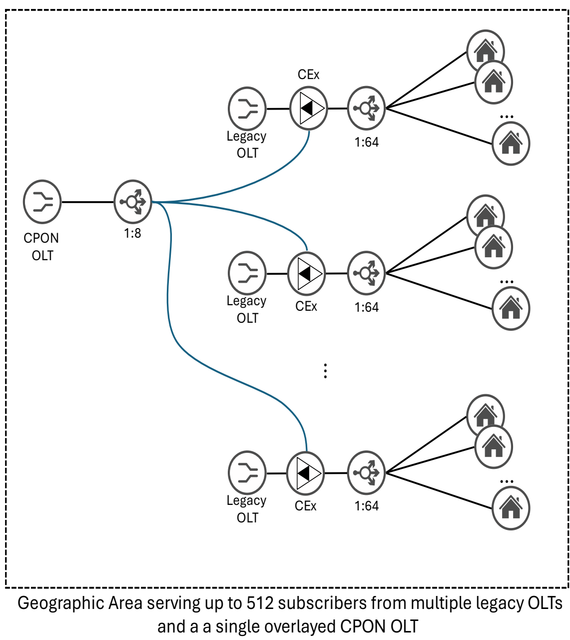

Figure 2: CPON overlay on a Residential PON

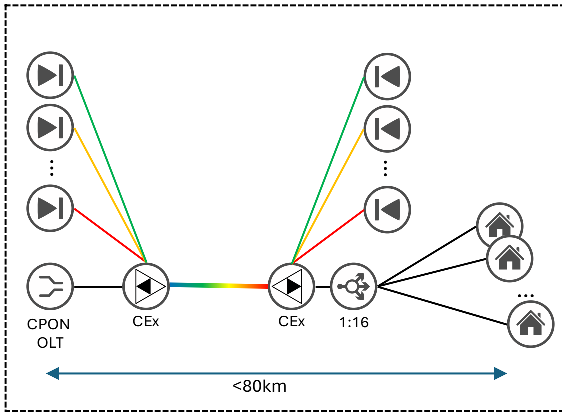

Geographic Area serving up to 16 subscribers from single CPON OLT overlayed with WDM applications

Figure 3: CPON in a WDM Network

Both scenarios introduce wavelength channel availability as a new dimension for planning. In the first scenario, this dimension should be minimized because legacy P2MP ODNs typically carry only traditional PONs which have fixed wavelength allocations based on international standards. However, the second scenario is likely to be deployed over longer fiber links that traditionally would be dedicated to DWDM point-to-point applications. This means that the operator will need to allocate wavelength usage on the long link in a way that can accommodate CPON's operating wavelengths. This long-distance architecture also reduces the available split ratio from 1:512 @10km to 1:16 @80km, meaning that fewer subscribers can be served from a single CPON OLT port and thus requiring multiple CPON OLTs to operate over the long-distance fiber link, and thus exacerbating the wavelength allocation concerns.

Introducing a tunable optical module in the CPON OLT and ONU could assist in solving this problem by making the CPON operating wavelength configurable. The historical perspective of tunability in PON systems, though, suggests that tunability will bring an unacceptably high cost. However, CPON introduces many new factors that might prove to make the higher cost acceptable or show that the cost of tunability in CPON does not conform to the historical trend. This is an area of continuing study within the CPON community.

4. Coherent PON System Cost Breakdown

To achieve 100G symmetric PON, we require the line rate to be increased to ~117 Gbps to compensate for Forward Error Correction (FEC) parity. A more detailed description of super-rating for FEC parity can be found in section 5.1.2. At that rate requirement, the anatomy of coherent PON transceiver can be designed to be low cost with fixed wavelength Distributed Feedback (DFB) lasers with no lockers, no Silicon Optical Amplifier (SOA) or Erbium-Doped Fiber Amplifier (EDFA), Silicon Photonics Transmitter [Quadrature Parallel Mach-Zehnder (QPMZ) structure] and Silicon Photonics Receiver (Heterodyne receiver with balanced photodetector), and low-cost DSP/PON ASIC. With true polarization muxing, the RF baseband electronics only needs a bandwidth of 15GHz since the line coded ~117 Gbps on 4 orthogonal lanes at baud rate of ~30G baud for In-Phase and Quadrature on X polarization and In-Phase and Quadrature on Y polarization. Figure 4 is an optical schematic of the CPON transceiver.

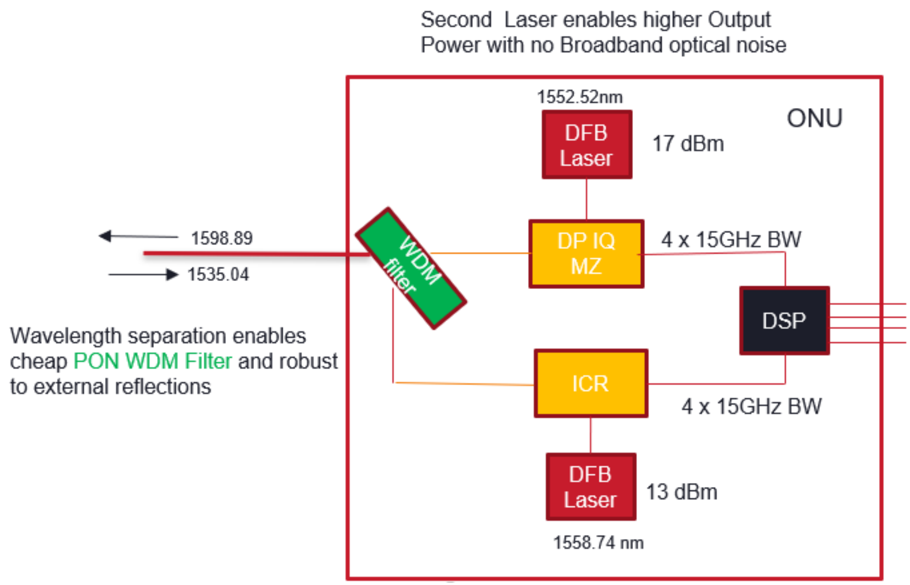

Figure 4: Anatomy of a Coherent PON Transceiver

The exact wavelengths have not been determined for CPON at this time. With CPON, it is possible to avoid the O-band (1260-1360) wavelengths that are already used by the early PON variations. For the solution proposed in this paper, the upstream and downstream wavelengths fall within the RF video band (1550nm-1560nm) which is available in many brown field PON deployments with access on a co-existence element port. Since IM-DD solutions need to use the O-band for the lower dispersion, they will conflict with the earlier generations of PON. Figure 5 shows how legacy PON has fully utilized the O-Band wavelengths.

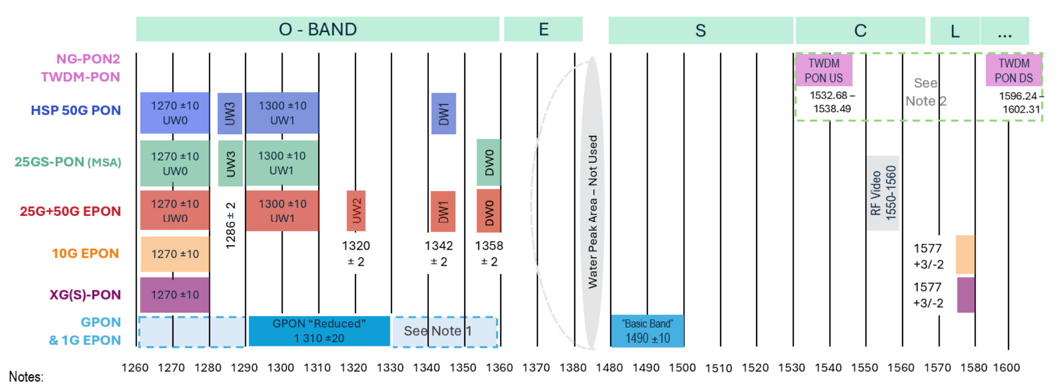

Notes:
1) Original GPON G.984.2 (2003) upstream lasers were wide ±50. G.984.5 (2007) provided US options: regular: ±50, reduced ±20 (shown), and narrow ±10 and Enhancement Bands 1 (not shown, in water peak area) and 2 (shown) for DS and US. 1G EPON IEEE 802.3ah (2004) aligned with the 1490 DS and 1310 US wavelengths as GPON with upstream lasers remaining wide ±50.
2) NG-PON2 (TWDM-PON, PtP WDM PON) G.989 series (2014). The TWDM PON DS and US bands each have 8 channels arranged in DS+US channel pairs and tunable ONT/ONU.

Figure 5: Wavelengths used in earlier PON technologies

In the literature there can be misleading statements about the relative cost of true 100G symmetric coherent versus the traditional IM-DD. It is true that traditional PON ONU transceivers at 10G are very cheap because only a Direct Modulated Laser and Photodetector is needed but at rates of 50G and higher the cost of IM-DD grows significantly with the addition of an Externally Modulated Laser, SOA and DSP ASIC in order to meet the link budget. Secondly, the current cost forecast for 400G ZR in QFSP-DD and for 100G ZR in QSFP28 with full-band tunable laser are not representative of the cost forecast for 100G CPON QSFP28 due to the use of fixed wavelength lasers and much higher volumes in the PON market relative to Metro Transport. For example, a fixed DFB laser is an order of magnitude cheaper than a full band tunable laser, which represents a significant fraction of the total module cost. Under these assumptions of similar volumes to existing PON technologies and the use of fixed lasers, we estimate that true 100G symmetric CPON ONU is only 20% more costly than 50G symmetric PON ONU. 100G CPON is therefore 50% cheaper on a relative $/bps basis versus symmetric 50G PON since 50G symmetric PON only carries 42 Gbps of Ethernet data while CPON (117 Gbps line rate) carries 100 Gbps of Ethernet data.

The low-cost coherent optics with a balanced photodetector described here has the powerful advantage to tune into a limited range of optical channels with minimal optical performance penalty and no additional cost. This attribute enables stacking different CPONs on different wavelengths on a PON network. The zero-cost limited tunability is estimated to be over 4x100GHz channels, implemented by tuning the thermal control on the DFB laser. If we add 4 CPON OLT ports to a 4x100GHz optical spectrum on a PON fiber, we can increase the average throughput of any ONU by a factor 4. This will additionally improve the CPON ONU relative at least cost $/bps versus symmetric 50G PON by a factor of ¼. Table 2 shows a possible CPON wavelength plan that can co-exist with legacy and provide 4 bidirectional channels of CPON.

Table 2: Possible CPON Wavelengths with 4 upstream (CU#) and 4 downstream (CD#) pairings

Channel	Central Frequency (THz)	Wavelength (nm)
CU1	193.10	1552.5244
CU2	193.20	1551.7208
CU3	193.30	1550.9180
CU4	193.40	1550.1161

Channel	Central Frequency (THz)	Wavelength (nm)
CD1	192.50	1557.3634
CD2	192.40	1558.1729
CD3	192.30	1558.9831
CD4	192.20	1559.7943

5. Performance Analysis

In addition to greater bandwidth, the Coherent PON standard can have some important additions that separate it from the earlier generations of PON technology. The performance analysis focuses on the Ethernet carrying capacity and latency for the upstream and downstream data. With this information, the use cases can be evaluated, and an operator can determine if performance is acceptable for the application.

5.1. Downstream Bandwidth

ITU-T based PON networks (i.e. GPON or XGS-PON) are based on SONET hierarchy of speeds. For 100 Gbps, a line rate of 99,532,800,000 bps would be the next logical step. In XGS PON, Ethernet Frames are carried in XGEM framing. This framing allows for segmentation and reassembly along with some additional information for encryption. The XGEM framing includes an 8-byte XGEM header and all or a segment of the Destination Address (DA) to CRC-32 Layer 2 Ethernet packet. In Ethernet, each packet has 8 bytes of preamble and 12 bytes of interpacket gap (IPG). Because of this difference, there is a 12-byte savings per packet between Ethernet and the ITU PON. In ITU PON, the framing sublayer (FS) has overhead for synchronization, physical layer OAM (PLOAM), and carrying the grants from the OLT to ONU in bandwidth maps (BW Map). Since the FS frame overhead is only every 125us, the bandwidth consumed by it is very small. At 100 Gbps, a large FS header would only consume 8 to 10 Mbps. This is a rounding error on a 100 Gbps PON. The per-packet savings of 12 bytes can be much more significant. Even with a very aggressive average packet size of 2000 bytes, the 12-byte savings adds over 594 Mbps of Ethernet capacity. Since 594 Mbps is greater than the 467 Mbps deficit from the SONET hierarchy speed plus 10 Mbps of FS overhead, the 100 Gbps Ethernet interface with an average packet size of 2000 bytes can be carried over the 99,532,800,000 bps ITU PON. If the average packet size is smaller, ITU PON capacity is increased compared to the Ethernet interface. If FEC wasn't required, the ITU PON would carry the 100 Gbps of Ethernet capacity implied in the name.

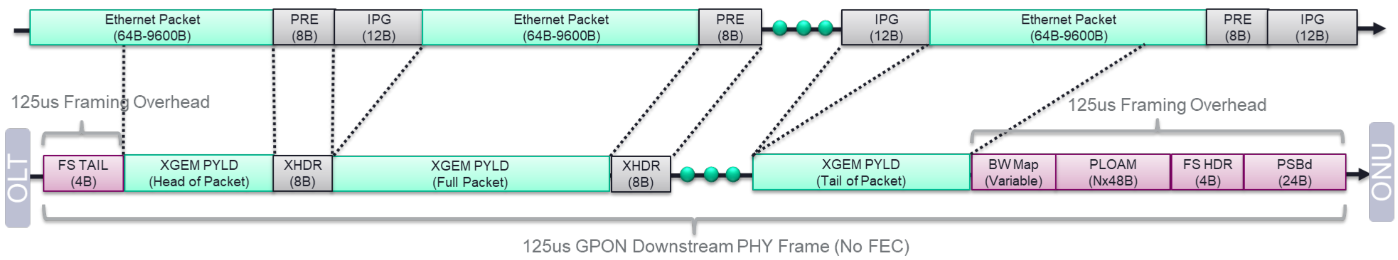

Figure 6: Ethernet to GPON Downstream Framing without Forward Error Correction

5.1.1. Sub-rating for FEC parity

The most significant reduction in the downstream bandwidth is due to FEC. In the previous generations of PON, the FEC parity bytes added to support FEC were sub-rated from the advertised line rate. When FEC is used, the parity took bandwidth away from the MAC layer causing the Ethernet capacity to decrease. Depending on the PON FEC generation, a reduction of 13% to 15% in the downstream capacity should be expected. Many operators are disappointed to find out the 10 Gbps PON only carries 8.7 Gbps of maximum size packets because of the bandwidth used for the FEC parity. For 1 Gbps/2 Gbps PON, FEC was optional and rarely used so it made sense to sub-rate when it was used. It was possible to use the full bandwidth or get a larger optical budget with lower capacity when with the FEC was enabled. With 10 Gbps PON, FEC was required by most operators to reach the desired 1:64 split ratio at 20 Km. In 25 Gbps and 50 Gbps PON, FEC is no longer optional, so the maximum bandwidth is 21 Gbps and 42 Gbps. It is desirable that the PON closely represent the carrying capacity of the matching Ethernet interfaces. For that reason, operators requested that 100G PON provide the ability to carry 100 Gbps of Ethernet frames. Unlike previous generations of PON, 100G Coherent should not sub-rate the PON for FEC parity.

5.1.2. Super-rating for the FEC parity

The FEC for CPON is not known as this time, so we will use the Low-Density Parity Check (LDPC) FEC from the ITU-T 50G standard for our super rating calculations in this paper. For every 1824 bytes of FEC payload, 336 bytes of FEC parity will be required to correct errors. With super-rating, the physical layer line rate can be increased to provide capacity for FEC parity without reducing the capacity from the MAC layer. The MAC Layer (Framing Sublayer in ITU terms) provides data at the SONET hierarchy line rate of 99,532,800,000 bps with the Downstream Physical Layer Synchronization Block (PSBd). It is identical to the earlier PON versions with the FEC disabled. In ITU PON, the downstream is framed into 125us blocks. This timing was required for 8KHz sampling of voice in early generations of telecom systems. For ITU PON, it is an important time reference for the time-of-day transport from the OLT to the ONU so the 125us time reference must be preserved. With Sub-rating, the FEC parity would be inserted into the 125us frame and XGEM payload capacity would be reduced. If we consider the ITU 50G PON FEC at 100 Gbps, the capacity will have 1,555,200 bytes every 125us with parity consuming 241,920 bytes (~15.5%) leaving 1,313,280 bytes. For super-rating, the parity is inserted into the 125us FS frame, but the line rate is increased by the exact amount of additional parity bits required. This allows the 125us frame time to be preserved and the true MAC rate to be achieved. In the 125us frame, the 1,555,200 bytes of payload need 286,608 parity bytes to be added. To absorb that capacity, the line rate is increased to 117,875,712,000 bps. At the physical layer input and the physical layer output, the 125us framing is preserved along with the ability to carry 100 Gbps. The higher line rate does have an impact on the optical budget. A simple calculation shows roughly 0.7dB of optic budget penalty for this increase. This is a small amount to recover roughly 15.5% of Ethernet capacity.

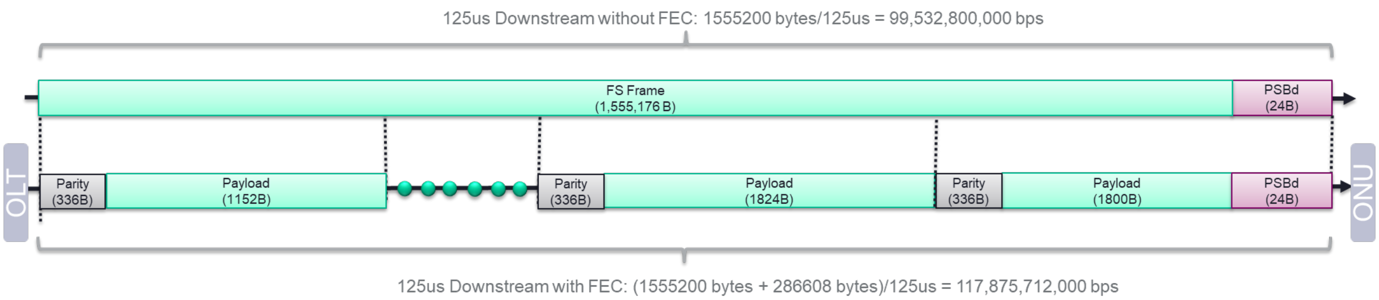

Figure 7: Super-rated Downstream Framing

5.2. Downstream Latency

The downstream latency in PON is very similar to a point-to-point link. The fiber flight time can be calculated with the rule of thumb to get 100 microseconds for 20km fiber. Coherent PON promises reaches of up to 80km. At 80km, the downstream flight time would be 400 microseconds. Bridging from Ethernet to PON at the OLT and ONU could add 10 to 20 microseconds to the latency. This paper doesn't consider the bridging latency since it is highly dependent on the vendor implementation.

5.3. Upstream Bandwidth and Latency

The upstream bandwidth and latency are more difficult to calculate. In addition to the framing and FEC parity overhead found in the downstream, the upstream requires bandwidth for Time Division Multiple Access (TDMA) bursting and polling for upstream queue status. The super-rating only compensates for the FEC parity overhead so the upstream will be reduced from the 100 Gbps Ethernet if the upstream overheads aren't minimized.

5.3.1. Upstream Burst Overhead

A shared TDMA upstream requires a per-burst overhead. The time required can be broken down into a few key components. Between bursts, there is time required for the ONU to power the laser ON/OFF along with a dead time between slots for any jitter in the slot timing. The start of the burst requires some number of bits to determine the signal level with automatic gain control (AGC) and additional bits for clock and data recovery (CDR) to get the frequency/phase alignment. In ITU PON, these overheads can be grouped as either a preamble sequence of bits (Laser ON, AGC, CDR) or a guard time (Laser OFF and slot jitter).

The recent PON standards have these as configurable values sent for the OLT to the ONU during initial discovery. These values often reduce as a PON technology matures. In the case of Coherent PON, it is still early to determine these values. For that reason, we will consider an aggressive and a conservative value. For units, we will assume upstream time slots (125us/9720=~12.8ns). For the aggressive numbers, we will use 2 time slots of preamble and 2 time slots of guard time. For the conservative numbers, we will use 8 time slots of preamble and 8 time slots of guard time.

Figure 8: Anatomy of an Upstream Burst

5.3.2. *Discovery*

New ONUs to the PON must be discovered by the OLT. The OLT grants a discovery slot to ONUs that are not registered. Since multiple ONUs could respond, the discovery slot has a random offset. To accommodate the random offset along with ONU latency, a 50-microsecond slot is required. Since the distance from the OLT isn't known, the OLT provides a guard band around the slot to accommodate the shortest or longest possible fiber distance. For most PONs, 0 to 20km is the default distance. With 100us flight time down and 100us flight time up, a 200us guard band is added to the discovery slot. In general, a discovery slot is 2xFlight_Time+50us. In most PONs, this results in a 250us discovery slot. The discovery slot is not very often. A discovery slot every 3 seconds is quite common. In terms of the bandwidth penalty, it is an insignificant 0.0083%. In terms of latency, a jitter of 250us can have a significant impact on applications with tight latency requirements. Coherent PON promises reaches up to 80km. If an operator wants to fully auto discover from 0 to the maximum distance, the discovery slot jitter impact will jump to 450us for 40km PON and 850us for 80km PON. Since many applications for Coherent PON are latency sensitive replacements for point-to-point Ethernet, the discovery slot should be addressed.

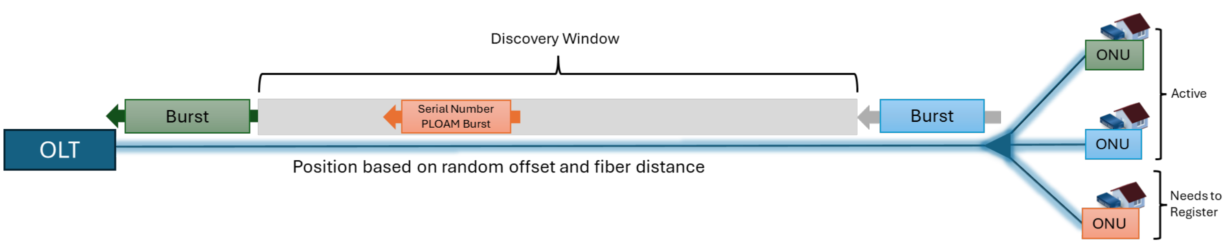

Figure 9: Traditional MAC Layer Discovery

5.3.2.1. *PHY Layer Discovery*

PHY Layer Discovery has been proposed as a solution for the discovery slot issue. The current PON standards use MAC Layer Discovery. The MAC Layer is paused to allow for a discovery slot. The pausing of the MAC Layer causes jitter in the upstream data stream. If the ONU could be discovered at the PHY Layer and not disrupt the MAC Layer, the jitter would be removed. An out of band channel in the PHY Layer would allow for parallel operation of the MAC Layer data path and a path for the PHYs to communicate. It is easy to imagine 3 options for carrying the out of band channel. The second signal could be added with Amplitude Modulation(AM), Frequency Modulation(FM), or wavelength variation. This paper won't explore these options, but the authors believe that one of these solutions will be possible to create the out-of-band channel. In all cases, the out-of-band channel can have a low speed with a long symbol time.

PHY Layer Discovery doesn't require a grant from the OLT. When an ONU is not discovered, it starts sending a short random offset data burst (beacon) on the out of band channel. The random offset allows for multiple ONUs to be discovered at the same time and resolve collisions. The beacon message contains the ONU's ID, and a timestamp based on the downstream time reference. When the OLT PHY receives a beacon, it measures the difference in the current downstream time reference and the timestamp from the ONU. With this information, the PHY Layer has the ONU ID and distance. The OLT MAC Layer can read the PHY and get the required information to add an ONU. PHY Layer discovery has the advantage of no MAC Layer disruption and the ability to discovery ONUs at any distance if the signal can be received.

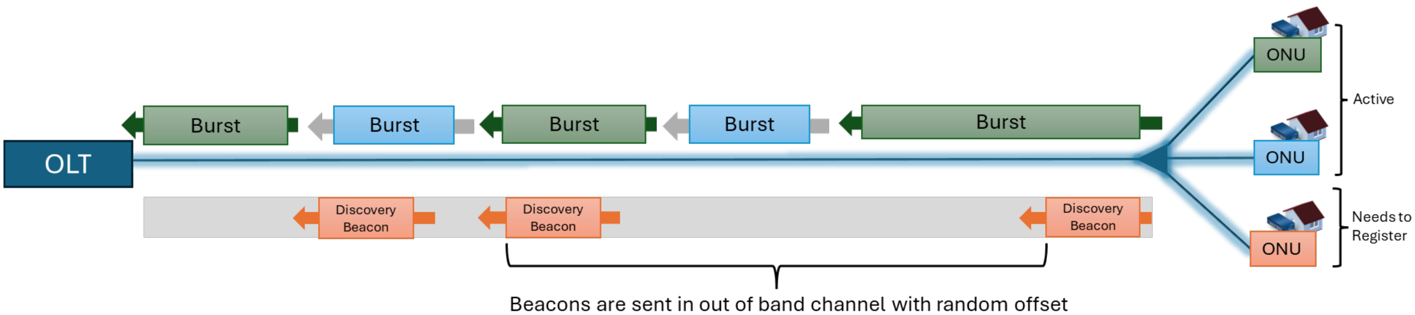

Figure 10: PHY Layer Discovery in Out-Of-Band Channel

5.3.3. Unsolicited Granting

The allocation of upstream slots can be very dynamic based on queue status or somewhat fixed based on a provisioned bandwidth. Unsolicited granting is a fixed size grant at a fixed interval. This method provides low latency but uses resources when not required. When the goal is the lowest latency and predictable performance, unsolicited granting makes sense.

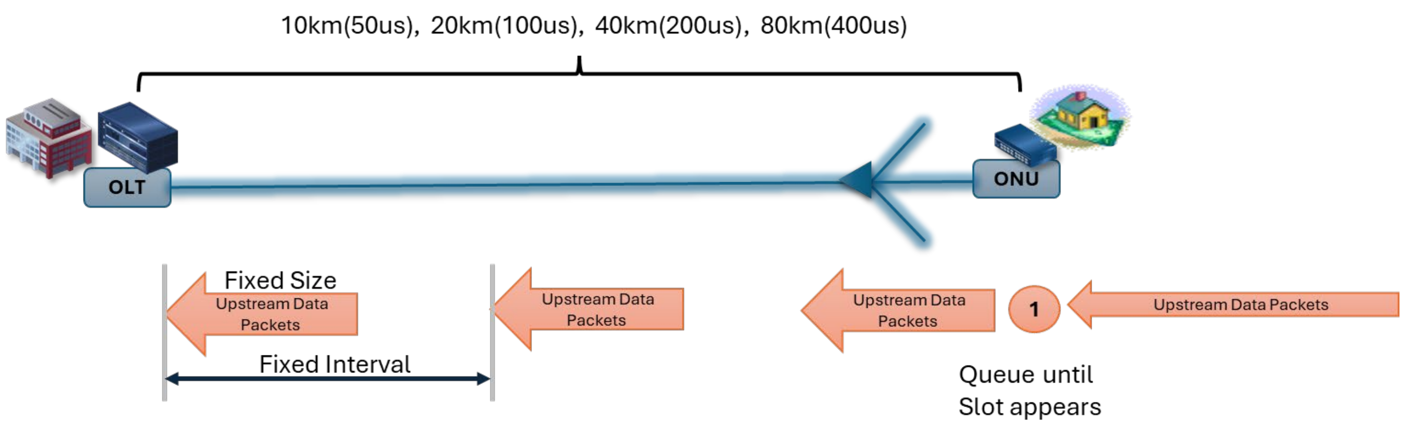

Figure 11: Unsolicited Granting Flow

For the simplest analysis, we will consider 100% of the PON as unsolicited granting and at the same speed. We will focus on a PON serving 4 ONUs with 25 Gbps each and another PON serving 10 ONUs with 10 Gbps each. All ONUs have the same priority. Of course, it is possible to mix different speeds, different ONU counts, and priorities but these simple cases can provide a guide for the performance. With discovery moved to out of band and unsolicited only, the OLT granting can be a simple round robin. We created a spreadsheet model for the round robin unsolicited to show the bandwidth and latency for different burst overheads and burst sizes. If each ONU transmits a large burst in the round robin, the burst overhead will be a lower percentage of the bandwidth, but the round robin will take a long time between service time to a single ONU. If the ONU burst is small, the latency for the round robin will be small but the overhead will be a higher percentage compared to the data. If the goal is 100 Gbps Ethernet BW upstream, the table below shows the burst size required and the latency in the round robin. The Ethernet

IPG savings is the source of the extra bandwidth in the upstream. The table shows the effect of the aggressive and conservative burst overheads on the results. The burst size required to achieve 100 Gbps is directly tied to the overhead. The number of ONUs and bandwidth of each ONU is relevant. Fewer ONUs with higher bandwidth dramatically reduces the latency.

The CPON upstream latency with unsolicited granting can be calculated by adding the round robin loop time to the upstream flight delay. For a 20km PON and the 4x25 Gbps load, it is roughly 100us upstream flight time and 72us of round robin time with the aggressive overhead. In this case, a worst-case upstream latency of 172us could be supported. Table 3 shows the latency calculations for the two scenarios and the 2 overheads. The latency numbers show the value of the PHY Layer discovery over MAC Layer discovery. With MAC Layer discovery, the latency could increase from 172us by 250us to 850us. This increase is very significant for latency sensitive applications. If less than 100 Gbps Ethernet is needed, it would be possible to decrease the latency even further by granting smaller block sizes and reducing the round robin time. These results show that 100 Gbps Coherent PON with PHY Layer Discovery and Super-rating can provide a muxed 100 Gbps upstream of virtual point-to-point links. It is a viable replacement for point-to-point for almost all but the tightest applications.

Table 3: Unsolicited Upstream Latency for Full 100 Gbps Ethernet Bandwidth

# of ONUs	BW per ONU	Burst Overhead	Average Ethernet Packet Size	Burst Size required for 100 Gbps Ethernet	Round Robin Loop Time	20km PON Upstream Latency	80km PON Upstream Latency
4	25 Gbps	2 & 2 (Aggressive)	1500 Bytes	300,000 Bytes	72us	172us	472us
4	25 Gbps	8 & 8 (Conservative)	1500 Bytes	900,000 Bytes	192us	292us	592us
10	10 Gbps	2 & 2 (Aggressive)	1500 Bytes	300,000 Bytes	217us	317us	617us
10	10 Gbps	8 & 8 (Conservative)	1500 Bytes	900,000 Bytes	653us	753us	1053us

5.3.4. Solicited Granting

With solicited granting, the ONU is only granted the upstream when it has data to send. We This allows for a statistically multiplexing of the upstream to support oversubscription. The granting simply follows the need up to the service level agreement or the full capacity of the PON upstream. The cost for solicited granting is latency and polling bandwidth. The OLT grants a small slot to get the ONU's queue status periodically (shown as a DBRu below). After receiving a non-zero queue status, the OLT grants the ONU a large upstream slot to move data. Finally, the OLT sends the packets upstream. With the upstream burst, the OLT receives a DBRu with an updated queue value. Bursts continue until a DBRu shows a zero DBRu. Polling will start and continue until a non-zero DBRu is found.

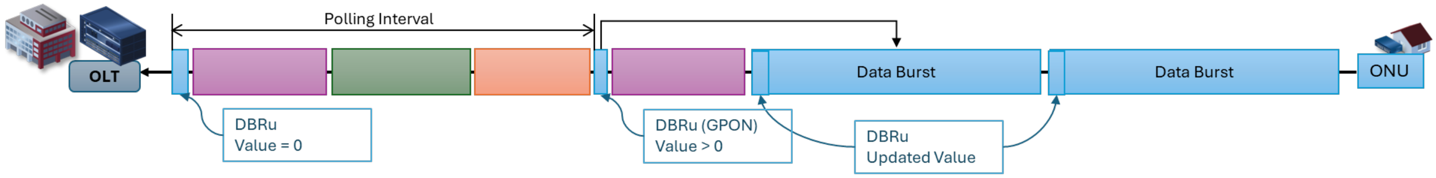

Figure 12: Solicited Granting Flow

Latency for solicited operation can be broken down into 2 scenarios: Start Latency and Continuation Latency. The Start Latency occurs when an ONU is idle, and packets arrive. The Start Latency is dominated by the flight time and the polling interval.

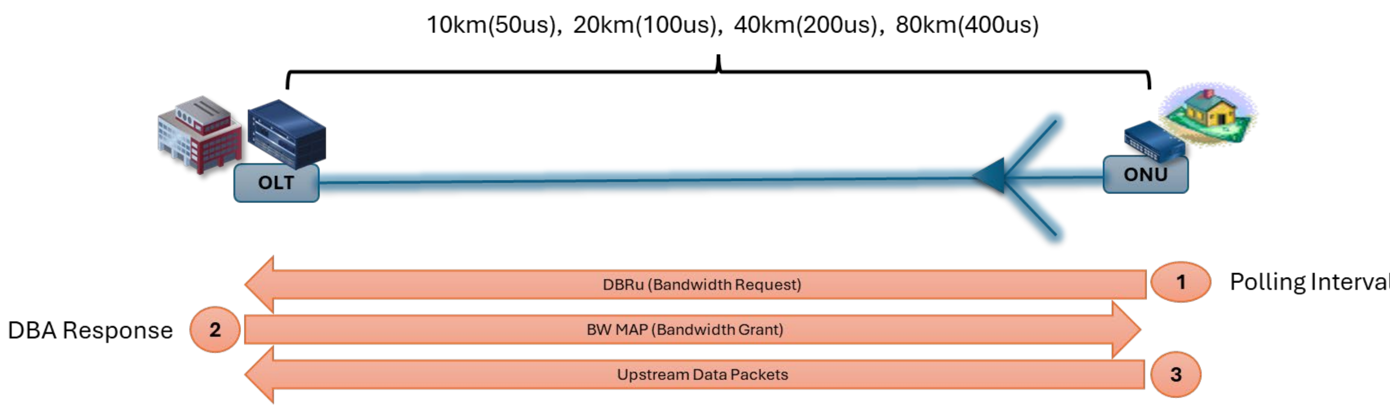

Figure 13: Solicited Mode Start Latency

The example calculations of the start latency show the impact of the polling and flight time. The 500us polling rate is an aggressive number to get the lower latency. Since solicited mode requires information to transverse the PON three times, the impact is tripled compared to unsolicited mode. In some cases, this effect is mitigated by mixing unsolicited and solicited but increasing the polling size. In this way, a portion of data is low latency, but higher bandwidth peaks can be supported with solicited mode. The ability to support both by configuration is a significant advantage of PON over a pure point-to-point link.

Table 4: Solicited Granting Start Latency Calculation

	20km PON	80km PON
Polling Interval	500us	500us
DBRu Flight time	100us	400us
DBA Response	125us	125us
BW MAP Flight time	100us	400us
ONU Response	50us	50us
Data Packet Flight time	100us	400us

Total Upstream Latency	975us	1875us

The Continuation Latency occurs when consecutive bursts are sent without polling. Like unsolicited mode, the round robin of multiple ONUs will have a big impact on the latency. Continuation Latency is dominated by the flight time, the number of actively transmitting ONUs, and maximum burst size per ONU. A simple example of 10 ONUs sharing a 100 Gbps was used to show the possible performance. Even with the conservative burst overhead, it is possible to reach 98.5 Gbps and a sub-millisecond worst case latency for both start and continuation. There are many other possible DBA modes of operation but these simple round-robin, single priority solutions are good representations of the worst-case performance for a loaded network.

Table 5: Solicited Granting Continuation Latency Calculation for 10 ONUs, Conservative Overhead, 600KB block

	20km PON	80km PON
DBRu Flight time	100us	400us
DBA Response	125us	125us
Round Robin Time (10 ONUs)	488us	488us
BW MAP Flight time	100us	400us
ONU Response	50us	50us
Data Packet Flight time	100us	400us
Total Upstream Latency	963us	1863us
Total Upstream BW (one active)	98,573,089,596 bps	98,573,089,596 bps
Total Upstream BW (all active)	98,927,240,705 bps	98,927,240,705 bps

If 100 Gbps Coherent PON was used to replace a full 1:64 PON of ONUs, the performance is still 96.3 Gbps with a single ONU transmitting and 63 polling at 500us and the conservative overhead.

5.4. Four Wavelengths of 100 Gbps

With no ONU cost penalty, it is possible to specify CPON as 4 wavelengths tunable. With this capability, the fiber plant capacities shown above can be multiplied by 4. The CPON capacity on a single optical network could reach 400 Gbps of Ethernet data. Multiple channels allow for fewer ONUs to share a PON OLT port, resulting in higher efficiency and lower latency. In addition to the 4 wavelengths of CPON, traditional GPON, XGS, or 25GS could also share the fiber plant. Low bandwidth and latency insensitive services can stay on the legacy PON while CPON can be focused on the ONUs requiring strict latency or very high bandwidth.

6. Applications and Use Cases

6.1. One Network to Rule Them All (Advantages of a Common Service Activation, Management, and Operations System)

The coherent PON architecture described above provides significantly more bandwidth than existing PON systems, and, along with flexible DBA and the ability to coexist over the same physical infrastructure as an existing DWDM channels allows for both an expansion of existing use cases as well as new use cases previously requiring point-to-point (PTP) or a DWDM channel. A primary advantage to this approach is in providing the operator with a single network system to manage, provision, and operate. Field technicians require a single set of skills and tools to troubleshoot all services, whether they be residential internet access, commercial ethernet services, or mobile backhaul applications.

6.2. Fully Utilize the Fiber Plant Deployed for Residential

An operator that has already deployed a prior generation of PON technology will be able to reuse the network already constructed to support residential services to also provide network access for small cell and nano cell backhaul service as well as macro cell towers co-located within the PON fiber footprint. Where network slicing is deployed for mobility, these slices can be mapped to L2 service-flows replicating the QoS and routing requirements across the PON access, while the same infrastructure continues to provide highspeed internet access to new and existing customers.

Using coexistence elements, existing customers can continue to use 10G EPON or XGS-PON while high usage customers can be seamless migrated to the higher bandwidth CPON system just by replacing the CPE ONT and provisioning it with the required services. New customers can be directly added to the CPON with no impact to existing customers.

6.3. Enterprise – How Customers Would Use It.

Enterprise customers will benefit from the ubiquity of the fiber infrastructure that provides differentiated quality of service to support service level agreements, private networking features, and value-added services. Unlike traditional service delivery over point-to-point Ethernet circuits, branch offices and remote locations can be rapidly added to a customer's network via centralized provisioning and activated remotely. Additional capacity can be added via provisioning, either by a customer service agent, or through self-provisioning and service activation through a customer-facing service portal. The additional capacity of 100G CPON allows for a greater density of higher speed services to be provided over a single interface over the shared infrastructure. The flexible QoS mechanisms allow enterprise customers to use the same delivery method to deliver mission critical applications and services to remote facilities, provide data-center connectivity for geo-redundant operations and provide public facing services to customers and employees.

6.4. Wireless Backhaul

Wireless backhaul requires high throughput service with consistent latency. The requirements are based on the number of radios per tower and the radio technology used. A macro cell tower typically served with a 10G Ethernet connection point-to-point Ethernet connection today, would benefit from the flexibility and expandability of a 100G CPON network. Additional throughput can be added via provisioning calls, rather than a circuit redesign. Similarly small cells have the additional benefit of sharing the ODN infrastructure with businesses and residents in the local neighborhood. When compared to point-to-point, the shared nature of the PON ODN allows for rapid turn up and deployment of tactically placed small cell sites. The increased throughput and flexible bandwidth management that 100G CPON provides allows wireless providers and operators who provide connectivity to wireless providers to more fully realize the capabilities of their outside plant, by mixing best-effort, latency insensitive traffic, with higher priority cell data traffic and thus limiting construction costs and time to market delays.

6.5. DAA Architecture with CPON to Backhaul Nodes

Similarly, Distributer Access Architecture (DAA) requires high throughput, low latency and jitter-controlled service to transport traffic across a Converged Interconnect Network (CIN). Additionally, some DAA architectures based around Remote PHY (R-PHY) require tight control of timing to synchronize clocks for PHY layer processing at a centralized location. This requires the ability to pass PTP traffic across the network with low latency queues. By using CPON, as part of a hierarchical timing network, the need to locate additional Stratum one grandmaster clocks is reduced, thus limiting exposure to potential Global System for Mobile Communication (GSM) risks to the timing infrastructure.

6.6. Eliminate Software Defined Wide Area Network (SD-WAN)

SD-WAN has become a popular tool for enterprises to extend their private corporate network across the public Internet. But SD-WAN has limitations; since all traffic is best-effort, there is no mechanism to grant higher priority to some traffic over traffic. The enterprise also becomes subject to Distributed Denial of Service (DDOS) attacks launched not only against their own infrastructure, but that of their provider. Lastly, SD-WAN is subject to routing policies that are optimized for internet traffic flows, rather than the most efficient routing between the customers sites.

CPON provides both the throughput and Quality of Service (QoS) requirements for enterprise grade private networking, as well as providing customized network slices with defined QoS and Service Level Agreement (SLA) parameters between sites. Operators can use CPON to provide access to Layer-2 and Layer-3 Ethernet Virtual Private Network (EVPN) services that can extend across the operator's network topology, providing optimized routing between customer sites. Furthermore, these Virtual Private Networks (VPNs) can be monitored by standard Ethernet and IP OAM protocols such as Y.1731 to measure and report on key performance indicators such as packet loss and latency that may affect a singular service in near real time.

6.7. Residential – Virtual Point-to-Point Plus High BW

As distributed workforces and work-from-home continues to be a common business model, the ability to provide higher speed connections to either the public internet or to private corporate resources continues to be important. Furthermore, some applications, such as remote medical imaging, small business point-of-sale processing, and school-from-home can benefit from prioritized QoS services to increase productivity and provide better customer and employee satisfaction. In these cases, a customer may require a virtual point to point circuit between a centralized hub site and an employee's home or other locations that may not be considered a typical business location. CPON provides both the throughput for these applications, and the ability to provision multiple services to the same location without a truck roll or having to design new circuits. Like the SD-WAN use case above, the different services retain their own QoS and security profiles. The available throughput of 100G CPON allows for the higher priority virtual point-to-point service to minimally impact the best-effort residential service.

6.8. Use Higher Link Budget to Serve Multiple PON ODN Networks (Higher Split Ratio)

The higher split ratio of up to 512 allows a smaller set of fibers to serve large MTU buildings in either a commercial office or residential apartment building. This application also allows for a reduction in power and HVAC load by enabling Passive Optical LAN service, replacing multiport Ethernet switches with passive optical splitters in building wiring closets and then extending fiber either directly to desktop ONUs or to ceiling mounted access points with integrated ONUs.

Even without a high concentration of multiple tenant unit (MTU) buildings, the higher split ratio still affords additional flexibility for new construction or for the subdivision of existing buildings into additional units. For example, an existing ODN with a more moderate split ratio, say 64 to 1, could allow additional units or buildings to be served later, simply by adding splitters as needed to meet the customer demand.

6.9. Lower Speed PON (10G/25G) or DOCSIS®/DSL Off of a Overlayed CPON From the Main PON ODN

The ability for 100G CPON to coexist with existing PON or optical technology is an intrinsic characteristic of the system. 100G CPON is designed to be backward compatible with existing PONs, utilizing different wavelengths, so that an operator can make use of their existing investment in their optical distribution network (ODN), as they install new CPON OLTs across their footprint. This means that there is no need for a flag-day, where all customers migrate from the old technology to the new. Rather, the older lower speed PON or hybrid fiber coax network will continue to operate with CPON acting as ships-in-the-night across the same infrastructure. Customers, or network infrastructure can be moved from one system to another simply by provisioning services on a new ONU and physically attaching it to the existing ODN.

6.10. No Need to Pull Additional Point-to-Point

As noted above, CPON alleviates the need for most point-to-point services. This saves fiber and allows for quick service delivery if the customer premises is already on-net, and allows for additional services, whether point-to-point or multi-point to be added later as customer demand changes. Having a relatively excessive amount of throughput available allows the CPON network to better meet future throughput demands. From a plant management perspective this brings efficiencies in documentation effort as fewer physical circuits must be added and removed from plant records. Operationally it allows for out-of-band testing on a distinct service that uses the same physical delivery. Lastly, the use of common facilities and end user equipment for point-to-point and multipoint service delivery means that field technicians can leverage common practices and utilize expertise gained from each customer, rather than treating each point-to-point as a special circuit requiring unique installation skills and troubleshooting methodology.

In cases where true point-to-point links are really required, CPON may coexist with point-to-point links over a DWDM system. By choosing channels that do not conflict with those defined for CPON service, the same ODN can be used to provide both CPON and point-to-point DWDM if desired.

6.11. Small and Medium Sized Business (SMB)

The SMB market is often a competitive one for most providers. These customers do not require, or have the budget, for the same level of complex services that enterprises often require, but they do often have service availability and other requirements that exceed those of a best-effort residential service. Often SMBs have symmetric upstream requirements, and as mentioned above, may benefit from some form of operator hosted private networking service. The ease of provisioning multiple services over the same infrastructure allows possibilities for a customer-facing provisioning portal where an SMB customer could order and activate new private services as easily as they can spin up compute resources on a cloud platform is a potential new revenue source for operators.

6.12. Enable "Stacking"

Wavelength "stacking" refers to the ability to provide additional channels of 100G CPON across the same ODN. This allows for the same OLT and ODN to provide more than an aggregate of 100G of service. ONUs could be dynamically assigned to one of four (or more) channels to reduce contention for resources, or to provide additional flexibility for network slicing for QoS or other considerations. Alternatively, in markets that support or require competitive access to an operator's ODN, one or more of the channels may be assigned to competitive operators, thus providing them with an independent domain for their own customers. The operator of the ODN would retain visibility over the network, all the benefits of using common practices and methods of procedure used for a single operator would still apply, but the competitive carrier would retain complete control of the services provisioned over their channels. Once a customer is "on-net" on the CPON ODN, they can be easily assigned between competitive carriers through provisioning and OAM messages that assign an ONU to the required channels.

6.13. CPON as Transport: Augment / Replace DWDM

Just as CPON enables operators to replace point-to-point Ethernet connections for customers, CPON can also be used as part of the core transport network to augment or replace existing uses of DWDM. DWDM while providing bandwidth is not flexible. There are only a finite number of channels that can be assigned across the system. Special colored optics must be used and moving a circuit requires a truck roll or CO visit to physically replace jumpers. Conversely, the throughput provided by CPON can in some cases be used to flexibly multiplex an arbitrary number of circuits, with the option to provide a more exact committed information rate to each service, if required. In the stacking scenario mentioned above, the aggregate throughput may be 400G, 800G or more, with transport services provisioned along with customer access services as required. No truck rolls are needed to move or groom these virtual circuit services. As noted, these same CPON services can co-exist with an existing DWDM plant for existing circuits that have yet to be migrated, or which have service requirements that could better be met with true DWDM.

7. Conclusion

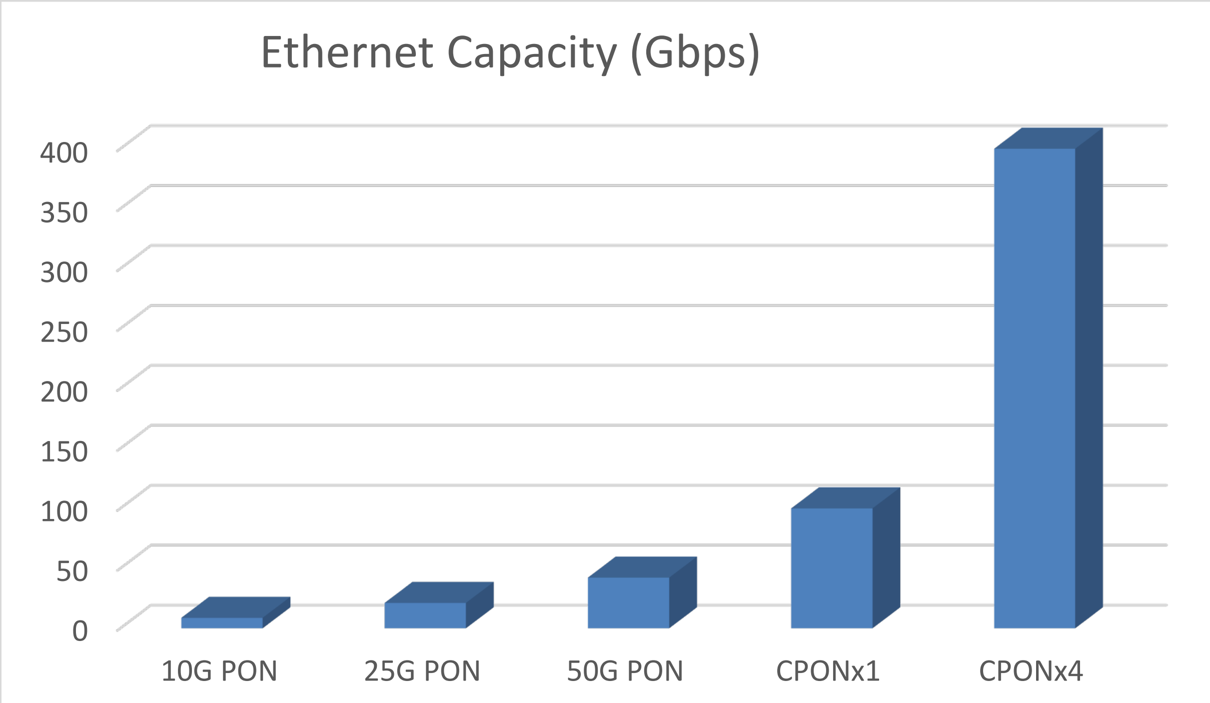

Figure 14: Downstream Ethernet Capacity by PON Type

We believe that Coherent PON can be a bold step forward in the evolution of access network technology. In this paper, we outlined a possible Coherent PON solution and the implications. It allows PON speeds to reach 100 Gigabits per second and a full 400 Gbps on a single ODN. Coherent technology allows for a larger optical budget, longer reach, and the ability to utilize the C-band for co-existence with legacy PON. Coherent PON can provide a cost-effective ONU and lower cost than IM-DD solutions on $/bps basis. The ability to super-rate and provide a full 100 Gbps of Ethernet payload allows for CPON to match the speed of a standard Ethernet interface. Physical Layer or out-of-band discovery can help remove a significant source of jitter and delay in current PON systems. With this performance, the PON access networks with CPON can handle many more uses cases that were reserved for point-to-point links. CPON could be applied to WDM fiber plants or access fiber plants to expand the services for an operator. As the standards for this technology are created, we hope that it can reach the full potential that we outlined in this paper.

Abbreviations

AGC	Automatic Gain Control
AM	Amplitude Modulation
ASIC	Application Specific Integrated Circuit
bps	bits per second
CDR	Clock and Data Recovery
CEx	Coexistence Element
CIN	Converged Interconnect Network
CMD	Chromatic Mode Dispersion
CPON	100 Gbps Coherent Passive Optical Network
CRC	Cyclic Redundancy Check
DAA	Distributed Access Architecture
dB	Decibel
DBA	Dynamic Bandwidth Allocation
DDOS	Distributed Denial of Service
DFB	Distributed Feedback Laser
DSP	Digital Signal Processing
EPON	IEEE 802.3 Ethernet Passive Optical Network
EVPN	Ethernet Virtual Private Network
FEC	Forward error correction
FM	Frequency Modulation
FSAN	Full Service Access Network
Gbps	Gigabits per second
GPON	ITU-T Gigabit Passive Optical Network
GSM	Global System for Mobile Communication
IEEE	Institute of Electrical and Electronic Engineers
IM-DD	Intensity Modulation-Direct Detection
IPG	Interpacket Gap
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
km	Kilometers
LDPC	Low Density Parity Check
MAC	Media Access Control Layer
MDU	Multiple Dwelling Unit
MTU	Multiple Tennant Unit
MULPI	MAC and Upper Layer Protocols Interface Specification
OAM	Operations, Administration, Maintenance
ODN	Optical Distribution Network
OLT	Optical Line Terminal
ONU	Optical Network Unit
P2MP	Point to Multipoint
PDU	Protocol Data Unit
PHY	Physical Layer
PLOAM	Physical Layer OAM
PMD	Polarization Mode Dispersion
PON	Passive Optical Network

PSBd	Physical Layer Synchronization Block Downstream
PTP	Point to Point
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
QPMZ	Quadrature Parallel Mach-Zehnder
QPSK	Quadrature Phase Shift Keying
SCTE	Society of Cable Telecommunications Engineers
SD-WAN	Software Defined Wide Area Network
SMB	Small Medium Business
SOA	Silicon Optical Amplifier
SONET	Synchronous Optical Networking
TC	Transmission convergence
TDM	Time Division Multiplexing
VPN	Virtual Private Network
WDM	Wavelength Division Multiplexing
XGEM	XGS Generic Encapsulation Method
XGS	ITU-T Ten Gigabit Symmetric Passive Optical Network

Bibliography & References

25GS-PON Specification:

- 25 Gigabit Symmetric Passive Optical Network, 2 November 2023, Version 3.0

CableLabs Specification:

- CPON Architecture Specification Version I01

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.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.9804.3: 50-Gigabit-capable passive optical networks (50G-PON): Physical media dependent (PMD) layer specification
- G.988: ONU Management and Control Interface (OMCI)
- G.9807.1: 10-Gigabit-capable symmetric passive optical network (XGS-PON)
- Y.1731: OAM functions and mechanisms for Ethernet based networks