If You Love Coherent, Set it Free:

Extending Coherent Optics to the Outside Plant

A technical paper prepared for presentation at SCTE TechExpo24

Venk Mutalik
Comcast Fellow
venk_mutalik@comcast.com

Amarildo Vieira, Principal Optical Engineer, Comcast

Scott Wood, Senior Access Engineer, Comcast

Dan Rice, VP Network Architecture, Comcast

Bob Gaydos, Comcast Fellow

Table of Contents

List of Figures

1. Introduction

"If You Love Somebody, Set Them Free" sang Sting in 1985, and as the popular saying goes on, "if they come back, they're yours." A technology personified by coherent optics is on the verge of being set free at Comcast. In this context, a technology as sophisticated as coherent optics — once used primarily in Core and Metro Networks — is now being used in Comcast to span headends and hubs in access networks and enable Distributed Access Architecture (DAA) for DOCSIS® networks and fiber to the home (FTTH). In previous papers we have discussed deploying single-fiber Coherent Optical solutions to maximize fiber utilization across the national footprint. Today initiatives such as Broadband Equity Access and Deployment (BEAD) and Rural Digital Opportunity Fund (RDOF) quicken the pace of deployment and call for extending the optical reach closer to the home.

Use of coherent optics in the outside plant (OSP) brings tremendous advantages to a network operator by extending reach to 100km or more and by providing 100Gbps or higher capacity across the entire network. However, life of equipment on the strand or in a pedestal in the outside plant is harsh - surviving and functioning in temperatures between -40C and +85C, while also extending reach, enhancing capacity, and maximizing fiber utilization is a complicated endeavor.

In this paper, we describe for the first time unleashing the power of coherent optics in the OSP environment. When this capacity is combined with approaches like coherent muxponder (CMP) or Remote Switch on a Pole, the entire network can converge and carry 100Gbps to 10Gbps optical traffic down to individual nodes, businesses, and homes. We describe key technology innovations in bi-directional (BiDi) coherent optics, contrast dual-laser and sub-carrier approaches, point out operational challenges, and present architectures that help converge optical links from the core to the home.

2. Why Coherent?

A quick recap of coherent optics and their advantages can illuminate their need for continued and proliferated use in all optical networks, whether in long-haul links, metro links, access links, point-to-point data center interconnect links, or even intra data center links.

Comcast's optical network illustrated below incorporates long-haul, metro, access, and data center links, interconnecting systems that carry residential and commercial services across the country. Conventional intensity modulated transmission—which relies on on-off keying—typically runs out of steam with transmission rates exceeding 10Gbps. The chromatic dispersion normally present in single mode fiber (SMF) at around 16ps/nm/km typically starts spreading and distorting these pulses quadratically and limits reach. Thus, extending data rates to 25Gbps limits transmission to around 13km and around 800m for 100Gbps. Getting out of this natural limitation requires electronic compensation techniques that can provide limited relief or optical compensation techniques that inherently limit performance due their own loss and complexity.

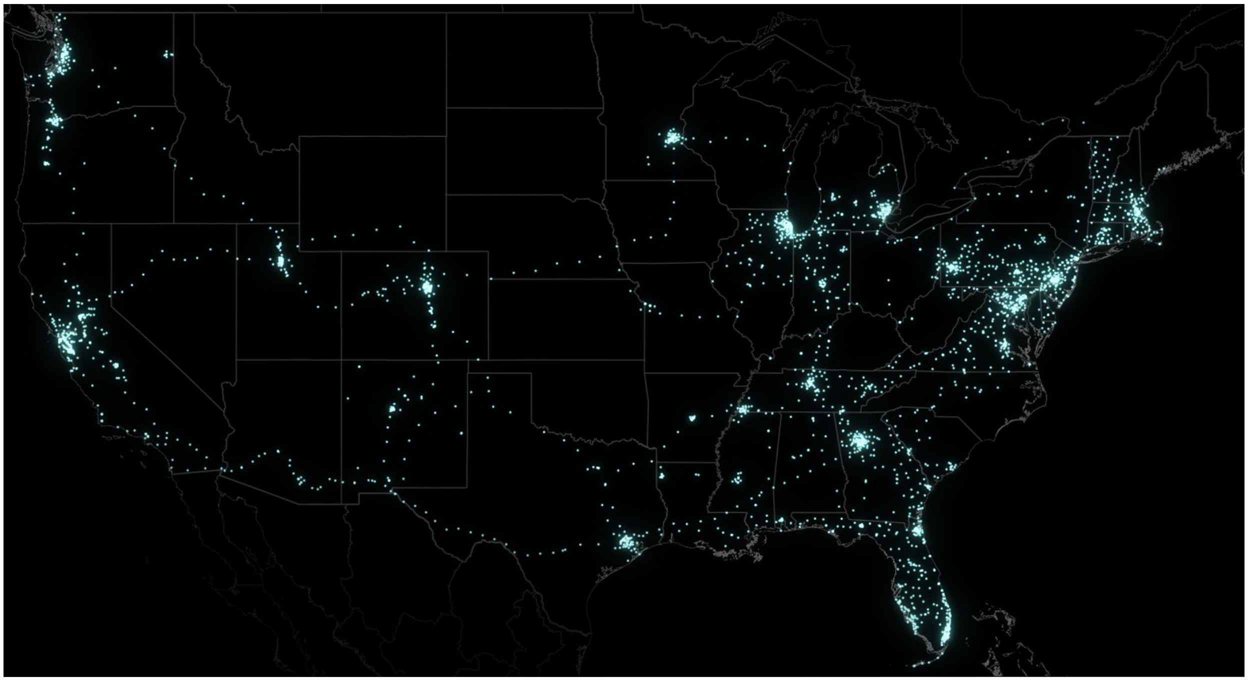

Figure 1 – Illustrating the Coast-to-Coast Comcast Network

As traffic requirements grow, the need for higher capacity and longer reach becomes critical. Even more important is the need to provide greater than 10Gbps capacity deeper and deeper into the network.

Fortunately, coherent optics provides a way out. Unlike intensity modulation, coherent optics look at intensity and phase of light utilizing a local oscillator in the receivers. While this complicates receivers, the tradeoff is well worth it in terms of capacity and reach [1]. One hundred Gbps transmission can exceed tens of thousands of km, while 400Gbps transmission has a reach greater than 500km with commercially available pluggable optical modules. The main reason for this ability is because coherent optics use mature digital signal processing (DSP) techniques in the optical modules to offset chromatic dispersion and further use optical spectrum effectively by increasing the complexity of modulation to take better advantage of optical signal-to-noise ratio (OSNR) available in optical links. In this context, optical transmission starts to resemble RF transmission, which is quite familiar to cable engineers well-versed in these tradeoffs. It goes without saying that the more things change, the more they remain the same.

2.1. Optical Spectrum, Capacity and Reach

In a previous paper [2] we have described the similarities in coherent optics and RF transmission as two sides of the same coin. While distortions due to chromatic dispersion, signal overload, and micro reflections are cumbersome to deal with, various ways of mitigating them may be thought of. However, random thermal noise is an intrinsic characteristic of the network and fundamentally limits capacity. Therefore, sophisticated techniques are used to take advantage of available signal-to-noise ratio and various mitigation techniques to approach the maximum amount of capacity that can be reached for given physical limits.

One such technique used in RF transmission involves the use of multiple levels of transmission rather than just on-off keying as illustrated below. Commonly known as quadrature amplitude modulation (QAM), a single signal can code multiple bits of information, thereby drastically reducing spectrum requirement, conserving signal integrity, and maximizing capacity for a given signal-to-noise ratio.

Coherent optics today can scale up to QAM 16 or even higher constellations delivering 400Gbps within a spectrum of 60GHz.

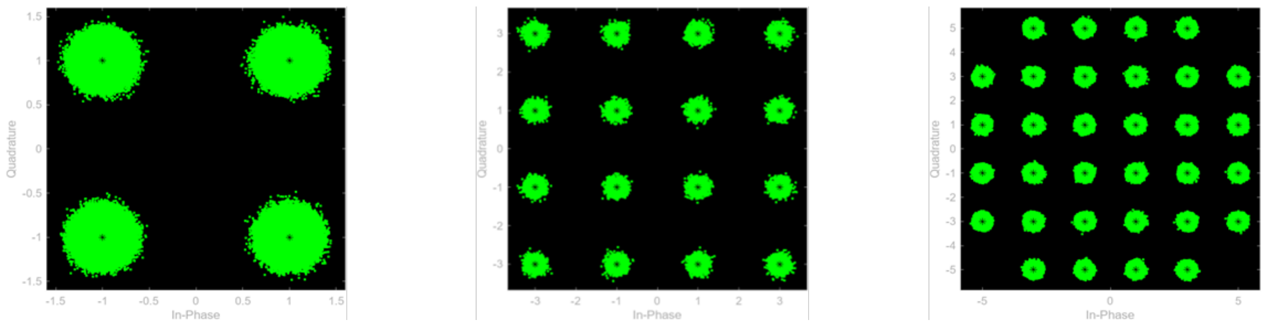

Figure 2 – Illustrating QPSK, QAM 16, and QAM 32 (L - R) Optical Constellations

The figure below illustrates that 100/200/300/400Gbps can be so spectrally efficient that they fit within +/-30GHz spectrum that is common for traditional 100GHz spaced optical filters.

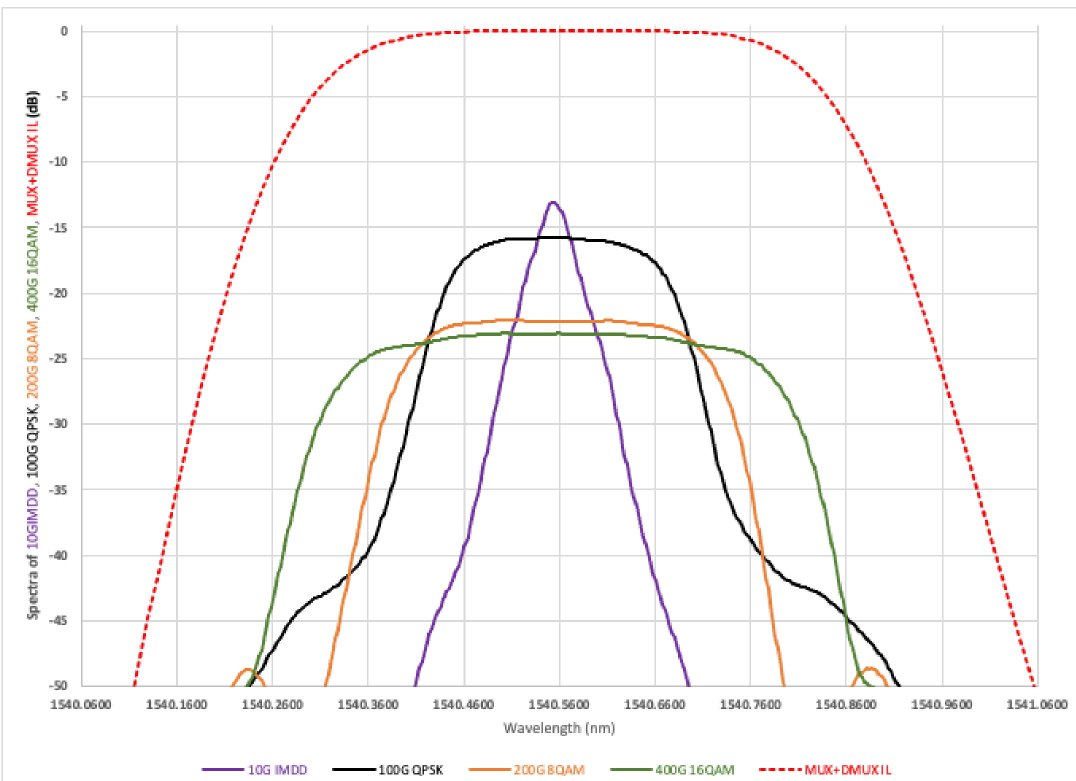

Figure 3 – Illustrating Optical Spectra of 10Gbps to 400Gbps Optical Transmsion

Since spectrum width of transmission signals critically determines reach, the ability to compress large quantities of data within limited spectrum and advanced DSP techniques has the added benefit of long reach. Thus, while 10Gbps on-off keying is limited to 80km without resorting to any DSP, 100Gbps with quadrature phase shift keying (QPSK) transmission and 30GHz spectra can reach over 10,000km, while 400Gbps with QAM16 transmission can reach greater than 500km.

But even the best of signal processing and optical spectral limiting techniques cannot compete with the ease and affordability of widely available 10Gbps pluggable optics commonly called small form-factor pluggables (SFPs). Fortunately, in the last few years, design and manufacturing technologies have size-reduced coherent optics into small pluggable devices and are discussed next.

2.2. Pluggable Optical Modules

Even a few years back, coherent optical modules were relatively big transponder modules that required elaborate chassis and called for sophisticated techniques to install, provision, and manage. Thus, they were primarily used only in long haul networks with their pervasive use being quite limited. On the other hand, 10Gbps intensity modulated modules were not much bigger than a 1Gbps SFP, which consumed less than 2.5W and could tune across the entire C-Band (conventional band) from International Telecommunication Union (ITU) 14 to ITU 61, with a reach of 80km. Therefore, much of 10Gbps had moved from traditional transponders to plugs a couple of decades back.

With the advent of miniaturization and sophisticated technology — such as silicon photonics and high quality and higher efficiency lasers as well as lower power DSP chips — coherent optics have also now been miniaturized into extremely compact, optical pluggables. All networks are networks in transition and the ability to do this highly prized, therefore the pitch of optical connectors in the front face of optical pluggables has remained the same across the years to simplify their coexistence in the plant.

The first coherent optics devices were quite big and sometimes needed mediation to connect to routers. In these optics, the DSP was spread out over a board while the optical circuits were hardwired in to the board together forming a line card that plugged in optical chassis. Many legacy systems and current transoceanic links still use this approach. Over time however, with the C form-factor pluggable type 2 (CFP2) form factor shown below, there was widespread availability of routers and switches that could handle coherent plugs directly, especially suitable for the vast majority of Metro networks. This attained greater levels of acceptance with the advent of 400Gbps transmission in a quad small form factor pluggable transceiver (QSFP) shown below. In general, introducing pluggables into routers also has the added feature of being able to build meshy networks improving reliability, cutting down on latency and creating a more elastic network overall.

Figure 4 – Illustrating Popular Coherent Optics Pluggable Form Factors

In addition of the CFP2 and QSFP form factors show above, there are many more form factors that are widely used and are summarized below.

SPEED	PLUG	Height	Width	Typ Depth	Front Area	Typ Volume	Front Panel
		H mm	W mm	D* mm	HW mm^2	HWD* mm^3	Devices/RU*
1G	SFP	8.5	13.4	56.5	114	6435	48
10G	SFP+	8.5	13.4	57.0	114	6492	48
25G	SFP28	8.5	13.4	58.0	114	6606	48
10G	XFP	8.5	18.4	78.0	156	12166	32
40G	QSFP	8.5	18.4	72.4	156	11293	32
100G	QSFP28	8.5	18.4	72.4	156	11293	32
400G	QSFPDD	8.5	18.4	89.4	156	13944	32
Other Options	OSFP	13.0	22.6	100.4	294	29471	16
	CFP2	12.4	41.5	107.5	515	55320	8

Figure 5 – Illustrating Critical Dimensions of Popular Optics Form Factors

The venerable SFP is the smaller plug of the bunch with a front face area of 8.5mmx13.4mm. The dissipation possible in this form factor is typically around 2.5W. The SFP28, which has the same form factor as the small form factor pluggable + (SFP+) allows for 25Gbps transmission. The deeper module enables the use of a small DSP that can perhaps push 25Gbps transmission to around 40km from the present day 13km limit. A form factor of this type can also enable fiber to the home (FTTH) type applications for 25G PON (passive optical network). There are many routers that can accommodate 48 SFPs in a single RU (rack unit).

A move away from the SFP gets us in the Quad SFP (QSFP) type form factor with its 8.5mmx18.4mm front area, allowing for much more complicated optical plugs. Most notably, the 100Gbps and 400Gbps coherent transmissions today use this form factor. With the ability to dissipate around 22W of power, this device can accommodate higher laser output power, an erbium doped fiber amplifier (EDFA) if needed, silicon photonics optical modulator (SiPho), and a fair-sized DSP that enables greater than 500km of reach: enough to power most metro networks through several re-configurable add drop multiplexers (ROADMs). These QSFP devices can have individual 100Gbps lane or quad 25Gbps lanes that ultimately enable designers to split off a 400Gbps capacity into more bite-sized 100Gbps or even 25Gbps lanes that can help multi-generational links in a network. There are many routers in the market that can accommodate 32 or more QSFP in a single RU.

When systems go beyond 400Gbps, there is a choice of continuing the QSFP form factor or a move perhaps to the Octo SFP (OSFP) form factor. This form factor, with its 13mmx22.6mm front area, is quite big and perhaps can accommodate more heat dissipation. While some would be tempted by this feature, others might count on Moore’s, Dennard’s, and Coomey’s laws to enhance computation, reduce power, and reduce space respectively and allow QSFP form factor to continue thriving.

Finally, the CFP2 is a much bigger form factor with its 12.4mmx41.5mm that can dissipate larger heat. But more importantly the larger space allows newer types of coherent optics to be demonstrated in this form factor. And indeed, as we will discuss later in the paper, use cases that required single fiber BiDi (bi-directional) applications were first demonstrated in this form factor and are now moving towards the QSFP form factor.

2.3. Optical Disaggregation

Optical transmission systems typically comprise of optical transponders, optical line systems (OLS), routers that connect to the transponders, and the software that administers and manages the entire system. In traditional systems these were tightly coupled and were available as a combined package for ease of use. However, the pace of innovation in each of the elements happens at differing rates and the entire system was bogged down by the lowest common denominator. Thus, innovations of size reduction and complexity reduction in optics could not be taken advantage of or advances in streaming telemetry would be incorporated, as well as space and power reductions possible with a move to more efficient routers and line systems unable to be taken advantage of.

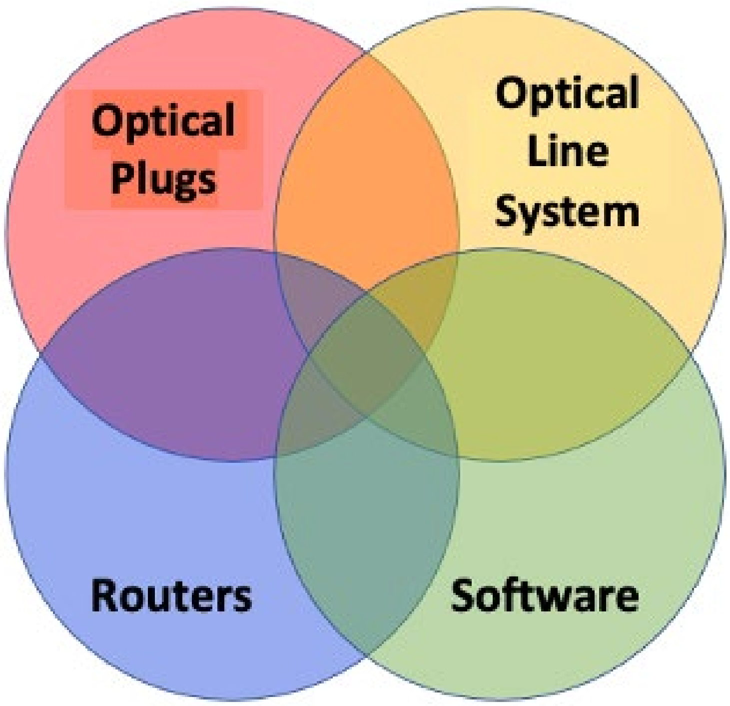

Figure 6 – Essential Elements of Optical Disaggregation

Attempts at innovation in the coherent optics space have taken off with the concept of optical disaggregation. Here, each of the individual components of an optical system are systematically delineated and looked at independently with a view towards performance, affordability, and innovations. As we have said earlier, optics currently span across long-haul, metro, data centers, and access. There are profound impacts to optical systems with innovations in each of these areas. While routers are becoming more efficient and faster for data center applications a push towards point-to-point links there has had a significant impact on optical technology moving away from traditional transponders to optical plugs, thereby reducing footprint and decreasing complexity. Streaming telemetry common in the access domain for the last several years with developments in distributed access architecture (DAA), which have now become commonplace with standards such as the gRPC network management interface (gNMI) that enable an array of folks to use artificial intelligence (AI) and machine learning (ML) to keep close real-time visibility on the vast optical network. Finally, the aforementioned move towards point-to-point networks and relative affordability of routers enables a more elastic network with lower dependence on ROADMs, which previously were a necessity.

2.4. A Quick Aside on Single Fiber Optical Links

In this context, all of the OLS or ROADMs in existence today require two separate fibers for their functioning. This is because all current ROADMs have single direction wavelength selective elements and optical amplification. This requirement of dual fibers for each optical link, while well-meaning, has a significant effect on lighting up new links. As we have explained in an earlier paper, [3] optical fibers are very expensive and are constrained assets. Constructing new fibers is expensive, time consuming, and prone to additional delays due to permitting processes. In many cases, obtaining additional fibers is simply not possible. If it is possible with leases, then they are likely not routed through favorable locations and could add to latency. Thus, the ability to light links over single fibers is crucial to the ongoing activity of network operators. Besides, the entire access network used by network operators is already only single fiber based, where the same fiber is used to convey upstream and downstream signals. Therefore, the ability to migrate to single fiber links for coherent optics where needed is a crucial tool in the arsenal of network operators and one which we will talk about in detail in later sections.

3. Why Remote?

In previous papers, we have described the use of coherent optics in access networks. [4,5,9] The ability to use coherent optics in access networks significantly enhances capacity and reach. For all links that require an end connection that exceeds 10Gbps or for any link that exceeds 80km in fiber distance, the use of coherent optics is often the only practical way to close the links.

3.1. The Sites that Serve

Major network operators like Comcast have many sites across the country that support long-haul, metro, and access systems. Large sites called primary sites typically support in excess of 100K household passings (HHP). Although a term such as household passing is used, it is generally taken to mean homes or businesses or any entity eligible to receive active connections. These primary sites can hold servers that virtualize cable modem termination systems (CMTS), major routers that connect to the internet, and metro connections that are connected to ROADMs and make their way to other primary sites as well as to other secondary sites that can form a ring before they fold back into the primary site.

Each primary is generally connected to several secondary sites typically serving up to 75K HHP. These secondaries also hold many routers that connect up to the virtualized CMTSs from primary sites. In addition, several business customers can also be supported by secondaries thru aggregation routers as well. Generally, primary to secondary connections are connected via redundant routes to ensure high availability for the larger numbers of HHP. Oftentimes, these secondaries are supported by battery banks and with generators and propane fuel, along with multiple fiber routes in and out of the facilities. Scoping, building and maintaining these facilities is a science and art all of its own.

As network operators expand into newer areas, secondaries are not readily available. Or even if they are, they can simply be too crowded to handle additional new equipment. To maintain a green footprint in far flung access points, Comcast has started deploying coherent optical links that can span up to 120km without any intermediate optical amplification. When bi-directional coherent muxponders (CMP) are deployed in these extension sites, up to 2400Gbps of capacity becomes available at each extension site over a single bi-directional fiber.

In practical terms, extension sites serve around 20K HHP, but the ability to have such a vast amount of capacity per extension site is a significant achievement that serves Comcast well and enables fulfillment of network traffic demand for residential and business users for years to come.

3.2. Introducing the Remote Enclosure

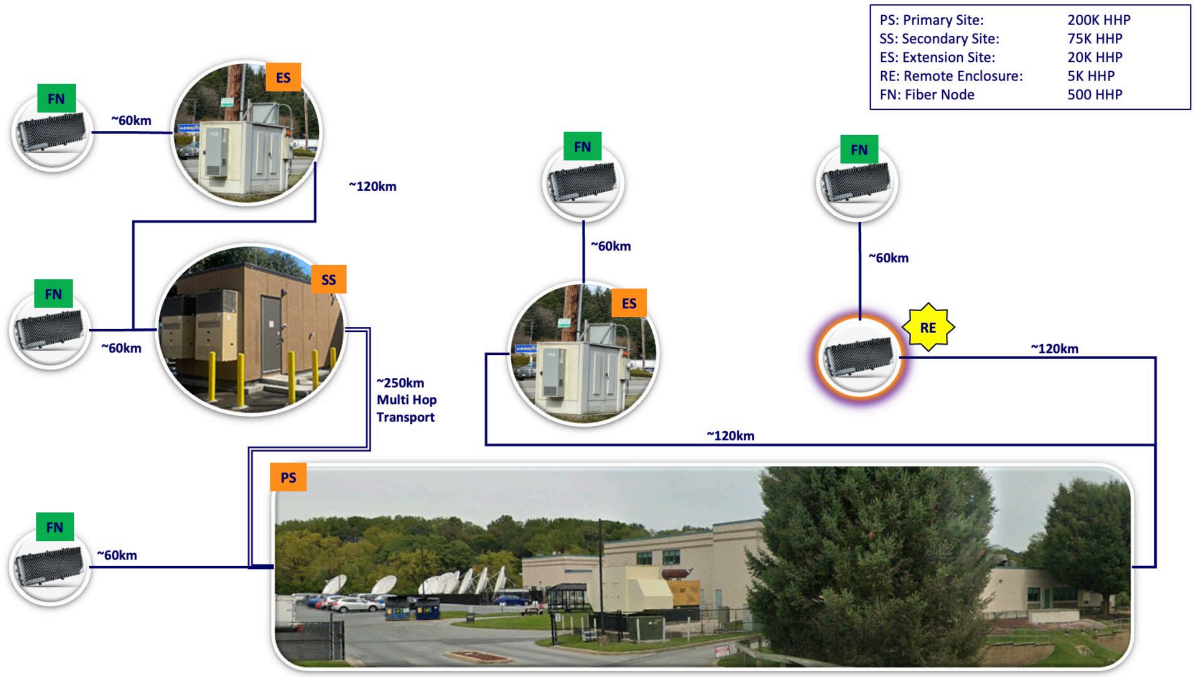

Figure 7 – A View of the Access Landscape

The picture above introduces a brand-new network element that could be a game changer for network operators. Powered by technological innovations and conceived to support far-flung remote operations with great capacity, this is a reimagination of a miniature extension site within the confines of a fiber node.

As we have seen before, setting up a new secondary is rather expensive and time consuming. In addition to securing real estate, there is the issue of securing permits for setting up a hub, enabling fiber pathways, and shoring up power supplies by battery banks, generators, and fuel. Setting up extension sites is a smaller version of setting up secondaries. In addition to securing real estate and ensuring power supplies, the construction of an extension site requires permitting processes that could substantially impact construction activity. Even in extension sites that are already available, space is rather limited and critical infrastructure rather sparse.

Many of the newer builds that require an extension of footprint are in rural and lower density locations. Furthermore, there is a need for many locations geographically distributed, but able to serve pockets of homes and businesses. Previously such construction was not easy, since coaxial reach was rather limited. Most optical nodes can serve areas of around a radius of a mile or two at most even with RF amplifier cascades. But optical fibers have much longer reach and thus enable a much more distributed build. In addition, new builds are skewed towards fiber construction and are generally FTTH and business connections. Presented below is an illustration on how remote coherent muxponder (rCMP) could be leveraged.

Figure 8 – Illustrating ways to use Coherent Optics to Extend Reach to address Urban Capacity and Rural Spread

Use of coherent optics in the Outside Plant (OSP) brings tremendous advantages to a network operator by extending reach to 100km or more and by providing 100Gbps or higher capacity across the entire network. It is a complicated endeavor to survive and function in temperatures between -40C and +85C, while also extending reach, enhancing capacity, and maximizing fiber utilization.

3.3. Challenges of the Outside Plant Environment

As indicated, we have previously talked about using coherent optics in environmentally stable locations. In those cases, we also have used the CFP2 based solutions. However, when we migrate towards the remote outside plant environments, a completely new set of technical, design, and operations issues spring up.

For starters, the outside plant (OSP) locations dictate temperature performance from -40 to +85C operation. In addition, the distances to be covered require exceptional monitoring of the end points, or else network availability could be seriously degraded. Operationally, provisioning, troubleshooting, and monitoring traffic across distances that can exceed 180km is challenging. Finally, there should be adequate provisions to track fiber cuts across the entire system. In this paper, we will go thru each of these individually and explain how we solved all these issues.

3.4. Technological Enhancements

We began this program envisioning a single fiber bi-directional coherent muxponder (BiDi CMP) in a CFP2 form factor for the controlled environment. After a short while it became clear that there was a great need for this system, especially when we looked at how effective this platform has been in resurrecting hurricane devastated Sanibel Island a few years back [5]. However, recreating this entire system for the harsh outside environment required many technical innovations.

3.4.1. Avoiding Spectral Overlap

First, a move to OSP required a move also to the QSFP form factor. This move was done with a view to embrace miniaturization, while keeping in view the long-term trend towards QSFP in the industry. However, the smaller QSFP DD form factor can limit heat dissipation and has limited space.

The main requirement of single fiber bi-directional transmission is to avoid common spectra for upstream and downstream transmission. While this can be done with two lasers, it can also be done with a single laser with subcarrier technology as well. These approaches are described in more detail in section 4. As described earlier, Comcast has already deployed optical filters in access plant that are 100GHz spaced. Since a full 400Gbps spectrum fits within the filters, both approaches easily fit within the filter passband. An advantage of this approach allowed us to double the capacity of the optical fiber from when we would have used a dual laser approach occupying separate 100GHz ITU channels.

3.4.2. Thermal Management

Industrial temperature (i-Temp) performance requires unique heat sink designs. For this reason, we selected an already qualified optical node that is used in current DAA architecture and supports RPDs (remote PHY device) and high output RF amplifier ports. This node can be deployed on a strand or into a pedestal that has sufficient inside space and casting volume to dissipate heat when stripped off of the RF components. The entire existing CMP was then size converted to accept one 100Gbps QSFP-DD and muxpond it out into 10 individual SFP28s that could then independently traverse an additional 80km to reach their node destinations.

word of caution might be in order here, QSFP-DD transceivers are currently used in data center type applications, so they are optimized for air cooling. However, inside a fiber node, there is no air circulation at all. Therefore, there is a need to heat sink these well to the chassis itself. However, special care must be taken to modify the heat sink capability of the QSFP-DD and match it to the internal node mass for enabling the QSFP lifetime performance. As part of this program, innovative heat sinking techniques that allow air cooled QSFPs to handle heat sink designs have been implemented.

3.4.1. Latency & Jitter

In these designs, we have still operated with mux-ponding and without any over-subscribing. This is important as there are now no jitter constraints on the network and even the most fragile of transmissions pass thru uneventfully across the temperature range. On the question of latency, the maximum amount of latency is, as it should be, based on the fiber transit time (with the muxponder and the QSFP adding negligible latency). Furthermore, the ability to virtualize the core means that the latency clock begins with the RPD location and not from the virtual cable modem termination system (vCMTS) location in the Figure 9.

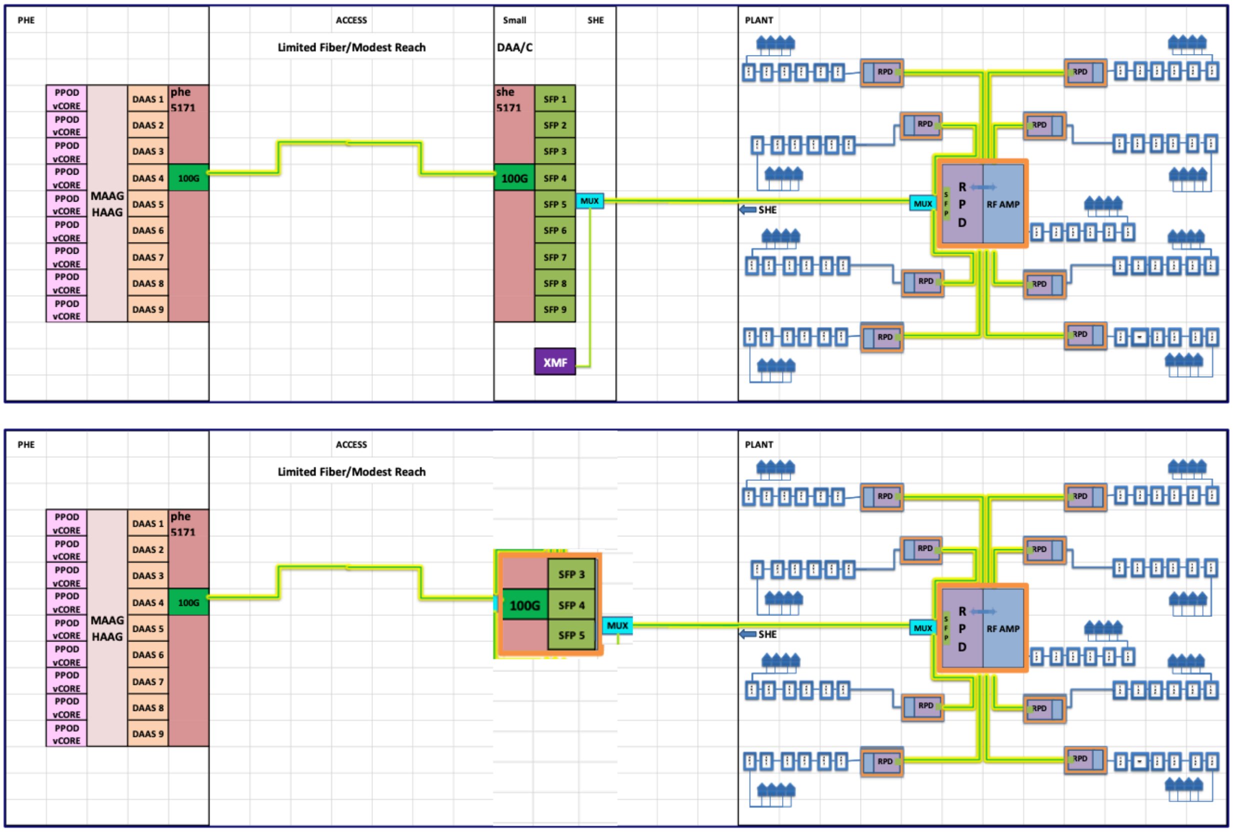

Figure 9 – Illustrating CMP and rCMP solutions for DAA

3.4.2. ROADMS and BROADMs

Even with the best of networks, there is often a need for fiber redundancy. For the bigger CMP that handles 40 RPDs and exceeds 20,000 HHP, a design that incorporated optical protection switching and optical amplification at the extremity suffices to provide always on redundancy. This required us to design Bi-Di optical line system, lovingly called BROADMs (bi-directional reconfigurable optical add drop multiplexer), although the very first implementation of this uses fixed optical multiplexers in view of the outside plant temperature ranges.

4. Single Fiber BiDi

Bi-directional (BiDi) fiber is a full-duplex communication system using a single (simplex) fiber. There are many BiDi solutions already available in the market, but they have been targeting very specific applications. That includes approaches using different downstream and upstream WDM channels or using different wavelengths such as 1550/1310nm pluggables or preventing frequency (or wavelength) overlap in a WDM channel. The goal of these approaches is to minimize the impact of back-reflections and fiber backscattering on the link performance. That is a reason to enforce the best practice of using angled physical contact connector (APC) connectors across the access optical plant from day one, as we move to coherent optics using QAM with higher order of modulation to keep up with capacity needs.

Traditional coherent optics devices are built with a single laser, acting as a transmitter and a local oscillator for the receiver, consequently, they always require two separate fibers to complete bidirectional transmission. However, when a single fiber bi-directional transmission is needed, this approach does not work due to the aforementioned fiber reflections and fiber backscattering. To get around this, the second approach below, where two separate lasers, one each of the transmitter and receiver are used to make the system largely insensitive to fiber reflections and to fiber back scatter. A third approach is now being introduced in Comcast, which has the potential of using a single laser, but with spectrum split into two separate parts one each for downstream and upstream. This novel approach uses a single laser, but also is largely immune to optical reflections and back-scatter and also fits within the optical passband of the Comcast deployed optical filters.

Figure 10 – Illustrating Dual Fiber and Single Fiber Coherent Systems

It is also important to notice that passive optical network, PON, is also a single fiber BiDi solution that uses different wavelengths for separating the Downstream and Upstream signals and also relies on time division multiplex (TDM) in the upstream path to allow customer premise equipment (CPE) devices to talk back.

4.1. Subcarrier Approach

Operation using a single carrier, with frequency overlap, is also possible, where a single WDM channel is using Tx and Rx frequencies (wavelengths) that overlap. In this case, power penalties due to optical return loss (ORL) of optical components need to be carefully accounted for, which is translated to reduced reach.

The approach taken to avoid penalties and maximize link performance was to use non-overlapping subcarriers. In this approach one 100 GHz dense wavelength division multiplexing (DWDM) channel is sub divided into subcarriers that can be allocated for Tx and Rx purposes. The number of subcarriers, the subcarrier width and the modulation order is selected to deliver a certain capacity and, depending on the Tx and Rx performance, define the supported link budget.

Figure 12 illustrates how the subcarriers parameters can be adjusted to deliver different data rates in BiDi mode. It shows the spectrum for 100, 200, and 400 Gbps data rates, using 4, 8, and 16 x 16-QAM subcarriers, respectively. Figure 13 show the example of a 100Gbps Single Fiber BiDi optical spectrum using 4 x 25 Gbps for downstream and 4 x 25 Gbps for upstream, employing 16-QAM modulation orders.

There are many advantages of using the subcarrier approach, but the most important is conserving DWDM channels being used in the system since the same DWDM channels is used for the Tx and Rx links. Such a solution could potentially double the fiber capacity if the approach is used in all DWDM ports.

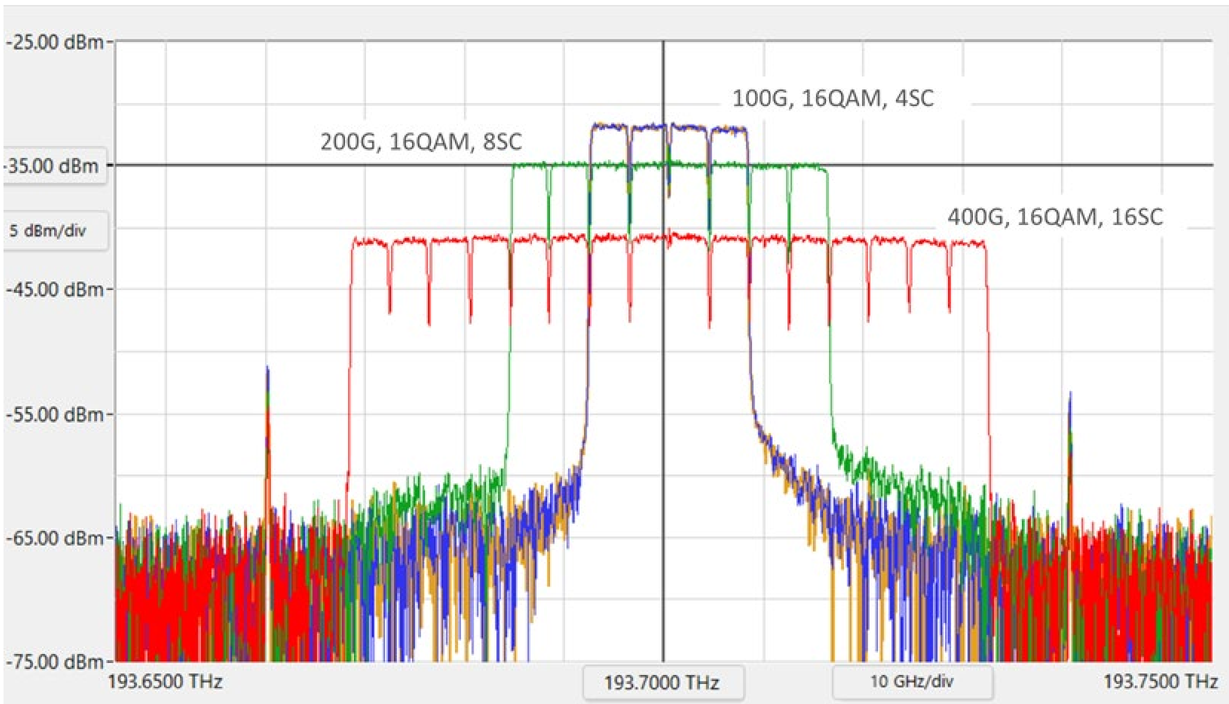

Figure 11 – Spectrum for 100, 200, and 400 Gbps data rates using multiple subcarriers

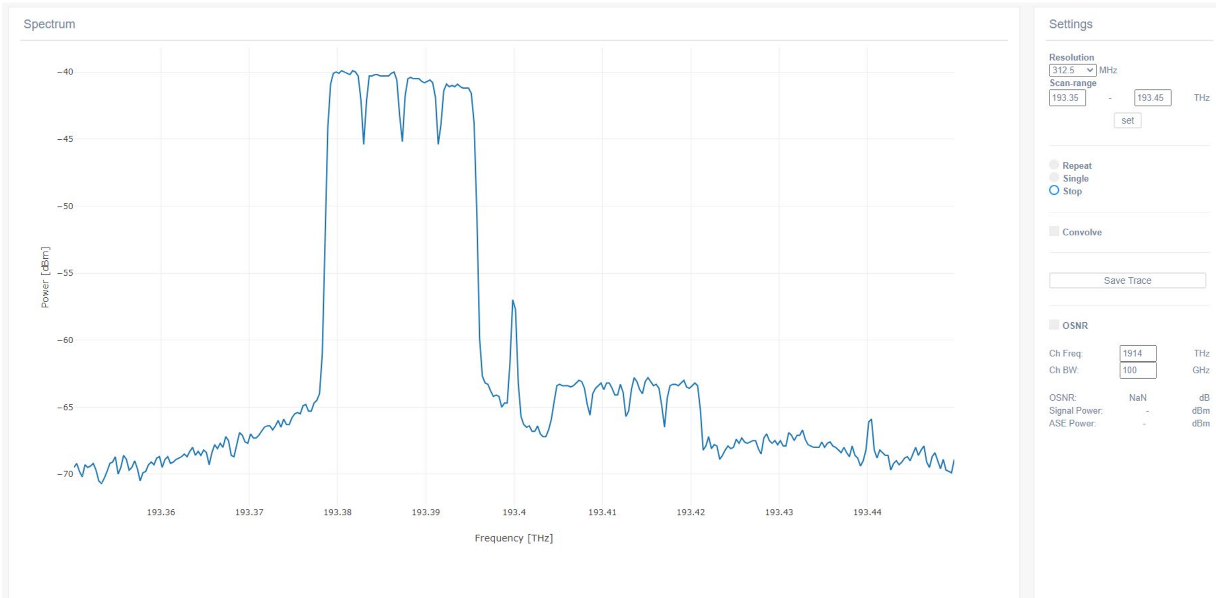

Figure 12 – Non-Overlapping BiDi Optical Spectrum: 4 x 25 Gbps Subcarriers DS & US

4.2. Reach

Another main advantage inherited from coherent optics is its long-reach capability thanks to the compensation provided by DSP. Link budgets in excess of 28 dB are easily attainable without optical amplification. Use of coherent optics in the outside plant (OSP) brings tremendous advantages to a network operator by extending reach to 100km or more and by providing 100Gbps or higher capacity across the entire network. That allows us to push the DAA system further down enabling site consolidation or even site elimination as we will discuss later in this paper.

4.3. Primary Challenges

It is a complicated endeavor to survive and function in temperatures between -40C and +85C, while also extending reach, enhancing capacity, and maximizing fiber utilization. Careful design of these considerations is a must for both pedestal and strand-based deployments.

Integration work with partners providing the main components of the CMP solution is key. That includes the switches, pluggable optics, and passive optics. A high level of collaboration is necessary to make that happen. The challenge for the node-based remote CMP solution (rCMP) is even more challenging, as we discuss in the next section.

There were some perception challenges about employing a technology used mostly in core and access networks from either a cost or power consumption point of view. However, when comparison is done in terms of bits/$ and bits/W, for the longer fiber reaches this technology enables, it is an easy win for the coherent optics solutions. The case is even stronger when you take in consideration the other benefits, such as time-to-market, facility power savings, and fiber conservation or fiber deployment deferral.

It is important to notice that CMP is one of the potential solutions in the operator's toolbox. Figure 6 shows the coexistence of different access technologies in this converged access network vision discussed here.

5. Remote Technology

The main driving force for the rCMP solution is to enable quick, low-cost deployment of different types of services while using the same infrastructure already built for the CMP solution. That is possible by mixing different generations of CMP solutions using CFP2 and QSFP-DD modules that were qualified and integrated to the solutions.

Although these modules are supposed to interoperate following standards and go through interop exercises through the industry, the truth is that there are still many integration challenges when using different vendors and different pluggable form factors in one solution. For that reason, there is a preference for bookending a link with devices from the same vendor and same form factor. In reality, there is a preference for using the same vendor in the same location or region for operational reasons to be further discussed in the remote operations, administration and maintenance (rOAM) Section 7.

Use of coherent optics in the OSP brings tremendous advantages to a network operator by extending reach to 100km or more and by providing 100Gbps or higher capacity across the entire network. It is a complicated endeavor to survive and function in temperatures between -40C and +85C, while also extending reach, enhancing capacity, and maximizing fiber utilization.

The main technology aspects related to the rCMP solution that need to be discussed here are capacity, I-Temp operation, and non-spectral overlap.

5.1. rCMP

The current CMP solution, which is deployed to extend the Distributed Access Architecture switch (DAAS footprint), is described in **Error! Reference source not found.**. The solution, which utilizes the same DAAS back-office infrastructure, employs 2 muxponders that aggregate and separate the traffic of 40 x 10Gbps SFP+ optical transceivers into 4 x 100 Gbps CFP2 or QSFP-DD used to interconnect the DAAS infrastructure system to the RPDs located in the OSP.

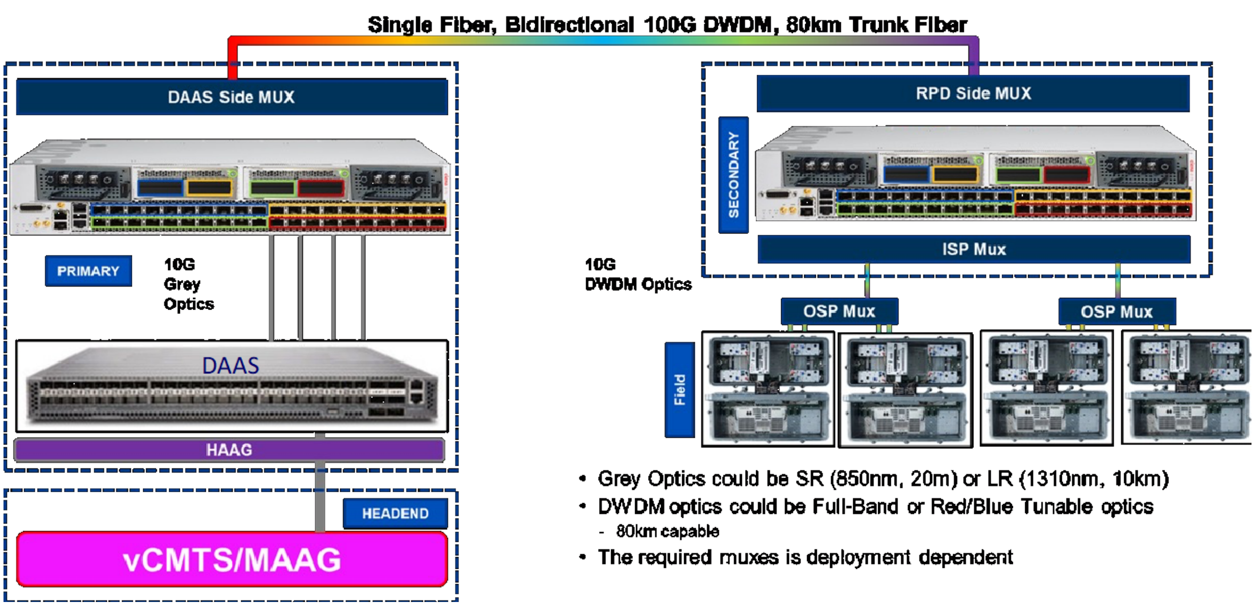

Figure 13 – Coherent Muxponder (CMP) Application Scenario for DAAS/RPD Connectivity

As we move the CMP deeper into the access network, there is a need to deliver higher capacity all the way to the node instead of the CMP. Based on our capacity needs and thermal analysis, the rCMP should be able to handle 100 Gbps to serve 10 x 10Gbps SFP+ modules to feed 10 optical nodes to deliver different types of services (DOCSIS, MetroE, PON, cellular backhaul, etc.), which covers most of the rural area deployments, the main case scenario for this solution, depicted in Figure 8. This picture shows the CMP feeding up to 4 rCMPs, and each rCMP feeding up to 10 Fiber Nodes.

The remote CMP is the remote and hardened version of the CMP to be used in the OSP. For that reason, it takes a Fiber Node Clamshell form factor shown in Figure 13.The main purpose of the rCMP solution is to drive capacity deeper in the network, enabling remote connectivity.

The main drive for the remote version of the CMP solution is to enhance the reach and enable facility reduction. The current rack mount version of the CMP already enables the DAAS and other services deeper in the network. The rCMP goes beyond that by allowing increased reach and number of hops to feed the RPDs, Metro Ethernet, and ROLTs in the field.

As already mentioned in the previous section, there are many business advantages in moving the CMP to the OSP besides extended reach that makes this solution very attractive to current and new deployments. The first is that CMP and rCMP do not need an environmentally controlled facility for their deployment, which is in line with the industry green initiatives by moving away from air-conditioned facilities. The rCMP goes further, since you don't need a cabinet as it can be installed on a pole or strand. On top of that, rights of way or real estate challenges, avoiding licensing, permits, and construction requirements, faster time to market and lower deployment costs are a welcome consequence.

As for challenges, we need to bring up the thermal management and remote device operation, administration and management, rOAM. Thermal management is nothing new for the industry with years of experience dealing with Fiber Nodes and RF amplifiers and also for 10G SFP+ optics. However, the main challenge is for 100G pluggable optics, particularly to the higher density QSFP-DD optics. Although CFP2 modules, due to their larger surface area and volume, are more suitable for OSP applications. QSFP-DD is the latest trend in the industry boosting high density, potential lower power consumption, and the content rich Common Management Interface Specification (CMIS). [6]

From a functionality point-of-view, the rCMP node is approximately one-quarter of the CMP switch. It still requires the same OMA features as the CMP, but it consists of a single 100Gbps module feeding up to 10 x 10Gbps services.

6. Convergence

One of the main driving forces behind the CMP solution is convergence on how the residential and commercial services are delivered by the operator's access network. Services include DAA, PON, Metro-E, and mobile telephony backhaul, which are currently delivered using 10 Gbps links but need to be increased to satisfy capacity needs and support new technologies. Speeds of 100 Gbps, which used to be confined to transport networks, are getting deeper and deeper in the network as described in this paper, and 400G Gbps is not far away. All that needs to be done while supporting legacy services, including analog DS and US optics. Convergence is key for operators since unifying how different types of services are delivered increases efficiency and reduces deployment time by simplifying and removing duplication.

This convergence, when combined with the Single Fiber BiDi implementation, enables us to bring together Metro-E and Access Networks onto one fiber, but requires a Converged Wavelength Plan. Figure 6 illustrates the Converged Access Network.

Such consolidation leads to great network simplification and efficiency, but consider the following:

- Careful wavelength planning needs to consider the impairment impact on other services. This is particularly important for legacy analog services which could be migrated to DAA if necessary. If that is not economically feasible, the best practice of keeping analog services separated from digital services is a legitimate recommendation.
- To get the most simplification from this consolidation process, a converged back-office infrastructure should be available as well. That may not need to happen day one, but should be a consideration
- Continuous real-time optical monitoring and fiber redundancy are a must since with many services aggregated in a single fiber. Keeping track of the fiber status and performance is key to minimize service interruption and preserve service experience.

To drive such convergence, a few enabling technologies to foster across the operation's organization are required. That includes consolidated optical passives, real-time monitoring tools and bi-directional EDFAs.

6.1. Extending Reach and Driving Technologies Deeper into the Network

Resulting from a collaborative initiative between the coherent optics and OSP hardware development teams, a solution has been developed where symmetrical 100Gb delivered on a single fiber into an OSP node housing can provide advanced services to all customers within 3000 HHP radius. The technology can be installed into standard, non-custom node housing, allowing to simplify the installation process. This evolution enables the providers to double the optical reach from their hubs by strategically installing an advanced coherent aggregation node.

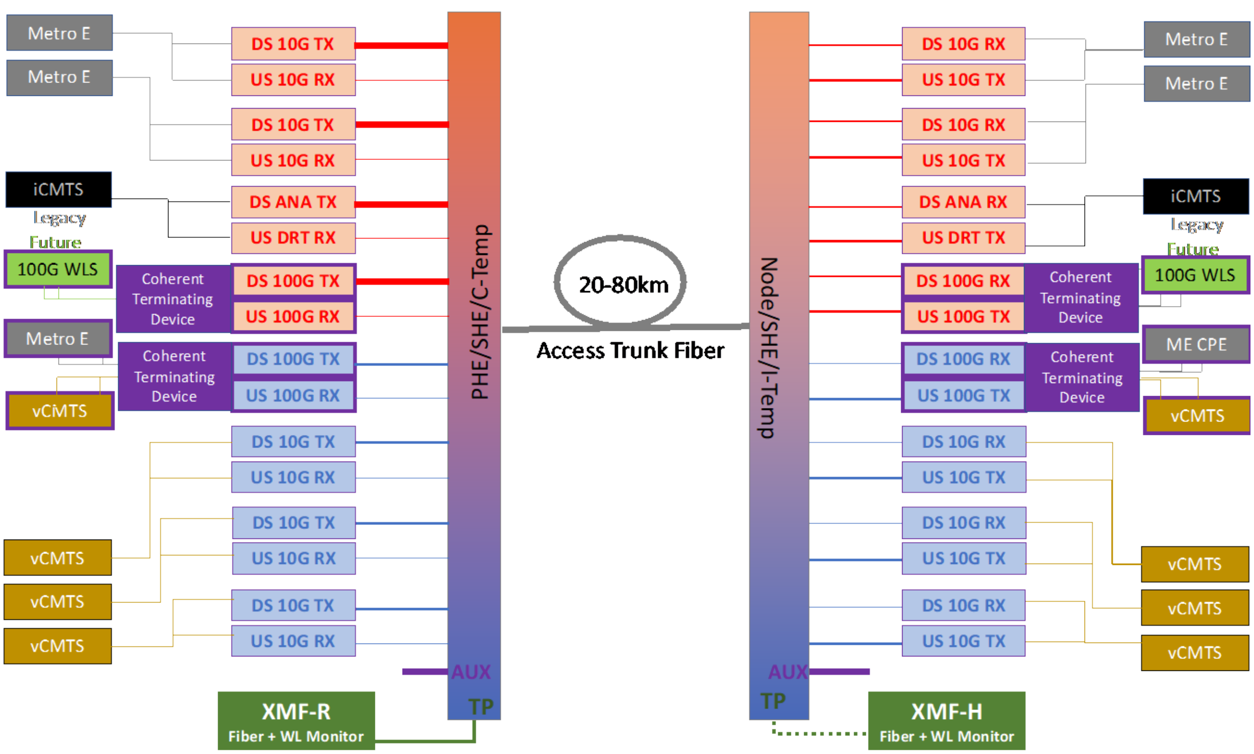

Figure 14 – Extend the Optical Reach with Coherent Node Housings and Drive 100G Deeper into the Network

In this paper, we are describing for the first time unleashing the power of coherent optics in the OSP environment. When this capacity is combined with approaches like CMP or remote switch on a pole, the entire network can converge and carry 100Gbps to 10Gbps optical traffic down to individual nodes, businesses and homes. We have described key technology innovations in BiDi coherent optics, contrast dual-laser and sub-carrier approaches, point out operational challenges, and present architectures that help converge optical links from the Core network all the way to the Home.

7. Remote Operations Administration and Maintenance

Sting also sang "Every Breath You Take," which in the current context supports operations, administration, and maintenance of remotely deployed coherent modules. The coherent platform was aggressively developed through cross-functional team input to ensure that once developed, it would allow for a repeatable and frictionless installation process and operational support. This was accomplished by focusing on interoperability with multiple open-source platforms which provide remote configuration and on-demand operational diagnostics metrics.

7.1. Fiber Optic Reliability

Using remote fiber monitoring platforms [4] will provide real-time fiber optic performance metrics, triggering notification to network support teams when a degradation occurs, or worse case, fiber cut has been identified within seconds of the detection. Deployment recommendations allow for quick and highly accurate identification of fiber segments, which have experienced degradation. As a result of segmentable isolation using the fiber monitoring platform, remediation technicians can be deployed to the area of interest in an unprecedented response time.

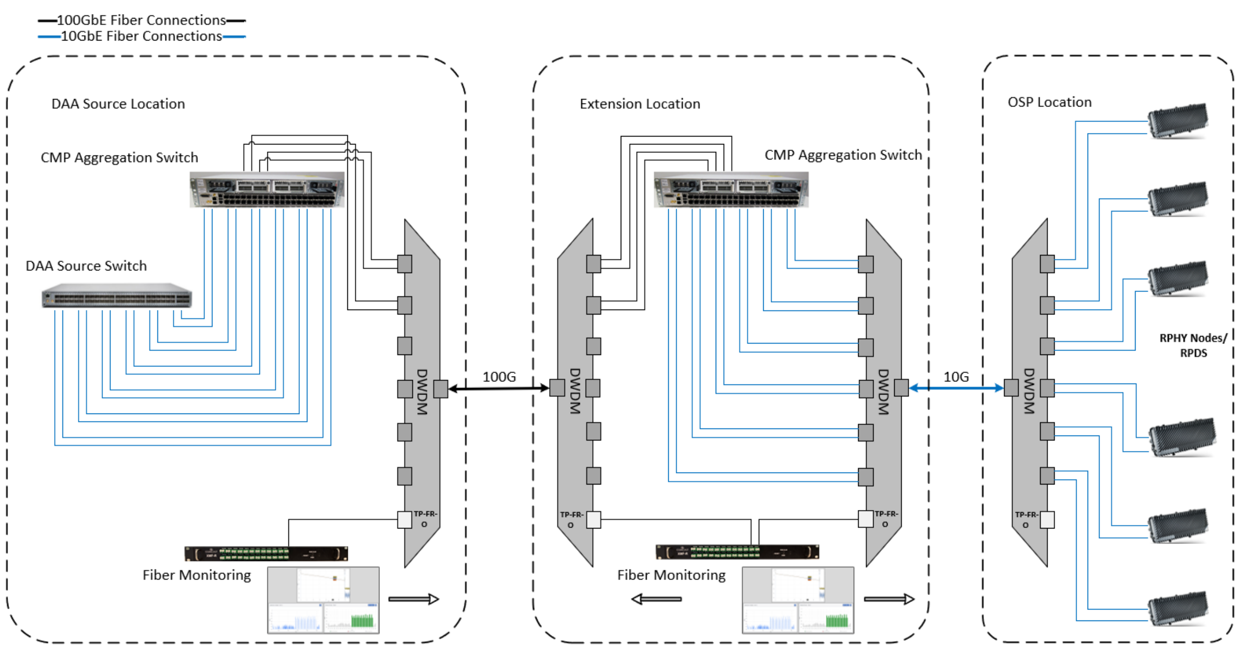

Figure 15 – In-Service Fiber Optic Health Monitoring

7.2. Continuous Optical Analysis

Part of the architecture interoperability was ensuring that the 100G BiDi was supported by our fiber monitoring platform, XMF-R. The platform is providing support teams health monitoring statistics surrounding the transmit and receive optical levels, OSA (Optical Spectrum Analyzer) to help identify OSNR (Optical Signal to Noise Ratio) impairments and OTDR (Optical Time Domain Reflectometry) analysis, showing overall fiber integrity and its effect on carrier health.

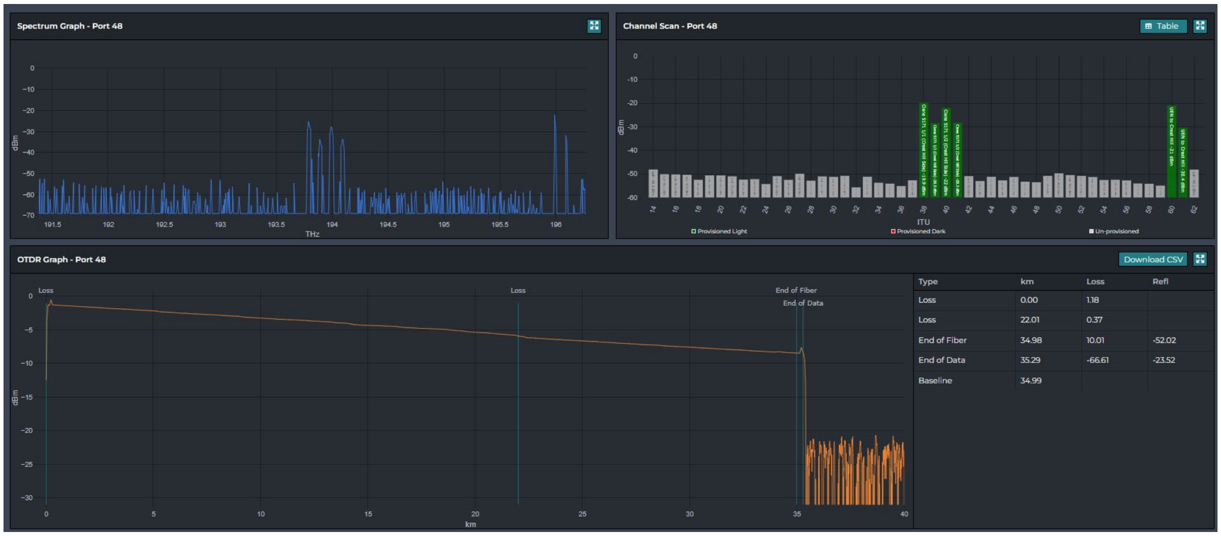

Figure 16 – Continuous Optical Analysis

This has been a tremendous catalyst to ensuring that timely degradation detection can be identified, and pinpoint troubleshooting can occur even before a remediation team is onsite. With the remote analysis and interpretation of the optical health statistics starting instantly, it provides a level of comfort as the optical reach has been extended in the Access Network to cover further distances. There was careful consideration into how to best take advantage or maximize the increased ‘windshield time’ the time the support technicians would incur by driving to and from to an incident site for repairs or even maintenance. The solution was to build in as much telemetry as possible to aid with intelligent data analysis.

7.3. Operational Health

This platform provides the ability to interact with open-source monitoring and troubleshooting platforms. One example would the Grafana visualization platform or some similar open-source agent. This allows for custom made views and analysis specific to the operators' requirements, as well as quick enhancements, evolutionary enhancements, and feature requests.

Figure 17 – Non-Specific Open-Source Grafana dashboard example

7.4. Administration

The optical platform provides advanced administration control through several different techniques, all of which are controlled, developed, and implemented by the support team. There is no need to stand up or build new infrastructure to manage the administrative tasks.

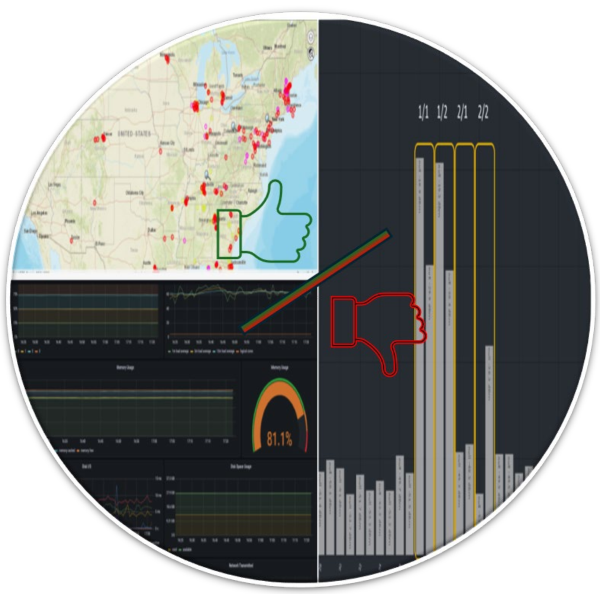

Figure 18 – Administrative User Access Control

7.5. Maintenance Standardization

Using automated maintenance techniques allows for consistent, repeatable, and identical deployment platforms. By utilizing either open-source or non-proprietary systems, this deployment model allows operators to use technology to ensure all deployed hardware is patched, updated, and running as effectively as possible.

The operator can also benefit from ensuring all new hardware deployed is done through automated processes, eliminating some artifacts from typically seen with some manual configuration methods. This is a powerful interface for the coherent optical host hardware to help deliver a level of confidence that all customer facing hardware is accounted for and managed seamlessly.

Deployment standards have been developed surrounding wavelength management which has provided the ability to quickly install configure and test up to twenty-four BiDi 100G circuits on a single fiber with no performance degradation.

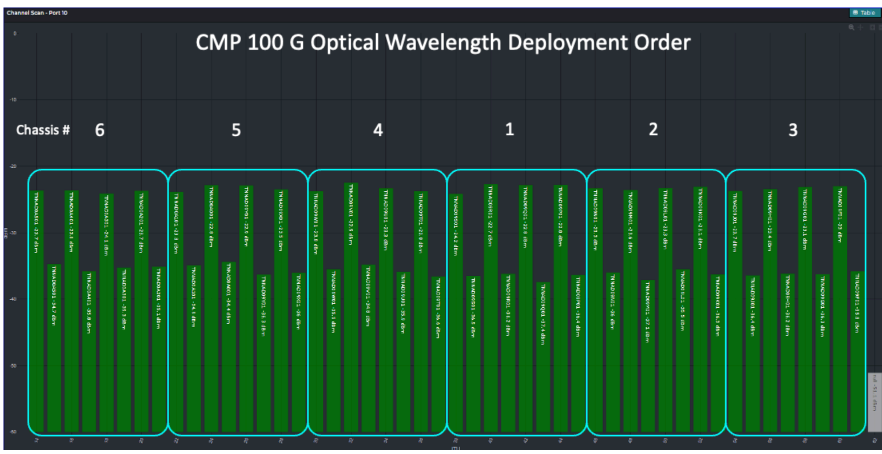

Install Order	TX Channel	Port #	TX Frequency
	CMP Chassis 1 Near Site		
1st	38	1/1	193800
2nd	40	1/2	194000
3rd	42	2/1	194200
4th	44	2/2	194400
	CMP Chassis 2 Near Site		
5th	46	1/1	194600
6th	48	1/2	194800
7th	50	2/1	195000
8th	52	2/2	195200
	CMP Chassis 3 Near Site		
9th	54	1/1	195400
10th	56	1/2	195600
11th	58	2/1	195800
12th	60	2/2	196000
	CMP Chassis 4 Near Site		
13th	36	1/1	193600
14th	34	1/2	193400
15th	32	2/1	193200
16th	30	2/2	193000
	CMP Chassis 5 Near Site		
17th	28	1/1	192800
18th	26	1/2	192600
19th	24	2/1	192400
20th	22	2/2	192200
	CMP Chassis 6 Near Site		
21st	20	1/1	192000
22nd	18	1/2	191800
23rd	16	2/1	191600
24th	14	2/2	191400
	CMP Chassis 7 Near Site		
26th	38	1/1	193800
27th	40	1/2	194000
28th	42	2/1	194200
29th	44	2/2	194400

Port #	TX Frequency	TX Channel
CMP Chassis 1 Expansion Site		
1/1	193900	39
1/2	194100	41
2/1	194300	43
2/2	194500	45
CMP Chassis 2 Expansion Site		
1/1	194700	47
1/2	194900	49
2/1	195100	51
2/2	195300	53
CMP Chassis 3 Expansion Site		
1/1	195500	55
1/2	195700	57
2/1	195900	59
2/2	196100	61
CMP Chassis 4 Expansion Site		
1/1	193700	37
1/2	193500	35
2/1	193300	33
2/2	193100	31
CMP Chassis 5 Expansion Site		
1/1	192900	29
1/2	192700	27
2/1	192500	25
2/2	192300	23
CMP Chassis 6 Expansion Site		
1/1	192100	21
1/2	191900	19
2/1	191700	17
2/2	191500	15
CMP Chassis 7 Expansion Site		
1/1	193900	39
1/2	194100	41
2/1	194300	43
2/2	194500	45

Figure 19 – Converging Direct detect and Coherent bi-Directional Wavelengths

8. Future Possibilities

Even with the many developments discussed in the paper, we have only entered the realm of coherent optics in the outside plant. As explained before, we have only used coherent muxponder options with a 100Gbps channel split into 10*10Gbps channels. However, it is possible to have 2*50Gbps or 2*25Gbps and 5*10Gbps or 4*25Gbps channels. The beauty of using such a method is that for the first time, we can provide greater than 10Gbps data rates across the vast reaches of the entire fiber footprint of Comcast, a capability not possible even a few years back.

Additionally, over time the muxponder option can be expanded to exercise modest oversubscription and provide a true switch on a pole or a pedestal (SOAP). Such a scheme is at the core of a PON system that also incidentally operates on single fiber BiDi links. This would truly provide higher capacity pervasively for when the need arises.

As a practical matter, access systems have traditionally borrowed from Metro and DCI realms for providing higher capacity and longer reach capabilities. However, in the current scenario, a number of innovations of the access domain could find their way into metro and DCI systems. Chief amongst these is the ability to use single fiber optical line systems or BROADMs. This feature could galvanize metro deployments, freeing us of the need to provision dual fibers for each link and double the inventory of fibers for other applications that may support other interesting application. Finally, having an optical system in the outside plant could help in a material way to bridge some of the pesky ultra long links with a cost-effective optical amplification way station.

9. Conclusions

In this paper, we have described for the first time unleashing the power of coherent optics in the OSP environment. We have shown that when this capacity is combined with approaches like a coherent muxponder (CMP), the entire network can converge and carry 10 Gbps to 100 Gbps optical traffic down to individual nodes, businesses, and homes approximately 150km away from primary locations.

In this paper, we have also indicated technologies that enable coherent optics to survive and succeed on the strand or in a pedestal in the harsh outside plant, where temperatures extend from -40C to +85C even, while extending reach, enhancing capacity, and maximizing fiber utilization. We have also described key technology innovations in BiDi coherent optics, while discussing the philosophy of optical disaggregation and its benefits in spurring innovation across routers, optical line systems, software and optical plugs that continue to fuel new solutions to networks operators.

Use of coherent optics in the outside plant requires care and precision in operations, administration, and maintenance and these topics have also been detailed in the paper.

Abbreviations

AI	artificial intelligence
APC	angled physical contact (connector)
BiDi	bi-directional
BROADM	bi-directional reconfigurable optical add drop multiplexer
C-Band	conventional band (1530 to 1565 nm)
CFP2	c form-factor pluggable type 2 (100Gbps)
CMIS	common management interface specification
CMP	coherent mux ponder
CMTS	cable modem termination systems
CPE	customer premise equipment
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture switch
DS	down stream
DSP	digital signal processing
DWDM	dense wavelength division multiplexing
EDFA	erbium doped fiber amplifier
FTTH	fiber to the home
Gbps	giga bits per second
gNMI	gRPC network management interface
gRPC	Google Remote Procedure Calls
HHP	household passing
Hz	hertz
ITU	International Telecommunication Union
ML	machine learning
nm	nano meter
OLS	optical line system
ORL	optical return loss
OSA	optical spectrum analyzer
OSNR	optical signal to noise ratio
OSP	outside plant
OSFP	octo SFP
OTDR	optical time domain reflectometry
PHY	physical layer
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
QSFP	quad small form-factor pluggable transceiver (40, 100Gbps)
QSFP-DD	QSFP – double density (400Gbps)
rCMP	remote CMP
ROADMs	reconfigurable optical add drop multiplexing
rOAM	remote operations, administration and maintenance
RPD	remote PHY device
RU	rack unit
SCTE	Society of Cable Telecommunications Engineers
SFP	small form factor pluggable
SFP+	small form factor pluggable +

SiPho	silicon photonics
SMF	single mode fiber
Tbps	terabits per second
TROSA	transmitter receiver optical sub assembly
US	up stream
vCMTS	virtual cable modem termination system

Bibliography & References

[1] Coherent Optics for Access Networks, 1st Ed., Zhensheng (Steve) Jia and Luis Alberto Campos, CRC Press, November 2019

[2] Fiber QAM vs. RF QAM, Demystifying Coherent Optics: Amarildo Vieira, Venk Mutalik, Dan Rice, John Chrostowski and Saif Rahman., SCTE Expo 2023

[3] Photon Avatars in the Comcast Cosmos: An End-to-end View of Comcast Core, Metro and Access Networks Venk Mutalik, Steve Ruppa, Fred Bartholf, Bob Gaydos, Steve Surdam, Amarildo Vieira and Dan Rice, SCTE Expo 2022

[4] It's 10PM, do you know where your wavelengths are? Continuous and Pervasive Monitoring of Optical Assets in the Access Domain, Venk Mutalik, Dan Rice, Rick Spanbauer, Simone Capuano, Rob Gonsalves and Bob Gaydos, SCTE Expo 2020

[5] Coherence After Turbulence: Coherent Links for Hurricane Relief in the Sunshine State, Venk Mutalik, Amarildo Vieira, Dan Rice, Eric Moore, John Chrostowski, Scott Wood and Elad Nafshi, SCTE 2023