Fifty Shades of Grey Optics:

A Roadmap for Next Generation Access Networks

A Technical Paper prepared for SCTE•ISBE by

Venk Mutalik
Executive Director
Comcast
1401 Wynkoop Street, Denver
860-262-4479
Venk_Mutalik@Comcast.com

Bob Gaydos
Fellow
Comcast
1800 Arch Street, Philadelphia
267-286-3214
Robert_Gaydos@Comcast.com

Dan Rice
VP
Comcast
1401 Wynkoop Street, Denver
267-286-3214
Dan_Rice@Comcast.com

Doug Combs
Architect
Comcast
1800 Arch Street, Philadelphia
267-286-3214
Doug_Combs@Comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

Introduction

Access network technology is evolving at an ever-increasing rate and the devices that are deployed in the outside plant may need to outlast multiple generations of change. In addition, as demand for bandwidth increases and the number of subscribers per node decreases, more access nodes are deployed, causing the cost per subscriber to increase. This requires an architecture that maximizes invested capital yet allows for flexible adoption of new technology.

This paper proposes a roadmap of 'Grey Optical Aggregation' (GOA) for the outside plant that will lower the cost of distributed access networks. As a rule of thumb, grey optics work well when distances are short and fiber is plentiful. Colored optics, like those used in DWDM systems, are best when the opposite is true. Also as a general rule, and comparatively, grey optics are inexpensive while colored optics are not. The challenge to the HFC industry is to optimize the use of each technology, deep in the plant, so as to maximize the fiber asset while keeping costs at a minimum. There are many ways to do this and hence there are many shades of grey optics.

The GOA architecture begins with lower cost grey optical nodes that are aggregated together at an RPD node location, allowing the subtended nodes to share the capacity of the 10Gbps DWDM Ethernet link. The roadmap culminates in a low powered, environmentally hardened 'Switch On A Pole' (SOAP) that multiplexes multiple 10Gbps grey Ethernet optics and leverages Coherent Optical links of 100Gbps and beyond to extend the headend into the outside plant as close to customer as possible. This allows the operator to pivot between or use multiple access technologies at the very end of the network easily.

In this paper, we begin with a description of grey optics and the benefits and tradeoffs relative to DWDM optics. We will then will describe the process of incorporating grey optics aggregation in DAA networks and demonstrate the benefits of the GOA architecture. We then discuss operational aspects of this new architecture and the various upgrade options. We then describe the SOAP architecture and provide a stable roadmap towards supporting ever growing demands of the future while utilizing multiple access network technologies such as PON, DOCSIS and Ethernet.

1. Distributed Access Architecture

Broadband access networks continue to experience substantial growth in High Speed Data (HSD) capacity demands year over year. Estimates indicate that the Compounded Annual Growth Rate (CAGR) for HSD downstream (DS) is around 35%; for the upstream (US), CAGR is around 20%. Additionally, Comcast is seeing increases in its HSD customer base as well. All in all, prescient analysis of this trend has led to the development of the Distributed Access Architecture (DAA), which has enabled the company to keep up with the capacity needs while also enhancing customer satisfaction. Operational benefits also accrue due to the architecture [1,2].

DAA architecture development began in Comcast around 2015. The architecture called for the elimination of RF Amplifiers in the Outside Plant (OSP) and the reduction of Households Passed (HHP) per node by driving fiber deeper into the network. Initially, this was achieved by using multi-wavelength optimized Analog Optical Transmitters in the DS and Digital Return Transmitters in the US. This called for an optimized wavelength plan and optical passives that mitigated optical non-linearities and fiber effects. Extensive architecture rules and play books were developed to determine node placements and other plant adjustments so as to prepare the network for the future. To date, over a million HHPs have benefited from DAA; simultaneously, the customer experience improved and operational expenses reduced.

Around 2017, Comcast began the move from Analog DAA to Digital DAA. This involved virtualizing the CMTS, and investing in high speed connections from the Primary Headend (PHE) to the Secondary Headend (SHE). At the secondary headend, white box Distributed Access Architecture Switches (DAAS) were placed with 10Gbps DWDM connections to individual fiber nodes. These fiber nodes had the same envelope as the previous analog DAA nodes and were placed at similar locations, as determined by the aforementioned DAA architecture rules and playbooks. These fiber nodes now have a Remote PHY Device (RPD) that takes in the 10Gbps signal and transmits an analog signal in the DS to the home. Similarly, in the return path, the RF signals are combined and fed to the RPD located inside of the node. An architecture of this kind provides high Modulation Error Ratio/(MER)/Noise Power Ratio (NPR) in the DS/US and enables higher orders of QAM modulation, thus leading to increased overall capacity. Additionally, better monitoring capabilities and dashboards enhance the already substantial operational benefits across the board.

In this paper, we describe the Digital DAA architecture in some detail, and explain new architecture variants that can substantially reduce cost, preserve upgrade options and help drive fiber deeper.

2. The Axioms

It is axiomatic that labor and infrastructure costs increase over time and electronics cost decrease over time. Furthermore, construction has various dependencies in the form of permits, crew availability, weather and materials availability. As such, there is an inherent strategic approach that needs to be taken in respect of new infrastructure deployment. Construction is an inherently slow and deliberate process.

Electronics, by contrast, change fast. With Moore's Law, computing speeds increase every 18 months. Dennard's Law predicts lower power consumption in later generation devices, and Koomey's Law predicts miniaturization in succeeding generations of electronics. Designers taking these benefits into account either offer lower cost or higher functionality -- or both -- by manipulating space/power/speed combinations.

It is worth noting that within 2 years of the initial analog DAA program, Comcast has evolved to the next generation of the DAA program, which is to harvest the benefits of speed/space/speed to virtualize and distribute the CMTS, and essentially put a substantial part of the CMTS functionality into the fiber node. Further work along these lines could move us towards Full Duplex DOCSIS (FDX), as well as high split or extended spectrum options. Additionally, there are convergence possibilities between 5G and broadband services in the edge.

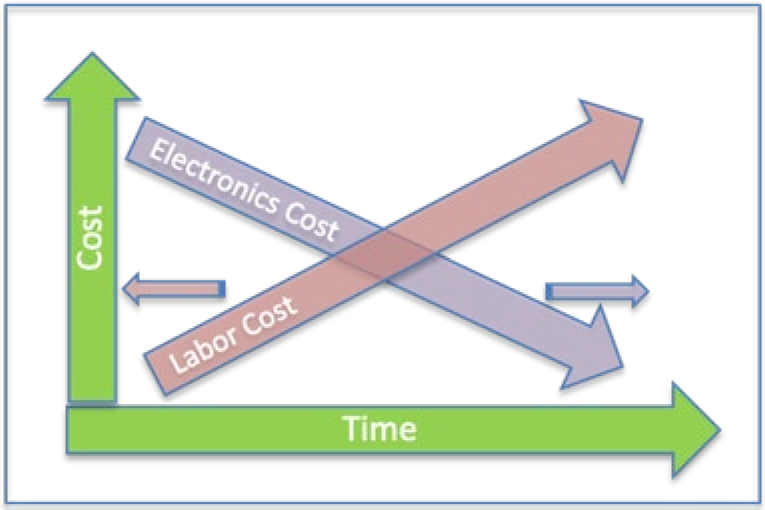

Figure 1 – Illustrating Cost and Time Metrics for Fiber Construction and Electronics Deployment

Driving fiber deeper is of strategic importance to many service providers as a way to further enable multiple services. That said, today's fiber deployments also involve significant electronics dependencies in the form of virtualizing the CMTS, and deploying RPDs and DWDM optics, which may be an overkill for current capacity needs.

There is therefore an urgent need to orthogonalize fiber/infrastructure deployments and electronics deployments. Doing so would enable service providers to construct fiber and accrue economies of efficiency by deploying electronics when the technology is mature and the capacity needs are manifest.

This paper addresses this issue and proposes an architectural concept of Grey Optics Aggregation (GOA), which will be described next.

3. Grey Optics

Figure 2 depicts the spectrum in an optical fiber. Typically it stretches from 1260nm to 1620nm, and is divided in 20nm bands called the Coarse Wave Division Multiplex (CWDM) bands. Various Comcast entities have different optical assets utilizing several parts of the spectrum. However, the wavelength range from around 1525-1570 nm is especially heavily utilized due to the wide availability of erbium doped fiber amplifiers (EDFAs). Therefore this band is further divided in many smaller bands of approximately 0.8nm width or 100GHz of optical spectrum – there are 48 such wavelengths (and sometimes with 50GHz spacing, in which case there are more than 96 such wavelengths) in that band alone. The wavelengths in this band are therefore called the Dense Wavelength Division Multiplex (DWDM) band.

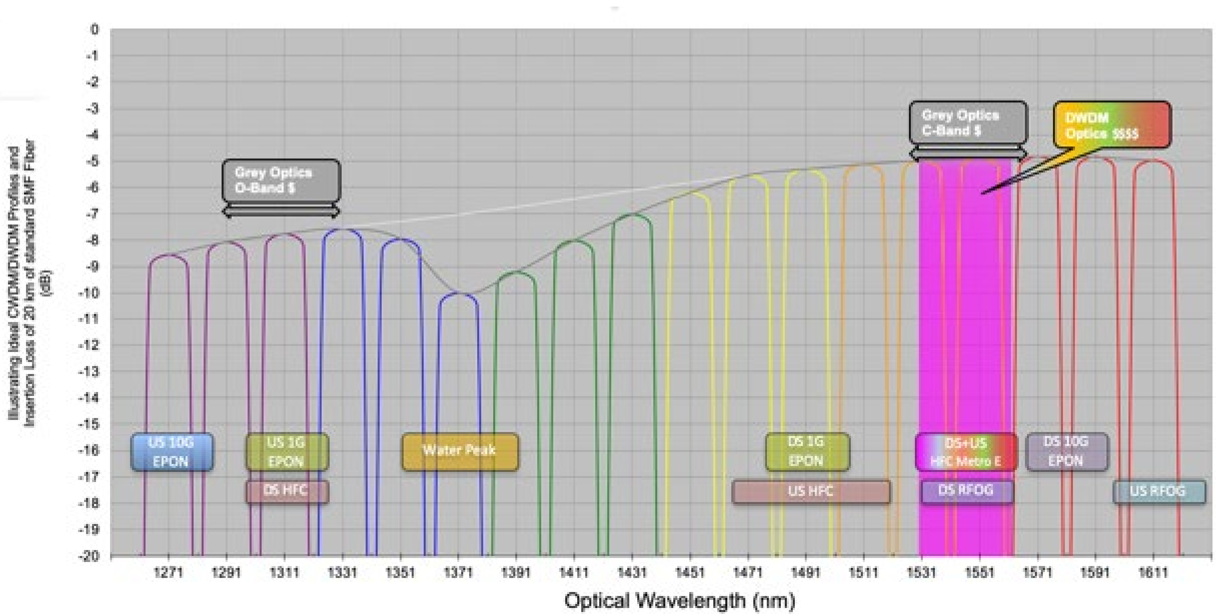

Figure 2 – Illustrating the Optical Spectrim in a single fiber

DWDM Optics have a high requirement to maintain the optical wavelength of their optics. Typically they have to maintain their wavelengths to within +/-0.1nm from beginning of life (BOL) to end of life (EOL). Designing and operating optical equipment with this stringent requirement is what increases the cost of DWDM optics. In common parlance these are also called Colored Optics. As difficult as the manufacture of DWDM Optics is, they are used routinely now for long haul optics and in trunk fibers. This is because the higher cost of DWDM optics is fully justified, in that the cost of fiber construction is even higher for those trunk links.

Grey Optics, on the other hand, do not maintain optical wavelengths in so tight of an optical range when the density is not required. Grey Optics wavelengths can move 10 - 50 times in spectrum location more than the Colored Optics. This enables significantly lower costs (~10x lower) for these, compared to the DWDM optics. Since each of these wavelengths is unconstrained, just one or very few of these wavelengths can be used in one fiber, which is why these are best suited for small distances in the edge, where new fiber construction for DAA type architectures deploy high fiber count cables.

4. Multiwavelength and Singlewavelengths Systems

Presented below is a simplified taxonomy of optical impairments in an optical fiber [3]. When multiple wavelengths (MWL) course through an optical fiber, and at high optical levels, many optical effects and non-linearities emerge and can impact performance. For this reason, Analog DAA, which uses up to 16 wavelengths, uses a carefully optimized industry standardized wavelength plan. In addition, launch optical power, fiber reach, receiver input power, choice of optical passives and transmitter design all have major impacts on performance.

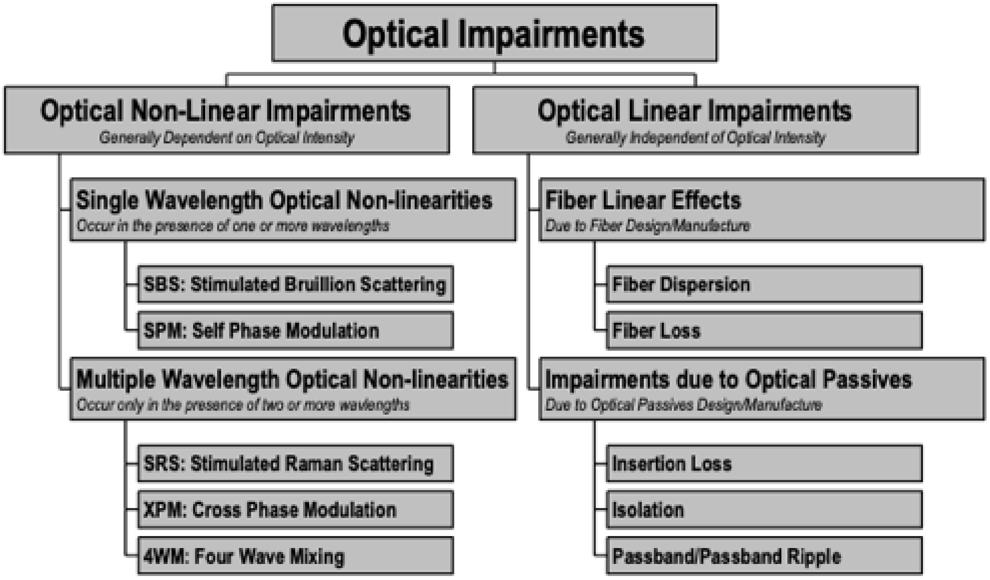

Figure 3 – Taxonomy of Optical Impairments

Multiwavelength systems are affected by multiple linear and non-linear impairments, which can include Shot noise, laser RIN, Receiver EIN, EDFA noise, optical passives ripple and fiber dispersion effects. In addition, impairments from fiber crosstalk, due to imperfect passives isolation, and optical non-linearities such as Stimulated Raman effect (SRS), Cross Phase modulation (XPM) and Four Wave Mixing (4WM), can occur. The resultant MER of DS wavelengths are relatively modest. Indeed, the ability to identify all impairments and manage their impact simultaneously is at the heart of making robust multiwavelength systems work.

Plant Composition vs. Plant Density Matrix

	Multiwavelength System 16WLs + Passives	Single Wavelength System C-Band
Short Reach <2km	MER ~40-42dB	MER ~46-48dB Well Suited for GOA
Long Reach >30km	MER ~38dB Used in Analog DAA	MER ~40-42dB

Figure 4 – Typical MWL and SWL Performance

However, when single wavelength (SWL) systems are used, and for a very limited reach, system design is much simpler and can be shot noise limited. As a result, much higher MER values are obtained. It is not the intent of this paper to describe MWL and SWL system design, but this important outcome is highlighted here because this result enables us to use analog optics without sacrificing DS performance when we next describe the Grey Optics Architectures.

Distributed Access Architectures

Hybrid Fiber Coax (HFC) has been the mainstay of Cable Networks and continues to offer significant benefits over other forms of last mile access, such as twisted pair, commonly used by telcos. Typical HFC systems use fiber nodes and RF amplifiers. While RF amplifiers help drive signals, they are numerous in number, require maintenance, consume power and vary over temperature. For this and other reasons, Comcast began to eliminate RF amplifiers and move to 'N+0' or "node plus zero" amplifiers. N+0, also called DAA interchangeably, reduces the number of actives in the plant by half, strategically drives fiber deeper into the network and reduces HHP/node, increasing capacity. Since the fiber is driven deeper into the network, more RF spectrum and the move to DAA is often associated with increasing the DS to 1 GHz or 1.2GHz (from 750MHz-860MHz) and to 85MHz in the US (from 42MHz) typically found in HFC. This increase in spectrum and the decrease in HHP/Node together give a very robust runway for future capacity growth, while the reduction of OSP actives and increased MER reduces maintenance and enhances the customer experience.

An RPD is a Remote PHY Device that accepts 10Gbps signals optically and converts the signals to conventional RF output. Since this conversion is done using advanced techniques such as Direct Digital Synthesis (DDS), the ensuing MER out of the RPD is very high -- as high as 48-50dB with full 1.2GHz loading. This is headend-grade signal quality, right at the node, and is the primary draw for a move to DAA and to incorporate RPDs in fiber nodes.

5. Basic RPD DAA Architecture

A strawman architecture of RPD DAA is presented in Figure 5. Typically, a conventional HFC 'Parent' node is divided into many smaller DAA nodes. Normally, near the parent node location, an optical passive is placed that has up to 48 ports. Then fibers are drawn to each of the nodes which have an RPD with DWDM Small Form Factor Pluggable (SFP) modules connected to optical passive in the parent node location. The RPD consumes 10Gbps of data and creates RF signals for distribution in the DS. In the US, burst mode RF from all of the ports is combined and fed to the RPD, which takes the analog input, demodulates it, converts it into a digital stream, and sends back to the headend via the DWDM SFP. Although the picture below shows just 9 RPDs, it is typical to see ~12 RPDs on average and as many as 24 RPD DAA nodes in one parent node location.

The muxed optical signals are then de-combined at the SHE (secondary headend). The SHE comprises an optical passive that is a mirror of the one at the parent node location, which then feeds as many SFPs as were used in the node of the outside plant. Each of these SFPs are housed in the DAAS switch with at least as many ports as SFPs. Oftentimes these DAAS ports are aggregated into a switch called a Headend Aggregation (HAAG) switch, and the data is connected to the PHE via 100-400Gbps high speed coherent links that may traverse thru multiple ROADM (reconfigurable optical add drop multiplexers). At the PHE (primary headend), these links are then connected on a one-to-one basis to virtualized CMTS also called vCOREs or vCMTS that are located in server racks, or Pods called PPODs (Primary Pods) that mirror their OSP configurations.

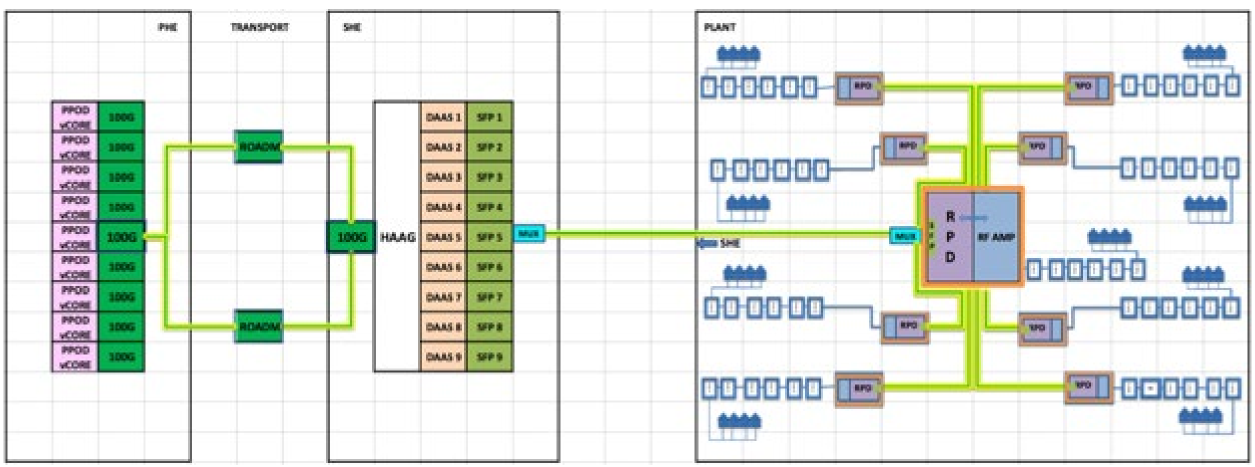

Figure 5 – RPD DAA Architecture

It is easy to see in Figure 5 that each RPD requires 2 SFPs, two pairs of optical passive ports, one DAAS port, access to the HAAG and 100-400 Gbps optics along with a connection to one vCORE. While RPD DAA was designed with about 80-100 HHP/Node in mind, it is common to see much lower densities of around 55 HHP/Node in the outside plant. When this happens, the number of RPD nodes increases proportionally, and along with the cost of RPDs in the OSP, one would also encounter significant costs in the ISP in the form of SFPs, DAASs, HAAGs and vCOREs, and the consequent increases in critical infrastructure (air/power and space) in the ISP locations.

Typically, DAA conversions are done on a secondary headend to secondary headend (SHE to SHE) basis, to enable the move from conventional CMTSs to vCMTSs. This helps the SHE to maintain just one architecture in its inside plant (ISP), and is a more sustainable way to support future capacity needs. Furthermore a move of this kind all but eliminates RF combining mechanisms prevalent in SHEs, and frees up the critical infrastructure for other uses. Perhaps most importantly, this is the only way to plan, build and commission the OSP. OSP construction is a complicated effort, with multiple moving parts from designs to walkouts to permits to traffic control and weather and seasonality. Aerial construction is comparatively simpler and less expensive than underground construction, which could be much more expensive (10-15x) and time consuming.

When density in the OSP turns out to be low, then the consequent increases in RPD node counts can extract substantial costs across both OSP and ISP. Even if the density on average is high, such as that in an urban or suburban location, there are still going to be several pockets where the density is much lower than optimal.

Electronics in the RPD that cost a lot today will cost substantially less in a couple of years. Tying up capital in electronics for small node sizes, where such capacity is not currently needed, would be much better spent on fiber construction activities, since fiber is an enduring asset that lasts for decades. Finally, technology itself is changing fast and the competitive landscape can call for FDX, ES, High Split, 5G front and back haul, EPON/FTTH or SOAP, or a combination all above.

There is thus a need to have an architecture option that enables stable, predictable and cost effective DAA implementation across various densities, and one that has multiple upgrade paths to accommodate capacity increases and transparent to technology changes.

6. Grey Optics Aggregation

Consider the architecture presented in Figure 6. From previous discussions, we know that trunk fiber is scarce with a long reach, and thus is suitable for long reach MWL DWDM 10Gbps optics. The plant fiber from the primary node location to all the child node locations is very short (0-2km), and more importantly plentiful, since it has been freshly laid out, and is thus suitable for SWL Grey Optics. Especially in low- to mid-density areas, where the number of HHPs/Node could be very low, this architecture provides a way to aggregate a constellation of DAA nodes.

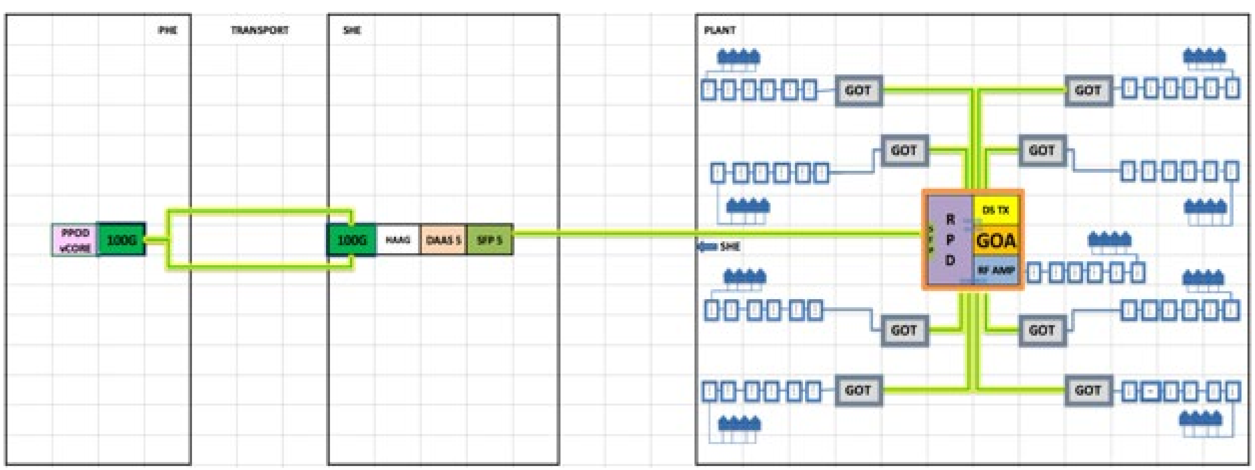

Figure 6 – GOA DAA Architecture

To do this, an RPD Node is modified to have some of its RF output connected to a DS Grey Transmitter. The cable industry has for a long time been using DWDM Grey Analog transmitters for US transmission. Furthermore, there has been a 4x miniaturization of DS DWDM transmitters over the last decade. In this case, both techniques are used together and transmitter bandwidth is enhanced to cover a range of 100MHz to 1.2GHz while maintaining its temperature performance over the industrial temperature (I-Temp) range (-40 to +85C). Furthermore, the wavelength control over the transmitter output is relaxed, which, when combined with the very small link spans, creates the perfect opportunity for a downstream Grey transmitter (SWL, analog) with robust MER and low cost. Typically, these transmitters have 10dBm of output power, which, when split 8 ways and needing to cover <2km, typically arrives at the constellation node receiver at around 0dBm.

We call this central node the Grey Optical Aggregating Node (GOA), and the constellation nodes the Grey Optical Terminating Nodes (GOT). With this terminology settled, and with a friendly reminder that "GOT" in this case does not stand for "Game of Thrones," we can now describe downstream and upstream transmission.

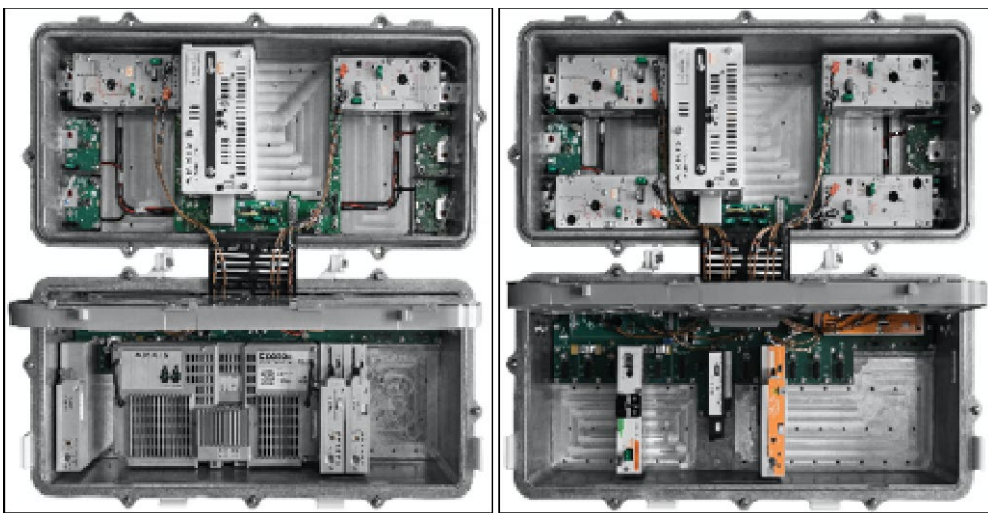

Figure 7 – Typical GOA Node (Left) and GOT Node (Right)

The GOT Nodes have analog receivers, just like in Analog DAA nodes. However, unlike the Analog DAA nodes, the GOT nodes include optical receivers with Optical AGC functionality. Although typical optical input coming is around 0dBm, as indicated earlier, the Optical AGC enables the receivers to put out a constant RF output, regardless of minor changes in received power. This concept will be described later in the paper.

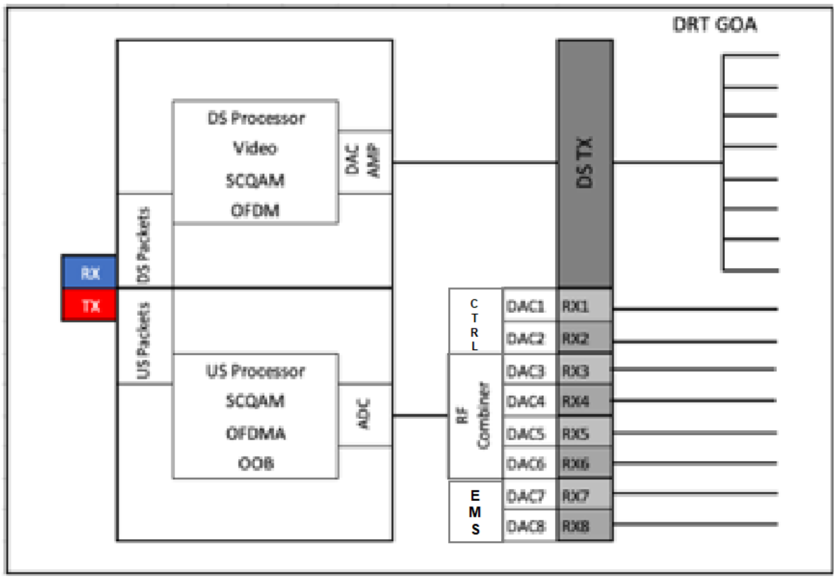

Figure 8 – Block diagram of the GOA Node

In the US, the GOT Nodes have digital return transmitters (DRTs) just like in the analog DAA node. However, unlike the analog DAA node, the GOT Node DRTs have Grey Optical SFPs, which, again are small-form pluggable modules, and a key reason these architectural advancements are even possible. This reduces the cost of the GOT node considerably, while enabling it to span the modest 0-2km links to traverse to the GOA Nodes. At the GOA Node, each of these DRTs are received in separate photo-diodes, not unlike the ones used in headends to receive the DWDM inputs for the Analog DAA system. However, the RF combining is considerably simpler, and these are then combined together with the RF from the GOA node port itself and presented directly to the RPD. In Analog headends, considerable critical infrastructure is given to RF combining. By minimizing and simplifying this, considerable power and space is conserved, and all 8 of the GOT nodes are thus combined in the GOA Node.

Because the RF combining in the US is done via the DRTs, the US system is exceptionally robust and high performing. Figure 7 shows a typical GOA node with two Quad GOA modules that help to provide a level of modularity while still gaining the benefits of aggregation.

Furthermore, a change in the PHY layer technology from the RPD to vCORE does not affect any of the GOT devices downstream of the GOA as long as the RF technology from the GOT node on down remains the same. This feature offers an element of flexibility in accommodating technology changes on the trunk route.

Finally, each of the GOT Nodes has a one-way element management system (EMS). This enables the entire set of critical information from the GOT nodes to be projected to the GOA node. The GOA Node itself has the I2C data bus enabled, so that the RPD is able to connect to and query all results from the GOA module, as well as measure and control the elements in the GOA node itself.

GOA Performance

With the GOA architecture now clear, we can now look at its downstream and upstream performance.

7. Downstream

In the DS, the critical elements in the chain are the RPD, the analog GOA-GOT link, and the performance of the Node itself. The RPD MERs are generally very high, but it is also dependent somewhat on the RF load required of the RPD. Typically the MER of the RPD is in the vicinity of 48-50dB for the 1 to 1.2GHz load. The RF portion of the Node itself contributes to the MER, especially at 1.2GHz loading and its high DAA RF levels (73.6dBmV TCP); for 1GHz loading, the impact of the node is very modest at the DAA levels (69dBmV TCP).

The MER of the GOA-GOT link is dominated by shot noise and laser RIN, since the input power to the receiver is generally quite high and the fiber distance is quite small (<2km and around 0.5km on average). This MER is around 46-48dB for the 1-1.2GHz load, depending upon the optical input to the GOT receiver. For this reason, we expect little meaningful impact to performance due to the GOA-GOT link, which has been borne out by several tests as shown in Figure 9.

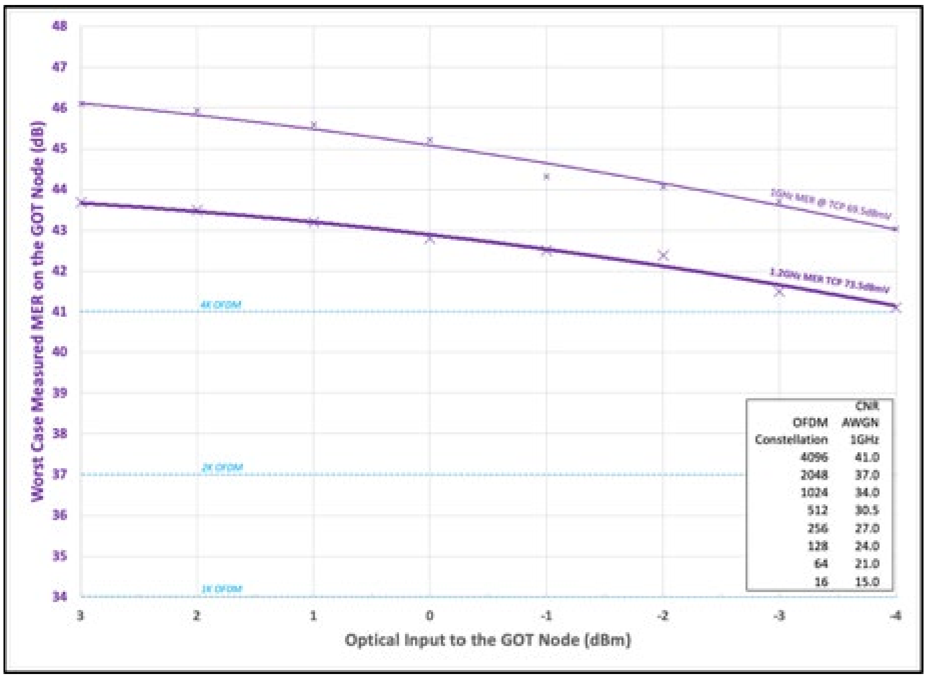

Figure 9 – Measured MER values for RPD, GOA-GOT links and node

Figure 9 shows the measured MER values of the entire chain of RPD, GOA-GOT links and the node itself. At input powers of around 0dBm, it is seen here that the performance of the GOA system is quite robust at 1GHz and at 1.2GHz (again, close attention to the TCP, unit-to-unit variations and temperature performance should be taken into account). In any case, the MER values have a healthy margin over the 4K OFDM QAM.

Figure 10 illustrates the measured RF output of an optical AGC receiver over the range of +3dBm to -4dBm and the Total Composite Power (TCP) of the RF output. It is seen that the receiver established good control over its output and held the node to its stated output power over very large optical input levels. In reality, the optical levels from node to node are much more modest, and this feature goes a long way to making the GOA architecture plug-and-play.

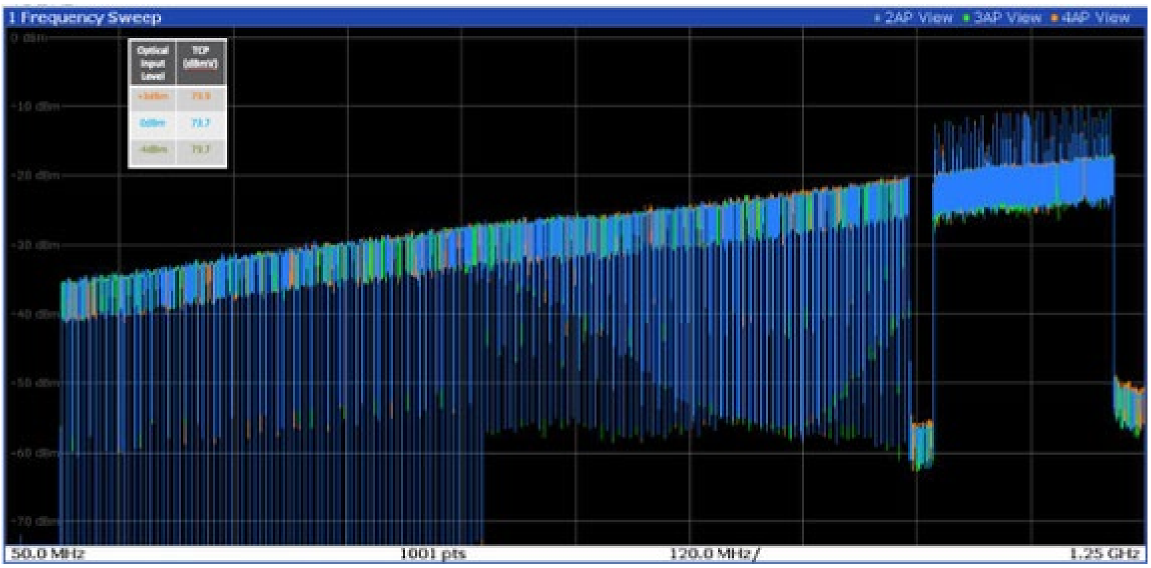

Figure 10 – Optical AGC in the DS Link

8. Upstream

In an earlier section, we explained how the upstream path of this system is based on DRT links. As is well known, the RF levels of the DRT link are independent of optical levels, as long as the link is functional. The optical noise power ratio (NPR) is independent of the optical links as well, again as long as the link is functional. Therefore one gets true plug-and-play link performance, as well as a high NPR, while using DRT links.

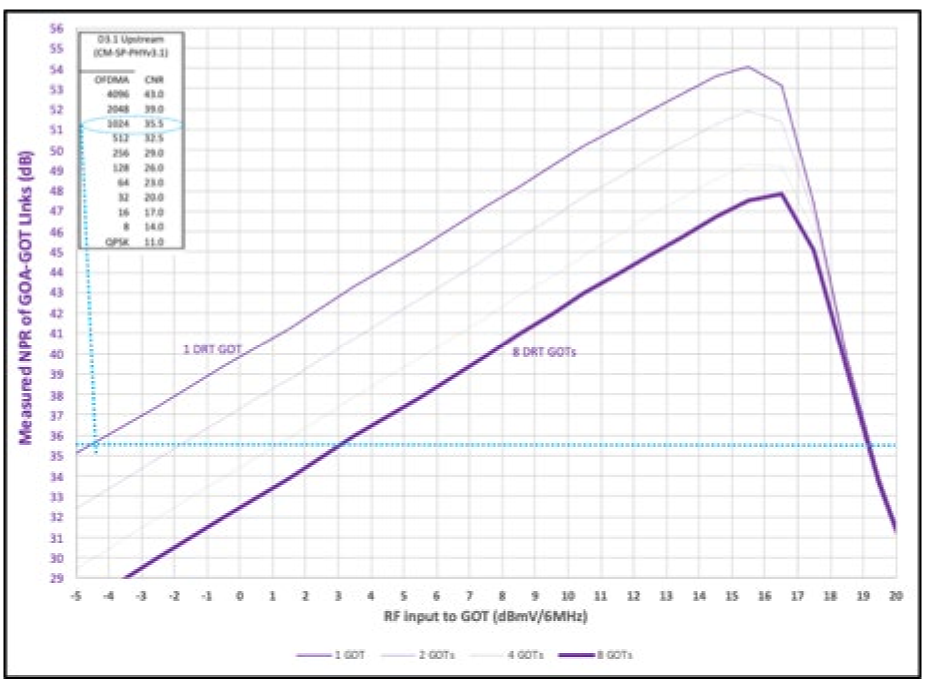

Figure 11 – GOA US Performance

Presented above is a measured NPR of GOA-GOT system. A typical DRT over an upstream mid-split (5-85MHz) would have an NPR curve with a peak NPR of ~54dB and a dynamic range ~23dB for 33.5dB, sufficient for OFDMA 1K Operation. When more DRT links are combined, there is a progressive reduction in the NPR values. With all 8 DRT links from the 8 GOT nodes combined, the NPR curve still has a peak of ~48dB and a dynamic range of ~15dB. Furthermore, at a typical 8dBmV RF input to the node, the NPR is ~40dB, has ~9dB protection from clipping, and a 6dB margin over the OFDMA 1K performance and could support more efficient higher capacity modulation.

Table 1 – DRT RF Levels in the US

GOT	GOA in	RPF Rfin
GOT 1	-4.0	7.8
GOT 2	-11.9	7.8
GOT 3	-15.0	7.8
GOT 4	-8.7	7.7
GOT 5	-3.7	-
GOT 6	-11.7	-
GOT 7	-14.3	-
GOT 8	-8.5	-
Median	-10.2	7.8
Max	-3.7	7.8
Min	-15.0	7.7

Several tests have been conducted over optical levels from -15dBm to -4dBm for various individual GOTs. Unsurprisingly, the RF levels and the NPR performance was completely unaffected. A complete system, built with GOA, multiple GOTs, Cable Modems (CMs) and traffic generators connected has been used to determine that the performance is just as solid over packet and frame losses, over the entire range of testing.

9. Additional GOA Options

GOA nodes with Analog Links instead of the DRT links were also evaluated. Constructing an Analog GOA was with an MDR (multi diode receiver), similar to the ones used in current day RFoG links to eliminate Optical Beat Interference (OBI) [4]. In our evaluations, the US links were at 1610nm; the MDR multiplexed the DS and US lightwaves and used just one fiber for the GOA-GOT links. Since the combining is only over 8 GOTs, instead of the traditional 32 encountered in RFoG links, the performance was quite good and matched that of the DRT links. In addition, the MDR-based approach offers wider BW, such as might be useful with a move to a high split or even FDX. However, and as expected, there is a lot of upstream RF level variation. Comcast has a very wide and diverse footprint and for this reason, we selected the DRT-based approach for our GOA architecture.

The GOA node was designed with only two RF blocks in the base. This is done to accommodate two GOA modules and the DS Tx and maintain the power envelope of the DAA node. Alternately, a single Quad GOA receiver module and 3 RF blocks could be configured in the base as long as the power envelope is maintained. This is particularly useful in moderate density environments where the additional RF block can cover more HHP and may be useful depending on RF coaxial cable routing.

Another variation that can be considered is to configure no RF blocks in the GOA node at all. We call this case the "Pure-GOA", in that the GOA does not cover any HHPs at all and simply performs the aggregating function. While this will reduce the power dissipation of the GOA node and might be a good option to streamline design as well, it comes at a cost. One additional GOA will require an additional strand or pedestal location, and will require a GOT node nearby depending on existing hardline cable routing. In many cases, the cost and complexity of trying to arrange for two DAA sized nodes in close proximity is not justified and so, we have chosen the GOA approach with RF capability.

Fiber management and ergonomics of the GOA construction is an important topic. Certainly, optical splitters and fibers can be accommodated in the fiber tray of a DAA node. However, doing so would require opening the GOA node and risk performance impairments to the GOT constellation for fiber maintenance issues. For this reason, we have uniform sized splice enclosures for each of the GOA and GOT nodes. An added benefit of doing so is the ability to upgrade GOT nodes to RPD nodes with relative ease and is described later in the paper.

It is possible that GOA architecture can accommodate RF amplifiers in the plant. This might be especially useful in underground plant, however any decision to use RF amplifiers has the unfavorable side effect of dramatically increasing the number of field actives. In our analysis, we noticed a doubling of field actives for N+1 type architectures relative to N+0 type architectures. This result must be carefully weighed with known cost of maintaining the additional field actives and the cost and ease of field upgrades to accommodate traffic over the life of the plant and adoption of technology such as FDX.

Finally, creating a headend architectural equivalent of the GOA can enable an easy 4x4 split of an existing node. Innovative options that utilize the essentials of GOA architecture to address legacy plant node splits while moving towards vCMTS and consolidating headend RF combining are worthy of further analysis.

GOA Operations

As impressive as the performance of the GOA-GOT links are, additional considerations must be kept in mind, in terms of operations, monitoring and ergonomics. Since the optics in the all the GOT nodes are Grey Optics, there is a benefit in their cost and simplicity in their procurement.

10. RF Levels across the Plant

Presented in Figure 12 is a schematic of one GOA, one GOT and all the RF levels in the system. The RF output levels of the GOA and GOT exactly match the RF levels of an RPD node or of an analog DAA node. As described previously, the combined effect of the GOA-GOT links with optical AGC ensures that all GOT nodes put out the same set RF levels and promote Plug-and-Play operation in the DS.

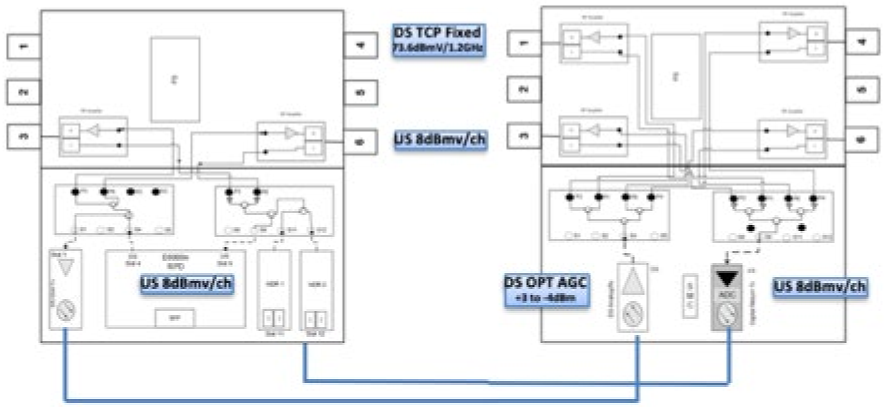

Figure 12 – GOA Operations: RF Levels in the Plant

For the Upstream, the RF levels for the mid-split continue to follow the set 8dBmV/Ch standard for both the GOA and GOT. The RF levels out of the GOA receiver modules are then combined in return combining boards and transformed so that a Unity Gain circuit is maintained from the GOT input to the RPD input, just as conceived in current RPD node standards.

The combined effect of the plug-and-play modules at the DS and US operations, and the modular GOA architecture, enables easy and simple upgrades in both even and uneven traffic growth conditions that will be discussed later in the paper.

11. Fiber Connectivity

Providing fiber connectivity from the GOA to GOT is depicted here. Typically there would be an optical passive with multiple ports (12 ports seems most reasonable, as typically more than one GOA exists within one Parent node domain) in a splice enclosure near to the GOA or a parent node location. The SFP in the RPD is connected to the SHE thru the splice enclosure passives.

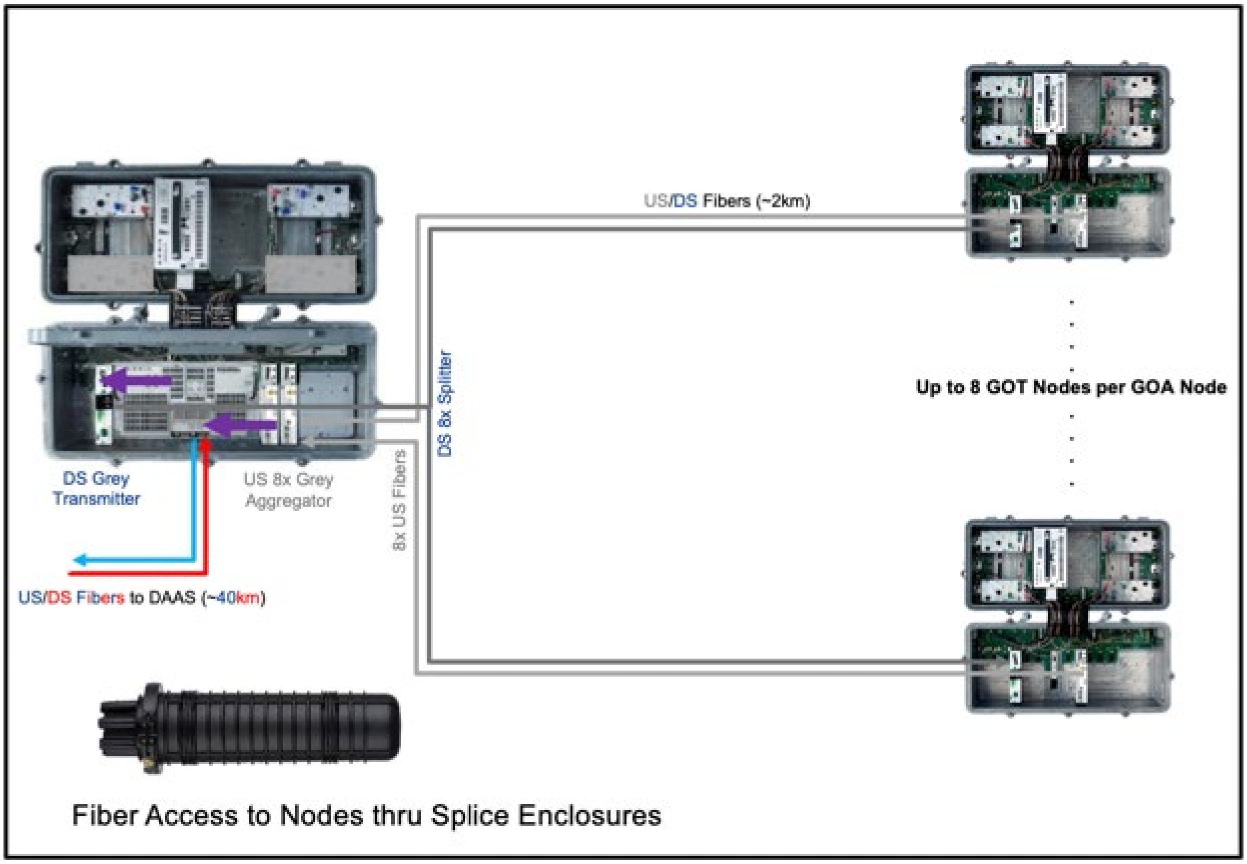

Figure 13 – Fiber Connectivity in the GOA architecture

The RF level out of the RPD is distributed to the RF ports in the GOA node and to the DS transmitter. The output of the transmitter is then brought to the splice enclosure and split up to 8 ways; fibers are then connected to the GOT node. Each GOT node receives the same copy of the DS. In the US, the DRT fiber is connected to the splice enclosure and routed thru to the GOA module. Fiber entry service cables and a fiber connection standard to identify GOT nodes connected to GOA modules have been architected to enable these connections between the splice enclosure and the GOA node.

12. Ingress Monitoring and Control

While the GOA architecture offers great performance, it is at its heart an aggregating network. As with all aggregation points, there is opportunity for any one single element to pump in noise and impact the entire system. In real life, this could have an unfortunate effect: the field is unable to know which GOT was the offending node. As a consequence, a visit to the GOA would be required to open the node and fiber pull (pinch) to isolate ingress and identify the offending GOT node. Then the node would be closed, the ingress hunted down and mitigated; a return visit to the GOA is likely required to get the all clear.

To avoid these trouble shooting challenges, the GOA architecture provides powerful tools to remotely identify ingress effects, with the help of the monitoring and control design decision logic indicated in earlier sections, and described in the next section in more detail. Each GOA-GOT link can be either shut down or attenuated by 6dB with an RPD command. This level of remote control gives one the ability to isolate ingress or noise remotely and enable the operations technicians to identify the GOT node affected and avoid opening the GOA needlessly.

13. Monitoring the GOA and GOTs

Because there are many remote elements connected to the one RPD in the GOA node, there is a great desire to obtain, collate, curate data from all the GOT nodes. This is effectively done by the one way EMS typically offered by the DRT-GOA link and picked up by the GOA module which is connected to the RPD. The following table provides a few representative parameters that can be monitored.

Table 2 - Example Table of monitorable and control parameters

GOA: Monitorable and Control Parameters	GOT: Monitorable Parameters
– General Module Information • Model Number, • Serial Number • Node/FTX/GOA/ • RFAs/PSU – Downstream Transmitter • **Optical Tx Power (dBm)** • **Laser Temp (C)** – RF Amplifiers • Amplifiers Bias • Ingress Control (On/-6dB/Off) – Powering • **AC Input Voltage (V)** • **DC Output Voltage (V)** – Digital Return Receiver(s) • **Optical Input Power (dBm)** • Ingress Control (On/-6dB/Off) • SFP Temperature (C) • GOA-SFP Model Numbers	– General Module Information • Model Number, • Serial Number • Node/FRX/DRT/ • RFAs/SMC/PSU – Forward Receiver • **Optical Rx Power (dBm)** – RF Amplifiers • Amplifiers Bias – Powering • **AC Input Voltage (V)** • **DC Output Voltage (V)** – Digital Return Transmitter • **Optical Output Power (dBm)** • SFP Temperature (C) • SFP Model Number

The ability to dashboard all this information can also help additional systems such as traffic analysis tools, to identify individual GOT nodes and understand their compositions for traffic allocation, which enables more effective capacity management in upgrade scenarios.

Upgrade Options

Figure 14 shows nodes used in analog DAA, RPD DAA, the GOA and the GOT. The point here is that GOA and the RPD nodes have a similar power consumption profile, and Analog DAA node and the GOT nodes have a similar power consumption profile. In a system that has many more GOTs than the GOAs, the power consumption is always going to be less than that of an RPD-only deployment. However, systems are designed with the full RPD power budget in mind to enable future upgrades if and when required.

Also notice that converting a GOT node to an RPD node is as simple as taking out the DS receiver, the DRT and the SMC card and replacing it with an RPD. Note that since the RF levels were all mirrored, it should be fully plug and play. These DS receivers and the DRTs are all relatively inexpensive Grey devices and can be reused elsewhere.

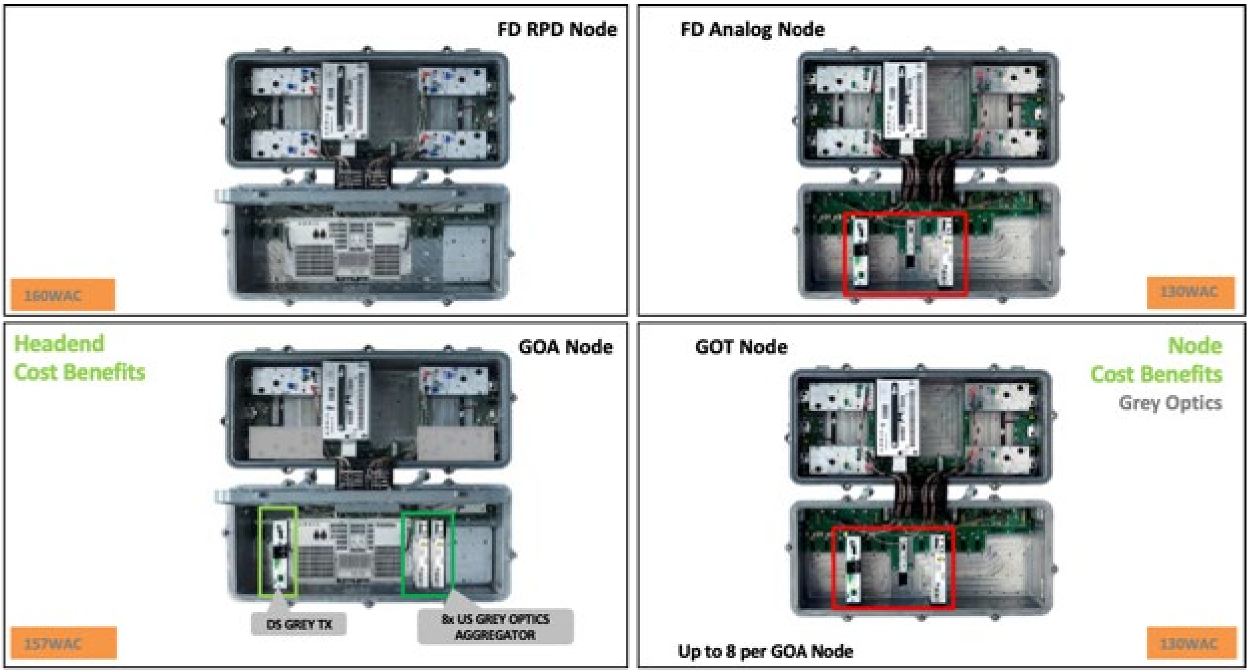

Figure 14 – RPD, GOA, Analog and GOT DAA Nodes in use

14. Uneven Traffic Growth

We begin with GOA architecture in a modestly dense location. However if one or possibly two GOT nodes experience huge traffic growth, such as might happen if an MDU were connected in a GOT node location, then that GOT node is converted to GOA. To do this the DS receiver, SCM and DRT are removed and replaced with an RPD. Then the fibers connecting the Receiver and DRT connect instead to the RPD. At the original GOA location, the fibers leaving the splitter for the GOT node and the incoming fibers going to the GOA are connected to the optical passives ports and linked back to the SHE. At the SHE, SFPs accepting the OSP are connected, a DAAS port is allocated and a vCMTS microservice is spun up in the PHE PPOD. Thus, deploying the compute and electronics just in time, taking advantage of Moore's Law improvements at the time the capacity is required.

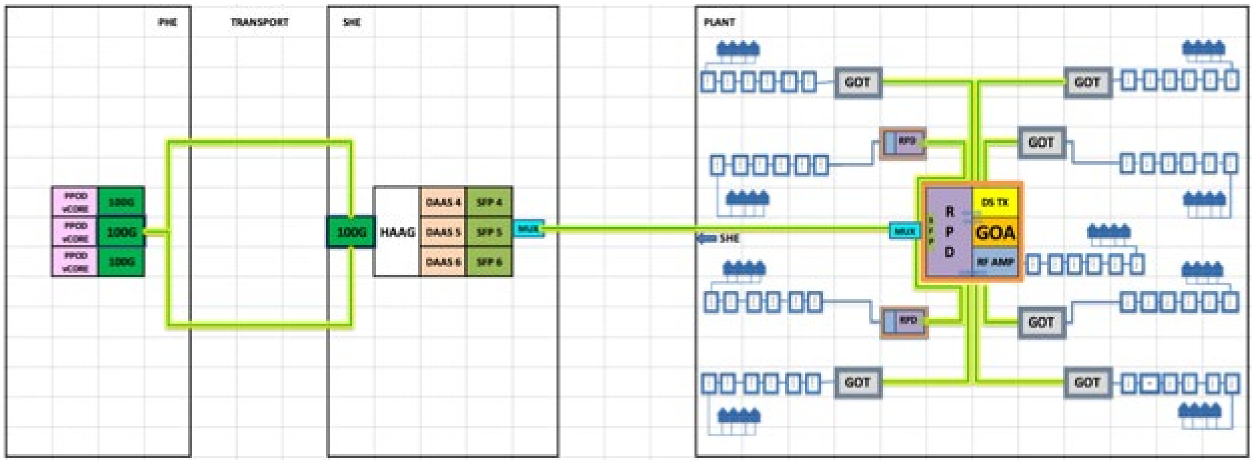

Figure 15 – Upgrade path for uneven traffic groth

It is also possible that a spare set of fibers from the GOA location to each GOT can be allocated and spiced during the original deployment, thus saving a visit to the GOA location when an upgrade is required. This is a time value cost and labor tradeoff, but if the original GOT locations are less likely to be upgraded, one may use the minimal cost approach and use the first revisit to the GOA system as an opportunity to add the additional optical equipment and invest in the labor to provide connectivity to the entire set of GOTs at that time.

It is important to note that the RPD that is deployed for additional traffic is success based. The intent is to protect initial capital outlay; the new RPDs have had the benefit of Moore's, Dennard's and Koomey's laws, collectively making them less expensive and more capable at their time of deployment.

15. Even Traffic Growth

Comcast has great visibility into keeping track of traffic on an RPD basis that extends all the way to GOTs. Thus the ability to perceive and predict creeping traffic increases is well understood.

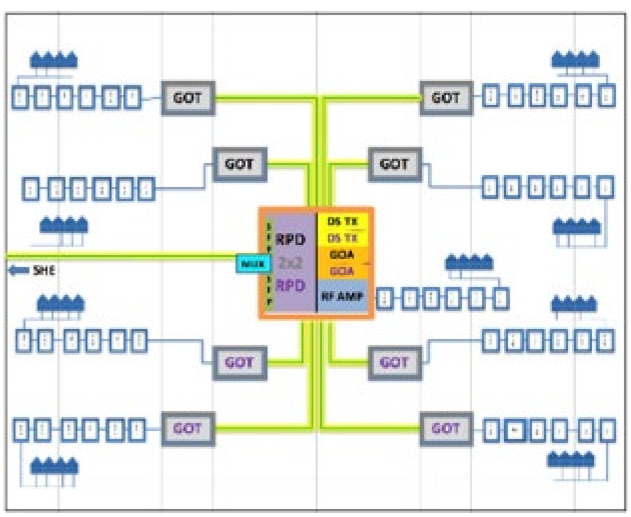

Figure 16 – Illustrating Even Traffic Growth

When critical alarms appear, an elegant way of enhancing traffic overall would be to add an additional DS transmitter and disaggregate the two Quad GOA modules. Then a 2x2 RPD can be deployed with distinct connections to the DS transmitters and quad GOA modules to double the capacity for the entire GOA system. An important part of this type of upgrade is that technicians need not visit the GOT locations at all. A 2x2 RPD, when deployed, also fits in with the Moore's, Dennard's and Koomey's laws, and will likely be more efficient and cost effective when required.

16. Full Duplex in GOA Architecture

The 10G initiative in the industry led by Comcast envisions the ability to use Full Duplex (FDX) type of transmission in the RF plant. In this architecture, the band of frequencies from 85MHz to 684Mhz can be used for both US and DS transmission at the same time on the RF plant. At the heart of this technology is Echo Cancellation of DS and US signals so as to preserve sufficient MER/SNR for adequate demodulation. Additionally, the concept of interference Group (IG) and transmission group (TG) are invoked in CMTS scheduling algorithms to further enhance overall capacity in the network.

A move to FDX is fairly easy in the GOA architecture once the underlying technology options have been standardized. In the simple case, each of the GOT nodes would be converted to FDX RPD nodes and the fiber links established as described in the previous sections.

There is however additional work in the industry to accommodate RF Amplifiers in an FDX environment. The idea is to distribute echo cancelling in the RF amplifiers and use appropriate IG and TG groups to maintain thruput capacity. In this approach, each of the GOT nodes could be treated similarly and each GOT enclave be part of an IG or TG. We note however that this approach is still under development as part of our FDX design activities.

Economics of the GOA

When one looks at the make of the various nodes in discussion, it is easy to see that the GOA node may be more expensive than a regular RPD node due to the addition of the DS transmitter and the GOA aggregation receiver modules, even if a couple of the RF blocks are taken out. The RPD node, in turn, is more expensive than an analog DAA node, after accounting for the DWDM SFPs in both nodes and due to the more mature components used in the analog DAA node. The GOT node is less expensive than the analog DAA node since it uses Grey optics for the DRTs, which are generally less expensive than the DWDM counterparts and do not need optical DWDM passives to enable their transmission. Therefore the OSP cost is less for the GOA and GOT combination than a straight deployment of any other individual architecture.

From an inside plant perspective, the RPD architecture provides great benefits of virtualization and white box switching over the more conventional analog DAA architecture. Better monitoring and visualization of the network is also possible here. The RPD architecture also saves much needed critical infrastructure in the primary and secondary headends. Since the GOA architecture utilizes the same RPD plant -- but quite a bit less of it -- there is quite a bit of GOA savings in the inside plant as well, relative to all other architectures as shown when comparing Figures 5 and 6.

While it is true that the RPD architecture provides smaller service group sizes, the same virtue becomes vice in low and moderate densities.

17. Density and Aerial/Underground Plant

Outside plant is a complex entity; an adequate understanding requires years of experience. In every case, touching the plant requires permitting, construction, and traffic management, among other challenges. For this reason, DAA build guidelines prescribe expected RF levels, performances and practices that are conducive for the long term plant viability.

For purposes of this discussion, we classify plant in terms of density and composition. We selected 4 separate systems with varying densities and plant compositions. We then designed the system as RPD or GOA and compared attributes of the design.

The plant may have higher or lower density, and it is measured in terms of HHP/Mile or as HHP/Node. There is a good correlation between the two, as shown in Figure 16. However, a note of caution is in order for the HHP/Mile parameter. Sometimes this parameter can have very high values if measured in a predominantly rural area with one or two MDUs. In this case, it skews the numbers, since the relatively small mile footprint of the MDU sometimes gives can suggest that the plant is much more dense than that of the plant median. For this reason, the HHP/Node could be a better number, however even that is prone to design guideline changes that might come but infrequently. For example, dropping the RF levels of nodes could reduce HHP/Node, or reducing drop levels could increase the HHP/Node metric, but in either case the HHP/Mile number is unchanged.

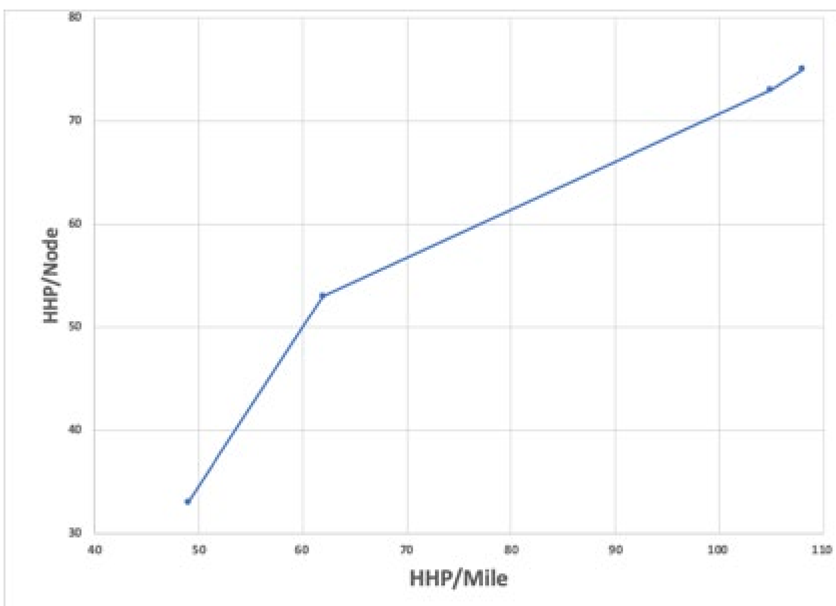

Figure 17 – Mapping Density in Cable Networks

Figure 17 depicts the four systems with their HHP/Mile numbers drawn against the HHP/Node numbers. While the trend of higher HHP/Mile leading to higher HHP/Node is unmistakable, there is considerable variation. It is best to use multiple metrics to understand the OSP, as well as people who are skilled in the art, to initiate and complete designs that involve complex and interdependent parameters.

Like most service providers, Comcast's footprint is diverse, in terms of mixed aerial and underground plant. Aerial plant is both easier to build and considerably less expensive than the underground plant. As rule of thumb for DAA, in aerial plant construction labor cost and material costs are similar, but in case of underground, construction and labor is generally higher than material cost.

In Figure 17, for the two low density plants, one each was selected in a predominantly aerial area while the other was in a predominantly underground area. The same is the case with the two high density areas, where one each was in aerial predominant and underground predominant areas.

18. Economics and Discussion

Presented below is a graph of the total cost in arbitrary units of the design of various RPD and GOA designs in low density high aerial, high density low aerial, high density high aerial and low density low aerial areas. The design was first done using standard DAA guidelines. Note here that the construction costs are similar for RPD and GOA architectures. This is because both of them are high performing fiber deep nodes and have identical reach. The significant reduction in cost is due to the OSP and ISP material cost. This includes lower cost DWDM SFP and Passives, and fewer RPDs, lower cost GOT nodes, DAAS, transport elements and vCOREs. In this example we have not used the added benefits of Critical Infrastructure reductions and is a subject of evaluation internally.

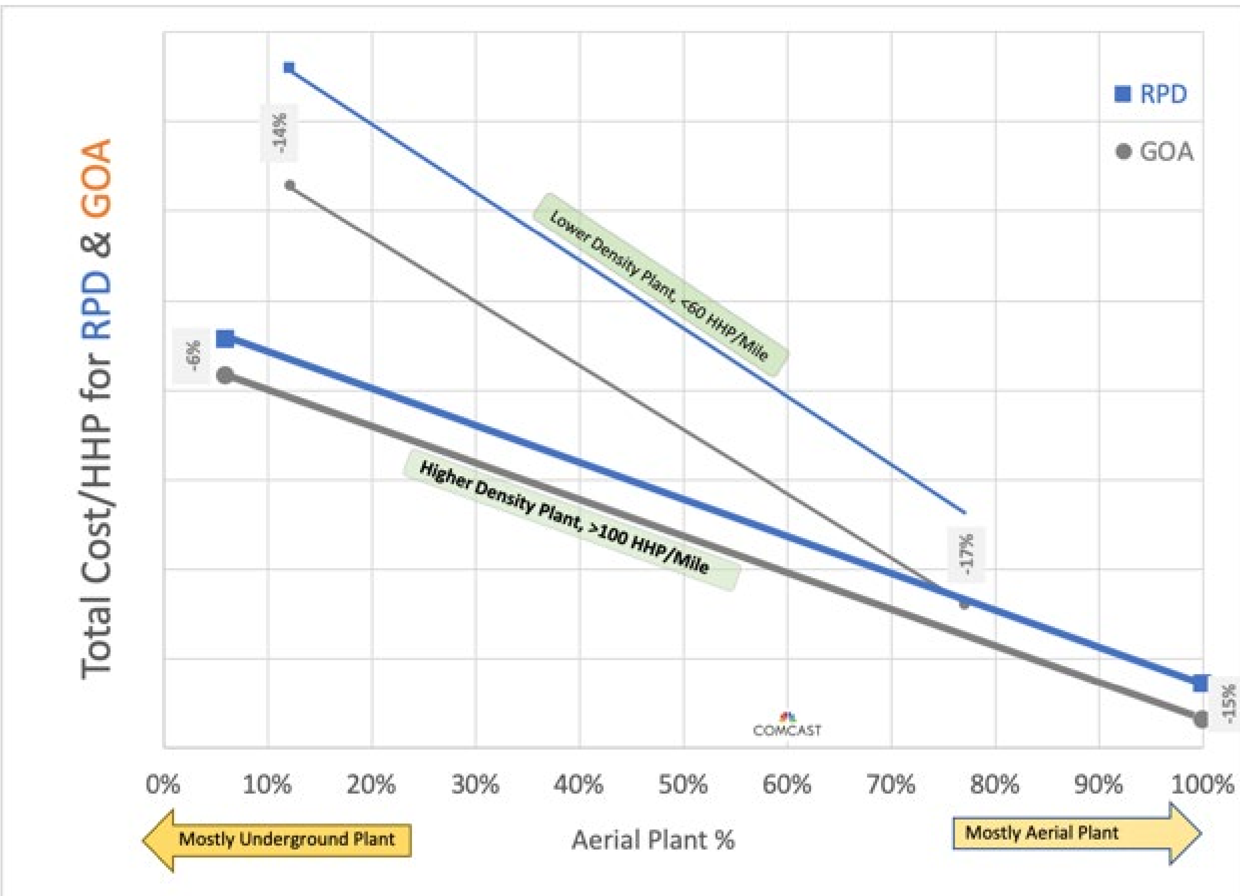

Figure 18 – Total cost of DAA for RPD and GOA architectures

It is seen here that the cost of RPD and the GOA both decrease as the density increases and also as the plant is predominantly aerial. It is also seen that while the GOA is always less expensive than the RPD, its attractiveness grows at lower densities.

Table 3 – Analyzing GOA Cost Reductions over RPD

Plant Composition vs. Plant Density Matrix		
	Plant Composition Mostly UNDERGROUND (>85%)	**Plant Composition Mostly AERIAL (>80%)**
LOW Plant Density <60HHP/Mile	**Total GOA Cost -14%** **GOA Material Cost -41%** **GOA Construction Cost 0%** Mostly UG Plant => Construction Costs High Low Density =>Material Costs High	**Total GOA Cost -17%** **GOA Material Cost -36%** **GOA Construction Cost 0%** Mostly AR Plant => Construction Costs Low Low Density => Material Costs High
HIGH Plant Density >100 HHP/Mile	**Total GOA Cost -6%** **GOA Material Cost -27%** **GOA Construction Cost 0%** Mostly UG Plant => Construction Costs High High Density => Material Costs Low	**Total GOA Cost -15%** **GOA Material Cost -29%** **GOA Construction Cost 0%** Mostly UG Plant => Construction Costs Low High Density => Material Costs Low

The summary quad chart above shows that in all quadrants, the benefit of GOA is manifest. However in three of the 4 quadrants there are double digit reductions in cost. This enables the deployment of fiber deeper into the network for the same capital investment by reducing material cost until it is required in the future gaining the benefits of the improvements in silicon technology. The combination of high Aerial plant with lower densities is especially suitable for GOA. In the underground plant with higher densities, construction cost dominates, which may make it advantageous to deploy RPDs on day one.

GOA Trial

Even though the science, technology and economic trends are favorable, it is always best to run a limited trial to verify the numbers and to gain valuable experience that may help to further fine tune the architecture as we move forward. To this end, a trial in the late summer/early fall timeframe is planned. Our partners in the Central Division have selected an area that enables us to build and stress the GOA architecture in multiple ways, which will help us to further understand operational issues, verify performance numbers, do further economic analysis and provide firmer basis for future deployments.

19. Trial Area Selection and Basic Illustrative Rules

As part of the trial, the following basic apriori rules have been established as illustrated in the table below. If a node were to have more than 100HHP, then it will be constructed as an RPD node as the node density justifies this option. Furthermore, we limit the number of HHP/GOA to no more than 256HHP with no more than 8GOTs per GOA. These limits will be refined as we proceed thru trail.

Table 4 – Preliminary GOA Design Rules

Max HHP/GOT = 100 ~ 128
- If HHP/GOT exceeds 100, consider putting in an RPD-DAA Node just for it

Max HHP/GOA = 256 1.2/85 ~ 304 860/42
- Recommend additional GOA once the HHP/GOA exceeds 256

Max Subs/GOA = 128 ~ 152
- Assuming average 50% Penetration

Max CPE/GOA-RPD or vCMTS = 256 ~ 304
- Assuming Up to 2 two-Way CPEs/Sub (XB6+X1, 45M/26M all over Comcast)

Max GOT+GOA = 8 GOTs + 1 GOA
- GOA system would have Max of 34 RF Ports

Ideally, a trial of the GOA architecture would encompass moderate to low density areas, require multiple GOA nodes within one legacy node footprint, and enable GOA nodes to have the variety of GOT nodes per GOA node, including a full complement of 8 GOT nodes per GOA, and have a variety of HHP/GOA. Fortunately, the first trial of the GOA architecture attains to reality with precisely such attributes.

Figure 19 shows a geographical sketch of the trial site. Currently, a legacy node of 551 HHP is split up into 21 DAA nodes, averaging 26HHP/DAA node, for which fiber construction has been done. These 21 DAA nodes are now divided within 3 groups, each fed by a GOA node. The first GOA node has 240HHP and supports 8 GOT nodes, the second GOA supports 202HHp and 7 GOT nodes while the third GOA node supports 109HHP and 3 GOT nodes.

This spread of GOT nodes and HHP/GOA will give us a good opportunity to study various aspects of the architecture. As part of the trial, it is proposed to analyze and track a matrix of predicted vs. actual performance, operational issues related to deployment and maintenance, historical traffic for the legacy node vs. the actual in the GOA architecture, and ISP cost saving in respect of critical infrastructure.

	HHP	NODE
240 HHP	19	**GOA A**
	18	GOT A1
	27	GOT A2
	20	GOT A3
	40	GOT A4
	38	GOT A5
	33	GOT A6
	23	GOT A7
	22	GOT A8
202 HHP	24	**GOA B**
	26	GOT B1
	39	GOT B2
	15	GOT B3
	29	GOT B4
	19	GOT B5
	29	GOT B6
	21	GOT B7
109 HHP	24	**GOA C**
	31	GOT C1
	28	GOT C2
	26	GOT C3

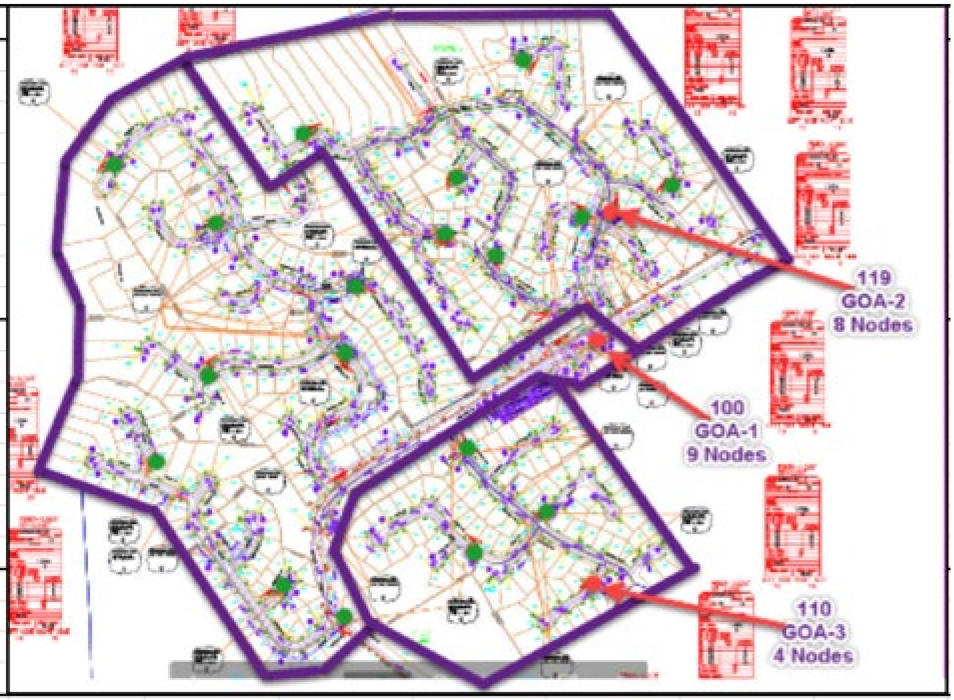

Figure 19 – GOA Trial Area Description

Over time, this will also be used to analyze ingress monitoring and mitigation as well as the process of upgrading individual GOT nodes to RPD nodes as traffic dictates. Furthermore, we can track the operational benefits of fiber predominant architecture and compare it to known trends of already deployed analog and RPD based DAA links.

Switch On A Pole

The GOA roadmap culminates in a low powered, environmentally hardened 'Switch On A Pole' (SOAP) that multiplexes multiple 10Gbps grey Ethernet optics and leverages Coherent Optical links of 100Gbps and beyond to extend the headend into the outside plant as close to customer as possible. This allows the operator to pivot between or use multiple access technologies at the very end of the network easily such as Remote PON OLT or wireless nodes.

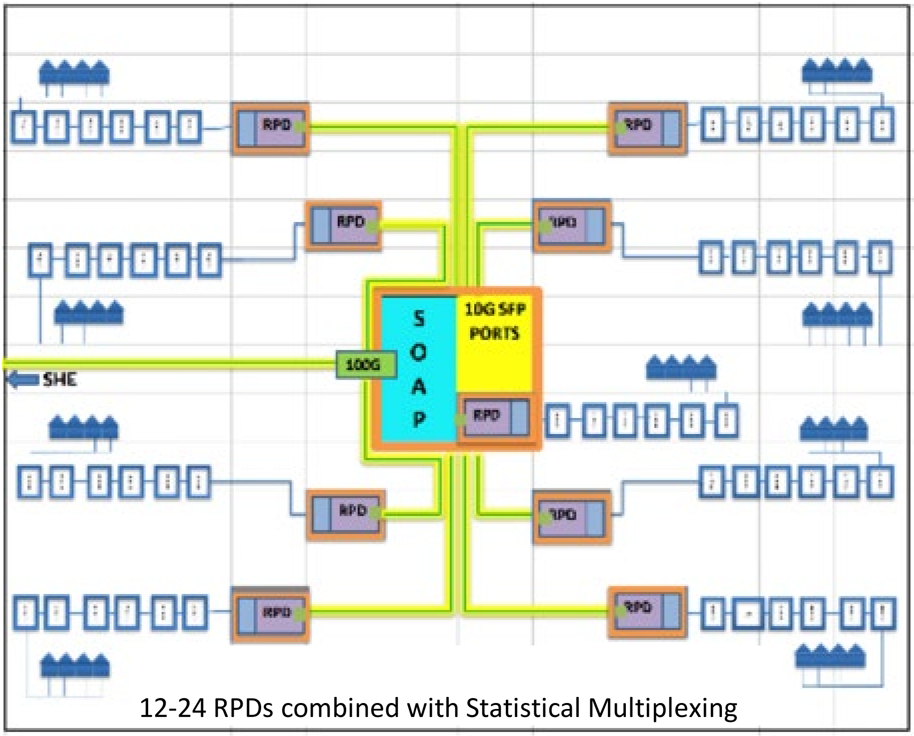

Figure 20 – Illustrating the Switch on a Pole (SOAP) Architecture

As previously mentioned, GOA is especially attractive in low to moderately dense areas. In higher density areas, after factoring in the labor and depreciation for the hardware, we will probably begin with RPDs instead of the GOA. There are still ways to optimize the costs of this solution using grey optics. One of the biggest "costs" of doing DWDM is the opportunity costs of the fiber and the wavelengths in the fiber. The fiber from SHE to node is one of the most valuable assets a cable operator owns and therefore maximizing its potential is paramount to long term shareholder value.

Cable operators have utilized statistical multiplexing, aka "over-subscription" techniques for decades. Rather than use individual wavelengths to address RPDs, a method of multiplexing at layer 2 is required and that is accomplished by pushing a switch deeper into the network using cost optimized commodity switching silicon solutions hardened for industrial temperatures.

The simplified picture above shows only 8 RPD nodes connected to a 100G link, however with the over-subscription model, ideally, the switch would support between 12 and 24 grey optical ports that would feed access nodes, each requiring a pair of fibers. In addition the switch would support a pluggable, dense, coherent capable DWDM wavelength that is 100Gbps (or more) capable of reaching lengths of 40-80km (or more). While this technology has been in existence for quite some time, meeting the low power and harsh temperature requirements of an access network node has been difficult. Recent innovations in opto-electronics are allowing for I-Temp capable 100G Coherent pluggables to be economically manufactured. Merchant silicon network chips have also matured and allowed multiple non-traditional network device manufactures to include switches in their products. In other words, node manufacturers can now use plug and play long haul, coherent optics, grey short haul optics, and a switch chip with their existing housings and power supplies.

By having an Ethernet switch deep in the network, cable operators can pivot between access network technologies quickly with little cost. In fact the switch allows the MSO to offer a multitude of services to consumers in the area. A SOAP device could support low cost, remote OLTs, 5G radios, RPDs, and direct Ethernet switches in an oversubscribed mode. In other words, if only a few customers were using a new technology while the bulk were on another legacy technology, there would be no need to dedicate an entire wavelength to each service leaving one empty and one full. Eventually, the traffic will shift from one to the other. A grey optic in the switch allows for lower deployment costs of the new technology which allows for a faster innovation cycles.

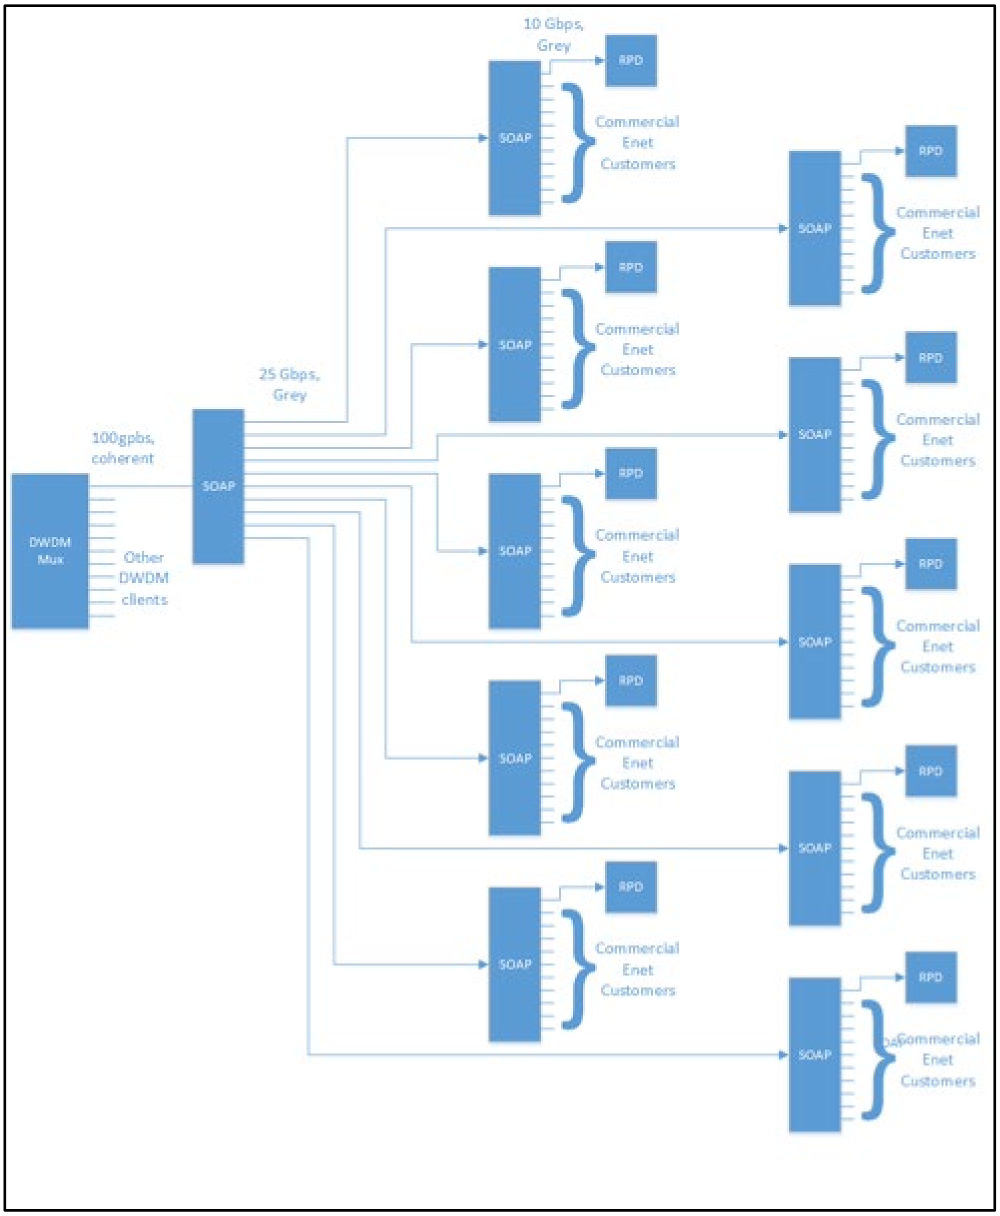

Figure 21 – Cascaded SOAP and the Converged Network

SOAP devices can be cascaded without adding expensive optics or consuming multiple fibers and wavelengths, as is the current case with direct Ethernet services. For example, in order to accommodate 100 direct Ethernet customers today, 200 wavelengths across multiple fibers would be required from the SHE to the neighborhood. Also, multiple fiber splices and muxes would be required. Alternatively, 11 SOAPs could be deployed using one as a parent which aggregates 10 deeper in the network, which in turn aggregate 10 customers a piece. The result would only consume 1 long haul optic, and 2 wavelengths and still support 20 RPDs. Combining residential and commercial services in such a way would allow for rapid commercial service growth without consuming very much optical capacity.

Conclusion

In this paper we proposed a roadmap of 'Grey Optical Aggregation' (GOA) for the outside plant that will lower the cost of distributed access networks deployment -- while still providing the goodness of DAA performance and the potential for lowering operations costs. Just as importantly, the pivot to GOA enables a more effective use of trunk fibers from the SHE to the field, and conserves critical infrastructure at the headends, while driving fiber deeper into the plant.

We described several GOA options that demonstrate the flexibility of the architecture, as well as upgrade roadmaps for uniform and localized traffic growth. Economic analysis of the architecture in aerial and undergound areas indicates savings across the board, but most especially in low to moderately dense areas. A field trial of the GOA in 2019 was described which will help in refining deployment and operational rules.

We described the evolution to an innovative SOAP architecture that leverages a low powered, environmentally hardened 'Switch On A Pole' (SOAP) that multiplexes multiple 10Gbps grey Ethernet optics and leverages Coherent Optical links of 100Gbps and beyond to extend the headend into the outside plant as close to customer as possible. An innovative cascaded-SOAP architecture allows the operator to pivot between or use multiple access technologies at the very end of the network easily.

Acknowledgements

It is now our pleasure to acknowledge the help, support and encouragement of Tony Werner, Elad Nafshi and the senior leadership at Comcast. Tom Bach's insights into operational aspects of GOA made this architecture more robust. Thomas Carroll of the Central Division has earned our gratitude with excellent selection of trial site and for his support. Finally we thank the folks at ARRIS/CommScope, especially Brent Arnold, for their support in the development of GOA technology.

Abbreviations

AGC	Automatic Gain Control
AP	Access Point
BOL/EOL	Beginning of Life/End of Life
bps	Bits per Second
CAGR	Compound Annual Growth Rate
CMTS	Cable Modem Termination System
CWDM	Coarse Wave Division Multiplexing
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
DOCSIS	Data Over Cable Service Interface Specification
DRT	Digital Return Transmitter
DS	Downstream
DWDM	Dense Wave Division Multiplexing

EDFA	Erbium-Doped Fiber Amplifier
EIN	Equivalent Input Noise
EPON	Ethernet Passive Optical Network
ES	Extended Spectrum
FDX	Full Duplex DOCSIS
FTTH	Fiber To The Home
FEC	Forward Error Correction
GOA	Grey Optical Aggregation
GOT	Grey Optical Terminating Node
HFC	Hybrid Fiber-Coax
HHP	Households Passed
HD	High Definition
Hz	hertz
IG	Interference Group
ISBE	International Society of Broadband Experts
ISP	Inside Plant
MER	Modulation Error Ratio
MWL	Multiple Wavelength
NPR	Noise Power Ratio
OBI	Optical Beat Interference
OSP	Outside Plant
PON	Passive Optical Network
PHE	Primary Headend
PPOD	Primary Pods
RIN	Relative Intensity Noise
RPD	Remote PHY Device
SOAP	Switch On A Pole
SCTE	Society of Cable Telecommunications Engineers
SFP	Small Form Factor Pluggable
SHE	Secondary Headend
SRS	Stimulated Raman Scattering
SWL	Single Wavelength
TG	Transmission Group
US	Upstream
vCMTS	Virtual Cable Modem Termination System
vCORE	Virtual CMTS Core
XPM	Cross-Phase Modulation
4WM	4 Wave Mixing

Bibliography & References

1. *Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access*, Nagesh Nandiraju et. al., SCTE EXPO 2016
2. *Aboard the Technology Wave: Surf Report*, Rob Howald et. al., SCTE EXPO 2016
3. *When Wavelengths Collide Chaos Ensues: Engineering Stable and Robust Full Spectrum Multi-wavelength HFC Networks,* Venk Mutalik et. al., SCTE Cable-TEC EXPO 2011
4. *Cable's Success is in its DNA: Designing Next Generation Fiber Deep Networks with Distributed Node Architecture,* Venk Mutalik et. al., SCTE EXPO 2016

