THE YIN AND THE YANG OF A MOVE TO ALL FIBER: TRANSFORMING HFC TO AN ALL FIBER NETWORK WHILE LEVERAGING THE DEPLOYED HFC ASSETS

Venk Mutalik, Marcel Schemmann, Zoran Maricevic, John Ulm
ARRIS Group

Abstract

Many forward looking thinkers in the the Cable industry believe that a move of HFC to an all fiber network is necessarily a move to a Passive Optical Network (G/EPON). However such a change would be quite disruptive and would bring with it challenges of maintaining two separate systems - one for the vast majority of customers now served by HFC and the other for those to be served by PON networks - for a very long time. For starters, only a small percentage of the total network could be converted to all fiber each year, even if all available resources are dedicated to this conversion activity alone. Furthermore, HFC traditionally has been cost effective because of large service group sizes per different offerings (HSD, VOD, VoIP), while PON link budgets generally limit one to far smaller service groups thus leading to much higher startup cost.

What if there was a better way to go all fiber? This ideal way would enable the HFC and the all fiber systems to coexist, with the same equipment at the home and in the headends. It would enable MSOs to continue deployment and activation with the same set of tools and personnel. Finally it would provide a substantial increase in bandwidth, capacity and reliability to last the next several decades.

Recent advances in RFoG technology enable just that. Now Cable MSOs can seamlessly migrate to an all fiber system while increasing capacity, conserving critical infrastructure and enhancing reliability.

Why is now the right time? Up until now, a wider adoption of RFoG was severely limited by a particularly deleterious effect called Optical Beat Interference (OBI). Even very modest amounts of OBI have a severe effect on not only the upstream throughput but on the downstream throughput as well. Recent advances in technology have enabled the complete elimination of OBI thereby unlocking the true potential of fiber.

In this paper, we provide critical insights into the innovations that enable OBI Free RFoG transmission. We will the discuss intrinsic capabilities of what we call Hybrid PON (HPON) technology, explain how this technology works with existing HFC analog and QAM video and D3.0 and D3.1 signals while also being completely transparent to the myriad of traditional PON standards such as the 10G EPON, 1G EPON, GPON and XGPON1.

INTRODUCTION

The broadband industry is fully aware of Nielsen's Law, which has popularized the "Billboard Internet Speed". Nielsen's Law of Internet Bandwidth simply states that a high-end user's connection speed grows by 50% per year. This means it doubles roughly every 21 months. It turns out that this "Law" has held fairly consistently for three decades now, having started with 300-baud phone modems in 1982. Nielsen's Law is important because the highest data speed offered is a determining factor for sizing the network.

The research of Dr. Thomas Cloonan, CTO of ARRIS [1,2] combined with Nielsen's Law is captured in Figure 1. The chart shows the growth from 1982 through the present and projected to the year 2030. This data, referred to as Cloonan's Curve, also reflects the historical 50% CAGR as does Nielsen's Law. Based on this, the highest Billboard Tier will out grow HFC capacity within the next 5-8 years. These predictions have caused cable operators to seriously consider when to start a migration to Fiber To The Premise (FTTP).

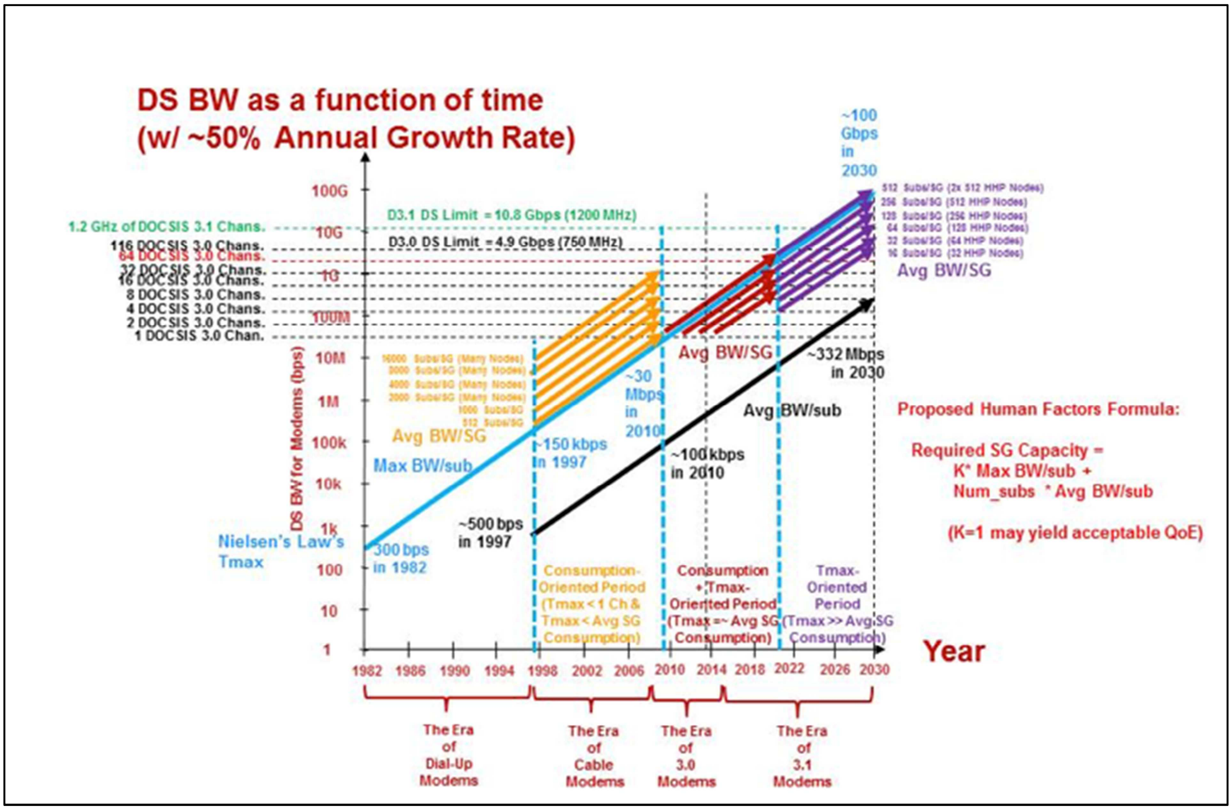

Figure 1: DOCSIS HSD Downstream Traffic Engineering predictions with 50% CAGR in future

As networks reach these critical inflection points many MSOs consider a move to all fiber as a good investment, unlocking the potential of the fiber and the ancillary benefits associated with increased reliability and lower power consumption such networks bring to bear. Competitive pressure from Telcos such as the FiOS and persistent demands from new housing developers all add to the crush of all fiber deployments.

Many today however equate a move to all fiber also with a move to binary-modulated PON based architectures. In reality, there are multiple options, ranging from the aforementioned binary-modulated PONs such as GPON and EPON and a new Hybrid PON (HPON) architecture that is fully compatible with existing HFC and completely transparent to any of the RF physical layer options including GPON/EPON.

We begin this paper with a discussion of all things involving a move to all fiber. While Greenfield builds are increasingly all fiber, the case of converting an existing HFC plant to all fiber is also then considered. The process of converting HFC plant to all fiber plant can be a long duration process, involving newer construction techniques. There will most likely be multiple logical stop points that are sufficiently fiber deep to let demand accumulate before a final push to all fiber is undertaken at the last mile. Since the process of going all fiber is a long duration event, it is critical to have a PON technology that is compatible with the current infrastructure at the headend, in the field and at the homes, has a long shelf life and a simple migration path. Finally we discuss HPON technology and how it satisfies all the requirements given above and aids the migration to all fiber.

<u>TRANSITIONING TO AN ALL FIBER NETWORK IS A JOURNEY NOT AN EVENT</u>

We begin with an interesting thought experiment: if indeed "money is no object", how long it would take to transition a nation to an all fiber infrastructure?

With a population of roughly 320 million, USA has approximately 130 million households. Great majority of those (~120 million) are reached by cable. These are referred to as "households passed" (HHP) or "homes passed" (HP). Pockets of HHP are served by HFC nodes, to a tune of approximately 400 HHP per node, thus there are a total of around 300,000 nodes in the country. An HFC node is oftentimes equivalent to an autonomous Service Group (SG) whose members share into services comprising HSD, telephony or VoD, for example.

While in a typical coax architecture of N+6 there are around 25 actives per node, a rule of thumb in the industry is that there is a fiber node per 5 plant miles in a typical Cable plant. Comcast for example, [3] with 628,000 plant miles and 125,000 confirms this rule of thumb. Furthermore, the typical distance from the hub to a node is around 10 miles. If these statistics are expanded and projected to a national level, there are about 1.5 million of coax plant miles. (300,000 nodes in the nation times ~5 plant miles per node results in ~1.5 million coax plant miles).

The buildup of this amazingly-capable communication infrastructure did not happen overnight – it took roughly 40 years (1950-1990) to complete the first wave of predominantly coaxial-cable plant, followed by ~20 years (mid 1990s to present) of intensive laying of fiber, predominantly for the trunk fiber lines, the ones connecting hubs to nodes. This historical perspective lets us arrive to the construction rate of approx. 37,500 miles of coax plant built per year, in the first 40 years of cable TV network buildout [3,4].

Another approach to take is to look at what a typical fiber construction crew can build per day / week / month, and extrapolate from there. In an urban setting, a crew may do a mile per day of an aerial plant, or it may do half-that if an underground plant. (A separate splicing crew would follow the construction crew to do all the fiber splicing and connecting). These assumptions give us about 2.5-5 miles per week, at 5 workdays per week. If we look at a blended aerial/underground rate of ~4 miles per crew per week, for a typical construction season that lasts 6 months or 26 weeks – it would result in ~100 miles per crew per year.

To complete a coax network to an all-fiber network transition over, say a period of 20 years, it would take ~75,000 miles of rebuild per year, and require ~750 fully-dedicated construction crews to do so.

Verizon FiOS deployment for example took roughly 8 years to build fiber to roughly 18 million homes [5,6]. Using our per-node statistics from before, the 18 million HP number corresponds to ~45,000 nodes and represents roughly 300,000 miles, constructed over 8 years. From these two numbers we get a construction rate of 37,500 miles of fiber route per year, which, interestingly enough, is exactly the same as the estimate of rate with which the first wave of coax construction took place! So if the Verizon FiOS buildout rate were projected onto the whole nation, it would take around 40 years to fiberize the nation.

During the transition, a legacy HFC network has to be operated, side-by-side, next to an all-fiber network. Setting up, connecting and maintaining the network and provisioning of services doubles – if HFC were to remain HFC and FTTP deployments exclusively go the way of a traditional PON technology. It is therefore essential to secure a technology that enables a harmonious coexistence of HFC and FTTP while not excluding the options of deploying traditional PON services for residential or business customers as needed. In this paper, we introduce just such a technology, which we refer to as Hybrid PON (HPON).

Many MSOs have considered:

1. Capping coax plant growth and switching to fiber-only for the new builds
2. Success based fiber deployment for businesses that need it and are willing to pay for it
3. Selective buildout to a few very high-demand users, who need and can pay for higher service rates. As a result of this, in addition to satisfying the high bandwidth users, their former service group members get a fairer access to the newly available bandwidth

A “Fiber-to-the-Curb” (FTTC) approach enables all the three scenarios described above. With a “mini-node” (or ONU) located at the curb - say at a typical tap location - fiber route will pass by virtually all potential customers, so the whole area is described as a “fiberhood”. All of the standard triple-play services are available at all the homes-passed, by the virtue of the existing coax drops. The new HPON architecture and technology to be described in this paper will enable very high capacity bandwidth offerings over standard coax [15].

For those customers who require more, a separate or additional fiber-drop is drawn from the mini-node to the customer premise, to provide a dedicated high-speed link, of a multiple Gigabit Symmetrical data rate. But this step needs to be taken only as needed, with options to connect any and all of the homes-passed [17]. This approach provides high capacity overall, preserves the option of serving high data customers with fibers and obviates the immediate need for fiberizing the last 200 feet or so of coax - which would be a huge cost and time saver – and upgrading it as part of a “success based” schedule.

An interesting technological option becoming more economically available nowadays comprises “coring” the existing hardline coax cables by removing the inside conductor and the dielectric foam, and then using the outside conductor as the conduit for the fiber optic cable. Such hardline cables typically terminate at the curb [14]. This approach provides for a fast and no-mess construction, with considerably lesser chance of blowing through a neighborhood electrical or a gas line.

TRAFFIC ENGINEERING FOR THE NEXT TWO DECADES

Research by co-author J. Ulm [7] of several large North American MSOs reveals that only a fraction of a percent of their customers actually subscribe to the Top Billboard Service Tier.

The study provides a breakdown of subscribers per Service Tiers. This is shown in Table 1. Anywhere from 76% to 94% of the subscribers fell into the most popular Common and Economy Tiers. The remaining subs fell into the Performance Tiers.

Subscriber Distribution	MSO #A	MSO #B	DS Speed	CAGR
Top Tier, Billboard speed	0.2%	0.1-1%	300 Mbps	50%
Performance Tiers	24%	5-18%	75	32%
Common, Popular Tiers	64%	50-72%	25	26%
Economy Tiers	12%	15-40%	5	15%

Table 1: Subscriber Mixes across the Service Tiers

In addition to the subscriber distribution across the different service tiers, the study also uncovered that the service tiers were all growing at different rates! While the Top Billboard Tier was following Nielsen’s Law with a 50% CAGR, the Common Tiers had a growth rate of half that. But this actually makes sense. The operators are the ones actually controlling the growth rates. The Top Billboard rate is for bragging rights, yet only impacts a fraction of the customer base. The 50% CAGR results in CPE equipment being obsolete for that tier every 2-3 years. A small price to pay (or may be not) for top bragging rights.

Meanwhile, operators try to get as many years as possible out of the CPE investments for the majority of their customers in the Common and Economy Tiers. The lower CAGR for these tiers means that CPE equipment will remain viable for 5-8 years before being relegated to customers in lower tiers, making for a very financially successful business plan.

Figure 2 shows the predicted growth rates for the various Service Tiers over the next 20 years. If we consider 10Gbps to be the approximate capacity limit of HFC utilizing DOCSIS 3.1, it shows that the Top Billboard Tier still runs out of capacity in 6-8 years. However, the Performance Tiers have a 12-15 year runway before exhausting HFC capacity. Finally, the Common and Economy Tiers (e.g. 76%-94% of subs) are happily content with HFC capacity for 20+ years. In fact, if their growth rates stall then the majority of HFC subs might live on HFC "Forever".

The other interesting fact from this study is that traffic engineering for the next decade will be dominated by the Service Tier burst rate as defined by the DOCSIS Tmax parameter. This means that once below ~500 subs per SG, further SG size reductions have diminishing benefits. This is very important in an FTTP strategy. Many PON architectures are optimized for 32 or 64 ONU, while traditional DOCSIS systems handle eight times as many devices. What this means is that the PON architecture will require 8 times as many OLT ports as needed for CCAP.

HPON – AN FTTP TRANSFORMATION

A cable operator considering a significant FTTP investment to every subscriber in the near term may be providing capacity at significant expense for a need that might not materialize for decades, if ever as shown in Figure 2. The Hybrid PON (HPON) architecture allows the operator 'success based' steps towards FTTP, with more modest near term investment. Network capacity grows as it is needed and targeted at only the subscribers that actually need it.

A good first step is to push fiber a bit deeper (e.g. down to 250 Homes Passed) and upgrade the HFC plant to 1GHz or 1.2GHz and maybe move to an 85MHz upstream. While DOCSIS 3.1 will work in the existing HFC plant and give operators significant benefits; this upgrade helps enhance D3.1 capabilities. For about 20% of the cost of pulling FTTP to all subscribers, the operator now has an HFC plant with the equivalent capacity of 10G/1G EPON!

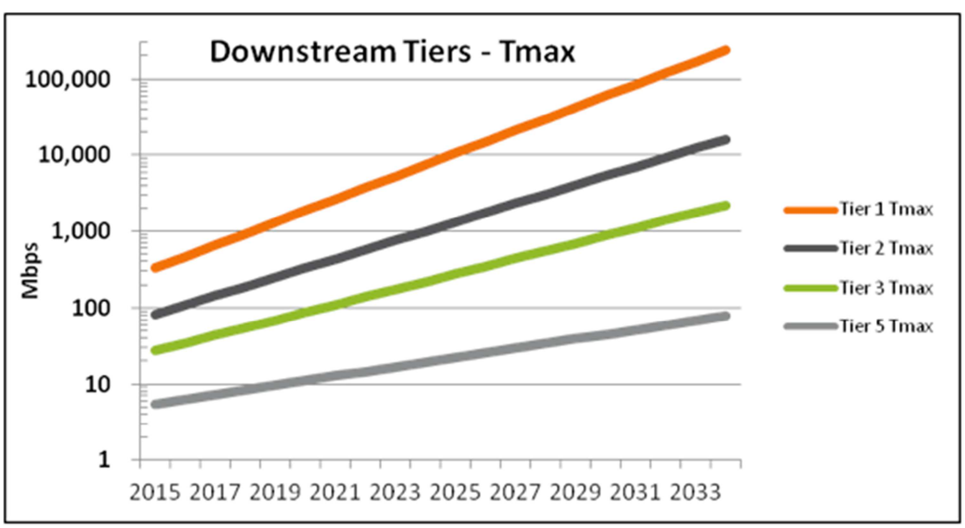

Figure 2: Multi-Service Tiers Growth

As was shown in the traffic engineering analysis in Figure 2, this relatively small investment in HFC has effectively doubled its capacity from ~5Gbps for 750MHz to 10Gbps for 1.2GHz D3.1; and will still be used for 20+ years by the majority of the Operator's subscribers. Plenty of time for the Operator's to recoup their investment.

Eventually, the Top Billboard Tier will out grow HFC, and the operator will need to move that small handful of subscribers to FTTP. With the HPON strategy, every time fiber is pushed deeper, sufficient dark fiber is pulled with FTTP everywhere as the end game. In this case it is pulled from the nearest optical node to the premise. These are all costs that would have been part of a full FTTP upgrade anyway, but have now been deferred in time. Once the first FTTP user is lit up, then the rest of the neighborhood is ready for FTTP service, with only the final fiber drop needing to be pulled. Also, operators can push optical nodes deeper at this time and eliminate even more of the HFC active components.

The HPON strategy allows operators to incrementally enhance all of their existing HFC plants in parallel and roll their FTTP evolution on a success driven basis. Doing a massive full FTTP migration would have forced the operator to this one serving area at a time and stagger this over a very long period. This means some areas may be under-served with HFC when they have Top Billboard Tier customers that need FTTP sooner.

COMPARING HPON AND TRADITIONAL PON CAPACITIES

Part of the attractiveness of HPON is that it is agnostic to the underlying technologies being transported over the FTTP, and will support a simultaneous mix of these technologies. So, one customer can get 10G/10G EPON if that's the service needed while other customers get OBI-free D3.1.

Note that HPON-based D3.1 has some interesting inherent advantages. Because the downstream and upstream are carried on separate wavelengths, their spectrums can overlap. In the downstream, spectrum might range from 54-1218MHz. This allows all legacy signals (e.g. STB OOB, FM music band) to be carried in their original spectrum and avoid the headaches caused moving the upstream split on existing HFC.

HPON-based D3.1 has even greater benefits in the upstream. With OBI-free operation and enhanced SNR (which is detailed later in the paper), the upstream can operate freely at 204MHz with up to 4096-QAM modulation. This is almost a 2.5G PHY rate and nets almost 2 Gbps of data capacity. Most importantly, the data rate is sufficiently high to offer a Gbps upstream service tier. This is something that 1G EPON, GPON and 10G/1G EPON cannot claim.

The bottom line, if we normalize HPON-based D3.1 to traditional PON PHY Rate terminology, it would be roughly equivalent to a 6G/2.5G PON. Taking a look at other possible D3.1 configurations, the downstream and upstream tables show how its capacity would align with various traditional PON technologies.

EPON SCHEDULER EFFICIENCIES

We now look at how the standard CMTS scheduler in an OBI free HPON architecture might match up with an EPON scheduler over a wide range of parameters including large Service Groups.

Previous research into EPON efficiencies [8,9] shows that the EPON Scheduler efficiencies are affected by several key parameters: ONU count, LLID per ONU, and the Grant Cycle time. These analyses were done for a very limited number of ONU/LLID since most people deploying PON were looking at using 32 or maybe 64 ONU per PON.

Downstream Spectrum	Nominal Data Capacity	PON Equiv
32 '3.0' + 96MHz OFDM	2 Gbps	GPON, 2 x 1G EPON
32 '3.0' + 2x192MHz	5 Gbps	2 x GPON, ½ 10G
32 '3.0' + 4x192MHz	8.7 Gbps	10G EPON, XG-PON1, NG-PON2
12-24 x 192MHz	~20-40Gbps	NG-PON2

Table 2: Mapping HPON D3.1 Downstream Options to PON Equivalents

Upstream Spectrum	Nominal Data Capacity	PON Equiv
85 MHz OFDMA	750 Mbps	1G/1G 10G/1G EPON, GPON
RFoG 204 MHz OFDMA	1.8 Gbps	EPON w/ 10G/1G co-exist, XG-PON1, NG-PON2 (2.5G)
RFoG 500 MHz OFDMA	~5 Gbps	EPON w/ 10G/1G co-exist
RFoG 1.2 GHz OFDMA	~10Gbps	10G/10G EPON, NG-PON2

Table 3: Mapping HPON D3.1 Upstream Options to PON Equivalents

The size of the PON SG was generally limited by the optical power budget, which is very sensitive to the distance. So maybe the service provider would have 32 ONU per PON at 20km.

Some recent advances may allow operators to deploy PON over longer distances and with larger fan-out. Technologies like Remote OLT and PON extenders can decouple the trunk link budget from the ODN fan-out. Conceivably, PON can support 128 or 256 ONU per SG as well.

However, PONs are a pure polling system whose overheads increase with SG size and Grant Cycle Time (GCT). While 10G EPON upstreams promise to make data transmission 10 times faster, the truth is that the PHY overheads such as laser turn on & off times may be the same as 1G upstream or just marginally better. Early PON systems also tended to be focused on best effort data services. As more latency sensitive services like voice and gaming are deployed, then Grant cycle times must be increased and additional polling overhead is incurred. The impact of all these overheads for large SG is an area for further research and will be detailed in subsequent papers.

While talking about upstream EPON efficiencies, it is also important to note the impact of 1G/10G upstream coexistence. It turns out that "10G EPON" actually comes in two flavors: 10G/1G and 10G/10G EPON. They both have 10 Gbps downstreams, but their upstream PHY rates are 1 Gbps and 10 Gbps respectively. The wavelengths for the two flavors are very close so they share a common optical receiver in the OLT. This means that the scheduler must treat them as a single SG and only allow one to transmit at a time.

The problem this creates is very similar to that seen in the WiFi world when 802.11g was introduced and had to co-exist with 802.11b devices. When the slower devices transmit, it takes a proportionately longer time to send its data. When a 1G upstream transmits a packet, it takes ten times longer to transmit as the same size packet from an ONU with a 10G upstream.

Some operators may want to deploy 10G/1G initially because they are more cost effective. Later, 10G/10G devices are deployed into the field. However, the co-existence issue might cripple upstream capacity. With a 50/50 traffic mix between 1G and 10G upstreams, there might be less than 2 Gbps of upstream capacity available. The impact of 1G & 10G coexistence is another area for further research and will be detailed in subsequent papers as well.

WHAT IS IN A NAME? PASSIVE OPTICAL NETWORK OR POWERED OPTICAL NETWORK

Cable operators requirements tend to have long distances and limited trunk fibers. While typical distances from the headend to the node are around 20 to 40 km, some of the longer links approach 60 to 80 km. As networks have grown manifold since the trunk fibers were laid out many decades back, the trunk fibers are always at a premium.

Typical Passive Optical Networks (PONs) serve anywhere from 16 HHPs to 128 HHPs with an average of around 32 HHP. There is therefore a deep relationship between the trunk distance traversed from the headend to the splitting location and the number of fiber splits that can be supported by a PON network. For example, a typical 24 dB PON network can support 20 km with 32 splits or 10 km with 64 splits or 2 km with 128 splits.

Because of the typical Cable operator distance requirements and the scarce trunk fibers, it is nearly always the case that PONs will always have remote OLT (Optical Line Terminator) or PON Extender locations very close to the fiber splits. These remote OLTs or the PON extenders typically perform OEO type of an operation to enhance the splits and the reach and many times also do Wavelength translations and aggregation to more fully utilize truck fibers. Therefore none of the commonly accepted PONs are really passive in the true sense and are really Powered Optical Networks.

	GPON	*1G-EPON*	**XGPON1**	*10G/1G-EPON*	*10G/10G-EPON*	
Spec	*ITU-T G.984*	*IEEE 802.3bk-2013*	*ITU-T G.987*	*IEEE 802.3bk-2013*	*IEEE 802.3bk-2013*	*units*
PON Rate	2.488G / 1.244G	1.25G/1.25G	10G/ 2.5G	10G/1.25G	10G/10G	*Gbps*
DS λ	1480 - 1500	1480 - 1500	1575-1580	1575-1580	1575-1580	*nm*
US λ	1310 ±20	1310 ±50	1260 - 1280	1310 ±50	1270 ±10	*nm*
Split Ratio	1:128	1:64	1:128	1:64	1:64	
Max Reach	20 (A, B, C) 40 (B+, C) 60 (C+)	PX10: ≥10 (1:16) PX20: ≥10 (1:32), ≥20 (1:16) PX30: 20 (1:32) PX40: 20 (1:64)	20, 40	PRX10: ≥10 (1:16) PRX20: ≥10 (1:32), ≥20 (1:16) PRX30: 20 (1:32) PRX40: 20 (1:64)	PR10: ≥10 (1:16) PR20: ≥10 (1:32), ≥20 (1:16) PR30: 20 (1:32) PR40: 20 (1:64)	*km*

Table 4: Various Traditional EPON Standards

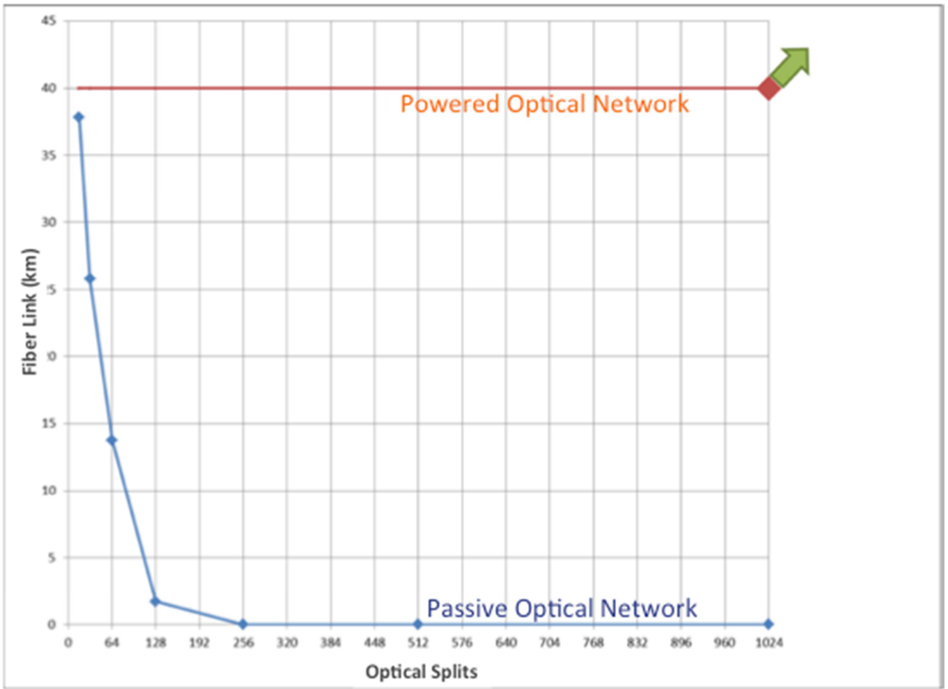

Figure 3: Illustrating Fiber Link vs. Fiber Splits for the passive and Powered Optical Networks

Typically, however this is not a serious problem. This is because if this were a Greenfield design, it is always possible to provide power for the remote OLT locations or to the PON Extender locations. Brownfield designs will already have power due to existing Coax network. While remote OLTs typically require large amounts of power and specialized Cabinets with heat exchangers, the PON Extenders, typically require more modest amounts of power and can typically be strand mounted within existing optical node housings. However it is always the case that typical PON systems require cabinets in one way or another to accommodate fiber splicing.

The Hybrid PON architecture described in this paper uses technology that eliminates OBI and facilitates the complete reuse of existing DOCSIS infrastructure. In addition, it enables data capacity that exceeds 10G EPON in the upstream and downstream while it simultaneously enables vast numbers of fiber splits covering up to 1024 HHP. Yet since the system has minimal processing power that is characteristic of traditional PONs, it requires very minimal power. The ability to reach 512 HHP for example requires less than 75W, which is about half of the power consumed by a typical 4X4 Node in deployment today.

The significance and benefit of a move from a processing intensive PON OLT operation to a minimal processing transparent HPON architecture cannot be overemphasized. In case of the traditional PON architecture, each iteration of speed increase would have to result in a change in the OLT and the PON extender as well as the ONU at the house. And we have seen that letting myriad forms of PONs to be in existence could lead to efficiency reductions. However in the HPON architecture, the downstream is completely transparent to the modulation PHY formats and the upstream is completely OBI free and a well-designed system is also transparent to the upstream PHY formats. Similar is the case for the ONU unit in the house. Therefore a vast majority of the devices remain in the network unaffected as improvements in speed; capacity and throughput get worked out over the next several decades.

OPTICAL BEAT INTERFERENCE

Optical Beat Interference (OBI) is a profound issue in RFoG reverse path and affects the system in debilitating ways [10].

In a traditional HFC system such as the one shown above, since the CMTS scheduler allows multiple CPEs to burst at the same time at different RF frequencies, there are no collisions in the upstream. This is not the case when the HFC system gets replaced with a traditional all optical RFoG system.

In an RFoG system, upstream bursts from the CPEs connect to their own Optical Network Units (ONUs). The ONUs detect the burst and then turn on. When they thus operate, they are all nominally on the same optical wavelength. In reality each laser is at a slightly different wavelength due to manufacturing tolerances of the lasers.

If these wavelengths are appreciably different, or if near identical wavelengths burst at different times, there is minimal interaction along the upstream. However, when the near identical wavelength ONUs simultaneously burst they create strong interaction in the receiver, this is called Optical Beat Interference (OBI).

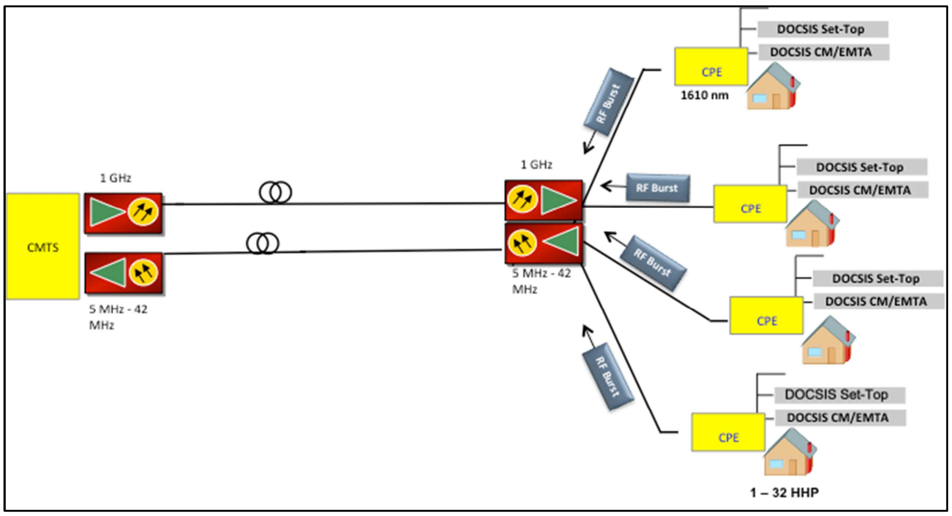

Figure 4: Illustrating a typical HFC Network

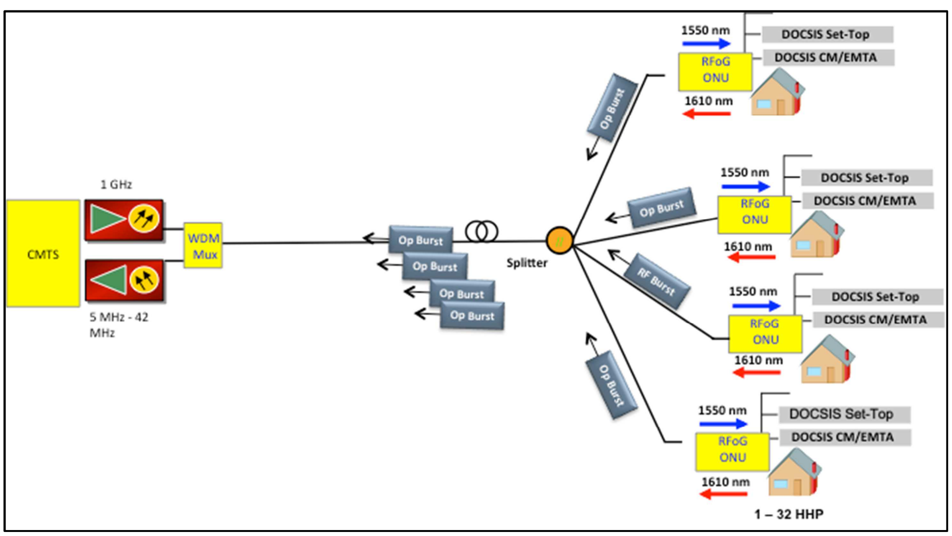

Figure 5: Illustrating a typical all fiber traditional RFoG network

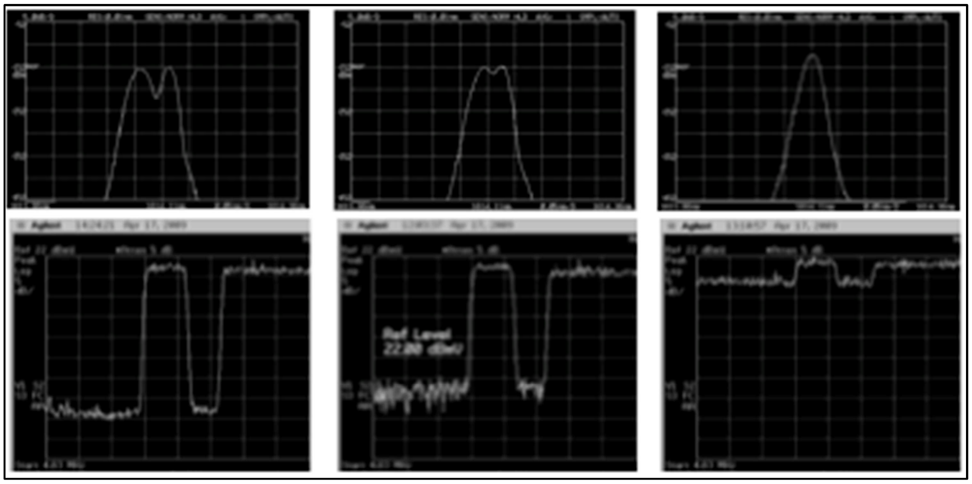

Figure 6: Illustrating OBI as a function of wavelength difference

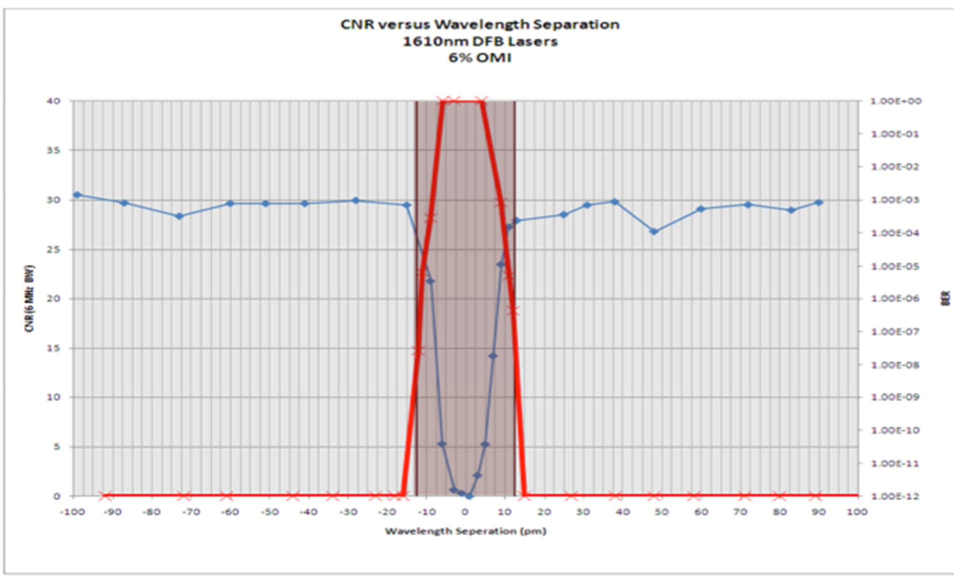

Figure 7: Illustrating SNR and BER under OBI conditions

OBI is a result of the heterodyning of the two (or more) closely spaced wavelengths present on the same detector. Heterodyning results in the down-mixing of the optical frequencies of the two or more lasers into the RF domain, appearing as wideband noise.

The resultant noise has the ability to impact the signal integrity of upstream communication channels, most notably on the receiver where the OBI has occurred, but can also affect the upstream DOCSIS service group via the introduction of noise into the RF combining network presented to the CMTS.

In the above shows a measured result of SNR and BER as a function of optical wavelength separation of two ONUs. When the ONUS wavelengths are far apart, the SNR and BER are good. However, when the wavelengths are within around 0.0125 nm of each other, there is a precipitous drop of SNR and a consequent rise of BER, corresponding to the picture shown above depicting the optical and RF spectra above.

It is important to realize that the 12 pm is around 10GHz of optical frequency, where typical 1610 nm lasers operate at around 200THz. Based on these test results it is seen that an OBI event wipes out the entire upstream SNR. In the case where there are more than 2 ONUs that can be simultaneously operational, such as when there are 4 DOCSIS 3.0 bonded channels, the net effect is that the an OBI event between any two ONUs wipes out the transmission of any other ONUs that were transmitting at the same time, even if they did not really contribute to the OBI. While it will be shown that this effect is already very severe in a 4 channel DOCSIS 3.0 case, it will be much worse when the number of bonded channels increase such as with the 85MHz D3.0 or with the 204MHz D3.1, where dozens of ONUs can simultaneously operate.

OBI probability depends upon the probability of two or more ONUs operating at the same time, and the probability of two or more ONUs having wavelengths within the OBI zone. The probability of two or more ONUs being on simultaneously is based on the Utilization of the upstream and is given by an application of the Binomial distribution, while the probability of the wavelengths being within the OBI zone is given by a ratio of the OBI zone relative to the entire region over which the optical wavelengths may be distributed. In modern laser designs, the wavelength range is typically less than 2.5 nm, with the OBI region be around 0.025 nm, the probability of OBI occurring is pretty high.

In reality however, the probability of the OBI occurring is much higher that the already high probability show above because of start up drift of a laser.

Since the ONUs only turn on when there is a burst of RF, the laser within the ONU is off until such time as it perceives a burst of the RF. Then, when the laser turns on, it starts absorbing the bias current applied to it and puts out optical power. All through the initial time of its turn on, until it reaches steady state, the temperature of the laser chip keeps increasing, with a consequent increase in the optical wavelength. All of this happens in a time scale of about a 100 to 200 us, and is fundamental to the operation of the laser. Since the time scale is so very short, it is not therefore possible to correct for the fast wavelength drift. The laser ultimately settles to a steady state value of about 0.5 nm higher than where it initially began operation.

The net result of this start up drift, which is a characteristic of the dynamic upstream system, is that the probability of OBI occurrence is increased by about two orders of magnitude, rendering the upstream system vulnerable to severe OBI for all meaningful upstream utilizations and a significant reduction in upstream throughput. What this means is that eliminating OBI needs to be treated as a strategy and cannot be solved by simplistic and opportunistic changes in optical wavelengths.

Eliminating OBI therefore requires the ONU wavelengths to be separated by at least 0.5 nm, preferably around 1 nm so that they never overlap. The only practical way of doing this is to DWDM like optics in the 1610nm region. Such a solution besides being expensive is also operationally impractical. This is because there would be fewer than 20 wavelengths possible in a typical SCTE specified 1610+/-10nm band. Therefore, there would have to one receiver per 20 ONUs and furthermore, the operator would have to map each household to one of the ONU wavelengths. The larger numbers of receivers thus needed would also increase the thermal noise in the system thereby drastically reducing the dynamic range of the upstream in addition to adding to the overall system cost. Finally, sparing costs would be prohibitive as would the system upgrades would be unintuitive and expensive.

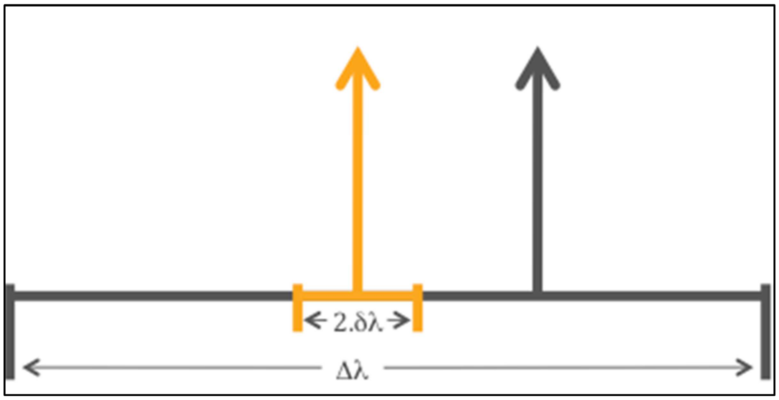

Figure 8: Evaluating the probability of OBI under static conditions

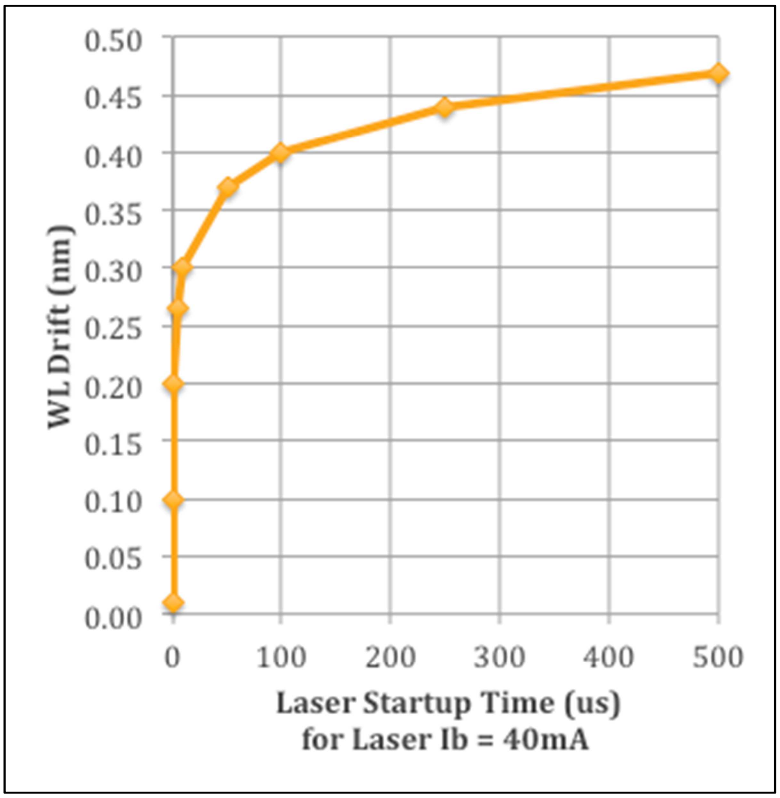

Figure 9: Illustrating the start up drift of an uncooled laser

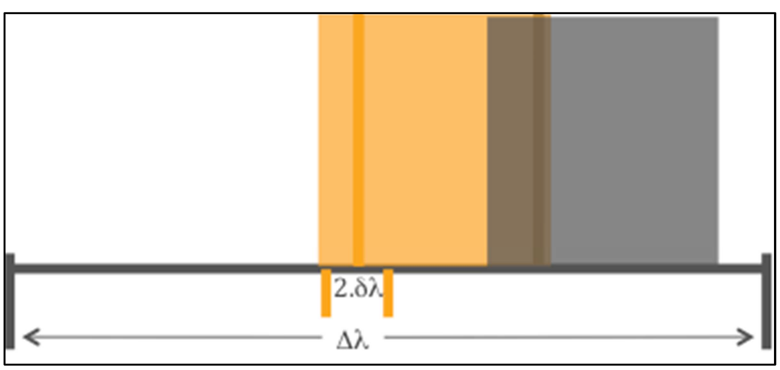

Figure 10: Evaluating the probability of OBI under burst mode conditions

THE ACCELERATED RFOG TEST

Since the dynamic nature of the wavelength drift is the dominant factor in eliminating OBI, it is easy to miss the exceedingly deleterious OBI effects if the test were done in a very simplistic static manner Therefore it is essential to evaluate OBI and potential elimination methods in the dynamic environment of a CMTS test bed.

The figure above represents a typical CMTS test bed. 32 CPEs feed 32 ONUs, which are combined in a splitter and go thru 20km of fiber. These are then connected to a receiver and a transmitter thru a WDM at the headend. A QAM256 channel is combined at the output of the receiver to act as a probe channel to be able to measure MER, BER and constellation diagram. Any impact on the probe channel is solely due to the OBI effects (recall that when OBI occurs, entire upstream is splattered by the OBI induced noise). The forward transmitter is then connected to the other end of the WDM and receives its HSD signals from the CMTS. The CMTS also receives the burst signals from receiver thus completing the loop. Both the CPEs and the CMTS are run thru a traffic generator, in this case thru a ByteBlower. An optical spectrum analyzer is connected thru an optical splitter as are RF spectrum analyzers to measure the upstream RF spectrum and the probe MER, BER and constellation. One can monitor the Code Word Error Rate (CER) before and after FEC on the CMTS as well as the frame loss on the ByteBlower as a measure of system performance.

When the CMTS and the ByteBlower are activated, the CPEs and the ONUs burst in the upstream creating OBI, which impacts the optical, RF and the CMTS domains.

Presented in Fig. 12 is a representative optical spectra of 8 ONUs with varying wavelength separation. After the CMTS is activated, the OSA is put on a max hold so we could clearly see the effects of startup drift.

Presented in Fig 13 is a representative optical spectra with Max hold on an OSA after just 2 minutes of the CMTS being activated. The smear of the optical wavelengths indicates the effects of the startup drift and the reason and the cause of exceedingly high amounts of Obi observed in real RFoG upstream systems. It is for this reason that opportunistic wavelength movement schemes of Obi mitigation do not suppress Obi effectively. And we have already discussed that meaningful wavelength separation is expensive, cumbersome and difficult.

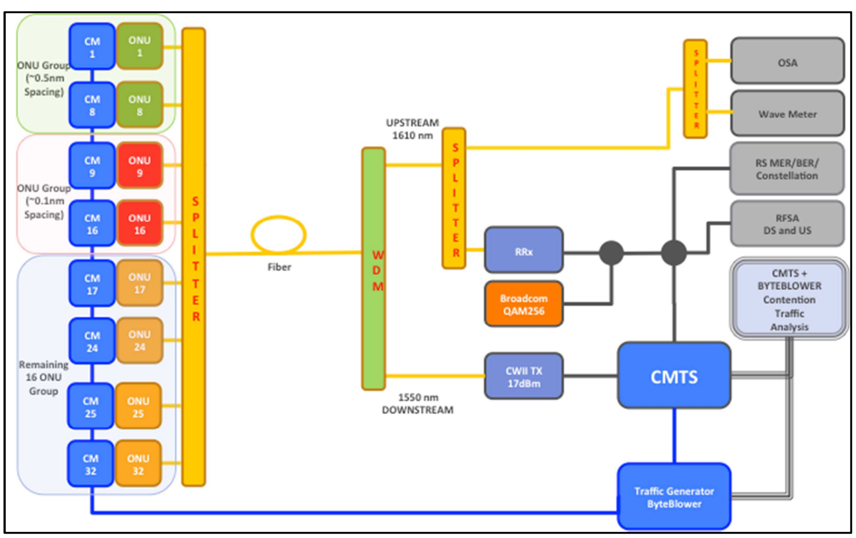

Figure 11: Comprehensive test bed to evaluate the effects of OBI

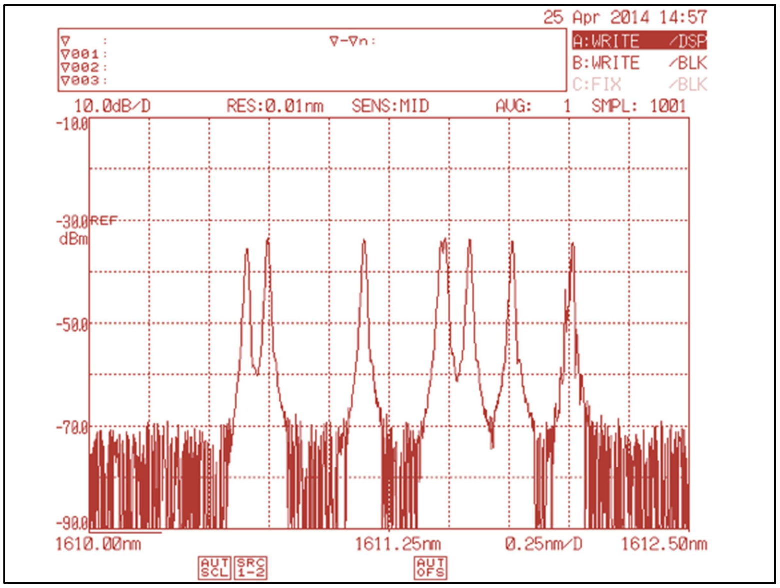

Figure 12: Measured optical spectrum under static conditions for traditional RFoG

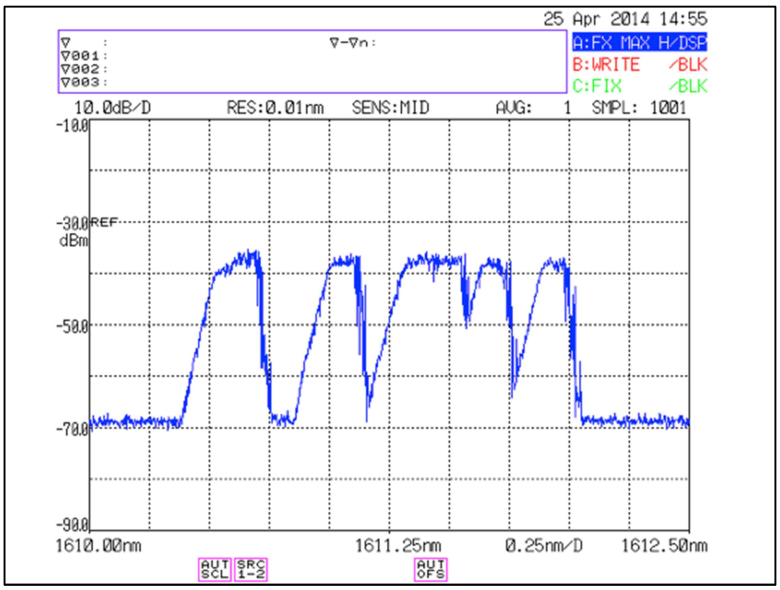

Figure 13: Measured optical spectrum under burst mode operation on Max Hold for traditional RFoG

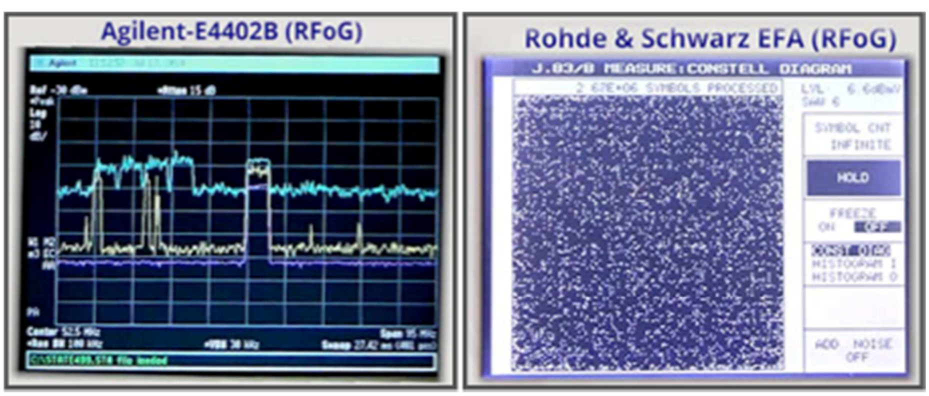

Figure 14: Measured Upstream RF Spectrum & constellation of probe channel with OBI in traditional RFoG

Presented above is the RF spectrum of upstream and a constellation diagram of the probe channel of a typical 32 ONU test bed described earlier for a utilization of 50% after a 2-minute test. It is seen that the spectrum is dominated by the OBI noise and the constellation is completely dark.

With such poor spectrum, the frame loss at the ByteBlower exceeds 20% and both the UDP and the TCP traffic at the upstream are severely affected. Since the TCP traffic at the upstream is severely affected, the TCP traffic of the downstream is affected as well.

The practical customer impact of such severe degradation is a significant impact on an Ookla speed test and a loss of significant voice quality on a VoIP connection. Over the Top services such as Netflix are affected as well due to accumulating TCP acknowledgements that are unable to make their way back to the Netflix servers in a timely manner.

Tests show that even modest amounts of utilization produce significant OBI and these contrive to affect a precipitous drop in a customer measured Ookla speed test as show above.

ELIMINATING OBI

We have previously seen that eliminating OBI at the ONUs is not possible with simplistic and opportunistic wavelength shifting methods. Eliminating OBI by using wavelength selective ONUs is expensive, cumbersome and impractical. An analysis of eliminating OBI by implementing a TDM CMTS scheduler is shown in the next section. Besides severely limiting throughputs this approach is no guarantee against OBI due to non-DOSCIS channels and due to malicious users. It is therefore essential to find a new way to eliminate OBI that is ONU and CMTS independent and immune to malicious users.

A complete elimination of OBI unlocks the full potential of the optical fiber and enables up to 10Gbps of upstream capacity and up to 20Gbps or more of downstream capacity with appropriate optical amplification. An active splitter configuration enables such OBI elimination with a very modest power requirement of about 5W.

Furthermore, the active splitter overcomes the upstream and downstream splitter losses thus enabling long distances and high splits as well as enabling the aforementioned high capacities. The large throughputs can also allow larger service groups which can help eliminate substantial upstream and downstream RF combining and splitting in the headends thus saving valuable headend space and reduce power consumption of headends.

With the same test configuration described earlier, but with the passive splitter replaced by the active splitter, the CMTS test described earlier is repeated to verify OBI elimination.

It is seen in Fig 17 that the RF spectrum is clean and the probe constellation is clear, thus establishing the OBI elimination even in the burst mode CMTS test bed.

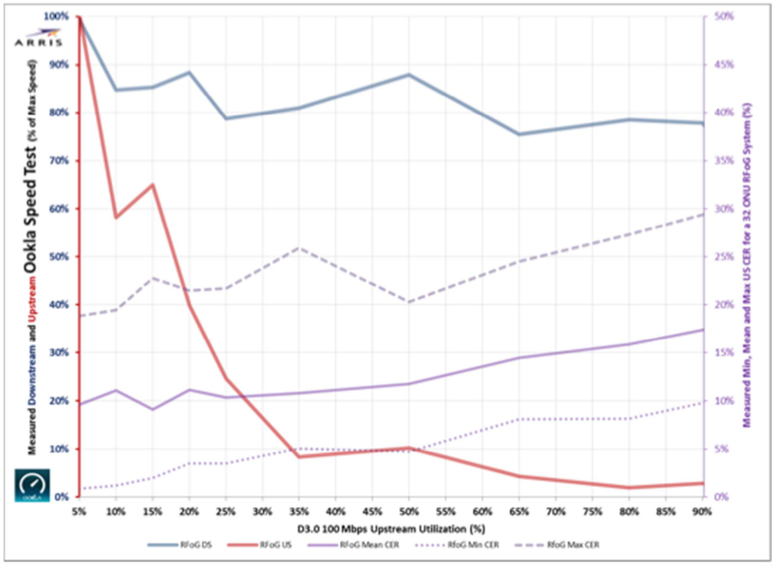

Figure 15: Measured speed test results under OBI for traditional RFoG

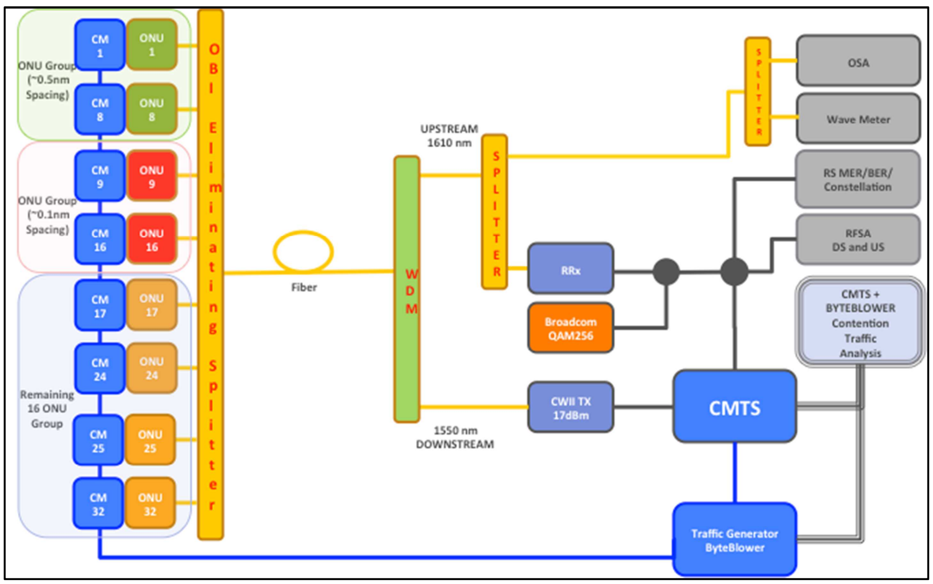

Figure 16: Comprehensive test bed to evaluate the effects of OBI with the OBI eliminating active splitter

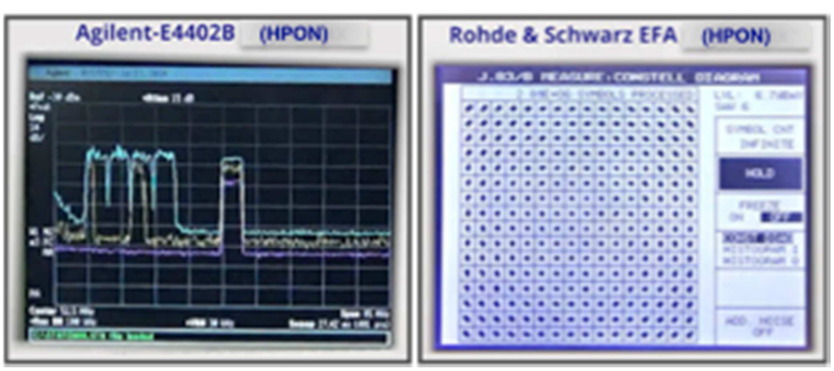

Figure 17: Measured RF Spectrum of the Upstream and the constellation of the probe channel under OBI conditions on an active splitter with OBI eliminated

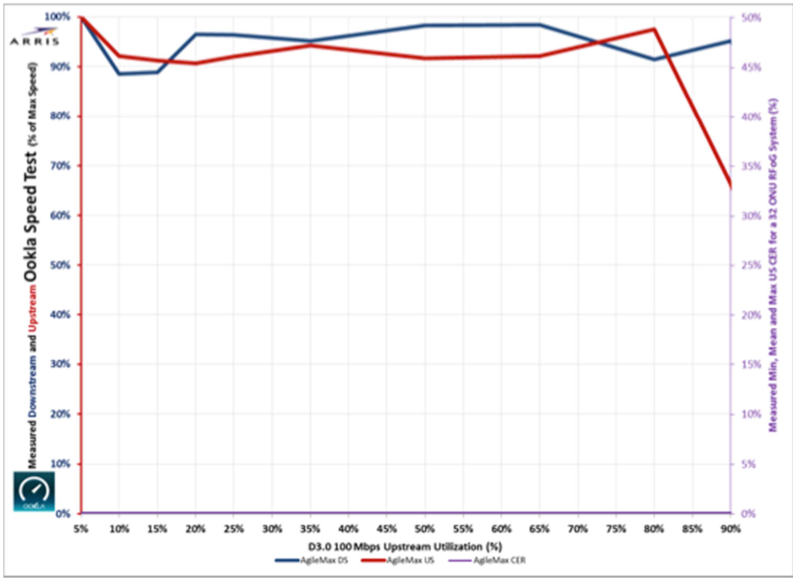

Figure 18: Measured speed test results under OBI for with the OBI eliminating active splitter

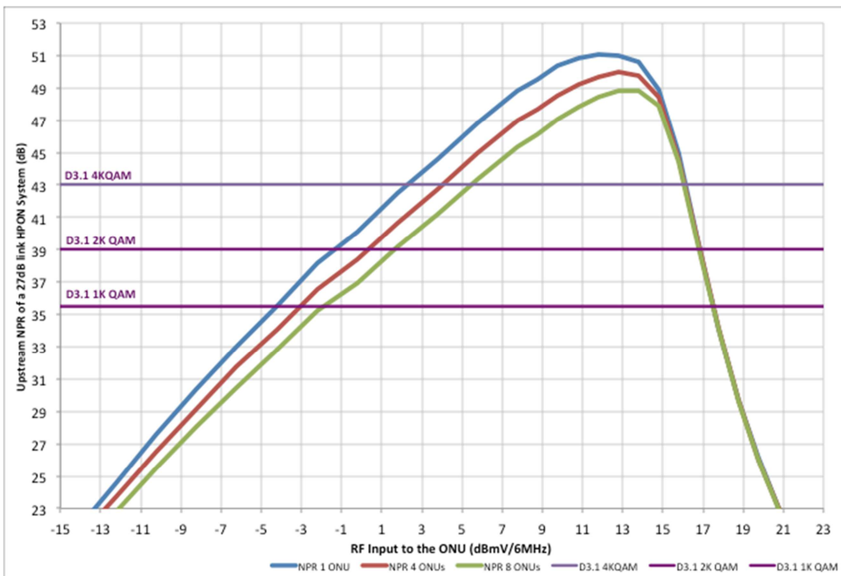

Figure 19: Measured NPR of a 25km, 27dB Link Budget 204 MHz HPON upstream system

Repeating the Ookla speed test confirms that the upstream and downstream capacity is unaffected. The Vonage connections on the VoIP system are also clear.

To further establish the complete elimination of OBI and the multiple benefits of the active splitter concept, NPR tests at 85MHz were done with progressively higher numbers of ONUs always on a 30dB RFoG upstream reverse path.

As can be seen from the measured data above, the NPR of the active splitter RFoG system is excellent with minimal noise side degradation with multiple ONUs on simultaneously [11]. This important test illustrates the fact that complex DOCSIS 3.1 modulation formats may be used for upstream transmission and that even if dozens of ONUs simultaneously operate as is expected in the DOCSIS 3.1 environment, the transmission rate is unaffected. Similar testing in the downstream with present day ONUs and lower cost Direct modulated transmitters over similar distances yield better than 41dB MER, sufficient for 1K QAM operation to 1.2GHz.

Since HPON upstream is in the 1610 nm band and the downstream is in the 1551 nm band, the Active Splitter allows a completely transparent path for all the other remaining wavelengths to pass through, thus enabling true RFoG and traditional GPON/EPON coexistence. An optimized multi wavelength plan enables high capacity upstream and downstream wavelengths to be used thus conserving trunk fiber. With this in place, one can deploy high performance, high capacity OBI free RFoG along with GPON and EPON for a truly Hybrid PON architecture [12,13].

OBI MITIGATION ATTEMPTS IN THE CMTS

We have seen in previous sections that eliminating OBI at the splitter location is effective and efficient. This gives the operator ultimate freedom to choose any best of breed ONU and CMTS from any vendor. This approach is vital to a viable and successful deployment of the HPON architecture.

As we have shown earlier, other PHY layer OBI mitigation attempts at the ONU are proprietary, cumbersome, impractical, expensive and do not scale with larger Service Group sizes.

Single TX CMTS Scheduler OBI Mitigation

Perhaps the most common method of mitigating OBI recently has been by using a specialized CMTS scheduler that only allows one DOCSIS cable modem to transmit at a time. This is conceptually simple and has had some success in the very early days of RFoG when most of the operators used DOCSIS 2.0.

While this approach eliminates the possibility of OBI from DOCSIS modems it does not address the possibility that OBI might occur when another service such as Legacy Settop Box (STB) Out-of-band (OOB) signals transmit at the same time as the DOCSIS modem. This is completely out of control of the CMTS scheduler.

It turns out that STB OOB traffic is relatively light so it would have minor impact on the DOCSIS traffic and show up as a slightly increased error rate. However, as RFoG becomes more successful and the DOCSIS utilization increases there is now a risk that the STB OOB traffic may be starved if at least one of the DOCSIS upstreams is being used every time the STB OOB tries to transmit. While this problem may not have been noticed yet, it is something that the operators need to be cognizant about.

In a DOCSIS 3.0 environment, multiple cable modems may transmit simultaneously. Therefore the opportunity of OBI is considerably more and approaches near certainty for increasing upstream channels, increasing modems and increasing utilization. Continuing this approach to mitigate OBI by the CMTS Scheduler in DOCSIS 3.0 environment allows only one cable modem may transmit at any one time. In addition to being a specialized software that is unique to RFoG, an operator who has been reliably running a CMTS Scheduler for years across millions of HFC subscribers must now use different software, even if it is from the same vendor. And as we will show, this approach is exceedingly inefficient and has very limited throughput, defeating the purpose of its development.

DOCSIS 3.0 CMTS Scheduler Efficiencies

The Active Splitter based OBI solution allows operators to use the identical CMTS Scheduler that they are using across millions of other subscribers. These schedulers have been optimized to deliver maximum performance across multiple channels with guaranteed Quality of Service (QoS).

Meanwhile, a Single TX OBI CMTS scheduler (a.k.a. TDM CMTS scheduler or the

“RFoG aware Scheduler) has made a conscious decision to limit transmission across all DOCSIS channels to a single modem. This means that the scheduler is starting to trade-off between total capacity and QoS. In the early RFoG days with very few DOCSIS upstream channels and limited utilization, this was an effective trade-off. However, as shown below, as the DOCSIS utilization increases, and the number of bonded RF channels increase from 4 to 8 and above this method becomes exceedingly inefficient and self-limiting.

Our analysis first looks at scheduler behavior for DOCSIS 3.0 systems. We assume that any DOCSIS 2.0 modems have been removed as they can significantly degrade performance of a Single TX OBI scheduler approach with multiple upstream channels. An example of a Single TX scheduler in a D3.0 system with four US channels would take a 1000 byte packet and fragment this into four simultaneous transmissions from this modem of 250 bytes plus DOCSIS overheads.

Obviously, sending nothing but extremely large packets would result in decent efficiencies. However, this would not be indicative of real world behavior, especially if it was a large Service Group with hundreds of modems. Research has shown that the upstream traffic is very bi-modal, with pre-dominantly small packets around 64 bytes and large packets near the Ethernet maximum of 1522 bytes. Some analysis of live broadband systems shows that a mix of 70% small packets and 30% large packets is reasonable.

Note that with this mix the average packet size is about 500 bytes, and that 91% of the data is actually carried by the large packets, even though they are only 30% of the total packets. Our results for a DOCSIS 3.0 system are shown in Fig 20. The number of upstream channels are varied from one to eight channels with each channel being 6.4MHz and 64-QAM.

Looking at the results in the Figure, the first observation is that both schedulers have the same efficiency for a single channel. This makes sense since nothing special needs to be done to avoid OBI. Also, the efficiency of ~88% is the result of the small packets being mixed in. Had there only been large packets, then the efficiency would have been in the 98% or 99% range.

As the number of DOCSIS 3.0 upstreams are increased, a standard CMTS scheduler with a splitter-based PHY Layer solution can schedule freely across all channels and won’t see any efficiency degradation. It remains constant no matter how many upstream channels are deployed.

A Single TX CMTS Scheduler becomes more inefficient as channel count increases. If you think about a 64-byte packet fragmented across an eight-channel system, there is only 8 bytes of payload in each channel and the DOCSIS PHY burst Overhead starts to dominate.

At two upstream channels, the difference between the schedulers is a relatively small 13%. As the system capacity increases to four upstream channels, the Standard CMTS scheduler now provides ~40% more capacity than the Single TX OBI scheduler. Finally, at eight upstream 3.0 channels, the Standard CMTS Scheduler is providing twice the capacity of the Single TX OBI Scheduler. This would map to ~200 Mbps for a Standard CMTS Scheduler vs. ~100 Mbps for the Single TX CMTS Scheduler.

In addition to these DOCSIS data rate efficiencies, it should also be noted that a Single TX CMTS Scheduler would not be able to use any contention based Request (REQ) mechanisms either. This means that for large SG it may not scale as well and have increased latencies compared to a Standard CMTS Scheduler.

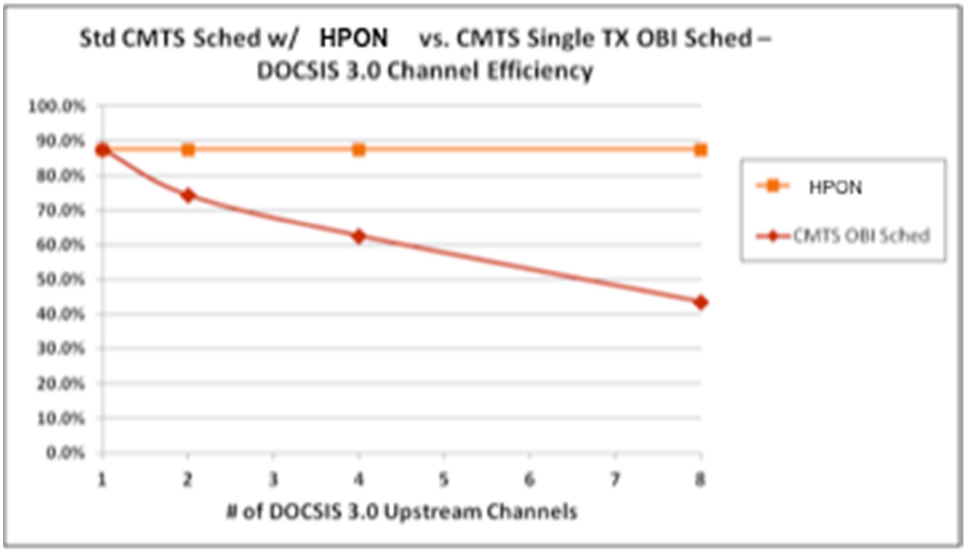

Figure 20: CMTS Scheduler Efficiencies – Std Sched over HPON, Single TX OBI Sched

DOCSIS 3.1 CMTS Scheduler Efficiencies

While the DOCSIS 3.0 results are very intriguing, perhaps the most important factor in scaling is the ability to deploy DOCSIS 3.1 capabilities over HPON architecture. D3.1 brings many potential technology improvements, but perhaps the most significant is the wide channels enabled by OFDMA. The first D3.1 products can potentially support up to a pair of 96 MHz OFDMA channels on a 204 MHz upstream.

In addition to the wider channels, the HPON-based D3.1 systems will have much improved SNR, which also enables 4096-QAM modulation in the upstream.

From an OBI perspective, the D3.1 OFDMA channels represent a serious challenge, as there may be dozens of different modems transmitting simultaneously on different subcarriers. To provide a handle on the potential impact, it is valuable to understand some basics of the D3.1 OFDMA channel.

The DOCSIS CMTS Scheduler is completely based on Mini-slots, for both D3.0 and D3.1. All transmit assignments, a.k.a. Grants, are a multiple number of Mini-slots. For a D3.0 ATDMA channel, Mini-slots are defined as a unit of time (e.g. 32 byte times) and is thus one-dimensional. For a D3.1 OFDMA channel, a Mini-slot is two-dimensional in both time and frequency.

The D3.1 Mini-slot is defined to be 400KHz wide in frequency. This maps to 8 subcarriers if using 2K FFT or 16 subcarriers if using 4K FFT. In the time dimension, the Mini-slot can vary from 120 to 360 microseconds. This maps to 6 to 18 symbol times for 2K FFT or 6 to 9 symbol times for 4K FFT. The symbol depths were chosen to allow interleaving in the upstream.

In D3.1, a Frame Time is a collection of Mini-slots that are all synchronized and transmitted at the same instant of time. The total number of Mini-slots per Frame Time is now a function of the total upstream spectrum available to the OFDMA channels. This is shown in Table 5.

The number of payload bytes per Frame Time can vary significantly. If the objective is to minimize the total bytes per Frame Time to reduce the possibility of OBI, then that leads to the selection of 2K FFT with only six symbol times per Mini-slot. The last row of Table 5 shows how many payload bytes are available in a Frame Time for this configuration and various spectrums. Note that this is best case for OBI and could increase up to a factor of six larger than this.

The number of payload bytes per Mini-slot can vary as well, and might fall in the range of 50 to 400 bytes. For the scenario above (i.e. 2K FFT, 6 symbols per Mini-slot), a small 64 byte packet will consume 2 Mini-slots and a large 1522 byte Ethernet packet might be ~30 Mini-slots. As can be seen from Table 5, a D3.1 system can potentially have dozens of simultaneous transmitters.

To eliminate OBI in D3.1 with a Single TX CMTS Scheduler becomes extremely problematic. If a modem has a single 64-byte packet to transmit, they must consume the entire Frame Time in order to eliminate OBI. Even transmitting large 1522 byte packets, or concatenations with 4500 bytes or 9000 bytes is small compared to the total bytes available in a Frame Time, especially the 33KB Frame Time for 204MHz @ 4096-QAM. The CMTS Scheduler efficiency results with the same packet size mix of 70% small, 30% large packets are shown in Figure 21.

D3.1 Mini-slot Efficiency	42 MHz	65 MHz	85 MHz	204 MHz	204 MHz Enhanced
Total Mini-slots per Frame Time	~82	~140	~190	474	474
Modulation	1024-QAM	1024-QAM	1024-QAM	1024-QAM	4096-QAM
Bytes per Frame Time	4483	7862	10,774	27,608	33,178

Table 5: DOCSIS 3.1 Upstream OFDMA Frame Parameters

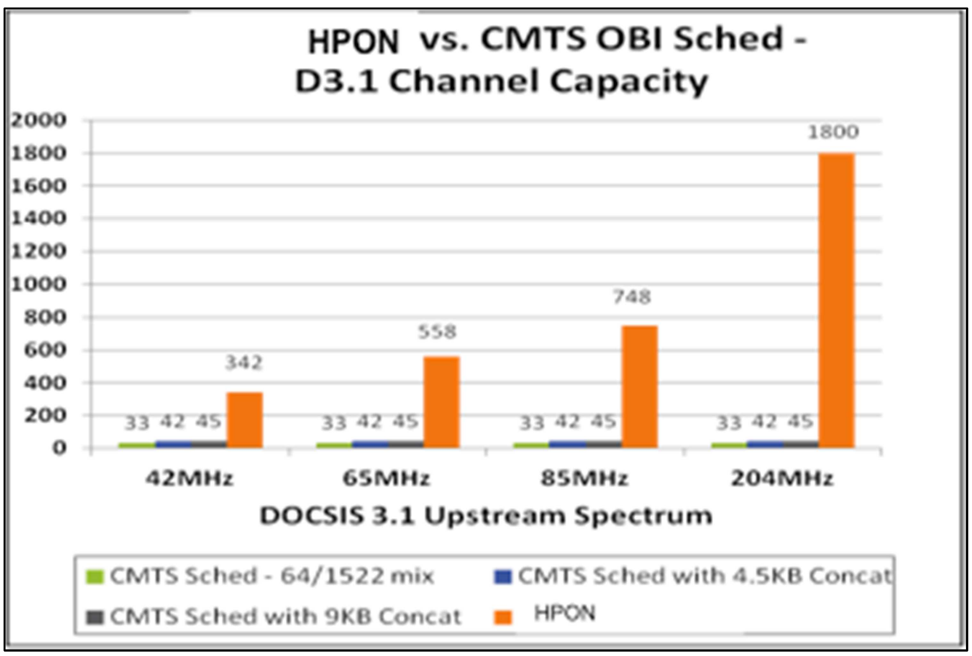

Figure 21: D3.1 Scheduler Efficiencies – Stud Sched over active splitter, Single TX OBI Sched

As can be seen in Figure 21, the Single TX CMTS Scheduler cannot even come close to filling a 42MHz OFDMA channel. In fact, it gets higher capacity using D3.0 ATMDA whose Mini-slots are shorter in time.

Meanwhile, the Standard CMTS Scheduler operating over HPON cannot only achieve full D3.1 OFMA utilization, it can also take advantage of the improved upstream SNR to operate at 4096-QAM. With HPON enabling overlapping upstream and downstream spectrum, then 204MHz operation becomes a reality and HPON-based D3.1 can deliver a 1.8Gbps upstream data rate (almost 2.3Gbps PHY rate). This is more than enough to offer a true 1Gbps upstream service tier to cable customers.

FUTURE OF BINARY AND ANALOG OPTICS

At 1550 nm RFoG/HPON is more suitable for SMF-28 fiber access networks than binary modulation. This surprising result is due to the twin effects of the superior capability of QAM schemes to take full advantage of available bandwidth and convert it to higher throughput relative to the binary formats and the effects of fiber dispersion on links where fundamental limits can cap NRZ binary modulation formats to around 20 Gbps throughput at 25 km and fewer at greater distances. This paper indicates that HPON when used with the DOCSIS 3.1 standards can support greater than 40 Gbps and distances up to 40 km on a single wavelength.

Attainable SNR of the optical link is a function of frequency and distance [11, 16]. 40 Gbps throughput can be obtained with around 4 GHz of bandwidth.

For a binary system the throughput was simply given as:

$$\text{bin\_bps}(L) := 10 \cdot \sqrt{\frac{120}{L}}$$

This is plotted as markers for the 3 fiber lengths at their respective Nyquist bandwidths. Whereas the binary modulation permits operation at higher bandwidth (compared to a dual sideband AM modulated signal around a carrier) it does not provide more throughput.

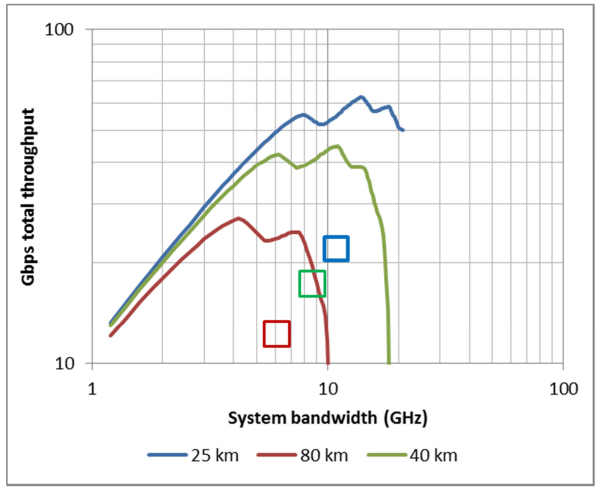

Figure 22: Illustrating the Capacity of an RFoG link with OFDM/LDPC vs. Binary Modulation

The addition of an EDFA in the headend, as needed for RFoG, increases the available downstream bandwidth in the system from 10 Gbps (PON) to more than 40 Gbps (RFoG). As shown in the figure even more throughput can be available with wide bandwidth receivers. However this paper is limited to ONUs that can be realized with currently available low cost components and the discussion is therefore be limited to around 5 GHz receiver bandwidth. It is further concluded that due to dispersion a binary transmission format cannot reach the same throughput as an RFoG system with EDFAs, neither at 25 km reach nor at 40 km reach.

RFoG Receiver performance

Estimated capacity is shown below as a function of receiver power with 4 GHz receive bandwidth at 25, 40 and 80 km. At -6 dBm receive power almost 40 Gbps is achieved, the performance saturates above –3 dBm to 40 Gbps.

These estimates are for a multi-carrier RF modulated signal (OFDM with LDPC) with a modulation index set to the same value as for current QAM256 operation. With peak-average power reduction (PAPR) the throughput can be increased by around 10% beyond these results. Margins needed for current DOCSIS 3.1 impementations can reduce throughput by around 10%.

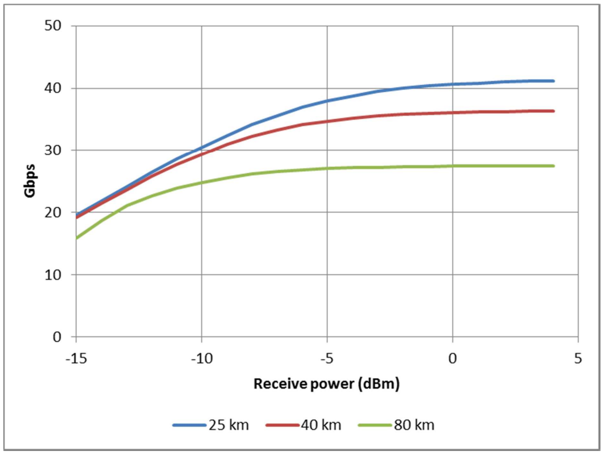

Figure 23: Illustrating System Throughput in Gbps as a function of receiver input power

RF amplification of the received signal to 4 GHz or higher is readily available with low-cost commercial MMICs. Thus the receive side of the ONU can be realized with a subset of the components used in a 10 Gbps PON ONU. The required components can satisfy reduced bandwidth and gain requirements. The transmit side of the ONU that will be proposed here can be based on a regular directly modulated DFB laser with 1 GHz of bandwidth. Given the availability of wideband lasers and amplifiers that is not significantly more complex than RFoG ONUs with less upstream bandwidth.

Modulation Formats and Throughput

In the access part of the plant the fiber loss budgets are relatively small compared to long haul telecommunication. In RFoG the number of forward wavelengths is also small compared to long haul communication systems. By using optical amplification in the access plant, permitted for one or a few wavelengths on a fiber, the SNR can be increased so much that analog modulation formats, in spite of their shortcomings, readily obtain 4x the throughput of binary modulation formats given practical receiver bandwidths available today using low cost components. Binary modulation formats with optical amplification could permit larger service groups, however due to traffic engineering, the size of service groups needs to be limited unless much more than 10 Gbps forward capacity can be offered, that is not available today in low cost binary data processing receiver and electronics and its application is limited by fiber dispersion. Therefore the binary modulation formats do not benefit from optical amplification in the access. It can thus be concluded that, unlike in long-haul telecommunication, analog modulation formats are superior to binary modulation formats in the access plant. This should come as no surprise, in most short distance bandwidth limited systems analog modulation formats are used (DSL, CAT-6 10 GbE cables, WiFi and of course DOCSIS).

While it may sound contradictory to classify a fiber system as a bandwidth limited system it is fiber dispersion and the OE conversion (receiver) and associated processing electronics that do pose a practical bandwidth limitation in fiber systems. In fact this bandwidth limitation is acknowledged by 40G PON manufacturers resorting to 4 wavelengths each operating at 10 Gbps rather than attempting to run one wavelength at 40 Gbps. Reality is that aggregate throughput levels are expected to increase to a level that far exceeds the baseband bandwidth of low cost receiver electronics that are used with simple on-off keying on standard SMF fiber. Current low-cost analog technology on the other hand can already provide a transparent pipe with the required throughput capability today.

FUTURE UPSTREAM CAPACITY

In the upstream SNR is evaluated accounting for laser RIN, laser power and loss budget, the upstream bandwidth and the modulation index that can be allowed per channel given a total bandwidth, an additional dynamic window to accommodate uncertainty in transmitter OMI setting and the optional use of PAPR (Peak to Average Power Reduction) in the transmission format.

The effective modulation index just under clipping induced BER is 20% for regular RF modulated systems and around 32% for RF modulated systems with PAPR (Peak Average Power Reduction) methods applied.

The return bandwidth can vary, up to 1 GHz should not significantly affect return laser and driver cost, this primarily affects the choice of diplex filters. The attainable throuput is analyzed [11] and provides up to 10 Gbps upstream capacity. In a 204 MHz return system more than 1.5 Gbps of upstream capacity can be attained. It should be noted that in an RFoG system the forward and reverse bandwidths could be overlapping such that 1 Gbps symmetrical operation can be offered in a system with an 85 MHz split in the in-home network.

CONCLUSIONS

We began this paper with a premise that a move to all fiber need not necessarily be a move to one of the traditional binary PON architectures. Since a move to all fiber is a long duration process spanning several decades maintaining disparate end-to-end systems could be cumbersome and expensive.

In this paper, we presented an innovative Hybrid PON architecture that enables a true coexistence of traditional HFC and traditional PONs (such as GPON and EPON) along with an OBI Free high capacity HPON system that is capable of supporting analog, broadcast video, D3.0 and D3.1 formats.

We also discussed the true nature of OBI and its debilitating impact on tests that affect customer perception - such as the Ookla Speed Test. We also discussed the various significant inadequacies of some approaches that seek to just mitigate OBI rather than fundamentally eliminate it and the glaring inefficiencies of others that could eliminate OBI but at a significant loss of throughput that defeats the purpose of their deployment.

A deployment of the OBI Free HPON system as described in this paper maintains all options to deploy traditional binary PON architectures for business customers and very high data consumers, while enabling a well-ordered migration to FTTH by supporting traditional HFC as well. Thus HPON satisfies the anticipated growth in traffic demand while providing MSOs with plant and equipment investment protection well into the future.

REFERENCES

1. Predictions on the Evolution of Access Networks to the Year 2030 & Beyond; T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, S. Chari, The Cable Show NCTA/SCTE Technical Sessions Spring 2014
2. Neilsen's Law vs. Nielsen TV Viewership for Network Capacity Planning, Mike Emmendorfer, NCTA 2014
3. Comcast hybrid fiber and coax plant: http://business.comcast.com/docs/default-source/brochures/the_comcast_network.pdf?sfvrsn=0
4. Time Warner Cable fiber plant: http://www.carrierethernetnews.com/articles/828971/time-warner-cable-business-class-enhances-ethernet/
5. Verizon FiOS FTTP deployment: en.wikipedia.org/wiki/Verizon_FiOS
6. www.bloomberg.com/news/articles/2015-02-05/verizon-to-sell-landlines-towers-in-two-deals-for-15-6-billion
7. Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning; John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, J.P. Fioroni; SCTE Cable-Tec Expo 2014
8. 10 G-EPON: Drivers, Challenges, and Solutions; Glen Kramer, ECOC 2009
9. Performance of 10 G EPON; Rajesh Roy, Glen Kramer, Marek Hajduczenia, Henrique J. Silva, Topics in Optical Communications, IEEE Communications, November 2011
10. Beat It! Handling OBI In RFoG Systems; Ayham Al-Banna, Larry Spaete, Amarildo Vieira, Sebnem ZorluOzer, Marcel Schemmann, Tom Cloonan, Dean Stoneback, Venk Mutalik, NCTA 2014
11. Partnership for Extended Capacity: DOCSIS 3.1 with RFoG; Venk Mutalik, Marcel Schemmann, Ayham Al-Banna, Zoran Maricevic, NCTA 2014
12. When Wavelength Collide Chaos Ensues: Engineering Stable and Robust Full Spectrum Multi-Wavelength Networks; Venk Mutalik, Fernando Villaruel, Marcel Schemmann, Kwang-Yi Wu, SCTE 2011
13. Value of Transparency in the Evolution of the HFC Networks; Venk Mutalik, Marcel Schemmann, Brent Arnold, SCTE 2013
14. Coax Core Ejection technology: http://deepfibersolutions.com/applications/
15. Ayham Al-Banna et. al., "Measurement-Based EoL Stochastic Analysis and DOCSIS 3.1 Spectral Gain," in Proceedings, The NCTA Cable Show Spring Technical Forum, May, 2015.
16. Mike Emmendorfer et. al, "Estimating Downstream Performance and DOCSIS 3.1 Capacity in CAA and DAA Systems," in Proceedings, The NCTA Cable Show Spring Technical Forum, May, 2015.
17. Tom Cloonan et. al., "Lessons from Telco and Wireless Providers: Extending the Life of the HFC Plant with New Technologies," in Proceedings, The NCTA Cable Show Spring Technical Forum, May, 2015