SCTE CABLE-TEC EXPO'15
OCTOBER 13-16 NEW ORLEANS

SCTE Society of Cable Telecommunications Engineers

HFC Transformation to FTTP: The Role of RFoG, PON and Hybrid solutions

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

John Ulm
CTO- Network Solutions
ARRIS Group, Inc.
22 Topside Rd. Unit 1, Moultonborough NH 03254
(978) 609-6028
john.ulm@arris.com

Venk Mutalik, ARRIS Group, Inc.
Engineering Fellow and
VP of Technology and Strategy
ARRIS Group Inc.
15 Sterling Drive, Wallingford, CT 06492
+1(203) 494-6556
Venk.Mutalik@arris.com

Table of Contents

Title **Page Number**

Introduction 5
HFC to FTTP Capacity Planning 6
1.1. How Long before HFC Capacity runs out? 6
1.2. Traffic Engineering Fundamentals 6
1.3. Selective Subscriber Migration Strategy 9
1.3.1. Network Capacity Modeling of FTTP Migration 11
1.3.2. Economic Impacts of Selective Subscriber Migration 13
1.3.1. Selective Subscriber Migration Summary 14
DOCSIS® 3.1 Overview – Extending the Life of HFC for Decades 15
2.1. DOCSIS 3.1 Overview and Benefits 15
2.2. DOCSIS 3.1 Capacity Examples 15
2.3. DOCSIS 3.0 to 3.1 Migration Example 16
Hybrid PON (HPON) – A Revolutionary FTTP Breakthrough 17
3.1. HPON Overview 17
3.2. For Fiber To The Home (FTTP) Transition: DOCSIS or EPON/GPON? 17
3.3. HPON and the Role of DOCSIS 18
3.4. Mixed HFC and HPON DOCSIS 3.1 Operation 18
3.5. Stacking Up: D3.1 over HPON Capacities to Other PON Architectures 19
3.6. HPON Topology Options 21
3.6.1. HPON Topology Migration Example 22
3.6.2. HPON Topology Migration Example – Remote Devices 25
3.7. DOCSIS 3.1 and RF Performance over HPON 28
3.8. D3.1 over HPON Summary 31
10G & 1G EPON, GPON on HPON: Scaling Considerations for Residential Use Case 32
4.1. EPON Burst US Structure 32
4.2. Extending 10G EPON Capacity Analysis 33
4.3. 1G EPON, 10G EPON and GPON Efficiency 34
4.1. EPON: 1G and 10G Coexistence – Control Overhead Impacts on Efficiency 34
4.2. 10/10 & 10/1 Coexistence compared to GPON and D3.1/HPON 36
4.3. Residential Applications present Traffic Engineering Challenges 37
The role of FTTP and Hybrid PON – Other Considerations 39
5.1. Economics of HPON 39
5.1. Fiber Trunks, Wavelengths and OLT/CCAP Ports 40
5.2. HPON and Energy Considerations: Energy 2020 – Annual kiloWatt-hour/HP 41
HFC to FTTP Transformation: The Role of HPON – Conclusion 43
Abbreviations 44
Bibliography & References 46

List of Figures

Title	Page Number
FIGURE 1- CLOONAN'S CURVE AND NIELSEN'S LAW	7
FIGURE 2- NQOE FORMULA OVERVIEW	7
FIGURE 3 – TRAFFIC ENGINEERING WITH DOCSIS 3.1	8
FIGURE 4 – INDIVIDUAL SERVICE TIER GROWTH YEAR OVER YEAR,% SUBSCRIBERS AND TMAX FOR THE VARIOUS TIERS	10
FIGURE 5 – CAPACITY NEEDS OVER TIME, 128 SUBS PER SG	12
FIGURE 6 – CAPACITY NEEDS OVER TIME, VARIOUS SG SIZES	12
FIGURE 7 – EXAMPLE 1 GHZ/85 MHZ UPGRADE COSTS VS. FTTP	13
FIGURE 8- HFC UPSTREAM AND DOWNSTREAM CAPACITY	16
FIGURE 9 – HFC & HPON SPECTRUM OVERLAY	19
FIGURE 10 – DOWNSTREAM CAPACITY: HFC & HPON VS. XPON	20
FIGURE 11 – FIBER TO THE TAP EXAMPLE	22
FIGURE 12 – HPON FTTP MIGRATION: BASELINE	23
FIGURE 13 – HPON FTTP MIGRATION: FTTP FOR SELECT TOP CUSTOMERS	23
FIGURE 14 – HPON FTTP MIGRATION: FTTP FOR MORE PERFORMANCE CUSTOMERS	24
FIGURE 15 – HPON FTTP MIGRATION: SG SEGMENTATION AS NEEDED	25
FIGURE 16 – DISTRIBUTED ARCHITECTURE: CONVENTIONAL WISDOM	26
FIGURE 17 – HPON AND DISTRIBUTED ARCHITECTURE: SHARED REMOTE DEVICE	27
FIGURE 18 – HPON AND DISTRIBUTED ARCHITECTURE: SG SEGMENTATION AS NEEDED	28
FIGURE 19 – HPON US MER FOR VARIOUS RF SPECTRUM	29
FIGURE 20 – HPON US CAPACITY – OFDMA VS. SC-QAM	29
FIGURE 21 – HPON DOWNSTREAM RF PERFORMANCE	30
FIGURE 22 – HPON DOWNSTREAM RF CAPACITY, OFDM V SC-QAM	30
FIGURE 23 – EPON BURST TRANSMIT OVERHEADS	32
FIGURE 24 – 10G EPON UPSTREAM CYCLE TIME & ONU IMPACTS	34
FIGURE 25 – 10G & 1G CONTROL OVERHEAD EFFICIENCY	35
FIGURE 26 – 10G & 1G UPSTREAM COEXISTENCE CAPACITY	35
FIGURE 27 – 1G EPON, 10/10 & 10/1 COEXISTENCE, AND GPON	36
FIGURE 28 – D3.1 AND EPON/GPON CAPACITY	37
FIGURE 29 – AVG TX BURST SIZE, TRAFFIC EVENLY DISTRIBUTED	38
FIGURE 30 – AVERAGE TX BURST SIZE FOR HEAVY USERS	38
FIGURE 31 – RELATIVE COSTS FOR EPON & HPON SYSTEMS	39
FIGURE 32 – FIBER TRUNKS, WAVELENGTHS AND OLT/CCAP PORTS REQUIRED	40
FIGURE 33 – RELATIVE ENERGY COSTS FOR HFC, EPON AND HPON SYSTEMS	41

List of Tables

Title	Page Number
TABLE 1 – 2014 SERVICE TIER MIX, RATES & GROWTH	9
TABLE 2 – 2014 SERVICE TIER MIX, RATES & GROWTH	11
TABLE 3 – MAPPING D3.1 TO PON EQUIVALENTS, DOWNSTREAM CAPACITIES	20
TABLE 4 – MAPPING D3.1 TO PON EQUIVALENTS, UPSTREAM CAPACITIES	21
TABLE 5 – 10G EPON UPSTREAM EFFICIENCIES	33
TABLE 6 – 10G EPON UPSTREAM CAPACITIES	33

Introduction

Over the years, Hybrid Fiber-Coax (HFC) has continually evolved to push fiber deeper. Eventually, it will reach the point where it becomes a Fiber to the Premise (FTTP) architecture, but this may take decades at an economical pace. FTTP is happening today in new Greenfield deployments, yet there are significant operational challenges to make this transformation in existing HFC infrastructure, a.k.a. Brownfields.

So, how long will HFC last? What strategies can be deployed to help extend the life of HFC through this transition period? These are some of the considerations the industry needs to tackle. The paper gives an insight into the network capacity requirements over the next 10-20 years. This analysis then introduces a strategy called Selective Subscriber Migration which move the top Premium Tier subscribers to FTTP to create a mixed HFC/FTTP system. With these traffic engineering improvements and the introduction of DOCSIS 3.1, the life of the HFC may be extended for decades.

The next issue to address is the appropriate FTTP technology to facilitate this transition. Radio Frequency over Glass (RFoG) is one technology under consideration. Until recently RFoG has been perceived as a transitional technology for FTTP on the way to Passive Optical Network (PON) everywhere. However, recent RFoG advances that *completely* eliminate Optical Beat Interference (OBI) have unlocked the full DOCSIS 3.1 potential over the FTTP architecture. This along with traditional binary PON systems being *completely* transparent makes the new RFoG a transformational technology that supports both the legacy DOCSIS RF infrastructure along with traditional binary PON, to become a truly Hybrid PON (HPON) architecture.

HPON provides operators with a choice of either PON or DOCSIS technologies over FTTP. The use of standard PON technologies such as EPON and GPON are well known, but the benefits from DOCSIS 3.1 over FTTP is a brand new phenomena. The paper explores what this brings to the operator’s table and why an operator might consider this path.

In addition to FTTP architectures, HPON also enables other Fiber Deep topologies such as Fiber to the Curb (or Tap), Fiber to the MDU and even Fiber to the Deep Node (N+0). By leveraging the extremely high bandwidth of existing coax as the final drop cable, eliminating the costs of pulling fiber over the drop cable, and sharing ONU costs across multiple homes; these other Fiber Deep topologies with DOCSIS 3.1 (D3.1) over HPON provides operators with very interesting and cost effective alternatives to pure FTTP.

Finally, with D3.1 capabilities unleashed over HPON, the paper takes a close look at how D3.1 capacities over HPON align with various PON technologies. With this knowledge, operators are better prepared to decide when it is best to use D3.1 over HPON or traditional PON technologies or a combination of both over HPON.

Our conclusion is that the Hybrid PON FTTP architecture provides the best of both worlds and gives operators the greatest flexibility with minimal operational impacts in transforming their HFC into a true FTTP architecture on their own terms over the next couple decades. HPON is a revolutionary new FTTP breakthrough.

HFC to FTTP Capacity Planning

1.1. How Long before HFC Capacity runs out?

This fundamental question is the stuff that keeps us up at night. But is it really the right question to ask? Rather, we need to dissect this into more basic questions:

1. When will the migration of existing HFC to FTTP absolutely need to start?

Some may argue that it is already here. Competitive pressures to offer 1 Gigabit (1G) services have required some cable operators to offer this service today over FTTP networks. However, this is really just a short term gap while we wait for the deployment of DOCSIS 3.1. Once D3.1 is available in the near term, HFC can truly offer 1G services, too. However, as we will see later, at current growth rates this might buy us 5-10 years before the Top Tier hits the HFC limits.

2. How long will it take to migrate ALL HFC subscribers to FTTP?

This question is an economic one depending on how aggressively an operator wants to invest in their outside plant. An analysis by Venk Mutalik [1] shows that this will still be in at least 20-40 year window to get ALL HFC subscribers migrated to an FTTP network. If you think about it, an operator moving 5% of all subs to FTTP each year (which is aggressive and expensive) would require 20 years. Based on historical spending on plant upgrades, this seems overly optimistic.

3. How can an orderly transition to FTTP over this extended period?

Given that the transition will take multiple decades, then this becomes a critical question on how to extend the useful life of HFC for several more decades. This will require the adoption of D3.1 technologies and the use of an intelligent strategy for moving subscribers to FTTP. This paper proposes a strategy called Selective Subscriber Migration.

4. Finally, what are the appropriate technologies to use over the FTTP architecture?

One possibility is traditional binary PON technologies such as EPON and GPON along with their future evolutions. However, recent advances that have eliminated OBI in RFoG systems mean the DOCSIS 3.1 over FTTP is now a viable alternative as well. New Hybrid PON (HPON) technology allows the choice of either or both.

1.2. Traffic Engineering Fundamentals

Significant work in the network traffic engineering space has been done by Tom Cloonan [2]. Some of his graphs have become infamous and are called “Cloonan’s Curves”. These graphs incorporate an observed phenomenon called Nielsen’s Law. Nielsen’s Law roughly states that the highest offered internet speed will increase at an annual growth rate of 50%. This chart is shown in Figure 1.

With the migration to DOCSIS 3.1, the capacity of HFC plant is roughly 10 Gbps. As can be seen by the circled region on the chart, the expected 50% growth hits the 10G ceiling around the year 2024, less than a decade away. Remember, this date assumes continued 50% growth rates the whole time with a corresponding migration to all IP Video as well.

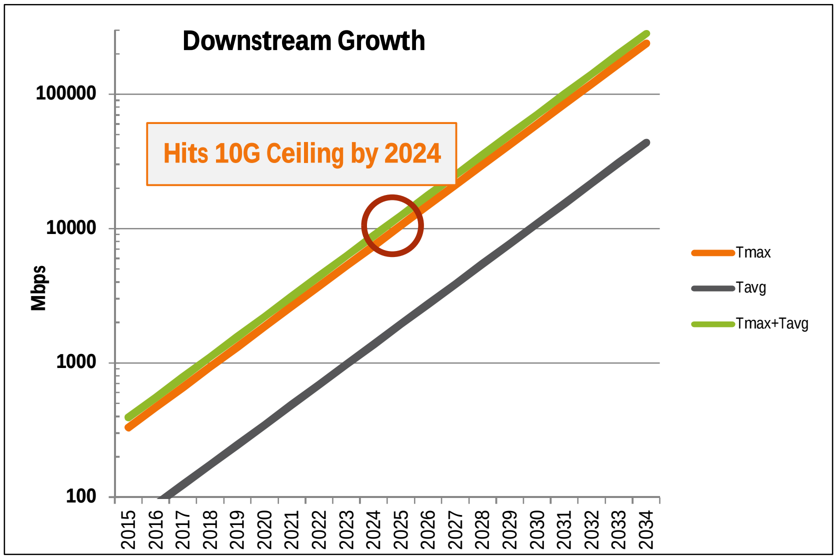

Figure 1- Cloonan's Curve and Nielsen's Law

At first glance, this may be disconcerting to cable operators who might think that HFC will have run out of capacity by that date. The reality is that this only represents the Top Billboard Tier which is typically less than 1% of all subscribers. So this date is the answer to the first question above on when the HFC to FTTP migration must begin. Obviously an operator may decide to start the migration ahead of this time.

Now let's try to understand what happens to the other 99% of the subscribers. That requires a deeper dive into the Network Capacity Traffic Engineering fundamentals. A detailed network capacity model is described in [3]. The overview of Mike Emmendorfer's Network Quality of Experience (NQoE) formula is shown in Figure 2. The NQoE formula goals include:

- Achieve Max Service Tier even during busy periods
- Allocate appropriate amount of network resources
- Configurable to accommodate any data network
- Accommodates Estimates of Service Tier and Traffic Growth Rates
- Achieve Max Service Tier Though Next Network Capacity Adjustment

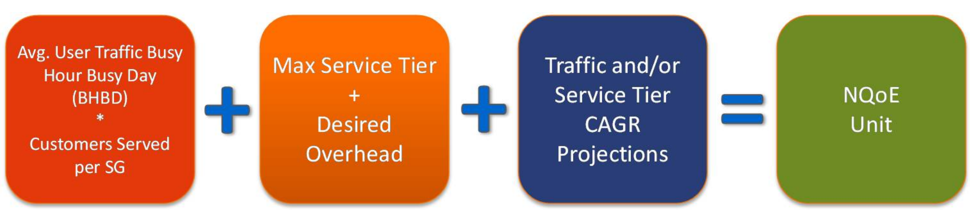

Figure 2- NQoE Formula Overview

While the detailed formula is extremely complex, the simplified version below has been found to work quite well in most situations:

$$\mathbf{C >= (Nsub * Tavg) + (K * Tmax\_max)} \quad \textbf{(Eq. 1)}$$

– C is the required Bandwidth Capacity for the service group
– Nsub is the total number of subscribers within the service group
– Tavg is the average bandwidth consumed by a subscriber during busy-hour
– Tmax_max is the highest Tmax offered by the MSO
– K is the "magical" QoE constant (larger values of K yield higher QoE levels)…
    - K values for typical scenarios will fall in the range **K = 1.0-1.2**

The first component in Equation 1 represents the average static traffic load and is a function of the number of subscribers per Service Group (SG) and the average bandwidth per sub at busy hour. The second component of the equation is the headroom required for good Quality of Experience (QoE). Tmax is the Maximum Sustained Traffic Rate parameter for DOCSIS Service Flows. Tmax_max is the highest Tmax across all Service Flows. It should be large enough to support a burst from the highest offered service tier. Many operators may choose a QoE constant, K, equal to 1.2 to give themselves an additional 20% cushion.

With this equation in hand, let's take a look at several example traffic engineering scenarios that may happen over the next five years. For a present day HFC scenario, assume a top service tier of 300 Mbps (i.e. Tmax_max) with 500 subs per SG and Tavg = 400 kbps. This scenario requires 200 Mbps for the static traffic load and 360 Mbps for QoE headroom for a minimum capacity of 560 Mbps. So an operator might deploy 16 DOCSIS 3.0 channels (96 MHz) to support this. This is shown on the left of Figure 3.

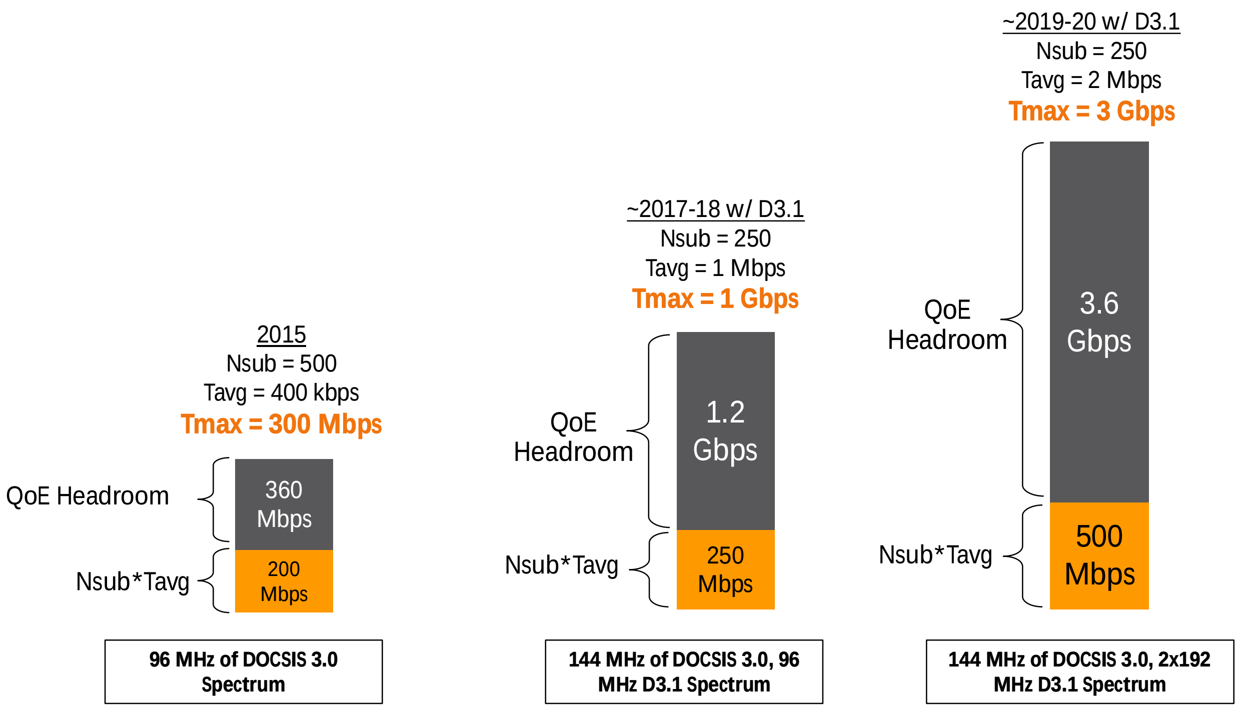

Figure 3 – Traffic Engineering with DOCSIS 3.1

A couple years down the road, the second scenario supports a max service tier of 1G with 250 subs per SG and Tavg = 1 Mbps. As shown in the middle of figure 3, this scenario requires almost 1.5 Gbps. This might be achieved by bonding 24 DOCSIS 3.0 channels (144 MHz) with 96 MHz D3.1 OFDM channel.

By the end of the decade, operators may try to max out D3.1 CPE capabilities and offer a 3 Gbps service tier. By then, Tavg might be 2 Mbps. This scenario requires at least 4.1 Gbps of capacity, which might be achieved by bonding 24 DOCSIS 3.0 channels (144 MHz) with a pair of 192 MHz D3.1 OFDM channels.

Looking at the three scenarios, the DOCSIS spectrum has soared from 96 MHz to 240 MHz to 528 MHz by the end of the decade. To control this spectrum growth, should an operator consider splitting SG? The answer to this is, “No”. Splitting a service group only impacts the average static load. As can be seen in Figure 3, this becomes a smaller and smaller component of the traffic engineering. During this stage, it becomes more important to increase the HFC spectrum (e.g. from 750 MHz plant to 1 or 1.2 GHz plant). But even that may not be enough.

1.3. Selective Subscriber Migration Strategy

What other options does an operator have to manage this bandwidth growth? This requires a more detailed look at the other service tiers in addition to the Top Billboard Tier. This was first discussed by the author in [4]. This paper looked at various service tier breakdowns from several major North American MSOs. A representative sampling of a service tier breakdown is shown in Table 1. The Top Billboard Tier for this sample in 2014 was 300 Mbps and less than 1% of the subs took this service.

Looking at the other service tiers, roughly 14% were in the “Performance” Tier @ 75 Mbps with the majority of subscribers in the Basic Tier (65% @ 25 Mbps) and Economy Tier (20% @ 5 Mbps).

2014 Service Tier Levels	% of Subs	Tmax (Mbps)	Tmax CAGR
Top Tier – Billboard rate	1%	300	50%
Performance Tier	14%	75	32%
Basic Tier	65%	25	26%
Economy Tier	20%	5	15%

Table 1 – 2014 Service Tier Mix, Rates & Growth

Perhaps the most interesting aspect of this study was in the growth rates for the different service tiers. While the Top Billboard Tier was growing at the famous Nielsen’s Law 50% CAGR (Compounded Annual Growth Rate), the other service tiers had a significantly lower growth rate. The lower the performance of the service tier, the lower the CAGR.

While this difference in CAGR among service tiers initially made us scratch our heads, it actually makes much sense. It turns out that if all service tiers grew at Nielsen’s 50% CAGR, then every installed cable modem would be obsoleted within 2-3 years of introduction. That level of investments would be staggering. Since operators are the ones that control the CAGR for each service tier, they effectively control how long the cable modem technology stays viable in the field. Note that the Economy Tier could still be using DOCSIS 2.0 modems from over a decade ago. In 2014, a 16-channel bonded cable modem was probably used for the Top Billboard Tier of 300 Mbps. In a few years, once 1G service is available on HFC, the

300M service with its 16-bonded channel modem becomes the Performance Tier. A few years after that, it will be relegated to the Basic Tier.

What does this mean from an HFC to FTTP migration point of view? Figure 4 shows the Tmax growth for each service tier over a 20 year window. As shown previously, the Top Billboard Tier hit the ~10 Gbps HFC limit by 2024. What about the other service tiers? The Performance Tier Tmax does not hit 10 Gbps until 2033. By this time, these subs will need to be moved to FTTP. Meanwhile both the Basic and Economy Tiers are well under the 10 Gbps limit 20 years from now. In this example, 85% of the HFC subs will be capable of staying on HFC for more than two decades. And this assumes that growth rates continue unabated. Within 10 years, HFC infrastructure will be able to deliver an Ultra-HD video stream to every eyeball on the plant, so some other new application will have to drive the growth engine after that.

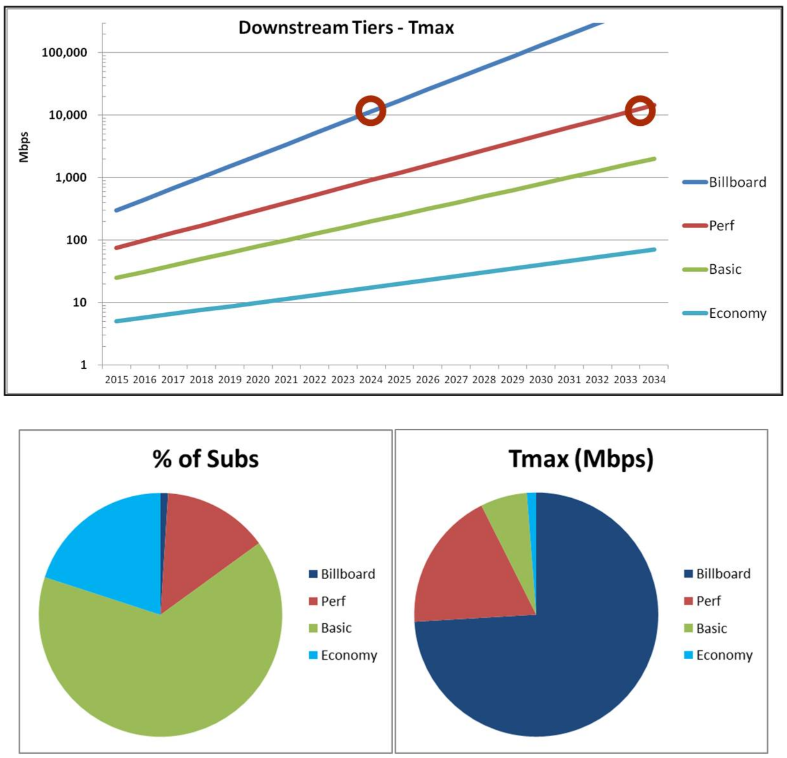

Figure 4 – Individual Service Tier Growth Year Over Year,% Subscribers and Tmax for the Various Tiers

Note that the approximately 10 Gbps HFC limit assumes that DOCSIS 3.1 has been deployed and legacy MPEG spectrum has been recovered. If an operator chooses to stay with DOCSIS 3.0 technology, then their best Tmax would be around 1 Gbps using 32-channel bonded modems. From Figure 4, the Performance Tier now needs to move to FTTP by 2024 and the Basic Tier (i.e. 65% of subs) need to be on FTTP by ~2030. This drastically alters the HFC to FTTP migration plans.

Taking a closer look at the impact of this Selective Subscriber Migration strategy, Table 2 shows where each service tier might be by the end of the decade. With a 50% CAGR, the Top Billboard Tier might be at

3G service rate. The Performance Tier with its approximately 32% growth now reaches ~500M service rate; while the Basic Tier has grown to 100M and the Economy Tier is around 10M.

~2020 Service Tier Levels	% of Subs	Tmax (Mbps)	Tmax CAGR
~~Top Tier – Billboard rate~~	~~1%~~	~~3000~~	~~50%~~
Performance Tier	14%	500	32%
Basic Tier	65%	100	26%
Economy Tier	20%	10	15%

Table 2 – 2014 Service Tier Mix, Rates & Growth

From our previous traffic engineering example in Figure 3, the Top Billboard Tier still fits within the DOCSIS 3.1 cable modem capabilities (i.e. 2x192 MHz OFDM channels), but now requires the operator to have 4.1 Gbps of DOCSIS capacity to offer this service tier.

With the Selective Subscriber Migration strategy, the Top Billboard Tier would be migrated from HFC to FTTP. Note that this tier is typically less than 1% of the total subscribers, so a 250 sub SG might only have 2 or 3 subs in this tier on average that need to migrate.

Now that the Top Tier has been removed, the traffic engineering can be re-calculated for the remaining HFC subs. The static load is basically unchanged since so few subscribers have been removed. However, the QoE portion of the formula has been drastically reduced since the top service rate (Tmax_max) is now 500 Mbps instead of 3 Gbps. This means that the operator now only needs **1.7 Gbps** of DOCSIS capacity **instead of** the previous **4.1 Gbps** before the subscriber Migration. This maps to **a savings of 250-300 MHz** of spectrum using DOCSIS 3.1 OFDM channels. By migrating the Top Billboard Tier to FTTP, the operator has effectively extended the life of the HFC for the remaining subscribers.

1.3.1. Network Capacity Modeling of FTTP Migration

A network capacity model of this service tier example is shown in Figure 5 assuming 128 subs per SG. This particular model migrates subscribers to FTTP starting with the highest available service tier as DOCSIS capacity exceeds 10 Gbps. The red portion of each bar is the QoE element driven by Tmax_max. The orange portion of each bar is the static load.

As can be seen in Figure 5, Tmax dominates in the early years. The 50% CAGR on the Top Billboard Tier is evident in the growth through 2023. In 2024, the Top Billboard Tier (i.e. <1% of subs) is moved to FTTP and there is a drastic reduction in DOCSIS capacity that is required. The growth rate is now the slightly lower Performance Tier. The Performance Tier is fine on the HFC thru 2028 but needs to migrate to FTTP by 2029. By 2029, there is only 15% of subs that had to be migrated from the HFC to the FTTP.

At this point, with the highest tiers moved to FTTP and continued growth in Tavg, the static load has now become the dominant piece of the network capacity formula. What starts to happen in the year 2031 in Figure 5 is that Basic Tier subs start to migrate to FTTP in order to reduce the static load. Hence, the orange component starts to drop as there are fewer and fewer subs left on the HFC. Figure 5 also assumes a fixed SG size of 128 subs. Once the static traffic load starts to dominate, it now becomes desirable to split SG size which will reduce the static traffic load. Figure 6 provides four charts corresponding to SG sizes of 256, 128, 64, and 32 subs.

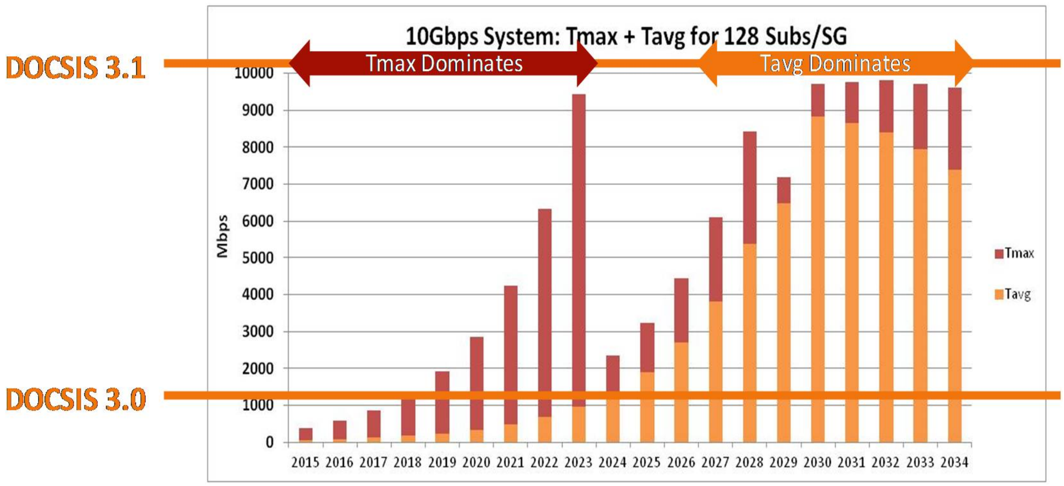

Figure 5 – Capacity Needs over Time, 128 Subs per SG

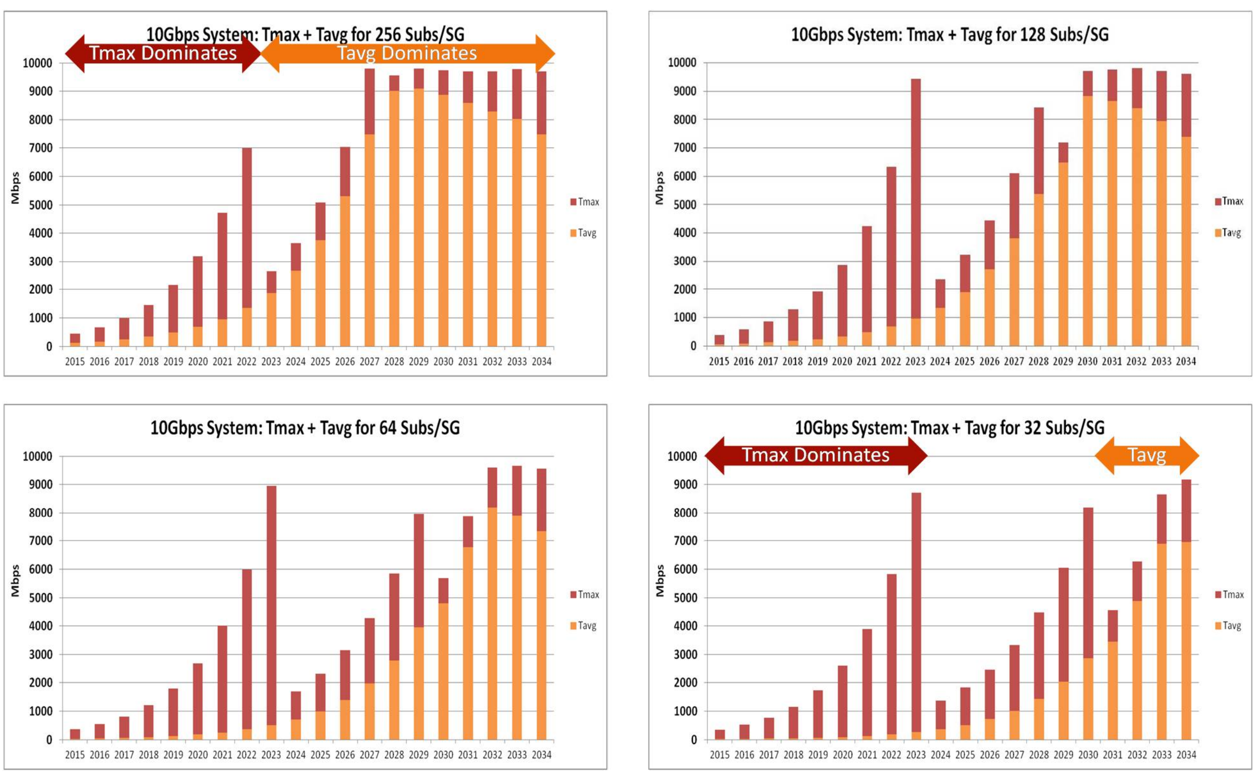

Figure 6 – Capacity Needs over Time, Various SG sizes

As can be seen across the four charts, SG size has little impact over the next 8-10 years on determining when the Top Billboard Tier needs to migrate to FTTP: 32 subs per SG is only 1 year after 256 subs per SG. However, SG size is a big factor on when the static load starts to dominate. For 256 subs per SG, the static load becomes the dominant portion by the year 2023. For 32 subs per SG, the static load does not become dominate until the following decade.

What can we learn from these charts regarding our HFC plant strategies? First, migrating to DOCSIS 3.1 is important. Figure 5 shows the relative HFC limits with both DOCSIS 3.1 and 3.0. Having DOCSIS 3.1 capabilities greatly extends the life of the HFC. Second, the near term focus should be on increasing spectrum for DOCSIS 3.1. This might mean upgrading a 750 MHz plant to 1002 MHz or even 1218 MHz plant. To offer Gbps downstream services over the HFC, the operator should also consider an 85 MHz upstream split at this time as well. Given existing asymmetric traffic loads, the 85 MHz upstream should match well with a 1 GHz downstream. As the operator looks further down the line, the static load will start to dominate and SG splits will come into vogue once again. As operators migrate the highest tiers to FTTP, they should keep in mind that they will eventually need to do some SG splits on HFC as well.

1.3.2. Economic Impacts of Selective Subscriber Migration

The previous section showed some of the traffic engineering benefits of the Selective Subscriber Migration strategy; now let's investigate the economic impacts of this approach. Figure 7 shows example plant upgrade costs for a suburban case study with a serving area of almost 1000 homes passed (HP).

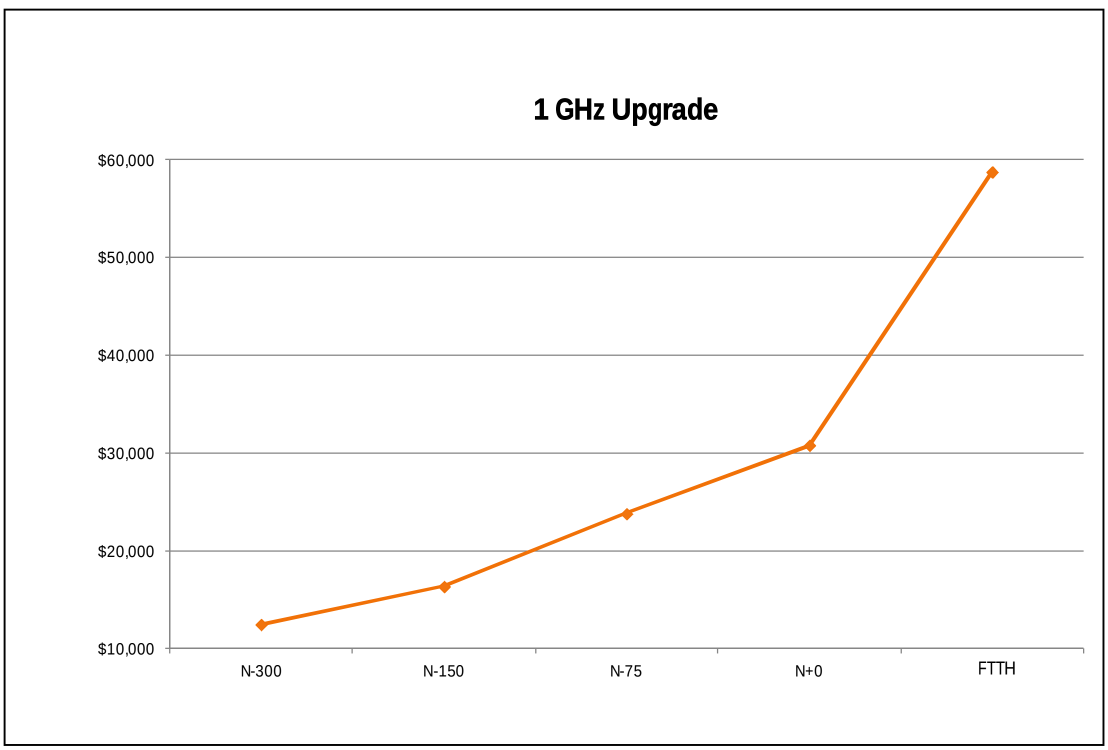

Figure 7 – Example 1 GHz/85 MHz Upgrade Costs vs. FTTP

A full FTTP upgrade is compared to various HFC upgrades to 1 GHz/85 MHz. The HFC options show progressively deeper fiber. N-300 has no more than 300 HP on any leg and is typically N+1 or N+2 with a very limited number of outlying homes at N+4. N-150 and N-75 continue to increase nodes and reduce SG size. These options all use existing node & amplifier locations. Finally, the N+0 upgrade is almost a complete rebuild of the HFC with nodes put in new sites as needed. The N+0 upgrade averages about 60 HP per node.

As is shown in Figure 7, the plant upgrade costs skyrocket as fiber goes deeper. The ~$30K upgrade cost of N+0 is more than twice that of the N-300 upgrade. The ~$60K cost of FTTP is double the cost of N+0 and is five times more expensive than the N-300 upgrade. A key reason on why the FTTP is much more expensive is that a significant portion of the fiber installation is associated with the last drop cable over the last couple hundred meters.

With the Selective Subscriber Strategy, an operator only needs to do the N-300 HFC upgrade in the near term at substantially less money than either N+0 or FTTP. The N-300 upgrade provides essentially the same spectrum as N+0, so this satisfies the short term needs when Tmax dominates. With the money saved, a handful of Top Billboard Tier customers can be given FTTP connections. Over the next decade, the Performance Tier can be gradually migrated to FTTP. When this happens, the fiber will also be pulled to enable a fiber deeper HFC migration to N-75 or even N+0 when needed over ten years from now. This approach allows operators to grow slowly as needed and spread plant investments over a lengthy time window, yet still be prepared for fiber deep SG splits when needed a decade from now.

1.3.1. Selective Subscriber Migration Summary

In summary, selectively Migration subscribers from HFC to FTTP starting with the highest service tiers, combined with DOCSIS 3.1 and 1 GHz/85 MHz upgrades to maximize HFC capacity, can provide a sensible HFC to FTTP transition and relieve pressure to reclaim legacy spectrum. It not only saves money, it adds decades of life to the HFC plant for 80% to 95% of the total subscribers by being able to support Gbps services to the masses. And if entertainment and Ultra-HD is all that these Basic & Economy Tiers require, then maybe these subs can reside on HFC forever.

DOCSIS® 3.1 Overview – Extending the Life of HFC for Decades

2.1. DOCSIS 3.1 Overview and Benefits

DOCSIS 3.1 [5] is a key element in this strategy to extend the life of the HFC for decades. Some of the key underlying D3.1 technologies include: OFDM, LDPC Forward Error Correction (FEC), Multiple Modulation Profiles in the downstream, and Time and Frequency Division Multiplexed (TaFDM) CMTS Scheduler.

DOCSIS 3.1 provides these important benefits:

- DOCSIS 3.0 backwards compatible; operates in existing HFC plants no changes
- Ultra-wide, variable width channels:
    - 24-192 MHz DS, 6.4-96 MHz US channels
- Higher modulations yield Increased spectrum capacity
    - DS to 4096-QAM (16,384-QAM optional), US to 1024-QAM (4096-QAM optional)
    - Bps / Hz gains: 40% - 75% DS; 66% to 100% US
- New spectrum availability -
    - Optional future spectrum of 1218 MHz DS, 204 MHz US for 10+ Gbps DS, 1.8 Gbps US
    - Robust OFDM + LDPC Leverages Roll-off region in existing plants (~1 Gbps possible)
- OFDM + LDPC and TaFDM maximizes existing upstream (e.g. ~250 Mbps in 42 MHz)
- DOCSIS 3.1 MAC enables bonding across 3.0 SC-QAM + 3.1 OFDM

The DOCSIS 3.1 specification also requires that the first generation D3.1 cable modems must support two 192 MHz OFDM channels downstream and two 96 MHz OFDMA channels in the upstream. That means these D3.1 modems, once deployed in the field, will be capable of providing capacities of 5 Gbps DS and 1.8 Gbps US.

2.2. DOCSIS 3.1 Capacity Examples

DOCSIS 3.1 greatly increases the potential capacity of HFC. This is shown in Figure 8 for several different HFC plant spectrums. Today's DOCSIS 3.0 cable modems are limited to 32x8 configurations. The 32-bonded downstream channels enable just over 1 Gbps of capacity. The 8-bonded upstream channels provide about 200 Mbps of upstream capacity.

From an HFC plant perspective, total capacity for today's HFC is actually the combination of both the DOCSIS 3.0 channels and the MPEG Video QAM channels. This is represented in Figure 8 with the 3.0+QAM bars.

For a 750 MHz HFC plant, the downstream capacity goes from ~4 Gbps for 3.0+QAM to ~7 Gbps for DOCSIS 3.1. A 1 GHz HFC plant sees the gap grow from ~5 Gbps for 3.0+QAM to almost 9 Gbps for D3.1. Finally, DOCSIS 3.1 can provide over 10 Gbps of downstream capacity over 1218 MHz of spectrum.

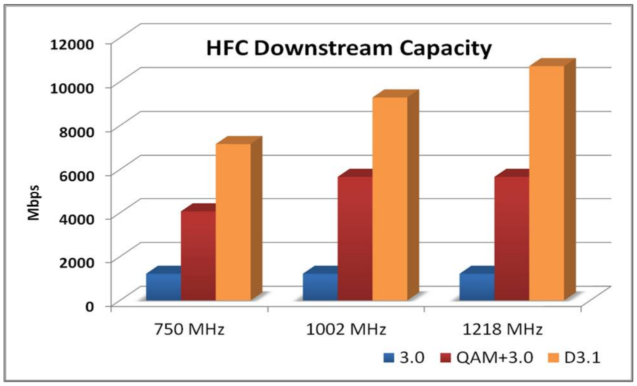

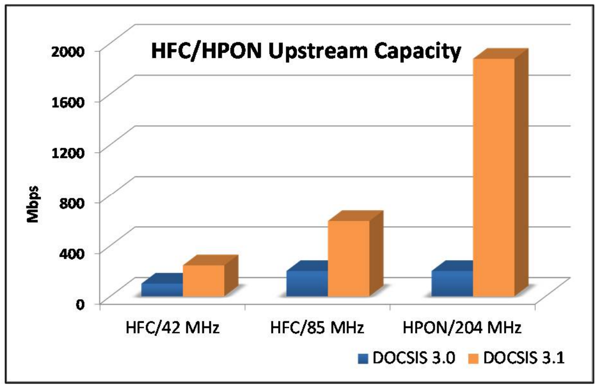

Figure 8- HFC Upstream and Downstream Capacity

2.3. DOCSIS 3.0 to 3.1 Migration Example

It is important to understand how an operator might migrate from 3.0 to 3.1. In the first step, no HFC plant changes are needed. DOCSIS 3.1 can be introduced into existing plants, providing capacity gains with improved spectral densities. The challenge becomes finding available spectrum for D3.1. In the downstream, D3.1 provides a bonus in that it can operate in the roll-off region. For example on a 750 MHz plant, an OFDM channel could be placed from 750 to 900 MHz. An analysis of an actual 870 MHz plant showed that there may be as much as 1 Gbps of capacity in the roll-off region, but this may vary substantially from HFC plant to plant.

At this point, the operator has some D3.1 upgrade options that they may choose to pursue. The first option is to expand the existing HFC spectrum. It is suggested that the downstream be extended to at least 1002 MHz. 1 GHz upgrades are straightforward; cost effective; and have been done for years. Some operators may consider going to 1218 MHz but this will introduce some additional challenges, especially considering power and tilt as well as potential MoCA interference. When upgrading the HFC downstream spectrum, the operator may also consider increasing the upstream split to 85MHz. This will help future proof the HFC from an upstream capacity perspective.

The second D3.1 upgrade option is to migrate select subscribers to HPON. This will give these HPON subs immediate access to expanded spectrum (e.g. 1218MHz downstream, 204MHz upstream) while not requiring any immediate changes to the existing HFC.

Every operator has unique circumstances that may vary from plant to plant. Which D3.1 upgrade option is selected and in which order will be very dependent on each situation. It may be that many operators will pursue both options in parallel.

The final piece in this D3.1 migration is for the operator to enable all IP video so legacy MPEG spectrum can be reclaimed and the entire HFC spectrum utilized by DOCSIS. The IP video deployment should leverage the latest Multicast adaptive bit rate (ABR) protocols to make the most efficient use of capacity.

The above migration sequence will allow operators to grow their DOCSIS capacity on HFC from 1 to 2 to 5 to 10 Gbps over time.

Hybrid PON (HPON) – A Revolutionary FTTP Breakthrough

3.1. HPON Overview

What exactly is Hybrid PON, a.k.a. HPON? HPON contains a new innovative fiber splitter technology that 100% eliminates OBI for RFoG wavelengths. It requires very minimal power, on the order of 150 mW per drop connection. While this is no longer a purely passive plant, it is not different from many PON installations that require PON extenders or Remote OLT at much higher power consumption.

While minimal power is needed for RFoG wavelengths, HPON is still completely passive and compatible with Ethernet and PON technologies: 10G Ethernet, EPON, 10G EPON, GPON, NG-PON2. So even if the RFoG wavelengths lose their power, PON and Ethernet continue to operate.

HPON is standards compliant on both ends of the network. HPON is completely backwards compatible with today's RFoG ONU and RFoG Headend Optics. An operator may use any vendor's RFoG compliant ONU or Optics. Because HPON eliminates OBI, an operator is free to choose any vendor's CMTS/CCAP with its traditional US scheduler. HPON enables full D3.1 performance with OBI-Free environment. Certain other RFoG solutions needed specialized CMTS 3.0 scheduler which handicaps performance; and is not usable with DOCSIS 3.1.

What exactly is "Hybrid" about HPON? There are multiple meanings to the word Hybrid:

- Hybrid HFC & FTTP: Supporting Legacy HFC services Over FTTP
- Hybrid D3.1 & Traditional Binary PON (e.g. EPON, GPON, NG-PON2)
- Hybrid DOCSIS Over both HFC & FTTP
- Hybrid DOCSIS 3.0 and DOCSIS 3.1
- Hybrid Passive & Powered
- Hybrid Asymmetric & Symmetric Applications

3.2. For Fiber To The Home (FTTP) Transition: DOCSIS or EPON/GPON?

Up until recently, EPON or GPON seemed to be the only reasonable long term FTTP choices. The DOCSIS over RFoG alternative was hampered by Optical Beat Interference – OBI as discussed in [10]. With HPON, DOCSIS 3.1 over FTTP becomes a viable long term option.

Note, it is not an either/or choice for the operator as HPON supports both EPON/GPON **AND** OBI-Free D3.1. The operator can support DOCSIS &/or EPON/GPON as needed, whichever is best suited to the service needs. For example, an operator might deploy symmetric 10G EPON over HPON for Business Services while D3.1 over HPON for Top Tier residential customers.

As operators consider DOCSIS or EPON over FTTP, there are certain philosophical considerations between them. EPON leverages Ethernet ecosystem. It supplies more than abundant bandwidth capacity up front and offers symmetric capabilities. From a MAC perspective, it kept a KISS principle and relies on small SG and polling for access.

The DOCSIS philosophy is to fit seamlessly into HFC infrastructure, being spectrum friendly. It supplies bandwidth capacity as needed – 'Just in Time'. This was evident with 3.0 as the number of bonded channels grew over time while always being backwards compatible. HFC and hence DOCSIS has had an asymmetric focus on residential applications. The MAC is full featured to provide guaranteed services to very large SG. Early DOCSIS days saw SG sizes >1000 modems.

HPON supports both DOCSIS and PON, so which should an MSO choose? EPON and GPON have been around for years and is well known; so let's first explore what this new OBI-Free D3.1 over HPON capability now brings to the table.

3.3. HPON and the Role of DOCSIS

HPON support for DOCSIS over FTTP brings many potential benefits to the operator. First and foremost is that it leverages the existing DOCSIS/HFC Infrastructure. This allows both CCAP and DOCSIS CPE investments to be reused in an HPON world. D3.1 over HPON supports legacy MPEG Video services. This means operators can reuse legacy STB investment in the field.

HPON unleashes D3.1 capabilities to the full extent, providing PON-like Gbps data rates for both downstream and upstream directions. Initial D3.1 modems will have 5 Gbps DS, 1.8 Gbps US capacities to start. This will enable true 1G Upstream services, unlike 1G EPON, GPON or 10G/1G EPON which lack sufficient QoE upstream capacity.

By leveraging the DOCSIS MAC capabilities, D3.1/HPON supports existing SG sizes AND distances which are significantly larger than traditional PON. DOCSIS is designed to handle 80 km distances with potentially 1000 modems, while traditional PONs are limited to 20 km and 32-64 ONU. This conserves trunk fibers & wavelengths as well as CCAP ports.

Looking to the future, D3.1 OFDM technology in an OBI-free environment offer the potential of 40 Gbps DS, 10 Gbps US on single wavelength.

3.4. Mixed HFC and HPON DOCSIS 3.1 Operation

In a Selective Subscriber Migration strategy, there may only be a couple Top Tier subs on the FTTP in a serving area. From a Headend infrastructure equipment perspective, it seems wasteful and expensive if an entire CCAP or OLT port must be dedicated this small number of customers. DOCSIS 3.1 over HPON can overcome this hurdle by reusing the same CCAP port that is being used by the HFC plant.

Figure 9 shows an example of how the HFC and HPON spectrum can overlap and be shared from a single CCAP port. This example assumes that most of the existing 750/42 MHz HFC spectrum is being used for 3.0 and legacy QAM services. This might include 24-32 3.0 channels. A 96 MHz 3.1 OFDM channel is placed on the HFC from 738 to 834 MHz so it only replaces two QAM channels and leverages the 750 MHz roll-off. This is enough DOCSIS capacity to offer 1G DS services and 100M US services (within 42 MHz).

Because HPON is full FTTP, it can support 1218 MHz downstream. The CCAP port can put two additional 192 MHz OFDM channels from 834 to 1218 MHz. This spectrum can then be sent down both the HFC and HPON. HFC modems will only use the 96 MHz OFDM bonded with 3.0 channels and ignore the top 2x192 OFDM. The HPON modems can bond across all OFDM and 3.0 channels as needed. This could enable a 2.5G or even 3G service in the downstream.

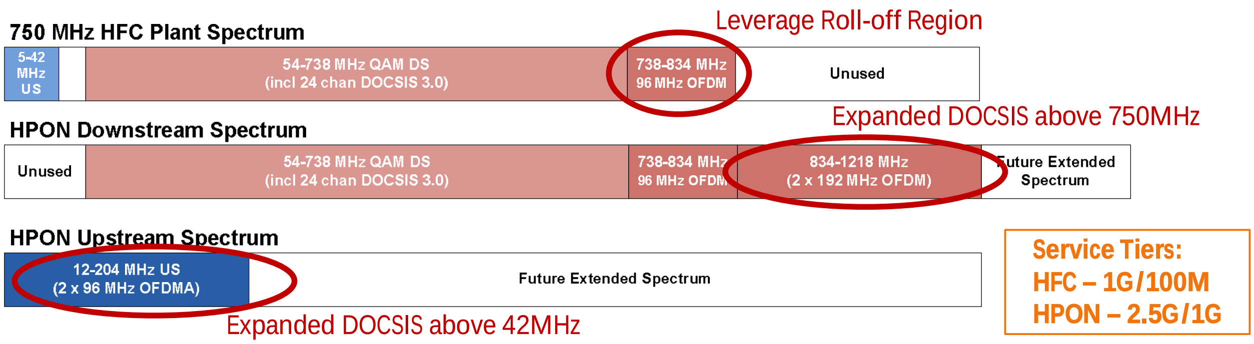

Figure 9 – HFC & HPON Spectrum Overlay

Another significant advantage of HPON is the isolation between downstream and upstream spectrum; each with its own dedicated wavelengths. This provides the operator with a cost effective operational mechanism for migrating select customers to a D3.1 204 MHz upstream with true 1G upstream services while keeping the vast majority on existing HFC. Over time, this capability can also enable Extended Spectrum RFoG with significant bandwidth capacity enhancements in both upstream and downstream. At the Headend, the 204 MHz HPON upstream can be combined with the 42 MHz HFC upstream and use the same CCAP port. The 42 MHz spectrum is shared between HFC and HPON while 42-204 MHz is available to HPON 3.1 modems. It is also noted that HPON provides improved upstream Signal to Noise Ratio (SNR) and reduces upstream noise funneling from ingress in the home which should make 4096-QAM modulation a reality in the upstream.

The overlapping spectrum has some additional benefits. Because the downstream spectrum can stay 54-1218 MHz, it can continue to support legacy services such as STB in the lower spectrum.

The bottom line with HFC and HPON spectrum overlay on the same CCAP port is that a small number of subscribers can be moved to FTTP cost effectively. No additional hardware is required, just licensing of additional D3.1 OFDM channels. For a PON migration, moving a small number of subscribers to FTTP might trigger the installation of an entire OLT where there may have been none before.

3.5. Stacking Up: D3.1 over HPON Capacities to Other PON Architectures

Operators have many different potential network options available to them and their competitors, so it is important to understand how these various technologies stack up against each other. A comparative chart of downstream capacities is shown in Figure 10. Note that PHY Layer Rates are after encoding and FEC (if used).

Copper based infrastructure has made significant progress over the years and VDSL2 and G.fast are the current state of the art. Figure 10 shows some estimates for these copper solutions. They 'only' provide hundreds of Mbps of downstream capacity to the user, not Gbps as in the other solutions.

The traditional PON technologies include GPON, 10G EPON and XG-PON. GPON provides almost 2.5 Gbps DS while 10G EPON and XG-PON provides ~8.7 Gbps of DS capacity. "10G" is a bit of a misnomer as it loses about 13% of capacity to the FEC. NG-PON2 was not included in this chart as it is a multi-wavelength technology and this focuses on what can be delivered to a user with a single wavelength.

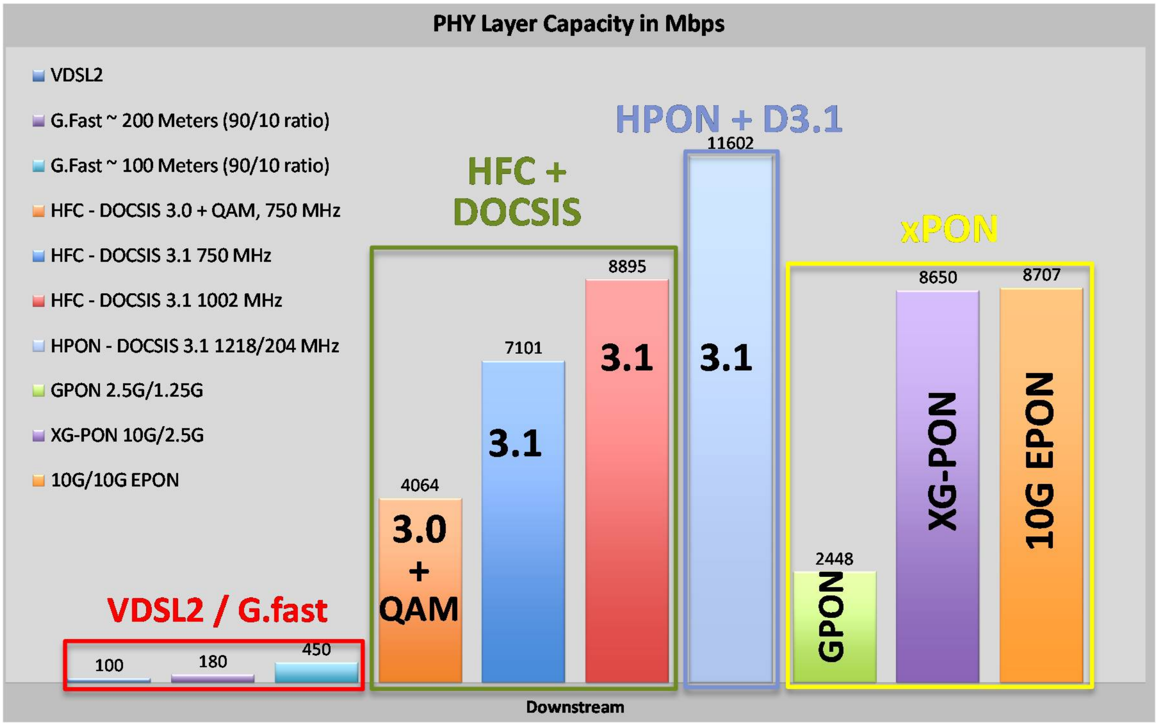

Figure 10 – Downstream Capacity: HFC & HPON vs. xPON

For DOCSIS on HFC, the chart shows the capacity for a 750 MHz plant with 3.0+QAM (4 Gbps); 750 MHz with D3.1 (7 Gbps); and 1 GHz plant with D3.1 (8.9 Gbps). Note that a 1 GHz plant with D3.1 is basically equivalent to a 10G EPON downstream capacity.

Finally, D3.1 over HPON provides almost 12 Gbps of capacity in 1218 MHz. This is 33% more downstream capacity than 10G EPON.

Downstream Spectrum	Nominal Data Capacity	PON Equiv
30 '3.0' + 96MHz OFDM	2 Gbps	GPON, 2 x 1G EPON
30 '3.0' + 2x192MHz	5 Gbps	2 x GPON, ½ 10G
30 '3.0' + 4x192MHz	8.7 Gbps	10G EPON, XG-PON1, NG-PON2
12-24 x 192MHz	~20-40Gbps	NG-PON2 (multiple λ)

Table 3 – Mapping D3.1 to PON Equivalents, Downstream Capacities

Table 3 shows a mapping of downstream capacity for various DOCSIS configurations into traditional PON systems. A DOCSIS system with 30 3.0 channels bonded with 96 MHz 3.1 OFDM channel provides roughly 2 Gbps and is roughly equivalent to GPON and is double 1G EPON. A 2x192 MHz OFDM with

3.0 channels now provides almost 5 Gbps, which is twice GPON but slightly more than half of 10G EPON. As the number of OFDM channels grow over time, just as 3.0 channels grew, a 4x192 MHz OFDM bonded with 3.0 channels is equivalent to 10G EPON downstream. Finally, work in our research labs shows that Extended Spectrum D3.1 can achieve up to 40 Gbps DS over a single wavelength. This downstream capacity would be equivalent to NG-PON2 which would require 4 wavelengths for the same capacity.

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

Table 4 – Mapping D3.1 to PON Equivalents, Upstream Capacities

Table 4 shows the upstream capacity mapping. An 85 MHz D3.1 HPON system upstream capacity is roughly equivalent to 1G EPON, 10/1G EPON & GPON with usable capacity in the 700-800 Mbps range. The 204 MHz D3.1 system equals XG-PON 2.5G US. Later in the paper shows that it also matches 10/10 + 10/1 EPON co-existence under certain traffic conditions.

3.6. HPON Topology Options

Up to now, discussion has focused solely on a migration from HFC to FTTP. As operators start to consider delivery of multiple Gbps services to every home, then a PON ONU needs to be in the premise (i.e. FTTP) as copper drop cable technology has limited bandwidth and prevents a FTTC approach with PON. However, things now change with DOCSIS over HPON. Coax is a great drop cable technology that can support more than 10 Gbps to each home. This now opens the door to looking at other fiber deep architectures besides FTTP.

Looking at where best to utilize the existing coax drop cable, D3.1/HPON could deploy a Fiber to the Curb or Tap architecture. This is depicted in Figure 11. New deployment technologies are now available that allow fiber strands to be economically blown into coax conduit. Figure 11 shows each HPON ONU at a Tap location driving coax drops to four homes. This approach saves the cost of pulling fiber drops to each home and shares the cost of ONU across multiple homes.

Another D3.1/HPON topology would be Fiber to the MDU. The ONU could be located in the basement or supply room and leverage existing coax distribution throughout the building. Alternately for a larger MDU, the fiber could be pulled to every floor where a single ONU serves the entire floor via coax.

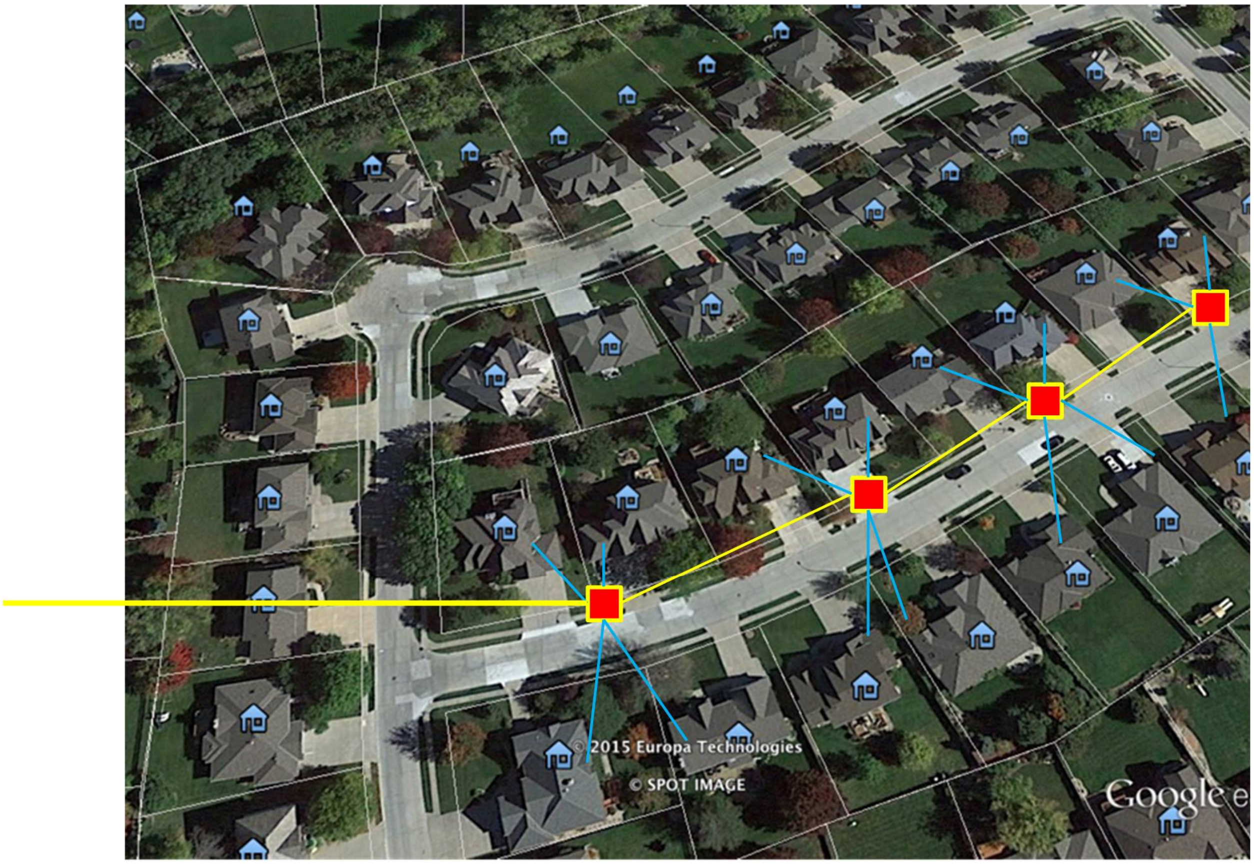

Figure 11 – Fiber to the Tap Example

3.6.1. HPON Topology Migration Example

An HFC to HPON migration example is provided to better understand the various Topology options.

3.6.1.1. HFC Baseline

Figure 12 shows the baseline of an existing HFC Plant. No changes are necessary to the plant.

3.6.1.1. Step 1 – FTTP for Select Top Customers

Step 1 is shown in Figure 13 where select customers are migrated to FTTP using HPON. This could be done with either DOCSIS or PON. In the figure, a business is connected with 10G EPON and a Top Billboard Tier user gets a D3.1/HPON FTTP connection to their home.

As operators pull fiber to these Top customers, they will most likely pull additional dark fibers as well. This will enable a future fiber deep migration along this path. As can be seen in Figure 13, the yellow boxes represent former amplifier locations that can now become Fiber Deep nodes and upgraded to 1 GHz/85 MHz.

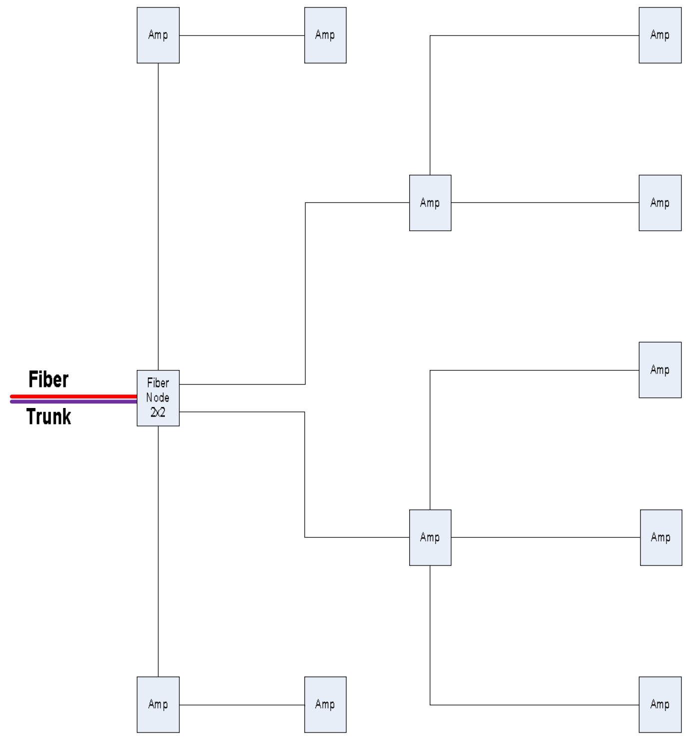

Figure 12 – HPON FTTP Migration: Baseline

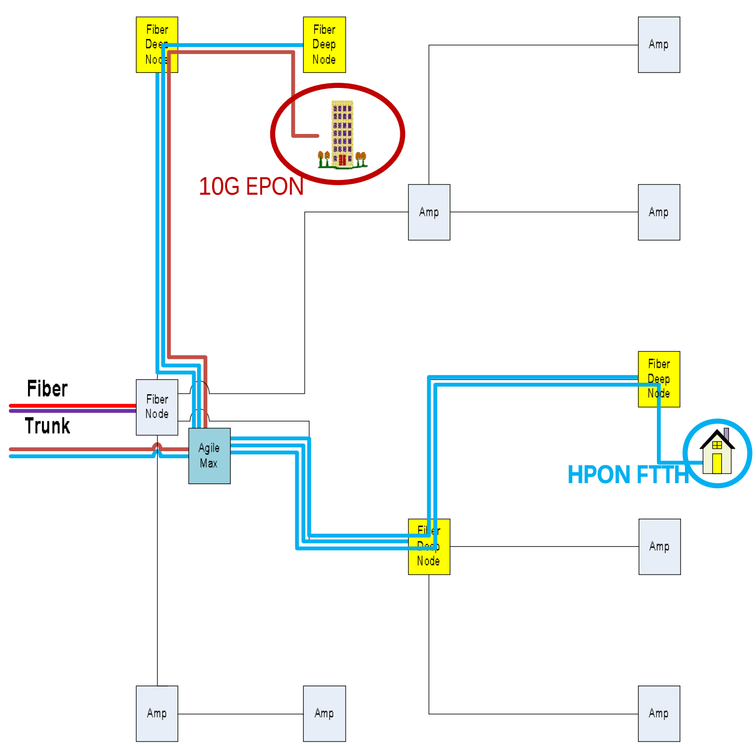

Figure 13 – HPON FTTP Migration: FTTP for Select Top Customers

HPON provides an additional benefit for Fiber Deep deployments. Traditionally, every Fiber Deep node would need its own set of Headend optics and require a separate wavelength on the fiber trunk. By leveraging RFoG optics, the HPON system can act as an aggregator for Fiber Deep nodes, and they can reuse the same optics being used for the D3.1/HPON FTTP home. This makes Fiber Deep more economical.

3.6.1.2. Step 2 – FTTP for More Performance Customers

Step 2 is shown in Figure 14 where more Top customers are migrated to FTTP using HPON. This example shows additional 10G EPON users, along with D3.1 users. The figure also depicts D3.1/HPON being delivered to an MDU as well as Fiber to the Curb being shared by several homes.

As the Performance Tier is migrated to FTTP (e.g. 5% to 15% of subs), then most of HFC plant will be covered by Fiber Deep. In the figure, there are only a couple stray amplifiers left without fiber before the entire HFC can be converted to Fiber Deep. Note that all of the Fiber Deep nodes are still sharing the same single set of RFoG optics.

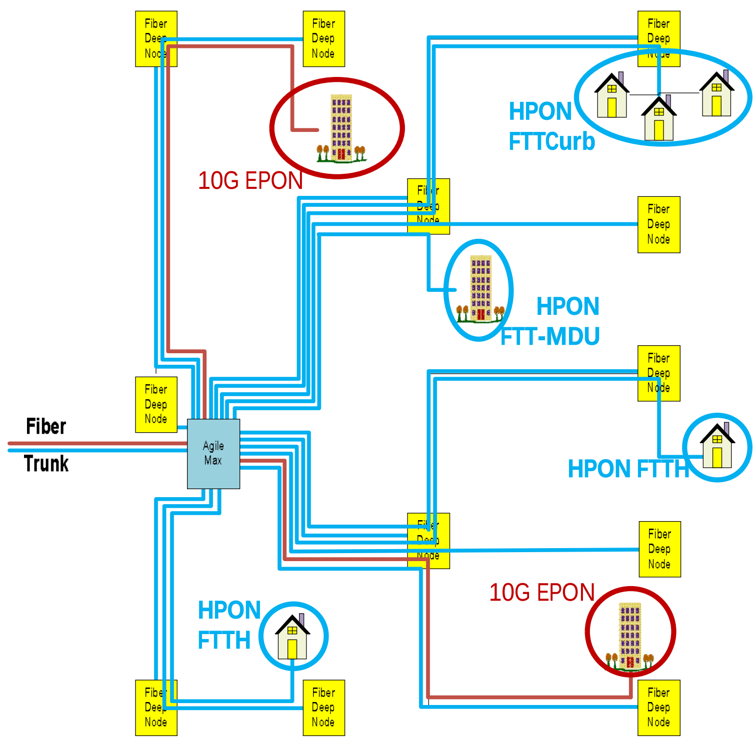

Figure 14 – HPON FTTP Migration: FTTP for More Performance Customers

Step 3 is shown in Figure 15. At this point, the Fiber Deep HFC has been completely built out and the Top Tier customers moved to FTTP. Eventually, the static traffic load will increase as in Figure 6 and the operator will need to split SG. At this point, all SG segmentation is localized to the HPON splitter and multiple wavelengths came be sent down for the different SG. In Figure 15, each color represents a distinct service group. SG segmentation becomes simple. This is analogous to node segmentations done today.

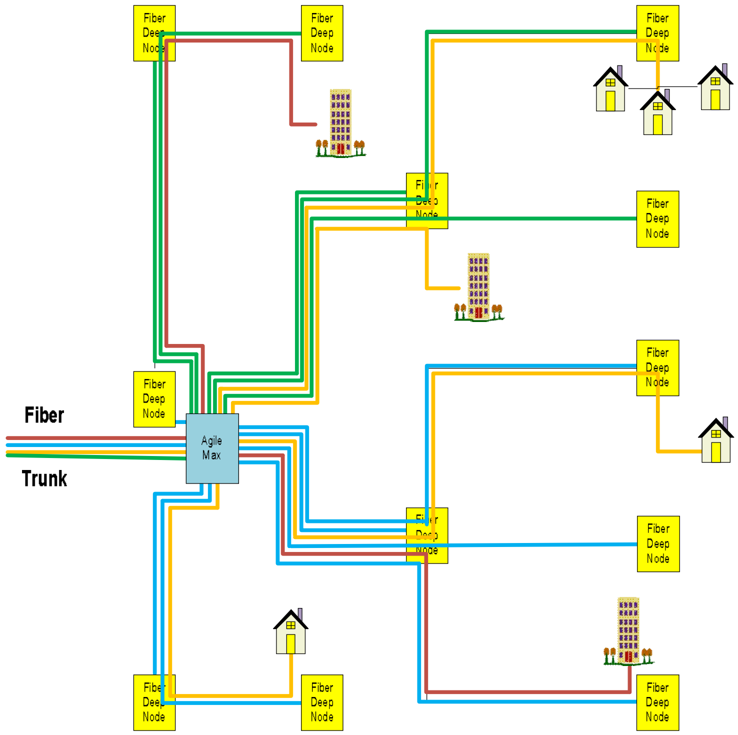

Figure 15 – HPON FTTP Migration: SG Segmentation as Needed

3.6.1.3. Step 3 – SG Segmentation as Needed

With this strategy, the operator only needs to deploy as many CCAP ports and Headend optics as is warranted based on subscriber demand; and then grow these over time as demand requires. This is exactly the DOCSIS philosophy.

3.6.2. HPON Topology Migration Example – Remote Devices

The HPON architecture has a primarily passive Outside Plant with its reduced operational expenses while maintaining a traditional centralized Headend architecture. An alternative approach is the Distributed Access Architectures (DAA) where intelligent devices such as Remote PHY, Remote MAC+PHY and/or Remote OLT are pushed out to the nodes in the plant. But HPON and Remote Devices are not mutually exclusive.

3.6.2.1. Remote Devices – Conventional Wisdom

A key motivation for Remote PHY Devices (RPD) and Remote MACPHY Devices (RMD) is the elimination of long analog AM optic fiber links enabling higher D3.1 capacities. The deployment of RPD/RMD is often considered with Fiber Deep upgrades as well.

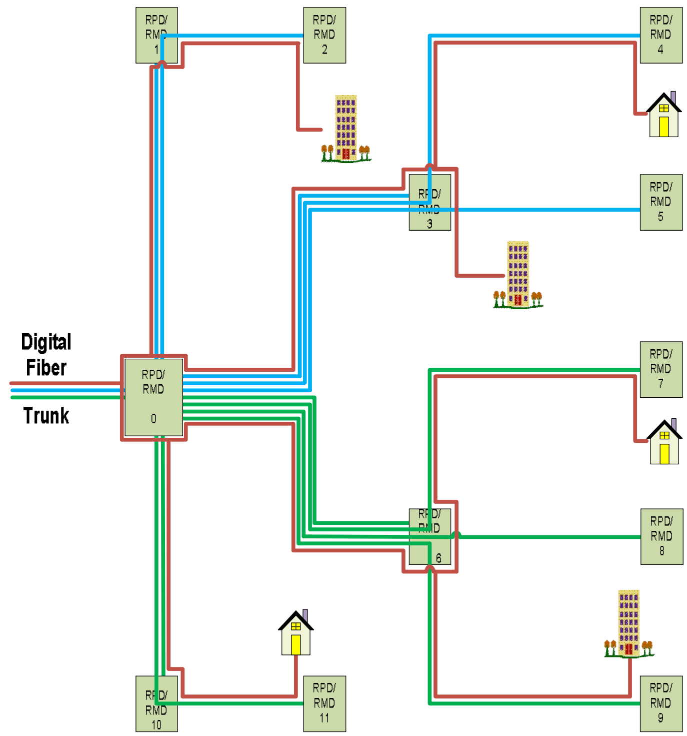

Figure 16 – Distributed Architecture: Conventional Wisdom

Conventional wisdom today places the RPD/RMD at the Fiber Deep Node location. Starting from the baseline example in Figure 12, the distributed architecture might appear as in Figure 16. Note that there are now 12 Remote Devices in the serving area, and each might only be serving ~60 HP or only ~30 subs. Based on previous traffic engineering results, Remote Devices will have excess capacity for another decade or two.

3.6.2.1. Shared Remote Devices using HPON

HPON enables an alternate distributed architecture with shared Remote Devices. This is shown in Figure 17. This figure shows that HPON migration steps 1 & 2 have been completed. Top customers have received FTTP and the Fiber Deep nodes are aggregated using HPON.

The difference with Figure 14 is that the previous connection from the HPON splitter to the Headend optics and CCAP over the Fiber Trunk has now been replaced with a connection to a single RPD/RMD remote device that is logically placed near the HPON splitter. The Remote Device must now contain short distance AM optic modules that support distances less than a kilometer. This particular example shows a single 2x2 RPD/RMD that can support 2 SG.

The key benefit here is that it only requires a single Remote Device compare to a dozen devices required in a typical distributed system shown in Figure 13. HPON ONU is significantly less complex than RPD/RMD devices which save the operators significant costs and power at every Fiber Deep node location.

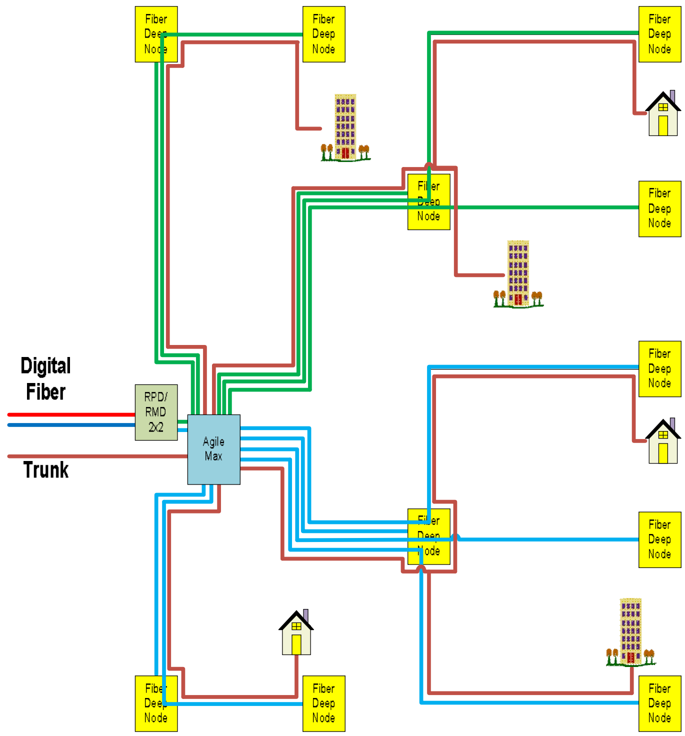

Figure 17 – HPON and Distributed Architecture: Shared Remote Device

3.6.2.1. Shared Remote Device using HPON: SG Segmentation

Eventually, the time will come where the SG size needs to be split. In a shared Remote Device scenario, the additional resources can be added at the same location as the original Remote Device. This device might be upgraded from a 1x1 or 2x2 RPD/RMD to a 4x4 or 6x6 or 8x8 device. Since this upgrade will occur many years in the future, this will be done with much newer technology thanks to Moore's Law and give the operator substantial cost and power savings per SG. This example is shown in Figure 18.

This approach does not preclude adding other Remote Devices in other locations. For example, maybe there is a neighborhood hotspot or an MDU that deserves its own Remote Device. Figure 18 shows an additional RPD/RMD being added for an MDU location.

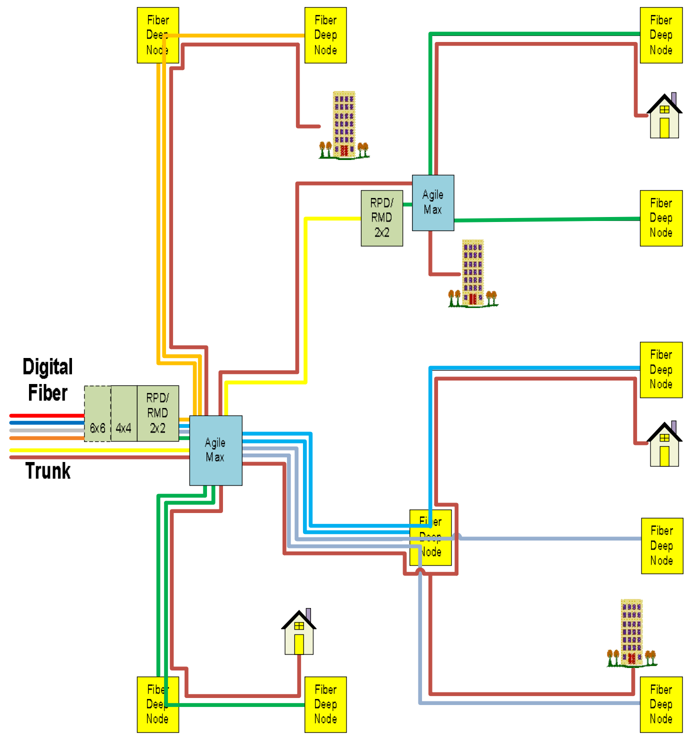

Figure 18 – HPON and Distributed Architecture: SG Segmentation as needed

3.7. DOCSIS 3.1 and RF Performance over HPON

To verify the potential of DOCSIS 3.1 over HPON, several lab measurements were done to analyze the RF US performance. Figure 19 below shows a set of MER curves for various RF spectrum loads on a first generation HPON splitter system over 20 km. For this system, a reverse transmitter was modified to go up to 1.2 GHz in the upstream, which then fed the HPON splitter. When looked at in the context of what DOCSIS 3.1 needs by way of SNR, one can easily see that a 200 MHz spectrum can easily support 4K QAM, with very good SNR for higher frequency spectral load.

When the above SNR graph is converted to the capacity available, as indicated in the Figure 20 below, it is seen that the capacity available is a monotonically increasing function of bandwidth, and at the 1.2 GHz upper limit, provides for almost 10 Gbps of upstream data throughput. At a more modest RF bandwidth of 200 MHz, the HPON system provides 2 Gbps of capacity. This is compared to SC-QAM technology @ 64-QAM which is what 3.0 uses today. By way of comparison, current 42 MHz DOCSIS 3.0 4 channel bonded system provide only 100 Mbps of throughput.

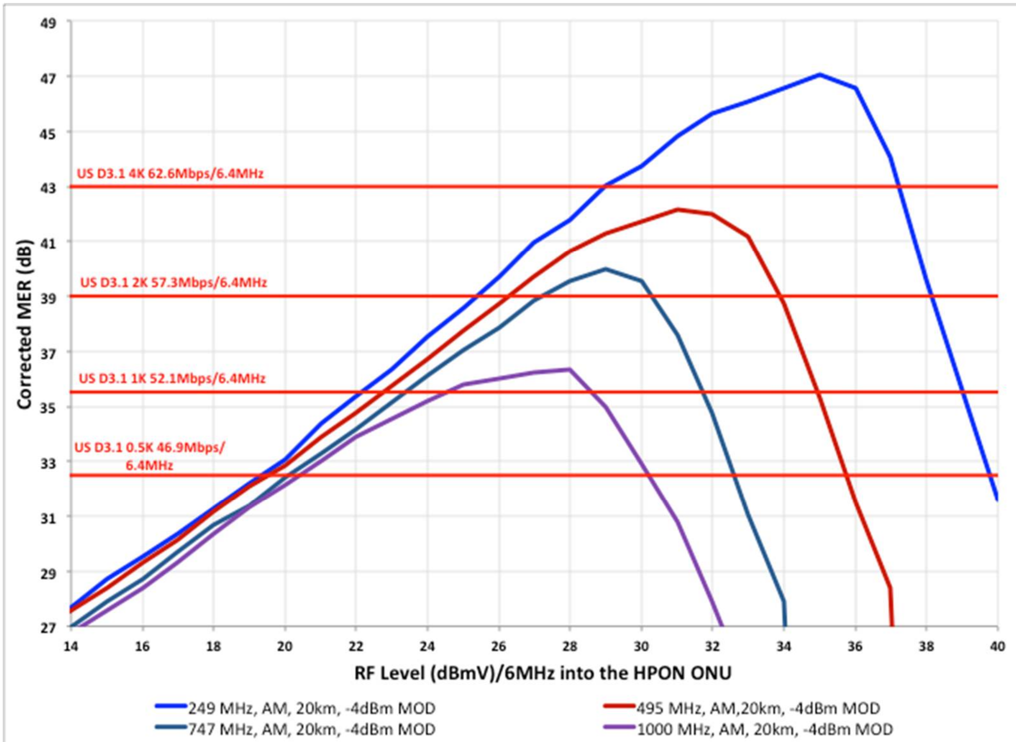

Figure 19 – HPON US MER for Various RF Spectrum

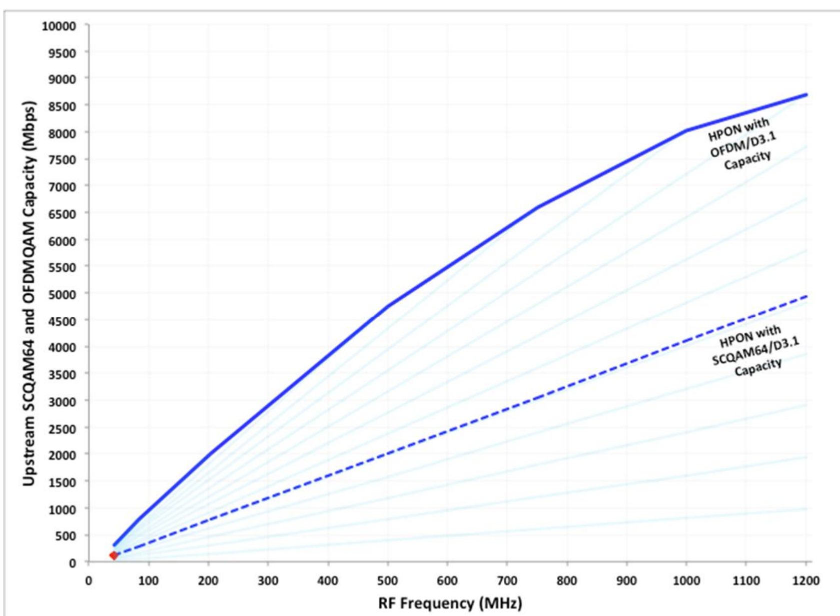

Figure 20 – HPON US Capacity – OFDMA vs. SC-QAM

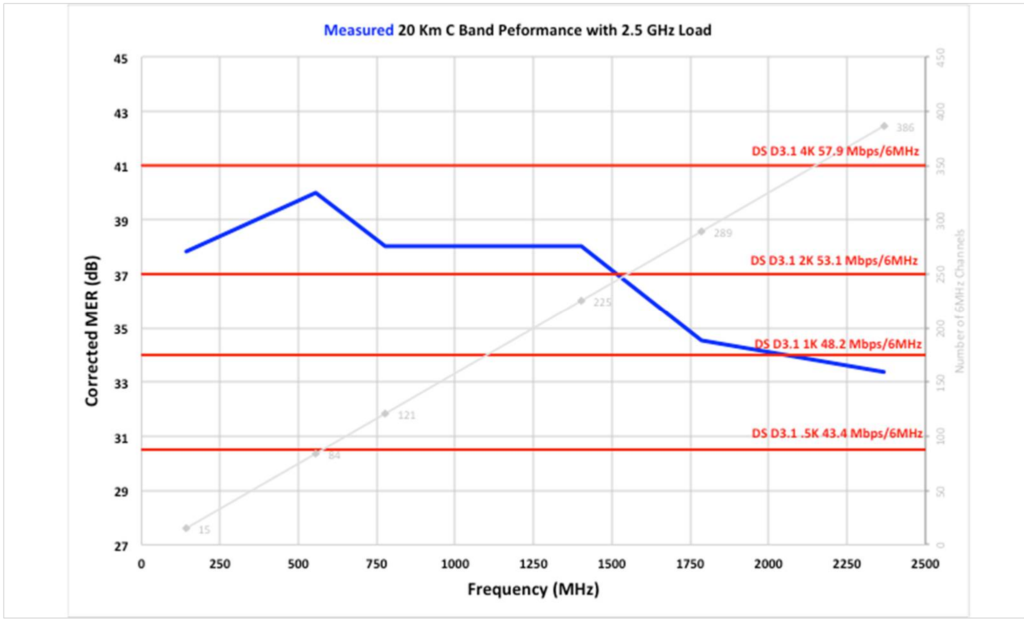

Figure 21 – HPON Downstream RF Performance

Figure 21 is focused on the downstream; a DML transmitter was modified to 2.5 GHz of spectral load and over 20 km of fiber, at the DS receiver produced a corrected MER that enabled 2K QAM for much of the spectrum and 1K and 0.5K QAM for the remaining portion.

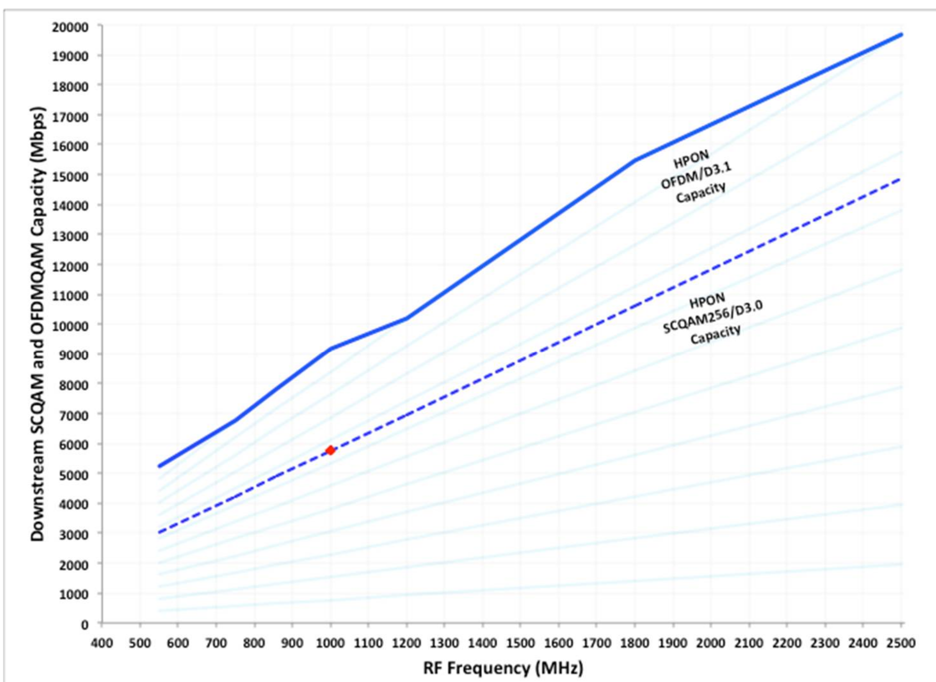

Figure 22 – HPON Downstream RF Capacity, OFDM v SC-QAM

When the downstream capacity is now computed for various spectral loading, one can see from Figure 22 that the HPON capacity can approach 20 Gbps for a 2.5 GHz spectral load. The figure also shows how OFDM capacity compares to SC-QAM channels @ 256-QAM.

3.8. D3.1 over HPON Summary

By enabling OBI-free DOCSIS 3.1 over HPON, a whole new world of options opens up to operators. HPON unleashes D3.1 capabilities to offer PON-like Gbps services in both upstream and downstream. It leverages the DOCSIS infrastructure making it very cost affordable for incremental investments for a gradual HFC to FTTP migration. It also opens up new potential HPON topologies such as FTTC, MDU and N+0.

Since DOCSIS supports large SG, it enables fiber and wavelength conservation in the plant and allows the CCAP port costs to be amortized over a larger number of users. Having significantly fewer CCAP ports also helps with headend space and power considerations.

10G & 1G EPON, GPON on HPON: Scaling Considerations for Residential Use Case

A better understanding of the benefits of DOCSIS 3.1 over HPON has raised similar questions on the scalability of traditional PON systems to handle larger SG, and in particular larger residential SG. Can EPON scale to hundreds of ONUs & thousands of LLID as well?

4.1. EPON Burst US Structure

To better understand PON upstream capacity, let's take a look at the EPON Burst upstream structure as shown in Figure 23. This shows the various overheads associated with each upstream transmit burst. Of particular note is the laser turn on and turn off times at the ONU, and the Automatic Gain Control (AGC) and Clock Data Recovery (CDR) times required by the OLT receiver. It turns out that for a 1G EPON upstream, the total burst overhead is in the range of 1.5 to 2.1 microseconds. This maps to an overhead of *188 to 264 bytes* for every transmit burst.

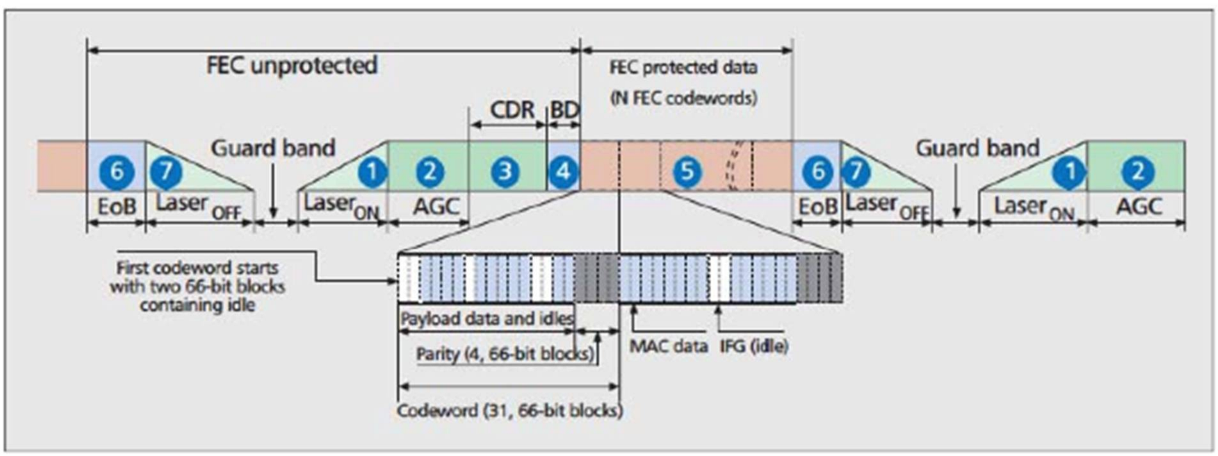

Figure 1. *Internal structure of the upstream channel data burst in 10G-EPON.*

Figure 23 – EPON Burst Transmit Overheads

As EPON evolved to its 10G upstream, the TX burst overheads was reduced, but not by a factor of ten. For a 10G upstream, the TX overhead may vary from 0.6 to 1.6 microseconds. This maps to an overhead of *764 to 2000 bytes* for every transmit burst.

A survey of industry literature by Glen Kramer, et al, uncovered [6, 7] that shows how the EPON upstream is impacted by the number of ONU and LLID and Grant Cycle Time. The Grant Cycle Time is the frequency of the OLT polling of each LLID in the ONU. This results in a 64 byte Report message being sent in the upstream. Tables 5 and 6 show some results from these papers for a 10G upstream.

ONUxLLID	1 ms	2 ms	4 ms
32	85.00%	86.05%	86.57%
64	82.91%	85.00%	86.05%
128	78.72%	82.91%	85.00%

Table 5 – 10G EPON Upstream Efficiencies

ONUxLLID	1 ms	2 ms	4 ms
32	8.47 Gbps	8.59 Gbps	8.65 Gbps
128	7.78 Gbps	8.24 Gbps	8.48 Gbps

Table 6 – 10G EPON Upstream Capacities

As can be seen for the parameters tested, efficiencies varied from ~79% to 87%. It is noted that the FEC accounts for 13% overhead. This means that the TX burst overhead varies from 0.5% to 9% based on these input parameters. This shows that EPON TX Burst overhead is very sensitive to ONU, LLID and Grant Cycle Time.

4.2. Extending 10G EPON Capacity Analysis

Looking at the previous traffic engineering for the Selective Subscriber Migration strategy, an operator might only need a SG size of 250 subs for the next 5-7 years. A large SG size would minimize OLT ports and fiber trunks required. However, each ONU might also have 4-8 LLID associated with it too. This implies that the product of ONU x LLID might go up to 1024.

The Kramer analysis was recreated and then ran it for a wider range of parameters. The results are shown in Figure 24. It shows a set of curves with different Cycle Times that fall off rapidly with increasing ONU x LLID. For example, 512 LLID (e.g. 64 ONU with 8 LLID each) with a 1 msec Cycle Time (needed for voice, gaming & MEF applications) has capacity of only ~5 Gbps.

Given this sensitivity to TX burst overhead, a closer look at the parameters was taken to determine a reasonable set for further testing. While an ONU might support 8-16 LLID, many will not be active and not require any polling. Based on DOCSIS experience, it seems that 4-5 active LLID per ONU would be reasonable.

The DBA scheduler in EPON also has the capability to poll each LLID at different intervals. Our analysis assumes that one LLID is needed for low latency applications with a 1 msec Cycle time, while another 4 LLID might have an average cycle time 4 msec. Since EPON allows multiple Reports per TX burst, our model assumes that there would be on average one TX burst per millisecond with an average of two Reports per TX burst.

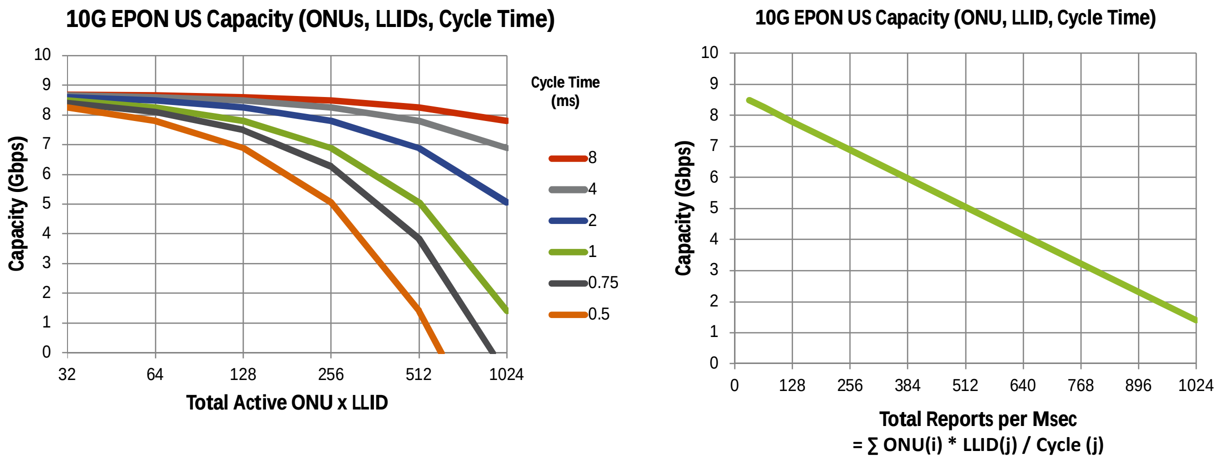

Figure 24 – 10G EPON Upstream Cycle Time & ONU Impacts

4.3. 1G EPON, 10G EPON and GPON Efficiency

10G EPON has a 10 Gbps downstream PHY rate, but supports two different upstream PHY rates: 1 Gbps and 10 Gbps. These are often referred to as 10/1 and 10/10 EPON. As a first step in our analysis, the control overhead efficiency is calculated and shown in Figure 25. The control overhead efficiency is basically the % of time available to transmit after the polling overhead. It excludes the FEC overhead for the 10G upstream. The efficiency is calculated for both 1G and 10G upstream, and for the min and max TX burst overhead. As can be seen, 1G US loses 28% to 36% capacity for 128 ONU while 10G US loss is in the 9-23% range for 128 ONU.

The chart also shows the efficiency for the GPON upstream. GPON is a synchronous system with only a 2 byte status report that is sampled every 125 microseconds. GPON efficiency is close to the 10G best case.

4.1. EPON: 1G and 10G Coexistence – Control Overhead Impacts on Efficiency

10G EPON supports the feature of simultaneously allowing 10/10 and 10/1 ONU to share the same OLT port. This is very desirable from an operator's perspective as they can deploy lower cost 10/1 ONU in asymmetric applications like residential while more expensive 10/10 ONU go to symmetric applications like business services. Other operators may decide to deploy cheaper 10/1 ONU today and then in the future deploy 10/10 ONU once they are more cost effective.

However, coexistence can have significant impact on upstream efficiency and capacity. Since the 10/10 and 10/1 share the same OLT port, only one can be transmitting at a given time. This is analogous to the 802.11 scenario where 11b and 11g WiFi devices coexisted in the same spectrum. The slower 11b devices took so much transmit time it left little capacity for 11g devices. 10G EPON concerns are potentially worse as the difference in speeds is now a factor of ten.

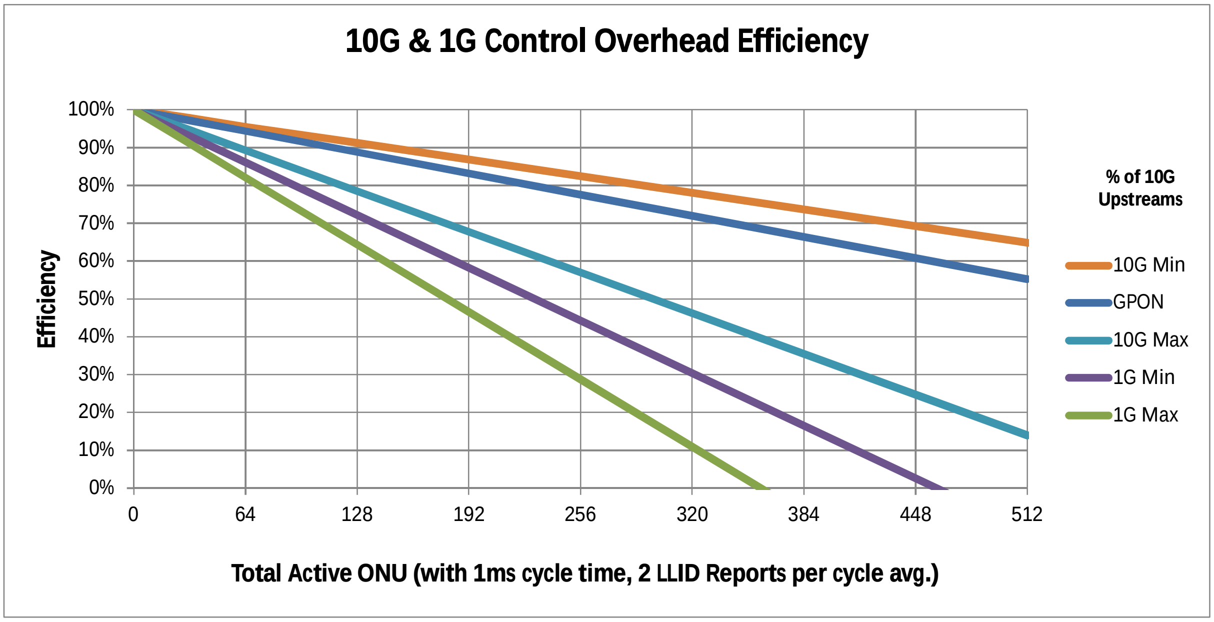

Figure 25 – 10G & 1G Control Overhead Efficiency

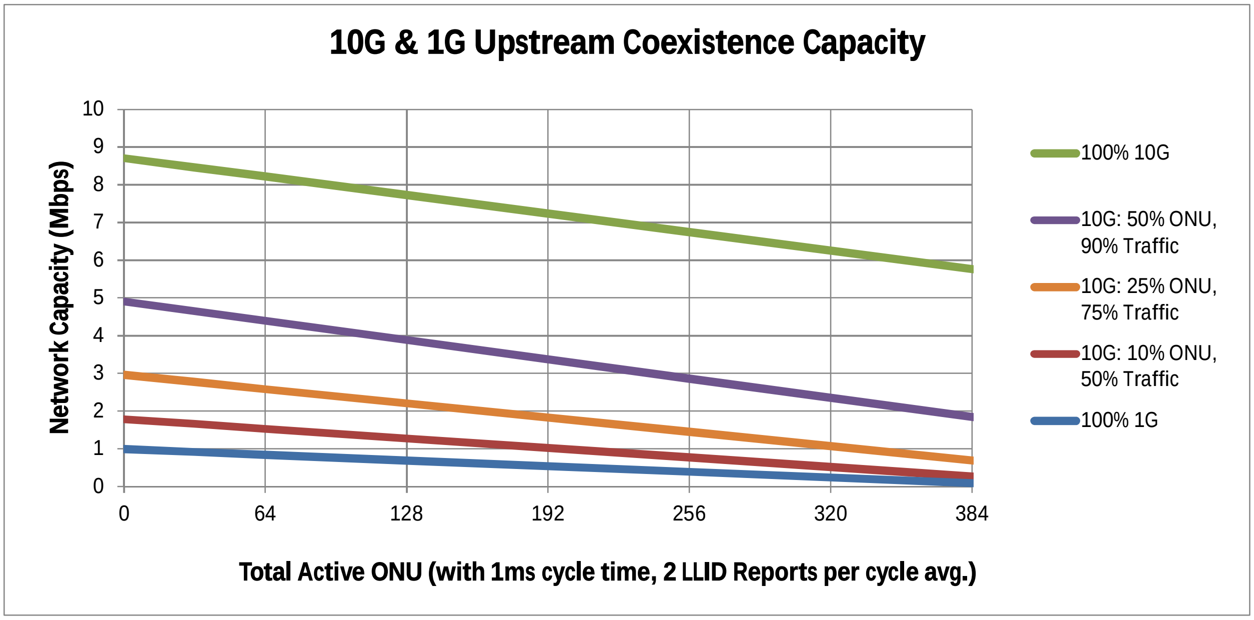

Figure 26 – 10G & 1G Upstream Coexistence Capacity

Analyzing the 10/10 and 10/1 coexistence scaling, there are two key factors that come into play. First, the control overhead is a function of the ONU mix. The efficiency becomes a blend dependent on the ratio of 10/10 ONU and 10/1 ONU. The second key factor then is the traffic mix between 10/10 and 10/1 ONU. It is assumed that 10/10 ONU will provide a higher upstream traffic load than 10/1 ONU.

With these factors in mind, three different scenarios were considered:

- Scenario 1: 50% of ONU are 10/10, 50% 10/1; Traffic Mix is 90% 10/10, 10% 10/1
- Scenario 2: 25% of ONU are 10/10, 75% 10/1; Traffic Mix is 75% 10/10, 25% 10/1
- Scenario 3: 10% of ONU are 10/10, 90% 10/1; Traffic Mix is 50/50

Figure 26 shows the results. For Scenario 1, network capacity is cut in half compared to 10G only US, even though 90% of the traffic is coming from a 10/10 ONU. For Scenario 2, network capacity is only one third compared to 10G only US. Finally in Scenario 3 where 50% of the traffic is coming from a 10/10 ONU, network capacity is less than 2 Gbps**,** marginally better than 1G only US.

4.2. 10/10 & 10/1 Coexistence compared to GPON and D3.1/HPON

With such significant degradation in capacity caused by 10/10 & 10/1 coexistence, it is useful to see how these scenarios fared when compared to GPON and to D3.1 over HPON. Figure 27 adds GPON to the mix and compares it with Scenarios 2 and 3 as well as 1G EPON US. As can be seen, GPON handles larger ONU count better than EPON. GPON capacity is competitive with these mixed 10/10 & 10/1 scenarios for large ONU counts.

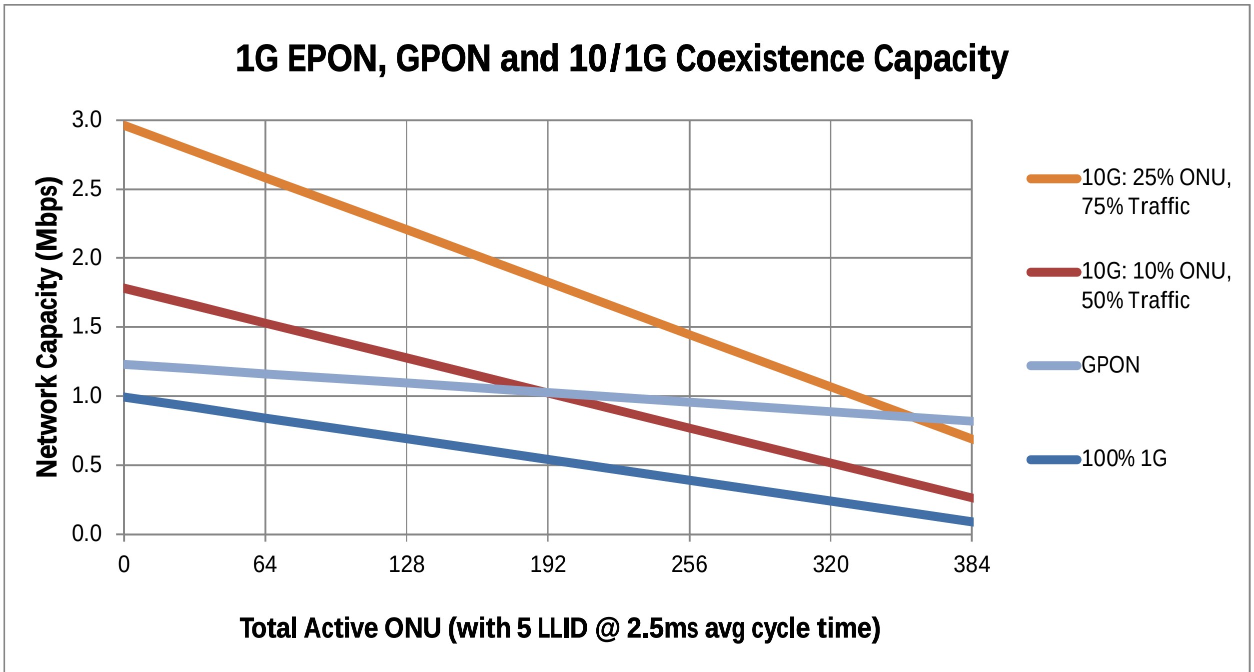

Figure 27 – 1G EPON, 10/10 & 10/1 Coexistence, and GPON

In Figure 28, DOCSIS 3.1 upstream capacity is added to the mix for both 85 MHz HFC and 204 MHz HPON networks. As can be seen, the D3.1 Network Capacity relatively independent of ONU count. D3.1/HPON outperforms 10G EPON Scenario 3 with 50% 10G US Traffic. D3.1/HPON is comparable to Scenario 2 (75% 10G US Traffic) for many ONU. D3.1 on 85 MHz HFC comparable to 1G US, especially for larger ONU counts.

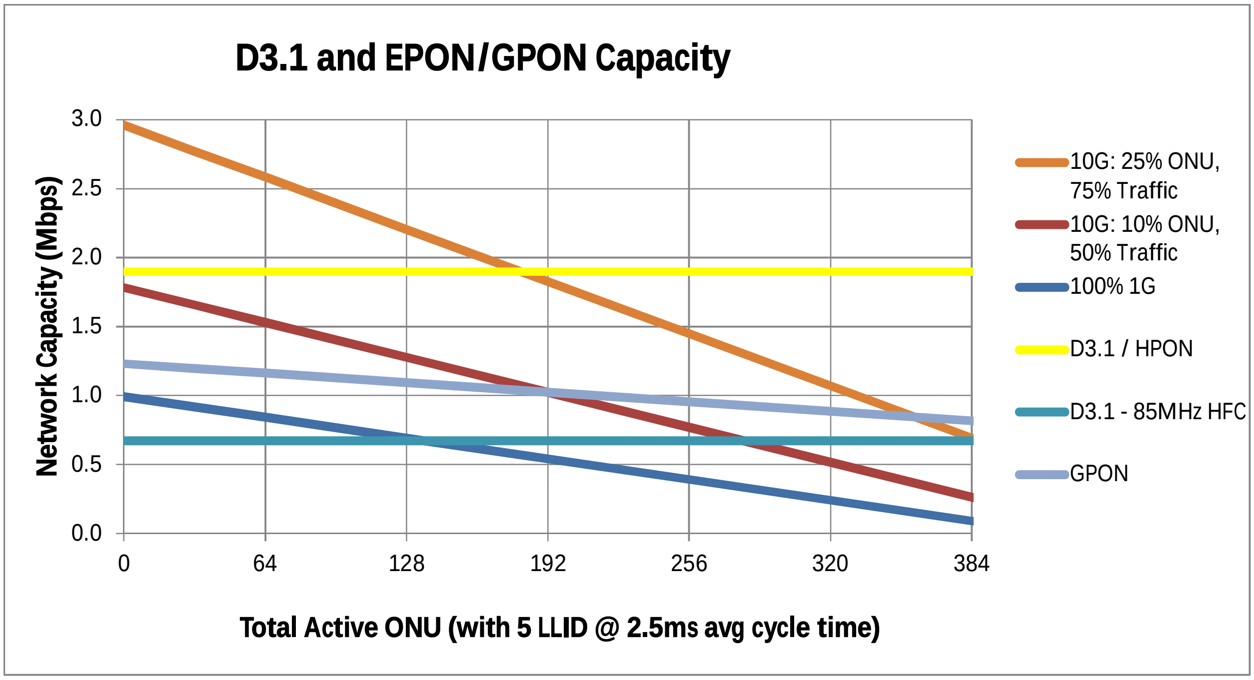

Figure 28 – D3.1 and EPON/GPON Capacity

4.3. Residential Applications present Traffic Engineering Challenges

As previously seen, 10G EPON has significant TX burst overheads, up to 764 to 2000 bytes. This means that the average TX burst needs to be sufficiently large to minimize the effect of this overhead. However, this may be problematic in a residential environment.

What is known about Residential Traffic usage today? Packet Size distribution is roughly 30% small packets (e.g. 64B), and 70% large packets (e.g. 1500B). There are a relatively small percentage of heavy users that account for majority of upstream traffic. Recent Sandvine data shows BitTorrent file sharing as the leading upstream application; with the remaining top applications related to real-time entertainment (e.g. Netflix, YouTube). Sandvine data also shows that traffic asymmetry actually increases during peak busy hours.

Here are some extrapolations from these observations:

- File sharing applications will be bursts of large packets from limited # of users with good probability of bursts of large packets together
- Real-time entertainment drives many small packets (e.g. IP Acks) from many users with little chance of bursts of more than a couple small packets together

Since video is driving the bandwidth growth engine, this traffic mix is not likely to change anytime soon. What is the impact of this packet distribution?

Figure 29 shows the average TX Burst size required for 100% utilization for upstream traffic load spread evenly across all ONU. Looking at 64 ONU with 100% 10G US, each ONU needs a 16KB average TX burst size each millisecond from every ONU to maintain 100% utilization. Figure 30 shows the average TX

Burst size for heavy users required for 100% utilization with a packet distribution based on the extrapolations above. It turns out that with 64 total ONU (100% 10G), of which 8 are heavy users, the heavy users need to have a 117KB average burst size every millisecond to maintain 100% utilization of the 10G upstream.

These results show that 10G EPON will need extremely large TX burst sizes in order to maintain its utilization, which becomes significantly worse when a packet distribution from a residential use case is factored in.

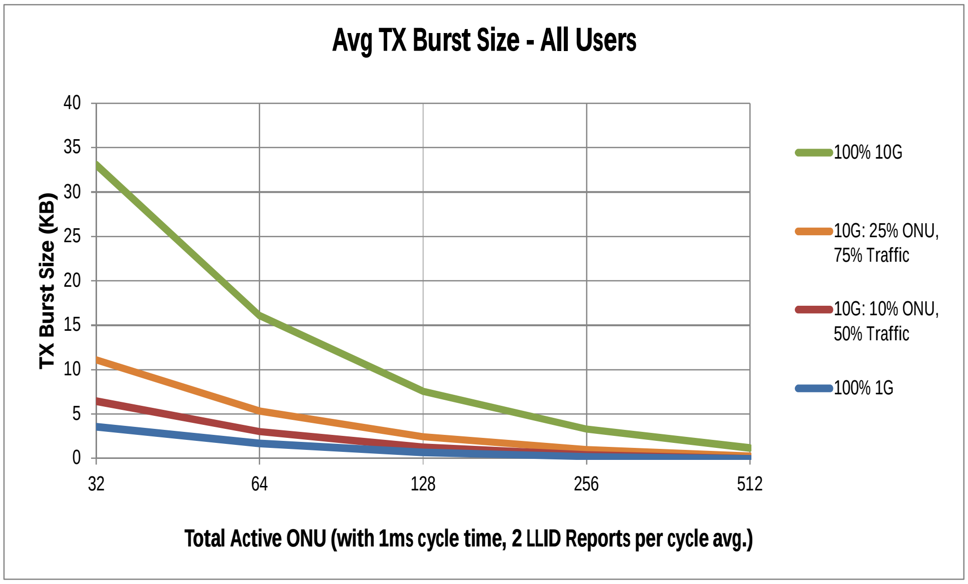

Figure 29 – Avg TX Burst Size, Traffic Evenly Distributed

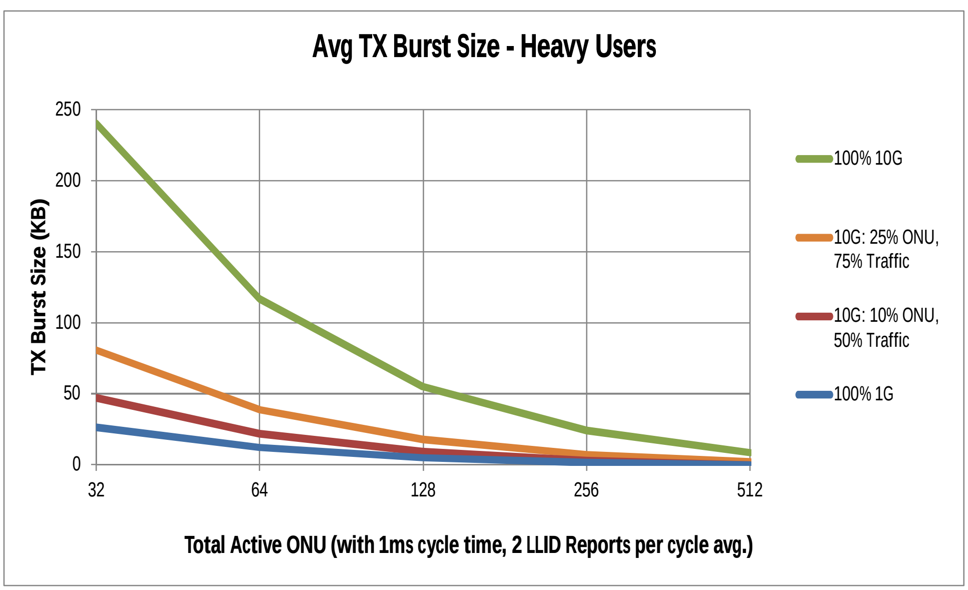

Figure 30 – Average TX Burst Size for Heavy Users

The role of FTTP and Hybrid PON – Other Considerations

5.1. Economics of HPON

A lot of the discussion so far has been on the capacity of an HPON system. It is also important to consider the economics of HPON. The following analysis includes total system costs including fiber deployment, ONU, and CCAP/OLT along with associated optics. The 1G EPON case was used as a baseline for a relative system cost comparison. The results are shown in Figure 31.

The top two curves compare 10/1 EPON costs to a D3.1/HPON FTTP costs. Both are assumed to have 1 user per ONU. Both are about 2½ times the baseline cost of 1G EPON. The D3.1/HPON costs are slightly less than 10/1 EPON as it can reuse existing HFC CCAP ports.

The bottom two curves show 1G EPON compared to D3.1/HPON FTTC costs with 4 user per ONU. They are very close in costs. The HPON FTTC approach generates significant savings by eliminating the need for a fiber drop to the end user and by sharing the cost of the ONU across 4 users. With HPON FTTC, an operator ends up with 10/1 EPON capacity at a cost of 1G EPON. This also highlights the HPON FTTC savings when compared to HPON or 10/1 EPON FTTP costs.

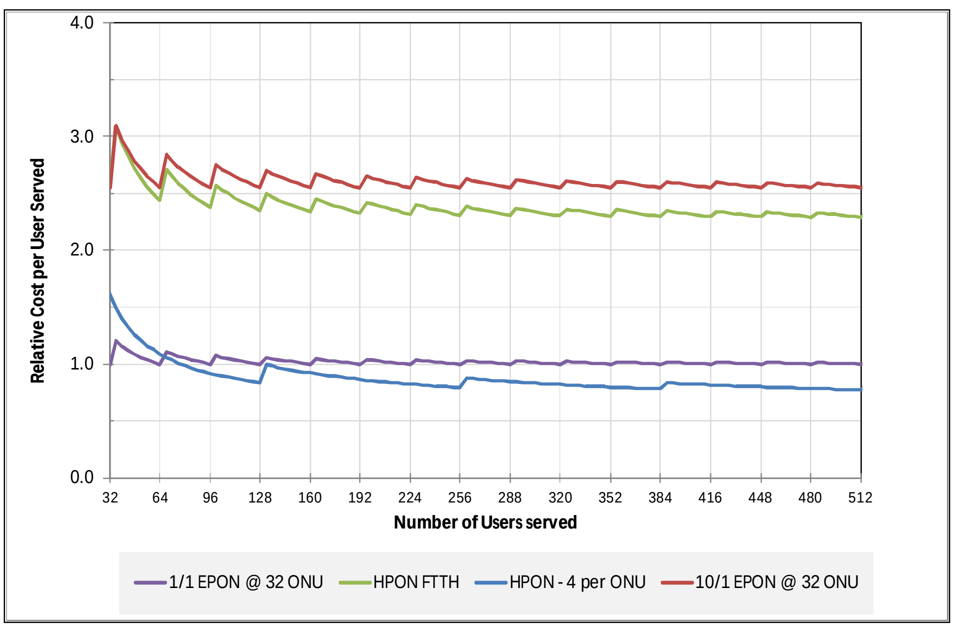

Figure 31 – Relative Costs for EPON & HPON Systems

5.1. Fiber Trunks, Wavelengths and OLT/CCAP Ports

For the HFC to FTTP migration, most operators will plan to reuse their existing fiber resources as much as possible and focus investment on pushing the fiber deeper towards the home. Many operators have limited fiber between their Headends and hubs to their serving areas, so both fiber count and wavelengths are critical resources. There are also Headend space and power considerations based on the number of OLT/CCAP ports required.

Figure 32 takes a look at the number of fiber trunks and/or wavelengths required, which also maps directly into the number of OLT/CCAP ports that are needed. A traditional PON system at maximum 20 km distances would typically have 32 users per SG/OLT port. This SG size is often limited by the fiber loss budget. For every 32 users in a serving area, another fiber trunk is needed as well as another OLT port. For 512 users in a serving area, the traditional PON system would need 16 fiber trunks and 16 OLT ports.

An alternative PON approach is to use a PON extender or Remote OLT technology. This will increase both the distance from the Headend as well as SG size. But the increasing SG size needs to be balanced against the capacity efficiency concerns discussed in the previous sections. Figure 32 assumes the extended PON can support 64 users per SG. This means a 512 user serving area would have 8 SG, need 8 wavelengths, and have 8 OLT ports.

D3.1 over HPON leverages the DOCSIS infrastructure and can support large SG. It might only need 1 or 2 SG for a serving area of 512 users. That means only 1 or 2 wavelengths and 1 or 2 CCAP ports are required. This saves the operator significant Headend space and power compared to PON approaches. At a later time when additional capacity is needed, then the SG can be split and additional CCAP ports and wavelengths added as needed.

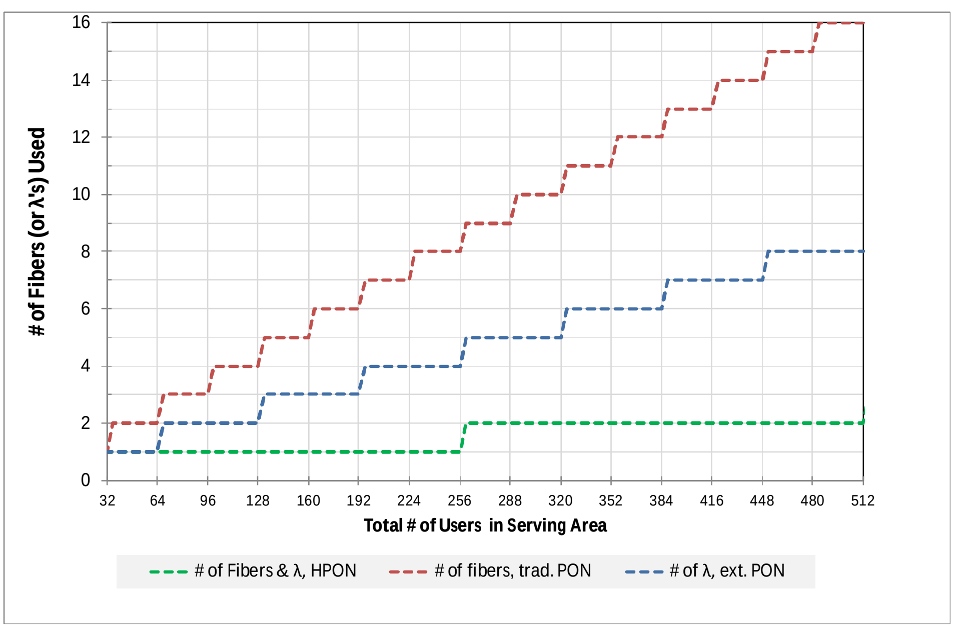

Figure 32 – Fiber Trunks, Wavelengths and OLT/CCAP Ports Required

5.2. HPON and Energy Considerations: Energy 2020 – Annual kiloWatt-hour/HP

In today's world, energy consumption is becoming increasingly more important. In reviewing the different architectures, power consumption is shifted between the Headend and the outside plant (OSP). For traditional PON, 100% of the operator's power is in the Headend, but in a distributed R-CCAP architecture, almost all of the power consumption is in the outside plant.

To be able to compare these different architectures, it is important to consider the total energy consumption as the key Metric. This must include BOTH outside plant AND Headend power impacts. Figure 33 takes a look the relative power consumption of various HFC, PON, and HPON alternatives. The power consumption is normalized on an annual cost per Homes Passed. The red portion of the bars represents the Headend power while the blue portion of the bars represent power consumption in the outside plant.

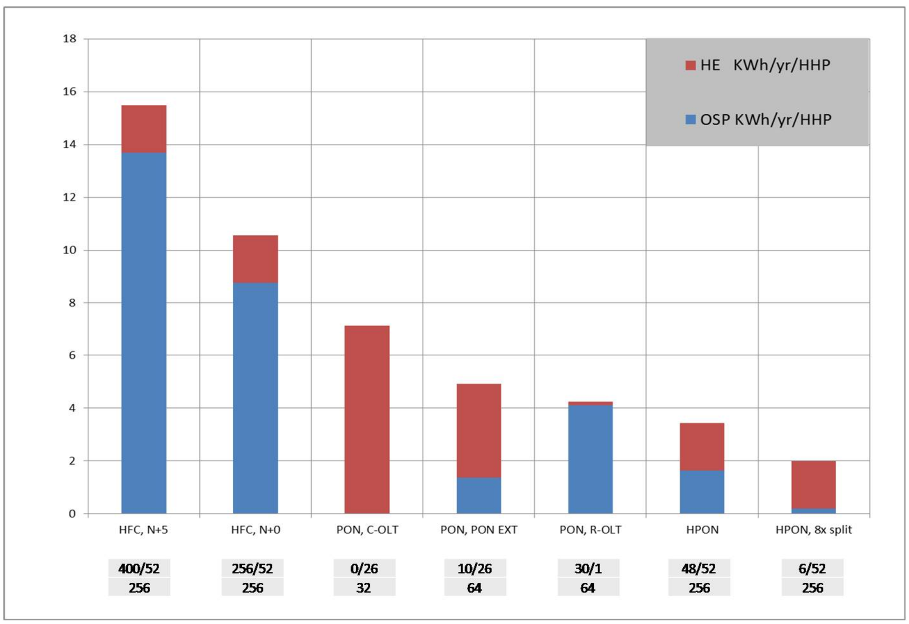

Figure 33 – Relative Energy Costs for HFC, EPON and HPON Systems

HFC systems are on the left in the figure. A typical N+5 system is the most power hungry of all the architectures. Most of the power is being consumed by amplifiers, actives and nodes in the outside plant. Next to that is the Fiber Deep N+0 HFC system. This reduces the N+5 power consumption by more than 25%, but is still high compared to the other alternatives.

The PON systems are next on the chart. The traditional PON has 100% of its power consumption in the Headend. It is about half of the H+5 HFC and 30% better than N+0. It is still relatively high because it is limited to 32 users per OLT port, requiring a large number of total OLT ports. A PON system with an extender continues to make improvements. By doubling the SG size to 64, the OLT port count and Headend power is cut in half. This is offset slightly with some additional OSP power for the PON extender. Finally, an estimate of a Remote OLT solution appears to provide the best total power consumption of the PON systems, but just marginally better than an extended PON.

Finally, the D3.1/HPON system is the two bars on the right. One is an HPON FTTP topology and the other is FTTC with 8 homes passed per ONU. The Headend power is the same as HFC, and leverages the fact that each CCAP port supports 256 users. For HPON FTTP, the OSP power consumption is close to the same as Headend power. Note that HPON FTTP power is less than 25% of the N+5 total power consumption and roughly half the total power consumption of a traditional PON system and better than Extended PON or Remote OLT systems.

The HPON FFTC solution is the most power efficient End-to-end (E2E) system. By sharing a single ONU across 8 homes, the OSP power consumption becomes negligible. The HPON FTTC system is the most cost effective from both a CAPEX and OPEX perspective.

Note that this analysis does not include the power for the ONU since that will often be powered at the premise.

HFC to FTTP Transformation: The Role of HPON – Conclusion

For existing plants, it has become clear that this is not a choice between HFC or FTTP. The transition is going to take many decades, so it is a matter of managing an on-going transformation. Our Network Capacity modeling has given us some insights into this. This has led to the Selective Subscriber Migration strategy to intelligently move Top Tier subscribers to FTTP. This approach can add decades to the life of HFC with 80% to 95% of all subscribers remaining put. It is also economically prudent, showing where and when is best to invest in outside plant.

A key piece of this strategy is the use of DOCSIS 3.1 on HFC. This can increase DOCSIS capacity by tenfold over 3.0 data rates. This is a critical element to make sure HFC remains useful through the FTTP transition period.

For the FTTP transition, it has long been thought that traditional PON was the only option. It turns out that a recent revolutionary breakthrough that completely eliminates Optical Beat Interference (OBI) has created a new option known as Hybrid PON, or HPON. HPON can simultaneously support traditional PON such as 10G EPON or GPON as well as OBI-Free DOCSIS 3.1 over HPON. This splitter based technology supports standard based components on either end of the network and is completely transparent. While EPON and GPON technologies are well known, the paper provided an in depth analysis of this new DOCSIS 3.1 over HPON option now available to operators.

HPON unleashes the capabilities of DOCSIS 3.1. Operating in a FTTP environment allows full use of the spectrum in both the upstream and downstream. Separate wavelengths allow spectrum overlap which enables the initial D3.1 modems to support 5 Gbps DS and 1.8 Gbps US, with higher data rates expected in the future. The downstream capacity of D3.1 over HPON is actually 33% more than 10G EPON. The 204 MHz upstream capacity is twice that of 1G EPON, 10/1 EPON and GPON. It enables operators to offer a true 1G upstream service which these other PON technologies do not.

By leveraging coax as a high performance drop cable, HPON also enables other fiber deep architectures besides FTTP. HPON supports Fiber to the Curb (or Tap), Fiber to the MDU (basement or floor) and even economical Fiber Deep nodes (N+0). An HPON architecture can also be used jointly with distributed architectures to provide the best of both worlds: a shared Remote Device to amortize cost and lowest cost and power Fiber Deep nodes.

Many of the advantages of using DOCSIS over HPON come from leveraging large SGs. Over the years, DOCSIS has been shown to scale nicely to many hundreds of modems and thousands of Service Flows. EPON efficiencies are very sensitive to the number of ONU, LLID and the Grant Cycle time. Given a reasonable number of LLID per ONU and Cycle times to support low latency applications, it will be hard to push an EPON system beyond 128 ONU.

10G EPON supports a coexistence mode that can support 10/10 and 10/1 ONU. While nice from an operational point of view, there are significant potential negative performance impacts. A scenario with 50% 10/10 ONU and 90% 10/10 ONU traffic will lose half its capacity to the slower 10/1 upstream. Another scenario with 10% 10/10 ONU and 50% ONU traffic gets less than 2 Gbps capacity, which is less than a 204 MHz D3.1/HPON system.

Economics and energy consumption are two key factors to be considered in determining the best solution path forward. In both cases, the HPON FTTC solution leads the way in both cost and power.

It turns out that an optimum solution for many operators is one that can simultaneously support a mix of both RFoG and PON over a shared Optical Distribution Network (ODN). This gives the operator total freedom to migrate subscribers between D3.1/RFoG and PON at their discretion as needs arise with minimal operational costs. They can always pick the best of breed of any technology.

Abbreviations

1588	IEEE 1588 Precision Timing Protocol (PTP)
1G	1 Gigabit
ABR	Adaptive Bit Rate
AGC	Automatic Gain Control
AM	Amplitude-Modulated
B	Bytes
Bcast	Broadcast
bps	Bits Per Second
CAGR	Compounded Annual Growth Rate
CCAP	Converged Cable Access Platform
CDR	Clock Data Recovery
CMTS	Cable Modem Termination System
COTS	Commercial Off The Shelf
D3.1	Data Over Cable Service Interface Specification 3.1
DAA	Distributed Access Architecture
DBA	Dynamic Bandwidth Allocation
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
E2E	End to end
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
FEC	Forward error correction
FTTC	Fiber to the Curb
FTTP	Fiber to the Premise
Gbps	Gigabits Per Second
GCP	Generic Control Protocol
GHz	Gigahertz
GPON	Gigabit-Passive Optical Network
HFC	Hybrid Fiber-Coax

HP	Homes Passed
HPON	Hybrid Passive Optical Network
I-CCAP	Integrated Converged Cable Access Platform
IEEE	Institute of Electrical and Electronics Engineers
KB	Kilobyte
L2/L3	Layer 2 and Layer 3
LDPC	Low Density Parity Check
LLID	Logical Link Identifier
MAC	Media Access Control interface
MACPHY	DCA instantiation that places both MAC & PHY in the Node
Mbps	Mega Bits Per Second
MDU	Multiple Dwelling Unit
MHz	Megahertz
MSO	Multiple System Operator
Ncast	Narrowcast
OBI	Optical Beat Interference
ODN	Optical Distribution Network
OFDM	Orthogonal Frequency Division Multiplexing
ONU	Optical Network Unit
OOB	Out Of Band
PHY	Physical interface
PON	Passive Optical Network
PTP	Precision Timing Protocol
QoS	Quality of Service
RF	Radio frequency
RFoG	RF Over Glass
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
TaFDM	Time and Frequency Division Multiplexing
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
US	Upstream
W	Watt

Bibliography & References

1. *"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 Shemmann, Zoran Maricevic, John Ulm; 2015 INTX NCTA Spring Technical Forum
2. "*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
3. "*A Side-by-side Comparison of Centralized vs Distributed Access Architectures"*, Mike Emmendorfer et. al., 2014 SCTE Spring Technical Forum, SCTE.
4. *Data-Over-Cable Service Interface Specifications DOCSIS 3.1* Physical Layer Specification; CM-SP-PHYv3.1-I05-150326, CableLabs.
5. *"Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning"*, John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo.
6. *Data-Over-Cable Service Interface Specifications DCA - Remote MACPHY; Remote MAC-PHY Technical Report* CM-TR-R-MACPHY-V01-150730, CableLabs.
7. "*Performance of 10G-EPON*"; Glen Kramer, Marek Hajduczenia, Rajesh Roy, H. J. Silva; IEEE Communications Magazine, Nov 2011.
8. "*How Efficient is EPON?";* Glen Kramer, Teknovus White Paper
9. "*On Configuring Logical Links in EPON*"; Glen Kramer, Teknovus White Paper
10. *"From OBI and SNR to Ookla and Netflix: How Network Impairments affect Customer Perceptions: The role of Leading and Lagging Indicators as We Evolve HFC to FTTP";* Sebnem Zorlu Ozer, Ph.D., Venk Mutalik, A. Vieira, Ph.D., J. Chrostowski; SCTE Cable-Tec Expo 2015