GIVING HFC A GREEN THUMB

A CASE STUDY ON ACCESS NETWORK AND HEADEND ENERGY & SPACE CONSIDERATIONS FOR TODAY & FUTURE ARCHITECTURES

JOHN ULM
ZORAN MARICEVIC

ARRIS™

TABLE OF CONTENTS

INTRODUCTION

With the SCTE Energy 2020 initiative in full swing, the cable industry is seeing vigorous interest in getting a handle on its energy consumption. This paper provides a case study on the combined energy consumption for both the headend facility equipment and the access network plant to understand the total impact of various network access architecture options.

After reviewing the network capacity planning for the next decade, the paper takes a look at a baseline case study of five different actual physical nodes and analyzes several possible HFC upgrade options and their relative power consumption. The upgrade options considered include fiber deep architectures such as Node+0 which includes Fiber to the Last Active (FTTLA).

This is then followed by a space and power analysis of some existing headends with older CMTS and Edge QAMs. The headend facility savings are shown from introducing a CCAP chassis and from including the benefits of integrating all of the narrowcast EQAM into the I-CCAP box.

In addition to this “business as usual” progression, there are several new architectures being considered for the near future. These potential new architectures include:

1. Remote PHY and Remote MACPHY
2. EPON FTTH (centralized OLT with and without PON Extenders; Remote OLT)
3. Distributed Node Architecture solutions

The distributed access technologies can significantly reduce headend energy consumption but push complexity into the plant and have a negative energy impact there. FTTH solutions can offer a completely passive outside plant but could require an increase in energy consumption in the headend and at the consumer premise. It is important to consider both headend and plant energy together to get the total picture on energy consumption.

Our paper takes a look at the space and power impacts of these various architectures and provides the operator with some guidance with regards to total energy consumption in selecting between these options. A companion paper [ULM_2016] takes a look at the economic considerations for several of these architectures.

NETWORK CAPACITY – PLANNING FOR THE NEXT DECADE

The Internet has been growing at a breakneck speed since its inception. And with it, we have seen a corresponding growth in dedicated network capacity. While Moore's Law is infamous in silicon realms, Nielsen's Law of Internet Bandwidth has become renown in the networking world. It basically states that network connection speeds for high-end home users would increase 50% per year. This law has driven much of the traffic engineering and network capacity planning in the service provider world. It has also led to much research on those topics.

Nielsen's Law and Cloonan's Curves

In [CLOONAN_2014, EMM_2014], this research was expanded to also include traffic utilization in addition to the network connection speed. In his chart below, known as Cloonan's Curves, Nielsen's Law is represented by the blue line in the middle. Since it is a log scale, the 50% Compounded Annual Growth Rate (CAGR) appears as a straight line. An interesting fact is that the graph starts in 1982 with a 300-baud phone modem. We are now in the fourth decade of closely following this trend.

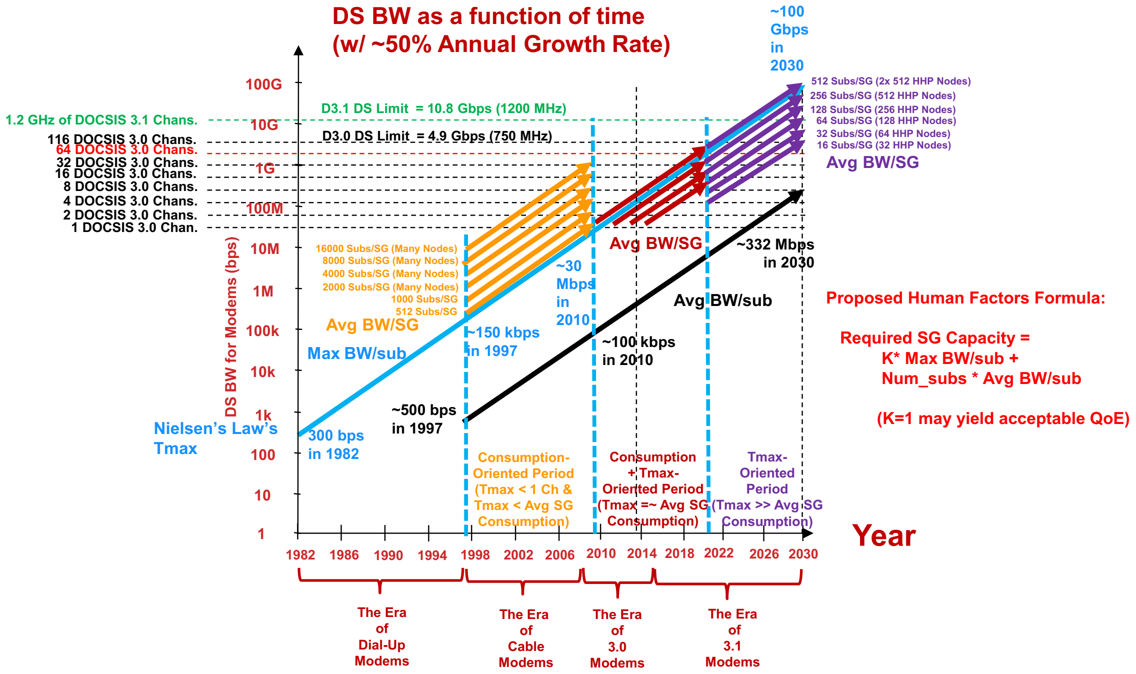

Figure 1 – Cloonan's Curves

Cloonan noted that the primetime average subscriber consumption (a.k.a. Tavg) has also been following this same basic trend as shown in the Figure 1. For service providers, an important metric is the traffic utilization in a Service Group (SG). The SG traffic utilization is a function of the number of subscribers (Nsub) times the average bandwidth per sub (Tavg) and is shown in a series of lines above Nielsen's line.

In the early DOCSIS days, many nodes were combined together and a SG might consist of thousands of subscribers. At this time, the SG traffic was an order of magnitude higher than the maximum network connection speed (a.k.a. Tmax after the DOCSIS parameter that dictates max network rates). Over time, the SG size has been shrinking and with it the ratio between Nsub*Tavg to Tmax. As shown in the chart above, the SG traffic eventually approaches that of Tmax. As SG sizes dip below 100 subs, then Tmax starts to dominate the traffic engineering.

We have been monitoring subscriber usage for many years now. The chart below shows Tavg, the average subscriber downstream consumption during peak busy hours, for a number of MSOs over the last six years. At the start of 2016, Tavg was approximately 850 Kbps. Over this six year period, Tavg has grown at ~45% CAGR. We are expecting that Tavg will break the 1 Mbps barrier sometime in 2016. The chart also maps out Tavg growth through the year 2020 assuming a 45% CAGR.

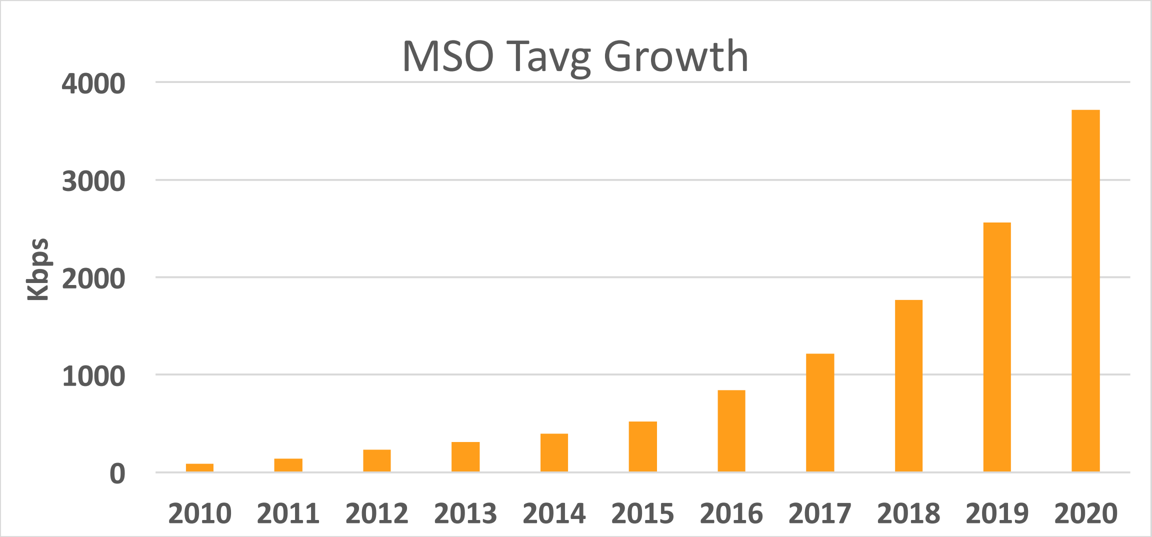

Figure 2 – Tavg, Average Subscriber Consumption

Interestingly, the upstream traffic is growing at a significantly slower rate. During the same six year interval, the upstream Tavg only grew at ~20% CAGR. The industry is seeing more asymmetric traffic with video being the driving application for downstream consumption [see EMM_2014]. At this point, there is about a ten to one ratio in traffic and still expanding.

Selective Subscriber Migration Strategy

As operators approach capacity planning, they are trying to understand how long the HFC architecture might last before they must migrate to a Fiber to the Premise (FTTP) network. To get an insight into this, the chart below zooms in on the Cloonan's Curve & Nielsen's Law over the next two decades. It predicts that top network speeds will reach 10 Gbps by ~2024 and pass 100 Gbps in the early 2030's. The initial DOCSIS 3.1 (D3.1) goal was 10 Gbps, so that implies that the HFC may hit its ceiling by approximately 2024!

At first glance, this is a scary proposition in that HFC networks might be obsolete in 5-7 years while it may take decades to build out an FTTP infrastructure. However, this is not the full story. As was shown in [ULM_2014], Nielsen's Law applies to the top speed tier which is only a very small percentage of the entire subscriber base, perhaps less than 1%. So the key question then becomes, "What happens to the vast majority of subscribers on HFC who are not in the top speed tiers (a.k.a. billboard tiers) and when?"

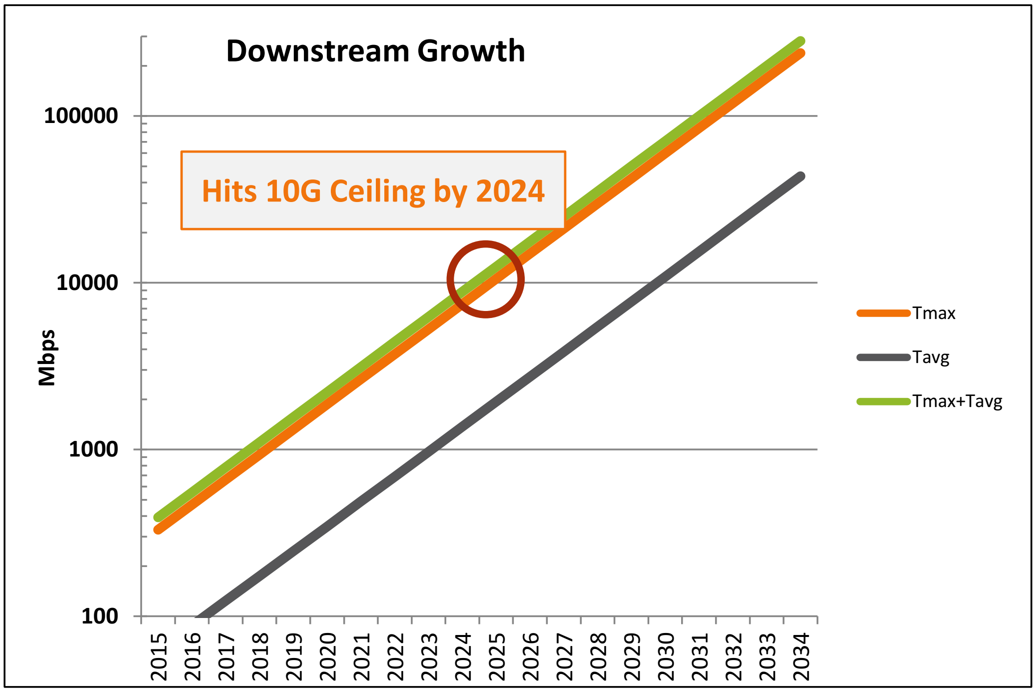

Figure 3 – Downstream Growth over Next Two Decades

The [ULM_2014] case study took a look at service tier evolution at a few MSOs. Table 1 lays out results from that study. Perhaps the key finding from this study is that the different service tiers are growing at different rates. While the top billboard tier continues to follow Nielsen's Law 50%, each subsequent lower speed tier is growing at a slower rate. Hence, the lower the service tier rate, the lower its CAGR.

Table 1 – MSO Case Study on Multiple Service Tier Levels

2014 Service Tier Levels on HFC	% 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%

Figure 4 maps out the various service tier growth over the next two decades. While the 1% of subs in the top billboard tier hit 10 Gbps in ~2024, the 14% of subs in the performance tier don't hit that mark until ~2032. Notice that 85% of subscribers in the flagship basic tier and economy tier stay below this mark for several decades.

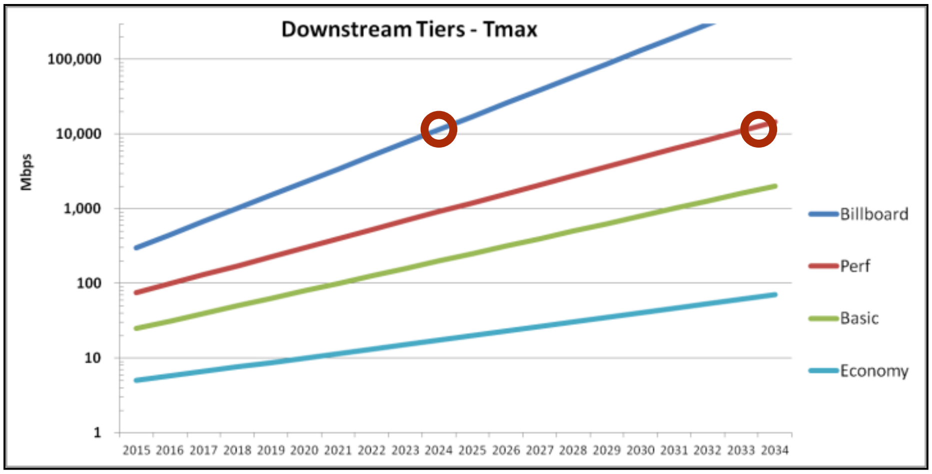

Figure 4 – Downstream Growth with Multiple Service Tiers

Data was input into the ARRIS Network Capacity model to take a closer look at the network traffic growth. Table 2 shows the Tmax migration used for each tier level over the next decade. Note that by 2021, the top billboard tier starts to exceed the capacity of the initial D3.1 modems that are being used today. And by 2026, this tier is forecast to hit 40 Gbps. This will require new technology, which might be a newer generation of DOCSIS (e.g. Extended Spectrum) or possibly a next generation of PON technology (e.g. 100G EPON).

Table 2 – Service Tier Migration for Network Capacity Model

MSO Case Study DS Service Tiers	% of Subs	Tmax CAGR	2014	2016	2021	2026
Top Billboard Tier	<1%	50%	300	675	5G	40G
Performance Tier	14%	32%	75	125	500	2G
Basic Tier	65%	26%	25	40	150	400
Economy Tier	20%	15%	5	10	20	50

It is important to note that 99% of the subscribers are still comfortably using today's DOCSIS technology on HFC a decade from now.

Some results from the ARRIS Network Capacity model are shown in Figure 5. It provides an insight into both Tmax and SG Tavg behavior. During the next 5-7 years, the Tmax component dominates traffic engineering as it is driven by Nielsen's Law. The bandwidth needed by the top billboard tier dominates compared to the SG Tavg.

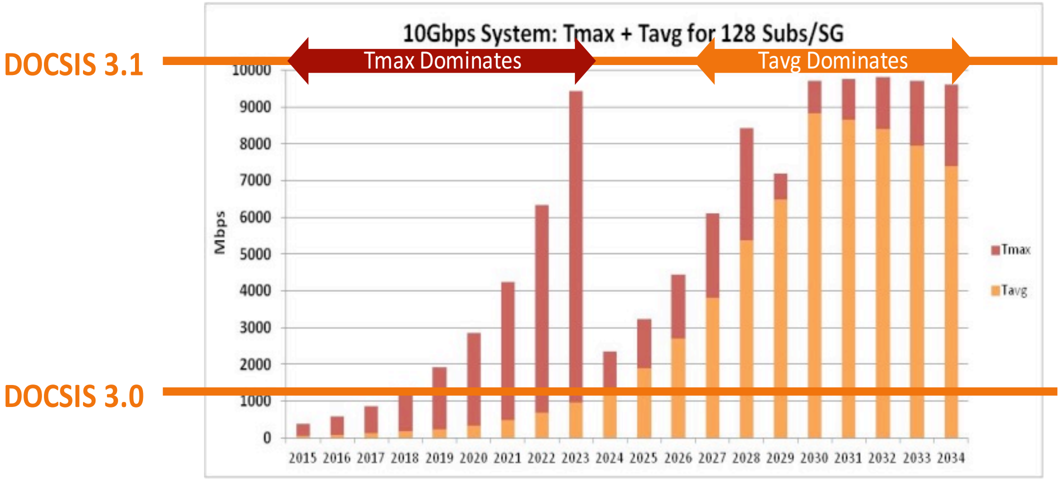

Figure 5 – Network Capacity Model Results

This leads us to a Selective Subscriber Migration strategy that will need to start in the next 5-8 years. By moving the top billboard tier to a Fiber Deep access network that is separate from the general HFC plant, there is a significant reduction in the required DOCSIS capacity. This reduction can be seen in year 2024, in Figure 5, after the top billboard tier is removed from the HFC network. The performance tier is then moved in 2029, in this example, for a smaller drop.

Note that the Fiber Deep access network might be any one of several FTTx options including: FTTP, Fiber to the Curb (FTTC), Fiber to the Tap (FTTT), Fiber to the Last Active (FTTLA), or Node+0 HFC. These options are discussed in detail in the next section.

Eventually, with the top tiers migrated to FTTx, the SG Tavg finally catches up and operators will need to consider reducing SG sizes again. The model in this example predicts that this will be roughly 10-15 years from now.

Another observation from this analysis is that D3.1 is a key technology to extend HFC life for decades to come, especially for the vast majority (e.g. 65-95%) that are in the flagship basic and economy tiers. Any brownfield FTTx transition may take decades, so D3.1 successfully gets operators through that window.

In summary, Selective Subscriber Migration strategy is a sensible approach to the topic of an HFC to FTTx transition. Moving top tiers to FTTx can buy HFC extra decades for 80-95% of subscribers in the flagship basic/economy tiers. Tmax dominates for the next 5-7 years, so it is more important to increase the HFC capacity to at least 1 GHz spectrum rather than split nodes. However, Tavg finally catches up 8-10+ years from now; and SG size reductions come back into vogue. Operators should push Fiber Deep enough to enable Selective FTTx for top tiers on demand and be prepared for the next round of SG splits.

And which FTTx is the best option is another interesting debate. DOCSIS continues to evolve with work on Full Duplex (FDX) and Extended Spectrum DOCSIS. Some of this research was highlighted in [CLOONAN_2016]. These new technologies promise to do for DOCSIS & cable what G.fast is attempting to do for DSL and twisted pair. Figure 6 shows some results from that paper for both FTTC and FTTLA systems. As can be seen, the system capacity can increase significantly as fiber is pushed closer to the premise.

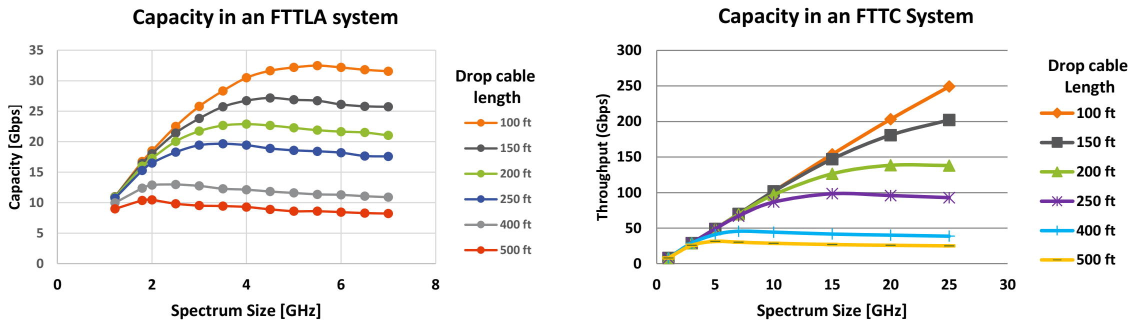

Figure 6 – Network Capacity Model Results

ACCESS NETWORK CASE STUDY

The network capacity planning shows that operators will need to evolve their existing Hybrid Fiber Coax (HFC) networks to remain competitive with FTTP service providers such as Google Fiber and Verizon FiOS [VENK_2016, VENK_2015 and ULM_2015]. For cable operators, they can utilize their existing fiber investments as a starting point to get a jump start compared to new entrants that must start their fiber installation from scratch. But the critical question for cable operators is how deep should they pull the fiber? They are presented with a toolbox of architectural choices to consider:

- "Business as usual" (BAU) – a node split where needed, and a refresh of the HFC field actives, with perhaps an upgrade to 5-85 MHz in the return and 104-1002 MHz in the forward
- Fiber Deep (FD) Node+0 (N+0) pushes fiber much deeper into the HFC and eliminates all of the active RF elements. There is an array of potential options including:
    - Traditional Fiber Deep Node+0 "FD N+0" which redesigns existing HFC (e.g. N+3 to N+6 with 3-6 actives after the fiber node) into "node as the last active". The typical way to do this is to rewire the coax plant in a way to minimize how many of these standard-size new nodes need to be added. Each new node may ultimately become its own service group, and in addition to the RF and optical modules, it may house Remote PHY Devices (RPDs) and PON OLTs
    - Fiber to the last active (FTTLA) is a variant of the Fiber Deep N+0 architecture. However, in this case the nodes are located precisely at legacy RF amp locations. These nodes then get aggregated into a properly-sized service group. This aggregation can be done by using an "active splitter / combiner", housed in a virtual hub, which is located precisely at the legacy node location to save on optics costs & space in the facility
    - Fiber to the curb (FTTC) or Fiber to the tap (FTTT) where fiber is run down the street but the existing cable drop cables are reused
- Fiber to the Premise (FTTP) – this is what is being deployed today with traditional PON systems as well as RFoG system

Collectively, these fiber deeper options are referred to as FTTx or Fiber to the "x", where "x" might be Premise, Curb, Tap, Last Active, or Fiber Deep node. For cable operators to build out any of the above architectures in today's brownfields, the new fiber construction begins from an existing fiber node; unlike the new entrants who must build the fiber construction from the central office / headend.

Each MSO will make changes to their own HFC plant to optimize for the attributes that they deem to be the most important. Different MSOs will likely prioritize the many attributes in different ways. For example, some MSOs may choose to optimize their network evolution by moving as rapidly as possible to end-state technologies of the future. These MSOs will likely move rapidly towards (passive optical network) PON or Point-to-Point Ethernet solutions. Other MSOs will choose to optimize their network evolution to reduce headend power and rack-space requirements by moving towards Fiber Deep architectures with Distributed Access Architecture sub-systems that remove functionality from the headend. These MSOs will likely deploy (Remote PHY) RPHY or (Remote MACPHY) RMACPHY sub-systems within their nodes. Other MSOs will want to preserve much of their current architectures while capitalizing on improved technologies.

In order to calibrate our conceptual thinking against reality, a set of five real-life HFC nodes was identified for evaluation, representing a diversity of implementations. These are representative of low, medium, and high densities, as measured by how many homes are passed per mile in each area. The five node areas, labeled A, B, C, D, & E possess other attributes of interest: miles of hardline coax plant, percentage of aerial plant, number of RF actives, number of homes passed per node, and HP/mile, as shown in Table 3.

Figure 7 shows the topology of one of the nodes: Node C. The headend (upper left) is fiber-linked to the node (center-left in pink), which RF-feeds into RF amps (blue triangles) RF splitters (blue circles), and taps (orange diamonds). Two 15A field power supplies provide enough power for the whole node area. Node C contains 3.5 miles of coax plant (excluding drop cables) with 21 actives and 398 Homes Passed (HP). So this might represent ~200 subscribers @ 50% penetration.

Node C will be used as a baseline example to show how the other architectures might be implemented.

Table 3 – Properties of 5 Node Areas Under Study

Node	**A**	**B**	**C**	**D**	**E**	**Overall**	**Average**
Plant Coax Mileage	4.2	6.2	3.5	2.5	1.9	18.3	3.7
% Aerial	20%	77%	97%	87%	91%	70%	70%
Total Active	21	30	21	19	14	105	21
Actives/Mile	5.0	4.9	5.9	7.6	7.4	5.7	5.7
Cascade Depth	N+3	N+3	N+3	N+3	N+2		N+3
Total Homes Passed	153	352	398	469	520	1892	378
HP/Mile	37	57	112	187	274	104	104

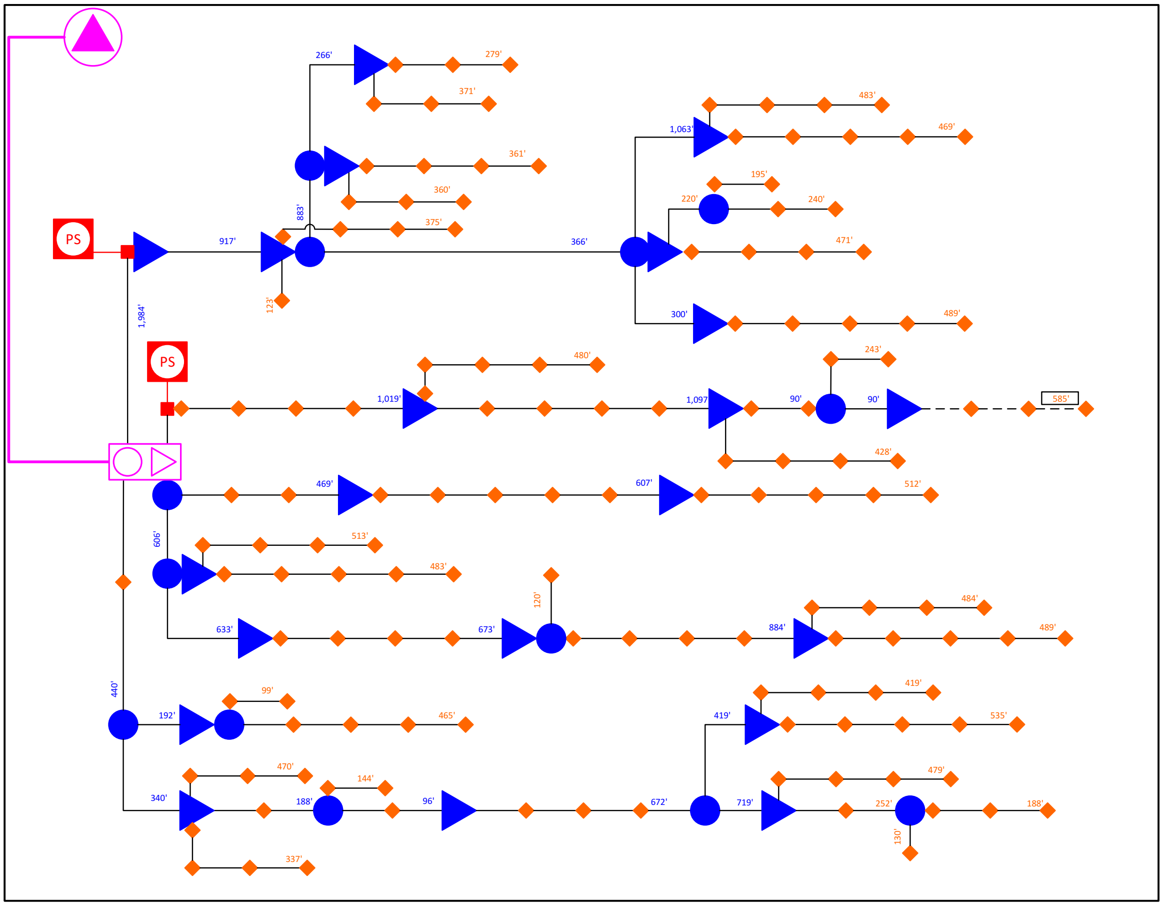

Figure 7 – Topology of the Node C Area

"Business as usual", as the name implies, applies no topology changes. The idea is to refresh all the actives, typically by replacing the existing RF modules with 5-85 / 103-1003 "e-packs". Taps are assumed to function to at least 1 GHz. Node segmentation can be done "in place" by converting this 1x1 node up to 4x4 node, with optical transport multiplexed over the same fiber. While the segmentation can drop the average size down to 100 HP (~50 subs), the distribution is often unbalanced between the RF legs.

Fiber Deep (FD) N+0 will eliminate all the RF amps and reconfigure the network in a way to deploy the minimum number of new nodes, possibly in a new location. Figure 8 shows one such implementation for Node C, where the total number of new actives is reduced, from the original 1 node and 21 RF amps down to just 6 nodes. Note that the new nodes might need augmented output power, e.g. 64 dBmV, to drive the additional coax to reduce the node count. This is one of many trade-offs to be made in a fiber deep design.

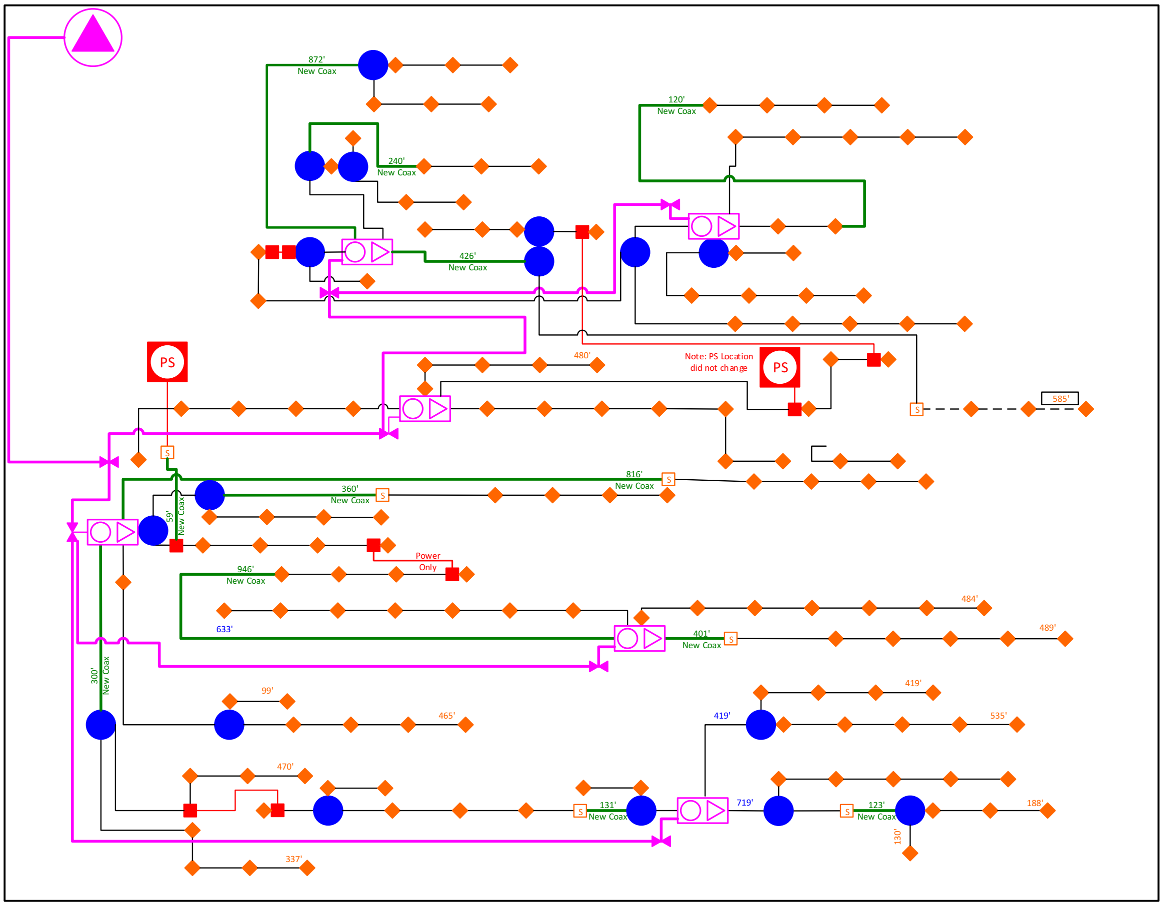

Figure 8 – Node C Area Reconfigured as Fiber Deep N+0

In addition to the new fiber required to feed those nodes, there is a need to add some coax plant, too. The new coax segments are shown in green. A significant redesign of the tap values and orientations is required, too. However, if an operator already plans to upgrade the taps to 1.2 GHz performance, then the argument is the tap rework may not be so onerous of an extra step. The additional new fiber to connect the new nodes is the reason this approach is called “Fiber Deep”. For FD N+0 in Node C, this step takes fiber to as close as 195 feet to the last tap, while the furthest tap is at 1,448 feet. On average, taps are 1,007 feet away from the fiber plant. The new nodes are also capable of housing Remote PHY Devices (RPDs) and PON OLTs if and when needed.

Fiber to the last active (FTTLA) is also an N+0 implementation. However, the number of actives is not minimized. Rather, the locations (and even the housings, if warranted) of the existing RF actives are preserved – and reserved for the last-active nodes. Figure 9 shows topology of such a network, if implemented for Node C. This results in 21 nodes for this design replacing the original actives.

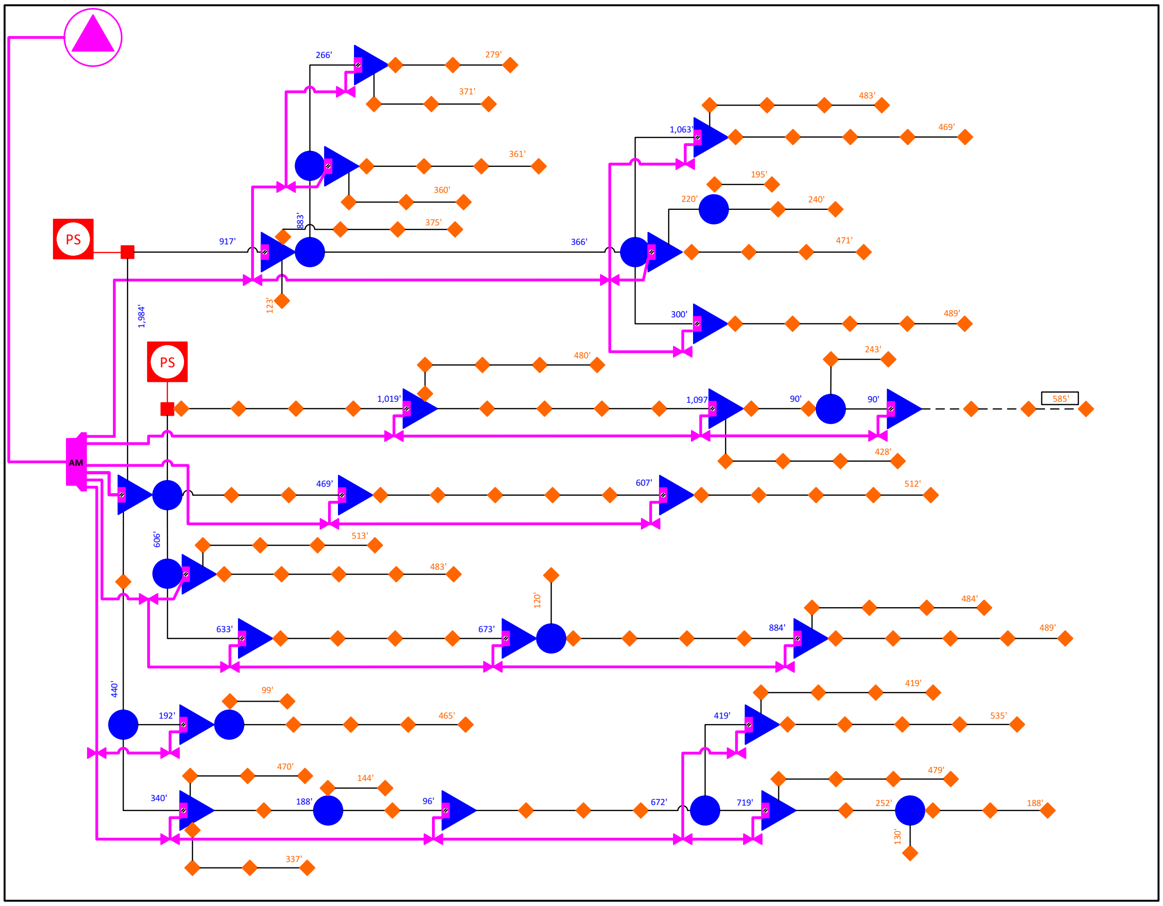

Figure 9 – Node C Area Implemented as an FTTLA in DNA Style

This approach is an even deeper Fiber Deep architecture. In case of FTTLA in node C area, fiber gets as close as 99 feet to the last tap, while the furthest tap is at 585 feet. On average, taps are 408 feet away from the fiber plant. Table 4 shows how much of the new fiber is required for the five areas under study.

For FTTLA, there is no need to touch the coax plant – hardline, taps, even levels for the existing services – so the whole plant upgrade investment is applied to getting fiber deeper, rather than spending part of it on reconfiguring the coax plant. This simplification and getting the fiber even deeper, however, are a trade-off against the number of actives required in the plant. Replacing the taps for 1.2 GHz is an option if an operator wants the additional capacity.

Table 4 – New Fiber Construction Required for FTTLA Implementation for the 5 Nodes

Node	**A**	**B**	**C**	**D**	**E**	**Overall**	**Average**
New Fiber Mileage	2.1	4.0	2.4	1.4	1.2	11.0	2.2
Aerial	0.6	2.8	2.4	1.4	1.1	8.2	1.6
Underground	1.5	1.2	0	0	0.1	2.8	0.6
New Fiber as % of hardline plant	51%	64%	67%	54%	62%	60%	

FTTLA may be favored by those that don't want to touch the taps and passives and put more of their investment dollars into pushing fiber much closer to the premise. FD N+0 is more feasible when the taps are being replaced anyways and the operator wishes to minimize the number of active elements in the plant. FD N+0 also has much fewer nodes which reduces overall maintenance costs as well as cable power losses. In reality, there is a spectrum of fiber deep choices between these two extremes that an operator can optimize for any given location.

FTTLA in particular aids the Selective Subscriber Migration strategy in a few ways. In this strategy described earlier, a small number of high performance subscribers are moved onto a separate FTTx network. In the near term, an operator might pull fiber to the last active only for the location associated with the high performance subscriber. In the Node C example with ~200 subscribers, perhaps two subscribers get the top billboard tier. The operator only needs to upgrade two actives to effectively put them on their own separate upgraded SG, leaving the other 19 actives alone. And while pulling fiber to these two actives, it may enable FTTLA for several other actives along the way. Longer term, the operator may want to start migrating the top tiers to FTTC or FTTP. Using the FTTLA as a launching pad gets them much closer to the homes (e.g. 408' to tap on average for Node C). Selective Subscriber Migration strategy can be implemented with FD N+0 as well. It just requires more work to upgrade the HFC around that node and the fiber is not quite as deep as FTTLA.

DOCSIS Full Duplex (FDX) may require a Fiber Deep system with no actives beyond the node. So from an FDX perspective, both FTTLA and FD N+0 will meet these requirements.

The Fiber to the Curb (FTTC) architecture effectively replaces all of the plant's hardline coax with a fiber overlay. So the new fiber mileage required would essentially be equal to the plant coax mileage from the first row of Table 3. The Fiber to the Premise (FTTP) architecture would require all of the FTTC fiber plus the drop cable for each subscriber. No picture is needed as these simply overlay the existing HFC coax with fiber.

ACCESS NETWORK ENERGY CONSIDERATIONS

HFC Upgrade Options

As a first order of business, the HFC upgrade options for the access network case study in the previous section are considered. Each of the cases was selected to provide a range of scenarios with varying number of Homes Passed (HP) per mile. This can dramatically impact the HFC design and its potential consumption. Figure 10 shows the power consumption for the various HFC upgrade options discussed in the previous section.

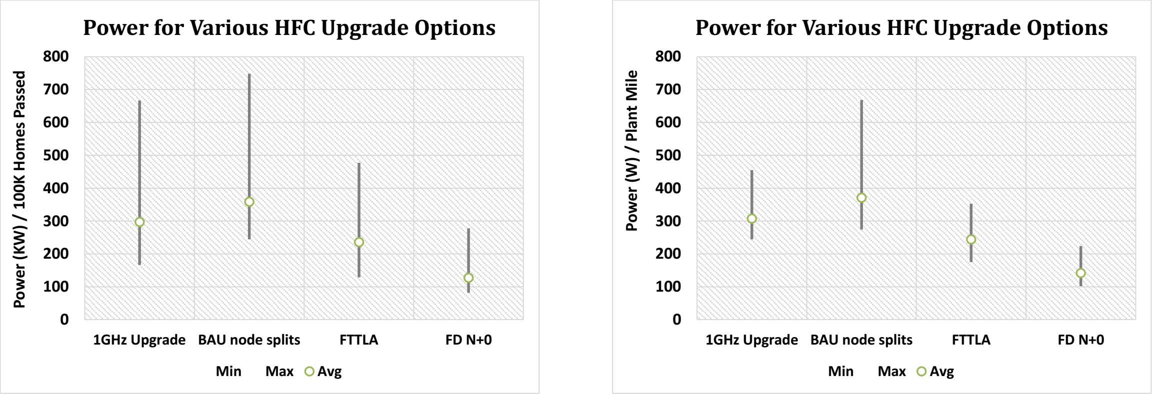

Figure 10 – Power Consumption for Various HFC Upgrade Options

The left hand side of Figure 10 looks at power consumption relative to Homes Passed (HP), while the right hand side evaluates power per mile of plant. As shown, the five use cases can create a wide variation. The W/HP might vary by a factor of 3X to 4X from use case to use case. The variation in W/mile is less, but is still between 2X to 3X but is the same general overall trend. Note that in all cases the average power consumption is typically much closer to the min value. Note that this analysis does not factor in power losses over the coax distribution. Preliminary estimates indicate that this should be on the order of 5% or less of the total plant power consumption.

The baseline case is a 1 GHz upgrade without changing the number of nodes. Making this a 1.2 GHz upgrade does not substantially change the power analysis. In the Business-as-Usual (BAU) node split scenario, the number of nodes increased roughly 3-fold on average. The additional new nodes resulted in slightly higher power consumption than the 1 GHz upgrade baseline, on the order of 20% for the average.

The fiber deep approaches, Fiber to the Last Active (FTTLA) and Fiber Deep Node+0 (FD N+0), both eliminate any active components following the fiber node. As such, both of these will accommodate DOCSIS Full Duplex (FDX) in the future. FTTLA is typically deploying nodes with just one or two outputs while FD N+0 is a more extensive HFC re-design deploying fewer nodes with 3 to 4 high drive outputs. The FTTLA approach provides roughly a 20% power savings on the average compared to the 1 GHz baseline HFC. The FD N+0 with “normal” output power provides more than 50% power savings from the 1 GHz HFC plant. Some operators might consider FD N+0 with high output amplifiers (e.g. 64 dBmV) which consume significant power. The power savings drops to around 35%. This case is not shown in the charts.

While the power consumption per HP and per mile are interesting data points, the power consumption per unit capacity (e.g. KW/Tbps) is also considered. An architecture’s total system capacity is a function of the network link capacity and the number of service groups. With DOCSIS 3.1, the network link capacity might vary from architecture to architecture; especially as the length of the cascade is reduced eventually to zero. For this analysis, FTTLA & FD N+0 might have a 5% to 10% advantage over the N+3 1 GHz upgrade (e.g. 4096-QAM instead of 2048-QAM modulation). This is a relatively minor impact.

The most significant impact on total system capacity is the number of unique SGs that the architecture can support. This is shown in Figure 11. It is assumed that the node in the baseline architecture has already been segmented, so there is only one SG for each node. For the BAU node splits scenario, the HP/node is reduced by a factor of three. It is assumed that these could also be 2x2 segmented, so the number of SG increases by a factor of six, compared to the baseline.

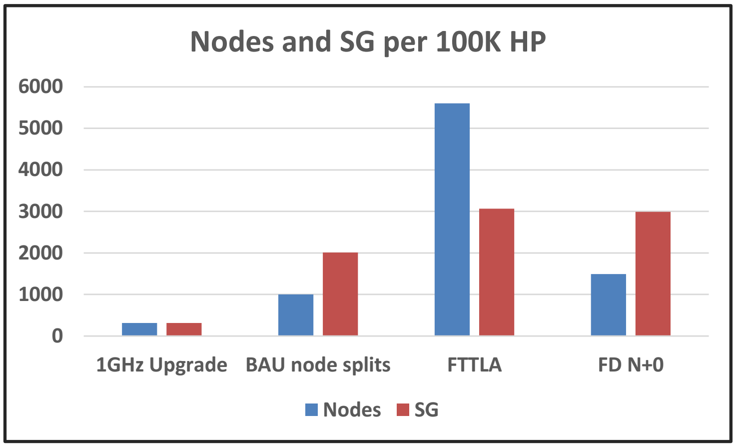

Figure 11 – Nodes and SG per 100K HP

The FTTLA scenario generates extremely small nodes with very few HP/node. Typically, many nodes may be aggregated to form a single SG. For this analysis, it is assumed that the minimum sized useful SG is 32 HP. So while FTTLA may have 18 times as many nodes as the baseline, it roughly has 10 times the number of SGs as the small nodes will be aggregated together.

FD N+0 strives to minimize the number of nodes compared to FTTLA. In this study, FD N+0 has almost 5 times as many nodes as the baseline. This is far less than FTTLA. It is expected that the FD N+0 architecture will support a large node housing with 4 outputs. It is assumed that this node can eventually be segmented up to 4x4. In reality, not every node will support 4 outputs, but this analysis assumes that there will be an equal mix of 3 and 4 output nodes. The net result is that FD N+0 supports a very similar number of SGs as FTTLA!

Given these inputs on SGs per 100K HP along with a network link capacity of 8.6 Gbps for FTTLA and FD N+0, the power consumption per Tbps can be calculated for the serving area. This is shown in Figure 12. The chart on the left shows all four HFC upgrade options. Notice that the baseline 1 GHz upgrade is significantly higher than the other three options. The chart on the right zooms in on the other three options.

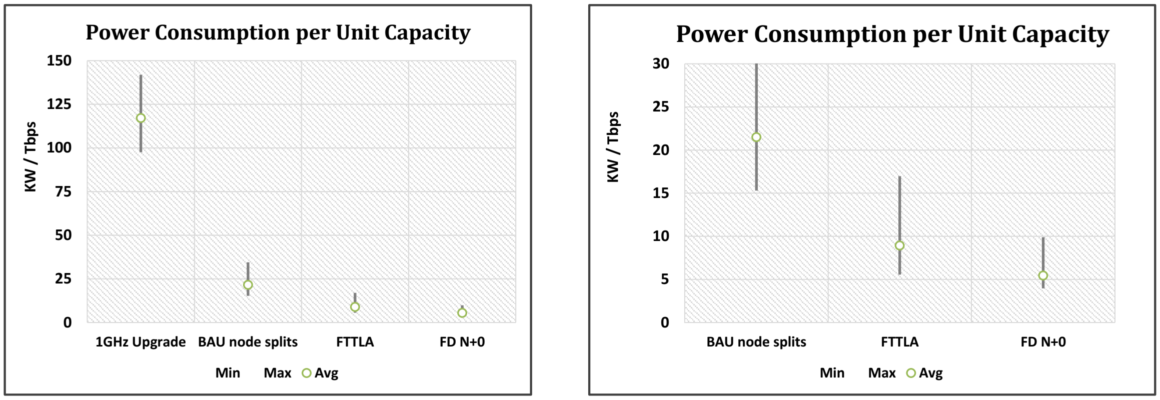

Figure 12 – Power Consumption per Tbps for Various HFC Upgrade Options

The BAU approach achieves about a 5X improvement in this metric. The FTTLA and the FD N+0 are even better. The FTTLA approach achieves about a 13X improvement while the FD N+0 with normal outputs achieves better than 20X improvement.

Distributed Architectures and PON Options

Distributed architectures are garnering a lot of interest recently. These include DOCSIS options such as Remote PHY (R-PHY), Remote MACPHY (R-MACPHY), and Remote CCAP as well as distributed PON options such as Remote OLT (R-OLT). These are relatively new technologies that are still in development. However, enough is known at this point that a preliminary power estimate can be made for some of these solutions.

R-PHY solutions could be applied to any of the HFC upgrade options discussed previously. The focus of this analysis is the relative power impact on fiber deep solutions such as FD N+0 and FTTLA. The R-PHY impact on the other HFC options is relatively insignificant.

For the FD N+0 architecture, it is assumed that every FD node supports a 1x2 R-PHY module. This module is called an R-PHY Device (RPD). In R-PHY solutions, multiple RPDs can then be aggregated in the DOCSIS MAC core in the headend facility to create a single DOCSIS SG. In this case study, there will typically be 5-6 RPDs per SG. Over time, additional MAC core resources can be applied to reduce the number of RPDs per SG. It is also possible that additional RPD capacity can be added to a FD node to virtually segment it into a 2x2 or even 4x4 RPD in the future.

In a conventional FTTLA system, putting an RPD into every node would increase the RPD count by a factor of 3 to 4. This would cause both the power budget and money budget to explode. An alternative approach for FTTLA utilizes the Distributed Node Architecture (DNA). In this architecture, many nodes are aggregated by a splitter element (passive in the downstream and active in the upstream to eliminate Optical Beat Interference – OBI). In a DNA system, the RPD can be located next to the splitter element and shared across many nodes. The RPD must also contain short distance optics that can drive a couple kilometers of fiber. Even with the added power from these optics modules, the sharing of the RPD can reduce the overall R-PHY power impact. Figure 13 shows a logical representation of a DNA system with a shared R-PHY. This distributes the power and the costs of the R-PHY across a larger homes passed population.

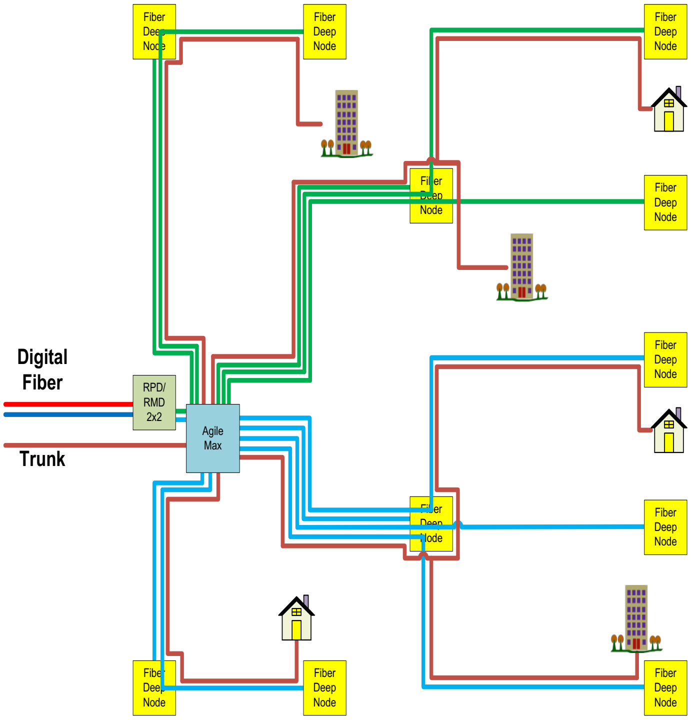

Figure 13 – Distributed Node Architecture (DNA) – FTTLA + Shared R-PHY

Figure 14 shows the power consumption per Tbps for FTTLA and FD N+0 systems, both with and without R-PHY. For FTTLA, adding R-PHY increases power consumption by 7% to 8%, while the FD N+0 plant sees almost 50% increase in power per Tbps. Remember, these are initial estimates and relatively new technologies so there will be improvements over time.

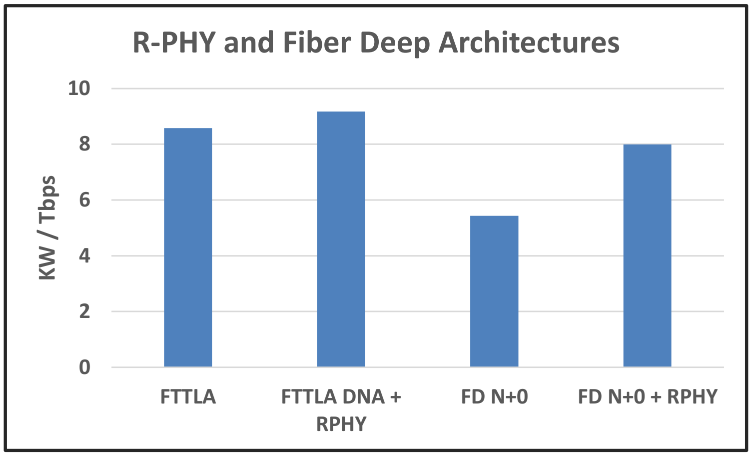

Figure 14 – Power Consumption per Tbps for R-PHY

Adding high power 64 dBmV outputs to the FD N+0 would add additional power as well.

At the time of preparation of this paper, accurate estimates for an R-MACPHY solution were not available. There are many other variables in R-MACPHY system that may affect power consumption depending on which functions are pushed to the node and what stays in the cloud. These options are detailed in the CableLabs technical report [CL_RMACPHY].

Some operators are considering FTTP systems for their future architectures. An FTTP system might be PON or RFoG. The general conception is that a “PON” is completely passive in the plant with no power consumption. This is not necessarily true. Many cable operators must run distances that are greater than 20 km. This long distance trend will increase as operators continue to consolidate hubs and headend facilities. A traditional PON system such as 10G EPON might only support 32 HP at 20 km distances. As distances increase, the fiber optic SNR budget is reduced forcing a reduction in the number of subscribers per OLT port.

There are two basic technologies to get around this problem. One is to move the OLT functionality into the field. This approach is called Remote OLT (R-OLT). The other approach, called PON Extender, is a physical layer repeater where OLT functions stay in the headend facility. This approach supports DWDM connection from the headend facility and converts to the standard PON wavelength in the field. Both solutions have a significantly larger optical budget to the home and can easily support 64 or even 128 HP per port. Also, both of these solutions will require some power in the field.

The power per HP for PON solutions is inversely proportional to the number of homes passed per OLT port. Figure 15 shows the plant power required for R-OLT and PON Extender solutions for several values of HP/port. Since the R-OLT has significantly more functionality, there is a corresponding increase in power. As with the R-PHY solutions, R-OLT technology is in its infancy and its power consumption should improve over time.

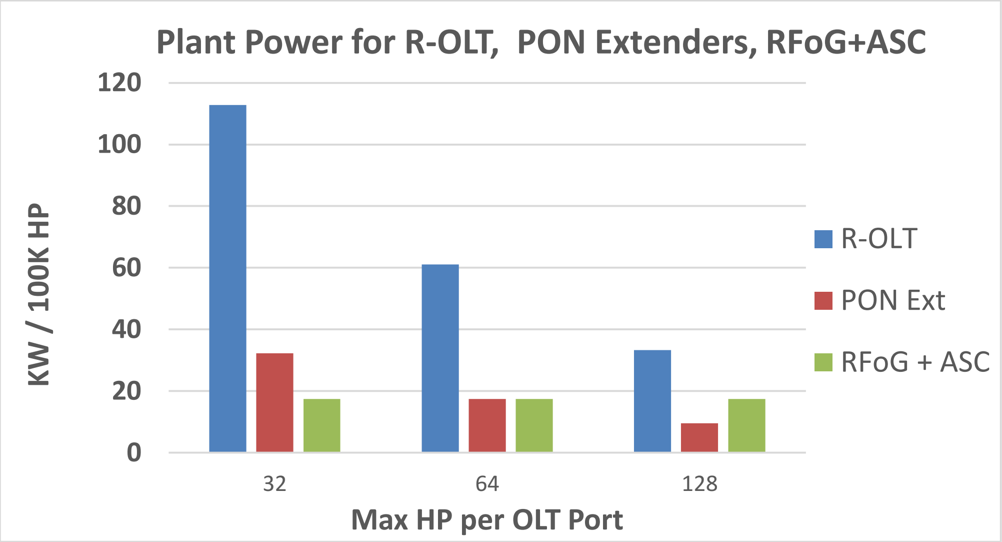

Figure 15 – KW per 100K HP for R-OLT, Pon Extenders, RFoG+ASC

Figure 15 also includes the plant power required for an RFoG system with active splitter combiner (ASC) technology. This has been described in detail in [VENK_2016, VENK_2015 and ULM_2015]. ASC completely eliminates Optical Beat Interference (OBI) which enables RFoG to utilize D3.1 and become a possible end game solution. As Figure 15 shows, its plant power consumption is comparable to the PON Extender with 64 HP per OLT port.

For operators, it is of interest to compare the power requirements of these PON solutions to HFC options. Figure 16 shows power per Tbps for the R-OLT and PON Extender with the HFC fiber deep options shown in Figure 14. Remember that this is only looking at the plant power at this point. The headend facility power will be discussed in the next section. This also does not include the ONU powering required at the premise to terminate the PON system.

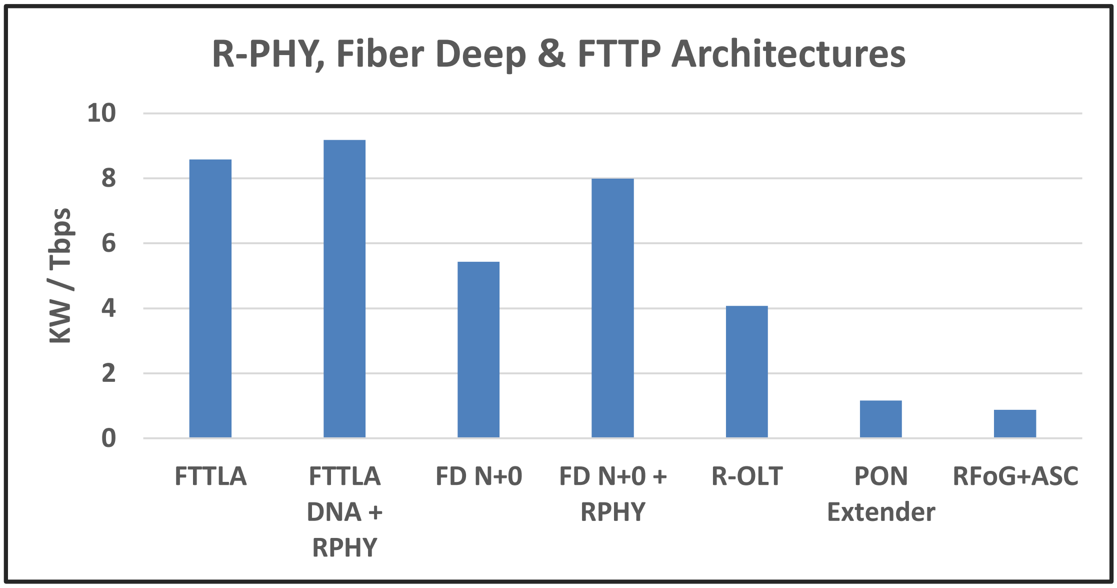

Figure 16 – Power Consumption for R-PHY, Fiber Deep & FTTP Architectures

To put this in perspective, everything shown on Figure 16 is more than 10X better than the baseline 1 GHz HFC plant. The PON extender is the best of these options but will pay a higher power price in the headend facility.

HEADEND ENERGY & SPACE CONSIDERATIONS

This study on headend space and energy considerations will only consider the components directly related to the transport of traffic over the access network. This includes the DOCSIS CMTS/CCAP as the major component (or OLT for PON systems), but also considers UEQAM, RF combining/splitting, and TX/RX Optical shelves in our analysis. This analysis does not include any other equipment that might be related to delivery of video services such as Digital Broadcast, VOD, or STB Out of Band (OOB).

CMTS/CCAP Space & Power – a Historical Perspective

DOCSIS has evolved dramatically over the years and that pace of innovation seems to be accelerating. In recent years, the CMTS transporting DOCSIS high-speed data has given way to a Converged Cable Access Platform (CCAP) that integrates the DOCSIS and digital video EQAM components. For any analysis, it is important to understand the baseline and what is being compared.

A good insight is the space, power, and capacity evolution of the ARRIS product line from the C4 CMTS, to the first generation E6000 CCAP, to newer generation technologies in the E6000. The year 2010 is used as our baseline and three generations of SG density & power are mapped out. This is shown in Figure 17.

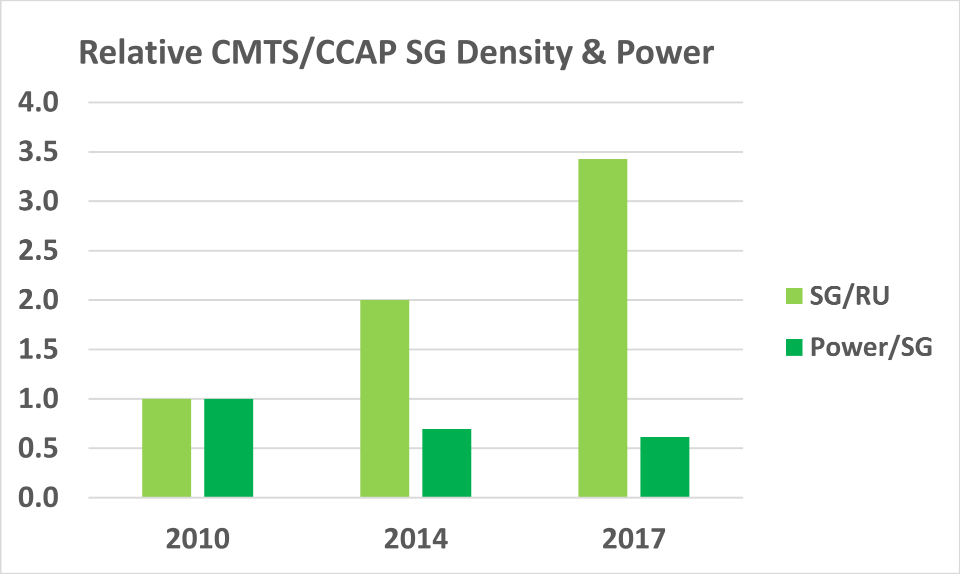

Figure 17 – Relative CMTS/CCAP SG Density & Power

As can be seen, the SG density has increased by almost 3½ times while the power per SG has almost been cut in half. This has helped significantly in allowing operators to keep pace with the broadband traffic growth.

The capacity per SG is another critical element to this picture. With DOCSIS 3.0, the capacity per SG has risen as the CMTS bonded more channels together. Today, D3.1 is available which enhances the capacity per SG even more. Figure 18 shows the resulting increase in capacity per SG along with the increases seen on the CMTS/CCAP Network Side Interfaces (NSI). These data points are all relative to the 2010 starting point. Note that the growth was so significant that this is drawn on a log scale. The capacity per SG will have increased ~60X over an 8 year period.

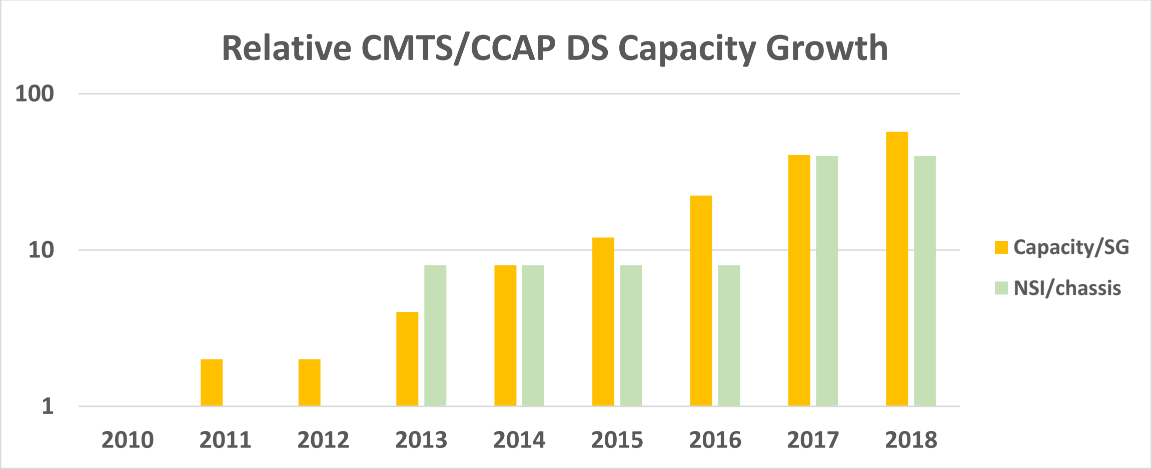

Figure 18 – Relative CMTS/CCAP Downstream Capacity Growth per SG

For our headend space and power analysis, a key metric to measure is the capacity density (e.g. Gbps per RU) as well as the unit capacity/power (e.g. Mbps/W). Figure 19 takes a look at how the CMTS/CCAP products have fared with these metrics since 2010 and compares them to the Tavg traffic growth that was discussed earlier in the paper. Again, note that this is a logarithmic scale.

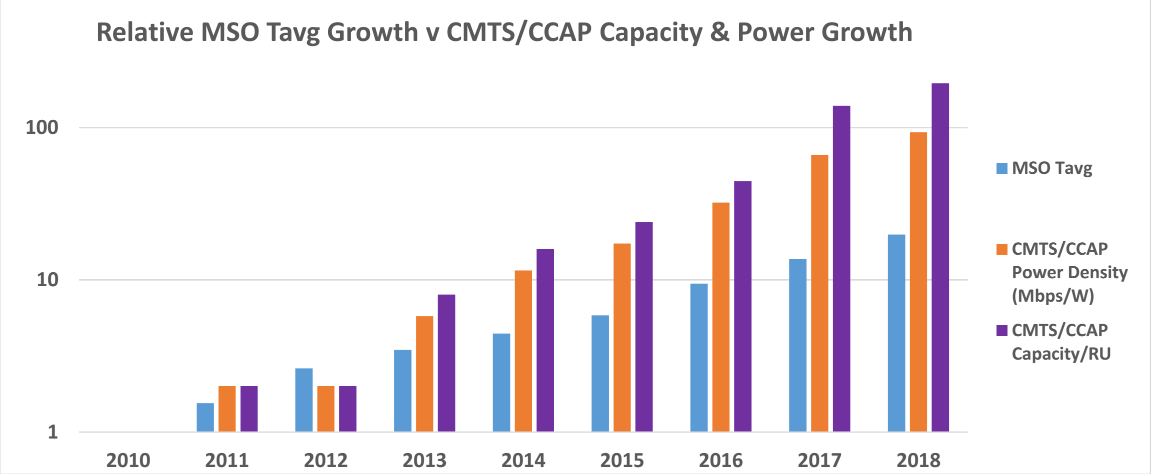

Figure 19 – Relative CMTS/CCAP Capacity & Power Growth vs. Tavg Growth

Over the 8-year window, the 45% Tavg CAGR has resulted in a 20X increase in the broadband data usage. Over that same period, the CMTS/CCAP capacity density has improved by a factor of almost 100X while the capacity/power density (Mbps/W) has increased almost 200-fold. The technological improvements in the CMTS/CCAP have far outstripped the growth in broadband traffic. For operators with older equipment in their headend facility, simply upgrading to the latest CMTS/CCAP products may give them a significant boost in space and energy savings.

Headend Space Consideration for Future Architectures

The CMTS/CCAP is just one of the access network components in the headend. A case study done by [ULM_2013] showed the possible space savings improvements that were possible with the newer technology. This trend has continued over the last several years as evidenced by Figure 19 above.

This paper extends that original analysis. By 2014, the state of the art technology enabled 112 SG in 4 racks of space for an average of 28 SG per rack. This was a significant increase over older technology that might only obtain 6-10 SG per rack. The 2014 baseline is shown in the first column in the table below.

Table 5 – Headend I-CCAP Space and Power Case Study

Headend I-CCAP Migration	**2014**	**2015**	**2015**	**2016-17**	**2016-17**
SG per Rack	28	32	42	96	128
kW per Rack	4.2 kW	3.0 kW	3.6 kW	-	-
W per ‘Chan’ Equiv.	2.66	2.03	1.89	~0.5	~0.5

Recently, the integrated EQAM (IEQ) functionality for integrated CCAP (I-CCAP) has been introduced. This eliminates the external EQAM which results in both a space and power savings. This is shown in the 2nd column. Note the significant reduction in power. An operator can potentially reduce headend access network power by 25% to 30% by upgrading their CCAP with IEQ capabilities and retiring the EQAM equipment. The power savings can actually be higher with retiring older EQAM equipment.

The operator may choose to re-use the IEQ space savings by adding more SG to the racks. This can result in an increased density to 42 SG per rack for a 50% increase in space density. The power per rack increases slightly but the Watts-per-SG metric drops slightly.

With the latest I-CCAP technologies, there will be a significant increase in SG per rack. Part of this is the improvements being seen in I-CCAP densities shown in Figure 19. But there are other factors at play as well. D3.1 capabilities such as Proactive Network Maintenance (PNM) allow the monitoring equipment to be removed. Elimination of the RF combining/splitting is another significant source of space savings. All in all, the SG density jumps to 96 SG per rack in roughly the same power footprint. This means another significant drop in the W per ‘Channel’ equivalent, where a channel is considered to be a 38 Mbps QAM channel.

If space density is a premium, then additional I-CCAP resources may be added to improve the I-CCAP space density to 128 SG per rack. The W per Mbps remains unchanged. So, in a several year span, the operator will have seen a 4.5X improvement in I-CCAP/CMTS space density and roughly 5X improvement in W/Mbps as seen in the last column of Table 5.

Some operators may be considering drastic headend changes such as the consolidation of multiple sites into a single location. This may strain the facilities beyond the I-CCAP improvements just discussed. New Distributed Access Architectures (DAA) have evolved such as Remote PHY (R-PHY), Remote MACPHY (R-MACPHY), and Remote CCAP that can give an operator additional space and power savings in the headend. Some example rack elevations are shown in Figure 20 to give the reader a relative sense of the headend space requirements for each solution.

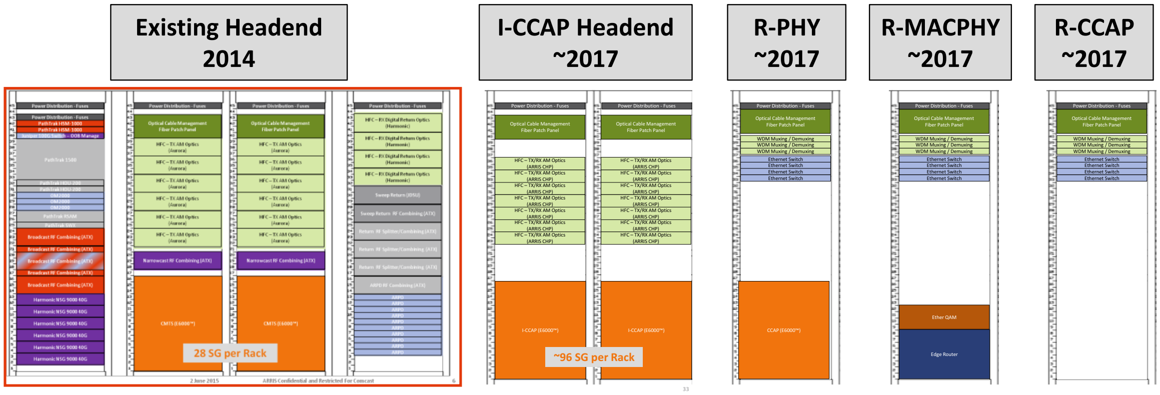

Figure 20 – Sample Rack Elevations for various Centralized and Distributed Architectures

Table 6 below shows the estimated space required to support ~200 SG. The SG Scale uses the 2014 system as a baseline to estimate potential improvements with each architecture. The more functionality that is pushed out into the access network, the more potential space savings in the headend.

Table 6 – Headend Space Density for Various Centralized and Distributed Architectures

Access Architecture (Central-CAA or Distributed-DAA)	**Time Frame**	**Space For ~200 SG**	**SG per Rack**	**SG Scale**
Older CMTS + Low Density EQAM	Pre-2014	20-50 Racks	4-10 SG	0.2-0.4X
High Density I-CMTS/UEQAM CAA	2014	~7 Racks	~28 SG	1X
I-CCAP CAA	~2016-17	~1½ Rack	~128 SG	4.5X
R-PHY DAA		~⅔ Rack	~250 SG	9X
R-MACPHY DAA		~½ Rack	~384 SG	14X
R-CCAP DAA		~⅕ Rack	~960 SG	34X

How much space an operator needs will be very dependent on what technology is currently deployed in the headend along with the current size of their SGs. Is the current headend utilizing 2012, 2014, or 2016 technologies? This will determine how much potential space savings there might be.

Some operators who already have their CMTS SG sizes under 200 subs, may only need to increase their SG count by a factor of 2 or 3 over the next decade. Other operators who may still have average SG sizes greater than 500 subs may be looking at an 8X increase. Every headend site needs to be individually assessed.

Headend Power Consideration for Future Architectures

As operators progress to the future, they are faced with a wide array of choices. From a headend perspective, there is business as usual with I-CCAP, there are new distributed architectures such as R-PHY and then there is the possible FTTP evolution with 10G EPON. Each has a varying impact on headend power consumption.

For I-CCAP/CMTS systems, a key variable in determining the required headend power is the mapping of nodes to DOCSIS SG. In older DOCSIS days, there were often many nodes mapped to a single SG. Over time this has reduced and many operators are now down to a 1:1 mapping of nodes to SG. However, there may be a round of fiber-deep HFC upgrades as suggested in the previous section that may cause the node-to-SG ratio to jump back up. Figure 21 shows the estimated power required for various node-to-SG ratios. The power is normalized to Kilowatts per 100K Homes Passed (HP).

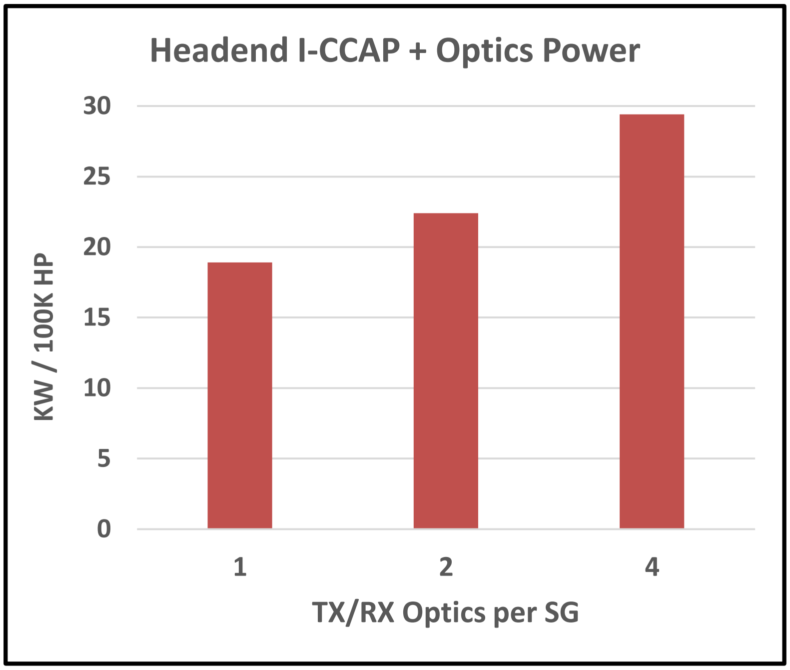

Figure 21 – Headend Power: I-CCAP + TX/RX Optics

As can be seen in Figure 21, the required headend power can vary by 50%.

The next architecture considered is the R-PHY system. This requires a CCAP with the DOCSIS MAC core, but with the DOCSIS PHY removed or disabled. In addition to the CCAP, there may be some Ethernet switches that provide an aggregation function that gathers many 10G Ethernet links and consolidates them into several 100G Ethernet channels.

One of the benefits of the R-PHY MAC core is that an operator can scale their MAC processing as needed. So as utilization increases, additional MAC core resources are added which in turn adds additional power in the headend. Another factor in headend power is the number of Remote PHY Devices (RPDs) that are connected to each DOCSIS SG. The more RPDs per SG, then the more Ethernet switching that is required. Figure 22 shows the impact of these two factors on headend power for an R-PHY system.

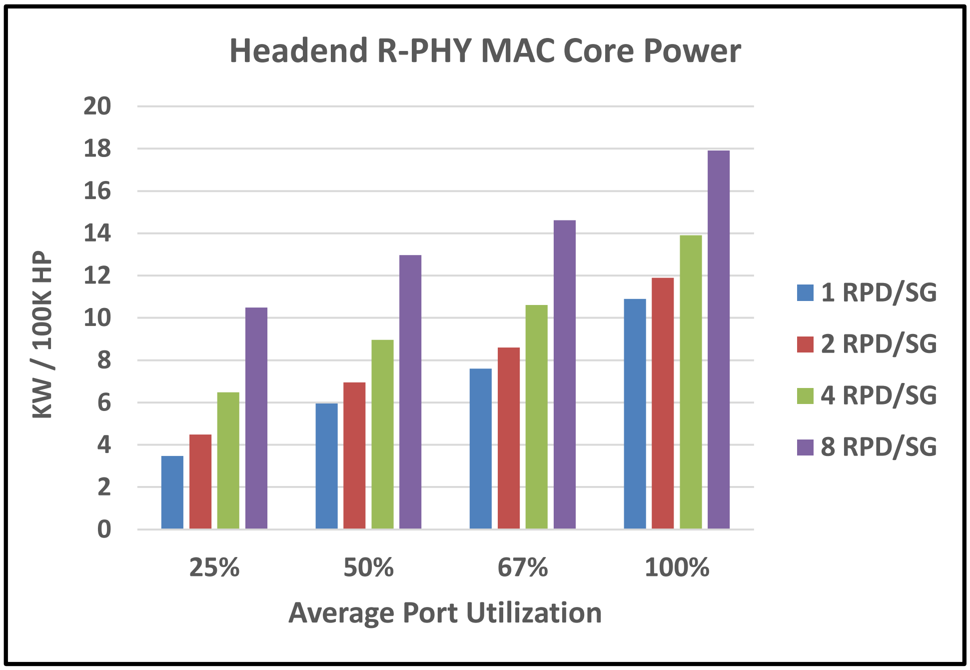

Figure 22 – Headend Power: R-PHY MAC Core + Ethernet Switching

Notice that power consumption can vary by a factor of six. The biggest jump in power occurs when going to 8 RPD per SG. There is also a steady climb in power with utilization.

PON systems can come in three different flavors. In a traditional PON system, there is an OLT in the headend and the access network is completely passive. The major factor in headend power is the number of homes per OLT port. This in turn is driven by the optical SNR budget, which is a function of the distance from the OLT to the homes. For short distances, an operator may configure 128 homes per OLT port. However, for longer distances, the optical SNR budget may limit the operator to just 32 homes per OLT port.

As discussed earlier, there are two ways that an operator can overcome this distance limitation. The first is using PON Extender technology. The OLT remains essentially unchanged except the optics is replaced with DWDM optics. A PON extender in the field then does a wavelength conversion and re-generates the PON wavelength. These PON extenders are typically placed close enough to the subscribers to support the 128 homes per OLT port. The PON Extender is a physical-layer-only conversion.

The other solution to the optic distance problem is to use a distributed architecture with a Remote OLT placed in the access network. Most of the OLT function has been moved out of the headend. However, the headend still needs Ethernet switches to aggregate the R-OLT links. The headend also needs to support controller, DPoE, and routing/switching functionality.

Figure 23 shows the relative headend power required for each of the three PON scenarios.

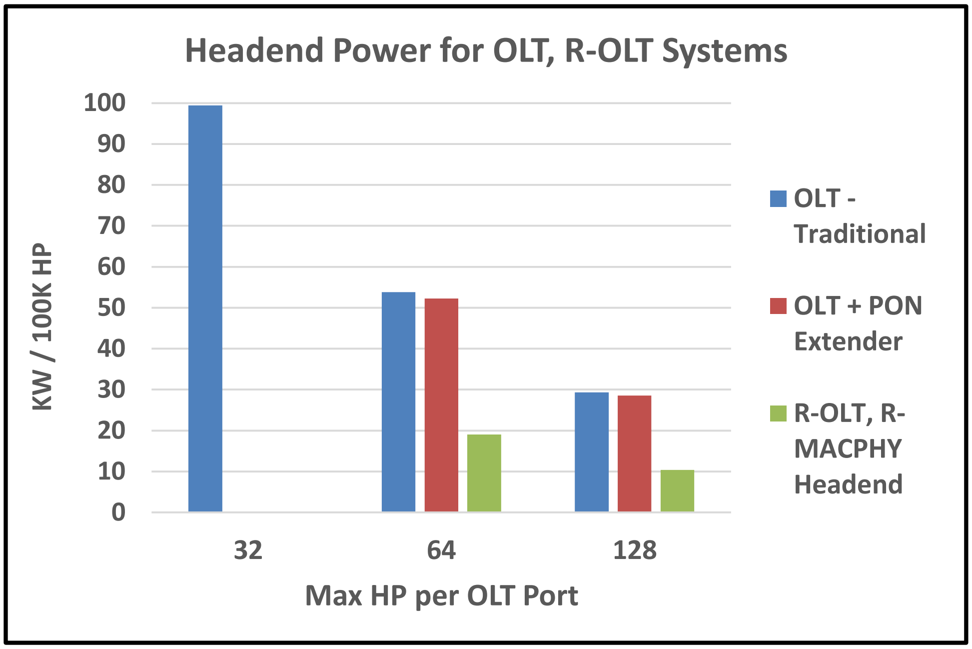

Figure 23 – Headend Power: OLT, OLT + PON Extender, R-OLT

As can be seen in the figure, the traditional OLT system pays a heavy facility power penalty in the headend if it must limit the number of homes per OLT port. If distances are short enough or PON Extender or R-OLT is used for 128 homes per port, then the PON's headend power is in the same ballpark as the I-CCAP and R-PHY MAC Core systems.

TOTAL ACCESS NETWORK ENERGY CONSIDERATIONS

Putting this altogether, Figure 24 gives us an insight into the access network's total system power consumption once the headend facility is combined with the outside plant.

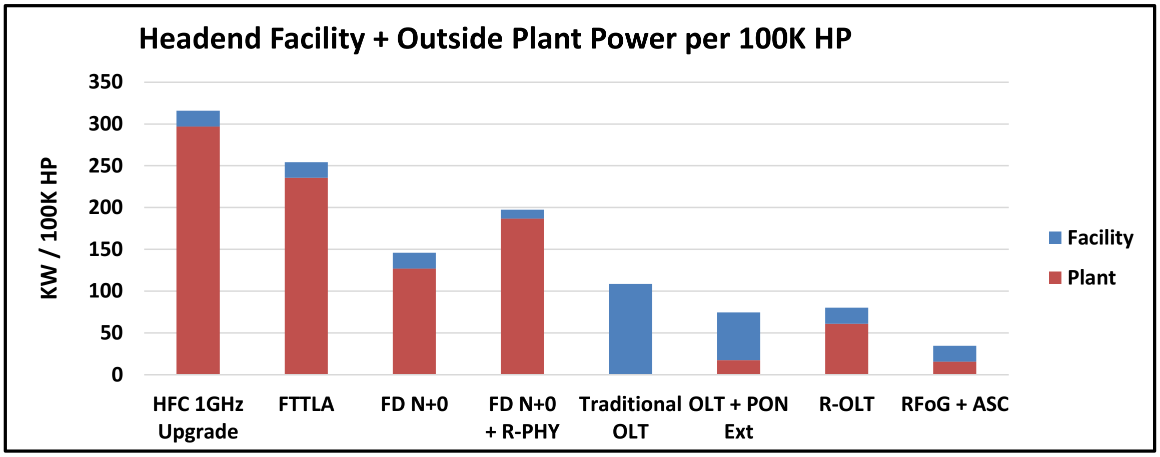

Figure 24 – Headend Facility + Outside Plant Power

For all of the HFC options, the power to drive the coax plant dominates. The Fiber Deep alternatives, FTTLA and FD N+0, reduce overall power consumption from the baseline 1 GHz HFC case. Besides cutting total power consumption by up to 50%, another significant advantage that these two provide is the big jump in total network capacity as was shown previously in Figure 11. The Fiber Deep architectures enable a 10-fold increase in the number of SG for less power. This means a corresponding more than 10-fold increase in total network capacity too.

The power for PON alternatives shown in Figure 24 comes in below the HFC options. But as shown in [ULM_2016], will this be enough to outweigh the economic costs of pulling fiber all the way to the premise? The traditional OLT has the highest headend facility power consumption of all the solutions. As shown in Figure 24, it is assumed to be limited to 32 HP per port due to fiber optic distance / loss budget. If fiber distances are relatively short, then KW/HP could potentially get cut in half. Using a PON extender to remove the distance limitation allows the OLT to support more HP per port. This provides a significant total power savings, especially in the headend facility. The R-OLT option comes in very close to the PON extender option, only with most of its power in the Outside Plant (OSP) rather than the facility.

For comparisons sake, Figure 24 also shows an RFoG solution with active splitter combiner (ASC) technology that eliminates Optical Beat Interference (OBI). RFoG with ASC provides the best of both worlds. It requires minimal plant power that is comparable to the PON extender. Yet, it leverages the DOCSIS headend facilities that handle much larger subscriber counts per SG with the resulting headend facility power savings.

Note that the other piece of the energy consumption that is not addressed in this paper is the Consumer Premise Equipment (CPE).

CONCLUSION

In summary, today's HFC outside plant power consumption dominates over headend facility power. Migrating to Fiber Deep FTTx technologies can result in a significant power savings for the outside plant. Since migrating to FTTP as an end game will be a multi-decade journey, this paper evaluated various options that might be a stepping stone along the way.

Selective Subscriber Migration strategy is a sensible approach for an HFC to FTTx transition. Moving top tiers to FTTx can buy HFC extra decades for 80-95% of subscribers in the flagship basic/economy tiers. Tmax dominates for the next 5-7 years, so it is more important to increase the HFC capacity to at least 1 GHz spectrum rather than split nodes. However, Tavg finally catches up 8-10+ years from now; and SG size reductions come back into vogue. Operators should push fiber deep enough to enable Selective FTTx for top tiers on demand and be prepared for the next round of SG splits.

To understand what the best option is to enable this migration, the paper analyzed in detail five very unique real nodes that varied from sparse rural node to a very dense urban node. Design work was then done on these five nodes for each of the following scenarios:

- “Business as usual” 1 GHz active drop in upgrade with node split as needed
- Fiber Deep Node+0 – FD N+0 and FTTLA
- FTTP

The results show that there is significant power consumption variations from use case to use case. On average, the FTTLA approach provided a 20% power savings compared to the baseline 1 GHz HFC plant. FTTLA minimizes the amount of re-work required to the existing coax plant while driving fiber closer to the home providing a great stepping stone to FTTP when or as needed. FD N+0 minimizes the number of nodes at the expense of some additional coax re-work but provides about 50% power savings compared to the 1 GHz HFC baseline. The reduced node count will also improve operational savings through reduced maintenance. In addition to the total power savings, these Fiber Deep technologies also increase the number of potential SGs by an order of magnitude, so the KW/Tbps metric increases up to 20-fold!

These Fiber Deep HFC systems are compared to distributed access architectures like R-PHY and FTTP architectures such as PON and RFoG. The R-PHY system effectively pushes functionality and hence power from the headend facility out to the plant with no net power savings. The PON systems have generally lower power consumption than the HFC system but require the largest investment with fiber to the premise. The PON power consumption is very sensitive to the number of homes passed per OLT port. Leveraging

R-OLT or PON extenders are useful methods of increasing HP per OLT port and providing more power savings.

Many headend facilities are strapped for space. This may be exasperated as operators continue to consolidate multiple hubs and headends into more centralized locations. As was shown, CMTS/CCAP technologies have made tremendous progress in both space and power densities. Just upgrading from older to newer technology can make a significant space savings. I-CCAP systems will continue to improve and are expected to show a 4-5X space density improvement over 2014 technologies. If this doesn’t provide enough space savings for operators, then various distributed architectures can be deployed to provide additional space savings from 9X to 14X over the 2014 baseline.

As operators migrate to a more green, energy efficient world, there are a number of choices available to actually reduce their overall power consumption while keeping pace with the unrelenting growth in consumer data traffic.

ACKNOWLEDGEMENTS

The authors would like to gratefully acknowledge the assistance of Stuart Eastman and members of his team who helped immensely with analyzing, dissecting, and creating data and material for the use cases that are the backbone of this paper. We also would like to acknowledge Venk Mutalik for his inputs on FTTLA and other fiber deep architectural choices.

ABBREVIATIONS

ABR	Adaptive Bit Rate
ASC	Active Splitter-Combiner
BAU	Business as Usual
Bcast	Broadcast
Bps	Bits Per Second
CAA	Centralized Access Architecture
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
DAA	Distributed Access Architecture
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DNA	Distributed Node Architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DWDM	Dense Wave Division Multiplexing
E2E	End to end
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
FD	Fiber Deep
FDX	Full Duplex (i.e. DOCSIS)
FEC	Forward error correction
FTTC	Fiber to the Curb
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTP	Fiber to the Premise
FTTT	Fiber to the Tap
FTTx	Fiber to the ‘x’ where ‘x’ can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
GPON	Gigabit-Passive Optical Network
HFC	Hybrid Fiber-Coax
HP	Homes Passed
HPON	Hybrid Passive Optical Network

HSD	High Speed Data
I-CCAP	Integrated Converged Cable Access Platform
IEEE	Institute of Electrical and Electronics Engineers
IEQ	Integrated Edge QAM
LDPC	Low Density Parity Check FEC Code
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
N+0	Node+0 actives
Ncast	Narrowcast
NFV	Network Function Virtualization
NPV	Net Present Value
NSI	Network Side Interface
OBI	Optical Beat Interference
ODN	Optical Distribution Network
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OLT	Optical Line Termination
ONU	Optical Network Unit
OOB	Out of Band
OPEX	Operating Expense
OTT	Over the Top
PHY	Physical interface
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
RFoG	RF Over Glass
ROI	Return on Investment
R-OLT	Remote OLT
RPD	Remote PHY Device
R-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RX	Receive
SDN	Software Defined Network
SG	Service Group

SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
TaFDM	Time and Frequency Division Multiplexing
Tavg	Average bandwidth per subscriber
TCO	Total Cost of Operation
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
UHD	Ultra High Definition
US	Upstream
VOD	Video on demand
WDM	Wavelength Division Multiplexing

RELATED READINGS

- **Powering PON with HFC - A Hybrid for a New Generation** – In the past 10-15 years, fiber-to-the-premises (FTTP) networks have been deployed in many regions of the world. This paper compares the total end-to-end costs and throughput of the most common types of PONs and demonstrates how the HFC node can be used to enable cable operators to deliver HFC and fiber-to-the-premises (FTTP) services simultaneously from the same node.

- **Looming Challenges and Potential Solutions for Future Distributed CCAP Architecture Systems** – This paper focuses on the analysis of three sub-classes of Distributed CCAP Architectures (DCAs). The authors describe these DCAs, compare them by operational costs, ease-of-use, infrastructure compatibility and design simplicity. Readers will gain a deeper understanding of the advantages and disadvantages for each approach.

- **A Comparison of Centralized and Distributed Access Architectures for PON** – This paper defines and compares two classes of access architectures that will emerge this decade for Passive Optical Network (PON). It proposes the adoption of multiple wavelength technologies for 10G EPON DPoE systems and examines three types of CAA DPoE systems.

MEET ONE OF OUR EXPERTS: John Ulm

John Ulm holds the position of Engineering Fellow, Broadband Systems for ARRIS within the Network Solutions CTO group. In this role he investigates Advanced Technologies for Broadband Systems including strategic technical directions for multiscreen services and bandwidth expansion. Recent activities include research into next generation CCAP architectures; HFC to FTTx migration including Hybrid PON (HPON), distributed access architectures such as Remote-CCAP and Remote-PHY; next generation protocols including DOCSIS 3.1 and NG-EPON; and Multi-screen IP Video solutions including multicast-assisted ABR.

John's three decades in the Broadband industry began as designer, architect, and MAC protocol developer at LANcity, pioneering the industry's first cable modem systems. He was primary author for the Cable Industry's DOCSIS 1.0 and 1.1 specifications that drove early cable modem success. He also spent time as Network Processor architect for Nortel and as senior technical consultant to the Broadband industry with YAS Corp.

John holds a BSEE and MSEE from RPI and has a multitude of papers and patents to his name.

REFERENCES

(1) [AYHAM_2014] “The Spectral Efficiency of DOCSIS 3.1 Systems,” Ayham Al-Banna et. al., SCTE Cable Tec Expo ’14, SCTE, 2014

(2) [AYHAM_2015] “Measurement-Based EoL Stochastic Analysis and DOCSIS 3.1 Spectral Gain”; Ayham Al-Banna et. al., The NCTA Cable Show Spring Technical Forum, INTX 2015

(3) [CLOONAN_2014] “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

(4) [CLOONAN_2015] “Lessons from Telco and Wireless Providers: Extending the Life of the HFC Plant with New Technologies,” Tom Cloonan et. al., The NCTA Cable Show Spring Technical Forum, May, INTX 2015

(5) [CLOONAN_2016] “Using DOCSIS to Meet the Larger Bandwidth Demand of the 2020 decade and Beyond” Tom Cloonan, Ayham Al-Banna, Frank O’Keeffe; The NCTA Cable Show Spring Technical Forum, INTX 2016

(6) [CL_D3.1_PHY] Data-Over-Cable Service Interface Specifications DOCSIS 3.1 Physical Layer Specification; CM-SP-PHYv3.1-I05-150326, CableLabs

(7) [CL_RMACPHY] Data-Over-Cable Service Interface Specifications DCA - Remote MACPHY; Remote MAC-PHY Technical Report CM-TR-R-MACPHY-V01-150730, CableLabs

(8) [EMM_2014] “Nielson’s Law vs. Nielson TV Viewership for Network Capacity Planning,” Mike Emmendorfer, Tom Cloonan; The NCTA Cable Show Spring Technical Forum, April, 2014

(9) [EMM2_2014] “A Side-by-side Comparison of Centralized vs Distributed Access Architectures”, Mike Emmendorfer et. al., 2014 SCTE Spring Technical Forum, SCTE

(10) [EMM_2015] “Comparing the DOCSIS 3.1 and HFC Evolution to the FTTH Revolution” Mike Emmendorfer; 2015 SCTE Spring Technical Forum, SCTE

(11) [EMM2_2015] “Estimating Downstream Performance and DOCSIS 3.1 Capacity in CAA and DAA Systems”, Mike Emmendorfer et. al, The NCTA Cable Show Spring Technical Forum, INTX 2015

(12) [EMM_2016] “An Economic Analysis of Brownfield Migration CTTH vs. FTTH”, Mike Emmendorfer, The NCTA Cable Show Spring Technical Forum, INTX 2016

(13) [EMM2_2016] “A Comparison of Centralized vs. Distributed Access Architecture for PON”, Michael Emmendorfer, Sebnem ZorluOzer; Technical Forum, INTX 2016

(14) [ULM_2013] “CCAP Case Study: Enabling Converged Video + Data thru Space & Power Savings” John Ulm, Saif Rahman, Jorge Salinger, Jeff Finkelstein, 2013 SCTE Cable-Tec Expo

(15) [ULM_2014] “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

(16) [ULM2_2014] “Scaling Traditional CCAP To Meet The Capacity Needs Of The Next Decade” John Ulm, Tom Cloonan; 2014 SCTE Cable-Tec Expo

(17) [ULM_2015] “HFC Transformation to FTTP: The Role of RFoG, PON and Hybrid solutions”, John Ulm, Venk Mutalik; 2015 SCTE Cable-Tec Expo

(18) [ULM_2016] “Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies”, John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

(19) [ULM2_2016] “Giving HFC a Green Thumb: A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures” John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

(20) [VENK_2015] “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

(21) [VENK_2016] “Cable’s Success is in its DNA: Designing Next Generation Fiber Deep Networks with Distributed Node Architecture” Venk Mutalik, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

(22) [ZORLU_2015] “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