The Scheduler and the Tap: The Odd Infrastructure Couple

A 100 Gbps Coaxial Future Story

A Technical Paper prepared for SCTE by

L. Alberto Campos
Fellow
CableLabs
858 Coal Creek Circle, Louisville CO, 80027
303 661 3377
a.campos@cablelabs.com

Lin Cheng
Lead Architect
CableLabs
858 Coal Creek Circle, Louisville CO, 80027
303 661 3335
l.cheng@cablelabs.com

Zhensheng (Steve) Jia, CableLabs

Jing Wang, CableLabs

Chris Stengrim, CableLabs

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

1. Introduction

The death or end of coaxial cable transport has been predicted in the past on more than one occasion to give way to fiber-to-the-home. The resiliency of coaxial transport, however, has proven to be quite enduring. Coaxial cable itself has not yet been used to its full potential and operators have demonstrated that with the always improving operational practices in addition to the robustness and flexibility provided by the evolving transport technologies, there are still effective means to get value out of our coaxial infrastructure investment. A key factor in this equation is that under several operator starting point scenarios, fiber-to-the-home (FTTH) deployment requires significant investment and coaxial based evolution alternatives can still address most subscriber requirements.

As it appears that coaxial based transport in the access will still be with us for quite some time, it behooves us to think how coaxial transport would look like 10+ years from now. This paper explores how a coaxial based future could evolve in the long run, what would be the implications on topology, spectrum, and technologies that we could see taking shape under the pressures of deployment costs, service offerings, operational efficiency, reliability and market competition.

Today demand for capacity and consumption patterns from subscribers is quite varied. Unfortunately, in deployments, where serving area sizes until recently consisted of about 400+ households per fiber node, there has been little flexibility to design based on the traffic characteristics that we experience. Instead to maintain acceptable levels of customer experience, we have been designing these serving groups driven by peak user speeds and competitive forces. The diversity of our subscribers' usage patterns indicates that peak-capacity based design can leave a good amount of resources and investment on the table. Service elasticity and resulting in network elasticity to flexibly support our subscriber diversity becomes more relevant in our environment that not only caters to residential subscribers but also small businesses and wireless connectivity sharing the same coaxial infrastructure.

1.1. Capacity Improvement Mechanisms

Our industry has leveraged 3 ways of improving HFC capacity. First is by improving efficiency, second through segmentation and third by increasing the amount of spectrum we use. We have been using all three tools available at our disposal as we have evolved our cable networks.

In cable, we have used several efficiency improvements tools. We have improved the efficiency of how we carry video leveraging more efficient compression techniques. We have introduced profiles in our DOCSIS 3.1 systems so that if a CM is closest to an amplifier or a fiber node, this CM could transmit at greater speeds/efficiencies and not at the lowest common denominator as was the case in earlier implementations. Cable has done a very good job at cleaning the HFC plant, improving channel conditions so that we could carry higher order modulations, reaching up to 4096 QAM in the upstream and up to 16384-QAM in the downstream. Distributed architectures have helped in that they eliminated noise and distortion contributions of the analog optical link, thereby facilitating higher modulation transport. Figure 1 shows a 16384 QAM DOCSIS constellation example.

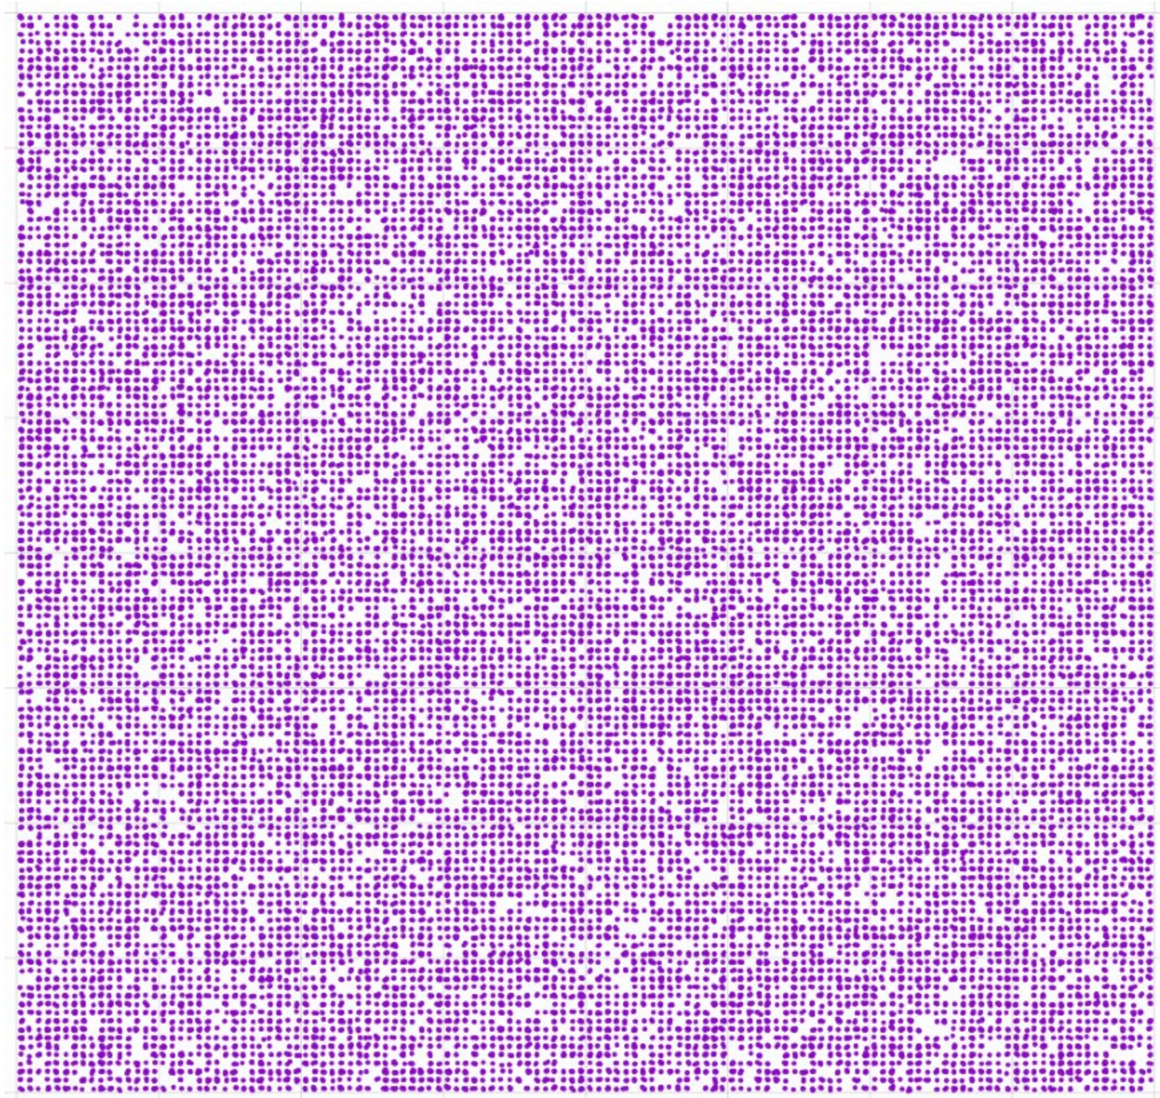

Figure 1 - Highest Efficiency 16384-QAM DOCSIS Downstream Constellation

We have an RF domain peak capacity limitation imposed by the amount of spectrum available. Still by dividing this RF domain into smaller serving areas or segments the amount of capacity per user can be increased. Segmentation in cable has been used on an “as-needed-basis” for quite some time, relying on node splitting to address capacity shortages. The growth in demand for capacity has reached a point in which a surgical approach to segmentation may not always be sufficient as this increase in capacity is more widespread and the rapid growth in demand may require longer term solutions. These solutions consisted of an HFC architecture migration from the original Node+4/Node+5 500 HHP architectures to the smaller Node+2 to Node+0 architectures. In addition to increases in aggregate capacity, they also bring improved performance and higher reliability. Figure 2 shows a Node +0 network migration from a legacy 500 HHP node.

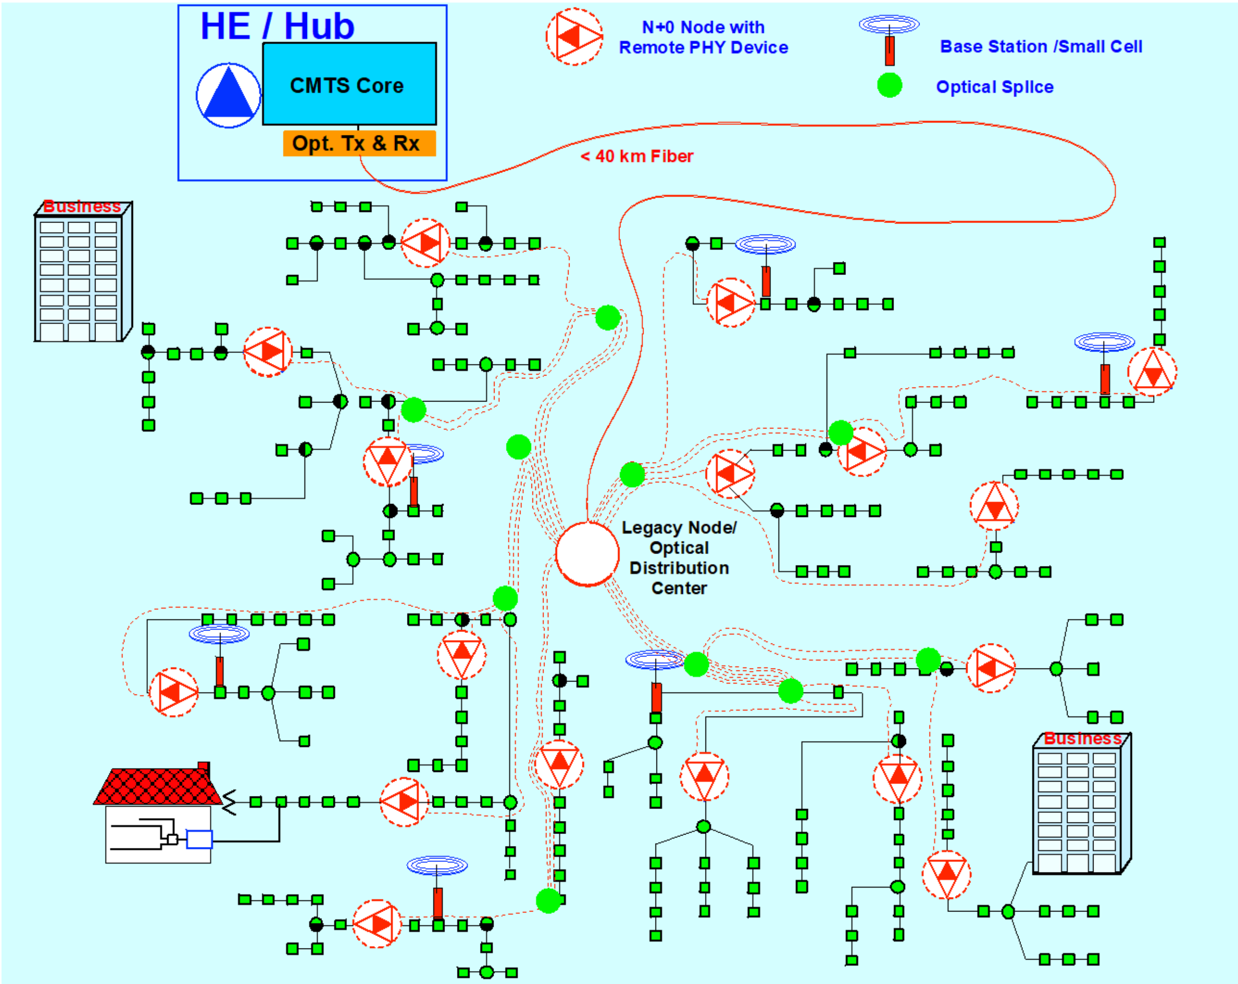

Figure 2 - Legacy/Original Node Segmented Into Node+0 Child Nodes

The changes that have been taking place regarding segmentation and efficiency improvement are quite transformational. The area that perhaps has shown a steady improvement is spectrum enhancement. In the past, it has been somewhat limited by regulatory forces but also by demand, operational complexity and ultimately by cost.

We mentioned earlier how spectrum could enable peak service capacity. In the past peak capacity has been increased by aggregating/bonding more channels within the overall available spectrum. Now however, we have reached the point that the maximum number of channels used by CMs is quickly approaching the maximum amount of spectrum we have available. Therefore, in order to increase service tiers, we need to increase our coaxial RF spectrum. DOCSIS 4.0 addresses this need but at a service tier CAGR of 25% to 40%, we need to start thinking about longer term options. We will discuss the paths we could take to address this demand.

2. Ultimate Coaxial Spectrum Resources

In the early days of cable, only video was carried over the network, video content was limited and the demand to support higher frequencies was also limited. Support to carry the highest practical frequency was also impacted by the attenuation of coaxial cable, the need for multiple amplification stages and the noise and distortion amplification introduced. In this early all-coaxial transport, cable networks evolved from maximum frequencies of 250 MHz to 350 MHz, 450 MHz up to 550 MHz. It was in the 1990's when the transformation of cable networks into a Hybrid Fiber Coaxial (HFC) architecture took place. At that time the HFC network typically consisted of fiber serving areas of 500 households passed (HHP). Instead of having up to 30 amplifiers in cascade, many fiber nodes serving areas were upgraded to 4 to 5 amplifiers in cascade before reaching the furthest subscriber. This network transformation also included the transport of digital video and bidirectional transport to carry data and video at a highest downstream frequency of 750 MHz. With IP-data becoming the dominant use of cable networks, DOCSIS system capabilities evolved along with plant frequency upgrades. DOCSIS versions included 860 MHz, 1002 MHz, 1.2 GHz and with DOCSIS 3.1 and DOCSIS 4.0 specifications 1.794 GHz became the next high frequency target.

The questions that we want to ask ourselves are: How high in frequency can we leverage our coaxial infrastructure? What are the challenges that we need to consider and are there potential approaches to address these challenges?

Coaxial transport relies primarily on the transverse electromagnetic (TEM) mode of propagation. In coaxial cable, TEM mode is radially symmetric and propagates along the direction of the center conductor (Figure 3).

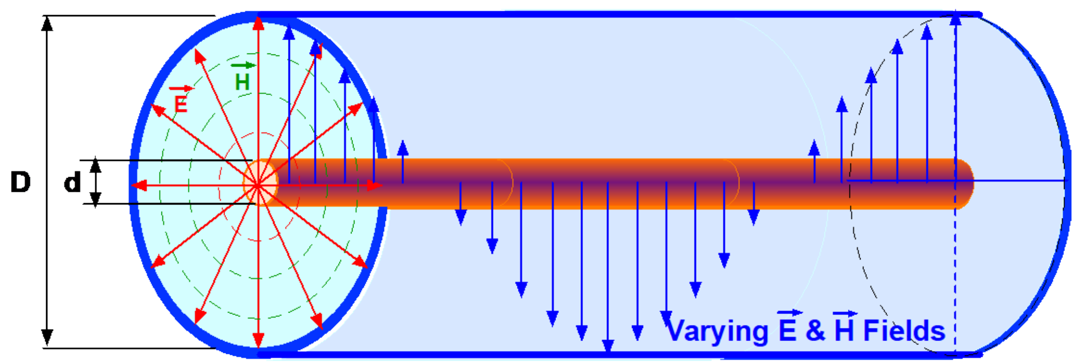

Figure 3 - TEM Propagation Mode in Coaxial Cable

As the frequency increases and the wavelength approaches the dimensions of the radius of the coaxial cable, other modes of propagation are excited. The transverse electric TE_{11} mode is the first to appear. When these modes are allowed to propagate with the TEM mode, they interfere with each other. The frequency at which TE_{11} appears is called the cut-off frequency (f_c) which for TE_{11} mode is given by:

$$f_c = \frac{c}{\lambda_c} = \frac{c}{\pi\left(\frac{D+d}{2}\right)\sqrt{\mu_R \varepsilon_R}}$$

where c is speed of light, D is the inner diameter of the outer conductor, d is the outer diameter of the center conductor and μ_R and ε_R respectively are the relative permeability and relative permittivity of the dielectric material. Table 1 shows the estimated cut-off frequency for different cable types we use in our industry.

Table 1 – TE11 Coaxial Cable Cut-Off Frequencies

Cable Type	Cut-Off Freq
RG6	29.1 GHz
RG11	18.6 GHz
0.5"	11.5 GHz
0.625"	9.3 GHz
0.75"	7.7 GHz
0.875"	6.6 GHz

As you can see with 0.625" and 0.5" hardline cables which are common in the distribution portion of the network, have cut-off frequencies of 9.3 GHz and 11.5 GHz providing significant opportunities for higher bandwidths. This opens up attractive potential resources we could leverage. Nevertheless attenuation, implementation, operational costs, service tier and aggregate capacity are also factors that need to be considered.

3. Legacy of Analog Video Transport

Transport of analog video was required to transmit at a very a high carrier-to-noise ratio. The downstream plant required very clean maintenance and other services had to coexist well with analog video. The channel spacing used by analog video was adopted by newer data services, and out-of-band emissions of the new tenants of cable's coaxial spectrum were also tightly controlled. There have been numerous decisions that have been made in cable based on analog video transport, not only in the design of services that share the same spectrum as analog video, but also decisions in the transport infrastructure itself so that it can optimally support analog video.

The original cable industry plant was designed around one type of service, mainly the delivery of broadcast analog video services. A Cable Television (CATV) service provider frequency-multiplexed a lineup of analog video channels from a central location such as a hub or headend. It transmitted the video signals to subscribers connecting to the coax network within a fiber node serving area. In order to have suitable reception of analog video, each home had to ideally receive the video channel signal at about the same target power level. Cable accomplished this with an RF distribution network where taps coupled RF energy out of the hardline into drop ports to connect to the subscribers' homes via drop cables. Each successive tap following a fiber node or an amplifier, has a specific coupling loss to the drop port so that even after the attenuation of coaxial cable, the power reaching the end-device, such as a set-top-box, is about the same for all subscribers. Figure 4 shows a schematic representation of a coaxial segment with taps of decreasing values.

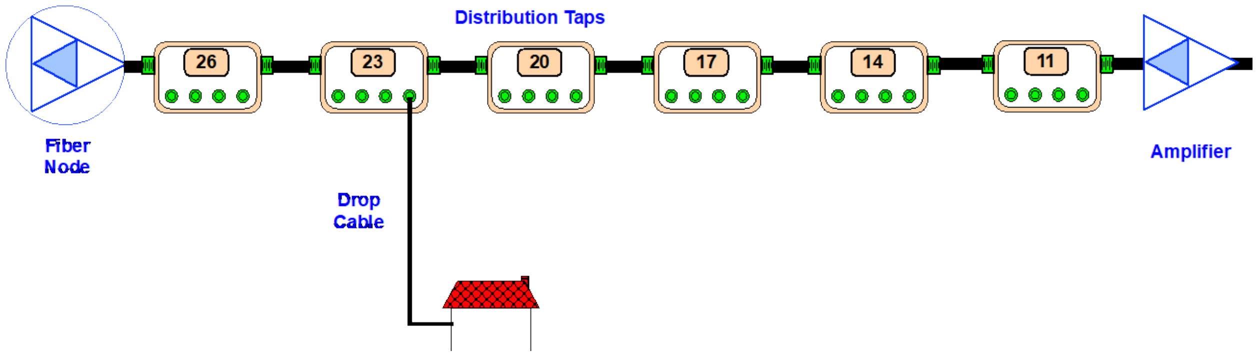

Figure 4 - Fiber Node-to-Amplifier Coaxial Segment With Tap Values To Provide Similar Power Levels Per Video Channel

Over the relatively narrow band that the upstream occupies, differentiation in channel conditions can depend on the amplifier cascade value and specific impairments which predominantly impact one or more cable modems. There is also some frequency dependence when comparing the edges of the upstream band with the middle of the upstream band. The downstream also imposes larger frequency-dependent behavior due to cable attenuation, however this behavior is typically ameliorated by an up-tilted downstream signal and by the decreasing tap values which effectively equalizes the power levels for all customers along the coaxial segment.

4. Evolving HFC Environment

While early in Cable, the focus was on having every analog video channel be received by every Set Top Box (STB) and TV receiver at approximately the same power level, an allowance to deviate from that philosophy of operation has been enabled with the introduction of channel conditions' dependent profiles in the DOCSIS 3.1 specification. Another trend worth noting is that more traffic in cable networks has steadily moved from broadcast to unicast. This has been due to the way services are delivered such as video on demand as well as IP video services. Therefore, the need of having one stream having to be received by all receivers is disappearing.

About a decade ago, the transition to digital video broadcast started in the United States which, except for low power TV broadcast stations, has now been completed. This transition also impacted "must carry" rules. When analog is no longer transmitted, there is no longer the requirement to carry broadcast channels at the over-the-air frequencies. For example, Channel 2 does not need to be carried at 50 MHz in our cable networks when only carrying digital video channels. For all practical purposes, we live today in an all-digital world. Not only video but also voice and data are carried over bit streams. We are no longer required to optimize our networks to distribute video signals so that their service endpoints are at about the same power levels. This is a brand-new ball game; the change in conditions allows us to break free from the restrictions of analog video distribution and everything it entails. The consequences of analog video distribution are not just limited to reclaiming coaxial RF spectrum by replacing analog video with the more efficient digital video. The change in conditions is significantly more encompassing—. Think about revisiting all the network design decisions that have been made since the early days of community antenna television.

4.1. A Step Beyond DOCSIS Profiles

As we move to an all-digital transport, our criterion needs to shift from the optimal transport of analog video to the optimal transport of bits. No longer are we required to deliver every analog channel at about the same power level. Now we need to focus on how to transport the highest number of bits across our entire spectrum.

When we introduced the DOCSIS 3.1 version of the specification, instead of treating all CMs the same, the concept of transmission profiles was incorporated. This meant that CMs that could perform at a higher efficiency level due to a more benign channel would be placed in a profile where higher order modulation would be used. Figure 5 highlights the variation in CNR of a population of CMs and an example of how they could be grouped in different efficiency buckets.

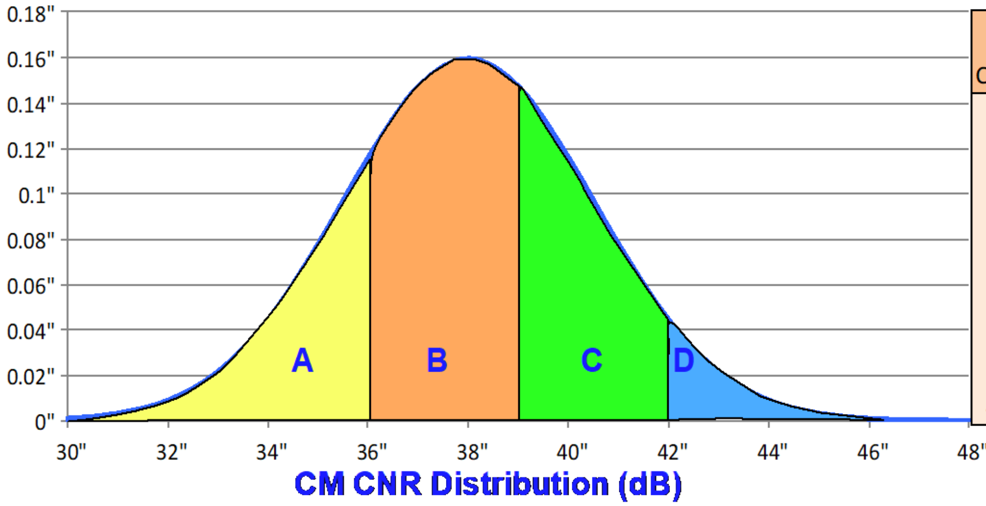

Constellation	Minimum CNR (dB) DS	US
16/QAM	15	17
64/QAM	21	23
128/QAM	24	26
256/QAM	27	29
512/QAM	30.5	32.5
1024/QAM	34	35.5
2048/QAM	37	39
4096/QAM	41	43

Figure 5 - CM CNR Distribution and Modulation Order Thresholds

The assignment of profiles is adaptable to the current conditions of a CM. Today, a CMTS is specified to support up to 16 downstream profiles and 7 upstream profiles. These profiles are defined within the downstream channel with subcarrier granularity and in the upstream with minislot granularity since the modulation order is defined on a per minislot basis. CMs could go to a higher or lower efficiency profile depending on the channel conditions and leveraging the ranging response mechanisms and MER and codeword error metrics.

4.2. Higher Frequency Off First (HFOF)

In a DOCSIS OFDM downstream, resources exist in symbols (time) and subcarriers (frequency). In the upstream, minislots consisting of symbols and subcarriers are also allocated in time and frequency. As we use higher frequencies, the variation of CNR versus frequency is more evident. The attenuation in hardline and drop cables versus frequency becomes the dominant factor determining the MER and CNR in a CM. Figure 6 shows coaxial attenuation versus frequency and Figure 7 highlights how such losses may impact CNR at higher frequencies on a CM attached to a coaxial segment.

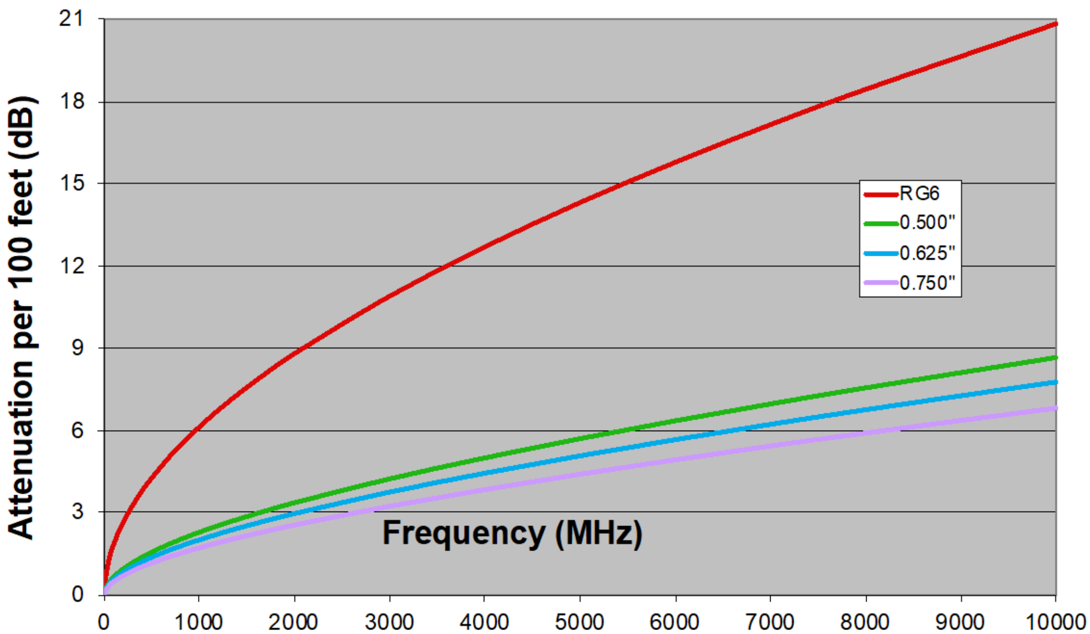

Figure 6 - Hardline And Drop Cable Attenuation Versus Frequency

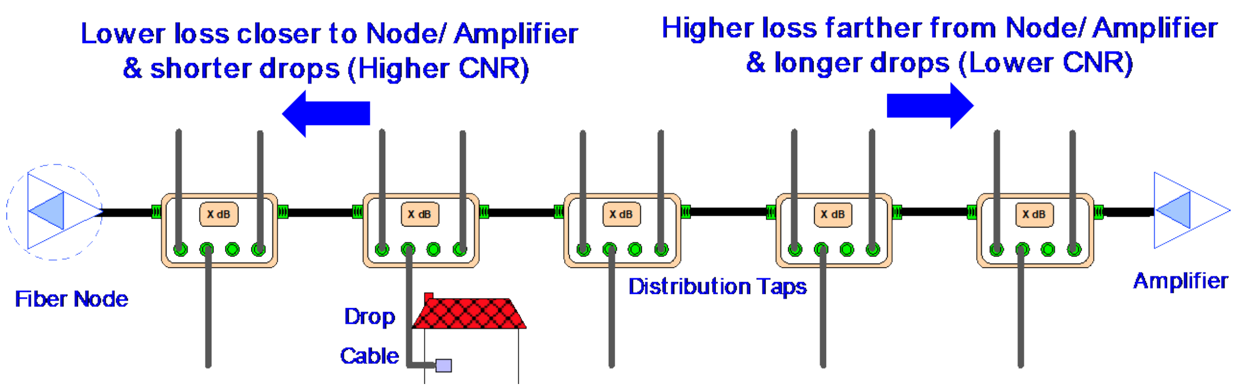

Figure 7 - CM Higher-Frequency CNR Based on Topology Location and Drop Length

Figure 7 shows that, at higher frequencies, CMs that are closer to the fiber node and with shorter drops will enjoy better CNR while CMs that are further out from the node and with longer drops cables will have lower CNR. This effect is due in large part to the fact that Total Composite Power for the transmitted signal from any Node/Amplifier is limited to "reasonable" levels on the order of ~70 or so dBmV, so the received power at higher frequencies is reduced by the combination of limited transmit power at the higher frequencies and increased attenuation at the higher frequencies. This qualitative assessment is quantified using the topology example of Figure 8. This coaxial segment consists of 4 taps in cascade and 12 CMs attached to these taps through drop cables. The drop lengths have been selected on purpose with great diversity of lengths so that we can better observe the significant effect that drop cable attenuation can have on performance. In this analysis, a hardline feeder cable of 0.5" and an available bandwidth of 11 GHz are assumed.

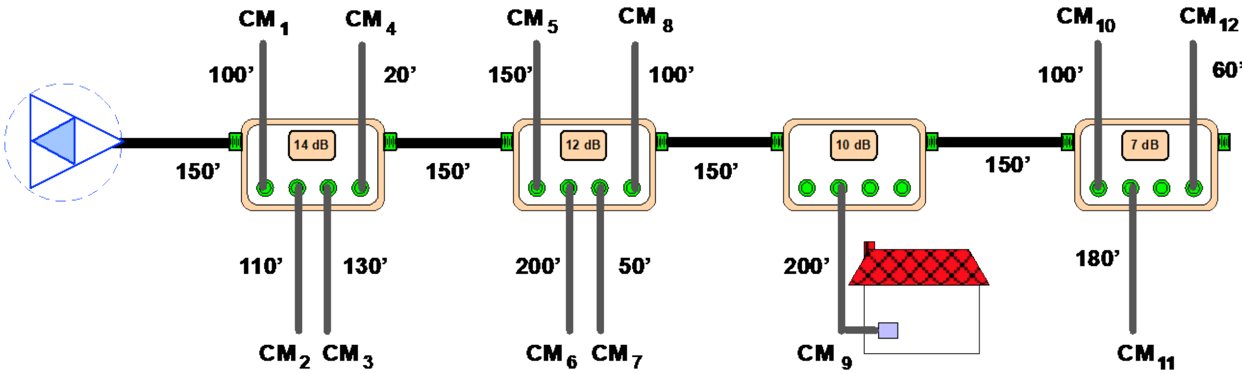

Figure 8 - Sample Coaxial Segment For Ultimate Capacity Estimation

Table 2 shows the total capacity that the CMs within the sample coaxial topology in Figure 8 could have if each CM has access to the entire 11 GHz spectrum.

Table 2 – Full Bandwidth Capacity Of CM Within Sample Topology

Modem	Full Bandwidth (Gbps)
CM_1	134.3592
CM_2	129.5459
CM_3	118.1095
CM_4	146.1113
CM_5	105.6579
CM_6	75.2112
CM_7	145.609
CM_8	134.6757
CM_9	74.7192
CM_{10}	134.053
CM_{11}	85.2012
CM_{12}	144.8434

As you can see the capacity that each CM could obtain varies widely from a lower 74.7 Gbps value to a higher 146.1 Gbps rate. In general, CMs that are close to the Node/Amp and CMs with short drop lengths tend to experience much higher capacities than the others. The average aggregate capacity is 112.14 Gbps which results in an average capacity per CM of 9.345 Gbps if it would be equally shared. If we assign spectral efficiency according to the downstream CNR table in Figure 5 and assuming a 3 dB receiver noise figure, the resulting efficiency versus frequency for the different CMs within the sample topology is shown in Figure 9.

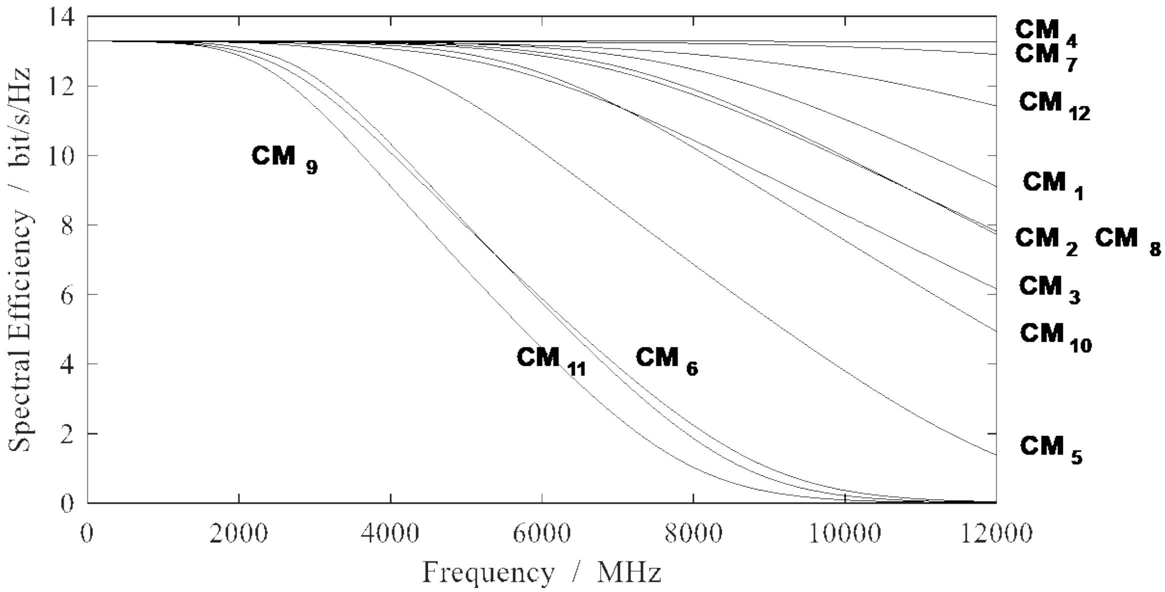

Figure 9 - Spectral Efficiency Versus Frequency Of CMs Within Sample Topology

As observed in Figure 9, CM_4 with a short drop and closest to the fiber node enjoys the best performance while CM_9, which has the longest drop and is located after 3 hardline segments, has the most limiting performance. The performance of CM_{12}, that is located in the last tap but has a short drop highlights the impact that drop cable can have as well as the fact that there is plenty of capacity left for the amplifier that sits behind the cascaded hardline segment.

If one would have the capability of flexibly assigning capacity on a frequency basis to the different CMs that enjoy different channel performance the overall aggregate capacity could be optimized. This mechanism is what we call "higher frequency off first" (HFOF) mechanism, which assigns the higher frequencies to the CMs that enjoy best higher-frequency CNR performance and leaves the lower frequencies for the CMs that have limited higher-frequency CNR performance. Figure 10 shows the allocation of capacity according to frequency bands. The capacity allocated to each CM is the same so there is some variation in the bandwidth allocated to each CM. An overall aggregate capacity of 142.81 Gbps is obtained which represents an improvement of 27% compared to the traditional approach calculated by averaging values in Table 2. (Note: If a particular CM is not utilizing its assigned spectrum, then the CMTS scheduler would be able to re-assign that spectrum to one or more other CMs. This intelligent scheduling would likely be typical in real-world scenarios).

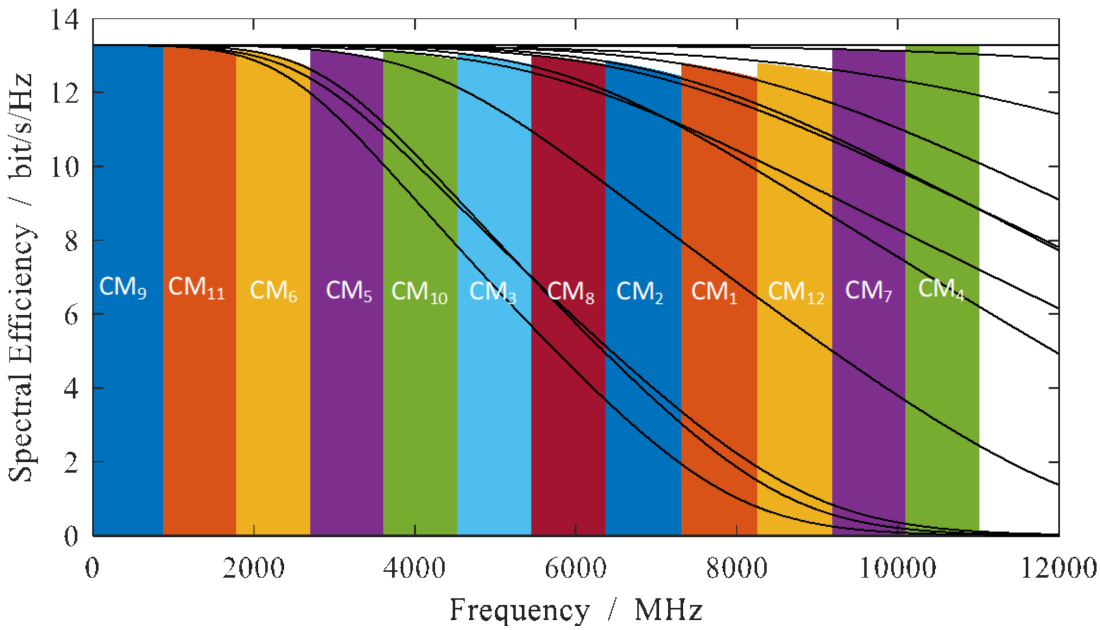

Figure 10 - CM Capacity Allocation Following HFOF Approach

A second simulation was conducted using 200’ coaxial segments for a total of 800’ total segment length. In that scenario using the traditional approach an average aggregate capacity of 105.888 Gbps was achieved or an average capacity per user of 8.824 Gbps. When leveraging the HFOF approach an aggregate capacity of 139.992 Gbps or 11.666 Gbps per user was obtained. This represents a 32% improvement of HFOF over the traditional approach.

4.3. Dynamic Range

The attenuation of hardline and drop cable versus frequency as shown in Figure 6 can be significant. In Figure 11, we combine the frequency response of all CMs along the coaxial segment. As you can see the loss across the entire 11 GHz bandwidth can be significant, but the loss across the portion of the spectrum allocated to each CM according to HFOF is bounded, resulting in a significant relaxation of dynamic range requirements. The blue and red segments on the CM loss curves stay within a limited loss range, highlighting the dynamic range benefits of the HFOF approach.

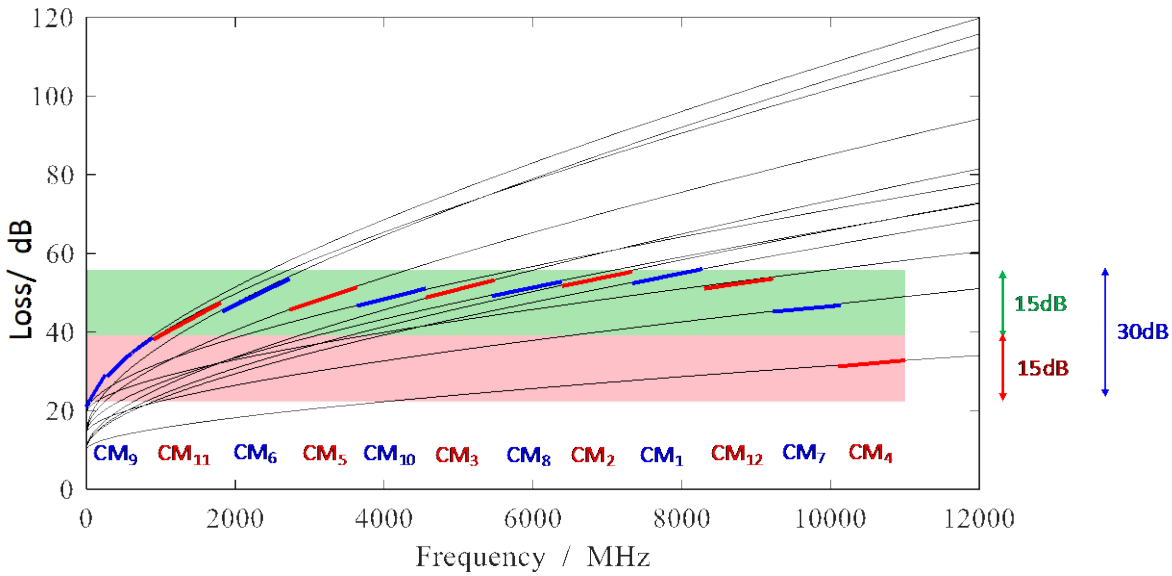

Figure 11 - End-to-end CM Attenuation Within Allocated Frequency Band

4.4. Implementation Implications Of Peak And Aggregate Rates

In the earlier sections, we have seen how a large amount of coaxial spectrum can be made accessible to CMs. In this section, we explore techniques in making this accessibility cost effective. In WiFi and mobile applications, we have systems with limited amounts of bandwidth available out of a diverse selection of spectrum bands. The accessibility to the many options of spectrum bands is achieved through the tuning capabilities of the receiver. How much bandwidth can a receiver simultaneously capture, process and aggregate, is an indication of the peak capacity a handset could reach.

A similar approach could be followed in cable where CMs could have accessibility to a wide spectrum while the bandwidth capture capabilities would indicate the potential peak bandwidth a CM could reach. Figure 12 shows an example of the cable analogy where 10.8 GHz of coaxial spectrum is available. This amount of spectrum is consistent with the cut-off frequency of 0.5" hardline cable. In this example a CM capture bandwidth of 1.8 GHz is assumed. The total amount of spectrum available for the downstream in this example is 10.2 GHz. Leveraging HFOF techniques and assuming a clean plant, a modulation order of 2048 QAM can be reached which leads to a 17 Gbps capacity per 1.8 GHz capture bandwidth and an aggregate capacity out of the entire 10.2 GHz of spectrum of approximately 100 Gbps assuming DOCSIS 3.1 level overhead. The CM capture range could be adjusted based on the target peak rates and implementation cost complexity criteria. Figure 12 depicts the scenario just described where some CMs, depending on where they are within the coaxial segment topology, are assigned certain frequency bands.

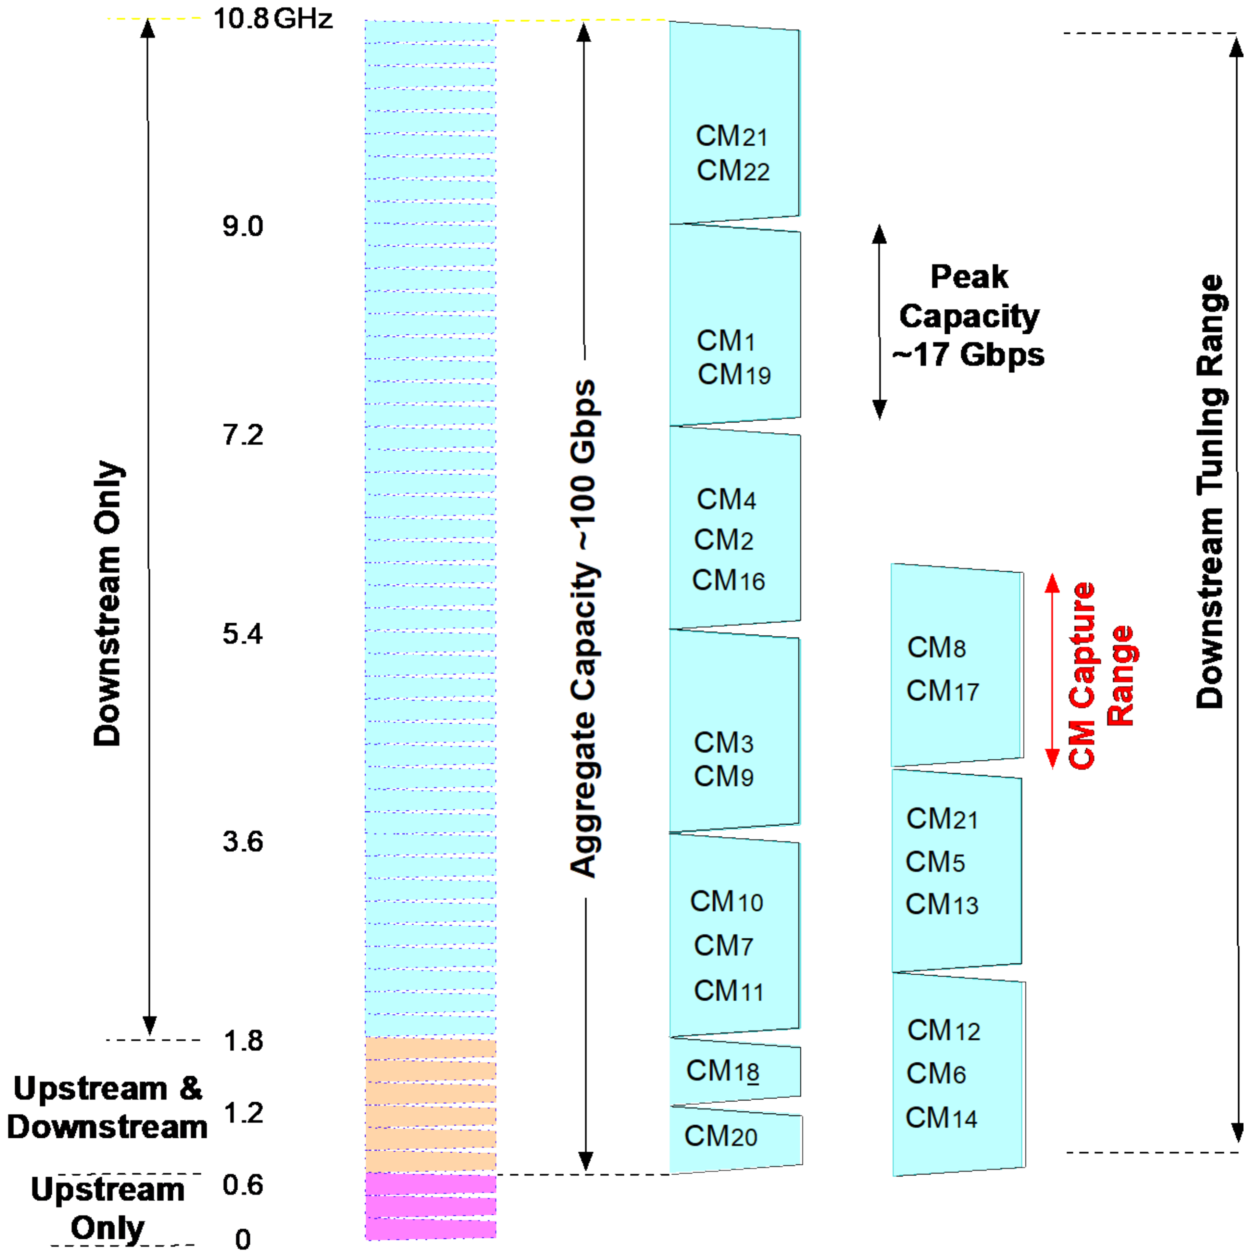

Figure 12 - CM Spectrum And Bandwidth Allocation Example

5. An Evolved Scheduler

In an environment such as the one depicted in section 4 where much higher frequencies are used, such as the case of coaxial spectrum above 3 GHz, subscribers' frequency-dependent performance becomes more noticeable. This behavior becomes highly dependent on which tap they are connected to and how long the drop cable is. In this environment, the importance of having an evolved scheduler that can allocate resources based on frequency and Modulation Error Ratio (MER) is very important.

We introduced profiles in the DOCSIS 3.1 specification which allows custom mapping of subcarrier modulation order versus frequency for groups of CMs associated to a profile. Figure 13 shows an example describing its implementation.

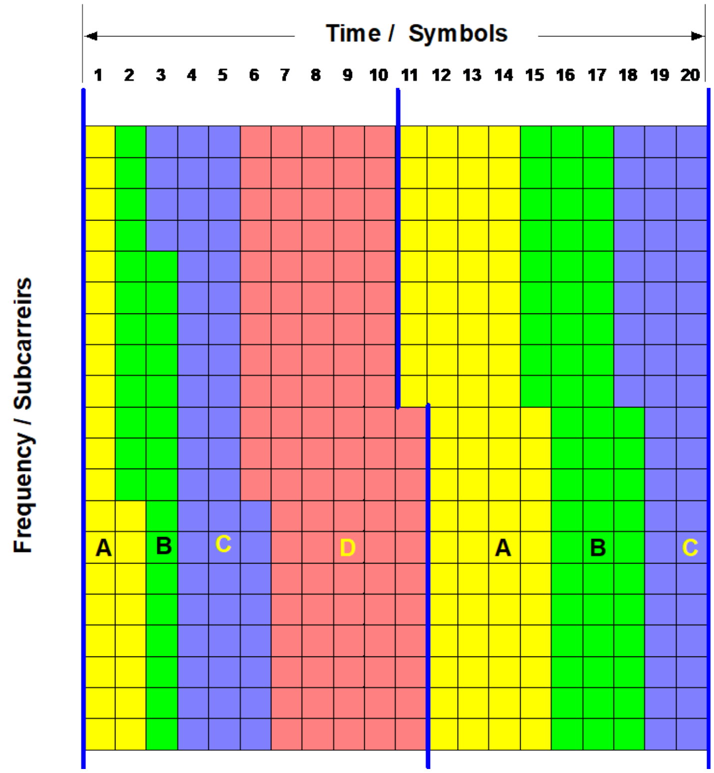

Figure 13 – Conventional Implementation of Downstream Modulation Profiles, A,B,C,D

The profile, however, covers the entire range of frequencies within a channel. In an environment with strong frequency dependent behavior, having the capability of limiting users under the same profile within a certain frequency range would be advantageous to more flexibly implement the HFOF concepts and improve overall system performance. If the channel could be flexibly split in two or three frequency segments as shown in Figure 14, one could optimize overall network performance.

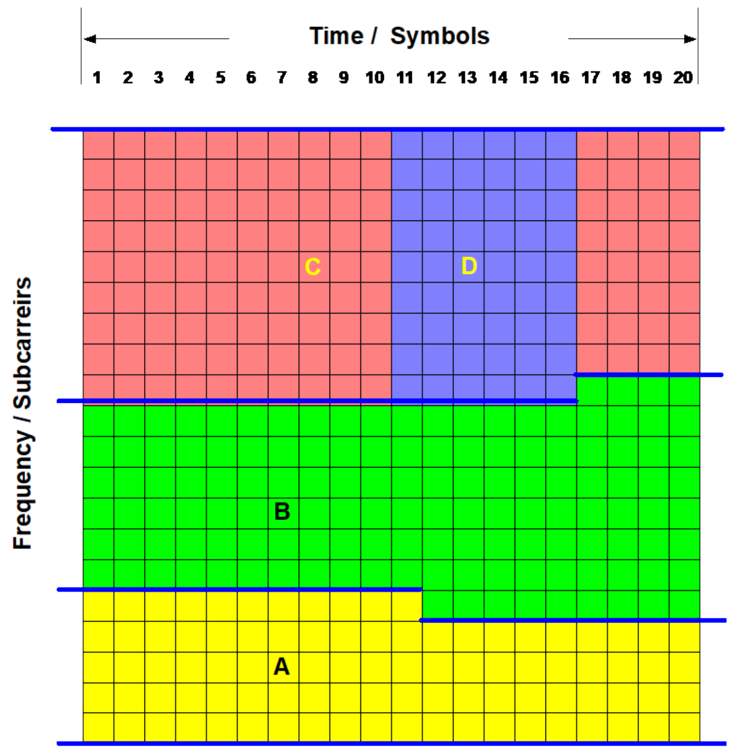

Figure 14 – Frequency Dependent Implementation of Downstream Modulation Profiles, A,B,C,D

Figure 14 shows a downstream frequency-dependent implementation of profiles. A similar capability can be implemented in the upstream

5.1. Additional Frequency Aware Scheduling Benefits

In addition to capacity optimization, having resource allocation control based on frequency and SNR provides other operational benefits. With efficient allocation, a user's capacity is less sensitive to the length of cables and location. Therefore, we could have a system implemented with more flexible coaxial reaches. As operators migrate to N+0 topologies, it is advantageous to split the original HFC node into as few child nodes as possible. To that effect operators can smartly reposition child node locations, not necessarily coincident with former amplifier locations leveraging longer coaxial segments.

Leveraging frequency dependent resource allocation of a next generation scheduler, you can have, for example, a 1.8 GHz coaxial segment with 6 taps supporting 9 192 MHz DOCSIS 3.1 channels. In such a scenario, you can have CMs connected to the first two taps allocated to use channels 5 through 9, CMs attached to the next two taps could be allocated resources from channels 3 through 7 and the CMs attached to the last two taps could be allocated resources from channels 1 through 5. Since the CMs that are farther away use the lower frequency channels one can afford much longer coaxial segments. Another way of looking at this is requiring only lower gain amplifiers resulting in overall lower power consumption.

6. Conventional Taps and Connectors

We mentioned earlier that in order to distribute analog video channels such that they arrive at approximately equal power levels to our subscribers, our coaxial networks were designed with decreasing tap values so that the attenuation of coaxial cable is somewhat compensated by the tap value. In this original approach, a tap in proximity and following a node or an amplifier would have a higher coupling loss than a tap that is farther away from the node or amplifier so that the impact of longer cable attenuation is compensated.

According to the old paradigm, every end device such as a TV set, a set-top box or a cable modem would receive about the same power level per channel. Operators would use RF distribution taps that consisted of a housing structure and a faceplate. Typically, the housing included ports to connect to the hardline cables or the rigid portion of the coaxial network. The faceplate, on the other hand, included ports that connected to the flexible portion of the coaxial network, the drop cables (Figure 15). The faceplate also included coupling and splitting circuitry to provide specific coupling loss values to the drop ports. These faceplates are removable and designed with different coupling loss values to reach the subscriber's premises at the target power level. If different coupling values or tap values are desired, the faceplate is replaced by another one with the desired coupling values.

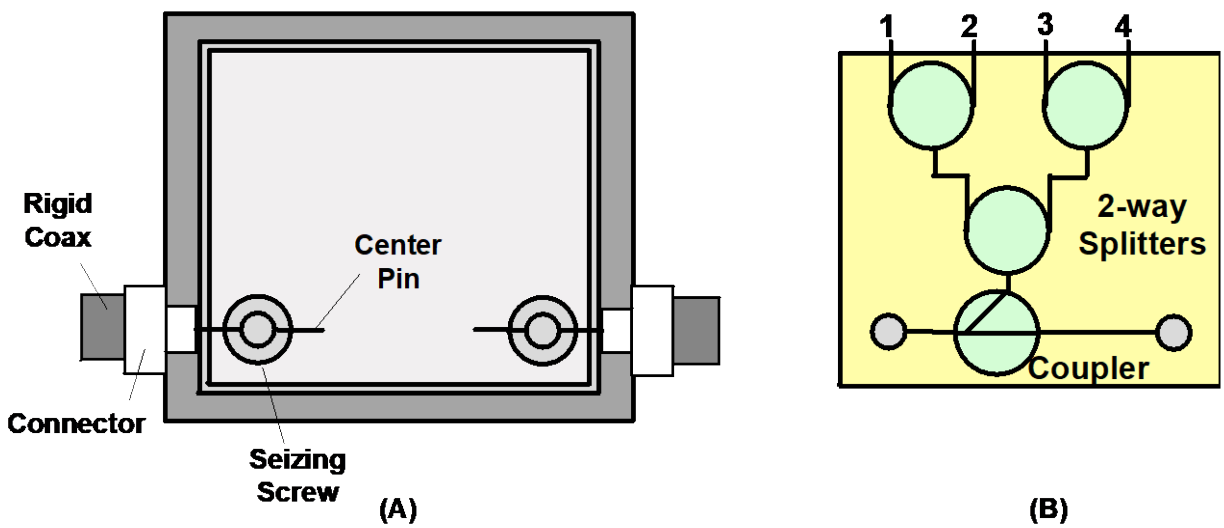

Figure 15 – Tap Housing (A) and 4-Drop-Port Removable Faceplate (B)

One reason to have the tap consist of two components, the housing and the faceplate, is so that faceplates with different tap values can be easily interchanged.

Another reason to have removable faceplates is so that during installation, technicians could have access to the internal structure of the tap. This would allow them to set/configure the tap to receive the center pin of an external KS connector attached to the hardline cable. The connector could attach from a vertical direction when used in pedestals with an underground coaxial distribution network, or from a horizontal direction when the transmission cables are inline as is the case of an aerial distribution network (Figure 16).

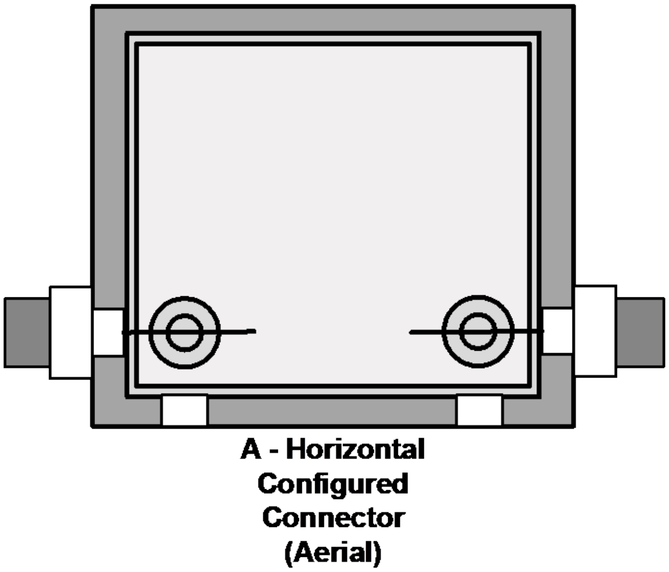

Figure 16 – Horizontal/Aerial (A) and Vertical/Underground (B) Configuration Of Traditional Tap Housings

With a faceplate removed, one can also verify that the long center pin of the KS connector is trimmed to the right length (different tap vendors require different center pin lengths and technicians adjust by manually cutting the center pin) and adjust the seizure screw to make sure a good contact is made with the center conductor of the KS connector.

A fourth reason is to change faceplates with a larger or smaller number of drop ports. This occurs when new customer premises are built and/or a greater number of ports are required.

One challenge that comes from having removable faceplates is that when the faceplate is removed the RF transmission to the elements downstream from the tap is interrupted. The industry solved this by including a conductive path that switches in place enabling an alternate path between the taps' input and output ports, therefore avoiding interruption of AC and RF transport. This alternate path is often implemented using a metal strip which has suboptimal performance at higher frequencies (Figure 17).

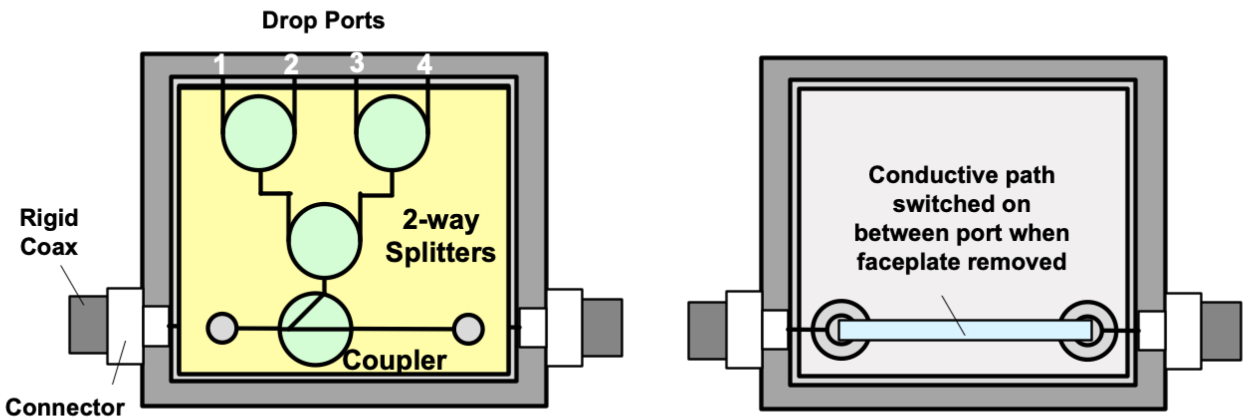

Figure 17 – Tap With And Without Faceplate Showing Conductive Path Switched On When Faceplate Is Removed

7. Revisiting Cable Distribution and Network Components Design

In today's digital age, if we want to control capacity that each subscriber could ultimately consume, we would ideally leverage digital tools that control resource allocation, mainly the CMTS scheduler. If possible, we should avoid using infrastructure means that impact resource allocation when digital means are available. It has been a long time since we moved away from using inline RF notch filters to control premium content access. We have evolved to using digital encryption tools instead. For resource allocation tasks we must also leverage as much as possible our digital domain tools.

Therefore, for the taps that are closer to the node or amplifier we can use a much lower value tap than what is conventionally used (Figure 4). Figure 18b, shows an alternate coaxial segment to the one presented earlier in Figure 4 and Figure 18a, with tap values adjusted so that subscribers leverage the channel conditions and performance they have available in their transmission medium to the node or amplifier.

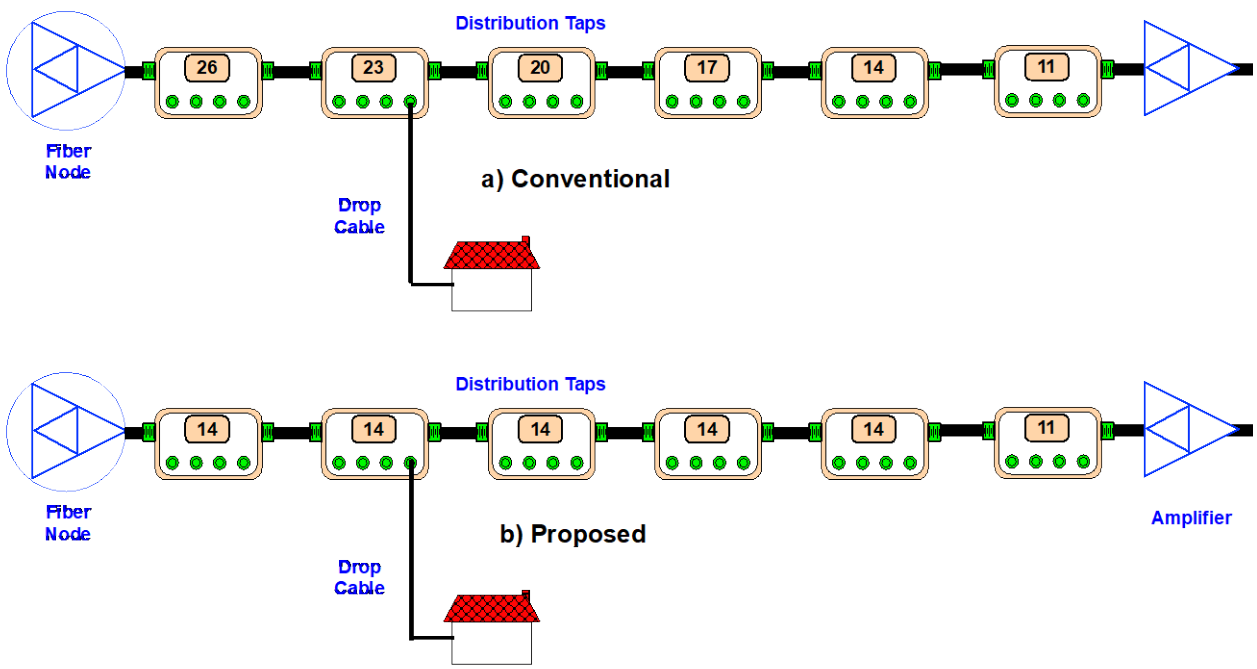

Figure 18 – Coaxial Segment With Adjusted Tap Values Optimize Subscriber Capacity

In Figure 18, the first four taps have decreased their tap values to 14 dB, the 5th tap remains at 14 dB and the last at 11 dB. At lower frequencies, the insertion loss of taps with values 17 dB or higher is dominated by the implementation or excess loss. The insertion loss value for the 14 dB four-port tap is still below 2 dB even after adding excess loss. At higher frequencies the tap implementation becomes more complex and there is a small, gradual excess loss that increases with frequency.

7.1. Single Value Tap

In Figure 19, all the taps for this 4-drop-port scenario, have the same coupling loss value of 14 dB. It does not represent a drastic change, even for the last tap where only a small capacity penalty is incurred. An exception can be made with the end-of-line "tap" when we are actually dealing with a splitter. Keep in mind that not all end-devices or cable modems need to have the same RF power level. The scheduler is in charge of controlling the capacity that the CMs receive, even though the RF receive power may vary among end-devices. Figure 19 shows the implementation of the coaxial segment using only one tap value.

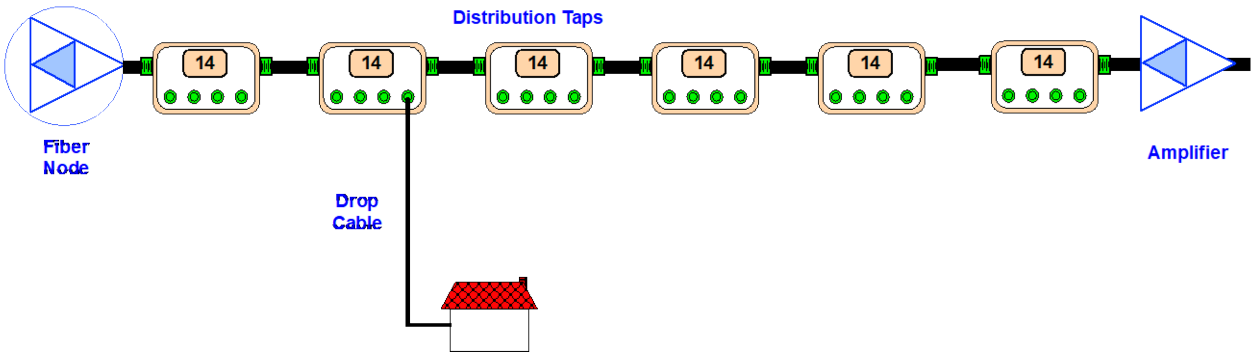

Figure 19 – Single Value Tap Segment For A 4-Port Tap

Using only one tap value for a type of tap significantly simplifies operations. The stocking of taps is much easier. Table 1 shows typical options for tap types and values commercially available.

Table 3 – Typical Tap Types And Coupling Loss Values in dB

2-port	29	26	23	20	17	14	11	8	4
4-port	29	26	23	20	17	14	11	8	
8-port	29	26	23	20	17	14	11		

The 14 dB tap coupling value shown as an example for a 4-port tap has not yet been optimized. Its optimization will depend on the spectrum bandwidth, lengths of coaxial cables and the number of taps that the operators are targeting in a coaxial segment. However, a suboptimal tap value should not significantly affect performance, according to our recent study.

With this in mind, it is safe to assume that there could be a single tap value for 2, 4 and 8-port taps in addition to an end-of-line splitter. This approach would result in the types of taps shown in Table 4.

Table 4 – Modified Tap Types And Coupling Values – Last Column For End-Of-Line Taps

2-port (dB)	11	4
4-port (dB)	14	8
8-port (dB)	17	11

The number of spare taps that a technician would need to stock has been greatly reduced. Without counting the end-of-line splitters, we are dealing with the implementation of a single-value tap for each number of port types.

In this new environment, can the tap performance be improved? One of the challenges when implementing taps is its port-to-port isolation. A simplified traditional tap circuit can be represented by a coupler followed by a splitter. We represent a 2-port tap by a coupler followed by a 2-way splitter, a 4-

port tap by a coupler followed by a 4-way splitter and an 8-port tap by a coupler followed by an 8-way splitter.

This configuration has a potential isolation issue between splitter ports. Isolation between drop ports is not optimal and degrades at higher frequencies. Figure 20 shows a design alternative for a 4-port tap.

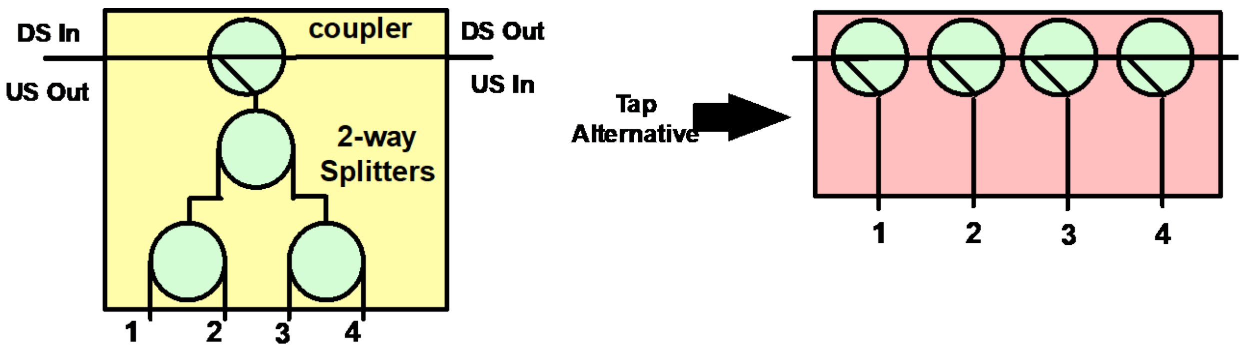

Figure 20 – High Isolation 4-Port Tap Design Alternative

7.2. Revisiting Need For Removable Tap Faceplates

Now that we have an approach to drastically reduce the inventory of taps, we need to ask the question again. Do we really need removable faceplates? Earlier, we discussed some of the reasons why we have removable faceplates. We will now explore the impact of not having them.

Properties of taps with removable faceplates include:

1) Changing tap values

2) Changing the number of drop ports

3) Switching the tap configuration between vertical and horizontal

4) Verifying proper length of center pin

5) Verifying proper contact of seizure screw

6) Maintaining connectivity to the elements downstream during faceplate removal

Changing tap values may no longer be necessary in an environment with a reduced number of tap types. The burden on inventory has been reduced with single-value taps for different numbers of ports. A reduced number of tap types also facilitates the development of a vertical-only housing and a horizontal-only housing. This provides the opportunity to explore using two types of tap housings, mainly horizontal and vertical connector entry tap housings (Figure 21).

Separate vertical and horizontal housings eliminate the need for technicians to mechanically configure a horizontal or a vertical connector entry option. The mechanism that enables dual connector entry options makes operation at higher frequencies more challenging.

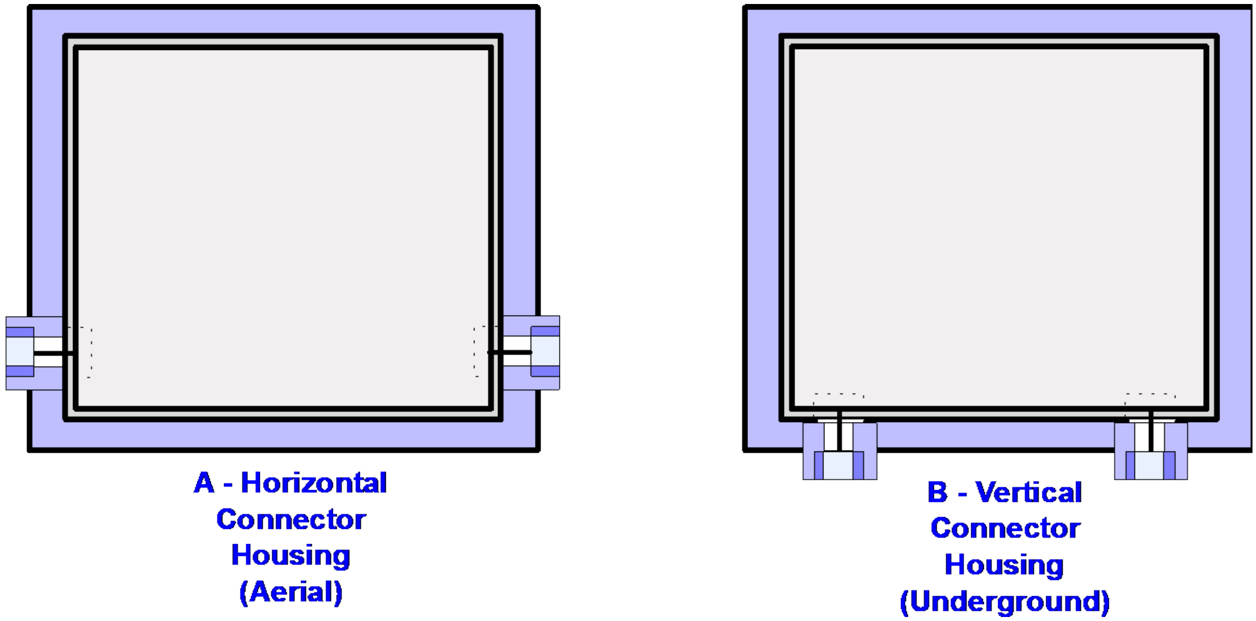

Figure 21 – Separate Housing SKUs For Horizontal (A) And Vertical (B) Taps

Specific connector entry (horizontal or vertical) implementation enables a mating tap connector permanently attached to the tap circuit board. This means that technicians do not need to trim the center pin at vendor specific lengths. It actually enables the use of a connector with fixed length center pin that screws to establish connectivity without the need for seizure screws. You can use a female connector attached to the tap and the board inside the tap and a male connector that attaches to the hardline cable (Figure 22).

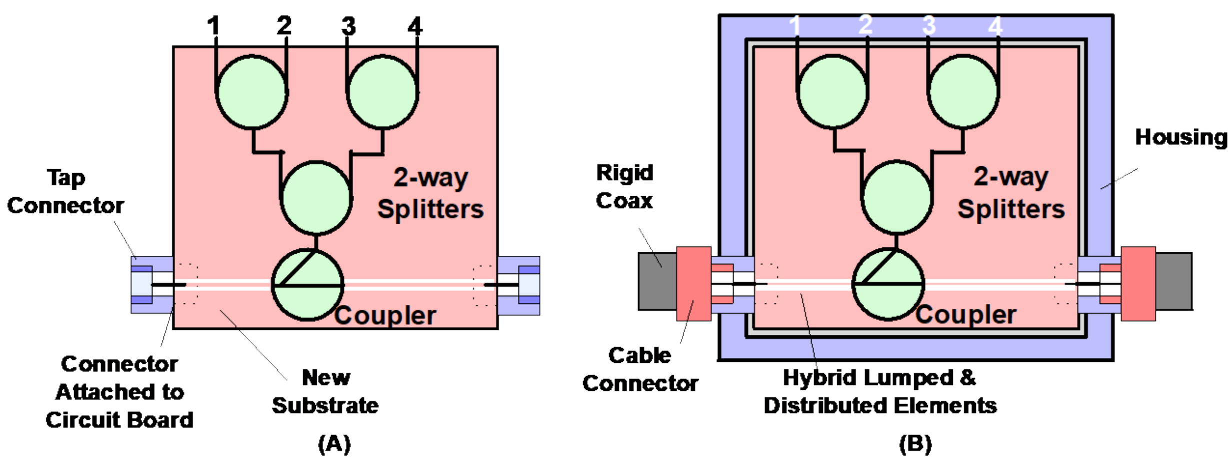

Figure 22 – Updated 4-Port Tap Design

There is no longer a need to verify contact and center conductor length by removing a faceplate. The tap can be implemented as a closed unit. Water leakage and radio frequency ingress problems would be drastically diminished. Challenges in implementing the switchable conductivity path when faceplates are removed are avoided. All circuitry can reside in one board and the assembly of a permanent connector that mates to a hardline connector would enable much higher frequency implementations.

One drawback of this proposed approach is that tap replacement results in a service outage for subscribers downstream from the tap. Reasons to require the changing of taps are diminished with the use of enclosed taps. In the rare event when a new premise needs connectivity. which was not anticipated during the original network design, and no spare drop port is available, you would need to replace the tap. Tap failure would also require a tap replacement, but in an environment where the taps are never opened this would be rare. In a fiber deep architecture, tap replacement is less of an issue as the network affected area is smaller. To replace a tap, two of these next generation hardline connectors are unscrewed to remove the whole tap, estimated at less than 1 minute of interruption.

As we move to higher frequencies, we need to be mindful of the natural tendency to hang on to traditional approaches. We have an inertia to continue using techniques that may no longer be the most efficient. Our environment is changing, we have been gradually pushing to higher and higher frequencies and the way we manufacture taps needs to evolve as well.

In the past, tap circuitry has been implemented using lumped circuit elements. However, as we look to support higher frequencies, incorporating distributed elements may be necessary. Similar questions must be asked regarding the tap circuit substrate. Fiberglass-based substrate, FR4, has been used in the past. Its low permittivity may result in higher loss and leakage at higher frequencies. Ceramic and PTFE (Polytetrafluoroethylene) based substrates should be explored. Their higher permittivity helps confine the RF energy and reduce leakage. While cable's traditional support of lower frequencies (< 1GHz) takes advantage of lumped element circuit components, hybrid lumped and distributed circuits next-generation designs could provide good performance at both higher and lower frequencies.

8. 100+ Gbps Experimental Setup and Demonstration

At this point, the discussion has focused on theoretical aspects of using higher portions of the coaxial spectrum leveraging simulations and modelling of a coaxial segment with cascaded taps. This section discusses experimental results obtained from an actual coaxial segment that has been built using cascaded taps linked by rigid coax cable. This network is actually a 50-ohm network, although the cables that were selected, have the exact attenuation versus frequency behavior of 0.54" rigid coax and RG6 flexible coaxial cable. This network leverages two types of transmitters. From 500 MHz to 3000 MHz a single-frame DOCSIS signal was composed in MATLAB and generated from an arbitrary waveform generator (AWG). This signal was configured with DOCSIS 1024 QAM subcarriers and the output was received in the last tap by a Rohde & Schwartz DOCSIS analyzer. A raw rate of 20 Gbps was estimated at the DOCSIS receiver.

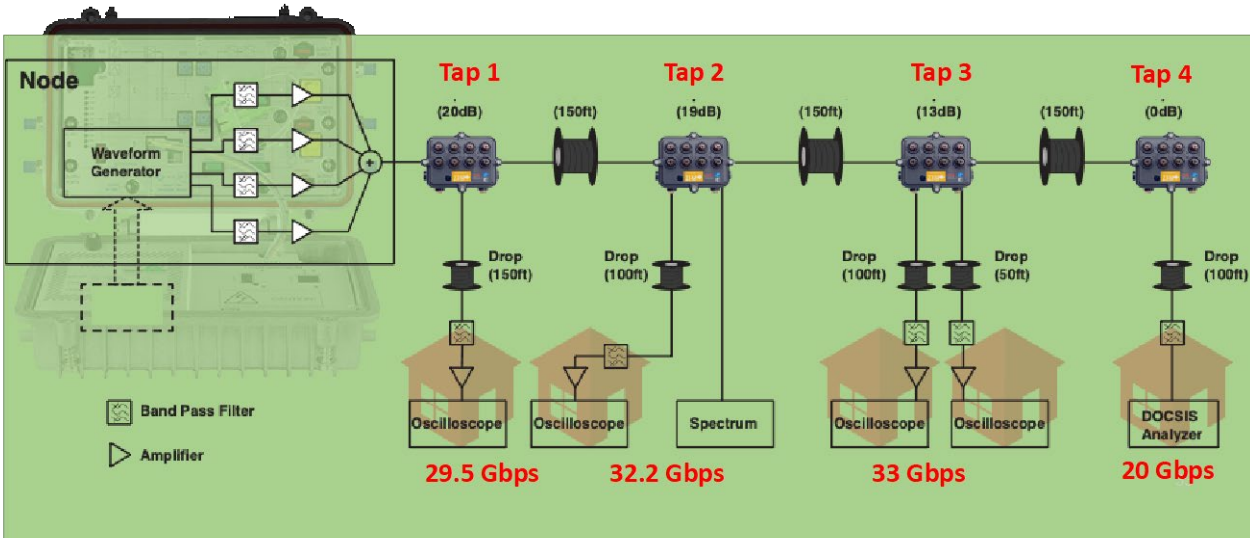

Figure 23 – Experimental Setup of > 100 Gbps Coaxial System

The rest of the spectrum covering frequencies of up to 12 GHz was occupied by re-designed DOCSIS-like OFDM symbols generated from AWGs (Figure 24).

Figure 24 – Spectrum Allocation Of Experimental Setup

A transmission band from 3 GHz to 6 GHz was assigned to Tap3, a band from 6.5 GHz to 9.5 GHz was assigned to Tap 2, and a band from 9.5 GHz to 12 GHz which was assigned to Tap 1 which enjoys the best CNR at higher frequencies. The portion of the spectrum between 6.0 GHz and 6.5 GHz was not used due to the aliasing signal from the AWG module operating at 12 GS/s. An optimized custom design could avoid that frequency gap. Capacity was estimated at 29.5 Gbps at tap 1, 32.2 Gbps at tap 2, 33 Gbps at tap 3 and 20 Gbps at tap 4. An aggregate capacity of 114.7 Gbps was obtained.

9. Node+1 Architectures

The analysis in previous sections examined a single coaxial segment. This capacity is therefore accessible to a Node + 0 architecture. Nevertheless, the use of higher frequencies doesn't have to be limited to N+0. In fact, as you further examine the performance of CMs that connect to the last tap and have a shorter drop length, such as in the case of CM_{12}, the spectral efficiency is quite high. An equivalent longer hardline segment with no drop connecting to an amplifier that follows can take advantage of an N+1 architecture. To highlight that higher cascade use, the spectral efficiency in CM_{12} is shown in red in Figure 25.

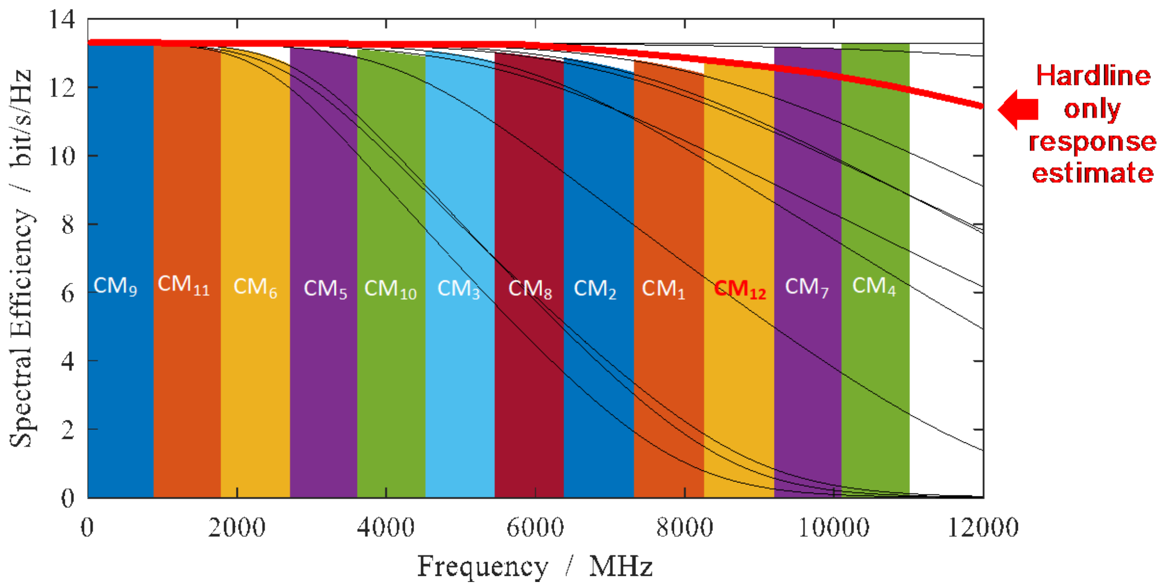

Figure 25 – Spectral Efficiency Available At End Of First Coaxial Segment N+1 HFOF Implementation

Figure 26 shows a N+1 concatenated coaxial segment with some sample scenarios regarding the bandwidths that could be implemented in the segment that follows the fiber node and the one that follows the N+1 amplifier. Figure 26 includes an aggressive scenario with 10 GHz bandwidth following the fiber node (Option 1), in addition to a more conservative 6 GHz bandwidth implementation (Option 2). Keep in mind that higher bandwidths are feasible if lower efficiencies are allowed or in shorter drop length scenarios.

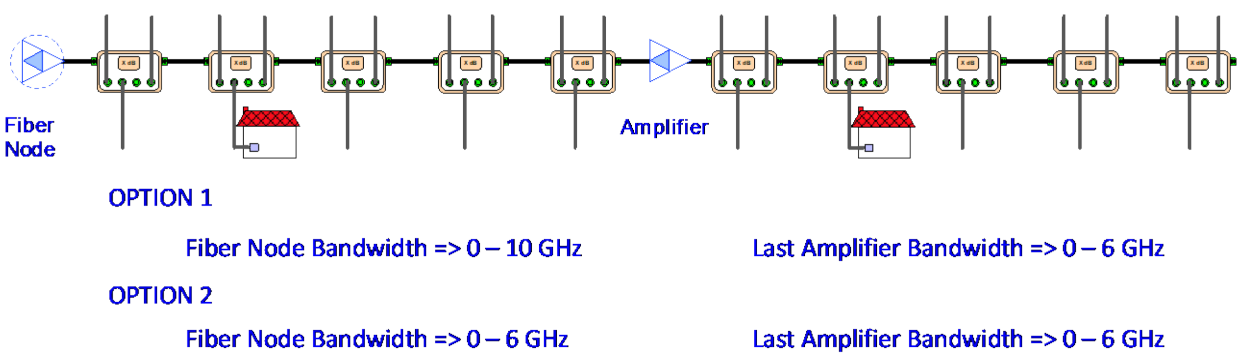

Figure 26 – Coaxial Segments in N+1 HFOF Implementation

9.1. Longer Coaxial Segments

As HFC networks evolve to fiber deeper. An alternative to Node+1 is to use longer coaxial segments. In order to avoid having too many optical nodes in a N+0 architecture, we have been pushing very high gains and power levels out of our optical nodes. We are doing this so that the furthest home can access all the resources that the CMTS or the RPD or RMD makes available. If we submit to the HFOF philosophy, you don't have to expect that every CM can handle the entire spectrum available. It is OK not to be able to consume all the resources. A consequence of this is that you could afford longer passive segments. We are calling this topology "N+0-Long". The intelligence of the evolved scheduler will help you with the appropriate resources based on the capabilities of each endpoint.

10. New Kind Of "Hybrid" Fiber Coax

As we have seen how we can still achieve significant capacity leveraging coaxial resources, in order to balance the cost-complexity of the end devices or CMs, one has to determine how much coaxial aggregate capacity and how much peak capacity are practical. Even though 100 Gbps is achieved as an aggregate from a coaxial serving area, the complexity of a CM capturing significant RF bandwidth has to be assessed when determining its practical peak rate. We estimate that 25 Gbps could be a practical peak rate target or 50 Gbps when stacking two receivers. An alternative to coax is the use of coherent PON or CPON. The emergence of coherent optical innovations in the access environment along with the cost reduction trends of coherent optical components, make CPON an attractive long term access solution. The cost of deploying fiber deeper and fiber-to-the-home varies significantly among operators and even within an operator, which may lead an operator to different coax versus fiber deployment strategies. This will depend; on the specifics of the starting point scenario to evolve towards FTTH, on the availability of conduits, on the cost of deploying fiber drops, on economical and operational aspects for extending the life and frequency range of coax and many other dependencies. For some operators, it may make sense to migrate directly to FTTH and for others leveraging the existing coax may make economical sense. The evolution of fiber deeper and FTTH could be made on an as needed basis. A technology like CPON which allows users to reach 100 Gbps on typical Hub-to-subscriber lengths, supporting split ratios of up to 512, could be leveraged to support subscribers demanding higher peak rates. There are many users that may not need peak rates above 25 Gbps for quite some time. Now that fiber is penetrating much deeper in cable networks, the high-end users requiring high peak rate services are a long optical drop away which could be implemented on a success basis. Operators can design the ultimate fiber-to-the-home network but only deploy it partially based on where the high-end customers are. This would result in a gradual transition towards FTTH depending on where the demand is. In some places, there may not ever be such a demand. A CPON network could feed "Extreme Coax" nodes, base stations, enterprise and residential high-end users. Figure 27 shows such a "Hybrid" network where CPON and an N+1 "Extreme Coax" network are jointly leveraged to address subscribers' long term demand of peak and aggregate data rates.

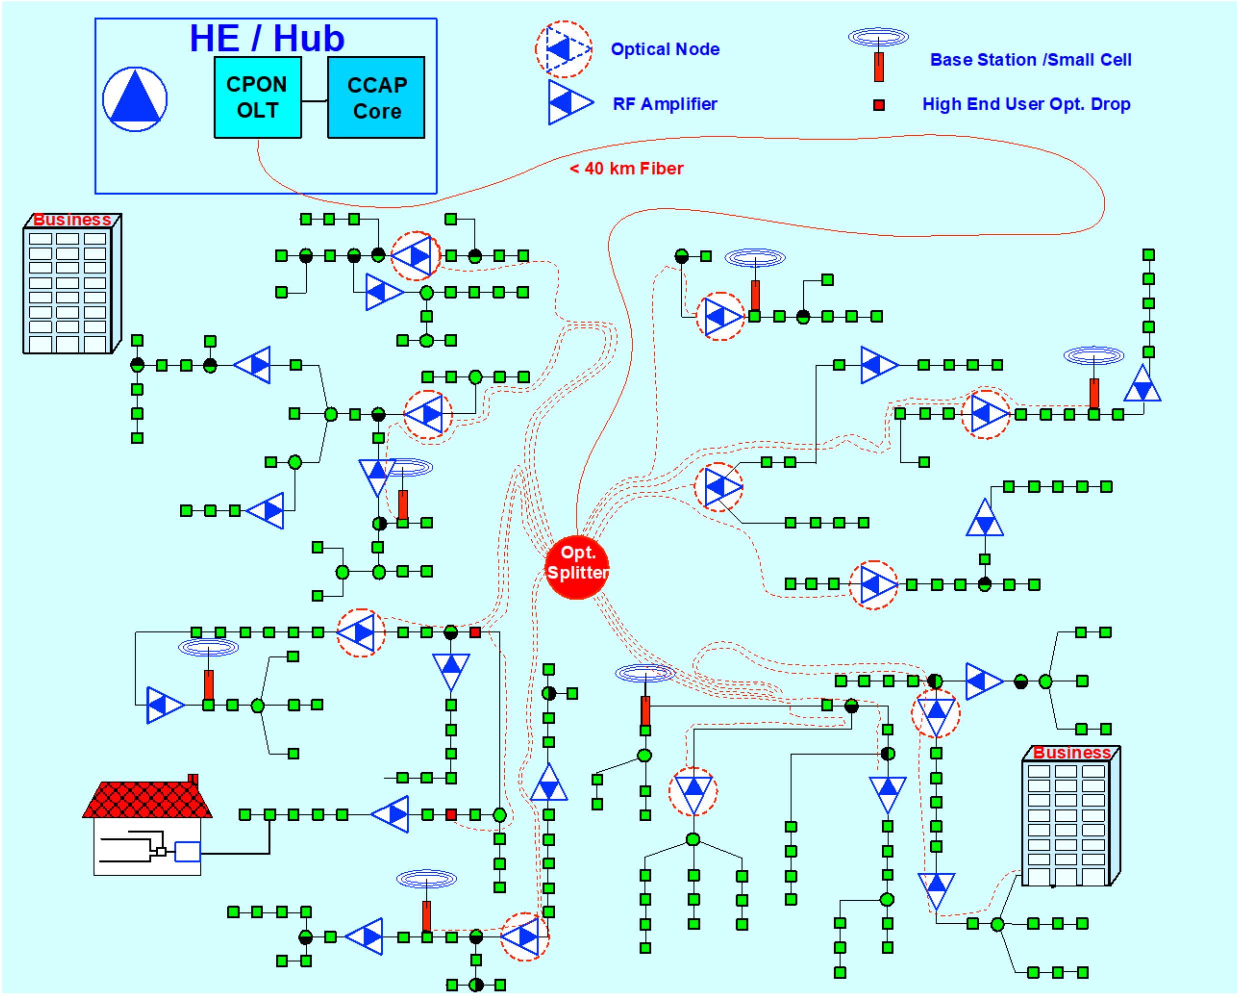

Figure 27 – Ultimate CPON/Extreme Cable "Hybrid" Fiber Coax Network

11. Conclusion

We have finally said goodbye to analog video and we need to fully embrace the digital era with all of its benefits, including the opportunity it provides in re-designing our coaxial network. The CMTS being the device that controls resource allocation still has plenty of room for improvement in this new environment.

As Cable entertains the support of 1.8 GHz, 3 GHz and even higher frequencies, the coaxial cable medium exhibits greater dependency in frequency. Having frequency-aware resource allocation provides great strategic advantage, helping enhance our data delivery capabilities over coax. Being free from analog video restrictions provides Cable the opportunity to drastically simplify the implementation of its coaxial infrastructure while preparing it to evolve to higher frequencies. Except for end-of-line taps, Cable can follow a single value tap for each tap type with the same number of ports. This reduction in inventory makes attractive horizontal- and vertical- specific taps, as well as taps without removable faceplates, avoiding many of the challenges in the evolution to higher frequencies.

A Higher Frequency Off First (HFOF) approach to allocating bandwidth has been proposed to optimize how we can use higher frequency resources as well as to facilitate the extension of coaxial segment lengths. This approach is not limited to N+0 architectures but can also be used with N+1 and higher cascade scenarios. Coaxial bandwidths greater than 100 Gbps have been demonstrated over a coaxial segment using HFOF and leveraging frequencies approaching the cut-off frequencies of the hardline cable (11 GHz). Balancing the capture bandwidth of the CM versus its tunability allows the optimization in the system's cost-complexity through peak versus aggregate rate assessment. The proposed HFOF approach also bounds the system's dynamic range. Frequency aware scheduling, HFOF, single value tap and high frequency tap redesign are key ingredients to this Extreme Cable approach. Together these concepts are powerful, but they could also be used independently and provide benefit to the evolution of our coaxial environment.

A new "Hybrid" Fiber Coax environment where CPON and Extreme Cable join forces to deliver data services is considered as a gradual, success-based transition to FTTH in the areas where it is needed.

Abbreviations

AWG	arbitrary waveform generator
CAGR	compound annual growth rate
CATV	cable television
CCAP	converged cable access platform
CM	cable modem
CMTS	cable modem termination system
CNR	carrier to noise ratio
CPON	coherent passive optical network
dB	decibels
DOCSIS	data over cable service interface specification
DS	downstream
FEC	forward error correction
FR4	flame retardant 4 circuit
FTTH	fiber to the home
Gbps	gigabit per second
GHz	gigahertz
HE	headend
HFC	hybrid fiber coax
HFOF	higher frequency off first
HHP	household passed
KS	klemmschrauben (clamp screw)
MAC	medium access control layer
MER	modulation error ratio
MHz	megahertz
OLT	optical line terminal
ONU	optical network unit
PHY	physical layer
PON	passive optical network
PTFE	polytetrafluoroethylene
QAM	quadrature amplitude modulation
RF	radio frequency

RG	radio grade
RPD	remote PHY device
RMD	remote MAC-PHY device
Rx	receiver
SNR	signal to noise ratio
STB	set-top-box
TE	transverse-electric
TEM	transverse-electromagnetic
TV	television
Tx	transmitter
US	upstream

Bibliography & References

DOCSIS 3.1 Physical Layer Specification. CM-SP-PHYv3.1-I18-210125, January 25, 2021, Cable Televisions Laboratories, Inc.

Microwaves101.com, https://www.microwaves101.com/encyclopedias/coax-cutoff-frequency

DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification. CM-SP-MULPIv3.1-I21-201020, October 20, 2020, Cable Televisions Laboratories, Inc.

DOCSIS 4.0 Physical Layer Specification. CM-SP-PHYv4.0-I03-201202, December 02, 2020, Cable Televisions Laboratories, Inc.

Z. Jia and L. A. Campos, "Coherent Optics Ready for Prime Time in Short-Haul Networks," in IEEE Network, vol. 35, no. 2, pp. 8-14, March/April 2021, doi: 10.1109/MNET.011.2000612.

J. Zhang, Z. Jia, M. Xu, H. Zhang and L. Alberto Campos, "Efficient preamble design and digital signal processing in upstream burst-mode detection of 100G TDM coherent-PON," in IEEE/OSA Journal of Optical Communications and Networking, vol. 13, no. 2, pp. A135-A143, February 2021, doi: 10.1364/JOCN.402591.

J. Zhang, Z. Jia, M. Xu, H. Zhang, L. A. Campos and C. Knittle, "High-Performance Preamble Design and Upstream Burst-Mode Detection in 100-Gb/s/λ TDM Coherent-PON," 2020 Optical Fiber Communications Conference and Exhibition (OFC), 2020, pp. 1-3.