PREDICTIONS ON THE EVOLUTION OF ACCESS NETWORKS TO THE YEAR 2030 AND BEYOND

TOM CLOONAN, CTO, NETWORK SOLUTIONS
MIKE EMMENDORFER, SENIOR DIRECTOR, SYSTEMS ENGINEERING
JOHN ULM, ENGINEERING FELLOW
AYHAM AL-BANNA, DISTINGUISHED SYSTEM ENGINEER
SANTHANA CHARI, SENIOR DIRECTOR, VIDEO ENGINEERING

ARRIS™

TABLE OF CONTENTS

INTRODUCTION 3
BACKGROUND 3
Traffic Capacity Trends for the Future 3
Bandwidth Pressures for the Future 4
ANALYSIS OF HFC PLANT EVOLUTION 5
Bandwidth Extrapolations & Node Split Efficacy In Future HFC Plants 5
USE OF ALTERNATIVE TECHNOLOGIES IN OR AROUND THE HFC PLANT 20
Overview 20
Improvements Provided By DOCSIS 3.1 21
Improvements Provided By Traditional Headend-based CCAP Systems 22
Improvements Provided By Distributed Access Architectures (DAAs) 23
Improvements from RFOG and Extended-Spectrum RFOG 26
Improvements from New Blends of Video Delivery Techniques 28
Improvements from New Blends of Video Compression Techniques 29
Challenges from New Blends of Video Resolution Techniques 31
Improvements Provided By EPOC 31
Improvements Provided By PON 32
Improvements Provided By Point-to-Point Ethernet 32
PREDICTIONS 33
Defining the Approach 33
Scenarios Impacting All MSOs 34
Scenarios Impacting MSOs Who Will Extend the Life-Span of Their Current HFC Plant 34
Scenarios Impacting MSOs Who Will Switch to New Technologies 36
CONCLUSIONS 37

INTRODUCTION

As the Cable Industry moves forward into the future, MSOs will experience profound changes in their HFC plants and in their headend equipment. These changes will be required to support the challenging bandwidth growth that is expected within all of the MSO service types, including DOCSIS HSD, DOCSIS IP Video, SDV, VoD, nDVR, and Digital Broadcast Video (SD, HD, 4K, and 8K resolutions).

These expected changes are forcing many MSOs to wonder about the ultimate evolution of their headends and HFC plants. Many questions abound. Will headend-based CCAPs be able to support the future bandwidth per RF port? Will headend-based CCAPs be able to support the number of RF ports needed after many rounds of nodes-splits? Will the power and rack-space requirements of future CCAPs fit within headend budgets? Will MSOs need to use the bandwidth expansion capabilities offered by DOCSIS 3.1? Will MSOs need to use the higher-order modulations offered by DOCSIS 3.1? Will Digital Optics be required to enable the higher-order modulation in DOCSIS 3.1? Will Distributed Access Architectures be required to provide the required scale? If so, which Distributed Access Architectures will likely be utilized? Will it be Remote PHY? Will it be Remote CCAP? Will RFoG offer needed benefits in the future? Will EPON and EPOC offer needed benefits in the future? When do the traffic engineering analyses predict that transitions will take place between these different technologies?

This paper will look into the crystal ball and use basic trend analyses and traffic engineering analyses to create predictions that attempt to answer all of these challenging questions.

BACKGROUND

Traffic Capacity Trends for the Future

By its very nature, traffic engineering is an imperfect science requiring many guesses and assumptions and approximations to be utilized. Sometimes, these guesses can be wrong. However, when implemented properly, traffic engineering predictions can be very useful in developing rough plans for future HFC networks and in the sizing of future HFC networking equipment. It may be valuable to explore some of the current predictions coming from recent traffic engineering work.

Using empirical data collected in the field, trending information on all service types can be analyzed in an attempt to predict the future. However, after doing this for most services, the authors believe that the most problematic service type within the MSO network of the future will undoubtedly be the DOCSIS HSD and DOCSIS IP Video service types. We will therefore focus on this particular service type within this section.

For the past several years, studies have indicated that DOCSIS HSD Downstream traffic has been experiencing a ~50% compound annual growth rate (CAGR). For almost 30 years, this growth rate has shown itself in the Maximum Downstream Sustained Traffic rates (aka the "Billboard Bandwidths") that service providers have offered to their subscribers, and it has also shown itself (with slightly more variation) in the Average Downstream Bandwidth Consumption rates that subscribers have consumed for at least the past 10 years. Upstream Billboard Bandwidths and Average Upstream Bandwidth Consumption Rates display much more variation, with CAGRs at different MSOs showing rates ranging from 10% to 30%.

There are some who believe that the Average Bandwidth Consumption rate of the future may slow as homes begin to receive adequate bandwidth to offer an HD video feed to every pair of eyeballs within the home. There have even been recent slow-downs in the Average Downstream Bandwidth Consumption rates witnessed by some MSOs that seem to validate this belief. Thus, it is very possible that the CAGRs may be reduced below 50% in the future. However, this potential slow-down in Average Bandwidth Consumption growth rate will likely be accelerated again as 4K (3840x2160 resolution) and 8K (7680x4320 resolution) UHD TV begins to gain acceptance in the future and be passed over DOCSIS networks. At a minimum, it is expected that 4K video feeds will likely be viewed on the main television screen in many homes within the next five years. These new video resolutions could lead to increases of up to 16x the required bandwidth of today's 1080p video resolution. In addition, new, yet-to-be-invented machine-to-machine applications of the future may also contribute to increased DOCSIS HSD Bandwidth Consumption demands.

Bandwidth Pressures for the Future

The MSO's competitive landscape has changed rapidly in the last five years- especially from Over-the-Top (OTT) video providers such as Apple TV, Amazon, Hulu, Netflix and others entering the On-demand video market. The resulting growth in OTT video usage has also led to a requirement for increased investment within the MSO's DOCSIS HSD network due to increased consumer bandwidth usage. Bandwidth growth from these OTT providers switching to 4K or 8K video will likely drive more bandwidth demands in the future.

Verizon FiOS and other FTTH providers that have rolled out networks over the past several years will also remain a threat to the MSO's triple play offering in the future. As an example, it was reported that Verizon will consider an upgrade to their FiOS network to the next generation PON technology that may be capable of offering 10 Gbps downstream and 2.5 Gbps upstream traffic. This would increase the bandwidths available in competing service providers, placing a new stress on the bandwidth demands within DOCSIS networks.

All of these trends will undoubtedly raise the height of the competitive bar that MSOs will have to try to match or beat over time. Most MSOs plan to counter these threats with smooth and sustained bandwidth capacity growth in their HFC network that matches the needs of their subscribers in a just-in-time fashion.

ANALYSIS OF HFC PLANT EVOLUTION

Bandwidth Extrapolations & Node Split Efficacy In Future HFC Plants

The existing HFC cable network infrastructure provides high levels of digital capacity (limited to ~6 Gbps to the home and perhaps ~100 Mbps from the home in pre-DOCSIS 3.1 systems, and limited to 10-15 Gbps to the home and perhaps 1+ Gbps from the home in post-DOCSIS 3.1 systems). The HFC network also offers great flexibility.

If an MSO uses this capacity carefully, it can likely compete well for many years to come. The cable industry is making investments in IP-based video delivery technology and expanding the high-speed Internet IP capacity as well.

The coaxial network is very nimble, and changes to modulation types, Forward Error Correction techniques, node sizes, spectral splits, and upper frequency limits may increase the spectrum allocation beyond the current levels in both directions. The capacity needed in each direction is projected into the future in this section. We will consider the Downstream and Upstream bandwidth demand trends, and then we will couple those trends with node-split changes that MSOs may opt to utilize. We will also consider the impact of new modulation types and Forward Error Correction techniques that are being provided by the DOCSIS 3.1 specification.

Node splits have long been a trusted tool used by many MSOs throughout the years to reduce the bandwidth demands within a Service Group. The basic idea behind the node split is that it divides the subscribers connected to a single Fiber Node into two groups (Group A and Group B), and the subscribers in Group A are re-connected to a single (smaller) Fiber Node and the subscribers in Group B are also re-connected to a different (smaller) Fiber Node. Thus, two separate Fiber Nodes (and the associated feeds for two separate Fiber Nodes) are required to support the bandwidth for the pool of subscribers, so there is a cost associated with the node split.

Node splits offer no change in the Service Group bandwidth requirements for Broadcast services—if the MSO needed 50 QAMs to support Broadcast Video prior to the node split, then the MSO will still require 50 QAMs to support Broadcast Video after the node split.

However, the principle benefit of the node split is associated with Narrowcast services (Switched Digital Video (SDV), VoD, DOCSIS HSD, DOCSIS VoIP, and DOCSIS IPTV). The benefit of the node split is primarily derived from the fact that, in the past, the Narrowcast bandwidth required on the HFC spectrum was roughly halved every time a node split was implemented. Node splitting oftentimes permitted MSOs to "free up" substantial amounts of Narrowcast spectrum whenever they performed the node split. For example, an MSO who had ~30 QAMs of Narrowcast could free up ~15 QAM channels (~600 Mbps) with a single node split.

Many MSOs envision the number of Narrowcast QAMs growing over time because they expect increased QAM counts associated with VoD (as VoD gains popularity and includes more HD video content), SDV (as SDV adds more HD content), and HSD (as users' Internet bandwidth demands continue to rise). Of these different Narrowcast services, the highest growth rate going forward will undoubtedly be associated with DOCSIS HSD (which may continue to experience a 50% CAGR in its Average Downstream Bandwidth levels as defined in the previous section). Without the use of node splits, an MSO who had 8 DOCSIS Downstream channels consuming some of the HFC spectrum in 2013 would likely experience DOCSIS channel growth (assuming a 50% CAGR) as shown below.

'13	'14	'15	'16	'17	'18	'19	'20	'21	'22
8	12	18	27	41	61	92	137	205	308

Table 1. Year & Number of QAM Channels Required For HSD Growth (No Node-Splits)

Obviously, this type of uncontrolled bandwidth growth would fully consume any reasonably-sized HFC spectrum within ten years. As a result of this anticipated problem, many MSOs plan to split nodes to control this Narrowcast QAM growth. The theory is that carefully-timed node splits can halve the QAM count associated with the DOCSIS HSD channels and can keep the total DOCSIS HSD QAM counts to reasonable levels. For example, if an MSO with a 750 MHz (115 QAM) plant wanted to keep the DOCSIS QAM count to levels below 80 channels (assuming the other 35 QAMs were reserved for Digital Video services), then the commonly-held belief is that, for the system shown above, node splits timed to occur in 2019 and 2021 would yield the desirable results shown below.

‘13	‘14	‘15	‘16	‘17	‘18	‘19	‘20	‘21	‘22
8	12	18	27	41	61	46 (split)	69	52 (split)	78

Table 2. Year & Number Of QAM Channels Required For HSD Growth (With Node-Splits in 2019 & 2021)

There are, however, two potential flaws in the above logic that may cause MSOs to experience undesirable issues with these planned future node splits. The first issue is that node splitting becomes more and more expensive each time a round of node splits is implemented—primarily because the number of Fiber Nodes that need to be split increases by a factor of two with each round of node splits. This may preclude MSOs from performing the node splits at the rates that they desire.

The second issue is that node splits do not help at all to reduce the QAM counts if the QAM counts are being driven primarily by the Maximum Sustained Traffic Rates (aka the Billboard Bandwidths) that the subscribers are being offered by MSO Marketing teams. These Billboard Bandwidths have also been growing at a 50% CAGR for quite a few years, and if that growth rate continues in the future, it will require that many DOCSIS HSD QAMs will have to be offered on the HFC spectrum within each Service Group *even if* a large number of node splits are implemented. In fact, even in the hypothetical case where numerous node splits are used to reduce the number of subscribers within a Service Group to only one subscriber, a high Billboard Bandwidth of, say, 4 Gbps would require 100 DOCSIS QAMs to be fed into the Service Group- even if the average bandwidth consumed by that single subscriber was only 100 Mbps. This is an interesting phenomenon that requires more study as MSOs begin to plan their node split strategies for the future.

In traffic engineering simulations carried out by the authors, preliminary results indicate that the amount of DOCSIS HSD capacity which might be required to provide an “adequate” Quality of Experience level to a Service Group can be roughly described by simple formulae. One form of the formula is given by:

Required Service Group Bandwidth Capacity = S*Tavg + Tmax (1)

Where S is the number of subscribers within the Service Group, Tavg is Per-Subscriber Average Busy-Hour Bandwidth, and Tmax is the Maximum Sustained Traffic Rate (Billboard Bandwidth) offering to the subscribers.

It should be noted that other more complex (and more accurate) formulae are also under study, but this simplified formula will be utilized extensively within this paper to make some key points. To make use of this formula, we will need to make some crystal ball predictions about the future values of S, Tavg, and Tmax. This is challenging, but we will turn to historical data to help make these predictions.

One can look back at the DOCSIS HSD Billboard Bandwidth (the Tmax values) and the DOCSIS HSD Average Downstream Bandwidth Consumption rates for a single subscriber (the average Tavg values) on a yearly basis. One can then assume that the trends of the past will continue into the future and then extrapolate the resulting curves into the future. This assumption may or may not be true, but we will use it as a baseline, and then we will also explore what might happen if the assumption does not hold true.

One can also predict the per-Service Group Aggregate Average Consumption Bandwidth (Tagg) levels for Service Groups of various sizes, which is simply the per-subscriber Average Downstream Bandwidth Consumption (Tavg) times the number of subscribers (S) within the Service Group. Obviously, Service Groups with different sizes (S) will have different Aggregate Average Consumption Bandwidths (Tagg). In general, Tagg is given by S*Tavg (which is the first term in Formula 1 above).

These predicted, future values are shown in Figure 1 for a traffic model that assumes the 50% CAGR of the past continues for both Tavg and Tmax values moving forward into the future.

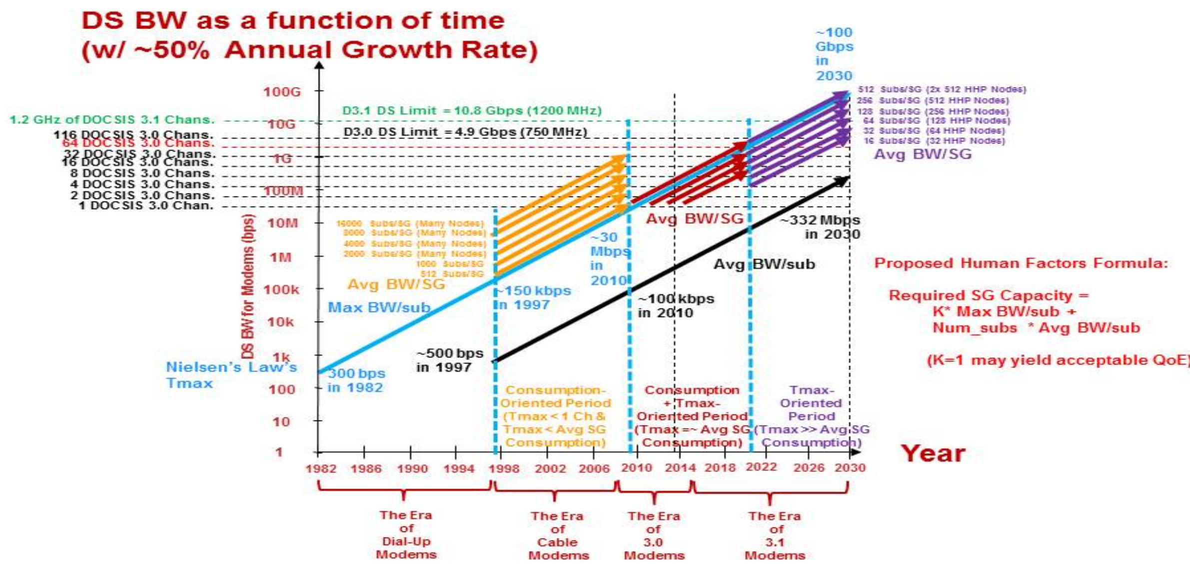

Fig. 1 DOCSIS HSD Downstream Traffic Engineering Predictions with 50% CAGR In Future

There is quite a bit of useful information provided by the plots in Figure 1. The plot has the Downstream Bandwidth displayed in a logarithmic fashion on the y-axis, and it has the years ranging from 1982 to 2030 on the x-axis. These years cover several different eras of modem service, including the Dial-Up Modem era, the Cable Modem era, the 3.0 Modem era, and the 3.1 Modem era.

The first plot to explore is the light blue plot, which shows Nielsen's Curve that identifies the expected Billboard Bandwidth (Tmax) values on a year-by-year basis. This plot illustrates that MSOs will likely have to provide higher and higher Tmax values to their subscriber pool on a yearly basis. If these trends continue, then Nielsen's Curve predicts that the Tmax value for a high-end modem may be on the order of 100 Gbps by 2030. One may wonder what applications could possibly require that kind of bandwidth 16 years from now. The honest answer is that we do not know what those applications will be—nobody does. But it is probably fair to say that the college students of today will help to invent those impossible-to-identify applications. It is also fair to say that we are probably in the same position as an Internet user of 1998. Back then, they probably had no idea about the types of services and applications that would ultimately be developed that would require high-end users to require 100 Mbps of service by the year 2014 (16 years after 1998). The 100 Mbps service level of today must have appeared astronomical to a 1998 user who was only receiving ~200 kbps of bandwidth capacity at the time. But if trends continue, the Tmax values are predicted to grow quite high within the 2020 decade.

The second set of plots (in Figure 1) that should be explored is the plot set illustrating Average Bandwidth Consumption rates as a function of time. The black plot illustrates the approximate average per-subscriber bandwidth consumed by a single subscriber during the busy-hour period of time (8pm - 9pm). This plot calculates an average value using contributions from both active and inactive users, so the average values (Tavg) are much lower than the Tmax values. In fact, the Tavg values for a single subscriber tend to be ~300 times lower than the corresponding Tmax values within a given year! However, the Tavg values in this chart are assumed to grow with a 50% CAGR, so the plotted values start in 1997 with a ~500 bps value and then grow to be ~100 kbps by 2010 and ~500 kbps by 2014 (a typical high-end number for today). The plot shows that if trends continue, then the average bandwidth per subscriber will grow to be ~332 Mbps by 2030.

While a plot for a single subscriber is interesting, it is not very useful for traffic engineering analyses looking at aggregate traffic patterns for many subscribers in a shared pool (like a Service Group). The single subscriber plot (shown in black within Figure 1) can, however, be scaled upwards in a linear fashion to yield the average aggregate bandwidth (Tagg) for a larger pool of subscribers that can share the DOCSIS HSD channels within a Service Group (SG). The orange, maroon, and purple plots contain that scaled up information for Service Groups of varying size that were utilized in past years or that will likely be utilized in future years.

The Service Group sizes shown in the figure include those with 16,000 subscribers per Service Group (in orange), which were utilized in the 2000 time-frame and could have corresponded to a Service Group containing ~24 Fiber Nodes of 2000 HHP each with a

33% take-rate on the DOCSIS HSD service. The Service Group sizes within Figure 1 also include those with as few as 16 subscribers per Service Group (in purple), which may be found in the 2020 decade if MSOs perform heavy node splits and create many 32-HHP Fiber Nodes with a 50% take-rate on the DOCSIS HSD service. It should be noted that two adjacent, parallel lines of color (orange, maroon, or purple) are exactly one node split apart from one another in terms of the bandwidth capacity required to support the Narrowcast services. (Note: The use of 32-HHP Fiber Nodes may be implemented by some MSOs in preparation for an ultimate transition to FTTH architectures in the 2030 time-frame. This transition to FTTH might make use of technologies such as RFOG, PON, or Point-to-Point Ethernet).

The family of curves illustrating the growing Service Group Average Bandwidth levels is divided into three different operating regimes, which in turn define three different periods of time in the life of the HFC plant. Within this paper, the three periods are labeled:

- The Consumption-oriented Period from 1997-2009 (illustrated in orange)
- The (Consumption + Tmax)-oriented Period from 2009-2021 (illustrated in maroon)
- The Tmax-oriented Period from 2021-2030 (illustrated in purple)

The Consumption-oriented Period occurred in the early days of DOCSIS deployment. It was characterized by a period of time when Service Groups were quite large (containing many Fiber Nodes) and each Service Group was serviced by a single DOCSIS Downstream channel. This approach worked quite well in its time, because the Tavg value of each individual subscriber was quite low, and it required the pooling of a lot of subscribers into a large Service Group to generate enough bandwidth to efficiently utilize the bandwidth capacity offered by a single DOCSIS Downstream channel (30-40 Mbps).

The Tmax values of the time were also extremely low; they were typically much lower than the Service Group's Aggregate Average Bandwidth Consumption (Tagg) values and much lower than the ~30-40 Mbps bandwidth capacity offered by a single DOCSIS Downstream channel. So Tmax values were practically negligible in the traffic engineering efforts of those days. As a result, the total bandwidth capacity required by a Service Group was dominated by Tagg value (which was equal to the product of the Tavg value and the S value representing the modem count).

During the Consumption-oriented Period, the Tmax values were much lower than the Service Group's Aggregate Average Bandwidth Consumption (Tagg) level. Thus, the Tmax values could generally be ignored during the Consumption-oriented Period, and Equation (1) could be reduced to an even simpler formula

Required Service Group Bandwidth Capacity = S*Tavg (2)

As a result, the main task of the traffic engineer in the days of the Consumption-oriented Period was to identify the Service Group's Aggregate Average Bandwidth (Tagg) for the pool of subscribers sharing the single DOCSIS channel and schedule Service Group "de-combining" activities to remove Fiber Nodes from Service Groups whenever the Service Group's Aggregate Average Bandwidth Consumption (Tagg) values began to rise to levels that were close to the ~30-40 Mbps capacities offered by the single DOCSIS Downstream channel. (Note: Oftentimes, this "de-combining" was triggered if the Service Groups Aggregate Average Bandwidth Consumption in the busy-hour grew to be ~70% of the capacity of the single DOCSIS Downstream channel).

Many MSOs transitioned from the Consumption-oriented Period to the (Consumption + Tmax)-oriented Period in the 2009 time-frame. This transition was typically implemented in conjunction with the MSO transition to DOCSIS 3.0 channel-bonding CMTS equipment. The (Consumption + Tmax)-oriented Mode of operation was characterized by the fact that a single DOCSIS Downstream channel was no longer adequate to support the rising Tmax values required by the subscribers within the Service Group. As a result, MSOs needed to utilize bonded Downstream channels to provide adequate bandwidth capacity to each Service Group. The (Consumption + Tmax)-oriented Period was also characterized by another interesting change—the calculation of the total required bandwidth capacity for a single Service Group could no longer ignore the impact of the Tmax value, because the Tmax values and the Service Group's Aggregate Average Bandwidth Consumption (Tagg) levels tended to be of the same order of magnitude. This seems to be true for Service Group sizes of 1024-HHP per Service Group (512 subs per Service Group), 512-HHP per Service Group (256 subs per Service Group), and 256-HHP per Service Group (128 subs per Service Group). As a result, the calculations for Equation (1) required the traffic engineer to consider both terms within the formula. Since the Tmax value and the Service Group's Aggregate Average Bandwidth Consumption level (Tagg = S*Tavg) tended to be about the same in magnitude, many MSOs modified this formula to produce a much simpler approximation formula given by:

Required Service Group Bandwidth Capacity = 2*Tmax (3).

The simplified and approximate formula in Equation (3) is roughly correct as long as the Tmax value is approximately equal to the Tagg (= S*Tavg) value, which is interestingly true for many MSOs right now. By definition, this constraint is roughly satisfied during the (Consumption + Tmax)-oriented Period. The (Consumption + Tmax)-oriented Period will likely exist for many years to come. During these years, MSOs will undoubtedly add more bandwidth capacity (channels) to the DOCSIS service tier to stay ahead of the Tmax growth and the Tagg (= S*Tavg) growth, and they may also perform node splits to periodically reduce the Service Group's Aggregate Average Bandwidth Consumption Tagg (= S*Tavg) level.

However, as a result of repeated node splits that reduce the Aggregate Average Bandwidth Consumption (Tagg) levels in a Service Group and as a result of expected increases in Tmax levels, there will come a point in the future when MSOs will likely transition from the (Consumption + Tmax)-oriented Period to the Tmax-oriented Period. At that point in time, the Tmax value will have grown to be much larger than the Service Group's Aggregate Average Bandwidth Consumption value (Tagg = S*Tavg). This is the period of time when the required Tmax values will begin to dominate the traffic engineering rules, and it is also the period of time when the Service Group's Aggregate Average Bandwidth Consumption level (Tagg = S*Tavg) may become practically negligible relative to the Tmax levels. This implies that the Tmax term in Equation (1) will be much larger than the (S*Tavg) term. Interestingly, node splits only impact the (S*Tavg) term of Equation (1) by reducing the number of subscribers (S) in the Service Group. Node splits do not impact the Tmax term at all. Thus, during the Tmax-oriented Period of the future, node splits will cease to have much impact on the required bandwidth capacity levels for a Service Group. This will have profound effects on the traffic engineering decisions that will be made by different MSOs in the future, with each MSO choosing a potentially different path.

To see why this is the case, it is best to consider Figure 2, which illustrates the Required Downstream Bandwidth Capacity required for different Service Groups of different sizes. (Note: Similar plots with similar dates can also be created for the Required Upstream Bandwidth Capacity, but due to space constraints, we will focus on the Downstream within this paper). This required bandwidth capacity value is calculated using the rule-of-thumb formula given in Equation (1) above. In essence, Figure 2 is the same as Figure 1, but Figure 2 shows the sum of the Tmax value and the Aggregate Average Bandwidth Consumption (Tagg = S*Tavg) value instead of showing each parameter individually. In Figure 2 the same three Periods of time are shown with the same Service Group sizes, but the bandwidth shown is the Required Downstream Bandwidth Capacity (which is equal to the sum of the Tmax and the Aggregate Average Bandwidth Consumption values (Tagg = S*Tavg)). When we compare Figure 2 to Figure 1, it becomes clear that the Service Group's Aggregate Average Bandwidth Consumption (Tagg = S*Tavg) values dominated during the Consumption-oriented Period. Within the figure, it also becomes clear that we are currently in an interesting period of time (defined as the (Consumption + Tmax)-oriented Period and shown in maroon) when the Required Downstream Bandwidth Capacity is driven by both the Tmax value and by the Service Group's Aggregate Average Bandwidth Consumption (Tagg = S*Tavg) value. Finally, it can also be seen that the future Tmax-oriented Period (in purple) may have Required Downstream Bandwidth Capacity values that will be dominated by the Tmax values, and the resulting required Bandwidth Capacity will be very close to the growing Tmax line (shown in light blue).

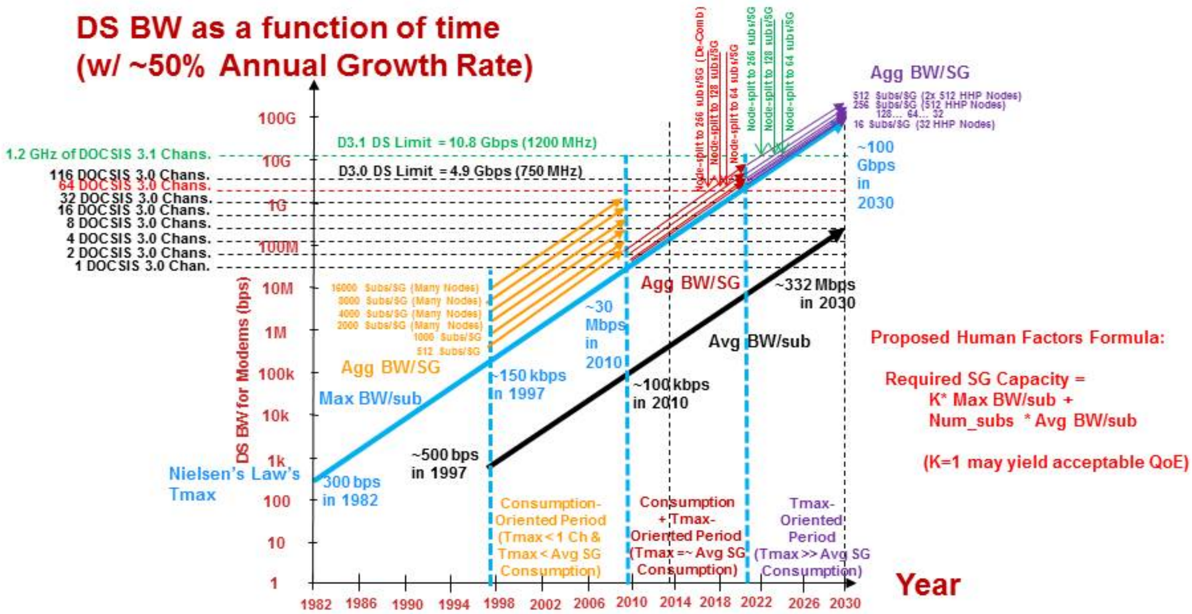

Fig. 2 DOCSIS HSD Required Downstream Bandwidth Capacity with 50% CAGR In Future

Each successive node split in the future will force an MSO's Required Downstream Bandwidth Capacity levels to drop from one curve to the curve directly below it. From the figure, it becomes clear that node splits had a very favorable effect on the Required Downstream Bandwidth Capacity during the Consumption-oriented Period and during the (Consumption + Tmax)-oriented Period, because the node splits tended to produce a ~50% reduction in the Consumption portion of the Required Downstream Bandwidth Capacity within the Service Group.

However, when MSOs move into the future and begin to operate in the Tmax-oriented Period, node splits will not be as effective as they once were. Successive node splits will yield less and less benefit as they produce smaller and smaller percentage reductions in the Required Downstream Bandwidth Capacity levels. This is primarily due to the growing dominance of Tmax values in the Required Downstream Bandwidth Capacity levels as Fiber Nodes are made smaller and smaller. Since the node splits only impact the (S*Tavg) value, they will likely not have much of an impact in the distant future assuming bandwidth growth trends and node split trends continue. For the assumptions used in our calculations above, one can easily calculate the percentage reductions for each node split (a single hop between adjacent curves), and these results are shown below in Table 3 below.

Node Split		**% Decrease in Req'd DOCSIS Downstream BW Capacity in the Service Group**
From (subs/SG)	**To (subs/SG)**	
16000 to 8000		49%
8000 to 4000		48%
4000 to 2000		47%
2000 to 1000		43%
1000 to 512		38%
512 to 256		32%
256 to 128		23%
128 to 64		15%
64 to 32		9%
32 to 16		5%

Table 3: Decrease in Required DS Bandwidth Capacity for Node Splits with Different Sized Fiber Nodes

Another way to look at the problem of diminishing returns on successive node splits is to study the amount of time that a node split "buys" for an MSO before another node split is required. To perform this study, let's assume that the MSO has decided that DOCSIS HSD will be given no more than 64 Annex B channels within the HFC spectrum because the rest of the spectrum is required for Analog, Digital Broadcast, SDV, VoD, and/or nDVR services.

Within Figure 2, the bandwidth capacities associated with several different numbers of Annex B DOCSIS channels are shown as dashed horizontal lines, with the number of DOCSIS channels associated with each line designated on the left-hand side of the figure. The bandwidth capacity associated with 64 Annex B channels is shown by the red, dashed, horizontal line. If the red line represents the MSO's ceiling that sets a limit on the number of DOCSIS channels, then every time the Required DOCSIS Downstream Bandwidth Capacity for a Service Group hits that ceiling, then some action (such as a node split) must be taken. Once the node split action is taken, the Required DOCSIS Downstream Bandwidth Capacity is reduced and "buys" the MSO some amount of time before the Required DOCSIS Downstream Bandwidth Capacity levels rise up to the ceiling again (requiring another node split). We can calculate the amount of time that each node split "buys" the MSO before another node split is required, and these results are shown below in Table below.

Node Split		# Months Each Node Split "Buys" the MSO Before Another Node Split Is Required (assuming a 50% CAGR in Tmax & Tavg)
From (subs/SG)	**To (subs/SG)**	
16000 to 8000		20.0
8000 to 4000		19.5
4000 to 2000		18.5
2000 to 1000		16.9
1000 to 512		13.9
512 to 256		11.2
256 to 128		7.7
128 to 64		4.8
64 to 32		2.7
32 to 16		1.5

Table 4: # Months before another Node Split for Node Splits with Different Sized Fiber Nodes

These diminishing returns on node split investments over time are concerning. It implies that at least for DOCSIS HSD services MSOs can expect to get smaller and smaller percentages of their DOCSIS HSD spectrum freed up by their successive node split activities if the conditions described above hold. It is even more concerning when coupled with the fact that successive node splits applied in a ubiquitous fashion across a market will become more and more expensive over time as the number of Fiber Nodes involved in the operation grows by a factor of two with each successive node split.

In addition, the average bandwidth utilization in a Service Group will become very low with more node splits occurring, because most of the bandwidth capacity within the service group will be added to simply provide the bandwidth capacity for the infrequently-used Tmax value. For the cases shown in Figure 2, the expected bandwidth utilizations within the different-sized Service Groups can be calculated, and these results are shown below in Table 5.

Service Group Size	Expected Average Bandwidth Utilization on DOCSIS Channels within a Service Group
16000 subs/SG	98%
8000 subs/SG	96%
4000 subs/SG	93%
2000 subs/SG	87%
1000 subs/SG	77%
512 subs/SG	63%
256 subs/SG	46%
128 subs/SG	30%
64 subs/SG	18%
32 subs/SG	10%
16 subs/SG	5%

Table 5: Expected Channel Bandwidth Utilization for Node Splits with Different Sized Fiber Nodes

Thus, as the Service Group size is decreased, the average channel utilization drops to extremely low levels due to the growing divergence between the Aggregate Average Bandwidth Consumption values and the Tmax values. Some MSOs may view this low channel utilization condition as being wasteful of the investment required to support the Required Bandwidth Capacity. One may wonder how MSOs will deal with this condition. One of the CCAP tools that can help is known as Output QAM Replication. It permits the channels within a CCAP chassis to be split in the digital domain and steered to more than one RF port on the chassis. This permits different service types (*e.g.*, DOCSIS HSD, VoD, and SDV) to have Service Groups of different sizes.

It should be noted that VoD and SDV are Narrowcast services that are quite different from DOCSIS HSD. VoD and SDV do not suffer from the complexities caused by the rising Tmax values in DOCSIS. VoD and SDV do not have anything like a Tmax value. With VoD and SDV service types in reasonably-sized nodes, any node split that halves the number of subscribers sharing the Service Group resources will also produce a similar halving of the Required Bandwidth for that Service Group. (Note: For smaller-sized nodes, this rule-of-thumb may break down).

In the Tmax-oriented Period for DOCSIS HSD, MSOs may find it interesting to explore having different Service Group sizes for VoD and SDV and DOCSIS HSD. As an example, the splitting of Fiber Nodes may be of high value for VoD and SDV, but may be of low value for DOCSIS HSD. If this is the case, then the MSO can use Output QAM Replication and permit the DOCSIS HSD streams for a single DOCSIS HSD Service Group to propagate to more than one RF port on the CCAP chassis, whereas the streams for a single VoD/SDV Service Group could propagate to only one RF port on the CCAP chassis. This approach could help the MSO save on the cost of enabling the additional DOCSIS HSD processing on the RF ports. Another approach (which will be described in more detail below) uses both dedicated bandwidth (DOCSIS bandwidth dedicated to a fiber node) and shared bandwidth (DOCSIS bandwidth shared between multiple fiber nodes) to solve the problems of having low bandwidth utilization levels within service groups after performing many node splits.

Due to the four reasons listed above (the reduced impact of the node splits on required bandwidth over time, the reduced impact of the node splits on intra-node split durations over time, the higher cost of the node splits over time, and the low channel utilizations that result from the node splits over time), some MSOs may decide that they will only perform node splits for DOCSIS HSD up to a certain point in the future. After that, they may find that the node splits are too expensive.

For example, if the MSO wishes to see at least a 30% reduction in the Required DOCSIS Downstream Bandwidth Capacity for any node split and that they must also see at least an 11 month period of time until the next node split is required, then the last node split that can be performed on a Service Group and accomplish these goal would be the node split that takes the Service Group size from 512 subs/SG to 256 subs/SG. In this case, node splits that reduce the size of the Service Group to values less than 256 subs/SG would be considered to be too expensive or too ineffective to be practical. This implies that many MSOs will likely stop their node splitting activities once their Service Groups fall to a certain size (on the order of 512 subs/SG, 256 subs/SG, or 128 subs/SG).

Due to the diminishing return on investment for node splits to smaller sized Service Groups, many MSOs will likely stop in that size range, placing an apparent limit on the window of time during which the HFC can provide adequate bandwidth capacity for the subscribers on the HFC plant. In other words, the ending of effective node splitting essentially defines a sunset time at which the HFC plant will cease to provide adequate bandwidth for all of the subscribers connected to it.

Many MSOs are trying to determine when this HFC sunset will occur. As can be seen in Figure 2, it is a function of many variables. How much total spectrum is supported by the HFC plant? How much of that spectrum will be dedicated to DOCSIS HSD services? What DOCSIS implementation (3.0 or 3.1) is utilized on the HFC plant? What are the signal-to-noise (SNR) levels on the HFC plant? Which modulation formats are permitted on the HFC plant?

Different MSOs will answer these questions differently. But two examples of the sunset are illustrated in Figure 2 (for the assumed growth rates defined for the figure).
The red arrows at the top of the figure locate dates at which an MSO would likely perform its last few node splits before the HFC sunset if they used DOCSIS 3.0 across only 64 QAM channels for its HSD services. These dates occur in the 2018-2021 time-frame and may permit the MSO to operate on the HFC network until ~2022. The green

arrows at the top of the figure locate dates at which an MSO would likely perform its last few node splits before the HFC sunset if they used DOCSIS 3.1 across an entire 1.2 GHz plant for its HSD services. These dates occur in the 2022-2025 time-frame and may permit the MSO to operate on the HFC network until ~2026, so four years of extra runway were effectively created by the transition from DOCSIS 3.0 to the more robust DOCSIS 3.1 modulation formats and spectral widths.

One may wonder how these dates can be extended. One obvious way is for the MSO to alter one or more of the terms within Equation (1). As an example, some MSO Marketing teams may have the option of slowing down the growth rates on Tmax values within that formula. This may or may not be possible depending on the Tmax challenges from the competition. But if it is possible, then this can have profound effects on the life-span of the HFC network. This is illustrated in Figure 3, which shows the same types of plots found in Figure 2, but the plots in Figure 3 are for a network where the MSO Marketing team has found it possible to limit the growth rate on Tmax values to a 30% growth rate beginning in 2014. The plots illustrated in Figure 3 also show the growth rate on Tmax values dropping to a 30% growth rate beginning in 2014 as well. (Note: These growth rate reductions can be seen by the flattening of the curves on the right-hand side of the figure).

As illustrated by the red arrows at the top of the figure, this change in the growth rate (from 50% to 30%) of both Tmax and Tavg has clearly extended the sunset of the 64-channel plant to ~2024 or so, and the change also makes node splits to small node sizes valuable again. The green arrows at the top of the figure are even more intriguing, as they indicate that the change in the growth rate of Tmax has greatly extended the sunset of the 1.2 GHz HFC plant deep into the 2030 timeframe. The change in growth rates has also made node splits to smaller node sizes more valuable again, and it also extends the HFC sunset deeper into the future. If MSOs are able to initiate this change in Tmax growth (and a similar change occurs in Tavg values), then a lot of value can be obtained.

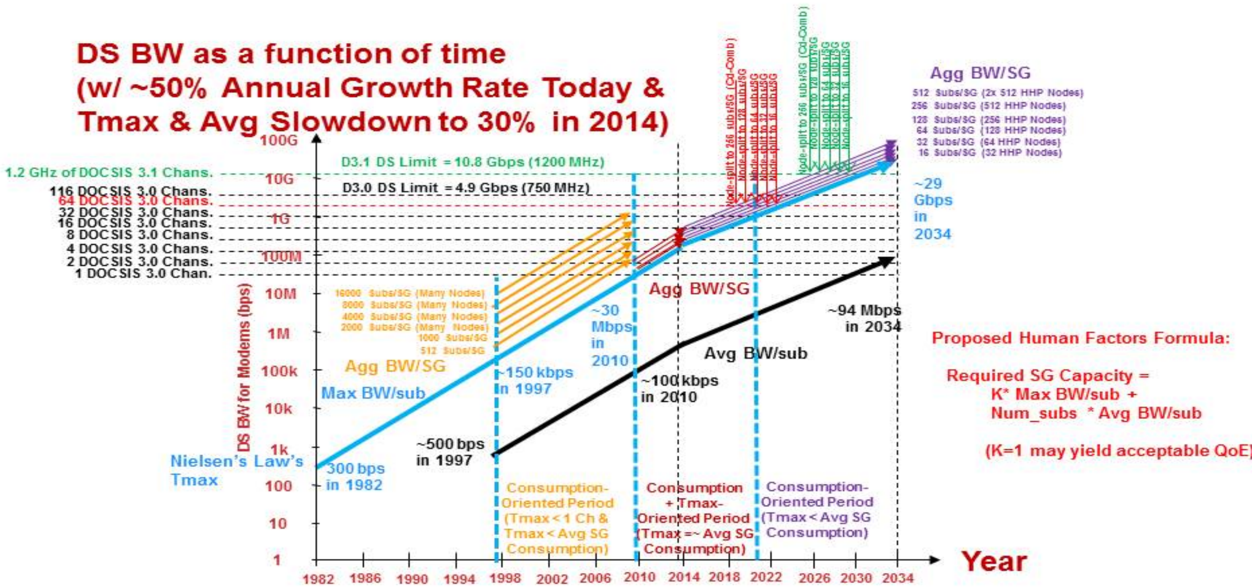

Fig. 3 DOCSIS HSD Required Downstream Bandwidth Capacity with 30% CAGR In Future

What if MSOs are forced by competitive pressures to keep raising their Tmax values at the current 50% growth rate? And what if the Tavg values continue to grow at their current 50% growth rates? If that occurs, then MSOs would be forced to live on the curves of Figure 2, and they would start to see their HFC plants running out of gas in the 2020s. What can MSOs do to extend the life-span of their HFC plant in this case?

One promising solution based on the use of dedicated and shared bandwidth will be outlined in later sub-sections. Other highly-promising solutions based on DOCSIS 3.1 and RFOG and PON and Point-to-Point Ethernet will also be outlined in later sub-sections. We will also explore several variants of another alternative technology known as Digital Access Architectures (DAAs) in later sub-sections. All of these alternative technologies, plus others, will be studied in the next section.

USE OF ALTERNATIVE TECHNOLOGIES IN OR AROUND THE HFC PLANT

Overview

There is a plethora of technologies being proposed that offer MSOs evolutionary or revolutionary changes to their existing HFC plant and spectrum to help accommodate the required bandwidth growth rates for all of their different service types in the future. In this section, we will explore several well-known approaches, and we will also define some new proposals. These approaches and proposals include:

- Increasing spectral efficiencies and spectral widths with DOCSIS 3.1
- Increasing spectral widths and Service Group capacities with Traditional Headend-based CCAP systems
- Improving SNRs with Broadband Compression Forward/ Broadband Compression Return (BCF/BCR) Distributed Access Architectures (DAAs)
- Improving SNRs and Decreasing headend power/rack space with Remote PHY Distributed Access Architectures (DAAs)
- Improving SNRs and Decreasing headend power/rack space with Remote CCAP Distributed Access Architectures (DAAs)
- Bridging to the future with RFOG FTTH
- Increasing bandwidth with Extended-Spectrum RFOG FTTH
- Improving bandwidth capacity with Dedicated DOCSIS/Shared Extended-Spectrum RFOG arrangements (based on bonding)
- Capitalizing on various blends of Switched IP Video Broadcast over DOCSIS/Nailed-Up IP Video Broadcast over DOCSIS/IP Video VoD over DOCSIS/ Broadcast Digital Video/ SDV/ VoD Digital Video/Analog
- Capitalizing on various blends of MPEG-2/H.264/HEVC compression
- Supporting various blends of SD/HD/4K/8K Video resolutions
- Increasing spectral efficiencies and spectral widths with EPOC
- Increasing bandwidth with PON FTTH
- Increasing bandwidth with Point-to-Point Ethernet FTTH

From the long length of this list, it should be apparent that the cable industry is entering an interesting period in the history and evolution of the HFC plant. With the existing HFC plant limited in bandwidth capacity, changes are required if subscriber demand for bandwidth continues in the fashion described in the previous sections. Should MSOs augment their HFC plant to accommodate the growing bandwidth demands? Should MSOs deploy new technologies to accommodate the growing bandwidth demands? If so, which technologies should be deployed? All are good questions. And all are challenging questions.

There is a large array of choices that MSOs will be making as they adapt their HFC plants over the next fifteen years. Each MSO will be making these decisions independently, and each MSO will usually have a set of constraints that are unique from all of the other MSOs. As a result, some bifurcation of the HFC market is likely to take place as we move forward into the future and different MSOs select different paths. However, the decisions for any particular MSO can only be made wisely if the MSO is equipped with accurate information on each of the available approaches in the above list. We will attempt to provide a brief and un-biased synopsis of each approach in the sub-sections below.

Improvements Provided By DOCSIS 3.1

DOCSIS 3.1 had been heavily discussed and publicized since its beginnings at CableLabs in the mid-2012 timeframe. Due to its importance for the future, the cable industry (consisting of both MSO and vendor communities) has dedicated many resources to the effort to create the new DOCSIS 3.1 specification in record time. As a result, the specification is now available and vendors are rapidly working to deliver DOCSIS 3.1-capable products within the next year or two.

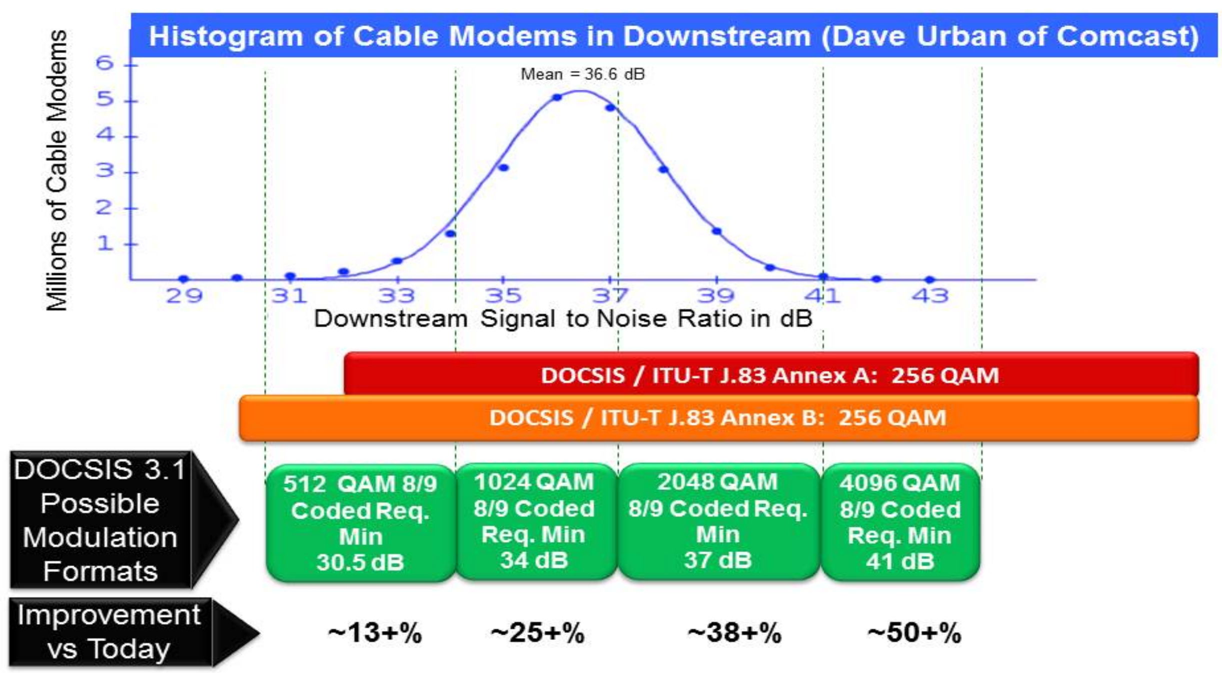

Fig. 4 Histogram of Modem Counts vs SNR and DOCSIS 3.1 Modulation Formats and Spectral Efficiency Improvements vs Today

DOCSIS 3.1 is a backwards-compatible augmentation to the DOCSIS 3.0 specification that promises better spectral efficiencies (more bps/Hz) and wider spectral widths for both the Downstream and Upstream paths. The specification provides improved spectral efficiencies via many techniques, including the use of Orthogonal Frequency Division Multiplexing (OFDM) modulation, the use of higher modulation orders (4096 QAM and higher), the use of more efficient Low-Density Parity Check (LDPC) Forward Error Correction, and the use of bit-loading to custom-fit the modulation orders to the varying SNRs across the spectrum of the HFC plant. Backwards compatibility is guaranteed by the fact that DOCSIS 3.0 and DOCSIS 3.1 channels can co-exist on the HFC spectrum. In addition, pre-DOCSIS 3.1 cable modems will work with DOCSIS 3.1 CMTSs, and pre-DOCSIS 3.1 CMTSs will work with DOCSIS 3.0 cable modems.

As a result of its power and flexibility and backwards-compatibility, many MSOs are looking to DOCSIS 3.1 to give them a boost that will extend the life of their HFC plant by (at a minimum) several years. The actual HFC plant life extension that will result from the use of DOCSIS 3.1 depends on many different factors, including the annual subscriber bandwidth growth rates, the number of node splits that are performed, the amount of investment that the MSO is willing to put into their plant to extend its spectral width, and the quality of the HFC plant (i.e. SNRs). To better quantify this last point, the bottom of Figure 4 illustrates the likely spectral efficiency gains that may be expected if an MSO transitions from DOCSIS 3.0 to DOCSIS 3.1. These gains are shown to be a function of SNR, and improved SNRs obviously lead to increased spectral efficiency gains. Overlaid on top of the chart is a histogram created by David Urban of Comcast, showing a distribution of sampled field data taken from a large number of modems in HFC plants today [URB1]. Obviously, improving the SNRs in HFC plants would lead to improved spectral efficiencies in the future.

Examples shown in Figures 2 and 3 illustrate that reasonable DOCSIS 3.1 investments (using full-spectrum DOCSIS 3.1 in a 1.2 GHz plant) could provide HFC plant life extensions of 5-6 years. Heavier investments that push the DOCSIS 3.1 spectrum to 1.7 GHz or higher could yield even longer extensions to the HFC plant life.

Improvements Provided By Traditional Headend-based CCAP Systems

As MSOs move forward, there is no doubt that node splits will be a natural part of the future HFC plant evolution going forward. Due to their importance, the benefits of node splits were discussed in detail within the first sections of this paper. In those sections, it was shown that there may be at least 3-4 rounds of node splits that many MSOs may perform in the next decade (before the value of node splits begins to diminish).

In addition, it is becoming clear that some MSOs will consider increasing the spectral widths on their Upstream spectrum and/or their Downstream spectrum.

CCAP chassis of the future will have to evolve to support these two changes. Well-designed CCAP chassis should be able to accommodate these changes with minimal hardware changes (such as the insertion of new cards into existing chassis). In addition, well-designed CCAP chassis should also be able to support the expected Network-Side Interface and backplane bandwidths that would be required for these node-splits and this spectral expansion. The authors have performed some studies of this topic in a companion paper [ULM1], and it appears that at least a subset of the existing CCAP chassis on the market today should be capable of supporting at least 3 rounds of node splits between now and 2020, while also providing DOCSIS 3.1 support for at least 1.2 GHz of spectrum on each Service Group. This capacity expansion should be achievable without requiring any significant change in chassis power or rack-space requirements. Further expansion may also be possible, but further innovation in RF connector density may be required to support that.

As a result of these facts, many MSOs will be able to rely on their existing Traditional Headend-based CCAP chassis to support their HFC network needs (DOCSIS and EQAM) for many years to come. And due to the benefits of Moore's Law, most MSOs will likely be able to utilize these Traditional headend-based CCAP resources without suffering from a need to increase their headend power and rack-space requirements. Using the plots within Figures 2 and 3, Traditional Headend-based CCAPs should be able to support MSOs (without headend power or rack-space increases) until ~2024 (with a 50% CAGR) or until ~2030 (with a 30% CAGR). If MSOs are able and willing to slightly expand the amount of power and rack-space required for CCAP chassis within their headends, then these dates can be extended even further into the future.

Improvements Provided By Distributed Access Architectures (DAAs)

As mentioned in the previous sub-section, most MSOs will likely be able to support their video and HSD services using Traditional Headend-based CCAPs. However, some MSOs may be planning to limit their HSD bandwidth capacity growth going forward (if they can), and as a result of creating that limit, they may foresee benefits to performing more node splits more rapidly than other MSOs. These MSOs may see a need to support more Service Groups than would be easily supported by a Traditional Headend-based CCAP chassis. In particular, they may run into issues related to the required power and/or rack-space for the large number of CCAPs that they may need to deploy in their headends.

There are two other types of MSOs who may also see issues in using Traditional Headend-based CCAPs. The first type consists of those MSOs who might be planning to use large amounts of Wavelength Division Multiplexing (WDM) on a small number of fibers feeding a large number of fiber nodes (which may be created by active node splitting activities of the future). The second type consists of those MSOs who might be planning to deploy (or have already deployed) long fiber runs to their fiber nodes. Both of these conditions (many wavelengths or long fiber runs) can lead to the introduction of many optical noise side-effects resulting from nonlinear effects within the core of the fiber.

There are numerous sources of these nonlinear effects within optical fibers. Some are due to changes in the refractive index of the medium with optical intensity. Others are due to an inelastic scattering phenomenon. Typical phenomena include:

- Self-Phase Modulation (SPM)
- Cross-Phase Modulation (CPM)
- Four-Wave Mixing (FWM)
- Stimulated Brillouin Scattering (SBS)
- Stimulated Raman Scattering (SRS)

In general, these nonlinear effects generate noise whose magnitude is increased with higher optical powers. As mentioned, these nonlinear effects are also increased as a result of more lambdas on the fiber and/or longer optical fiber runs. The changes that lead to this problem (more lambdas per fiber and longer fiber runs) are definitely expected in the future, and they will exacerbate the noise problems caused by nonlinear fiber effects.

Unfortunately, this implies that noise levels on the fiber will rise (and corresponding SNR levels will fall) at a time when many MSOs may want to deploy the higher QAM modulation orders that are permitted by DOCSIS 3.1. Since the higher QAM modulation orders (like 1024 QAM or 4096 QAM or even 16384 QAM) will require much higher SNR levels to operate with adequate bit error rates over the coaxial portion of the HFC plant, dealing with the aforementioned nonlinear noise issues within the fiber portion of the HFC plant may become a necessity in the future.

One technique for mitigating the effects of the nonlinear optical noise (and increasing the overall end-to-end SNR of the HFC plant) is to use a newer generation of lasers and fibers and receivers. As an example, the use of Externally Modulated Lasers (instead of Direct Modulated Lasers) will greatly help to reduce the dispersive effects of chirp, and that will generally help to reduce noise on the fiber. Studies have indicated that the correct usage of these components will permit operation to the higher modulation orders in most HFC plants.

However, there is another exciting technique for mitigating the effects of the nonlinear optical noise, which also helps to solve the problems of MSOs who have issues with the required power and rack-space within their headends. This technique employs Distributed Access Architectures (DAAs).

There are many types of DAAs being proposed for use in the future, and each proposal has its own sets of pros and cons [EMM1]. Due to lack of space, we will give only a brief description of three of them, which are shown in Figure 5.

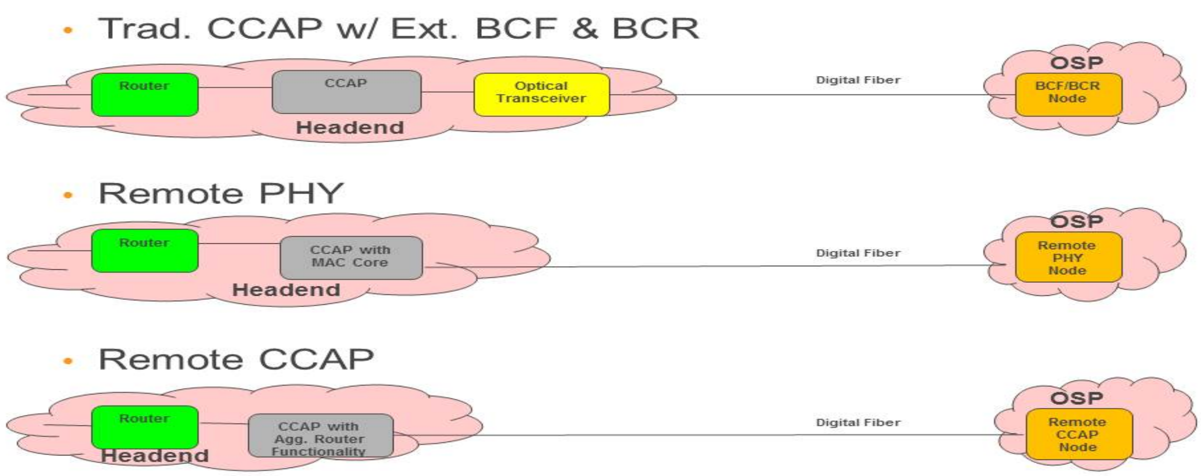

Fig. 5 Example Distributed Access Architectures

Broadband Compression Forward/Broadband Compression Return (BCF/BCR): This approach places new circuitry in the headend optical transceiver equipment and in the fiber node. It assumes a separate optical shelf receiving RF sources from analog video, Edge QAM, CMTS, CCAP, RF Out-of Band, and RF Test equipment. The Broadband Digital equipment receives multiplexed RF feeds in the headend and digitizes the spectrum before it is transported via Digital Optics to (or from) the Fiber Node. Key components of this process are the Analog-to-Digital Converters (ADCs) and the Digital-to-Analog Converters (DACs), which convert between analog and digital signal formats. The digital samples of the analog signal can be transmitted in the payload of standard Ethernet packets across the low-cost digital optic fiber system (which now carries digital Ethernet signals that are much more robust to optical noise than their amplitude modulated counterparts). This approach creates the least amount of disruption to the existing headend equipment, and it is also the only solution that allows for full transparency of the RF feeds over the outside plant. This type of solution is actually in use today for the Upstream direction- it is typically called Broadband Digital Return (BDR). Improvements in DAC and ADC technology speeds now permit this same approach to be utilized in the Downstream direction as well. In its most generic form, the Downstream approach would be called Broadband Digital Forward (BDF). However, if novel compression techniques are used to process the digital packet streams, then the resulting approach can be called Broadband Compression Forward (BCF) or Broadband Compression Return (BCR). BCF/BCR helps with the nonlinear optical noise problem, but it does not help with the headend power and rack-space problem.

Remote PHY (R-PHY): This approach separates the PHY (Upstream and Downstream) from the headend and places the full PHY layer (including the FEC, symbol generation, modulation, and DAC/ADC processing) into the Fiber Node. This requires that these functions be removed from the headend CCAPs, CMTSs, and EQAMs. The DOCSIS MAC processing remains in the MAC Core within the headend. This approach is slightly disruptive, as it requires many pieces of headend equipment (such as CCAPs, CMTSs, and EQAMs) to be modified. There are some similarities between this R-PHY approach and the Modular Headend Architecture (MHA) approach. But there are also many differences, such as the need to support Upstream MAC/PHY separation, the need to support new timing interfaces that work over Ethernet, and the need to add DOCSIS 3.1 support within DEPI and UEPI. However, this approach offers benefits as well. Remote PHY helps with the nonlinear optical noise problem, and it also helps with the headend power and rack-space problem. Another benefit of the Remote PHY approach is that it permits MSOs to continue to re-use their headend-based CCAPs as part of the solution. That represents a form of investment protection.

Remote CCAP (R-CCAP): This approach places the entire upper and lower MAC (Upstream and Downstream) and the entire PHY layer functionality (Upstream and Downstream) into the fiber node. In effect, this places all of the CMTS, Edge QAM, and CCAP functions into the Fiber Node and only requires a switch or router to remain in the headend. As a result, this approach is slightly disruptive. However, Remote CCAP helps with the nonlinear optical noise problem, and it also leads to the maximum amount of power and rack-space savings within the headend (even more than the Remote PHY approach). It is also possible that existing headend CCAPs (if appropriately modified) could be used to serve as dense Aggregation Routers (or PON OLTs) feeding the Remote CCAPs as well.

Improvements from RFOG and Extended-Spectrum RFOG

RF over Glass (RFOG) technology permits MSOs to transmit their standards RF signals (i.e. DOCSIS, MPEG-TS Video, and Analog) all the way to the subscriber homes over fiber. It requires a special ONU to be placed within each home, and the ONU is responsible for performing an optical-to-electronic conversion function (which is quite similar to the function performed by a typical fiber node). RFOG offers several benefits to MSOs. First, it permits them to begin transitioning their HFC plant into a Fiber-To-The-Home (FTTH) plant (which is likely to be the plant of the future). Second, RFOG eliminates the coaxial portion of the HFC plant, which can lead to improved SNRs and higher modulation orders. Third, RFOG can extend their DOCSIS 3.1 transmission system to spectral widths that exceed the 1.2-1.7 GHz spectral limits of typical coaxial distribution systems within the HFC plant.

The spectral extensions permitted by RFOG can be utilized in both the Downstream and Upstream directions, and it could potentially permit DOCSIS 3.1 systems operating over RFOG to provide much higher bandwidths. As an example, a shared, symmetrical 40+ Gbps DOCSIS 3.1 service to subscribers (as shown in Figure 6) is feasible with this Extended-Spectrum RFOG technology. This type of DOCSIS 3.1 bandwidth extension could greatly delay the ultimate HFC sunset deep into the future.

Interestingly, the resulting Optical OFDM DOCSIS 3.1 system could provide important mitigation against fiber dispersion, because OFDM is comprised of many narrow-spectrum subcarriers that are not likely to experience much dispersion.

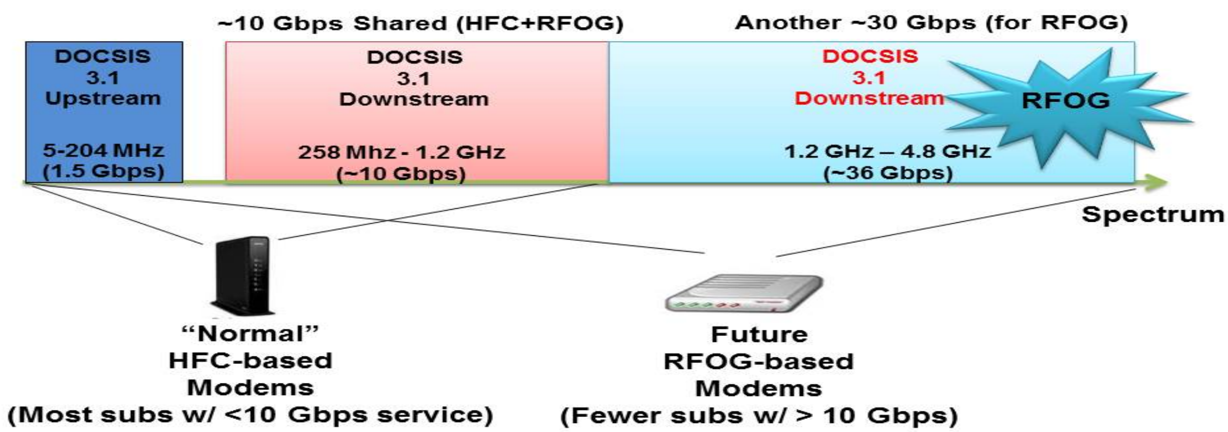

Fig. 6 Example Of An Extended-Spectrum RFOG System

An interesting attribute of Extended-Spectrum RFOG is that it can easily capitalize on the channel bonding capabilities of DOCSIS 3.1. Using this capability, the traditional spectrum below the 1.2-1.7 GHz limit can be channel-bonded with new RFOG-only spectrum that resides above the 1.2-1.7 GHz limit. This approach may provide a nice level of investment protection with the DOCSIS equipment of the past, with the new CPE devices that work above the 1.2-1.7 GHz limit being allowed to also re-use the bandwidth offered below the 1.2-1.7 GHz limit. It can thus permit MSOs to use the RFOG spectrum above the 1.2-1.7 GHz limit as a "boost feed," increasing the normal bandwidth into the subset of subscriber homes who request more Tmax bandwidth than that which is provided by the traditional spectrum below the 1.2-1.7 GHz limit.

Since only a small subset of subscribers may opt for the Extended-Spectrum RFOG technology in any given year, the use of this technology could also provide MSOs with a slow and incremental way to transform their HFC network into a FTTH network, a few subscribers at a time.

Improvements from New Blends of Video Delivery Techniques

Although High-Speed Data (HSD) is the fastest growing service within the MSO's HFC spectrum, MSO-managed video services still consume the largest percentage of the spectrum today. And those MSO-managed services produce a large amount of revenue, so they will undoubtedly remain a big portion of the spectrum for many years to come. How much spectrum will MSO-managed video consume in (say) 2025? That is an important question, because whatever spectrum is used by MSO-managed video cannot be used by the growing HSD services. (Note: The green horizontal lines in Figures 2 and 3 assume that HSD services are offered 100% of the available HFC spectrum. Obviously, this is not a likely scenario in the near or distant future). As a result, the seemingly un-stoppable growth in HSD bandwidth will eventually hit a limiting ceiling on the capacity that is available, and offering more spectrum to MSO-managed video services will force the HSD bandwidth to hit that ceiling sooner in time. In a sense, MSO-managed video and HSD will become fierce competitors for the precious spectral resources on the HFC plant as time moves forward into the future.

Thus, to accommodate the growing HSD bandwidth, MSOs will undoubtedly look to various technology paths that offer to squeeze the bandwidth of MSO-managed video into a smaller portion of the HFC spectrum. Some of these technology paths will force them to re-consider the technologies used to transport their MSO-managed video. Thus, the future will likely see different MSOs using different mixes of SD Broadcast Digital Video, HD Broadcast Digital Video, SDV, VoD Digital Video, and Analog.

Over time, Analog spectrum will be heavily reclaimed (if not entirely reclaimed). DTAs offer a good, low-cost technique for accomplishing that goal, but future Media Gateways with low-cost IP-set-top boxes may also provide similar low-cost alternatives. In the future, as SD Broadcast Digital Video becomes less popular and HD Broadcast Digital Video becomes more popular, SD Broadcast Digital Video will also begin to be turned off, yielding more reclaimed spectrum. SDV is another technique that can help to reclaim spectrum from the Broadcast Digital Video tier, whereby video streams are only transmitted over a Service Group if a subscriber is viewing that stream. This can save anywhere from 50-66% (or more) of the Broadcast Digital Video spectrum, permitting MSOs to offer up to 3 times more programs in a given amount of spectrum or permitting MSOs to reduce their spectrum requirements by 66%.

In addition to a transition away from Analog Video towards Digital Video, and in addition to a transition away from Broadcast Video and towards SDV, many MSOs are also looking to a transition away from MPEG-TS Digital Video delivery to IP Video delivery over DOCSIS. There are several reasons for this trend. First, IP Video delivery over DOCSIS will permit MSOs to capitalize on the higher spectral efficiencies of DOCSIS 3.1, which can yield a 50% improvement over the spectral efficiencies of traditional MPEG-TS delivery. Second, the wider effective channel widths created by DOCSIS channel bonding can help produce statistical multiplexing gains that can provide transport for an additional 20-30% of programs (or can provide commensurate savings in video service bandwidth requirements). Third, the use of IP Video can permit MSOs to capitalize on the bandwidth-saving benefits of Adaptive Bit-Rate (ABR) algorithms for all of their multiplexed unicast video streams. Fourth, more intelligent headend-based adaptive algorithms that take into account the nature of the video content and the nature of the video receivers can permit even more bandwidth savings [ULM2]. Fifth, IP Video over DOCSIS will permit MSOs to consolidate their headend servers and use common equipment for both their primary screen video delivery and their multi-screen video delivery. This may result in overall cost savings.

IP Video over DOCSIS may be phased in using many different tricks. The challenge is to find a way to shorten the "Simulcast window," which is the period of time during which MSO-managed video services need to be delivered in both their legacy MPEG-TS infrastructure as well as in the new DOCSIS infrastructure within a single Service Group. It may be used to transport VoD services (in lieu of the similar MPEG-TS-delivered VoD services). It may also be used initially to transport ethnic video services to a subset of the subscribers who access that ethnic offering. Over time, more and more programming will likely be delivered over IP Video, and less will likely be delivered over MPEG-TS. As more Digital Broadcast Video programming is moved onto the IP Video delivery infrastructure, MSOs will be able to utilize Switched IP Multicast over DOCSIS as well as Nailed-Up IP Video Multicast over DOCSIS.

Improvements from New Blends of Video Compression Techniques

As mentioned in the previous sub-section, MSOs will definitely be looking to various technology paths that offer to squeeze the bandwidth of MSO-managed video into a smaller portion of the HFC spectrum. Using these technologies will permit the MSOs to accommodate the rapid growth of HSD bandwidth on their HFC plant for a longer period of time. In addition to changes in the blend of Video Transmission Techniques, MSOs will also be looking very seriously at changes in the blend of Video Compression Techniques.

Many Digital Video Compression technologies are now available for MSOs to consider and utilize. These include the MPEG-2, H.264, and HEVC compression techniques. Proprietary statistical-multiplexing technologies that simultaneously compress multiple streams are also a key part of these compression algorithms.

	MPEG-2	H.264	HEVC
SD (480i30)	2.25 Mbps	1.25 Mbps	0.75 Mbps
HD (1080i30 or 720p60)	9 Mbps	5 Mbps	3 Mbps
HD (1080p60)	16 Mbps	9 Mbps	5 Mbps
UHD (4Kp60)	65 Mbps	35 Mbps	20 Mbps
UHD (8Kp60)	260 Mbps	140 Mbps	75 Mbps

Table 6: Video Bandwidth Requirements with Different Resolutions and Compression Techniques

The compression efficiencies of these various technologies continue to evolve and improve over time, but the comparative quality of each compression technique varies from technology to technology, from vendor to vendor, from month to month, and from video quality expert to video quality expert. As a result, it is always difficult to compare the efficiencies of these compression technologies using an apples-to-apples comparison that everyone agrees on. Nevertheless, one video quality expert recently gave an assessment of the video compression capabilities of these different technologies as they exist today. That assessment is captured in Table 6.

From the table, it is apparent that a move towards HEVC encoding would provide the best compression ratios. As a result, MSOs will likely want to eventually move in that direction. However, they will need to migrate in that direction from their current MPEG-2 equipment, and there are limits placed on them by the fact that the existing equipment deployed in the field may not support the improved compression technologies. As a result, this migration path may be a two-step path, moving from MPEG-2 to H.264 first, and then moving from H.264 to HEVC. (Note: The arrival of some 4K video content in the coming years will require that HEVC be instantly utilized for that content).

Challenges from New Blends of Video Resolution Techniques

While MSOs have a large array of available tools that they can utilize to help compress the amount of bandwidth required for MSO-managed video in the future, there are also some evolutions taking place that will unfortunately push the required bandwidth in the opposite direction- requiring more and more bandwidth capacity as time goes forward. The biggest threat in this area is the arrival of new and higher video resolutions for the future. SD video and HD video have been common-place on the HFC plant for many years. But the next round of video resolution improvements will produce Ultra High-Definition (UHD) video streams that require more bandwidth. 4K UHD content delivery (3840 pixels x 2160 pixels, progressive scan @ 60 frames per second, compressed using HEVC) will consume ~2 times the amount of bandwidth of today's HD content delivery (1920 pixels x 1080 pixels, interlaced scan @ 30 frames per second, compressed using MPEG-2). In addition, 8K UHD content delivery (7680 pixels x 4380 pixels, progressive scan @ 60 frames per second, compressed using HEVC) will consume even more bandwidth within the HFC spectrum.

Improvements Provided By EPOC

Ethernet PON over Coax (EPOC) is a developing standard that will provide a new technology available for MSOs to consider for use within HFC networks. In an EPOC system, Passive Optical Network (PON) technology is used to transmit signals over the digital fiber to the node. As a result, it helps mitigate against nonlinear optical noise issues. Within the fiber node, the PON signals can be converted into QAM-based OFDM signals for final transmission over the coaxial portion of the HFC plant. EPOC will likely utilize LDPC Forward Error Correction as well. As a result, the performance of EPOC is expected to be very similar (or identical) to the performance of DOCSIS 3.1 systems. Since both EPOC and DOCSIS 3.1 operate in an HFC network, the two technologies will likely compete for the interests of the MSOs in the future.

MSOs will likely make their EPOC versus DOCSIS 3.1 decision using the following logic. EPOC systems have the benefit of beginning to utilize a technology (PON) that may be a key technology for the future, but it has the disadvantage that it is not backwards-compatible with existing MPEG-TS or DOCSIS CPE equipment that is already deployed. On the other hand, DOCSIS 3.1 and CCAP offer backwards-compatibility to both existing DOCSIS CPE equipment and existing MPEG-TS CPE equipment. But DOCSIS 3.1 and CCAP also offer a future path to the PON systems of the future (since CCAP specifications incorporate PON functionality, even though no existing CCAPs actually offer it yet).

So the real question seems to boil down to whether MSOs will value a present-day ability to begin deploying equipment that they may be using in the future (for example

PON OLTs) or whether MSOs will value a present-day ability to be backwards-compatible with already-deployed CPE equipment (permitting investment protection and equipment re-use). Those who prefer the former scenario may choose EPOC. Those who prefer the latter scenario may choose DOCSIS 3.1.

Improvements Provided By PON

A Passive Optical Networks (PON) is a technology that provides a direct optical link between the headend and the subscriber home. The device in the headend is called an Optical Line Terminal (OLT), and the device in the home is called an Optical Network Unit (ONU) or an Optical Network Terminal (ONT). Many ONUs (or ONTs) can share a single FTTH optical feed from the OLT in the headend, so the bandwidth capacity provided by a PON is always shared by all of the ONUs (or ONTs) connected to the PON feed.

There are two incompatible PON technologies that have been defined in different standards committees: EPON (driven by the IEEE organization) and GPON (driven by the ITU organization). PON technologies of the future may include bandwidth capacities such as 1 Gbps, 2.5 Gbps, and 10 Gbps. Ultimately, 40+ Gbps bandwidths will also likely be provided. This is an overlay technology to the DOCSIS HFC delivery system, since it does not offer any form of backwards-compatibility to DOCSIS. PON will likely be used in Business Services and MDU environments first, but it will also find great utility in servicing Residential subscribers as well (once Residential subscriber bandwidth demands exceed those that can easily be provided by traditional DOCSIS systems).

PON may find a few competitors in the FTTH space. One FTTH competitor to PON is RFOG/Extended-Spectrum RFOG (which was described in a previous sub-section). If future Extended-Spectrum RFOG systems are created with low-cost DOCSIS 3.1 bandwidths exceeding 40 Gbps, then MSOs will definitely need to consider both approaches when doing comparisons for higher-bandwidth systems. Another FTTH competitor to PON is Point-to-Point Ethernet, which will be described in the next sub-section.

Improvements Provided By Point-to-Point Ethernet

Whereas PON is a FTTH technology that provides shared bandwidth services, Point-to-point Ethernet is a FTTH technology that provides dedicated bandwidth services. For a particular optical fiber in a Point-to-Point Ethernet system, there is one and only one subscriber connected to the fiber. The bandwidth capacities associated with Point-to-Point Ethernet will follow the Ethernet bandwidth curves. As a result, 1 Gbps, 10 Gbps, 40 Gbps, and 100 Gbps Ethernet are readily available today. 400 Gbps Ethernet is also becoming available (although it is still quite expensive). There is also work to extend Ethernet to 1 Tbps services in the future.

As a result of these higher bandwidth capacities, Point-to-Point Ethernet may become a popular access technology of the future for Business Services applications and MDU applications. It may also become popular for Residential services if/when the required Tmax bandwidth levels transmitted into each home exceed the bandwidth capabilities of PON or Extended-Spectrum RFOG. (Note: While it is difficult to predict the future, PON and Extended-Spectrum RFOG may peak out at 40 Gbps of bandwidth).

PREDICTIONS

Defining the Approach

Predicting the future is always a challenging task, and predicting the future to the year 2030 and beyond borders on the edge of fool-hardy, because new technologies and new subscriber demands can always develop without warning over a lengthy fifteen year period. As a result, those who attempt to make long-term predictions are likely to be wrong.

The authors expect that to be the case for many of the predictions that will be made within this section of the paper. Nevertheless, the authors are hopeful that there is still some value in laying out predictions about the paths that MSOs may follow into the future, because the prediction exercise requires one to examine the subscriber demands and the upcoming technologies that are likely to be available in the future that can accommodate those demands. Weaknesses and strengths of each technology can therefore be identified, and that information alone can oftentimes be useful. In addition, the creation of the predictions will hopefully stimulate good discussions that lead to better predictions in the future.

The authors will not describe a single path into the future, because it seems certain that the technologies used within HFC plants will likely be quite varied as MSOs move forward into the future. Market bifurcations will undoubtedly occur as different MSOs select different paths based on their different constraints, starting points, and biases.

Instead, we have decided to utilize the information collected in the preceding sections and attempt to describe the higher-probabilities scenarios that are likely to play out at various MSOs between now and 2030. These higher-probability scenarios will be presented in a list. The list will be divided into three sections:

1) Scenarios that will impact most MSOs (in general),
2) Scenarios that will impact MSOs who choose to extend the life-span of their current HFC plant, and

3) Scenarios that will impact MSOs who choose to switch to new technologies early as a means of supplementing the bandwidth on their current HFC plants (without incurring any costs of HFC plant upgrades).

We will utilize the terms “all,” “most,” “many,” “some,” and “few” to indicate the probability (from highest to lowest) that MSOs will follow a particular path.

Scenarios Impacting All MSOs

The higher-probability scenarios that will likely play out at most MSOs between now and 2030 include the following:

- All MSOs will begin to face growing demands for more and more bandwidth capacity for both their DOCSIS HSD services (due to standard growth) and their Video Services (due to a gradual shift towards UHD feeds), creating stresses on their HFC spectrum.
- All MSOs will be forced to transport larger and larger numbers of higher-bandwidth 4K UHD video streams within the next several years. Support for some 8K UHD video streams may also be required as we progress into the 2020 decade.
- All MSOs will be faced with an important decision to either extend the life-span of their current HFC plant to support the required bandwidth capacity (upgrading the HFC plant in an effort to delay any transition to new technologies) or to switch to new technologies early as a means of supplementing the bandwidth on their current HFC plants (without incurring any costs of HFC plant upgrades).

Scenarios Impacting MSOs Who Will Extend the Life-Span of Their Current HFC Plant

The higher-probability scenarios that will likely play out at most MSOs who will choose to extend the life-span of their current HFC plant include the following:

- Most of these MSOs will continue to take advantage of Traditional Headend-based CCAPs deep into the future.
- Many of these MSOs will converge their video EQAM functionality into their CCAPs in an effort to reduce power and rack-space requirements within the headend. Some of these MSOs will wait until actual power and rack-space issues develop (likely in the 2020 timeframe) before they begin this transition. Some of these MSOs will plan ahead and begin the convergence earlier.
- Most of these MSOs will perform at least 2-3 rounds of node splits in the next 6-10 years. This “business-as-usual” activity will help them accommodate Narrowcast bandwidth growth over the next decade. As they perform these

node splits, the MSOs will capitalize on the increased Service Group counts that will undoubtedly be provided by Traditional Headend-based CCAP boxes of the future.

- Most of these MSOs will take advantage of DOCSIS 3.1 as a means of increasing the available bandwidth capacity of their existing HFC plants.
- Many of these MSOs will likely come to realize that moving any video streams (whatsoever) from their current MPEG-TS delivery/MPEG-2 compression into IP Video over DOCSIS delivery/HEVC compression (with its improved DOCSIS 3.1 spectral efficiencies and improved large-channel stat-mux gains and improved ABR-based compression techniques) will be essential to permit UHD Video feeds to be transmitted over the HFC plant in the future. Any transfer of video content from MPEG-TS to DOCSIS will prove beneficial.
- Many of these MSOs will likely move quickly from MPEG-2 compression to H.264 compression, and then will move from H.264 compression to HEVC compression. Some will skip the intermediate step of H.264 compression and move directly from MPEG-2 compression to HEVC compression.
- Many of these MSOs will begin to slowly move towards IP Video over DOCSIS transport for their video in the next few years. But as those MSOs begin to move video streams to IP Video over DOCSIS deliver systems, they will eventually realize that they should expedite the process to get through the “simulcast window” more quickly, permitting them to eventually reclaim most of the bandwidth associated with MPEG-TS video. They will therefore expedite this process and attempt to retire MPEG-TS video transport as quickly as possible.
- Many of these MSOs will start using Switched IP Multicast for Linear IP Video transmissions to capitalize on the bandwidth efficiencies that can result from the use of Switched IP Multicast.
- Many of these MSOs will increase their Downstream spectrum to 1.2 GHz in an attempt to provide more HFC plant bandwidth capacity.
- Some of these MSOs will increase their Downstream spectrum to 1.7 GHz in an attempt to provide even more HFC plant bandwidth capacity.
- Many of these MSOs will change the split on their Upstream spectrum to be 85 MHz in an attempt to provide more Upstream bandwidth capacity.
- Some of these MSOs will change the split on their Upstream spectrum to be 204 MHz in an attempt to provide more Upstream bandwidth capacity. This may permit Tmax = 1 Gbps Upstream service offerings.
- Some of these MSOs may utilize BCF/BCR to improve SNRs and permit them to use the highest possible modulation orders within DOCSIS 3.1 systems. This may be more likely for MSOs who may be planning to multiplex a lot of wavelengths on a fiber or who may be deploying long fiber runs.
- Some of these MSOs may utilize Remote PHY or Remote CCAP or Remote PON architectures to improve SNRs (increasing modulation orders) and to reduce power and rack-space requirements in the headend. This may be more likely for

MSOs who may be planning to multiplex a lot of wavelengths on a fiber or who may be deploying long fiber runs. This may also be more likely for MSOs who may be performing rapid node splits. This may also be more likely in the 2020 decade, when MSOs may be out-pacing the abilities of CCAP boxes to support further node splits.

- Some of these MSOs may utilize Remote PON (a variant of Remote CCAP) to improve SNRs (increasing modulation orders) and to reduce power and rack-space requirements in the headend.
- Some of these MSOs may utilize RFOG as a technique to roll out FTTH technologies to homes without requiring them to change our existing headend and CPE equipment. This may be more likely for MSOs who are rolling out Greenfield deployments. This may also be more likely for MSOs who are planning to transition their HFC plants into FTTH plants in preparation for the introduction of new technologies (such as PON or Point-to-Point Ethernet) in the future.
- Some of these MSOs may utilize Extended-Spectrum RFOG as a means of increasing the bandwidth capacity within a Service Group as Tmax values rise in the future. This type of bandwidth (in excess of 10 Gbps) and this type of Extended-Spectrum RFOG technology may not be required until the 2020 decade.

Scenarios Impacting MSOs Who Will Switch to New Technologies

The higher-probability scenarios that will likely play out at most MSOs who choose to switch to new technologies early as a means of supplementing the bandwidth on their current HFC plants (without incurring any costs of HFC plant upgrades) include the following:

- Some of these MSOs may begin to use EPOC as an interim technology in preparation for an ultimate transition to PON.
- Some of these MSOs may begin to use PON as a technology to deliver higher bandwidths to their subscribers. They may transition to PON directly from DOCSIS, or they may transition to PON after using EPOC for a while.
- Some of these MSOs may begin to use DOCSIS 3.1-based RFOG as a technology to deliver higher bandwidths to their subscribers.
- Some of these MSOs may begin to use DOCSIS 3.1-based Extended-Spectrum RFOG as a technology to deliver even higher bandwidths to their subscribers.
- Some of these MSOs may begin to use Point-to-point Ethernet as a technology to deliver even higher bandwidths to their subscribers. The use of Point-to-point Ethernet becomes more probable as we move closer to the 2030 time-frame, because the dedicated bandwidths associated with Point-to-point Ethernet may not be required until the 2030 decade.

CONCLUSIONS

In general, the next fifteen years promise to be an extremely interesting time, and it is quite apparent that many changes will be taking place throughout the cable industry as MSOs prepare for the 2030 decade and beyond.

Fortunately, there are plenty of technology options available to MSOs that can help them maneuver their way through the transitions. Different MSOs will select different paths, but all MSOs will undoubtedly be focusing on increasing their bandwidth capacities to deliver high Quality of Experience levels to their subscriber base for all of the applications of the future (whatever they may be).

MEET OUR EXPERT: TOM CLOONAN

Tom Cloonan, has been the Chief Technology Officer or Chief Strategy Officer of ARRIS for more than a decade, responsible for directing the architectural work and future product planning for network products within the company. Prior to that, Tom was the CTO and co-founder of the start-up company Cadant, where he led the architectural work for the development of the C4 CMTS. In a previous life, Tom also worked at Lucent Bell Laboratories on various telecommunications and routing platforms. Tom has a BSEE degree, and MSEE degree, and a Ph.D. in Physics. His research has resulted in over 40 patents, over 80 published papers, one book and several book chapters.

REFERENCES

(1) [URB1] D. Urban, Comcast study

(2) [ULM1] J. Ulm and T. Cloonan, "Scaling Traditional CCAP To Meet The Capacity Needs Of The Next Decade," The NCTA Cable Show Spring Technical Forum 2014

(3) [EMM1] M. Emmendorfer, T. Cloonan, and J. Ulm, "A Side-by-side Comparison of Centralized vs. Distributed Access Architectures," The NCTA Cable Show Spring Technical Forum 2014

(4) [ULM2] J. Ulm, A. Eshet, and N. Radian, "Unmanaged ABR: How To Control Those Unruly Teenagers," SCTE Cable-Tec Expo 2013