FTTH Evolution of HFC Plants

Timothy J. Brophy, Ph.D, Business Development Manager
Steven Condra, Director, Architecture Technology
Martin Mattingly, Business Development Manager
Ron Hranac, Technical Leader
Leonard Ray, Marketing Manager
Cisco
tibrophy@cisco.com
condras@cisco.com
mattinm@cisco.com
rhranacj@cisco.com
leray@cisco.com
770 236 3101
www.cisco.com

Table of Contents

Abstract

The acceleration of data in the form of video, voice, and internet web page delivery has continued to drive the growth of the FTTH community. The belief that fiber as a medium over which all data consumption is both practical and achievable has been justified in multiple ways. The more pedestrian aspects of how that delivery infrastructure will be achieved, and the cost vs. capacity metrics are ultimately the most important ones to be answered to achieve the goals of ever-increasing adoption rates for this FTTH effort.

Broadband delivery of the so-called four-play services of voice, video, data, and wireless has been offered from traditional telephony service providers and the HFC community's multi-system operators (MSO) of the cable TV industry. For this latter group, the amount of fiber in their networks has steadily increased and the demarcation point between the fiber and the coax has moved ever-closer to the consumers' homes. Nevertheless, the drive to take fiber all the way to the home is hindered by multiple factors. When plant extensions or new (green-field) builds are required, fiber is the obvious choice for longevity and capacity considerations; however, deploying that fiber all the way to the residences and the ability to light up that fiber as the sole delivery medium is often constrained by existing infrastructure and operational considerations. Some recent initiatives in the broadband delivery marketplace have increased the opportunities for an end-to-end fiber to the home solution. Continuation of growth through commercial services deployments, and expectations of multiple data and video streams using adaptive bit rate for various devices from big screen televisions to smart phones are pushing the traditional MSO providers to evaluate and embrace FTTH technologies.

In this paper, we will examine the motivations and the practical considerations driving FTTH in the MSO stronghold of HFC infrastructure. The evolutionary options that allow the existing networks to gracefully and cost-effectively expand to a complete fiber solution will be described, and the capacities and costs of those evolutions detailed. The recent surge in efforts to include PON architectures and the potential to expand those from the commercial services arena to residential neighborhoods will be outlined. Tradeoffs detailing service and failure group size, effective split ratios and the options will be evaluated. Anticipated bottlenecks to the adoption of all-fiber networks—such as the advent of all IP networks, black-wire customers, and home network gateways—will be examined for their comparative role in the eventual deployment of fiber.

Ultimately, the current state of the deployment and evolution of HFC plants to become FTTH solutions will be explored, and the enablers and obstacles to that deployment identified. The market for driving fiber deeper in these networks is increasing and ready to begin the next stage of this deployment.

Executive Summary

Hybrid-fiber-coax networks have leveraged cost-effective shared media for the delivery of video content to a large number of subscribers for many years with particular effectiveness in the residential market segment. Today, under pressure to build and evolve those networks to support the rapid growth rates of video and data, operators are considering means to support capacities required by consumer and market trends but also promised by a wholesale change to FTTH infrastructure and topologies. The continued leveraging of the existing HFC plant is shown to be an effective and low-cost means to support even the most optimistic conversion to all-IP networks supporting advanced services, IPTV, and the aggressively growing data services within the residential community. Within the 10-year plan modeled herein, available technologies enable the continued data capacity expansion, support the transition to an IP-centric network, and define an extensible platform that is supported by the existing HFC network. A wholesale change to an FTTH architecture build out will occur when capacity demands finally outstrip the HFC plant's capabilities, but that will be sometime beyond the year 2020.

Introduction

The Hybrid-fiber-coax (HFC) architectures that are the mainstay of cable television networks today largely evolved as a low cost way to share what was originally a community antenna built to bring in weak terrestrial signals from the over-the-air broadcast TV stations. Those (often remotely located) antennas were connected to subscribers in the distribution area through lines of large diameter coaxial cable with regularly spaced trunk amplifiers to boost the signal. In the distribution portion of the network, bridger amplifiers tapped signals from the main trunk line and branched off feeder lines to individual streets. Down those feeder branches were line extender amplifiers and taps that fed drop cables (the in-home coax that is common today) to serve individual subscriber's homes. With the introduction of fiber optic transmitters and receivers in the late 1980s, the trunk and sometimes feeder portions of the plant gave way to fiber topologies. The optical receiver was housed in what is called the optical node, the optical to electrical (O-E) conversion point whose input was the optical fiber and whose (multiple) outputs were the radio frequency (RF) coax distribution plant so prevalent in the last 25 years. This fiber-to-the-node for transmission, coax-to-the-home for distribution is the HFC plant that still exists today.

The signals being sent through the HFC plant were the broadcast video and audio signals of the television channels being broadcast. The original over-the-air broadcasts were supplemented with local origination content and later by digital channels from satellites which were received, downconverted, transcoded, and remodulated to be compatible with the existing plant capabilities. As technologies continued to evolve through the next two decades, most transmission and distribution infrastructure kept pace by adding capabilities to handle ever higher RF signal frequencies and introducing bi-directional capability for data and voice. As the number of channels available and the number of subscribers of data services increased, the demand pull for more bandwidth fueled the evolution of fiber-deep networks, in which the O-E conversion point moved closer to subscribers' homes. This build out of fiber-rich infrastructure has not been uniform across the cable space, and today there are many instances of distribution plants with several RF amplifiers in cascade to boost signals in the feeder and drop networks; however, the evolutionary trend has been to push fiber nodes deep enough into networks so that few or no amplifiers are needed. These node and no amplifier (N + 0) topologies leave only the passive tap and drop network as coax distribution media.

Obviously, the move to fiber in the transport network greatly increased the capacity by use of multiple wavelengths and increased reliability by the elimination of many RF amplifiers which were subject to electrical failures of many kinds (powering, lightning strikes, etc.). So why have the HFC proponents failed to capitalize on the obvious advantages of an all-fiber network? In essence, their operational optimization focus has been dual purposed: to reach the highest capacity network possible ultimately (ala FTTH) but also to be optimized incrementally at each step of the evolution toward that goal. And while economic and technology improvements have continued to drive FTTH deployments, the HFC network has benefited from additional capacity through more efficient signal processing, protocol conversions, and transport efficiencies as well, and they have helped to fund the incremental optimization approach. Among these HFC improvements, conversion from analog to digital modulation schemes for coding efficiency gains, RF spanning for improved burst rate capacity for individual users, multicast streaming for improved efficiency of transport, and adaptive bit rate (ABR) technologies greatly enhanced network capacity in the existing topologies and were coincident with the continued explosion in demand by consumers and market competition. Consider that, as a specific example, the use of ABR has enabled a common user experience on multiple screens from the largest high-def TV monitors to small hand-held mobile devices without a significant increase in the total capacity requirements of the network. The sum of these techniques combine to improve the efficiency of each wavelength in a system. This thereby frees capacity for the content which is not in high demand, (i.e., that demanded by a few subscribers only—the so-called long-tail content). Ultimately, the improvement in efficiencies in broadcast and multicast content is such that the number of subscribers in a given service group is constrained by the amount of the long-tail programming or data services required.

To begin to lay out the rationale employed by many cable operators, we will look at some cost factors and some expected capacity requirements over time and relate those to the choices these service providers are faced with today. Should they prepare for future capacity requirements by adopting an infrastructure that can support the ultimate in bi-directional bandwidth from the outset, or continue to pursue incremental investment and capacity expansion of recent years? A casual observer might assume that continued node segmentation with its bandwidth doubling strategy and fiber deeper drive will continue over time until a small enough service group size is achieved. That, however, would miss potential advantages of jumping directly to the small, ultimate service group size (and thereby saving incremental split costs). Is the higher initial cost of that split justified? And how, if it is not pursued, can the less disruptive incremental approach compare in capacity and prepare that service provider for futures unattainable with present day network topologies and architectures? In the next section, some baseline costs of architectures are considered.

<u>Cost, Timing, and Capacity</u>

As shown in Figure 1, the costs for various HFC topologies compare favorably with those for some selected PONs. Here, we have considered a greenfield construction as a comparison, since the build out of a fiber deep upgrade to an existing plant would follow those costs. The various topologies of HFC architectures follow the nomenclature which lists the node (N) and the number of RF amplifiers at a given power level that follow the node in the coax portion of the plant. Hence, a N + 6, 52 out is a node (the end of the fiber run and the O-E conversion point) followed by a cascade of six RF amplifiers each with an output of 52 dBmV. The baseline construction and ducting cost of $561 per household passed includes the aerial strand portion of the fiber and coax, the ducts for underground construction (a varying percent depending on the modeled area) the pedestals and cabinets for housing the amplifiers, etc., and the installation of these components.

The DPON (DOCSIS PON) columns represent an FTTH version of an HFC plant that has no coaxial portion. This is an architecture compatible with the RFoG (RF over Glass) recommendations published last year under CableLabs' auspices. The optical network terminal (ONT) in such an architecture is in each subscriber's premises, and is fiber fed, with a combination of coax and Cat5 cable inside the home. A somewhat different way to parse the data is to look at the incremental cost to take fiber deeper in an existing HFC plant. When computed on this basis, the cost to take fiber to the curb (which here is essentially to the last active device) is about $200 per home passed from "average" cascades of three RF amplifiers. To extend that fiber all the way to the home, with the larger quantity of fiber cable, ducting (even if direct buried cable) and additional installation expenses costs ~$600 per home passed. This level of expenditure clearly requires additional time and/or higher average revenue per unit (ARPU) to yield an acceptable return on investment (ROI) for the capital and operational outlays. We note that the N + 1 architecture yields the lowest cost for all the architectures considered. We will revisit this point when evolution beyond the 10 year window is considered.

The cost analysis considered and shown incrementally above the gigabit passive optical network (GPON) column is not to be construed as applying solely to a GPON (or Ethernet passive optical network (EPON)) implementation. Indeed, the transition to an IP network in any scenario will require IP set top capability; however, set top boxes (STB) are already deployed in HFC plants by and large, and as the existing base of STBs turns over, newer models which incorporate IP or RF capability (a 'hybrid' STB) can be deployed for little, if any, incremental cost (and is typically borne by the subscriber when rented). It is also useful to observe that in a currently deployed HFC plant with no unused RF spectrum available to add channels, the duplication of content in both the RF and IP formats will be required. This will compare to a FTTH case, in which ALL homes need to have an IP STB for video delivery to legacy television, and so is virtually a complete rebuild of the customer premise equipment (CPE) end of the network. Furthermore, although it is not necessarily true for new entrants into a given market, an entrenched cable operator will already have operations support systems/business support systems (OSS/BSS) deployed and DOCSIS provisioning and management gear in place. If the cable operator were to adopt a PON and architecture, these systems would need to be replaced or supplemented, adding significant cost in purchase and training. An interesting recent development called DOCSIS Provisioning of EPON (DPoE) is a means to minimize these latter costs by using Ethernet transport combined with a DOCSIS control and management plane to provide Ethernet services initially targeted to commercial customers. This will be discussed briefly later in the paper.

Next we come to the capacity and timing issues that face an operator who wants to ensure that his network has adequate bandwidth to address subscriber needs as that network evolves to higher data rates over the next decade. Of course, these service providers are continuing to supply legacy services within their existing networks to current subscribers, and these services often include broadcast analog and digital tiers, some narrowcast digital channels (i.e., digital channels that only go to a subset of the service group), bi-directional data services, voice, some on-demand content (the largest of which is video on demand), and other lower capacity services. Of the available options for cable operators to provide additional bandwidth to serve their growing needs, address legacy customers, continue to leverage their previously-installed infrastructure, and begin to address the movement to a future network (irrespective of whether the medium is fiber or not) that employs all-IP transport, the use of IP video is preferred. This primarily results from two reasons: video is and will continue to be the application that is the most demanding in the network, and the efficiencies of IP video which can be leveraged to free up existing bandwidth in the plant. Just how efficient is IP video? To get a sense of the order of magnitude, refer to Figure 2 which shows a distribution of the content of today's networks in a 1 GHz spectrum, and a vision of how that content will be delivered in a network employing IP video in 2020.

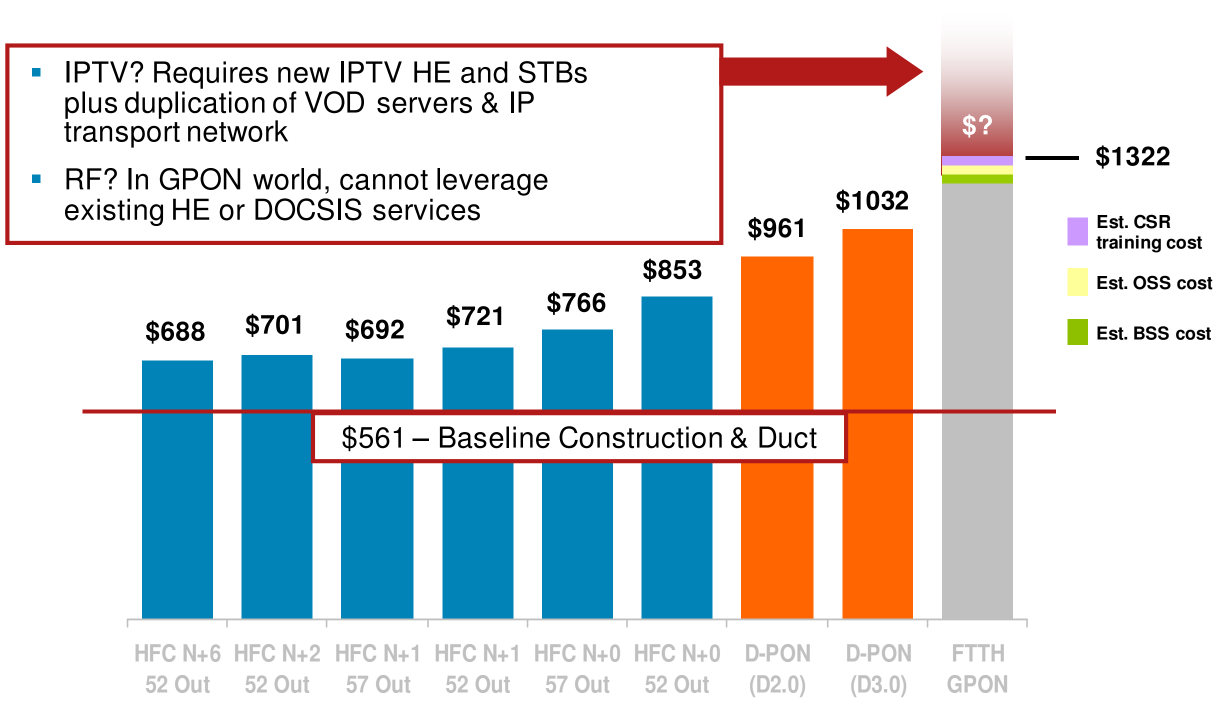

Figure 1. Greenfield installation cost comparisons for HFC and PON plants. (Baseline includes Aerial Strand, Ducts, HFC peds, and installation)

How did we achieve this remarkable improvement in bandwidth efficiency to yield such an abundance of spectrum for future growth? To fill in some details, we begin by defining the parameters upon which this network model was based, and the scenario for managing downstream bandwidth, legacy video, IP video, and bi-directional high-speed data (HSD). The baseline for today's network is a 400 homes passed node, with IP video growing from a modest 10% penetration today to a 50% penetration in 2020. Accompanying this IP video growth is HSD growth from an average of 32% today to 50% in 2020. These penetration rates (number of homes passed which subscribe to a given service) have the additional multiplier effect of a compound annual growth rate (CAGR) of 40% per year. These HSD growth and penetration rates are used for both forward (downstream) and reverse (upstream) data, although the starting points are different. All of these figures are consistent with what is being reported as deployed and realistic for North American cable operators. The basis for the HSD growth is computed as the total data capacity in a node serving area divided by the number of subscribers. This leads to 240kbps/sub in the forward and 100kbps/sub in the reverse, resulting from the differences in modulation (256 QAM vs. 16-to 64-QAM) and by limitations in the available upstream bandwidth (37 MHz in North America). The content provided to a home is accessed by a cable-ready big screen TV or a set top box or a personal video recorder which may take standard or high definition signals, cable modems, and the like. In today's network, typical households have an average of 2.4 tuners in use. This number is expected to grow, and aggressively (to ensure we do not underestimate bandwidth needs) we assume five streams will be required per household in 2020.

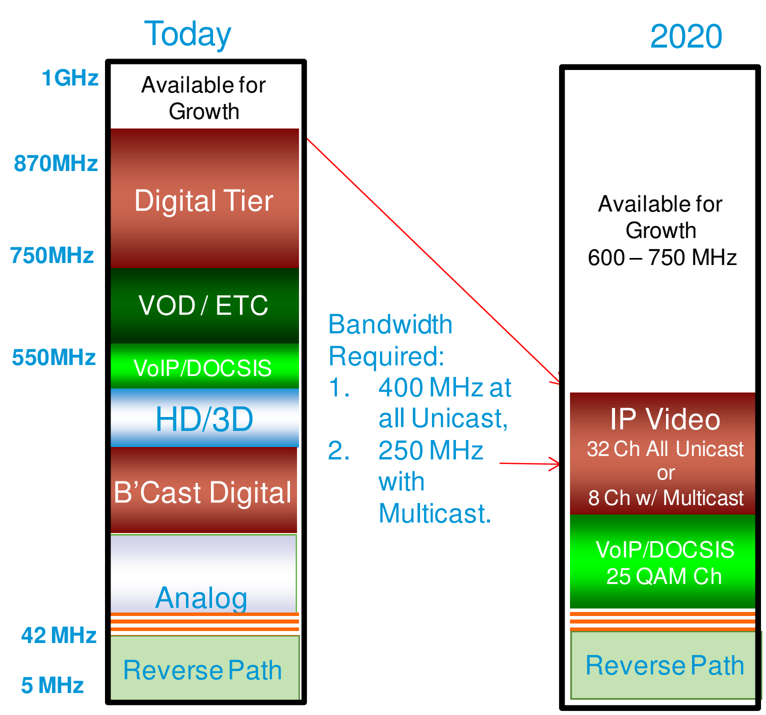

Figure 2. Application content in a 1 GHz plant today and in 2020.

One approach to bandwidth management being pursued by several cable operators includes analog reclamation, in which an analog TV channel is retired, and the freed 6 MHz of spectrum is replaced by content consistent with planning and local demand. That 6 MHz spectrum may be filled with 10-12 standard definition (SD) digital (256 QAM) channels, two to three MPEG-2 high-definition digital channels, video-on-demand (VOD) content, etc. The

MPEG-2 codecs in use today are being replaced by MPEG-4 (H.264) units, allowing at least four HD channels per 6 MHz slot. Furthermore, reductions in broadcast content means that only those video streams actually being viewed (or recorded) by a subscriber are being transmitted in the network.

In addition to the MPEG coding just mentioned, three additional techniques can dramatically reduce the bit rate required for transmission of video. The use of Variable Bit Rate (VBR) for video has advantages over Constant Bit Rate (CBR) transmission in that the content required by any individual video stream can vary in time to that required by concurrent (different) streams, and that even for an individual stream can vary in time depending on the relative motion from frame to frame. Live action sporting events require more frequent whole-screen changes relative to that required by a desk-bound news anchor, and so higher bit rates to maintain a given perceived picture definition. When encoded and sent on the same channel, the complementary minima and maxima for the individual streams can yield VBR gains as high as 50%; in practice, the statistical gains can be conservatively estimated at 30%.

Adaptive Bit Rate gains are different in that they are realized by considering the maximum resolution which can usefully be displayed on various screens to which the streams are directed. So a stream bound for a 50-inch high-definition television or an equivalent PVR requires approximately three times the bit rate as that of a pad device. Therefore, total bit rates required to individual subscribers can be reduced by identifying the device the video is targeting, and by extension to the distribution of devices requesting video in the service group as a whole.

Finally, the most significant reduction in required bandwidth results from the use of multicast streaming. When multiple subscribers request the same video, the initial subscriber’s stream is unicast, that is, directed only to that viewer. When the second and subsequent viewers request the same video, they are able to access exactly the same content without having to recreate the stream which would have increased the bandwidth consumed. These multicast gains are additive to those of VBR and ABR, and aggregate bit rates of just 25% of the already reduced VBR total are readily achievable provided the number of subscribers is statistically significant. A statistically significant size means that there are sufficient numbers of viewers (or requested streams) within a service group so that it is likely that multiple viewers request a (popular) stream. The implication here is that a strategy which continues to reduce the service group size by continually splitting the nodes or by jumping to a PON-friendly 32 household size initially may be foregoing the most useful and significant efficiency gain in bandwidth utilization. Also, in terms of the multicast gain, note that similar arguments can be made for the reduced bit rate streams used by secondary and tertiary screens within a home, but the fact that the bit rates are reduced lessens the gain of the multicast streaming.

Capacity Requirements Through 2020

Taking the baseline parameters and the bandwidth efficiency techniques to the level of modeling the capacity requirements for the ensuing decade involves projecting the application demand over that time frame and associating each application with a growth curve. For example, the HSD traffic is anticipated to grow at a rate of 40% per year in the forward and reverse. Alternatively, the demand for VOD traffic grows over the near term (~3 years) and then falls as the available IP video offerings and their network availability increase, thereby replacing the on-demand content with the managed services version of what is currently over the top video streams. Cisco has developed models for each individual application, and run many "what-if" scenarios to determine which realistically match with analyst application predictions, which match with customer inputs with regards to growth, and which are most useful from an equipment capacity perspective for bracketing the expected trends over time. The individual applications are summed together taking rates for screen types, percent of multicast (vs. unicast) content, and mix of applications into account to arrive at the curves shown below in Figure 3.

The three curves of Figure 3 show the bit rate required to support IP TV deployments consistent with the modeling assumptions stated earlier (400 HP/SG, 10-50% penetration, IP replacing VOD to some extent). The upper (blue) curve is consistent with recently published data, assuming a constant bit rate of 6 Mbps per stream. As is evident, the growth from 360 Mbps to nearly 2 Gbps is unsupportable if there is no spectrum in which to put the channels or no plan for dealing with these additional capacity requirements. The inclusion of variable bit rate alone (green curve) reduces the requirement by more than 30% to ~1.2 Mbps, even in an all unicast environment. Here though, because the service group includes 400 households, additional benefits will accrue from the use of multicasting. With field data we used to create and validate the model, this multicast VBR (red curve) results in a reduction to approximately 300 Mbps, a factor of 4x improvement from Unicast VBR and in excess of 6x compared to Unicast traffic using CBR.

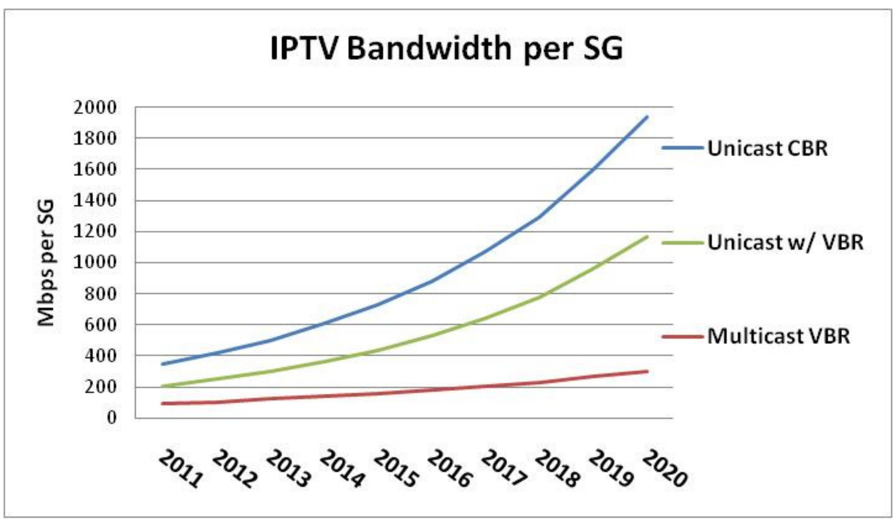

Figure 3. Required bit rate per service group to support IP video over time.

In Figure 3, we have considered ONLY the IP video requirements expected over the next decade. This is primarily forward path (downstream) traffic. The bi-directional nature of high speed data demand obviously will add to that; here the co-propagating downstream traffic is considered. Later we will look at the upstream traffic to consider overall network limitations. As noted earlier, the HSD demands for this 400 HP service group averages out to approximately 240 kbps per subscriber, presently has about a third of all households passed in a serving area as customers, and is growing at a compound annual rate of 40 percent. In Figure 4 we show the requirements for the service group with only data and showing CAGR values of 40, 50, and 60 percent. Some of these data may in fact be over the top video, in which case this is best effort unicast traffic.

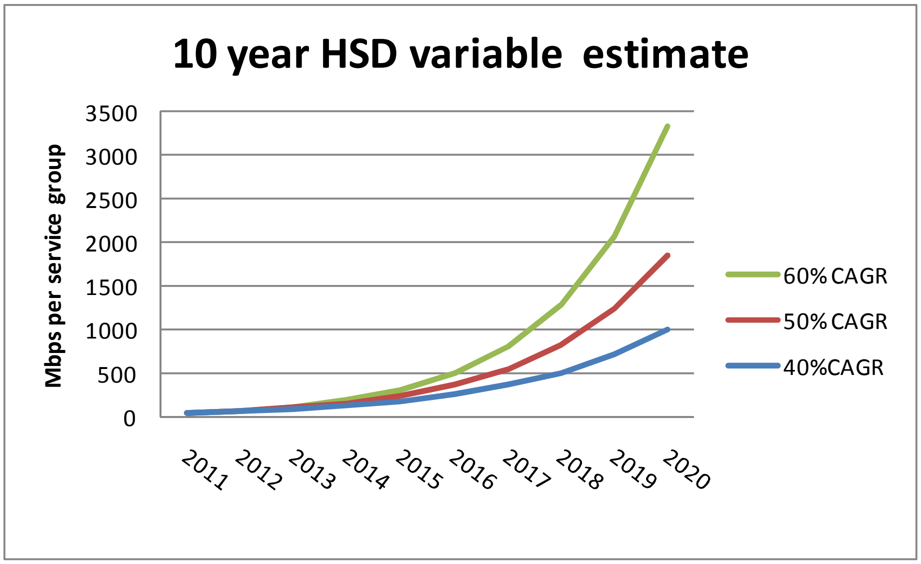

Figure 4. Required BW per service group for HSD only at 40, 50, and 60% CAGR.

In Figure 5, the aggregate bandwidth for both IP video and the downstream portion of the HSD load is plotted. Here, the data are shown as the equivalent number of 6 MHz QAM channels required. These are modulated at 256-QAM, common for downstream traffic in an HFC plant, and support approximately 40 Mbps per 6 MHz channel. At greatest efficiency, the 35 QAM channels use just 210 MHz of downstream bandwidth. This is well within the capability current HFC plants, showing its advantages and robustness for future growth.

Upstream High-Speed Data Requirements

For the upstream or reverse direction, the baseline model assumptions of 32 to 50 percent penetration over the 10 year period of the model hold, while the starting data rate is 100 kbps per subscriber. The 40% CAGR here leads to approximately 2.1 Mbps per subscriber or a cumulative 420 Mbps per service group. In contrast to the 256 QAM used in the downstream, an upstream 6.4 MHz, 64-QAM channel yields about 27 Mbps and thus a cumulative requirement of 16 channels of 64-QAM in the year 2020. This is shown in Figure 6.

The red line in Figure 6 shows the currently available 6.4 MHz, 64-QAM channels that may be grouped together and used as a single super channel (DOCSIS 3.0 only) for higher single user burst capability or may be used as individual channels separately addressed by different users (DOCSIS 1.0 through DOCSIS 3.0). Note that in 2016 the required number of channels exceeds the available number and this therefore becomes a bottleneck in a system with a 5-42 MHz return spectrum (common in North America). By contrast, a 5-85 MHz return spectrum, for which equipment is already available and which will support

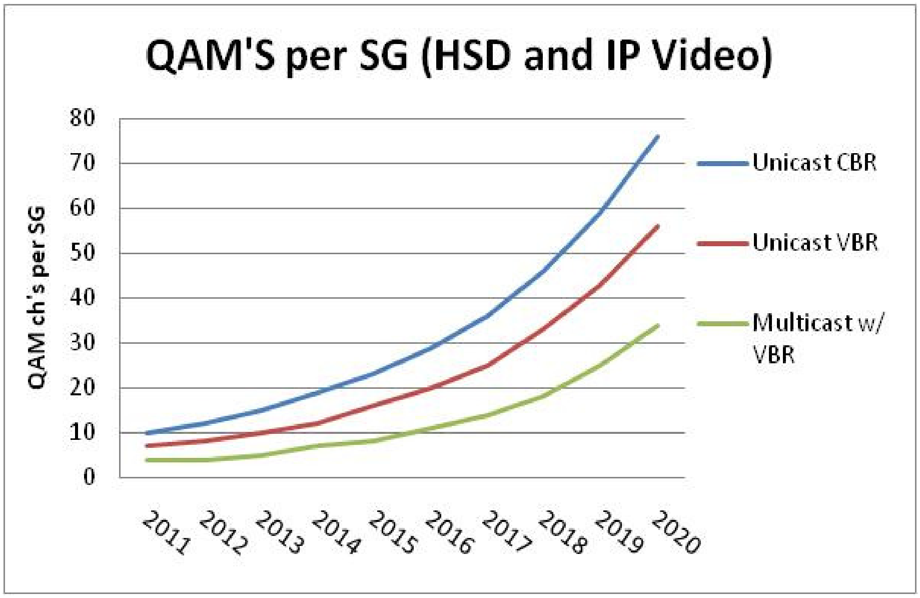

Figure 5. Aggregate IP Video + HSD requirements in 256 QAM channels

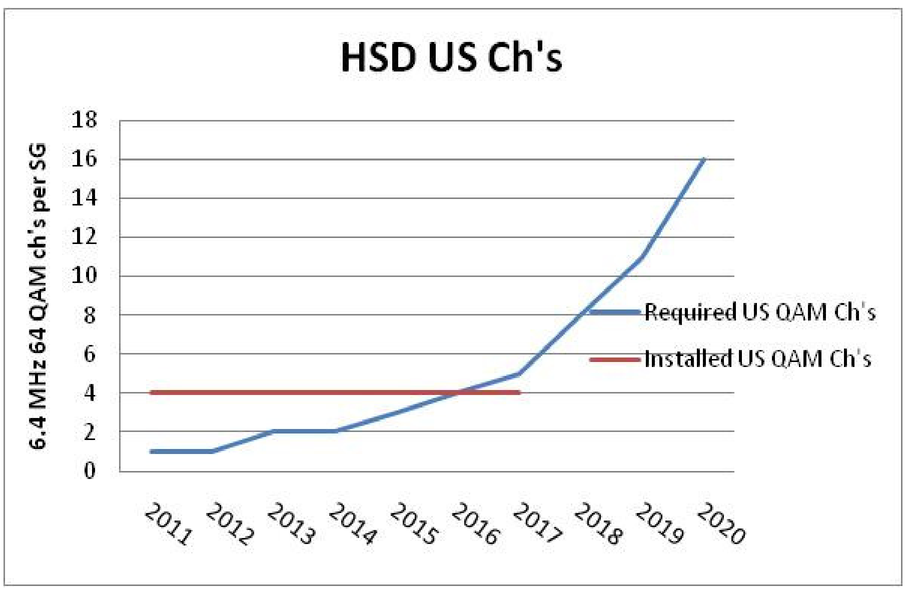

Figure 6. Upstream 64-QAM channels required (blue) and available (red) with a 42 MHz return

approximately 12 channels (each using 6.4 MHz of RF spectrum), does not become a limit or bottleneck until about 2019. There may be additional untapped capacity in this broader spectrum solution because the frequencies from 54 to 85 MHz can be modulated with the more complex and more capacious 256-QAM (~ 38 Mbps), and is not limited to the 64-QAM payload rate of 27 Mbps. In any case, when the bottleneck is reached, cable operators must take action to address the capacity limit, and the two choices implied previously are for a node split (this can be done independently for the reverse path and the forward path) or by adding additional spectrum (e.g., move to a 5-85 MHz or potentially a 5-200 MHz reverse spectrum).

Finally though, it is clear that continued demand growth beyond the modeled period here will result in bottlenecks and capacity limitations that will exceed even the flexibility of the low-cost HFC plant to address. Fortunately, an HFC topology allows for as smooth and painless a transition to a FTTH network as exists, and without the upfront capital costs of expanding network capacity far in advance of customer demand. Figure 7 shows what would be the penultimate network segmentation of an HFC plant, in which a serving area has been segmented such that each node (recall: the O/E conversion point in the HFC plant) serves a maximum of 100 to 125 subscribers (Figure 7a). This point is reached when either forward or reverse demand becomes a bottleneck in the 400-500 home serving area, and may be realized either with two incremental node splits or a single quartering depending on the particulars of economics existing at the time. In Figure 7b, one of the nodes has had its E/O electronics replaced by passive fiber 32-way splitters to accommodate the fibers to be run to the subscribers' homes.

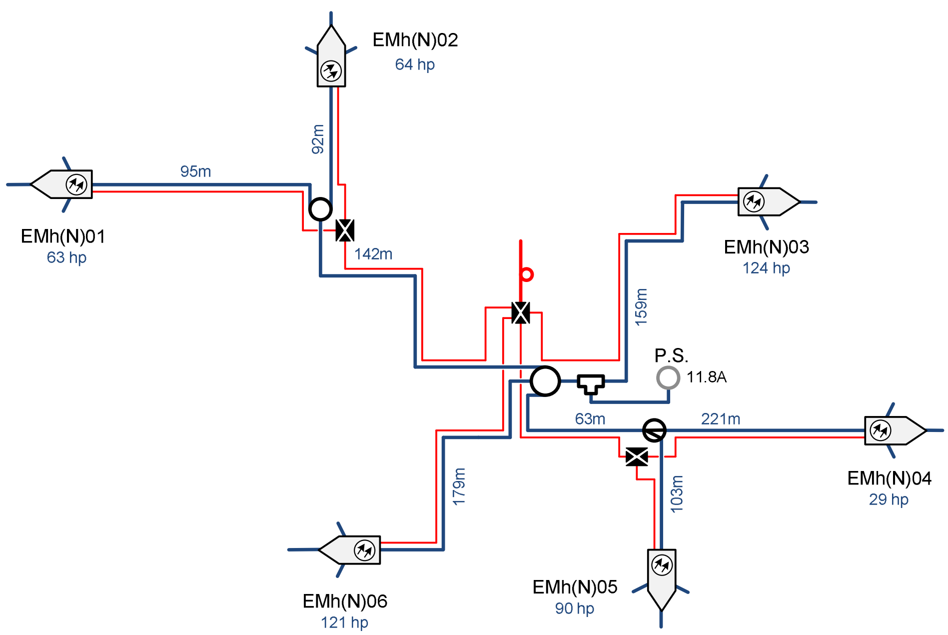

Figure 7a. An example of node serving areas that are at the HFC practical extent.

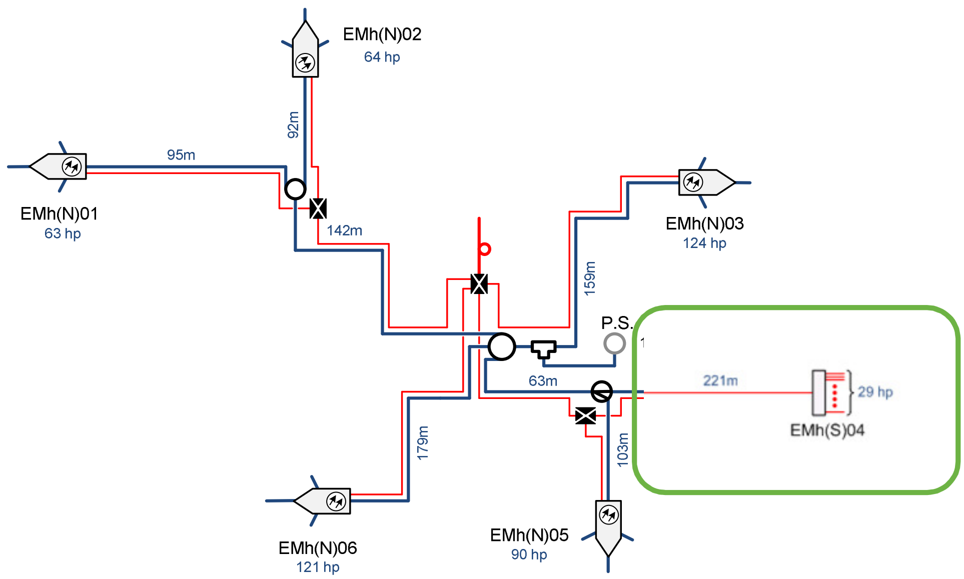

Figure 7b. One of the nodes of Figure 7a replaced with a splitter with fiber outputs for a FTTH rollout.

The node location(s) are straightforward to transition into an optical distribution frame (ODF) location and the same housings can be often be employed in the capacity of the ODF or simple splice housings if desired. With the final step in the process of replacing the coaxial distribution portion of the plant with fiber, a complete FTTH topology can be realized. As is usual in residential distribution plants, the majority of the sheath miles of fiber is in this portion of the network and a large percentage of it will be installed underground in the residential communities. By postponing the build out of this most extensive and expensive portion of the plant until either the demand for bandwidth exceeds that of the HFC network to provide it or until the coaxial cable reaches the end of its useful life, significant capital costs are avoided, the investment in the existing HFC plant is fully collateralized and exploited, and the network capacity is both adequate and readily matched to revenue-generating demand. The impressive bandwidth and lower plant operational cost of a passive network are not foregone, nor are they bought and paid for without a revenue stream to justify bearing the installation cost. To be sure, HFC plants are not monolithic, and variation in topologies exist. This is matched by variation in customer demand as well, and distribution of high revenue generating units is not uniform throughout a plant. The evolution and revolution of the network connectivity of HFC allows matching capacity to demand either when concentrated in a new neighborhood, or distributed among miles of existing plant serving well-established customers. Along the way, capital funds are spent where they are most needed and where they are most likely to realize a return on that investment.

Although it is beyond the scope of this paper, some current efforts to merge FTTH topologies with existing HFC networks are underway. The DOCSIS Provisioning of EPON effort seeks to use an Ethernet transport and distribution to serve commercial customers with either dedicated or shared fiber resources to fill their data and video demands. It somewhat complements the usage patterns of the residential networks it is deployed into, and physically branches off from that network as deeply as possible to most cost-effectively use the existing fiber infrastructure. Both the existing residential plant and any new commercial services accounts (e.g., data, voice, cell tower back haul) use the same, existing back office systems for billing, operations, and administrative controls and provisioning. This is likely to lead to ever more familiarity and adoption of "new" protocols as these services become more common.

Service providers using a traditional HFC plant embrace the brave new world of FTTH and the bandwidth and operational simplicity that provides. The path to realizing that vision can be rapid or slow, but will be cost conscious by necessity. Leveraging existing, working assets for their useful lifetimes or until they become a burden on the overall operational, administrative, maintenance, and provisioning costs of a network is a prudent path matching assets to demand, and expenditures to revenue.

Summary and Conclusions

HFC infrastructure and topologies were developed initially for one-way distribution of television signals carrying analog modulated RF channels in an optical environment. The shared nature of the media and the comparatively low cost of the equipment led to its widespread use. As customer demands for additional services increased, cable operators responded with increases in RF spectrum delivered to carry more channels downstream, and improved plant performance and protocols to permit high capacity and high quality bi-directional data networks. The ubiquity of video as content on web pages as well as from managed sources has continued to challenge service providers to match the capacity of their networks to rising demand and future expectations. Their huge investment in existing plant infrastructure, if it is to continue to yield returns, will need to be able to address this continuing and growing demand.

We have outlined here the costs involved in driving fiber deeper into the network and contrasted it to more traditional FTTH build costs which are similar to those of a greenfield build out. Next, we described techniques that are being and can continue to be effectively exploited to improve the efficiency of the existing RF spectrum for the delivery of video and data as the demand side pull becomes unavoidable. We have modeled the growth and capacity requirements over a 10 year period and matched that to techniques accessible to service providers using an HFC infrastructure, and shown not only adequacy, but continuing excess capacity to add new services as they become available. We looked at the reverse path and its unique bandwidth and capacity challenges and suggested timing and methods for when and how to deal with those limitations, and suggested an open-ended segue into a fully-functional FTTH network as capacity demands. We pointed out the strengths and flexibility that are the hallmarks of the existing HFC plant, and implied that the lower cost of leveraging the existing plant is a cost savings which does not impair an operator from a capacity and growth perspective both over the modeled period and beyond.

To reiterate some of the main points, everything comes down to cost, timing, and capacity. What capacity does an operator require, when is that capacity required, and what will it cost to provide that capacity? Node segmentation is a useful, cost-justified, and recommended step in leveraging the existing HFC plant and transitioning the network to provide higher capacity. This step is far less operationally disruptive and provides a bandwidth increase commensurate with the segmentation of the service group. An initial jump to a fiber to the last active approach only provides marginal increases in bandwidth per user at a comparatively high operational cost.

Currently available marketplace solutions offer an array of ideas spanning available technologies, topologies, capacities, and cost. Experienced architectural solution and equipment providers have a catalog of headend and distribution gear assembled in such a way as to meet growing capacity needs, be economical to deploy and operate, and provide the flexible groundwork to adapt to future evolution. Key among the advantages are improvements in

- capital and operational expense
- initial and ultimate service group size
- service disruption during installation
- average and burst data capacity
- power consumption

While segmentation is an effective beginning, a cycle of incremental network subdivision is a costly approach to arriving at an FTTH network. Maintaining a sufficiently large service group can share in both increased multicast gain and, for individual subscribers, a share of the increased aggregate bandwidth. The accelerating video traffic demands from network storage to multiple screens makes a multiple-upgrade scenario an unsustainable option for continued revenue growth.

Therefore, from a traffic balancing and content equalization perspective, continuing to segment fiber nodes is a sub-optimal strategy on the basis of cost, timing, and capacity. From an expense perspective, leveraging the existing infrastructure, signal delivery gear, and existing topologies is always less expensive than replacement, especially when there is (initially) no additional content for the new network to deliver. So long as the existing coaxial portion of the distribution plant has sufficient bandwidth and speed to meet the customer demand, it is unlikely that cable operators will remove perfectly good, serviceable drops with a new one (no matter how good) that a customer will not see or cannot acknowledge as superior. Holistic, end-to-end solutions bridge physical infrastructure from existing HFC to all-IP networks, and embrace the graceful evolution that preserves investment.