Hi Ho, Hi Ho, to a Gigabit We Go
Positioning the HFC Network for the New Gigabit Era

Phil Miguelez
Comcast

Abstract

The broadband landscape is rapidly changing driven by market competition from non-traditional providers, a resurgence in new home construction, and a vast array of new services to support the connected home environment. Traditional node splitting is not enough to keep up with bandwidth demands and high speed data rate comparisons to predatory fiber to the premise over-builders. The decisions that cable operators must make to meet these challenges include options that affect every part of the infrastructure network: facilities, OSP, fiber augmentation, and plant powering. Capacity planning which barely existed just a few years ago is now the major focus of every MSO.

This paper will compare and contrast the different network architecture options that are available today and developing network technology that will be ready to deploy in the very near future – DOCSIS 3.1 with RF bandwidth expansion, Fiber Deep cascade reduction, FTTH, and R-Phy / R-MacPhy distributed architectures. The paper will also explore the relative cost implications of each architecture design and the capability to migrate to the next generation as data capacity demands continue to grow and accelerate.

INTRODUCTION

Background

MSO HFC access network architectures have changed very little for over a decade. The arrival of DOCSIS 1.0 in 1997 and the subsequent enhancements leading to the release of DOCSIS 3.0 in 2006 helped to spur the last major network upgrades. During that time QAM video channel expansion dominated the growth of the downstream bandwidth while the still emerging HSD, IP telephony, and VOD services drove expectations of further bandwidth expansion requirements in the future.

GaAs technology allowed the development of 1 GHz nodes and amplifiers that fulfilled the potential needs of D3.0 access architectures. GaN technology followed providing the increased output levels necessary to convert 750 MHz or even older 550 MHz systems into 1 GHz networks with little or no expensive respacing of existing amplifier locations. HFC appeared to have all the tools necessary to compete with Direct TV satellite delivery of video content and the newly minted Regional Bell Operating Companies (RBOC's) that began construction of massive and expensive fiber systems to replace the twisted pair copper wireline networks that had been in place for nearly 100 years.

A negative consequence of the last rebuild cycle was that Wall Street increasingly took a dim view of further network upgrades without a clear perspective on the competitive or cost benefit. The economic recession that started in 2008 also dampened any incentives to expand the capabilities of the existing network.

Narrowcast and HSD Growth

Downstream broadcast channel growth has continued to expand fueled by broad adoption of High Definition (HD) video along with an array of sports, special interest, and ethnic diversity programming. At the same time the number of narrowcast channels has also steadily grown although at a slower pace. DOCSIS 3.0 provides for up to 32 bonded DS channels representing just over 1 Gb of digital capacity. Until recently almost all systems operated with only 8 to 12 NC channels. These channels carried voice, VOD, and high speed internet data sufficient to meet the tier rates guaranteed by the MSO based on the number of subscribers in the node service area.

High speed data usage has also been steadily increasing since the early days of audio modems. In 1998 Jakob Nielsen created a predictive law in which he stated that high end home users' bandwidth would grow by 50% per year. Subsequent data has validated what is now widely known as Nielsen's Law.(1) Figure 1 shows a plot of the projected data rates extended out for the next 15 years. The compound annual growth rate (CAGR) predicted by Nielsen has driven the steady BW expansion of the HFC network along with node segmentation and the gradual reductions in HFC amplifier cascades, resulting in lower numbers of homes passed per node. Every MSO has created their own set of CAGR charts reflecting factors specific to their network, customer base, or personal beliefs. These growth rate predictions will continue to drive investment for many years to come.

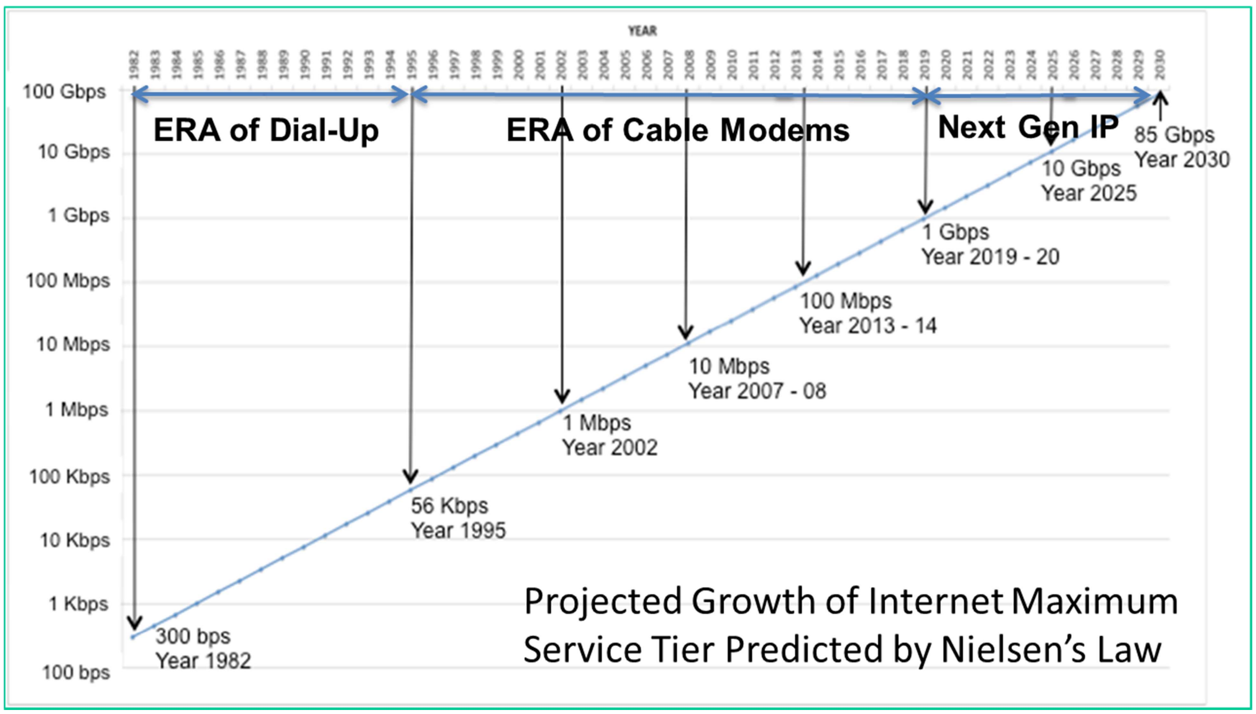

Figure 1 – Nielson's Law Predicts 50% Year Over Year Growth Rate

Downstream channel growth has historically been the largest driver of bandwidth in MSO networks. In order to create additional DS bandwidth without making changes to the outside plant operators have reclaimed analog channels, incorporated higher digital compression rates in QAM video channels, and installed Switched Digital Video (SDV) systems. SDV takes advantage of the fact that a number of channels have limited peak viewing time so these channels can be dynamically reclaimed when no one is actively watching.

Increasing narrowcast bandwidth per subscriber requires an additional set of network tools. Data consumption has continued to rise in line with predicted rates.

The primary driver of this consumption is the increasing popularity of social media and streaming video. Netflix, You Tube, and other popular video streaming services have had a major impact on the amount of data consumed by individual subscribers. Netflix alone accounts for roughly 35% of all internet downstream bandwidth.[2] Figure 2 shows the most recent data available on wireline data consumption.

	Upstream		Downstream		Aggregate	
Rank	Application	Share	Application	Share	Application	Share
1	BitTorrent	28.56%	Netflix	37.05%	Netflix	34.70%
2	Netflix	6.78%	YouTube	17.85%	YouTube	16.88%
3	HTTP	5.93%	HTTP	6.06%	HTTP	6.05%
4	Google Cloud	5.30%	Amazon Video	3.11%	BitTorrent	4.35%
5	YouTube	5.21%	iTunes	2.79%	Amazon Video	2.94%
6	SSL - OTHER	5.10%	BitTorrent	2.67%	iTunes	2.62%
7	iCloud	3.08%	Hulu	2.58%	Facebook	2.51%
8	FaceTime	2.55%	Facebook	2.53%	Hulu	2.48%
9	Facebook	2.25%	MPEG - OTHER	2.30%	MPEG	2.16%
10	Dropbox	1.18%	SSL - OTHER	1.73%	SSL - OTHER	1.99%
		65.95%		78.69%		76.68%

sandvine

Figure 2 – Q4 2015 Top Ten Peak Period Applications – North American, Fixed Access

The number of internet connected devices per home has also been a driving force in the amount of narrowcast bandwidth capacity needed per home.

THE NEW MARKET DYNAMIC

The Impact of Google Fiber

In 2010 Google announced a new business initiative – Google Fiber. Google's announced plan was to create a lottery of sorts to identify cities and towns that would be willing to provide access for new fiber deployment and a sufficient number of potential customers that would pre-subscribe for Google Fiber installation. The winning locations would be built out with FTTH PON and have HSD tier rates up to 1 Gb which far exceeded anything available from cable or other wireline broadband service providers at that time.

Google Fiber's actual competitive impact on the ground has been muted. After 5 years, the total number of video connected homes was recently reported to be only 53,000.[3] The number of internet data only homes is assumed to be higher and has been estimated at around 100,000 which is still quite low considering the number of years since the initial deployment. However, Google has had a major impact on the broadband market by raising the bar on subscriber expectations for digital delivery of high speed data rates and an even greater impact on the potential providers of broadband service delivery.

Google's size and name recognition allowed them to negotiate concessions with numerous municipalities in exchange for the investment that Google Fiber would provide to their community. The concessions included right of way easements and access to utility pole attachments which have always been costly barriers to new entrant providers. A consequence of the negotiations by Google was that they created a template that smaller FTTH broadband service provider start-ups could use to get equivalent deals when they entered a new market to compete with an incumbent provider. Within a few years FTTH organizations were springing up touting plans to "Gigafy America" and teaching startup companies how to sell and construct FTTH systems.

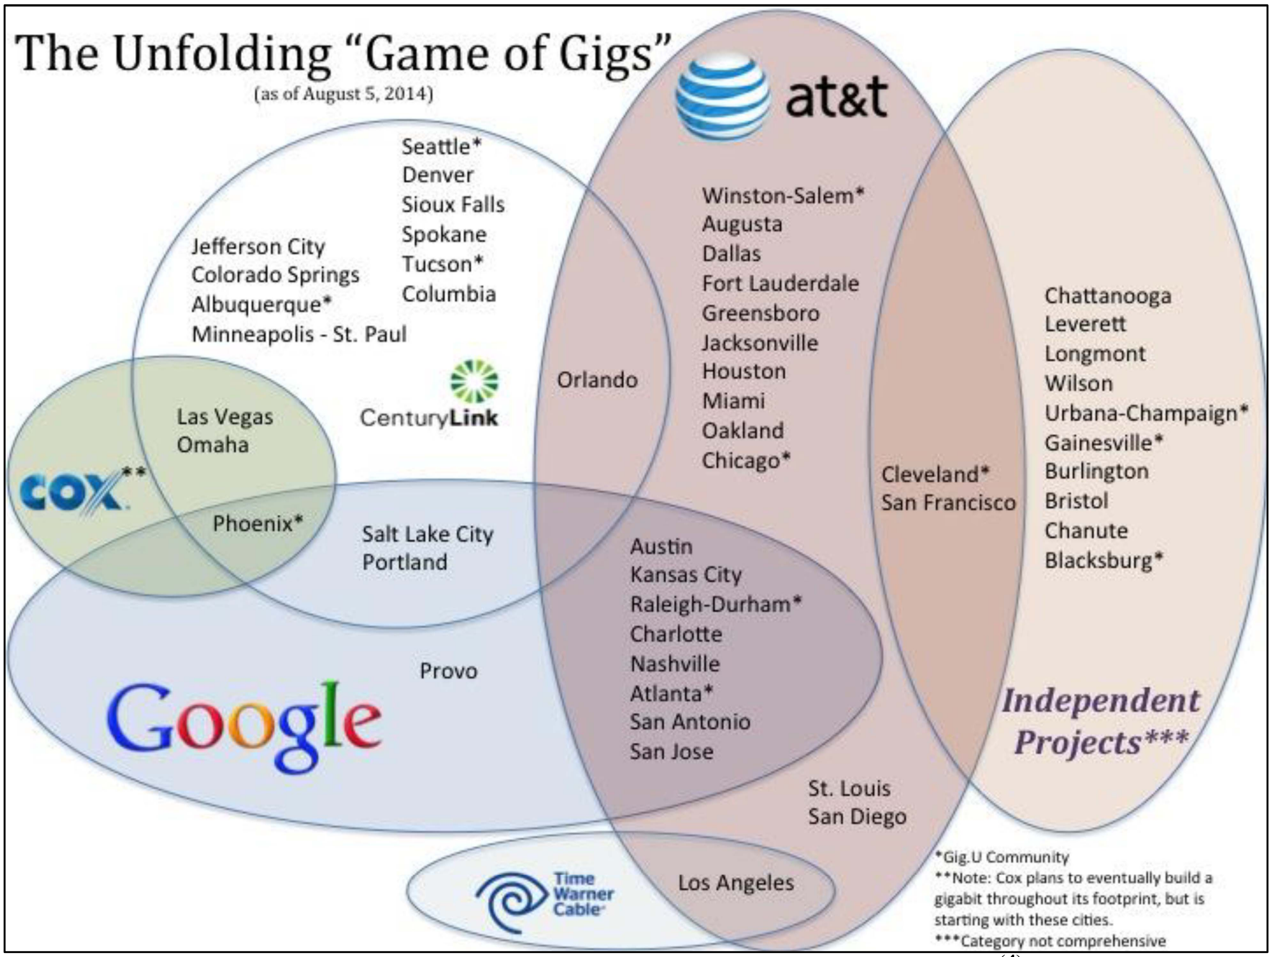

Figure 3 – Target City Deployments Announced by Major Over-Builders[4]

Figure 3 shows a graphical representation of the cities targeted by over-builders for Gigabit FTTH deployments as of August 2014. Many of the listed cities are Comcast, Time Warner Cable, or Cox Communications franchise locations. Additional cities have been announced in 2015.

More recently, AT&T has modified their direction announcing "Gigapower"[5] a fiber delivered 1 Gb internet data service. AT&T is now expanding their reach into targeted cities and residential areas wherever they have fiber assets. Verizon FIOS has taken a slightly different tact with clever ads such as the "Half Fast" campaign which is intended to point out that the cable upstream data capacity is much lower than the downstream data capacity. Verizon has followed up this marketing campaign by introducing symmetric (DS/US) data plans putting pressure on cable operators serving the same markets to expand their HSD tiers and add US capacity. Century Link, Hotwire, and countless smaller regional over builders are cherry picking MDU and gated community properties that provide 100%

penetration opportunities with long term contract guarantees. The market landscape is now filled with large and small competitors offering gigabit speed service and threatening to overtake the HFC cable space.

While it's true that the only immediate application for gigabit data rates today is the speed test, history has shown that new applications seem to always appear once the network BW is available to support them. One example is the number of smart devices in the typical home. A few years ago the network assumption was 3 to 5 devices on average and 5 to 7 devices in a high end user home. Today those numbers have jumped to an average of 5 to 7 connected devices per home and 10 to 12 for a high end subscriber.

Another pending driver is the Internet of Things (IoT). The number of internet connected devices is exploding. Smart home applications and monitoring devices especially video cameras are driving both US and DS bandwidth usage. Figure 4 shows a graphical representation of the various smart home devices available today.

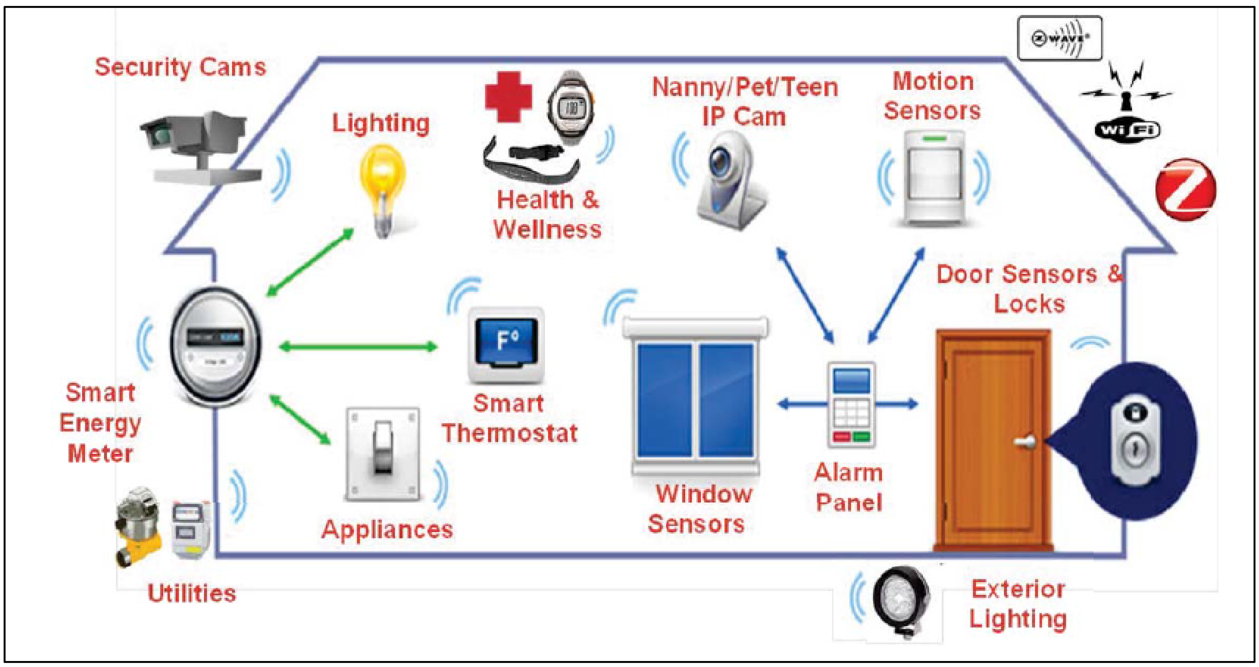

Figure 4 – Examples of the New Connected Devices That Make Up the Internet of Things

MEETING THE GIGABIT CHALLENGE

Realities of Existing HFC Networks

The DOCSIS 3.0 specification bounds the DS modulation format to 256 QAM in 6 MHz channels. With 32 bonded Single Carrier (SC) channels the usable data throughput is 1.216 Gb. Finding open channels or reclaiming existing channels in a typical 750 MHz system is a daunting task. Transitioning from analog video to all QAM is almost a mandatory first step. Replacing MP2 with MP4 modulation can also provide additional BW relief. Even these steps may not be enough considering that any increase in DS bandwidth will precipitate the need for added US capacity enhancement. This means added US channels must be allocated plus an appropriate guard band to adequately isolate the forward and return signal paths.

In higher density HP areas and especially urban serving areas, the number of targeted channels needed to satisfy multiple ethnic and special interest communities makes finding open channel space virtually impossible without an added RF bandwidth expansion. The same situation exists in low density and rural HP serving areas where many networks have been stretched to reduce the number of actives that must be maintained. This usually results in lower performance margins and reduced total bandwidth due to increased roll off at the upper band edge.

The digital capacity in an HFC system represents shared bits for all subscribers connected to a particular node. The typical homes passed node size has been steadily decreasing over the years. Physical node splits and node segmentation have reduced a node's serving area from >1500 HHP to the current average of 500 HHP. Even with this significant reduction, a few simultaneous shared users could potentially drag down the usable delivered data rates during peak time periods. Using historical CMTS concurrency rate assumptions the maximum guaranteed downstream data rate tier that could be supported with 32 bonded D3.0 channels would be 250 to 300 Mb at the current average node HP size.

The traditional method to increase narrowcast digital BW available per subscriber is to reduce the number of HHP per node. Node segmentation provides operators with a minimally disruptive method to significantly increase the delivered bandwidth per subscriber. The cost of the initial primary node segmentation is typically estimated at $20,000 since the majority of the expenses are usually limited to material costs rather than new fiber deployments.[6] By comparison, subsequent node split costs can increase almost exponentially due to fiber construction expenses, Headend expansion costs to support the additional service groups, and when calculated based on the fixed number of subscribers served by a particular node.

The high costs of node splits when evaluated as a means to compete with accelerating high speed data requirements led Comcast and other cable operators to examine alternative architectures.

D3.1 Availability in the HFC Network

In June 2012 CableLabs initiated a new specification effort to establish requirements for a major revision of DOCSIS. The DOCSIS 3.1 specification which was released in October 2013 defines a new modulation standard for HFC networks with a raw data capacity of 5 Gbps downstream (DS) and up to 1 Gbps upstream while maintaining the current 1 GHz RF bandwidth capabilities of existing D3.0 cable plant. The D3.1 specification provides an intermediate downstream RF bandwidth target of 1218 MHz which aligns with the deployment of near term achievable hardware that is already coming to market. The potential for 10 Gbps downstream capacity is achievable with an RF spectrum expansion to the full spec 1.794 GHz D3.1 DS BW allocation or eliminating the QAM broadcast channels. Figure 5 shows the subscriber data rates available for the various DOCSIS releases over the years.

Upstream capacity growth is provided by DOCSIS 3.1 with the implementation of RF bandwidth expansion to either a mid split of 85 MHz or a high split to 204 MHz. The high split option is required to achieve the 1 Gb raw capacity capability. More details on the D3.1 specification can be found on the CableLabs website.[7]

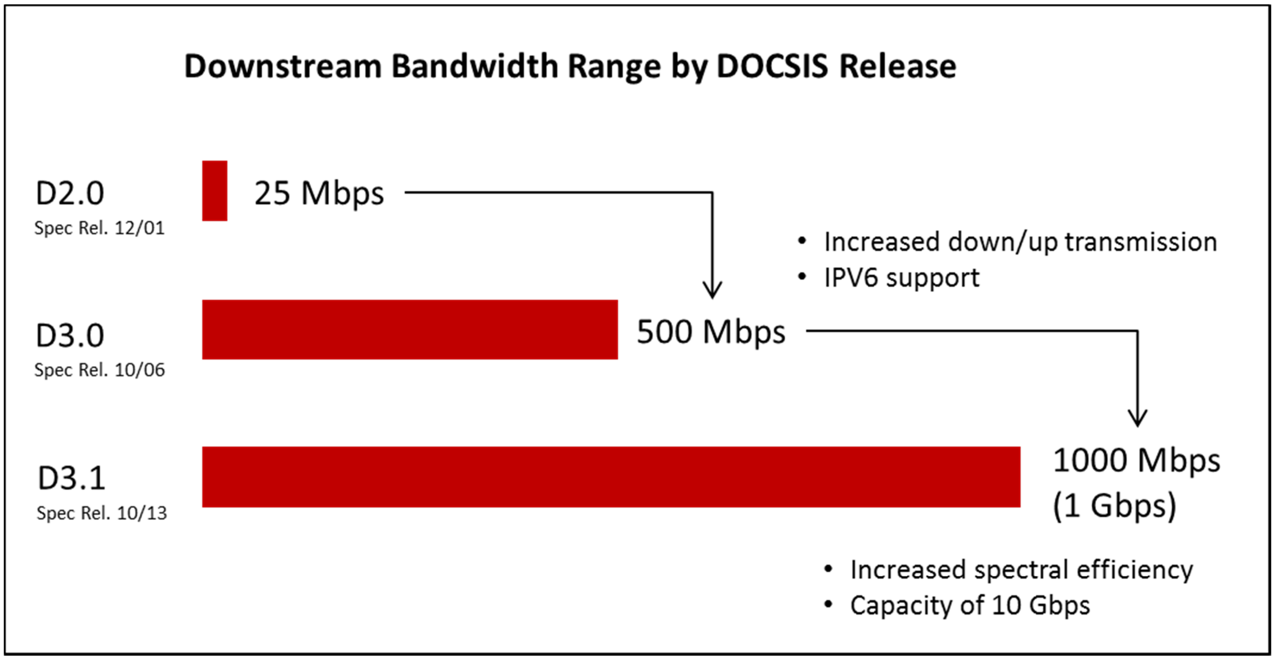

Figure 5 – DOCSIS Path to Gigabit Speeds

DOCSIS 3.1 initial silicon has been released and compatible cable modem samples have been in lab evaluation for more than 6 months. Preliminary field trials have been deployed beginning in Q3 of 2015. System level deployments have been announced and are expected to start in late 2016 / early 2017(8).

OFDM modulation and higher order subcarrier QAM formats allow DOCSIS 3.1 to increase the number of bits in a given channel block thus improving the available capacity. But this is still shared spectrum depending on the number of homes served by the local node. Subsequent sections of this paper will describe the architecture choices that determine the ultimate data capacity available per individual subscriber.

THE GIGABIT ERA
WHICH PATH TO CHOOSE?

Competition and HSD growth curves will continue to push MSOs to deploy gigabit and higher HSD capable networks or risk losing subscribers to a new entrant service provider with claims of a faster network. There are multiple solutions that a cable operator could employ to migrate the existing HFC network to gigabit per subscriber capability. The tradeoffs between different solutions usually are centered on the network migration costs and the expected useful lifetime of the network change.

The next sections describe the features of the possible migration options followed by details on the cost analysis of each migration path.

Fiber to the Home

The housing boom that occurred in the early 2000’s saw many developers insisting on a fiber ready or fiber to the home solution for their properties. Fiber was considered a future proof technology and home builders were able to charge up to $10,000 more for a fiber enabled home. When the housing market began to recover after the 2008-2009 recession a shift to MDU greenfield properties had occurred but the demand for fiber only entry into the building continued. In most cases contracts for these properties can extend for 5 to 10 years and many times the selected service provider has an exclusive contract making these MDU's the prime target for over builders and too important to ignore for MSOs.

Fiber to the basement (FTTB) has been a long standing MDU design solution for many years. Fiber is brought into the building to feed a standard HFC node. The output of the node is distributed over coax to the individual subscriber housing units. In cases where the developer demands fiber to the home service, RFoG is the deployed solution since it is transparent to the existing HFC back office and installed CPE equipment.

Unfortunately, FTTB and RFoG are limited to current HFC data rates and vulnerable to gigabit competitors. RFoG is not an ideal solution for other reasons. It's biggest weakness is optical beat interference (OBI). MDU's tend to be a worst case environment due to potential ingress sources and the fact that many of these new communities are premium subscribers with a higher than average number of connected devices that tend to increase the chances of OBI.

To counter the gigabit competitive threat a PON overlay must be deployed. The most commonly deployed PON formats today are GPON which provides 2.5 Gb DS / 1.25 Gb US and EPON which provides 1 Gb DS / US. These are PHY layer rates including packet overhead that will reduce the actual delivered data rate. The primary attraction of GPON and EPON is the use of lower cost uncooled optics. This made PON CPE pricing more affordable.

GPON is capable of supporting peak rates of 1 Gb symmetrical service while EPON's lower 1 Gb Phy rate means that packet overhead limits the peak rates to the low 900 Mb range. Speed test measurements can also be highly variable since the Ethernet switching chip sets used in commercially available modems, laptops, and WiFi equipment are limited to a maximum of 1 Gb. Measurements made in a home environment can typically range from 800 Mb to 940 Mb.

A challenge unique to cable operators is that DOCSIS HSD tiers are commonly guaranteed delivered rates rather than best effort. A 100 Mb tier subscription results in a verifiable 100 Mb delivered data rate. Offering 1 Gb FTTH service using GPON/EPON could put the operator in a difficult position since the delivered rate could not be guaranteed similar to DOCSIS tier rates. Also, MSOs deploying GPON/EPON have no competitive advantage compared to rival PON over-builder solutions.

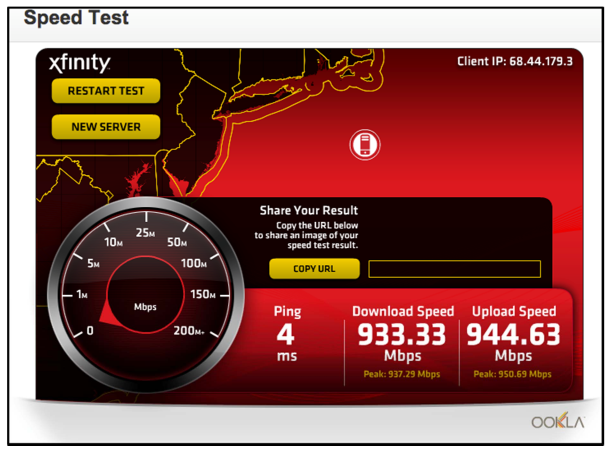

Figure 6 – Symmetric Gigabit Speed Test

10G EPON is an available solution that can solve the issues described above as well as provide expandable capacity for >1 Gb tier offerings. The cost of 10 Gb optics while higher than today's 2.5 Gb un-cooled optics are quickly decreasing based on growing volumes in North America and China. The Telco market is beginning to feel the pressure from emerging 10G EPON competitors. FSAN, the ITU standards organization that defined the GPON protocols has recently initiated a new XGS PON spec that is compatible with 10G EPON optics. Before long current GPON competitors will be raising the bar with 10 Gb service offerings.

FTTH construction costs are significantly higher than any other HFC alternative. While fiber cable and coax material costs are nearly equal the added material and labor costs associated with optical passives, fiber splicing, and connectors are very different. In building fiber to the home wiring is the major cost driver for this architecture.

Business as Usual (BAU) Node Splits

This is basically a strategy of delay allowing usage based node splits to keep up with increased data requirement hot spots until a clear view of the future migration path becomes evident. This migration tactic does not provide a direct path to gigabit per subscriber HSD capacity.

Historically, the number of node splits as a percentage of the total network has been consistent at roughly 2.5% to 3% for several years and remains very manageable. Since this level of node splits is simply a sustaining operation it does not protect against the accelerated growth rates predicted by Nielsen and similar CAGR estimates. It also does not provide any protection from competitive threats.

Traditional node splitting results in the highest network migration costs. Based on the 50% yr/yr capacity growth projections the number of node splits needed would go expediential within 4 to 5 years. HE and Hub capital equipment, rack space, power, and HVAC additions needed to support the increase in node splits will be a major driver to the costs of this migration approach.

1 GHz Drop-in RF BW Expansion

Another often discussed migration approach is to continue the traditional RF BW expansion strategy that has been used many times in the past. Changing out the RF electronics in the existing plant node and amplifier actives would provide additional RF spectrum that could be used for D3.0 channel bonding or to create room for a future channel block of D3.1 OFDM. Improvements in GaN RF hybrid and MMIC gain blocks allow the expansion to 1 GHz without the need to respace the location of the existing plant actives. All of the OSP actives would need to be touched in order to effect the DS BW enhancement. Therefore the opportunity exists to modify the upstream split at the same time thereby increasing the return path digital capacity. While the initial cost of this network upgrade method is very favorable compared to every other alternative, there are still a number of very real downsides.

1 GHz lasers, nodes, and amplifiers have been available for over a decade. 1 GHz taps and passives have been the standard for even longer. Every OEM manufacturer has obsoleted 750 MHz and 870 MHz versions of their actives portfolio and today only offers 1 GHz qualified products. Plug in accessories allow the gain and tilt to be adjusted such that the drop-in RF electronics module can be backward compatible with legacy equipment bandwidth designs or provide the necessary drive to obtain 1 GHz extended performance. In actual operation many cable operators have discovered that cascaded roll off factors at the high end of the RF spectrum limit the usable bandwidth to roughly 940 MHz. The frequency roll off is a cumulative effect caused by the RF response signature of the distribution taps and passives as well as some older HE lasers and OSP actives. The end result is a reduction in signal to noise performance that makes operating channels above 940 MHz extremely marginal. As a result, the cost of the BW expansion may not be as attractive with only a 70 MHz (for an 870 MHz legacy system) or 190 MHz (for a 750 MHz legacy system) frequency spectrum gain. This is especially true if the upstream split is modified at the same time further reducing the total DS frequency increase.

Another downside factor to consider is that while a 1 GHz drop-in frequency expansion increases the potential digital capacity of the system this is still shared bandwidth in a relatively large node serving group. Also, roughly 75% of existing multi-port nodes are unbalanced meaning that one leg of the RF distribution access link leaving the node has considerably more subscribers than the other RF legs. As a result, unless there is a significant reduction in the amplifier cascades on the unbalanced node leg, the effective data rate increase on a per subscriber basis will be relatively low. The drop-in approach will therefore still require future node splits to meet the competitive threat. These added node splits increase the capital expenditures required to support the new serving group additions.

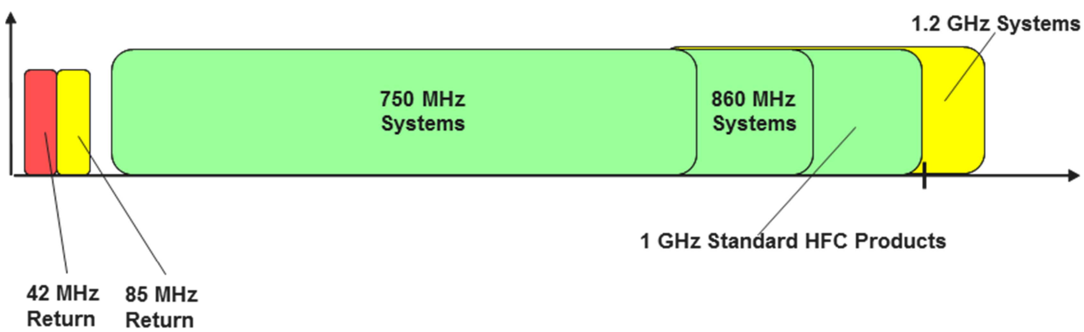

Figure 7 – HFC RF Spectrum

<u>1.2 GHz Drop-in RF BW Expansion</u>

A 1.2 GHz drop-in BW expansion will be several orders of magnitude more difficult and costly than 1 GHz. Increasing the usable bandwidth to 1.2 GHz means that every part of the network must be touched. CMTS, headend RF distribution, OHE lasers, nodes, amplifiers, taps, and CPE. The sheer scale of this level of changes without reducing the serving area of the nodes makes this architecture choice very questionable.

The higher coax attenuation at 1.2 GHz requires higher output levels and higher tilt to compensate for the increased loss. Node designs for 1.2 GHz have been available for a few years but these early products have all been designed using hybrid gain blocks that are essentially 1 GHz devices with the RF response re-tuned to meet the 1.2 GHz bandwidth requirement. The output level specified for these nodes is typically lower than a 1 GHz comparable model due to the increased channel loading. More recent development work geared to N+0 Fiber Deep requirements has generated higher power hybrid gain blocks including versions with more appropriate levels for an N+X cascade design.

The added channel loading for a fully loaded 1.2 GHz laser results in a 0.6 dB drop in laser OMI per channel which is equivalent to a 0.6 dB reduction in carrier to noise at the node receiver. The OMI reduction also requires an increase in node receiver sensitivity and gain to compensate for the lower modulation level.

RF Bridger and Line Extender amplifiers for 1.2 GHz are not yet generally available. Vendors supplying the European markets are several months ahead of their north American counter parts since they are seeing a higher level of interest from western European cable operators. Concerns about the increased levels needed to avoid re-spacing and the actual value proposition for a 1.2 GHz extended BW cascade have delayed the introduction of these products in the north American market.

FIBER DEEP N+0 ARCHITECTURE

Overview

Fiber Deep architectures have been discussed for many years and have different definitions depending on the author. The variations are largely based on the number of amplifier actives allowed after the node (N+small or N+0). N+small network designs provide a lower initial cost but eventually will require additional segmentation to fully optimize the digital capacity per subscriber. The end state of HFC Fiber Deep is N+0 which is the definition used for this paper.

An N+0 architecture comes closest to the ideal of a last touch HFC network design. A brief list of goals for this topology is the following:

- Centrally locating the node to provide balanced outputs
- Optimally reducing the node serving area to achieve the desired data rate tier capacity per HHP
- Eliminate all actives following the node to improve OpEx performance and provide a passive coax access link to the home
- Maximize the effectiveness of D3.1 and Distributed Architectures
- Provide the same DS data capacity as 10 Gb PON with the migration to D3.1 and 1.2 GHz bandwidth expansion
- Establish the FD node location as the Edge Network demarcation launch point for future system migrations such as Full Duplex DOCSIS or FTTH
- Extend the competitive life of the HFC network for many years to come

A Fiber Deep architecture is not without challenges. The previous section of this paper described the high touch impact of a 1.2 GHz network migration. Node + 0 increases the complexity due to the changes that need to be made to the actual HFC plant configuration.

Fiber Deep relies on new technology developments in order to be a successful and practical architecture. Depending on the homes passed density of the target system the number of new nodes required can range from 10:1 to as high as 16:1 compared to the existing N+X cascade design. The impact on the available rack space within the local Hub or HE can be dramatic. CCAP and higher density OHE platforms help to conserve rack space but the transition to smaller serving group's means shifting to 1:1 port count associations with lasers and nodes. There are very few Hubs with the space, power, and HVAC capability to accommodate the growth in equipment associated with this significant increase in nodes. Without a relatively near term solution new Hubs would need to be built to support the increase in FD nodes. The cost and time needed to acquire property, route connections and build out the infrastructure is prohibitive. This is the inevitable problem with continuing BAU node splits but has also been the main show stopper for alternative cascade reduction architectures, including fiber deep. The solution is to disaggregate and distribute the Layer 1 elements and some of the management functions from the CCAP platform out to the node. Distributed Access Architecture (DAA) solutions are being developed and evaluated today that will significantly reduce Hub density. These solutions could be available in the market starting the second half of 2017.

Other development efforts such as all IP transport and SDN/NFV will further reduce the current equipment density in the Hub and could allow the eventual consolidation of secondary Hubs into a master headend.

N+0 Node Size and Effective Reach

Up until the release of DOCSIS 3.1, fiber optic PON networks have had the potential advantage of providing the highest delivered data rates to an individual subscriber. In actual systems this was rarely realized since FTTH networks typically over subscribe the available data capacity to keep capital equipment costs as low as possible. Google Fiber, AT&T, and the multitude of new over-builders that are now challenging the established MSO networks have changed the paradigm by offering 1 GB data at a low monthly cost. DOCSIS 3.1 provides the digital capacity to compete with PON. Reducing the node serving area size allows an HFC network to optimize the data rate capacity available per subscriber. A target node size of 128 HHP delivers the same data rates as a 10G EPON port and positions the network for a future transition to PON when needed.

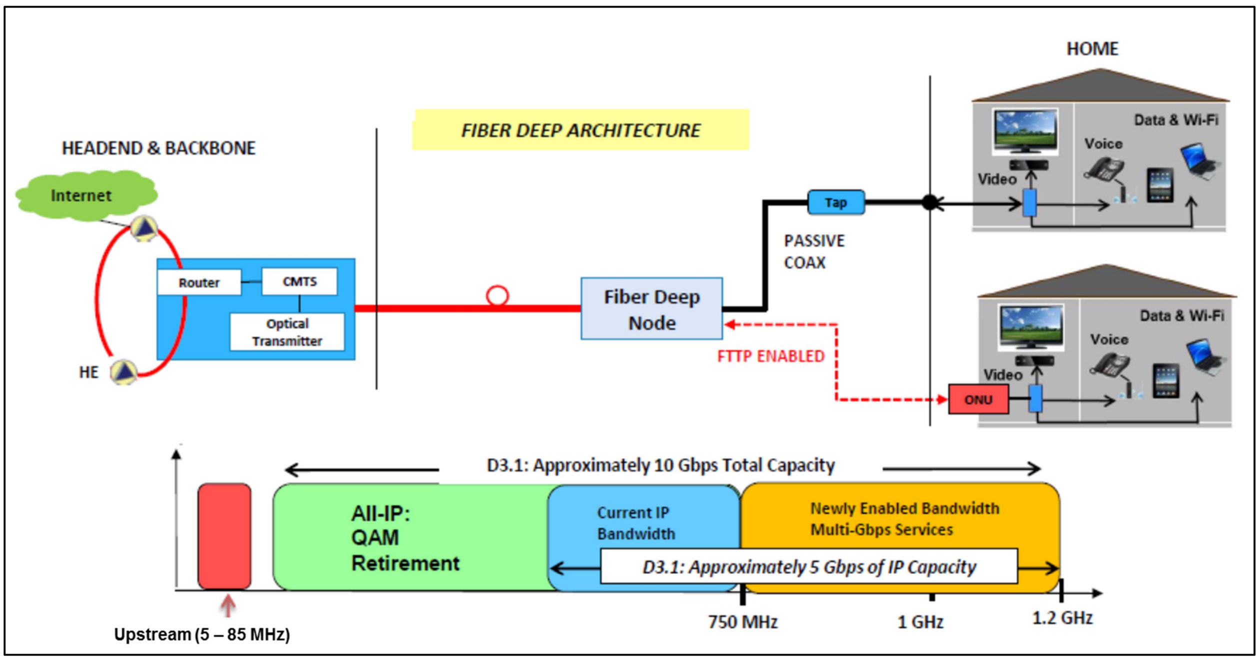

Figure 8 – Fiber Deep Architecture and Bandwidth Diagram

Figure 8 shows a pictorial of the Fiber Deep architecture design along with a bandwidth diagram depicting the current and future loading. The diagram also shows the potential for multiple services such as PON or P2P Ethernet connections emanating from the HFC FD N+0 node location.

As discussed previously, providing sufficient BW to take advantage of the full benefits that D3.1 offers requires extending the BW to 1.2 GHz. Coax plus passive losses at 1.2 GHz are considerably higher than a 1 GHz HFC network. By comparison, a DOCSIS 3.0 N+0 node design requires 60 dBmV output (analog equivalent power) and a tilt of 18 dB measured from 54 MHz to 1002 MHz. To achieve the same equivalent reach and performance at 1218 MHz the tilt line must be linearly extended 4 dB with an equivalent increase in the node output level. This seems like a daunting task but the limits of GaN semiconductor technology have not yet been reached and after a year and a half of development effort two major device suppliers succeeded in creating qualified hybrid gain blocks that are now commercially available. This accomplishment was only possible as a result of working in close partnership with multiple node design teams to solve technical hurdles related to power consumption and thermal capacity limitations of the node housing and cable plant powering.

The physical characteristics of an HFC node are constrained by OSHA regulations for the safety of the installers. Other standards are also closely watched such as the National Electric Code rules, Utility service provider requirements, and intra-service agreements which all define the physical spacing between different services hanging on the outside aerial plant or coexisting in underground conduits and access tunnels. The mechanical dimensions of the housing play an important role in determining the maximum power dissipation that the node can handle. Careful design including optimizing switching power supply efficiency has resulted in the high output Fiber Deep node consuming the same or lower power than a legacy 4X4 node today. Further improvements are still possible including efforts underway today to introduce predistortion techniques to the RF power amplifier stages in order to lower power consumption while maintaining performance.

At a target node size of 128 HP only very high density urban serving areas would possibly require a segmented node. In almost every other case a 1X1 centrally located four output node will accommodate the current legacy network densities. Unfortunately, even D3.1 mid split band plans will not resolve the asymmetric capacity of the available US BW so provision for 1X2 segmentation is incorporated in the node requirements.

Field trial designs have shown that at 64 dBmV node output the 128 HP maximum cap is an appropriate target. In reality, home densities vary from system to system and neighborhood to neighborhood so there is a wide distribution that occurs in most FD designs. The average node size based on field trial deployments is 64 HP. This value is used to estimate BOM equipment requirements in the HE and outside plant for all new fiber deep builds. At an average 64 HP node size the typical FD node ratio is a 12:1 increase over the existing legacy network.

Fiber Deep RF Distribution Components

The configuration change to 1.2 GHz N+0 requires modifications to every aspect of the network. Distribution Taps and Passives have been designed that maintain the current 1 GHz performance plus extend the bandwidth to 1218 MHz maintaining the same monotonic insertion loss response characteristic. A major cost saving feature of these devices is that the Tap faceplate is mechanically identical to the legacy housing allowing the tap housing body to stay in place and only the faceplate to change. In FD field trial designs 60% of the tap values must be changed. The ability to change the faceplate instead of cutting out and replacing the entire tap housing amounts to a considerable construction cost savings.

While creating the new 1.2 GHz faceplate designs the tap vendors also made improvements to the surge resistance and shielding effectiveness of these devices. These changes ruggedize the product against electrical storms and power line issues as well as improve ingress resistance to LTE and other interferers. Every major tap manufacturer has followed this same guideline and plans to obsolete and replace the current 1 GHz devices with 1.2 GHz standard product later this year.

The Mid Split Dilemma

The most difficult part of implementing any new architecture is making any change to the upstream that affects customer premise equipment. Data usage between downstream and upstream has always been asymmetrical. In the late 1990's and early 2000's as internet traffic was growing by leaps and bounds the ratio of DS:US traffic had narrowed to 5:1. In 2007 Netflix initiated video streaming and since then the popularity of over the top video caused the ratio to widen again to 10:1. Today, home monitoring and the increasing number of smart connected devices in the home is beginning to pressure the ratio back downward once again. DOCSIS 3.1 and Fiber Deep do not remedy the asymmetric upstream situation. They do add incremental bandwidth that can keep the DS:US ratio in check.

The dilemma with moving the upstream frequency split is the nature of the wide range of legacy set top boxes that are still in service today. Modern STBs have the capability to be frequency agile meaning that the low data rate communications pilot that the STB uses to validate the channel map of video signals that can be received can be moved across a wide range of frequencies and the STB will search, find, and lock in the new pilot channel. Different generations of STB's had different levels of frequency search capability. Almost all of these newer boxes can be reset via uploaded software.

One of the most widely distributed digital set top boxes is also one of the oldest in service, the General Instrument DCT2000 developed in the mid 1990's. This box has very limited memory size and therefore cannot be reprogrammed to increase its agility. The only solution in this case is to remove and replace the DCT2000's in service. This entails identifying and informing the affected subscribers, arranging to swap the box with a new model, and verifying that the changes are all in place before transitioning to the new mid split. Any boxes that have not been replaced will lose connectivity when the mid split is activated. The replacement process can take 3 to 4 months if everything goes well. The customer care aspects of a process like this must be handled very carefully and professionally to be successful.

Addressing MoCA

MoCA is the primary home networking solution for HFC CPE equipment. MoCA runs over the existing in-home coaxial cabling, enabling whole-home distribution of high definition video and content. MoCA 1.1 and MoCA 2.0 are the currently available versions. MoCA operates in 50 MHz wide channels starting at 1125 MHz (Ch D1) through 1650 MHz.

MoCA 1.1 is the most widely deployed version and provides 175 Mbit/s net throughputs (275 Mbit/s PHY rate). MoCA 2.0 supports two performance modes, Basic and Enhanced, with 400 Mbit/s and 800 Mbit/s net throughputs (MAC), using 700 Mbit/s and 1.4 Gbit/s PHY rates, respectively. MoCA 2.0 is backward compatible with MoCA 1.1.

MoCA allows frequency agile channels and most equipment performs a search for the optimum channel when initially configured. Due to coaxial losses and in-home frequency response characteristics, the lower channels tend to be the most utilized. Many MoCA 1.1 consumer off the shelf products only provide channel D1 support. The end result is that MoCA operates well with 750 through 1 GHz D3.0 in-home networks but the lower D1 and D2 channels will directly interfere with 1218 MHz D3.1 when deployed.

MoCA 1.1 and 2.0 MAC throughput rates will not support D3.1 targeted gigabit data rates. Next generation MoCA 2.5 and 3.0 specifications are currently in development with plans for a phased release schedule. Both will significantly increase the usable data rates providing true gigabit per subscriber rates. MoCA 3.0 includes a shift in the RF spectrum requirements that removes channel D1 and D2 allowing the full 1218 MHz D3.1 bandwidth to be used in the future.

Power Plant Considerations

Cable plant power supplies use a ferro-resonant transformer design that provides 60 or 90 VAC that powers the node, RF amplifiers and other actives connected to the coax access network. The legacy cable plant was originally designed with nodes located near trunk fiber access routes and local utility power locations which typically ran along major highways. This provided easy accessibility to the node in case of failures or normal maintenance. The cable power supply which powered the node and associated amplifier cascade was almost always co-located with the node. Additional power supplies were placed in the system if needed to overcome I^2R losses in long access cable spans or in serving areas with a large number of active elements. Cable plant power supplies are available in a wide range of amperage. The most common power supply size is 15 amps. The primary input power source for the cable PS is local utility AC power. MSOs have traditionally designed back up power systems in case of outages in the local utility AC mains.

There are many degrees of freedom when designing and implementing access network architecture changes. Changing the power grid is not one of them. Adding, relocating, or significantly modifying a supply requires high level interaction and permits from the local municipality and utility power provider. A change in the power grid can take a year to negotiate and implement. This is practically a show stopper for most system deployments. Therefore, for everything except new greenfield construction, any new equipment deployments must safely fit within the margins of the existing AC power capacity design limits. It is also critical to consider future application module deployments and their impact on total network power consumption.

As previously described, an N+0 architecture requires a significant increase in the number of nodes needed to cover the existing serving area. Additionally, Fiber Deep D3.1 capable nodes include higher output RF amplifier stages in order to reach the maximum number of homes possible at 1.2 GHz. Another tenet of N+0 design is moving the node to a central location within the serving area of homes passed so that the node has balanced output ports or as close as can be reasonably achieved. These changes are made without relocating the power supplies in the area or adding new supplies. In order to power the relocated nodes, coax feeder lines are added between the PS and closest node(s). AC power is also fed from node to node by bridging the access coax tap strings between adjacent nodes.

Construction Realities – Lessons Learned

The very first in-house N+0 design of a target node area yielded 20 HP per node. After passing around the smelling salts and consulting with some of the top tier industry designers the subsequent designs raised the node service area to 70 HP per node. Since then design guidelines and training have helped to keep the node designs within the established target window to meet internally established cost targets.

Balancing HP per node RF performance design targets with the power grid design constraints has been another challenge. The ultra-conservative criteria created to accommodate higher output power nodes while maintaining total amperage capacity margins caused design time to more than double. As FD Node product development progressed it was possible to confidently reduce the margin added for each node and bring design time back to normal levels.

There has not been a major construction cycle in the HFC network space for more than 15 years. The number of skilled construction crews has decreased dramatically as these craft people moved on to other work in different industries. Ramping up construction of a completely new architecture such as Fiber Deep N+0 has highlighted the issues of finding and training new crews. Retaining these crews is another challenge since AT&T and other over-builders are also ramping up their own projects to compete with cable.

An example of the initial skill level of the current crews is the number of coax splices performed per hour. During the peak of the 1990's upgrade cycle the average number of splices completed was 4.5 /hr. In 2015 during the build out of the first N+0 trials the average number of splices completed was only 1.2 /hr. Even bucket trucks have become a rare commodity these days with backlogs extending up to a year in some cases. The rollout of Fiber Deep will accelerate as construction crew training improves and new practices become standards. FD is a multi-year deployment effort.

The success of commercial services product offerings has created a business segment that is a growth leader in our industry. Business services customers were practically non-existent at the time of the last network upgrade cycle so establishing procedures to inform and work with these customers during construction is critical. Node cut in schedules must be determined based on the minimum impact to the customer. This takes a lot of extra planning and over communicating to the businesses that will be affected by the network interruption.

NETWORK MIGRATION COST COMPARISONS

Cost comparisons of the various network migration options were made based on a 10 year outlook. The comparisons assumed that in the case of BAU node splitting more FTTH deployments would be needed to counter PON over-builders. Another assumption was that D3.1 would provide a defense against FTTH competition. Therefore 1 GHz RF Active Drop-In (ADI) would provide some relief against having to build FTTH in order to retain subscribers. Even less FTTH builds would be needed in the case of Fiber Deep since the data capacity of FD was equal to PON competition. Figure 9 shows the relative cost comparisons between the four most likely deployment scenarios.

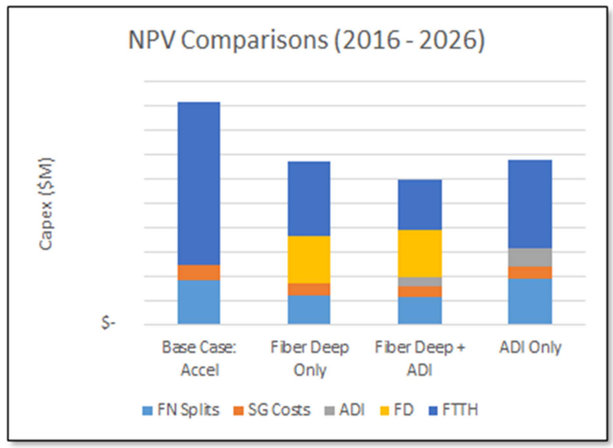

Figure 9 – 10 Year NPV Cost Comparison

Note: ADI includes the cost of subsequent node splits over the 10 year period.

The cost of Fiber Deep is almost the same as 1 GHz Active Drop-In although the expected lifetime of the plant built with ADI is significantly shorter. The best case scenario as seen in the graph is a combination of Fiber Deep plus some amount of Active Drop-In. The assumption in this case was that the ramp up of FD would take time and ADI could fill in the gap to enhance that would be under threat before FD could be built out in these locations. The worst case scenario was the base case of accelerated node splits.

WHAT COMES AFTER FIBER DEEP?

The first 12 to 18 months of Fiber Deep deployments will be built taking advantage of the higher density improvements coming from CCAP and new RF distribution modules in the HE and Hub. Older equipment will be cleared out to make room for optimally laid out racks supporting the new FD serving areas. Available rack space in the existing Hubs is not capable of meeting the 12:1 increase in nodes that complete conversion to FD will create.

To be successful, Fiber Deep must as quickly as possible change to a Distributed Access Architecture by changing out the analog optics with digital optics and adding an R-Phy or R-MacPhy module to the nodes. DAA will bring significant improvements in Hub density and provide the pathway for virtualization allowing even higher density results. Both R-Phy and R-MacPhy designs are in the lab today with early trials of both technologies being planned for later in the year.

A goal of the Fiber Deep N+0 design has been to strive for a last touch architecture. The introduction of DAA has always been known to be a necessary second phase of the Fiber Deep evolution and roll out plan. Nodes deployed prior to the changeover to DAA will be converted by removing the DS analog optical receiver and US digital return modules and inserting a DAA module which includes 10Gb digital optics. All of the Fiber Deep nodes qualified for deployment are designed to accommodate and provide for a simple change over to Distributed Architecture.

As stated earlier, DOCSIS 3.1 provides over 5 Gb of downstream data capacity but upstream usable data rates will remain asymmetrical for 1 Gb and higher tier offers. A new concept that has been demonstrated in the lab provides the capability to transmit DS and US data over the same frequency spectrum simultaneously. The benefit is an immediate expansion of the upstream bandwidth and the ability to carry multi-gigabit symmetrical US traffic.

This technology christened “Full Duplex DOCSIS” is an extension of principals used in DSL system designs. Full Duplex DOCSIS requires a passive coax network to operate effectively so N+0 access networks are ideal. Full Duplex DOCSIS was recently announced at the 2016 CableLabs Winter Conference.

CONCLUSIONS

Competitive threats from FTTH over-builders and the continuing growth of digital high speed data is driving the adoption and eventual deployment of DOCSIS 3.1 to provide expanded gigabit capacity and extend the usable life of the HFC DOCSIS network.

1.2 GHz Fiber Deep N+0 network design provides the DS and US bandwidth to maximize the benefits of DOCSIS 3.1. This implementation allows 10 Gb data capacity equivalent to 10G EPON FTTH downstream. The combination of N+0 FD plus D3.1 delivers the capacity to support gigabit per subscriber HSD tiers.

The incremental cost per subscriber of N+0 Fiber Deep is lower than most of the available alternatives such as traditional node splits, or FTTH. An N+0 network design provides a demarcation point at the node location that can be migrated to a distributed architecture or FTTH architecture or whatever path future technology innovation holds.

N+0 Fiber Deep deployments will accelerate and become standard practice as designers and construction crews gain experience and confidence in this new HFC topology.

The cable plant power grid is a fixed resource. For network operators the goal is to stay within the existing plant power capacity.

REFERENCES:

1. https://www.nngroup.com/articles/law-of-bandwidth/

2. https://www.sandvine.com/

3. Fierce Cable March 4, 2016 http://www.fiercecable.com/story/report-google-fiber-accounts-only-5100ths-1-us-pay-tv-market/2016-03-04?utm_medium=nl&utm_source=internal&mkt_tok=3RkMMJWWfF9wsRokuanIde%252FhmjTEU5z17%252BklW6Gyh4kz2EFye%252BLIHETpodcMScFiNLHYDBceEJhqyQJxPr3HJdQN18R7RhHnDg%253D%253D

4. Chart from Third Annual Report of Gig.U (August 2014)

5. http://www.bizjournals.com/atlanta/blog/atlantech/2015/04/source-at-t-ready-to-rollout-gigabit.html

6. Miguelez, Phil, Fred Slowik, Michael Aviles, "Lowering the Cost of Network Segmentation", SCTE Expo, Orlando, FL, October 17 – 19, 2012

7. www.cablelabs.org

8. Fierce Cable March 15, 2016 http://www.fiercecable.com/story/comcast-rolls-out-docsis-31-powered-advanced-trials-atlanta/2016-03-15?utm_medium=nl&utm_source=internal&mrkid=%257B%257Blead.Id%257D%257D&mkt_tok=3RkMMJWWfF9wsRokuK3Kce%252FhmjTEU5z17%252BklW6Gyh4kz2EFye%252BLIHETpodcMSsBiMLrYDBceEJhqyQJxPr3HJdQN18R7RhHnDg%253D%253D

Acronyms and Abbreviations:

ADI	Active Drop-In
BOM	Bill of Material
CAGR	Compound Annual Growth Rate
CPE	Consumer Premise Equipment
DAA	Distributed Access Architecture
DOCSIS	Data over Cable System Interface Specification
DS	Downstream
FD	Fiber Deep
FTTB	Fiber to the Basement
FTTH	Fiber to the Home
GaN	Gallium Nitride
GaAs	Gallium Arsenide
Gbps	Gigabits per second
GHz	Gigahertz
HD	High Definition
HFC	Hybrid Fiber Coax
HHP	Households Passed
HP	Homes Passed
HSD	High Speed Data
HVAC	Heating, Venilation, and Air Conditioning
IoT	Internet of Things
MoCA	Multimedia over Coax Alliance
MP2	MPEG2
NFV	Network Function Virtualization
OBI	Optical Beat Interference
OEM	Original Equipment Manufacturer
OHE	Optical Headend
OMI	Optical Modulation Index
OpEx	Operating Expense
OSP	Outside Plant
PS	Power Supply
QAM	Quadrature Amplitude Modulation
R-Phy	Remote Phy
RFoG	RF over Glass
SDN	Software Defined Network
SDV	Switched Digital Video
STB	Set top Box
US	Upstream
VOD	Video on Demand