What is 10G – The Technology Foundation

A Technical Paper prepared for SCTE•ISBE by

Dr. Robert Howald
Fellow
Comcast
1800 Arch Street
215 286 8037
robert_howald@comcast.com

Dr. Sebnem Ozer, Comcast

Robert Thompson, Comcast

Saif Rahman, Comcast

Dr. Richard Prodan, Comcast

Jorge Salinger, Comcast

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

The HFC architecture has proven over decades to be a most adaptable broadband network. One technological refresh after another, now coming faster than ever, has enabled the increased bandwidth and higher speeds customers demand, and with the reliability to power an always-on multi-screen, multi-device experience. Key advances in deeper fiber penetration, flexible and modular intelligent nodes, spectrum expansion, DOCSIS 3.1, Distributed Access Architecture (DAA), All-IP services, and virtualization have emerged or are emerging to drive this next generation of experiences.

While that amount of transformational change is an impressive undertaking, the HFC horizon is expanding once again with the industry-wide focus on the 10G initiative. A powerful capability being added to the DOCSIS family at the heart of 10G is DOCSIS 4.0 Full Duplex (FDX), formerly known as DOCSIS 3.1 FDX. FDX offers a massively expanded upstream spectrum and capacity on existing coaxial networks, introducing several key innovations that promise to carry the network long into the future. As a complement to the aforementioned capacity initiatives, 10G will achieve the promise of a world-class platform delivering all of the capabilities to power robust, real-time, broadband applications and services of today and envisioned (or not!) of tomorrow. A 10G network will be an intelligent, robust, self-healing, broadband, low-latency network that raises the bar on what a network can do, and how it can affect customers' lives for the better.

This paper will bring the reader from the ground up on these enabling next generation access network (NGAN) technologies that form the foundation of 10G. What were once emerging technologies are, in many cases, now being deployed. We will describe the progress and state of these access network roll-outs, rapidly coming iterations of these core technologies, and discuss the roadmap and anticipated milestones of development of these evolutionary next steps, and why they are important. Finally, we will provide a description of the technical foundations of DOCSIS 4.0 Full Duplex, including the network and home architecture implications for deployment, and hone in on some optimizations already in the works. We will share results from field trials that have taken place on the road to 10G deployment, including measured performance and comparison to the moels used by the CableLabs working groups to develop the FDX requirements. This will include the "Model 2" scenario, anticipated to be the most challenging for FDX, and truly exercise the technological capabilities. And, we will consider alternative 10G paths and the synergy between the options.

What is 10G?

Cable operators have steadily increased bandwidth and speeds to subscribers since the launch of the DOCSIS high speed data services (HSD) in the late 90s. The relentless march of compound annual growth rates (CAGRs) and the insatiable appetite for more bandwidth-intensive media, led by streaming video and increasingly complex interactive web sites and applications, has made consistent investment in the network simply the cost of doing business as a network operator. These trends have been so prevalent and reasonably predictable that operators of all stripes have been able to optimize their network investment and create efficient standard practices by which they enable more capability in their networks. As an inherent broadband medium of mixed technologies, HFC has been, and remains, uniquely positioned to adapt and continue to increase these capabilities to deliver on the experiences demanded by today's subscribers and businesses.

Cable MSOs have generally been conservative in touting their network capabilities and achievements, and technical terminology such as "DOCSIS 3.1" and "DAA" have become the parlance of the industry-wide advances that form the foundation of these noted initiatives and investments in the network. However, they are not particularly effective at telling the story about what the network can actually do and what it delivers for customers.

With the next technology leap forward, building upon initiatives like DOCSIS 3.1 and DAA, among others, the 10G initiative will illuminate what these new advances means to services and experiences delivered to the end user, in addition to telling the technology evolution story. So then – what is 10G? 10G sets clear and essential targets for the emerging cable broadband network – 10Gbps of intelligent, reliable, bi-directional capability that, when augmented with developments focused on reductions in latency, will vastly enhance the experiences of today. The 10G network will unleash an entirely new generation of applications and capabilities that will change everyday life in ways we can only project based on today's experiences – but also in ways we cannot yet fathom.

Trends in Network Architecture Evolution

Since the launch of the high-speed data (HSD) services, the need to service new bandwidth tied to the explosion of the Internet has been of paramount importance to operators. Over these past 2+ decades, broadband services have become equal to, and have arguably eclipsed, video services as most essential service offering. Compounding the problem statement for operators is that, while average bandwidth downstream to subscribers has grown consistently some 40-50% per year (and 20-35% upstream), the market-induced speed wars are forcing the addition of spectrum beyond that of what these per-user aggregate bandwidths might demand. Speed tiers amount to a minimum allocation of HSD spectrum to assure that the "speed test"– often lightheartedly referred to as the "killer app" for high speed broadband – is met with a sufficient probability of success to achieve an acceptable customer quality of experience (QoE). And, with new bandwidth tied to these speed tiers, operators must execute the balancing act of finding spectrum to allocate while negotiating the space with existing video services and the also-growing demands on that service, such as supporting 4k/UHD services and/or simulcasting IP Video services.

Figure 1 represents this problem statement. These types of bandwidth-vs-time charts and the numbers behind them are observed and projected obsessively by cable operators, as they perform that spectrum balancing act, and to support their existing and expected/future customer service demands. Such figures are exceptionally useful to derive the network breakpoints that inform operators of how much lifespan different network evolution stages offer.

Also, this perspective on growth helps to illustrate timelines that address particular network imbalances coming at them in the years ahead. Access network initiatives that include changes to construction and infrastructure inherently require significant upfront planning, and can only be managed and executed cost-effectively and transparently to end customers at a particular pace. We will discuss several of these initiatives in this section.

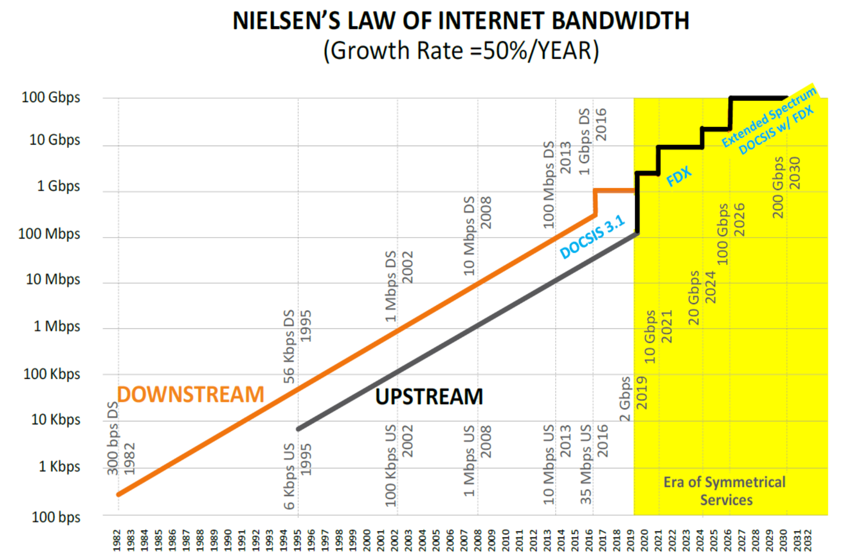

Figure 1 - Historical Compounded Annual Growth Rate (CAGR) and Relationship to Cable Eras of Broadband Evolution [T. Cloonan, Commscope (6)]

In addition to decades of focus on speed and bandwidth, new services that put a premium on decreased network latency are becoming key broadband drivers today. Today, the most notable application tied to low latency is the increasingly popular online gaming industry. Tomorrow, multiple applications in the Internet of Things (IoT) space, as well as virtual and augmented reality (VR and AR), are envisioned that also fall into this category. The history of Internet applications driving bandwidth demand has mostly been about media consumption and browsing behaviors. The bandwidth growth can be traced from voice to music to pictures to video to high quality video. Browsing in "human time" – website mouse clicks and hyperlinks – is not very demanding, when it comes to responsivity. It has generally been the download speed that has dominated the nature of the experience.

Going forward, however, network latency performance will move front and center as a new comparative metric. Gaming performance and cloud gaming platforms represent a practical beachhead for latency-sensitive applications, which are necessarily delivered in real or near-real time. Common examples include some IoT sensors, Augmented and Virtual Reality (AR/VR) games and experiences, Machine-to-Machine (M2M) and industrial applications – none of which are burdened by "human" response time constraints. End even online shopping and commercial stock trading qualify as low latency performance targets, where quantifiable financial benefits can accrue.

As games become more popular, more sophisticated, and taken more seriously, such as leagues that include financial implications to winning and losing, more attention is being given to providers and network attributes that impact the gamer. **Figure 2** is an example of such a study that is considered required reading for both game makers and network operators, and in order to properly service this exploding community. In controlled testing, the experience of the user and her/his perception of the fairness of the game is highly correlated with the network delay, and whether or not delay compensation removes the effects of delay variations.

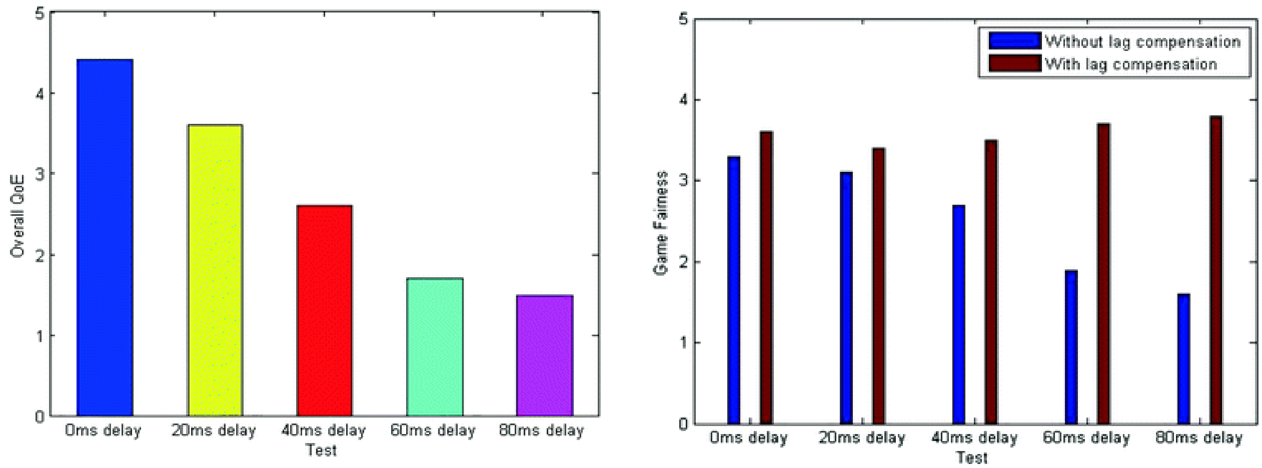

Figure 2 - Gamer Experience vs Network Delay (First-Person Shooter) [4]

In order to address some of the challenges described above, and to deliver on the promise of 10G, operators are embarking on perhaps the most ambitious set of simultaneous technology introductions ever into the access network. These initiatives are as significant and impactful, or more so, as any, going back to the introduction of fiber itself, to replace all-coaxial systems and thereby eliminate those glorious N+30 cascades. It is remarkable on two fronts – that so many technology options are occurring at once, and the pace at which the technology is changing.

Four of the major trends in the access network are Deeper Fiber penetration, Distributed Access Architecture (DAA), Network Function Virtualization (NFV), and the transition to All-IP. These trends are very meaningful in the context of FDX because, as we will describe, they are favorable to and in some cases very highly preferred complementary initiatives for the introduction of FDX. We briefly describe the what and why of these trends below.

Deeper Fiber

Cable operators have been pulling fiber deeper into their network and splitting nodes for many years, as one of the tools to manage CAGR. This involves shrinking the number of customers served by a single service group, in a manner very correlated to physical and virtual node size. Adding new spectrum and reallocating channels of existing spectrum are the other major tools available for managing growth. The same tools are leveraged for the upstream, but the flexibility that the operator has in the return path spectrum is much less. Increased average capacity per household passed (HHP) is the main benefit of splitting nodes and deeper fiber. Note that "node split" can refer to a physical split, or the segmentation of a single node such that each port is its own serving group. Such segmentation happens by way of additional wave division multiplexing (WDM) modules that isolate the RF legs from one another on a single fiber.

Many MSOs are taking a more proactive and/or aggressive approach to deeper fiber. Generally, as nodes continue to be split and the serving group size shrinks, the efficiency of these splits decreases. For example, rather than a 50/50 split of HHP or traffic loading, without additional design optimization, the split may end up being, for example, 65/35. In order to balance the split properly and maximize its effectiveness, the placement of a new node may need to change; doing so provides the opportunity to take fiber deeper, further eliminating some of the RF amplifiers in cascade.

Another approach to this scenario may be for operators to split multiple nodes at once, even though the capacity "alert level" may be many months away. In other words, when a wave of node splits is projected because capacity thresholds are expected to trigger a set of "hotter" nodes, often within a narrow time frame, it may make sense to perform preemptive splits. This makes best use of limited labor resources, improves the efficiency of what is a labor-intensive process, and allows the operator to get ahead of trends. It also means that updates to the inside plant, to support the physical plant upgrades, is also more efficient.

Lastly, several operators, projecting the long term trend of downstream CAGR and recognizing the unpredictability of upstream CAGR, have embarked on a move to go further than proactively splitting nodes and reducing RF cascades. This is often called "fiber deep" in the industry, and typically refers to evolving the HFC plant to its "natural" end state of a passive coaxial last mile, or N+0. This provides the long term runway consistent with "last touch" to the HFC plant, while simultaneously improving signal fidelity, increasing reliability, improving operational expense (opex) by eliminating amplifiers, and freeing up the last mile for additional spectrum and to efficiently re-allocate and utilize existing spectrum.

These trends can be summarized in diagrams as shown in **Figure 3.**

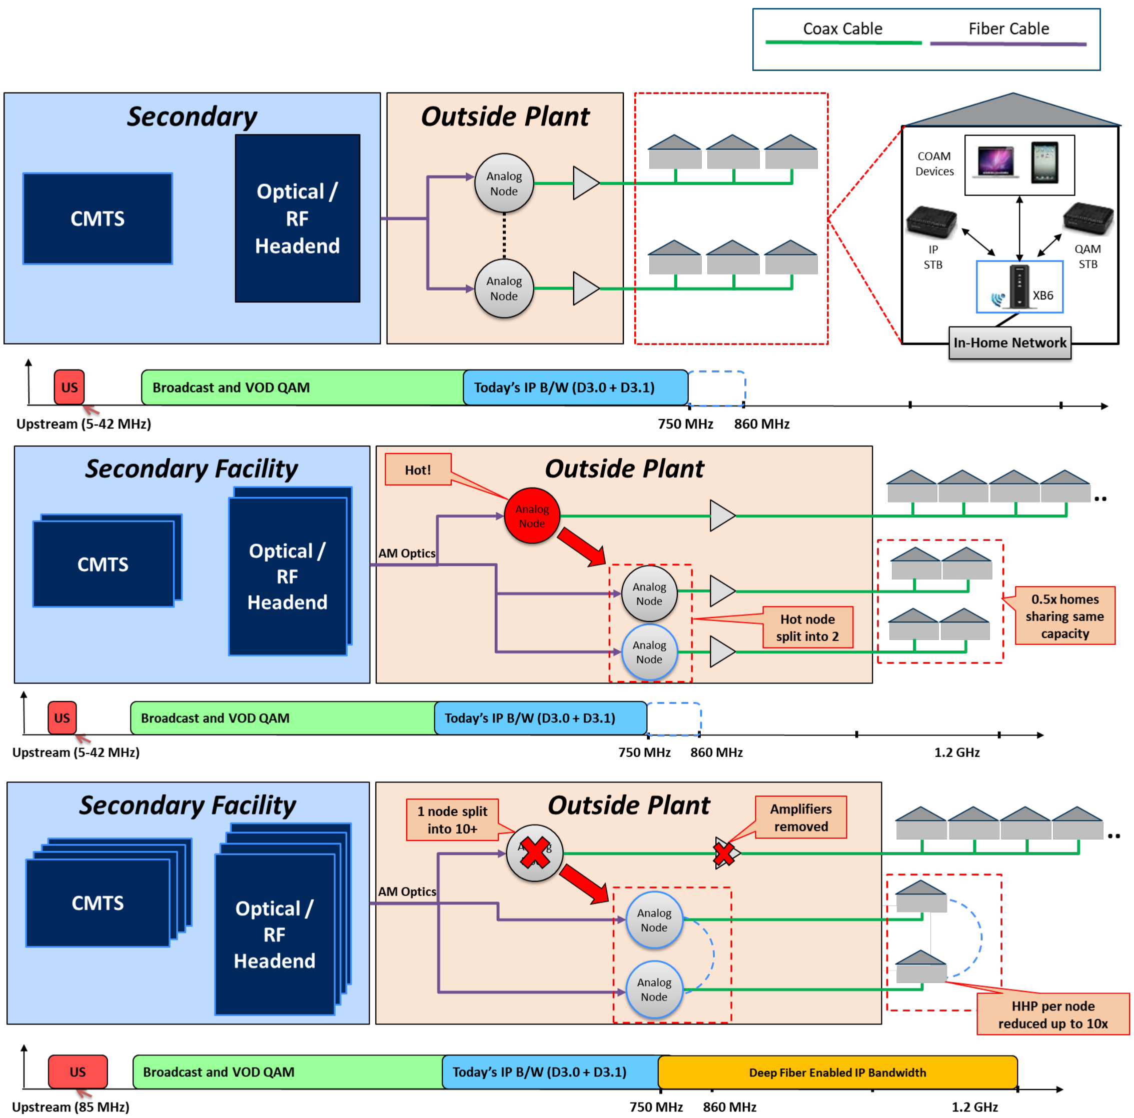

Figure 3 - Access Plant Migration Options and Tools – HFC Plant, Node Splitting, N+0

Comcast has been investing in a combination of "business-as-usual" (BAU) node split migrations. Additionally, and opportunistically, in light of the transition to DAA, we have been addressing these areas with a more holistic approach, for purposes of labor efficiency, capital investment and directional alignment to an end state vision network architecture. A DAA migration carries impacts to the network and fiber architecture that are broader than that of adding a new node via a node split, and as such are not well-suited to simply putting out the fire that is a "hot" (very actively used) node. Furthermore, the N+0 architectural step that fits in certain areas and scenarios increases the number of fiber nodes significantly – at the expense of amplifiers. In these cases, it is an opportunistic way to substantially move to a digital architecture in "chunks."

Comcast has been actively managing a holistic deeper fiber strategy, above and beyond historical node splitting, since 2015, and the optimal mix of fiber tools is reviewed annually. Comcast introduced DAA via "Remote PHY" – to be described below – in 2018.

Distributed Access Architecture (DAA)

Traditional HFC networks have long used cable-specific analog optics to carry signals downstream, and either similar analog optics for the upstream or proprietary "digital return" systems that act to digitize the return spectrum and replicate it at the headend.

Enter DAA, a term which encompasses typically three architecture variants that move an increasingly higher level of processing into the outside plant (OSP), often becoming a module in an advanced node platform. These modules are often developed to plug into mature / existing node housings that node suppliers have developed on top of for many years for HFC applications unrelated to DAA. The three variants are Remote PHY (RPHY), Remote MAC-PHY, and Remote CMTS. The pro/con lists among them are intricate - there tends to be some "religion" involved - and the debate about the "better" DAA architecture may never be resolved. Innumerable papers have been written and presented at conferences and shows about Distributed Access Architecture that detail these variants and the perceived differences between them. All of them bring the major advantage of DAA – the introduction of standardized, digital optical links, typically enabling 10GbE transport from the facility into the cable access network – and leverage the global ecosystem and massive volumes associated with Ethernet.

The transition to digital optics alone brings a host of benefits in addition to tapping into the global Ethernet ecosystem. These include:

- Improved reach and wavelength efficiency of digital fiber versus analog
- Fidelity gains (MER) that coincide with the removal of analog optical link degradation
- Physical scalability in the inside plant with the removal of RF cabling and combining networks
- Operational simplifications
- The vision of a last-mile access network that is (at last!) agnostic to the optical access layer

These MER gains and positive impacts to network operational costs have been observed in practice.

With DAA as a starting point, and considering RPHY or Remote MAC-PHY, it becomes natural to take a step back and look beyond this step, to consider where such an architecture leads. For starters, the CMTS function in the headend, having being separated from major DOCSIS-specific functionality, can now be revisited in the context of considering its new, less DOCSIS-specific, functional subset – which includes packet processing, switching, storage, and, to some extent, scheduling, in terms of compute resources. With today's compute power and the ability of software systems to deliver real-time services, a purpose-built DOCSIS machine is no longer necessarily required to provide this functional capability.

Instead, commercial off-the-shelf (COTS) servers and switches, including white-box switches, combined with software running the CMTS function from these servers, can be used to implement the CMTS function. This is referred to as Network Function Virtualization (NFV), and opens a door to a major shift in the roadmap of the CMTS, to its virtual instantiation: The vCMTS. Significant opportunities, flexibility, and benefits arise from the continuance of Moore's Law in delivering massive compute capability in increasingly physically dense footprints, with heavier reliance on software upgrades to adapt to new services and demands on the network. Comcast introduced its first DAA deployment hosted by a vCMTS in 2018.

Figure 4 conceptualizes the concept of DAA and its natural alignment with a path to vCMTS via NFV. As noted in the figure, and alluding to the access-agnostic last mile benefits of DAA, the concept can apply to both a DOCSIS and a PON / FTTH last mile, as well as potentially wireless last mile techniques – and bring synergy to the converged deployment of these technologies.

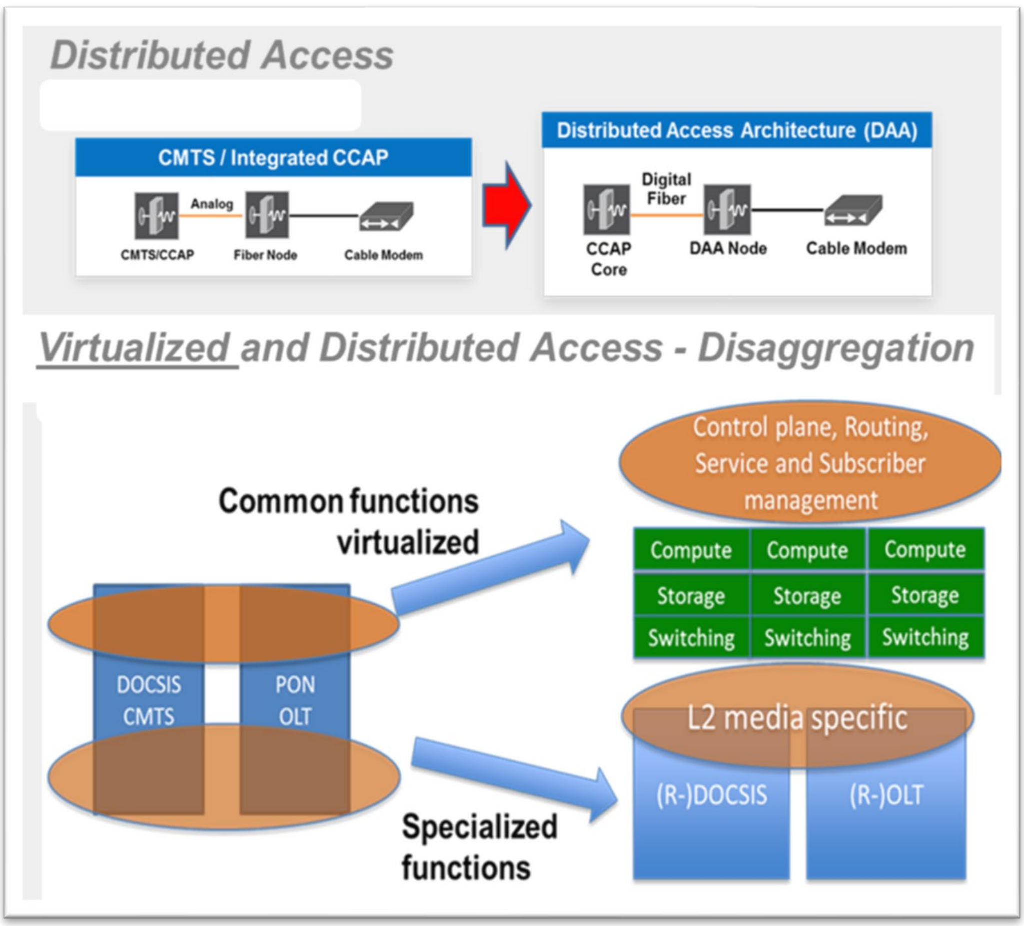

Figure 4 - DAA and Alignment to a Path Towards NFV

Transition to All-IP

Cable systems, of course, began life as one-way broadcast analog video networks. In the 1990s they evolved to introduce digital video, massively improving the spectral efficiency in comparison to analog-only, and paving the way eventually for hundreds of channels of HD content. Better video quality in the form of "Blu-Ray" quality HD, High Dynamic Range (HDR) video and 4kTV are all viable emerging advances in video.

Cable developed the legacy QAM video architecture to fit nicely within the 6 MHz slots around which analog video signals were deployed on the network. This architecture relied on fixed bit rates per 6 MHz, and the video content was groomed and aggregated accordingly to fit in these channels. The analog video legacy is the sole reason that these QAM channels were developed to this 6 MHz alignment. With the increasing shift of downstream spectrum to IP services via DOCSIS, the more efficient use of the spectrum via DOCSIS 3.1, and the "breaking" of the 6 MHz rules to allow for wideband OFDM, there is no longer a need to be tied as tightly to legacy QAM channel constraints, like 6 MHz channel widths. Nonetheless, millions of STBs are deployed that rely on QAM video – so the lifespan of QAM video, though it may be on its way to retirement, is still quite long.

In its place will not be a new and more efficient "video QAM" system, but a leveraging of the maturation of the IP network and DOCSIS growth on the plant to deliver video services. To be sure, there are major changes behind the access edge and in the home that are needed to deploy robust end-to-end IP Video. But transporting video via IP and over Ethernet has been part of cable operators' content delivery networks (CDNs) for many years, so there certainly are components in place and familiarity. Most operators have started down the path of IP Video in some form, be it in-home IP distribution to multiple screens, or VOD services delivered as IP, for example. Comcast began deploying all-IP services to customers in 2017.

DOCSIS 4.0 Full Duplex as a 10G Solution

Following on the heels of the very successful DOCSIS 3.1 initiative is DOCSIS 4.0 Full Duplex DOCSIS, or FDX. DOCSIS 3.1 tapped into network investments made over many years, bringing Gigabit speeds over HFC to the masses in one fell swoop of technology innovation. FDX enables multi-gigabit capability to and from the home over a coaxial last mile, massively increasing upstream capability compared to today's HFC. Having their origins in one-way broadcast networks, and through network investment and innovation, operators have turned these networks into best-in-class two-way broadband networks. With the inherent nature of consumer traffic being highly asymmetrical, cable operators optimized to meet this type of network utilization for data services.

Nonetheless, as emerging and future applications trend towards more upstream utilization – be it from today's games and HD security cameras to tomorrow's IoT sensors and telemedicine applications operators are swinging into action to get ahead of this next service mix transition with FDX, which allocates significantly more bandwidth on the cable upstream path without paying a penalty in downstream bandwidth, thereby staying inoculated against continued CAGR increases.

A brief summary of the key attributes of FDX is below.

FDX Fundamentals

Figure 5 identifies the fundamental trait of FDX with respect to cable spectrum – access to more upstream. It also makes clear a fundamental challenge of FDX: That the spectrum called out for new upstream capacity is also used for downstream transmissions, on the same cable. This flies in the face of decades of Frequency Domain Duplex (FDD) operation of the cable plant, where a portion of the spectrum was set aside for downstream, such as 54-1002 MHz, and another portion was set aside for upstream, such as 5-42 MHz in North America.

So how does this spectrum model proposed by FDX work?

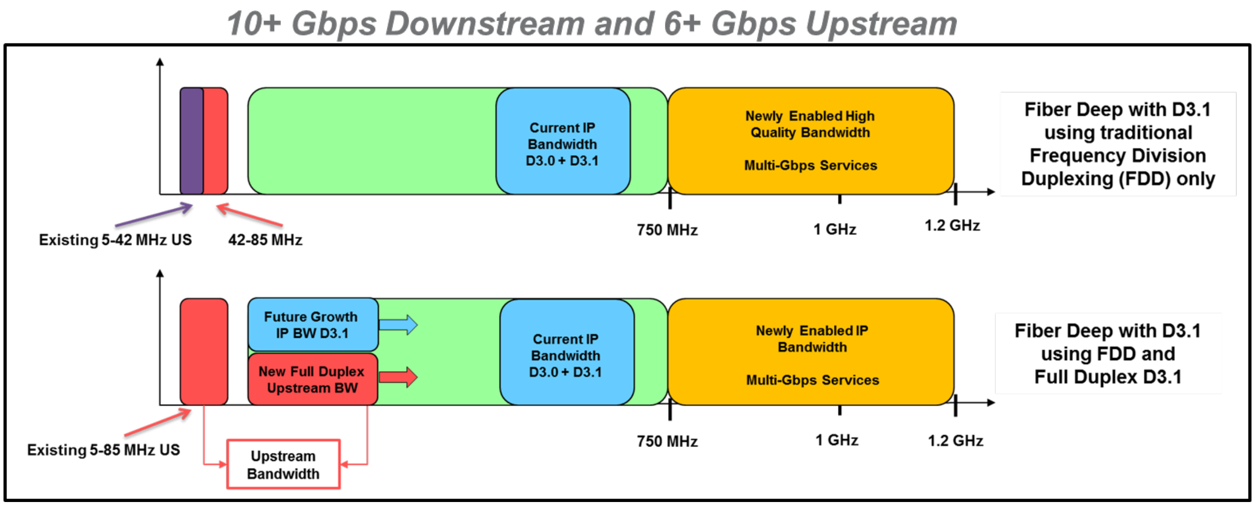

Figure 5 - Introduction of New Upstream Spectrum for Full Duplex DOCSIS 3.1 (FDX)

Key Innovations for FDX

First, it is important to recognize that in the original name, pre-DOCSIS 4.0, "Full Duplex DOCSIS 3.1" contains the term "**DOCSIS 3.1**." FDX defines significantly more upstream bandwidth, but it uses the same physical layer already described in DOCSIS 3.1 – in particular, 96 MHz blocks of OFDMA, using 25kHz and 50kHz subcarriers, the same range of adaptable QAM formats, and the same low density parity check (LDPC) for forward error correction (FEC). In FDX, six more of the 96 MHz blocks are added as potential new upstream in the 108-684 MHz band, when FDX is implemented as a complement to an 85 MHz "mid-split" system. This is shown in **Figure 6**.

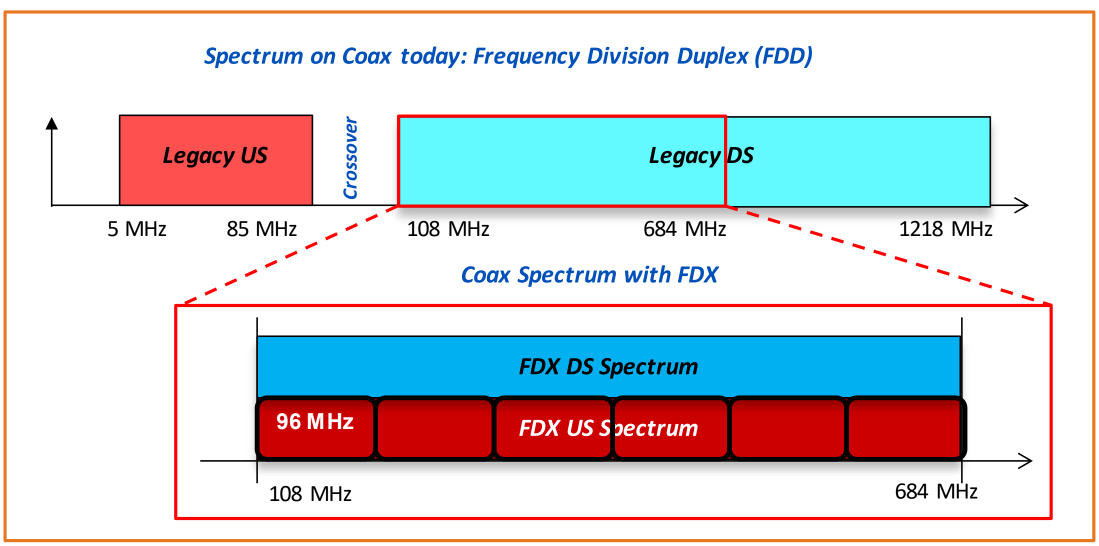

Figure 6 - OFDMA Blocks Enabled in the FDX Band

<u>And For Our First Trick….</u>

Of course, there is the issue of this 108-684 MHz being occupied by downstream carriers on the same coaxial cable. Simultaneous carriage of downstream and upstream traffic on the same carriers is not how a classic Frequency Domain Duplex (FDD) system works, nor how the coaxial spectrum has been implemented for decades, since its inception in the early '90s. However, in the xDSL telco space, and even far before that (considering simple voice telephony applications), technology had been developed to handle this type of impairment in the twisted pair environment. As xDSL became more sophisticated in achieving broadband speeds, these techniques and algorithms, which are known as Echo Cancellation (EC), became more sophisticated and capable, evolving along with the xDSL technology, right up to the G.fast versions of today. Therefore, Echo Cancellation is in practice quite mature in telco technology – not exactly the same as cable broadband requirements, but conceptually quite similar.

One significant architectural difference between cable and twisted copper pair is that the links from plant to home are the point-to-point links used by telcos. There are often multiple pairs to a home, but each is a point-to-point pair. By contrast, cable has operated a shared, point-to-multipoint architecture over the HFC network for decades. This is an important logical architecture difference that requires another innovation for FDX to be workable.

Based on the above, then, there are two new technical "tricks" that make FDX possible – broadband, high dynamic range Echo Cancellation and *Interference Group/Transmission Group* (IG/TG) classification and management. **Figure 7** represents these two new techniques required to make FDX work, from the perspective of the CMTS, and under the assumption in this case of a passive coaxial network (i.e. N+0). We will discuss FDX Repeaters in a subsequent section.

Often lost under the umbrella of "Echo Cancellation" are the isolation requirements between an RPHY Transmitter and RPHY Receiver. This is less about reflected "echoes" per se, and more about very sound RF design practices that route signals to their proper destinations with maximum fidelity and minimum interference from other signals traversing the same circuit board, and careful, dB-by-dB subsystem designs inclusive of all tolerances and variations expected in operation.

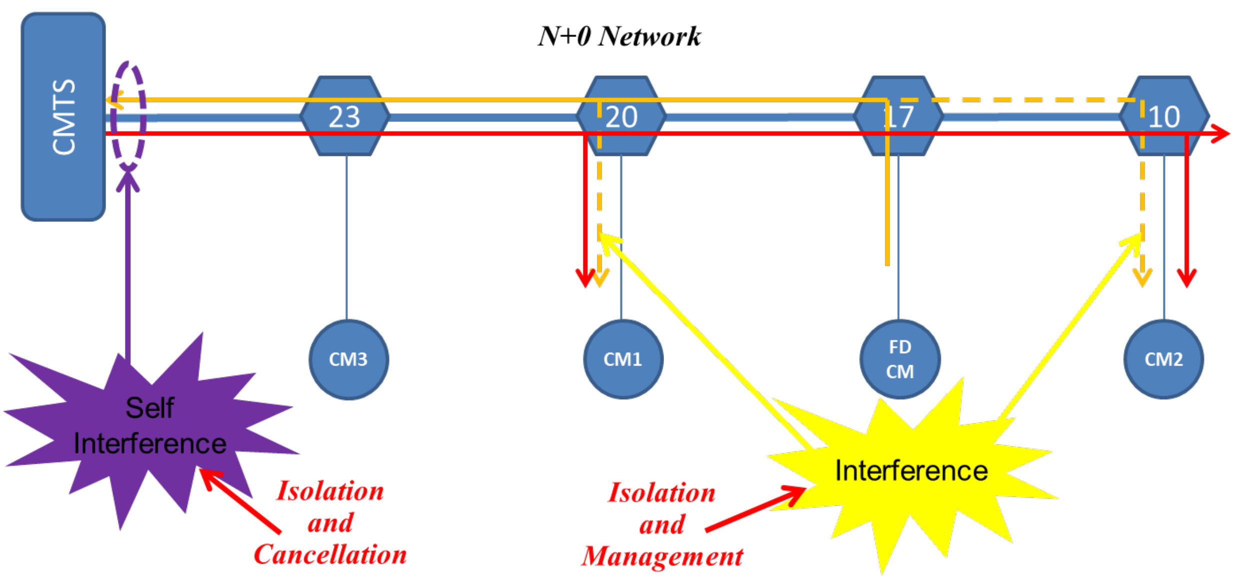

Figure 7 - No Longer an FDD System, FDX Requires Two New "Tricks"

To allocate new upstream traffic where downstream traffic already exists requires the upstream receiver to not be affected by the downstream signal that is sharing the band – in effect, it should not "see" the DS. To accomplish this, the downstream must be "subtracted" before the upstream receiver can process the upstream data. This is achieved through a combination of RF isolation of the higher transmit signal from the upstream receiver, with the ability to cancel any residual signal that cannot be isolated sufficiently. This element relies on the fact that the upstream receiver "knows" exactly the signal being transmitted – it's from the same chip – so can therefore subtract the artifact as it appears at the upstream receiver out. This basic function is represented as shown in **Figure 8**. (Note that **Figure 8** represents a non-DAA implementation.)

In practice, note that the function at the upstream receiver in the node can, and must, also measure the noise from the downstream transmitter (not shown in **Figure 8** explicitly), which can be as high or higher than the desired receive signal – enabling it to cancel out the noise. Note that transmit signals levels can be as much as 50-60 dB different, meaning the downstream transmitter must be isolated from the upstream receiver by this amount plus a further amount. This is to enable that there is sufficient MER to demodulate the data of a particular QAM format's fidelity requirement – no small task.

In addition to the signal path that leaks to the receiver and gets subtracted prior to the demodulator, the downstream transmitted signal, as it traverses the node and coaxial network, will see other imperfect (but within specification) RF interfaces. These discontinuities will create a reflection that shows up at the upstream receiver and must be cancelled. These are the so-called "Echoes" that give the EC function it name. As noted, "Echo Cancellation" has become loosely utilized lingo to describe the overall function of the isolation and cancellation required to clean the upstream receiver from the downstream signal that exists in the same band, and acts as interference to the upstream.

As noted, Echo Cancellation was developed and evolved in other industries – telco applications, as mentioned, but EC technologies are also an important part of wireless systems. What is unique to cable is the functionality required to satisfy a network that is very broadband, multi-octave, with high dynamic range. While Echo Cancellation has been around for a long time, with powerful algorithms that are both mature and well understood – these cable implementation requirements push the envelope.

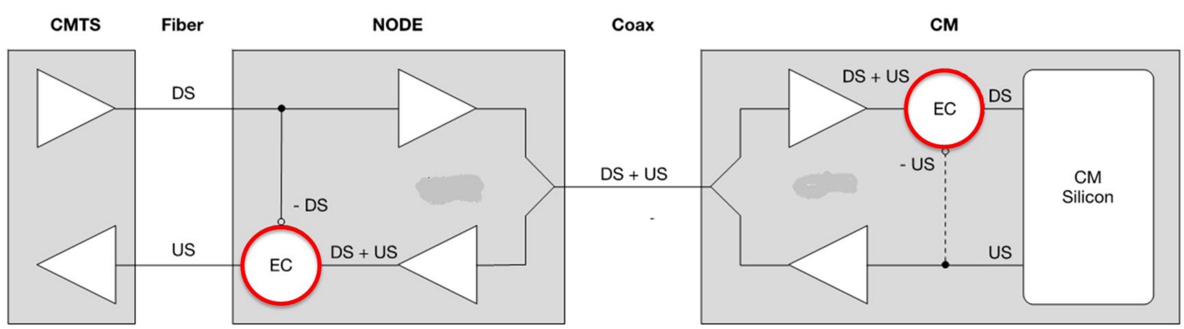

Figure 8 - Conceptual Basics of Echo Cancellation in a non-DAA Configuration

<u>And For Our Next Trick….</u>

As powerful as EC is, it relies on knowledge of the transmit signal and knowledge of the transmit noise (the former known a-priori, and the latter of which it can measure). Referring again to **Figure 8**, at the CMTS side, this poses no particular problem beyond the technical challenge of providing high enough dB rejection. However, on the home side, this can present another challenge. Recall that one of the distinguishing factors of the coaxial last mile, as opposed to the telco twisted pair last mile, is that it is a shared point-to-multipoint architecture. While a particular FDX modem can cancel its own signal and measure its own noise, as is done on the CMTS side, it cannot cancel its neighbor's transmission, because it does not know what the neighbor is transmitting. Therefore, for HFC, FDX must rely on more than just EC. It requires sufficient RF isolation neighbor-to-neighbor to prevent an FDX-enabled neighbor from interfering with another's downstream receive signal in the FDX band.

Unfortunately, it is known from today's tap specifications that sufficient isolation within a single tap and across taps cannot be sufficiently guaranteed to rely solely on RF isolation to protect against this interference. Furthermore, it is desirable not to have to embark on mass tap replacements to support FDX. It is also known that at least *some* taps, in particular commonly used high and some mid-value taps (29 dB, 26 dB, 23 dB), will have sufficient RF isolation between taps themselves. These relationships will all need to be determined as part of the FDX setup, or "sounding" process, whereby the dB isolation relationships of users on different ports of a coaxial network string are measured and sorted, for the purposes of dealing with this possible interference, and in order to effectively and efficiently operate an FDX system.

Rather than fully rely on RF isolation – although this is still a possible implementation option, with its own pros and cons – an alternative that makes use of the CMTS scheduler can be put into action. Fundamentally, the CMTS scheduler is a complex arbiter of user access to the network, allocating time slots for downstream and upstream transmissions to different users, to fairly and efficiently serve the aggregate demand. The CMTS does not need to pay particularly close attention to the various user's downstream and upstream access in time explicitly if it is fairly allocating bandwidth. That changes with FDX. With FDX, because a user on the same tap, for example, can transmit in the FDX band and create downstream RF interference in the FDX band for another user on the same tap, there must be an accounting for these relationships. During the "sounding" process, an FDX system determines these relationships – who can listen without interference while a particular FDX Transmitter is using the FDX Upstream.

These potentially co-interfering users are referred to – very cleverly, as engineers are apt to name things – as "Interference Groups," or IGs. A logical set of IGs is called a Transmission Group (TG), with the concept created because not every IG necessarily needs its own special treatment, from a scheduler perspective. It is largely dependent on the penetration of FDX users on a single node leg. This extra scheduler "step" must be added to make certain that the downstream receiver signal of a potentially vulnerable user is not overrun by an adjacent FDX upstream.

The above represents a simple description of the IG/TG concept. The FDX appendix to the D3.1 specification is written assuming that the potential future of FDX requires supporting a high penetration users. To enable fair and efficient access, these IGs and TGs can be granted sub-blocks of OFDMA spectrum, so that availability of a portion of FDX band is always possible for any IG or TG. The requirements have been written such that these sub-bands can be dynamically allocated as traffic demands shift with time.

Figure 9 shows and example whereby the entire 108-684 MHz FDX band is activated for FDX. Three TGs each have 192 MHz of downstream and upstream in the FDX band (roughly 2Gbps and 1.5Gbps of downstream and upstream capacity – highly dependent on architecture and conditions that set MER). This type of utilization often refers to the sub-band allocations as Resource Block Assignments, or RBAs.

Of course, the FDX band is not all of the *DOCSIS* spectrum likely to be available. For example, there could be another 192 MHz of non-FDX D3.1 spectrum in the downstream, for a total of 4Gbps. This is very MSO-dependent, as different operators have different near-term and long-term channel allocation plans. However, all MSOs expect to increase HSD bandwidth allocations in the coming years with whatever combination of D3.0 and D3.1 spectrum suits their particular needs. In addition, there is another 400 Mbps of upstream capacity available in the legacy mid-split band.

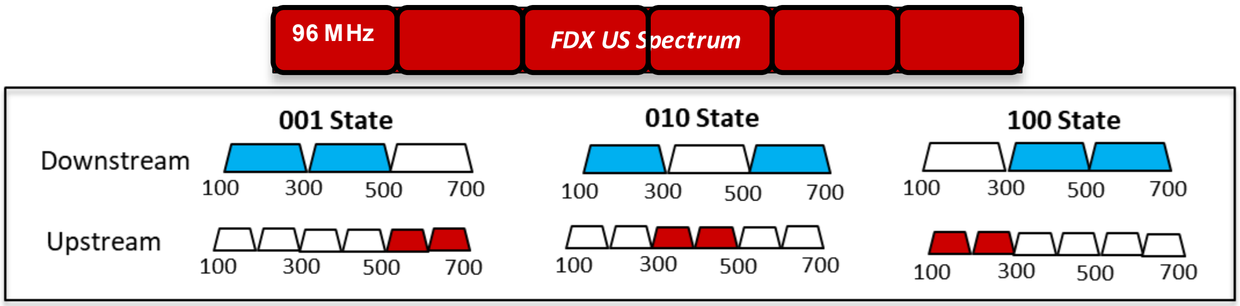

Figure 9 - OFDMA Blocks for FDX with IG/TG Relationships

Lastly, as described, the FDX IG/TG requirements have been written such that these sub-bands can be dynamically allocated as traffic demands shift with time. **Figure 10** represents this possibility, showing use of the FDX band allocation on the vertical axis versus time on the x-axis.

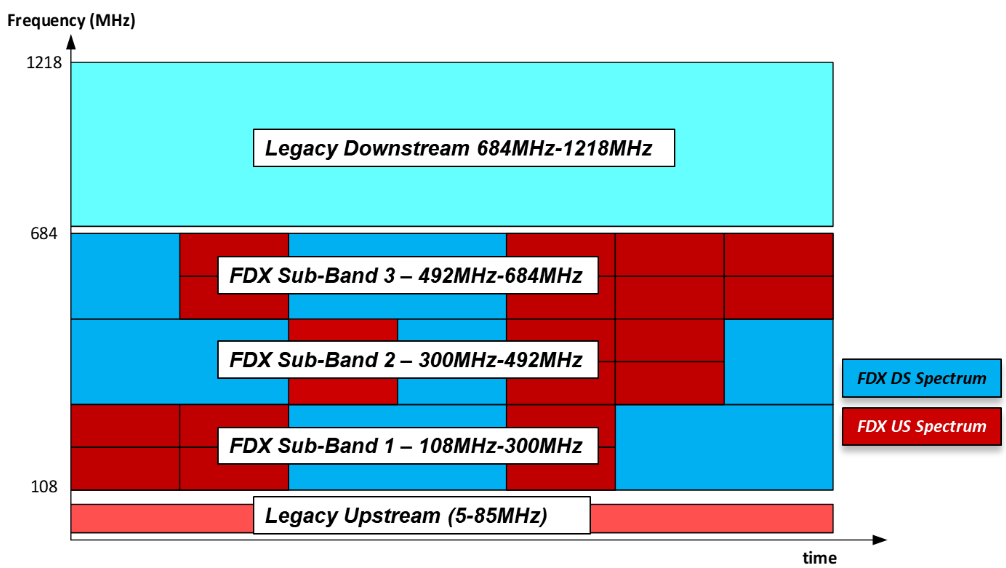

Figure 10 - Example of FDX Band Utilization vs Time Enabled by Dynamic Resource Block Assignments (RBAs)

Backward Compatibility and Coexistence

As with all generations of DOCSIS, backwards compatibility was an important part of the specification development. FDX is no different. However, there was recognition that there is a limit to what can reasonably be expected to be compatible, versus a lighter weight "coexistence" coexistence. Because FDX is fundamentally a DOCSIS 3.1 technology, the FDX band, such as the 108-684 MHz for the Mid-Split augment mode, can only be activated as FDX spectrum where downstream DOCSIS 3.1 spectrum is located. Furthermore, existing DOCSIS 3.1 devices must be made aware that they are participating in FDX, as this ensures they are effective participants in the IG/TG mechanism and will efficiently receive their data when downstream packets are sent their way.

Another way of understanding coexistence in the FDX context is to point out that the upstream FDX band will not be expected to operate in the same band as downstream legacy QAM carriers or DOCSIS 3.0 carriers. It is not a reasonable expectation that devices in the field of this vintage can be made aware of

FDX operation and participate in the protocol in that way. This is what is meant by a coexistence mode – use of FDX does not ban DOCSIS 3.0, lower DOCSIS versions, or legacy QAM video signals. That would not be a reasonable transition requirement for operators. But it also does not create too high of a burden for deploying FDX by making what would be impractical changes to these systems, undue burdens on the FDX system design, or unacceptable performance limitations of an FDX-enabled network.

Network Readiness for FDX

As has been shown over and over again with the evolution of cable technology, FDX does not rely on any extraordinary new HFC advances or access technology plans made of "unobtainium." Instead, FDX leverages emerging trends that operators have already deployed or are current emerging trends in the industry. Specifically, the following four items represent paths that simplify deployment of FDX. They do not necessarily represent must-haves, but they elegantly and efficiently fit into an FDX migration plan in way that is consistent with industry directions on many fronts.

1) Deep Fiber

   As previously described, one of the key technologies of FDX – the Echo Cancellation function – has been developed and optimized for xDSL systems that rely on point-to-point passive twisted pair links to a remote Digital Subscriber Line Access Multiplexer (DSLAM). The closest cable system parallel to this is a passive coaxial last mile, also known as N+0. This is the simplest implementation architecture for FDX. However, it is not a MUST requirement.

   Most if not all MSOs have plans to continue to take fiber deeper into their networks, and, in so doing, shorten amplifier cascades. In many cases, updating the spectrum split of these amplifiers is anticipated to manage upstream growth. It is feasible to have these newly updated and shortened RF cascades be FDX compatible in future generations of amplifiers, and in addition, through the use of EC, enable these amplifiers be software-configurable by split and/or FDX operation.

2) Distributed Access Architecture (DAA)

   Similar to Deep Fiber, DAA is a very desirable simplification and offers significant performance enhancements to an FDX system. Again, however, it is not a MUST requirement. However, when DAA is deployed, the upstream receiver, which for FDX can go as high as 684 MHz, is located at the node and preceded by the Echo Cancellation function, which may be a multi-stage EC. If the network were a traditional optical return path system (analog or digital return), new upstream transmitter development would be required to support this new bandwidth. This is not a complex technical hurdle – Downstream Transmitters, of course, extend to 1.2GHz. But it would also need to be augmented with complex new node development, since Echo Cancellation, as shown in **Figure 8**, would still need to occur at the node to suppress a high downstream RF transmission from the node's upstream transmitter input. Consider that the front end of an RPHY Upstream Receiver is an A/D converter that must be at least partially protected by EC. A digital return upstream transmitter also begins with an A/D converter and its performance is optimized by ensuring it is operating with maximum dynamic range allocated for the desired upstream signals.

   This EC function is best integrated with the sophisticated processing occurring in the receiver and integrated tightly with the surrounding circuitry in the node. With DAA via RPHY and

Remote MAC-PHY both gaining momentum, the expectation is that FDX will be implemented exclusively in DAA systems.

3) DOCSIS 3.1 Migration Plan

As mentioned, FDX upstream can only exist where there is D3.1 downstream. Therefore, in order to grow upstream speeds using the FDX band, there must be a commitment to a commensurate and increased allocation of spectrum in the downstream for D3.1. Since downstream speeds (which correlate with spectrum allocation) generally take precedence over upstream speeds (spectrum), this is not expected to be a significant limitation.

In other words, looking at a probable speed roadmap today, it seems unlikely that there will be a DS/US service tier that is "inverted", such as 1G/2G.

4) All-IP Services Plan

There are two aspects to this. First, an All-IP plan generally aligns with a growth plan for D3.1. This is because an important aspect of moving to all-IP is that there becomes a premium on using the available spectrum most efficiently, implying DOCSIS 3.1. Broadcast QAM will be on the network for some time, so any IP Video that is introduced is a simulcast and is technically redundant -- not unlike when operators simulcast standard definition/SD and high definition/HD versions of the same programming, and as far as spectrum utilization is concerned. While QAM broadcast limits the service and network flexibility, it is still spectrally efficient, as long as the video codecs are kept up-to-date and simulcasting the content in IP creates a spectrum cost which is an interim price to pay for making the all-IP transition.

Second, as we will discuss later, an FDX home is expected to become an All-IP home architecture in order to manage FDX deployment most cost-effectively and in alignment with the evolution of WAN and LAN connectivity in the home. This does not necessarily mean that the network itself is all-IP only, just that the specific user who requires FDX to meet his or her service demands will become All-IP to that home and within that home.

FDX Impacts to Network Components

This section describes how the new technology elements of FDX impact the architecture and components from end-to-end. **Figure 11** identifies the elements of the end-to-end network impacted by an FDX introduction, which we will discuss in the numerical order of the diagram.

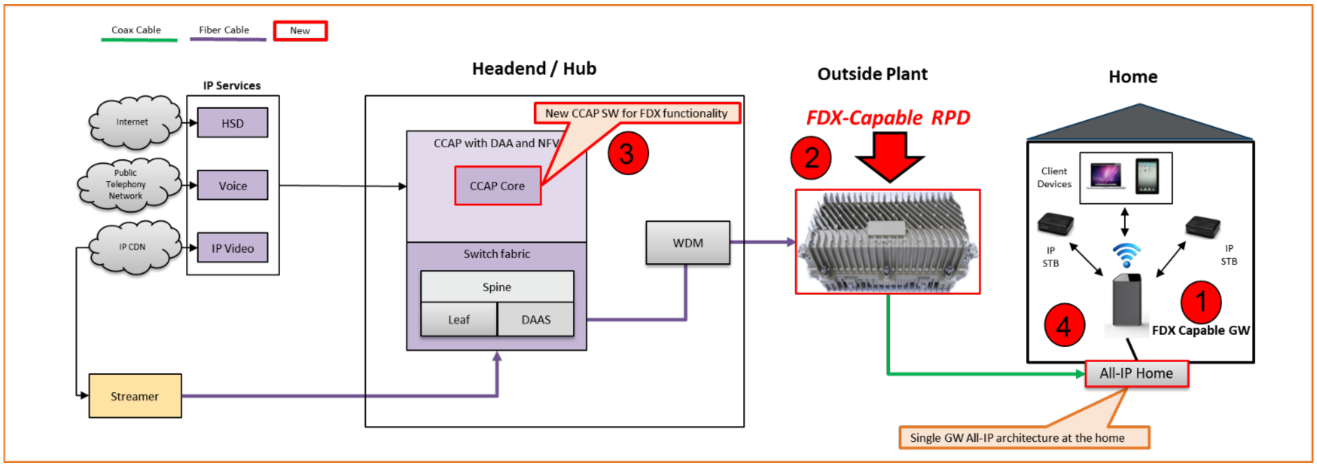

Figure 11 - FDX Impacts on End-to-End Network Components

<u>Consumer Premises Equipment (CPE)</u>

Perhaps the most straightforward item on this list of impacted subsystems are DOCSIS-based home gateways, built on System-on-a-Chip (SoC) ASICs that support the particular release of the DOCSIS standard for which they were designed. Because the Home Gateway tends to be in the home for many years, and the industry clearing house for DOCSIS standards development is a CableLabs consortium of operators, new standards are released in a pragmatic way, with agreed upon objectives by the service providers driving the objectives and requirements. The DOCSIS 3.1 standard was developed in 2013, and the FDX appendix to this standard was released in 2017.

As with DOCSIS 2.0, DOCSIS 3.0, and DOCSIS 3.1, support for DOCSIS 4.0 FDX requires new SoC development in order to implement the wider band upstream, the additional OFDMA blocks to fill it, and to enable the Echo Cancellation function. Support for the new scheduling aspects are also necessary, although this would not necessarily drive new hardware requirements, but would create differences in software. Indeed, existing DOCSIS 3.1 modems are expected to participate in FDX by accommodating FDX scheduling features in the FDX band related to IG/TG that enable FDX-capable gateways on the plant.

The RF design of the gateway around the SoC also becomes more intricate, as the first RF component behind the F-connector can no longer be the typical diplexer, separating the downstream and upstream. In the FDX band, they overlap, forcing a different approach based on splitters and couplers, in addition to diplexers as needed to manage legacy bands.

Because DOCSIS 3.1 enables a 204 MHz upstream and is capable of two 96 MHz OFDMA bands as a requirement, current DOCSIS 3.1 modems can be used with external components and software adaptions in early testing, in particular to evaluate development of critical technologies such as Echo Cancellation and IG/TG determination. This allows for early learnings around the core FDX-enabling technologies, de-risking deployments and providing an early look at potential optimizations.

<u>DAA Fiber Node</u>

The node in an FDX system must also support the expanded bandwidth, OFDMA blocks, and Echo Cancellation. Depending on the DAA "flavor," elements of the IG/TG determination and scheduler integration may also be supported. As with the home gateway, the internal RF design of the node is no longer an updated version of RF boards supporting downstream and upstream RF chains, independent of one another, with a particular FDD split separating them.

Also, and as with the gateway, the first RF component on the other side of the node's RF port cannot be a diplexer. It is one of the more intricate elements of the node design: To manage the downstream and upstream paths to support additional couplers and data converters needed for the EC function, while also supporting low path loss in the downstream for achieving desired output levels, and high dynamic range upstream receiver paths. These are not as much technology feats as they are an incorporation of RF best practices and optimizing design iterations.

An example block diagram is shown in **Figure 12.** In the figure, newly added items needed to support FDX are highlighed in red. The diagram represents a single node port or RF distribution and reive leg. While the same functions exist when there are multiple ports, system topologies and architecture have been observed to differ amongst node suppliers.

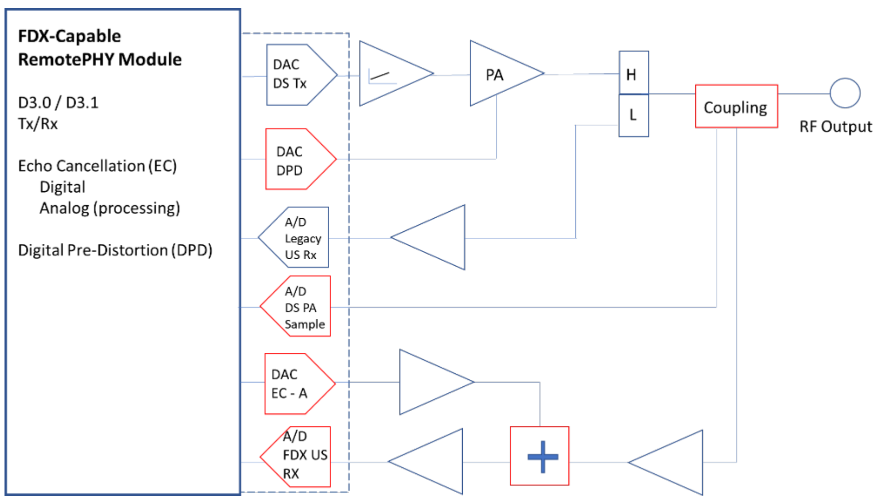

Figure 12 - Simplified FDX-Enabled Node Diagram (Single Leg)

It is possible that a DAA node deployed today in a traditional Frequency Domain Duplex (FDD) split can be an FDX node tomorrow. Such a node would become FDX-ready via a software upgrade and appropriately built-in hardware processing and RF capability. This would prevent future visits to a node, in order to enable FDX, by replacing modules and boards in the node. The complex aspects of this involve ensuring sufficient processing power, which is increased over a RemotePHY-implemented DAA node. The processing has to support up to six new OFDMA blocks as well as the EC function. It must also have the converters needed to execute EC, and it must have the hardware capability to switch in this functionality and bypass the diplex function when placed in FDX mode.

Lastly, of course, the software upgradeability must be in place, although this is likely established at the onset of the DAA deployment, in order to upgrade RPHY nodes with new SW, as necessary. It must achieve this flexibility within a power envelope consistent with the powering grid's limitations and upgrade appetite of a given an operator.

<u>CMTS Core</u>

As previously indicated in describing the new technical innovations needed for FDX, the CMTS gets involved in determination of the IG/TG relationships. From this determination, it uses that information as a basis for adapting the scheduler to enforce that transmitting and receiving users do not interfere with one another. It does this by scheduling them in appropriate time slots, which keep them apart, based on the IG/TG matrix.

Note that these IG/TG relationships are not static. New users can be added and subtracted, and devices may be moved around in a home or apartment, affecting the network characteristics. While these types of changes are not very dynamic, the system has to periodically re-execute "sounding" to establish state-of-state for the right IG/TG data to be used for the scheduler.

In addition, the network itself is also not static. Again, while it does not change dynamically, it can move slowly over time and temperature, and/or have maintenance done on it that effects the channel characteristics. It would not be expected that most of these affect dBs enough to alter IG/TG aspects, unless it is a repair to a major issue or, more commonly, adding on of a new subscriber. Note that the EC must be able to track these slow dynamics and adapt accordingly as well, and this is at the discretion and optimization of the system designers. The implementation must statically be sufficient to support the MER requirements that deliver a specified BER at particular levels and modulation orders defined in the specifications. However, there are not explicit transient behaviors or convergence time requirements called out in the specifications, so this is an area of potential optimization and differentiation.

Lastly, as shown in Figure 10 and repeated in **Figure 13** below, in its full-fledged implementation, at large penetration and scale, the scheduling function is expected to support RBA assignments and optimize RBA allocations that follow the needs of the traffic demands and service tier distribution among the IGs. Today's DOCSIS 3.1 schedulers are already quite sophisticated, and their capabilities are not yet fully leveraged. FDX will add the additional layer of the IG/TG no-go zones – at first, close to an additional "look-up table" to consider when allocating time slots – but which eventually will take on a more dynamic nature, as traffic patterns and expectations for use of the FDX band become better understood.

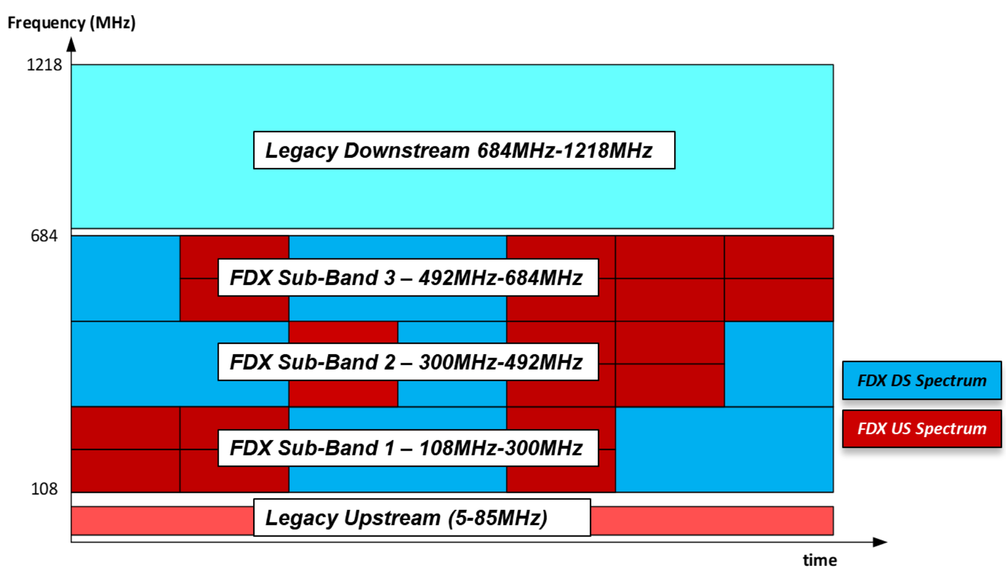

Figure 13 - FDX Dynamic RBA Scheduling

Home Architecture

Because an FDX home will transmit in what is historically the downstream band, legacy QAM STBs will see a very large amount of RF energy on their front ends. The upstream power in any 6 MHz band can be 40-50 dB higher than what would nominally be there in a downstream-only scenario. This band cannot be filtered away from QAM STBs easily, since the Upstream Transmitter must access the drop. Conceptually, homes with QAM STBs could have these boxes directly preceded by a filter, to protect their front ends from being overwhelmed with RF energy. However, channel allocations are not static, begetting filters that would ideally have different bandstop responses over time. Also, plumbing to support the DS and US out of band (OOB) would be a challenge, and home amplifiers – which do not exist for FDX are would be very complex to implement – supporting many of the deep-in-home STBs would be removed and cause level issues for these boxes.

Lastly, the FDX specifications are written, from a system engineering standpoint, assuming a point-of-entry terminating home gateway. Additional analysis in the working groups was done that considers a single gateway termination, but one that may be located deep in the home, at the end of higher RF loss. The advantage of this approach would be that a self-install kit (SIK) approach can be considered, which is a much more cost-effective way to deploy new devices in homes, and is highly preferred by customers than visits from field technicians.

These home install options are shown in **Figure 14**.

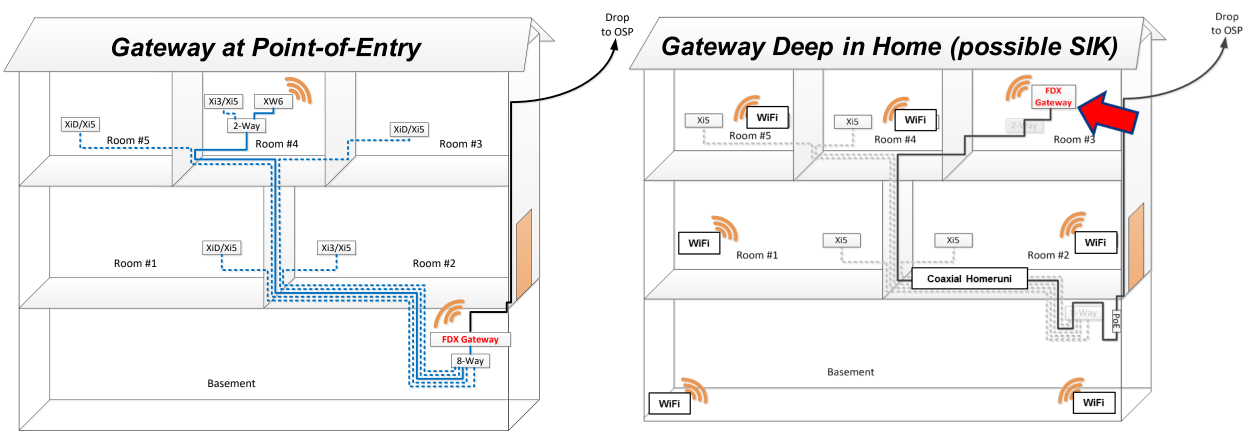

Figure 14 - FDX In-Home Architecture Options

With the above specifications and practical complexities in mind, it is expected that an FDX home will become an All-IP home. Customers who sign up for symmetric speed tiers that require FDX to deliver it will be converted to All-IP delivery at that time. An IP-transition initiative has already begun with several MSOs to take advantage of the many associated benefits that are unrelated to FDX.

<u>Other Network Components</u>

While not shown in **Figure 11**, as FDX development has progressed, enhancements and optimizations are being discovered along the way. Two areas receiving significant attention right now are in the areas of taps and FDX-enabled Amplifiers / Repeaters.

Tapping into Improved Specifications

Figure 15 shows what an IG determination might look like in a typical tap string of a passive coaxial network.

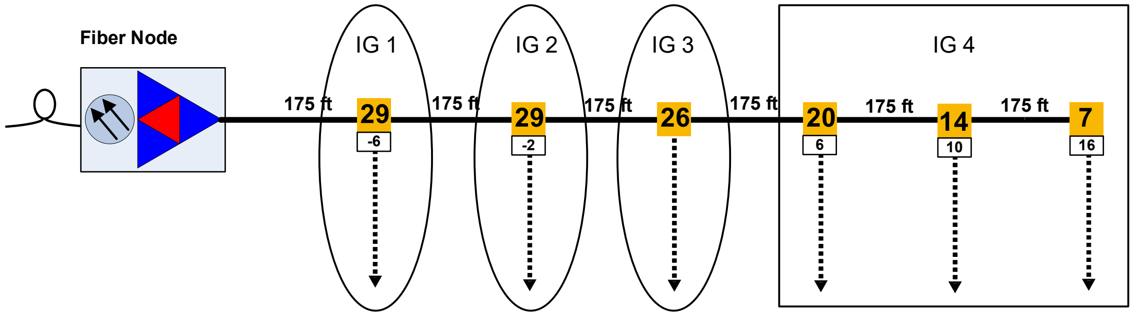

Figure 15 - Likely IG Determinations Using Typical Tap Specifications

An inherent trait of higher value taps is that they provide sufficient port-output isolation such that, across taps, each tap become its own rather small IG. This is an attractive scenario when considering a longer term future with potentially high FDX penetration. By contrast, lower value taps have less isolation between these ports, and thereby form a larger IG. This imbalance of IG sizes, while not concerning in early launches and likely for some period of time, could become an issue that impacts users in the large IG, in the future, because they are aggregated. This could place utilization constraints on users in this IG that are more burdensome than in the other IGs.

Based on this, one avenue of tap development that has been pursued as a potential improved FDX system of the future is that of of lower value taps with an improved isolation specification between output and drop ports**.** If the isolation could come from what is achieved in the higher value taps, then every tap could become its own interference group, forming balanced RF legs, from the perspective of FDX utilization and schedule management. Each MSO has its own view about what penetration of service tiers warrants concern about how granular an IG must be, in order to effectively support the offered tier. Furthermore, the economic justification for swapping taps – or tap faceplates, if this is an alternative – to enable this IG granularity would have to be weighed against rates of penetration growth, FDX band utilization statistics, and implementation of dynamic RBA features. Basically, it will be important to quantify how often access to the wire in the FDX band needs to be granted to these users for these emerging applications, and at what periodicity, since other subscribers in the IG cannot receive downstream D3.1 channels in the FDX band in that same time window.

A Repeat Performance

Now consider the case of an FDX Amplifier / Repeater in **Figure 16**. Multiple technology partners have put forth compelling ideas on how to implement an amplifier-enabling of FDX. The innovations already being developed for FDX can be applied and modified for the development of an FDX amplifier. And, of course, a major value proposition for the development of such an amplifier is so that FDX can be deployed in architectures other than N+0. This would broaden the scale of deployment of FDX and address operators that do not have specific plans or resources to migrate to a passive coax last mile.

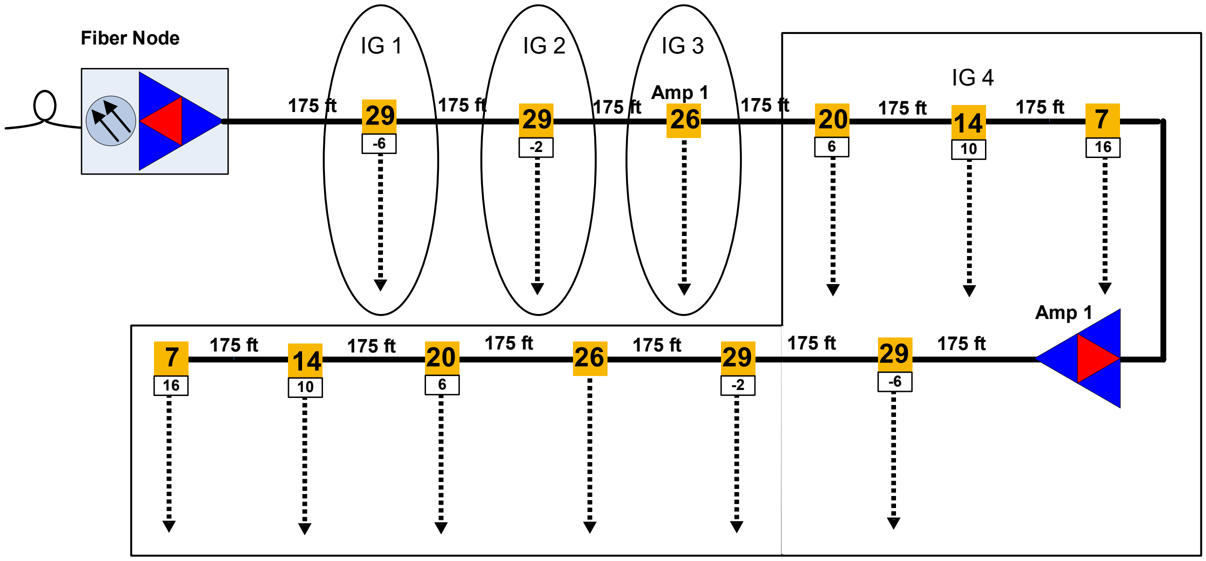

Figure 16 - FDX Repeaters Aggravate the Interference Group (IG) Situation

In Figure 16, however, rather than the tactical how-to of an FDX amplifier, we focus on the broader system implications of such a device to an FDX system. Again, using typical tap requirements, we observe that one of the major implications of the FDX Repeater is the expansion of IG4 in the diagram. Specifically, users downstream of the amplifier become part of the last IG of the tap string before the amplifier – an IG that is already larger than the preceding IGs in that leg. The limited and insufficient drop-output isolation characteristics which exist at the end of the first tap string are propagated into the subsequent tap string, downstream of the amplifier, along with the desired signal. This creates a much larger out-of-balance IG, and one that continues to get worse as amplifiers are added.

In Figure 16, taps that have better port-to-output isolation could resolve this issue. In the case of FDX Repeaters, to enable FDX in N+1 and beyond environments, it may be valuable to develop this tap performance or solve via alternative means such as a limit on minimum tap value. The latter, of course, increases the number of actives as a consequence.

As with the N+0 case, the need for small IGs may be some time off. However, with the use of repeaters, given their potential to significantly enlarge the IG, the time at which impacts become felt may be commensurately more near term. For example, for a given FDX penetration, more users per IG will exist on the larger IG. This ultimately limits the network efficiency and user availability of the D3.1 downstream channels and of the access to FDX upstream channels in the FDX band of that IG.

FDX Technology Road Shows

Inaugural Field Trial – Groton, Connecticut

Indoor and outdoor laboratories have been built and configured to test the boundaries of the early FDX node prototypes, and to evaluate whether the technology development is on the right path. Our controlled testing includes varying of network configurations, forcing negative test cases to observe the response to them, and finding the breakpoints using environmental laboratories normally reserved for quality assurance testing and survival of product. However, despite all of this ongoing data gathering, it is often impossible to truly emulate a difficult field environment. And virtually every service and product development effort complements successful laboratory test processes with an introduction to the field in trials, slowly evolving, to unearth remaining uncertainties. FDX is no different in how it is being approached, especially given that its technology innovations are based on DOCSIS – which is at the center of it all.

As such, Comcast, along with CableLabs and other vendor partners, conducted field testing for FDX in November 2018 in Groton, Connecticut. The field testing was conducted over multiple days, with the system under test being an N+0 last mile plant not yet connected to live customers. The primary objective of the activity was to characterize the network and assess the FDX node's ability to perform echo cancellation on the network response. In a related task, the tests observed the echo canceller's ability to handle dynamic changes, like those that might occur in normal operation due to weather impacts – such as wind and temperature variations, and some common maintenance activities on live plant.

Summarizing the field trial setup:

- Recently-built N+0 network with no customers yet connected. Access to drop cables and node distribution cable point for testing. This is the so-called "Model 1" as identified within the CableLabs specification development teams – single family suburban neighborhood. An example passive coaxial last mile is shown in **Figure 17**.
- Prototype proof-of-concept (non-production ready) FDX-capable nodes
- D3.1 modems capable of 2x96 MHz OFDMA upstream bands across a high split (5-204 MHz) upstream. External RF processing upconverted OFDMA blocks into selectable spectrum in the FDX band between 108-684 MHz.
- Array of network test equipment including a Vector Network Analyzer (VNA), Vector Signal Analyzer (VSA), and Vector Signal Generator (VSG)
- Three different setups, to execute:
    1. Test equipment characterization by Cablelabs
    2. Vendor 1 – Prototype node and modem
    3. Vendor 2 – Prototype node and simulated modem (VSG)

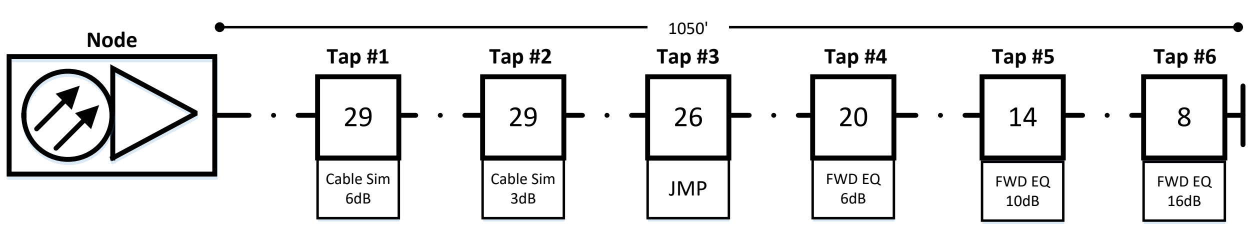

Figure 17 - "Model 1" Passive Coaxial Last Mile – Suburban Neighborhood

What was measured?

Comprehensive RF characterization was performed, with some measurements being replicated with different equipment and techniques to correlate results. This included typical network "S-parameter" RF characterization as shown in **Figure 18**, as well as measurements of micro-reflection response. This is the essential channel assessment, representative of the "sounding" process, and what the EC must be able to act on to enable FDX operations. Speed testing was also part of the sequence of measurements taken.

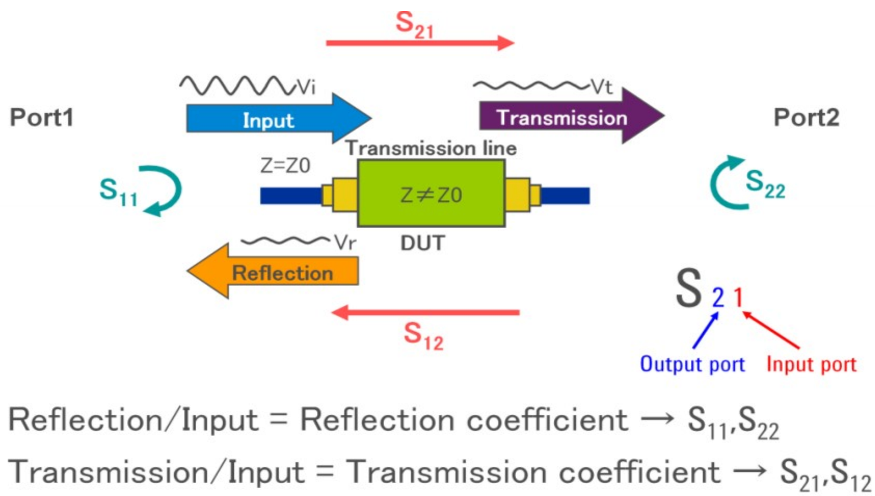

Figure 18 - S-Parameters

More specifically, the following measurements were taken:

- Plant frequency response: S_{11} at the node RF interface (Return Loss) and S_{21} (Insertion Loss) between the node and the upconverted D3.1 modem, under natural/static and dynamic conditions, in both time and frequency domains
- Static (slowly varying over time) and Dynamic (rapid variation over time) S11 measurements
- Performance evaluation of Echo Cancellers under natural (wind loading, temperature variation) and with maintenance conditions creating instantaneous discontinuities. **Figure 19** and **Figure 20** represent static and dynamic measurement examples. In **Figure 19** we can observe the reflection response (narrow spikes) associated with each return loss interface (tap connection) as seen from the node. In **Figure 20** we show the impact of a tap faceplate change on the reflection magnitude in the time domain. **Figure 20** was used to assess the ability of the EC to reconverge quickly from a poor transient RF environment to steady state conditions, following a common plant maintenance activity.
- Speed test, using iPerf, of downstream and upstream performance. Focus was on the upstream results, targeting Gigabit speeds concurrent with multi-Gigabit downstream speeds.
- Echo Canceller performance over a range of maintenance activities, including tap face plate removal and insertion, i-stop (ingress) tests, seizure screw loosening/tightening, and simulated inside-the-home changes (add/remove drops and terminators)
- Evaluate the above metrics across the entire FDX band from 108 to 684 MHz

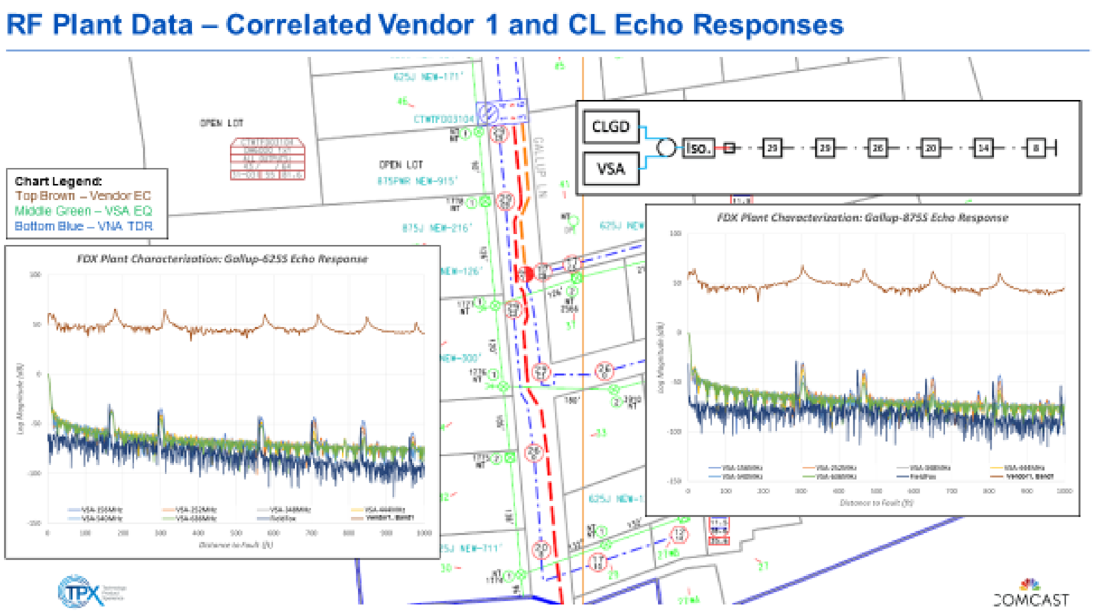

Figure 19 - Static Network Reflection Response

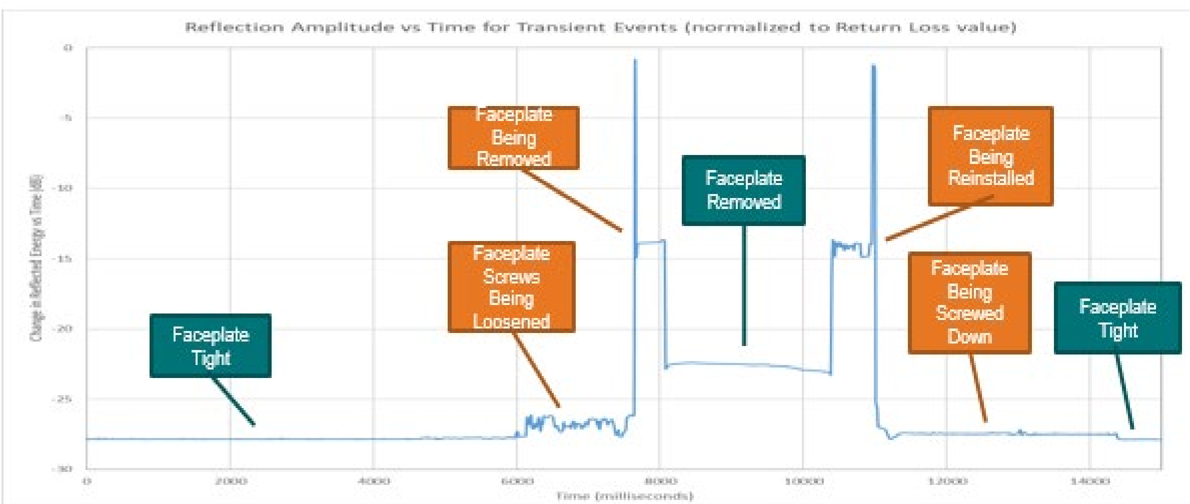

Figure 20 - Dynamic Reflection Effects of Plant Maintenance

First FDX Trial – Core Findings

i. Plant characteristics are matched well with predicted results

FDX system requirements were based on lab-characterization of representative N+0 networks, built in vendor-provided environmental labs, and using different tap models in each case. These setups were built as reference N+0 designs that covered the range of N+0 scenarios, from high density to low density. The CableLabs FDX "Models" include scenarios representing apartment complexes, suburban communities, and lower density outlying areas. This characterization was used for the development of the FDX system specification, because minimum values of tap return losses and maximum values of the insertion loss cascaded together, as a worst case design, would not be representative of the performance of a practical tap string, and would result in significant excess complexity, design margin, and cost.

Instead, these model systems were built and characterized over temperature to develop realistic yet practical and conservative values for the key parameters, from which system designers could develop FDX solutions.

Not surprisingly, the measured data in the trial plant was within the requirements set forth in the system specification that was originally derived from these reference models. This allows operators and solution partners to use simulated results with confidence.

ii. System performance was well-behaved during natural conditions of time and weather

Echo Cancellers from Vendor 1 and Vendor 2 were able to track wind gusts and temperature changes with no issues. Both represent slowly varying events, and in technical discussions with the solution partners, the time constants configured within the EC function that govern its rate of adaptability were much smaller than the rate of variation of plant metrics under natural, relatively static, plant operation.

Furthermore, and more important than the EC executing its role, the forward error correction (FEC) metrics showed no difference from steady state under these conditions. Thus, the customer experience would be unaffected, as desired, which is an essential outcome – and a promising one for a very first trial of this critical new DOCSIS feature.

Lastly, it was able to be shown that upstream speeds of 1 Gbps could be achieved while the downstream is transmitting in the same band. Again, for an initial trial, with first proof-of-concept technology, and no optimizations done for MER or OFDMA configuration, this is laudable performance from the solution providers.

iii. System performance recovered well during network-intrusive maintenance events

Multiple different types of "impulse" events that could occur during typical plant maintenance were performed to assess the stability and dynamic response of the EC function. These tests included performing the following maintenance-style activities:

- Face plate (remove/replace)
- Seizure screw check
- i-Stop (ingress) test

As **Figure 20** would lead one to believe, these events do not track error free through the plant trauma. Response magnitudes are instantaneously varying over 30 dB, for the case of the tap faceplate swap, for example. However, upon the completion of the maintenance activity, and with the steady state condition re-applied, the net effect was a very brief duration of roughly 2 seconds of elevated FEC errors. Not all FEC errors are necessarily noticeable to the customer, and this brief duration of FEC impacts is likely to be transparent to the majority of users, depending on applications being serviced.

Furthermore, since there is a full EC convergence and recovery, the brief hiccup is not likely to create additional activity into call centers. Brief network variations, for a range of other reasons and of longer duration than what is observed here, tied to many other network dependencies end-to-end, are commonly experienced by customers today.

Lastly, unterminated tap ports and simulated in-home activities had little impact on EC's ability to cancel. This is important because of relative ease of access to tap ports, frequent connect/disconnect cycles of field technicians, and the likelihood of "in-home" RF engineering that customers may undertake on to optimizie their home experience. And, regardless of what (assuming non-malicious intent) type of in-home engineering or wiring a customer chooses to do, they are unlikely to have an impact on the performance of the FDX system.

Second Field Trial – Denver, Colorado (week of 6/24)

A second field trial of Echo Cancellation performance was performed on a "Model 2" type network, an architecture that is meant to represent a multi-dwelling unit (MDU) – high densty apartments and condominiums. **Figure 21** shows and example design for such a scenario.

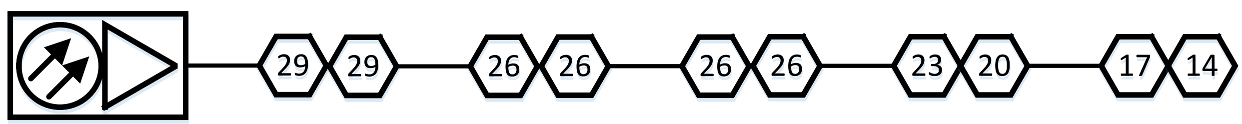

Figure 21 - "Model 2" Passive Coaxial Last Mile (MDU)

The Model 2 scneario, as mentioned, is one that is likely to provide a challenging environment for Echo Cancellation (EC) technology. These type of environments tend to degrade and become modified more rapidly over time for a variety of reasons – high turnover n the units, a single lock box, or perhaps several, that are continuously touched when changes are made and tend to be tightly packed, more limited access to the inside of the building, "hot tap" architectures (prone to use of more generic splitters off of tap ports) – to name a few. A network with a larger amount of higher reflections is anticipated in such a situation, and therefore the EC will be asked to work harder to cleanse the receive signals from the reflected transmitted counterparts that would otherwise interfere with performance.

The objective of the trial was to seak out a mature MDU site that has characteristics that you would expect from many years of operation in a cable infrastructure. A property was located in Denver that was planned for a re-wire as a maintenance clean-up as part of a contract renewal. The same team went to this site in advance of this rewire to characterize the network in this native form as a good example of "Model 2." **Figure 22** shows the lock box housing taps, splitters, and cables feeding the building, with the covers removed for test access. This example – a mature property with multiple years of customers additons and subtractions, upgrades, network adjuetments – is precisely the kind of scenario sought for this trial.

Figure 22 - MDU Lock Box for the "Model 2" Trial

Despite the concerns about taxing the EC, there was in fact no significant performance difference between Model 1 testing and Model 2 testing observed. **Figure 23** shows that, indeed, a noticable different reflection profile – in particular at the 400ft range on the figure – compared compared to the predictable discrete reflection observed in Figure 20 for the Model 1 scenario.

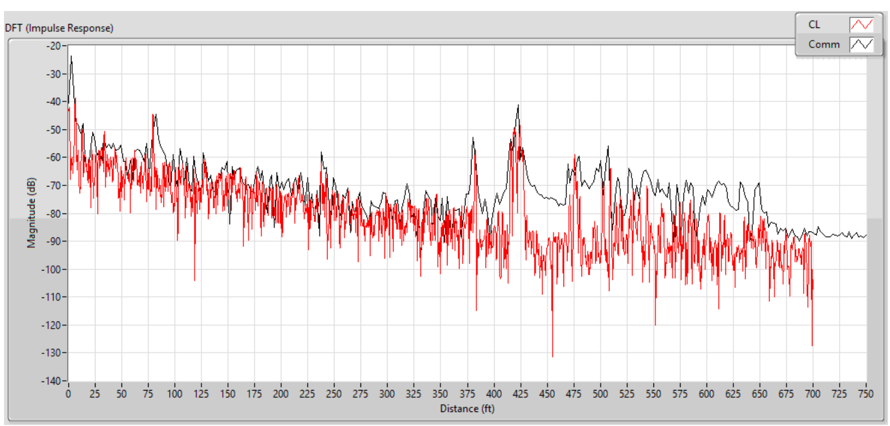

Figure 23 - Model 2 Trial: Network Impulse Response

Despite the concerns for the EC's ability to mange the environment, no performance differences were observed – meaning the EC design have sufficient margin in their designs to handle difficult cable environments. The same static and dynamic measurments described in the Model 1 trial were executed, and performance was the same or, in some cases, actually better than in the Model 1 case. These two trials provide significant de-risk value to the most critical technology introduced with DOCSIS 4.0 FDX.

Not Just Fast, But Instantly Fast

A 10G network is not only about higher speeds, but about supporting the best quality of experience (QoE) for services with various traffic and performance requirements tied to application needs. Higher access speeds via FDX, as well as faster Wi-Fi (i.e. 802.11ax) will enable higher throughput both in downstream and upstream directions, but, importantly, also deliver lower latency by improving congestion conditions. We describe some of the network and traffic considerations in this section, in recognition that it is more than just higher speeds that are necessary to provide better customer QoE.

Above and Beyond Speed: QoE Considerations

<u>Sharing Resources with "Elephant" Flows</u>

Different services and applications are aggregated in common DOCSIS service flows, sharing physical and logical resources. For example, if a subscriber tests the network latency under load, classic Transmission Contol Protocol (TCP) traffic will use these common resources aggressively, impacting latency and jitter tests.

A new set of "killer" applications will, as the history of the Internet has shown, based on an expectation of higher speeds:

- "Volumetric" 360° video
- 8K HDR video on the heels of 4K HDR
- Increasing video frames-per-second (fps) to 90+ fps
- Immersive virtual and augmented reality (VR/AR) experiences
- AR/VR with free-viewpoint or 6 DoF (6 Degrees of Freedom – 200Mbps to 5Gbps)

….and many others yet to be invented.

<u>Flows of Different Traffic Types having Different Performance Requirements</u>

Quality of experience for consumer-to-consumer, machine-to-consumer, consumer-to-machine, and machine-to-machine applications, and even applications within each category, depend on different performance metrics and values. For instance, most games do not require high throughput in the upstream direction but they require low jitter and packet loss. In First Person Shooter games, lag, as perceived by the user, affects the gamers' experience significantly, while some other games may be more impacted by packet loss. Lag perceived by the user is not only defined by the network latency, but also by lag compensation algorithms' performance against jitter, and by human reaction times. The ultimate gamer experience is when only the latter is a factor in the outcome of the game.

For game streaming with live chat, the amount of a viewer's read-ahead buffer vs. network latency variations define how the viewer's experience is effected. Latency may cause an unconfortable and even dizziness and nausea in a virtual reality application, with symptoms that are similar to motion sickness. The degree of "real-timeness," interactivity, and data volume among telemedicine applications (e.g. tele-education, tele-consultation, and tele-diagnosis) determine the tolerance to speed, latency, jitter and packet loss.

The key message is that speed alone cannot deliver all of the varying requirements for aggregated services.

<u>Varying Users' Perception</u>

Quality of experience assessment is a multi-disciplinary field encompassing engineering, individual and social psychology, demography, cognitive science and economics. The same application with the same quality of service may be perceived differently by various end users. Product choices vary among gamers, depending on whether they are hardcore gamers, enthusiasts, or causal gamers. Their knowledge and monitoring of network conditions vary as well. A hardcore gamer may prefer home devices and gateways designed specifically for gaming, procure a software add-on for routing optimization, and take advantage of other options available for gamers today. These solutions may be more important than increased speed for these customers.

A one-size-fits-all approach will not satisfy the expectations of subscribers with different objectives.

<u>End-to-End Performance and QoE</u>

Access network speeds increases such as provided by 10G will greatly change the HFC network capabilities, but the access last mile is but one part of the path between server and end user. It alone cannot be expected to improve the QoE for many services. Requirements for latency, jitter, packet loss, service availability, and device reliability must be met by the end-to-end architecture.

Low latency DOCSIS features target lower latency and improved consistency for the access network. However, home Wi-Fi improvements, edge services integration, routing optimizations north of the access network, and an SDN/NFV capable DAA, as an ensemble, will combine to yield the essential improvements necessary to meet the widening range of service requirments for high customer QoE.

For the operator, resource utilization and the efficiency in delivering these benefits to end users is a crucial metric.

DOCSIS Solutions for Lower Latency

Today, queueing delay, largely caused by bursty traffic (e.g. classic TCP traffic), and media acquisition delay in upstream (i.e. request-grant process) are the dominant factors in DOCSIS access networks' latency and jitter performance. TCP Acknowledgment ("ACK") prioritization and suppression in DOCSIS 2.0, along with new QoS features, buffer size control in DOCSIS 3.0, and Active Queue Management (AQM) and Light Sleep mode in DOCSIS 3.1, helped to improve average idle round trip time (~order of 10ms). However, round trip time (RTT) under load can still be 100s of milleseconds (ms) or more.

Most D3.0 modems have static default buffer sizes for different service flows. D3.1 modems support default buffer sizes computed for a given target latency – 250ms for pre-FDX DOCSIS and 50ms for FDX – and for a service flow's maximum sustained rate. These buffer sizes can be configured using configuration TLVs for both D3.0 and D3.1 modems. However, as many different services are aggregated within High Speed Data (HSD) service flow, it's impossible to define a static buffer size that's optimal for different traffic types, transport layer protocols with different congestion control algorithms, and varying network conditions that have different RTT values outside of access network. For example, home Wi-Fi and CMTS-to-server network conditions can affect the overall TCP traffic behavior.

Mandatory AQM in D3.1 cable modems and vendor specific downstream AQM in CMTS' aim to reduce buffer bloating issues by adjusting packet drops at a rate proportional to observed latency. AQM drives the queue delay to a target value (e.g. 10 ms default value). AQM provides more dynamic management and helps to improve classic TCP behavior by dropping packets when congestion starts building up. However non-queue building (NQB) traffic still shares the queue with queue-building (QB) TCP traffic, and different transport layer protocols, network parameters, and varying end-to-end network conditions can cause high jitter.

One of the main objectives of network QoS functions is to isolate and orchestrate overlayed logical network functions from the underlayed physical resources for different traffic types. From DiffServ Code Point (DSCP) differentiation to Virtual Private Network (VPN) overlays to network slicing, this same objective helps to accommodate very different traffic and performance requirements over shared physical network resources.

For example, when a classic queue-building TCP application and a non-queue building UDP application share the same queue in the access network, the UDP packet latency, jitter and packet loss performance will be affected by the queuing behavior of the TCP traffic. Applying buffer control to limit the buffer size, or applying the same AQM algorithm for very different traffic flow conditions, can improve performance of one flow while directly degrading the other flow. If isolation and orchestration of network functions can be supported through different queue management and forwarding schemes, services with different traffic and performance characteristics may be supported per their specific requirements.

New DOCSIS 3.1 Low Latency (LLD) features support isolation and orchestration of overlay logical network functions without reducing one traffic flow's latency at the expense of others flows' quality of service.

Low Latency DOCSIS and Its Impacts to Network Components

LLD requires classification of low latency NQB traffic (based on an ECN capable bit and/or DSCP) so that it can be separated from the QB traffic in the same service flow (**Figure 24** and **Figure 25**).

Non-queue-building traffic is traffic that underutilizes the link (e.g. low rate UDP traffic), or that can scale without building up a noticeable queue (e.g. Low Latency Low Loss Scalable (L4S) TCP flows). For instance, gaming traffic characteristics comply with a low latency traffic definition. Different buffer sizes and AQM algorithms are then applied to low latency and classic queues, while they are coupled to balance the congestion across the two queues. A queue protection scheme applied to the low latency queue detects QB flows that are falsely classified as low latency flows, and redirects them to classic queue. The aggregated service flow is rate-shaped with weighted inter-queue scheduler by the CMTS.

Other LLD features include MAP timing improvements. Early simulation and implementation of dual-queue approach show RTT less than 10ms for the 99th percentile of packets can be achieved for low latency traffic, without degrading QoE for other services.

A proactive grant scheduling (PGS) approach that eliminates media acquisition delay in upstream is also specified, with the objective of less than 1ms RTT for the 99th percentile of packets.

Today, there are already routers/gateways with adaptive QoS features, allowing end users to isolate gaming traffic from other in-home traffic. However, having the ability to isolate, per traffic, QB and NQB characteristics in a transparent way, and integrating this with the braoder access network resources, provides more effective overall traffic characterization, and also enables it to be applicable to the broader Internet ecosystem.

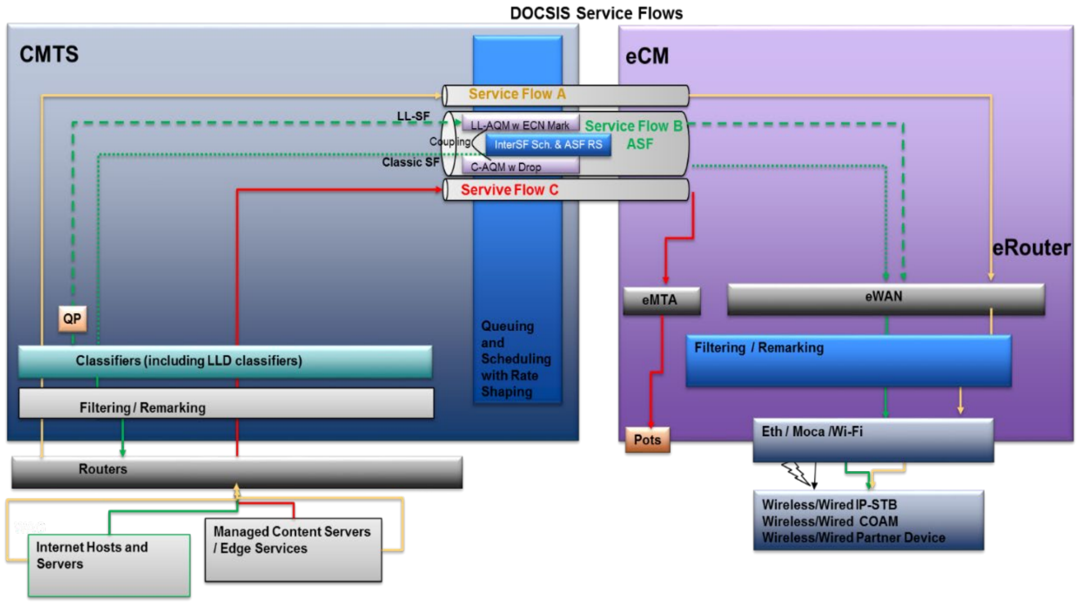

Figure 24 - DS Low Latency DOCSIS (LLD) Architecture

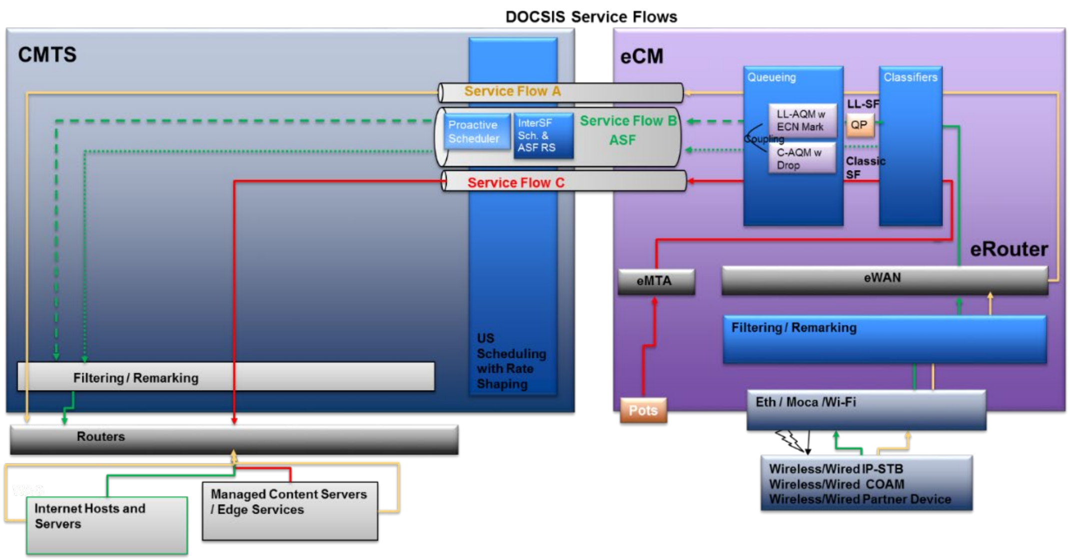

Figure 25 - US Low Latency DOCSIS (LLD) Architecture

Tieing Together Network Evolution and LLD

An end-to-end system (**Figure 26**) that optimizes networking performance per traffic characteristics and requirements may be extended by integrating the new LLD features with the FDX-enabled DOCSIS 4.0 architecture shwon in Figure 12.

<u>Consumer Premises Equipment (CPE)</u>

The following main LLD features are applicable to home gateways and modems:

- Config/registration TLV changes and LLD Classifiers
- Dual Queue with Coupled AQM and Queue Protection
- Latency Histogram Reporting (Calculation, MIB and TFTP file reporting)and QoS MIB updates

DAA Fiber Node

FDX-enabled RemotePHY will be a big step to deliver on the promise of 10G within a distributed access architecture. If DAA fiber node is Remote-PHY, there is no additional impact, except that perhaps some PHY parameters (e.g. interleaver parameters) will benefit from being adjusted for latency improvements.

CMTS Core

The following main LLD features are applicable to purpose built CMTS or vCMTS core:

- Aggregate QoS Profile and LLD classifiers merger
- ASF rate shaping and weighted interqueue scheduler for upstream and downstream
- Dual Queue with Coupled AQM and Queue Protection for downstream
- Latency Histogram Reporting (Calculation, MIB and TFTP file reporting)and QoS MIB updates
- Shorter MAP interval and turnaround time
- Upstream Proactive Grant Scheduler

Home Architecture

As discussed previously, MSOs can take advantage of the many associated benefits of an All-IP home to support current and emerging services and applications for a better QoE. An 802.11ax (Wi-Fi 6)-based home network provides improved airtime usage, and QoS and FDX-based DAA networks outfitted with LLD features can open the door for applications that are both high in bandwdith demand and latency-sensitive.

Embracing LLD in a Wider Ecosystem

An effective low latency customer experience requires and ecosystem of tight collaboration among service and application developers, home and access networking and device vendors, service providers, business and policy groups, standards bodies, and internet and open networking organizations. The end-to-end solutions must be transparent, and support different traffic and performance requirements without degrading any conforming traffic, and be applicable to a broader range of Internet ecosystem. However, there are still many areas requiring further study for performance and security assessments. There will also need to be changes in testing, operations and service assurance, as the priority on speed transitions to include the other key performance metrics more broadly than they are considered today, such as latency and jitter. The evolution to 10G will take this challenge head-on.

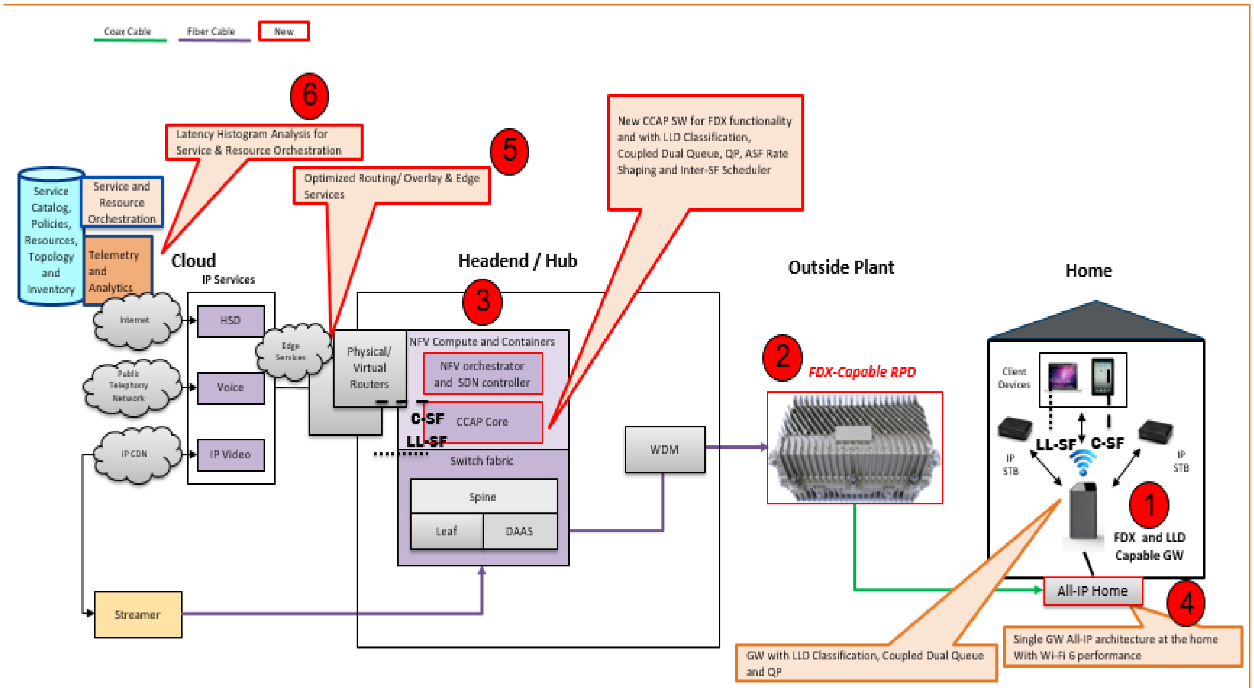

Figure 26 - DOCSIS 4.0 FDX End-to-End System

Et Tu, Extended Spectrum DOCSIS?

The 10G initiative is based on an understanding among operators of the capabilities of the HFC broadband network in the coming years, and the emergence of some of the key technologies that will power that future. Many were described earlier in this paper. However, "10G" is not explicitly prescriptive on the "how" of 10G, other than it will involve the evolution of plant and continue to rely on DOCSIS innovations. This is purposeful, because while all MSOs operate an HFC network, differences understandably exist among them, with respect to logical paths associated with 20+ years of technology iterations and rates of adoption, built on top of what was once dozens of cable operators, now largely consolidated.

Different tools in the toolbox, or a different sequence in the use of those tools, has inevitably occurred over time. A good example of this was the industrial divergence of paths to HD growth. This ramp saw differing priorities among operators who chose to deploy DTAs, SDV, and/or 1 GHz expansion. Different approaches or sequencing of technology phases is likely to occur also for 10G. While this paper focused on the development and progress of the FDX initiative, the industry is also hard at work looking at enablement of the full DOCSIS 3.1 specification to 1.794GHz (shorthand of 1.8GHz) as part of DOCSIS 4.0, and perhaps even higher.

Recall, the DOCSIS 3.1 specifications identify this extended spectrum to 1.8 GHz as an option. However, there has not been significant work done here, as the working group at that time did not prioritize the development of the system requirements for operating in this extended band. Consideration for beginning to develop the specifications has begun across operators and vendors at CableLabs under the logical name of "Extended Spectrum DOCSIS" or ESD (not to be confused with Electrostatic Discharge!).

An area where ESD goes beyond the DOCSIS "optional" 1.8 GHz Downstream is to consider an upstream beyond the 204 MHz "High Split". The definition of the FDX upstream band provides an obvious candidate for the grid to follow, enabling a natural synergy in the system designs of FDX and ESD that strengthen the ecosystem between operators who may prioritize FDX and those who may prioritize ESD. It is a new "FDD" Upstream that creates an ESD-based approach to 10G. This spectrum-oriented synergy can be represented as show in **Figure 27**.

As shown in **Figure 27**, a system design that supported both FDX functions and a 1.8 GHz downstream edge can have a common number of downstream OFDM and upstream OFDMA blocks. They are simply allocated differently, taking advantage of different configurations of features and non-DOCSIS processing within the multi-mode chip.

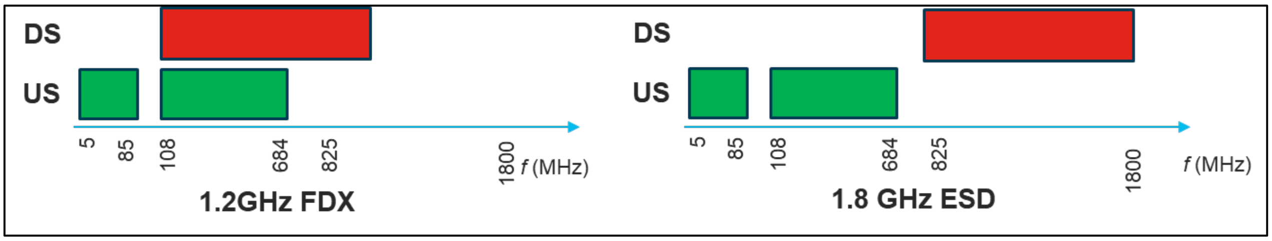

Figure 27 - Spectrum Synergies of FDX and ESD

An additional area of synergy possible is in the use of the EC function being developed for FDX. One of the significant negative aspects of an FDD system, when implemented using an ESD upstream, is the guard band between upstream and downstream – which is typically 20-25% of the upstream band edge spectrum. For a 684 MHz upstream band edge, this gets very large and wasteful. In a traditional FDD system, this translates to 100-150 MHz in practice, which is a significant inefficiency of a most valuable asset. Instead, the development of EC for FDX can be leveraged in ESD systems to significantly reduce the guard band through the use of the EC action to reject signal, acting as a more efficient alternative to the RF filter's ability to provide signal rejection.

Conclusion

The cable industry has rallied around 10G, as both operators and technology partners have gone all-in developing the requirements and critical innovations necessary to bring FDX to fruition. Multiple vendors have built FDX-capable DAA node proof-of-concepts, and several have graduated to be successfully proven in field trial environments. With the progress of FDX in the field, coupled with the DAA steps that align with more focused, deeper fiber strategies, have caused the innovation cycle of 10G to continue, as operators look for the best ways to use the new technology and align it to their network upgrade strategy. Examples include evaluation of new tap performance and technologies to do so, the development of FDX amplifier / repeaters and accompanying system engineering of such a network, and the evaluation of extended DOCSIS spectrum, such as to activate the 1.794GHz "optional" additional spectrum called out in the DOCSIS 3.1 requirements, which in the future will become part of DOCSIS 4.0.

These next generation steps are on the backs of the currently emerging set of core technologies that play a major enabling foundational role in 10G – Deeper Fiber, DAA, Virtualization, and All-IP migration. To varying degrees, and with time frames and pace aligned according to their individual starting points and demand projections, nearly all MSOs are building out their forward-looking capacity growth using some combination of these technologies.

Complementing the focus on capacity and speed, the industry has zeroed in on optimization of the network for latency-sensitive applications. This will include existing DOCSIS tools that have largely been underutilized; new development focused on pattern classification of service flow traffic, to determine more efficient and granular treatment flows; and predictive algorithm development to optimize the DOCSIS MAP process – all of which will make the network more responsive for all users.

10G – your time has nearly arrived! One more time, we can step back and marvel at the incredible ride that the adaptability of the HFC network has allowed us to take.

Abbreviations

AR	Augmented reality
BAU	Business as usual
CAGR	Compound annual growth rate
CDN	Content delivery network
COTS	Commercial off the shelf
CPE	Customer premises equipment
CMTS	Cable modem termination system
DAA	Distributed access architecture
DOCSIS	Data over cable service interface specification
DSL	Digital subscriber line
DSLAM	Digital subscriber line access multiplexor
DTA	Digital terminal adaptor
EC	Echo canceller
ES	Extended spectrum
FDD	Frequency domain duplex
FDX	Full duplex DOCSIS
FEC	Forward error correction
FTTH	Fiber to the home
HDR	High dynamic range
HFC	Hybrid fiber coax
HHP	Households passed
HSD	High speed data
IG	Interference group
IoT	Internet of things
LAN	Local area network
MER	Modulation error ratio
MTM/M2M	Machine to machine
NFV	Network function virtualization
OFDM	Orthogonal frequency division multiplexing
OFDMA	Orthogonal frequency division multiple access
OSP	Outside plant
PON	Passive optical network
QoE	Quality of experience
QAM	Quadrature amplitude modulation
RBS	Resource block assignments
RPHY	Remote physical (layer)
SoC	System on chip

SDV	Switched digital video
STB	Set-top box
TG	Transmission group
vCMTS	Virtual CMTS
VR	Virtual reality
WAN	Wide area network

References

[1] CM-SP-MULPIv3.1-I18-190422- DOCSIS 3.1 MULPI specifications

[2] *The Fiber Frontier* – Greg White, Karthik Sundaresan, Bob Briscoe Cablelabs, DOCSIS Research & Development, Feb 2019.

[3] *"The Future of Operations: Building a Data-Driven Strategy", by* Sebnem Ozer, Sinan Onder and Nagesh Nandiraju*,* SCTE Journal of Network Operations, Feb. 2019.

[4] *Lag Compensation for First-Person Shooter Games in Cloud Gaming* - Zhi Li, Hugh Melvin, Rasa Bruzgiene, Peter Pocta, Lea Skorin-Kapov, Andrej Zgan. https://www.researchgate.net/publication/325353977

[5] *Low Latency DOCSIS: Technology Overview* – Greg White, Karthik Sundaresan, Bob Briscoe Cablelabs, DOCSIS Research & Development, Feb 2019.

[6] *Predictions on the Evolution of Access Networks to the Year 2030 & Beyond"* – T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, S. Chari; The Cable Show NCTA/SCTE Technical Sessions, Spring 2014.

Acknowledgments

Comcast would like to thank its technology partners without whose support the FDX field trials would not have been possible. Teams from Commscope/Arris, CableLabs, and Maxlinear Corp all provided stellar support and invested significant time and energy into make these trials a success.

The Comcast Next Gen Access HQ team would also like to thank our generous and hospital colleagues in our Greater Boston and Mountain West Regions who were instrumental in helping us to find locations, manage logistics, manage customers, and make personnel available for these important trials.

