Roaring Into The '20s With 10G

A Technical Paper prepared for SCTE•ISBE by

Dr. Robert Howald
Fellow
Comcast Cable
1800 Arch St., Phila, PA 19103
215 286 8037
robert_howald@comcast.com

Robert Thompson, Comcast

Sebnem Ozer, Comcast

Daniel Rice, Comcast

Larry Wolcott, Comcast

Dr. Tom Cloonan, Dr. Ruth Cloonan, John Ulm, CommScope

Jan Ariesen, Technetix

SCTE•ISBE CABLE-TEC EXPO® 2020 VIRTUAL EXPERIENCE » OCTOBER 12-15

2020 Fall Technical Forum SCTE·ISBE • NCTA • CABLELABS®

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

It seems longer, but 10G unveiled less than 2 years ago! As exciting as the vision is, implementation began well before. 10G has a multi-element roadmap, with its foundations part of MSO initiatives for years. Some are in networks now, while others are in phases of development, proof-of-concept, and trial. Together, they paint a holistic picture of innovation leading to transformative digital experiences for customers.

In this paper, we explore key components of 10G, their synergies, and how they enable new experiences for residential and business consumers. We examine the payoff in speed, capacity, latency, advanced services, reliability and security – all while reducing carbon footprint. In particular, we touch on:

-The symmetric Gigabit foundation of DOCSIS® 3.1 technology

-Advanced DOCSIS 3.1 features – narrow subchannels, shorter cyclic prefixes (CPs), frequency interleaving, diagnostics, and Profile Management Application (PMA)

-Proactive Network Management (PNM) advances in data capture, analysis, localization, and tools for DOCSIS 4.0 technology

-Low Latency DOCSIS capability and implications to gaming and emerging requirements of IoT, M2M, and AR/VR

-Metrics and insights of fiber rich densification at scale, enabling a flexible last mile, increased reliability, lower Opex, and elimination of disruptive node splits

-Experience operating the largest Distributed Access Architecture (DAA) deployment based on remote-PHY (RPHY) and a High Availability (HA) Ethernet switch fabric

-A virtualized core leveraging real-time compute and rapid innovation cycles in the data center space, yielding significant advantages in facility consolidation and costs

-Fiber-rich symmetric networks to power growth in small-to-medium business and enterprise services

-DOCSIS 4.0 features, capabilities, progress, and steps leading to multi-Gig symmetric and full 10G

-The complementary nature of DOCSIS 4.0 Full Duplex (FDX) and Frequency Division Duplexing (formerly Extended Spectrum DOCSIS) FDD

Across Expo Tech Sessions, deeper dives illuminating these and other innovations are explored and described by Comcast subject matter experts focusing on this strategic program.

2. Scratching the Surface

It is hard to believe, but the specification development for DOCSIS 3.1 technology began 8 years ago, and service was launched, at Comcast, now over four years ago. Much has been written about the DOCSIS 3.1 foundation changing to OFDM/OFDMA, coupled with a range of higher order modulation formats, including up to 4096-QAM downstream and 1024-QAM upstream, supported by more powerful Low Density Parity Check (LDPC) Forward Error Correction (FEC), and 1 Gbps services.

There are additional features that received less initial fanfare but are receiving more attention now. Figure 1, Figure 2, and Figure 3 highlight these features.

In Figure 1, we compare the common North American upstream band edge of 42 MHz, to the 85 MHz enabled by DOCSIS 3.0 specifications, and to the DOCSIS 3.1 maximum, 204 MHz, often referred to as High Split. The choice of 204 MHz was selected to achieve 1 Gbps of upstream speed and capacity, recognizing that 1024-QAM @10 bps/Hz of raw throughput put that into reach. As such, with DOCSIS 3.1 technology, 1G/1G can be a product offering on the road to 10G.

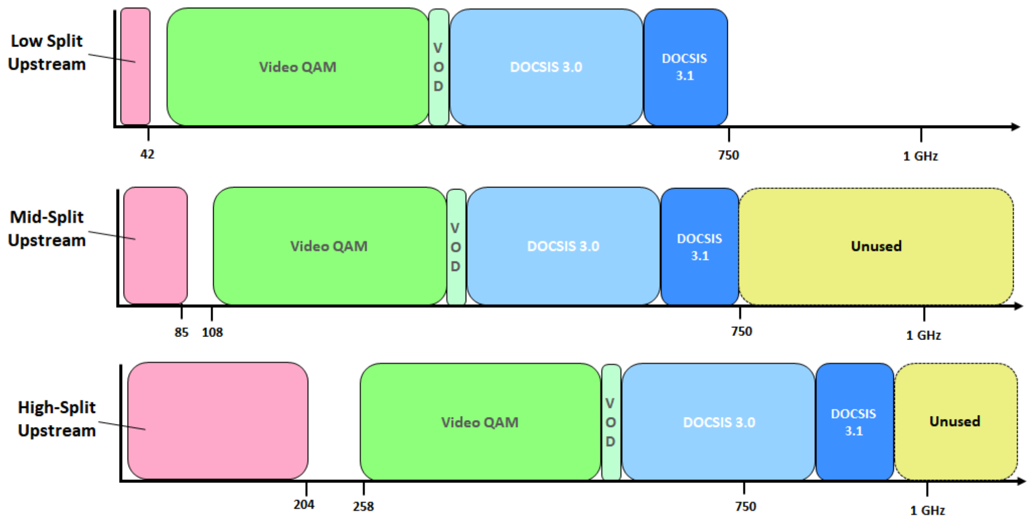

Figure 1 - Frequency Split Options Focused on Increasing Upstream Bandwidth

In Figure 2, the justification for defining Multiple Modulation Profiles (MMPs) is shown. DOCSIS 3.0 is limited to 256-QAM, and encumbered with a less effective FEC. With the combined effect of 1) better FEC (a lower SNR required for a given QAM format), 2) improved plant fidelity over time with deeper fiber and shorter amplifier cascades, and 3) migration to DAA, Cable Modems (CMs) are capable of much better bandwidth efficiency than 256-QAM. The dB attributable to the above can be used for capacity gain rather than just better performance margin of 256-QAM. Figure 2 shows why this is effective – many CMs that can support higher QAM efficiency than 256-QAM.

Lastly, note that it was determined during this study that this national SNR distribution also exists down to the service group (SG) level. Because of this, it makes sense for there to be different modulation profiles defined within a single SG, and these profiles cycled through based on their traffic demands on the time axis. This is shown in Figure 3, where sets of devices (color-coded) that exhibit common MER metrics have a time and a frequency slice of the downstream broadcast carrying their traffic.

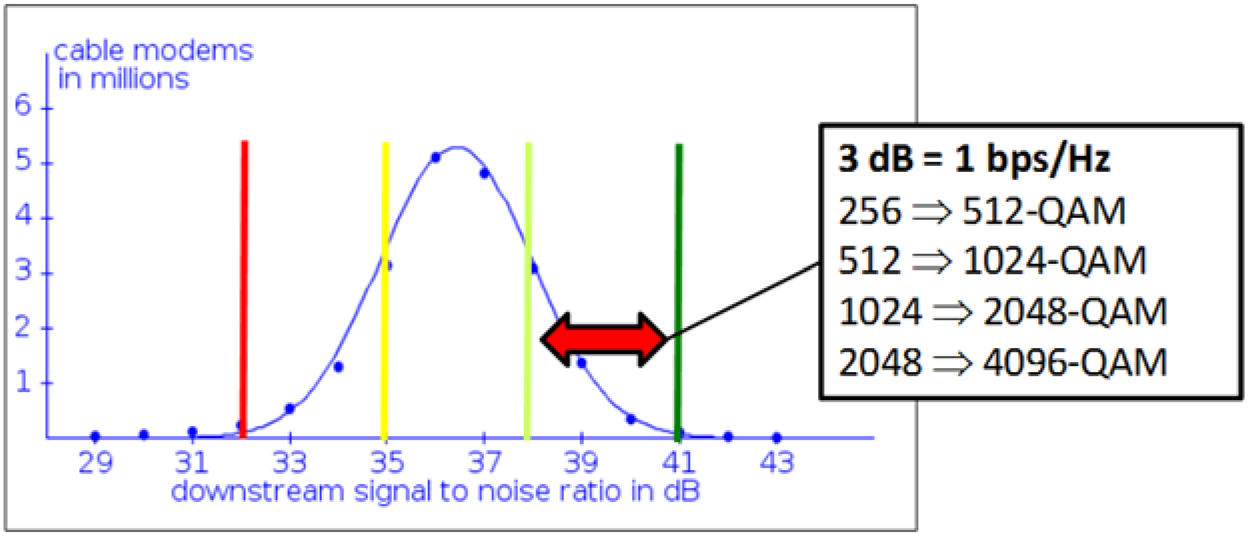

Figure 2 - Available SNRs Enable More Efficient Modulation Formats and More Capacity with DOCSIS 3.1 [8]

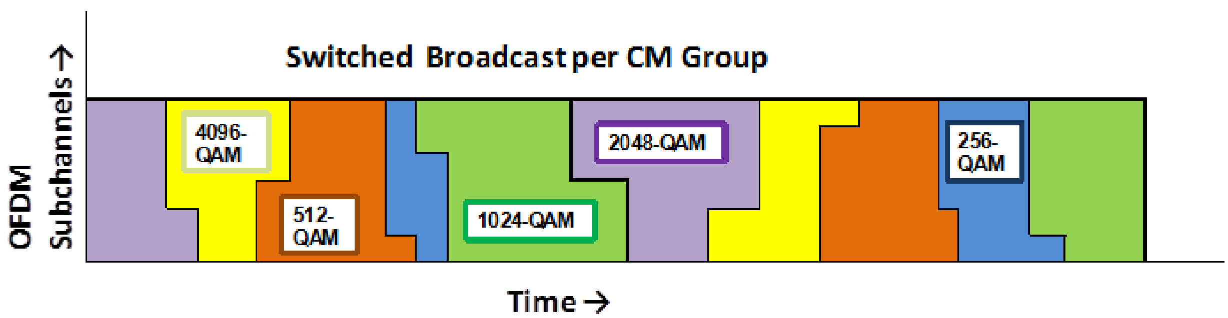

Figure 3 - Time Slicing of the QAM Profiles Associated with Grouping CMs of Like Metrics

2.1. Automated MMP: The Profile Management Application (PMA)

Efficiently managing the modulation profiles of DOCSIS 3.1 has proven to be a key component in successfully achieving the capacity promises of the technology. As DOCSIS 3.1 was deployed, it quickly became clear that dynamic and autonomous profile management is key to a holistic capacity plan [13]. Similar technology for DOCSIS 3.0 was recently deployed during the COVID pandemic to increase upstream capacity [14].

A PMA system leverages network data, such as RxMER per subcarrier in downstream OFDM channels, along with traffic, codeword, and signal level statistics to make optimization decisions that increase the capacity of the downstream, as shown in Figure 4.

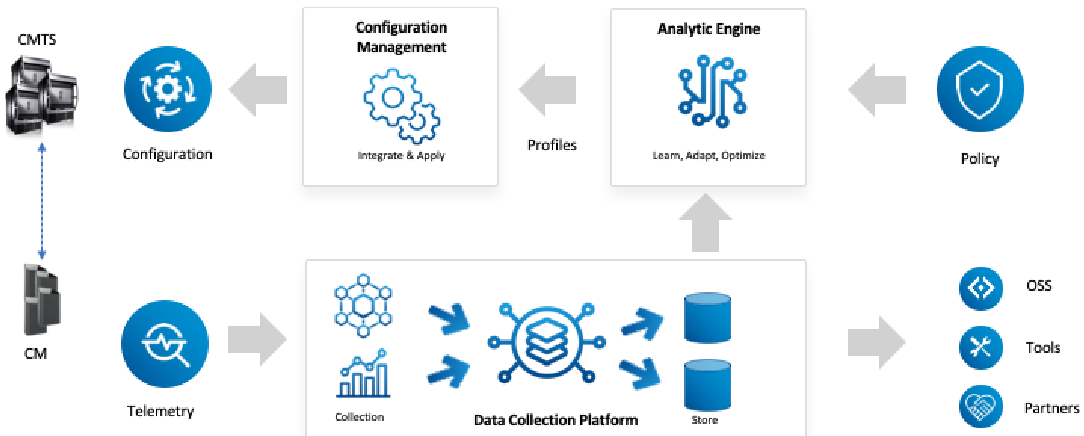

Figure 4 - Modulation Profile Management Platform

This same cloud platform has been used to optimize the capacity of DOCSIS 3.0 upstream channels. The optimization is based on three factors: upstream SNR, correctable and uncorrectable codeword errors, and signal levels. The metrics can be used by the analytics engine to adjust FEC and modulation for the data, Unsolicited Grant Services (UGS), and station maintenance Interval Usage Codes (IUCs). For DOCSIS 3.1 upstream transmissions, as in the downstream, optimization can be done by the analytics engine, which designs different modulation levels across the spectrum, tailored to channel fidelity.

This platform has enabled us to increase our access network capacity by over 36% in the downstream signal direction, and 20% in the upstream. The 36% is relative to a 256-QAM baseline and means PMA has brought the average efficiency close to 2048-QAM (37.5%). There is significant untapped gain in the upstream at this point, since 20% is comparing only different DOCSIS 3.0 settings. These results are tracked in operations as shown in Figure 5 and Figure 6.

Figure 5 - Downstream OFDM Profile Management Operational Dashboard

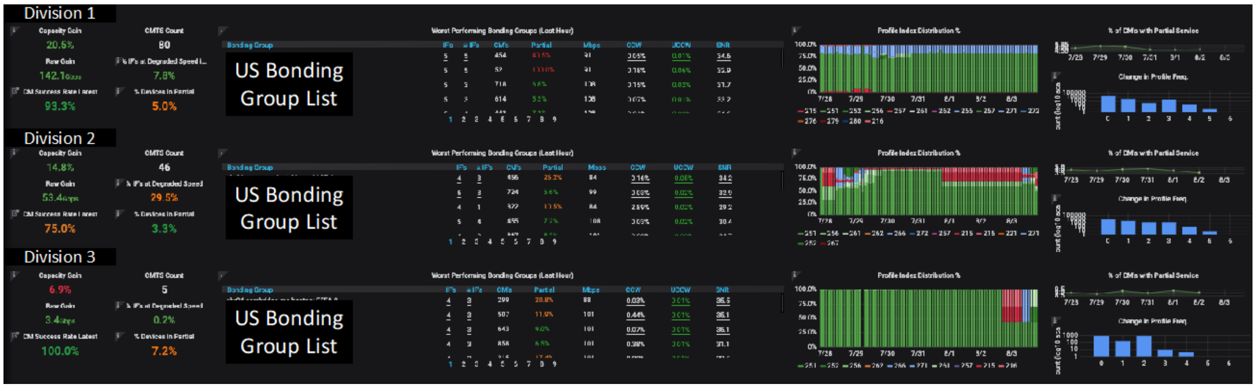

Figure 6 - Upstream D3.0 Modulation Profile Management

Beyond capacity gain for existing channels, this solution also enables spectrum to be used that was not previously possible. In some cases, spectrum is being used for OFDM downstream transmissions that are 100 MHz beyond common HFC frequency boundaries, driven by the pandemic-driven increases in traffic. In the upstream, a 6th DOCSIS 3.0 channel is being added below 15 MHz -- a notoriously noisy region that is suddenly more plausible when a PMA tool can be put to work to select the best FEC and modulation mix for a quality customer experience.

These PMA tools also help to ensure network robustness. For example, in the downstream spectrum, interference from mobile 4G Long Term Evolution (LTE) carriers is known to cause instability and loss of capacity without a system to effectively manage it. In the upstream, a variety of different noise sources [10] can be mitigated by a modulation profile management system. While these solutions mitigate the impact of the noise while maximizing the customer experience and capacity, the actual noise itself remains. It's still important to dispatch technicians to remediate the underlying issues. Combining profile management with PNM provides the best of both worlds. Technicians can be deployed to a specific network location that needs attention, and the customer is completely un-impacted and unaware through it all.

Similar techniques are being developed for upstream OFDMA channels, to optimize modulation profiles for challenging spectrum areas. The upstream signal path is a skinny, scrappy portion of total available capacity, routinely trampled by interference from FM radio, cascaded filter distortions, or the lower frequency noise common today, not to mention spurious/impulse noise. For DOCSIS 3.1, not only the modulation, but the pilot pattern, CP, and Roll-off Period (RP), along with subcarrier and mini-slot frame size are important attributes to optimize for robustness and capacity. These are important, for example, to achieving 1Gbps speeds in a high split configuration.

As product requirements grow beyond symmetrical Gbps to multi-Gbps symmetrical services, DOCSIS 4.0 Full Duplex (FDX) will be used. This technology will also benefit from OFDMA and OFDM modulation profile management. FDX introduces some new concepts that can also be managed as virtual network functions to optimize use of a shared spectrum resource, including Interference Group (IG) detection and Resource Block Assignment (RBA) usage.

2.2. Additional OFDM/OFDMA Specific Benefits

Two other features of DOCSIS 3.1 are worth mentioning. A key mechanism of how OFDM signals deal with linear distortion that results in Intersymbol Interference (ISI) is through a simple delay between OFDM transmission symbols. For frequency response distortion, it is also a simple mechanism – a single complex multiplier applied to each subcarrier. Compared to the complexity of tapped-delay line SC-QAM adaptive equalizers for linear distortion compensation, these are simplifications that come directly from the OFDM signal structure.

The time delay inserted is determined by understanding the delay spread of the channel, and the idea shown in Figure 7. The delay is known as the "guard interval" and is implemented by the addition of a CP of redundant information that is easily generated as part of the IFFT algorithm that creates the OFDM symbol. The CP is wasteful overhead, but necessary for clean transmission – the SC-QAM parallel being the spectrum shaping factor "alpha" that amounts to bandwidth overhead beyond the minimum Nyquist bandwidth to transmit a QAM symbol of a particular symbol rate.

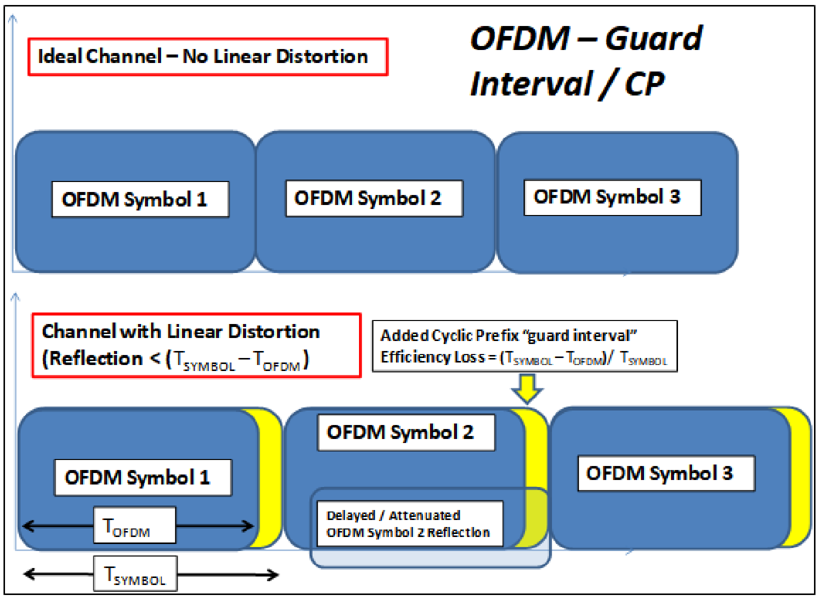

Figure 7 - The OFDM Cyclic Prefix (CP) Guards Against Channel Distortions and Can be Optimized to Channel Characteristics for Maximum Efficiency

Another DOCSIS 3.1 feature possible only because of OFDM is the frequency interleaver function shown in Figure 8. Each subcarrier is carrying a QAM symbol; a set of subcarriers form a codeword. However, when frequency domain interference strikes, only some of the OFDM subcarriers may be impacted. It is an OFDM advantage that subcarriers without interference continue to carry capacity, whereby an SC-QAM signal may be wiped out completely.

What the frequency interleaving attribute offers is the ability to spread the interference out such that it is not concentrated into a contiguous set of subcarriers or in a single codeword. By shuffling subcarriers prior to transmission, and then un-shuffling at the receiver, the tones with interference can be spread across the entire allocated OFDM band. They will then fall across multiple codewords, whereby each codeword will now have less interference than the single impacted codeword without interleaving, improving overall error correction performance.

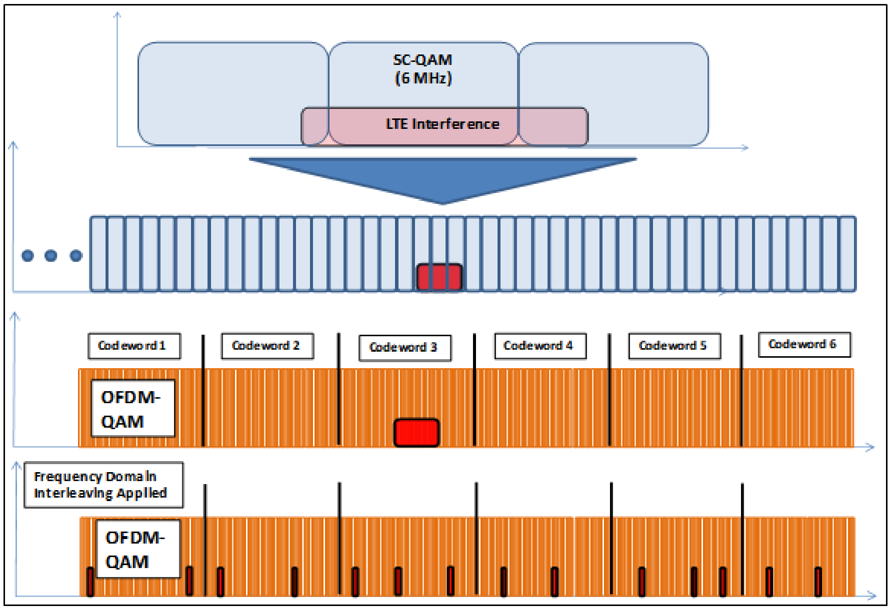

Figure 8 - OFDMs Narrow Subcarriers Make It Effective Against Wideband Interference with the Aid of Frequency Interleaving

3. Executing on the Fiber-Forward Future

Cable operators have been pulling fiber deeper into their network and splitting nodes for years, as one tool to manage the Compounded Annual Growth Rate (CAGR) of bandwidth usage and consequent capacity expansions. Node splitting reduces the number of customers in a service group by approximately two, doubling average capacity available to each home. Adding or reallocating channels of existing spectrum are other key tools to address CAGR. Spectrum flexibility, however, is very limited in today's upstream.

A "node split" can refer to a physical split, or the segmenting of an existing node by populating it with additional modules supported by new CMTS ports in the Headend. As nodes continue to be split and the serving group size shrinks, the efficiency of these splits decreases and the cost increases, because there are more physical splits rather than segmentation. A more forward-looking approach is to split multiple nodes at once. When many splits are projected over several years in an area, it may make sense to perform pre-emptive splits for labor efficiency and get ahead of trends as opposed to revisiting the same area repeatedly.

A version of "multiple node splits at once" is to be more methodical about the end architecture, as the node splitting occurs in these high growth areas. An approach we've taken in such areas replaces haphazard node splits with disciplined designs that optimize HFC design around a node with zero trailing RF amplifiers (N+0) end state, shown in Figure 9. Millions of homes passed are connected via this method today, and these networks will not have to endure disruptive node splits for many years, and possibly never.

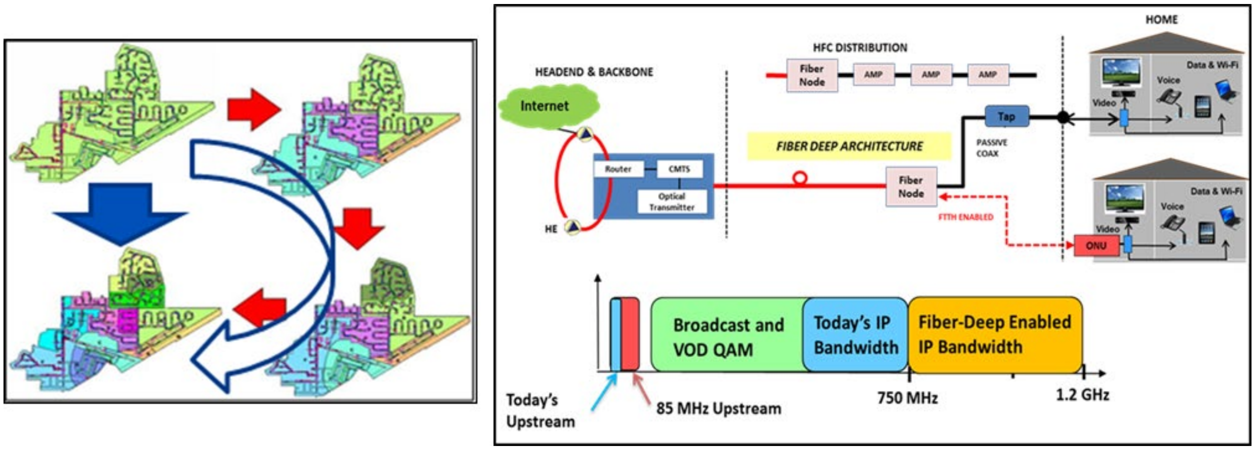

Figure 9 - Node Splits vs. N+0 Architecture

The N+0 architecture provides long-term capacity, improves End of Line fidelity, increases reliability, and adds substantial new spectrum currently limited by RF amplifiers in traditional HFC topologies. Operational performance is carefully monitored. Figure 10 is an example of comparative trends in Trouble Call (TC) performance between N+0 and standard HFC networks in a Comcast Division actively building N+0. Furthermore, the economics are favorable when compared to repeated node splits over time, especially in high growth areas of the network [7].

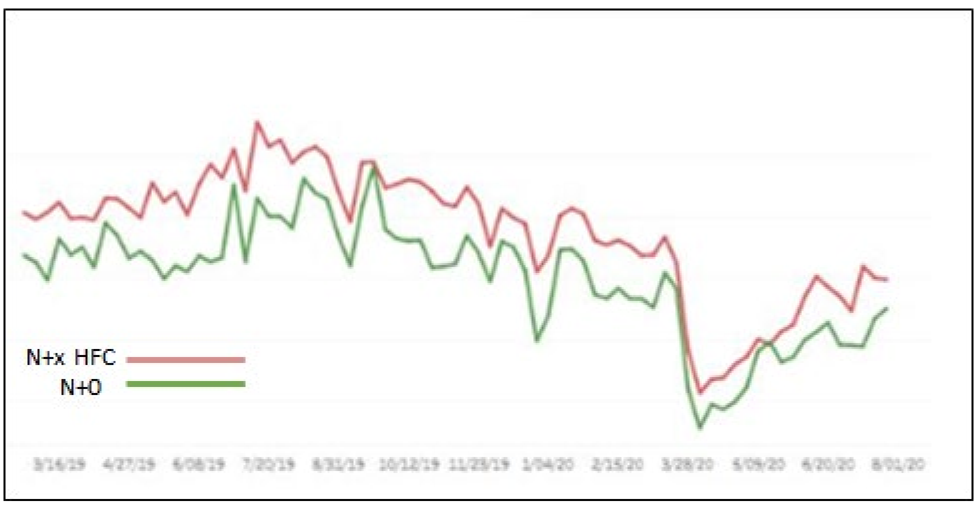

Figure 10 - Reduction of Trouble Calls in N+0 (green) v. Traditional HFC (red) Infrastructure

3.1. Remote PHY Based Distributed Access Architecture (DAA)

Traditional HFC networks have long used cable-specific analog optics to carry signals downstream; for the upstream, either similar analog optics or proprietary "digital return" systems. Our DAA deployments are based on Remote PHY device (RPD) technology and include all new N+0 systems because of the massive scale efficiencies, among other benefits, to the network and facilities. RPHY is by far the widest DAA flavor deployed. A brief summary of DAA benefits:

- Improved reach and wavelength efficiency of digital fiber versus analog
- Fidelity gains (MER) due to the removal of analog optics
- Physical scalability in the inside plant with the removal of RF cabling and combining networks
- Operational simplification, network availability, cost reduction
- Introduction of the global Ethernet ecosystem into the optical access network
- The vision of a network agnostic (at last!) to last-mile access technology
- Path to CMTS virtualization (vCMTS)

Regarding the last bullet, with DAA, a major part of the DOCSIS-specific CMTS functionality is moved into the plant. The subset that remains – packet processing, switching, storage, scheduling -- can be revisited from the perspective of today's compute power and the capabilities of real-time software. Commercial off-the-shelf (COTS) servers and switches can be combined with software to implement the CMTS function virtually. Our first DAA deployment, hosted by a virtual CMTS, or vCMTS, was introduced in 2018.

Figure 11 compares an N+0 network based on traditional AM optics and I-CCAPs to a DAA-based vCMTS solution. In the diagram, DAAS stands for Distributed Access Architecture Switch, and H-AGG stands for Hub Aggregation (also switching / concentration). Note the shifting of the CMTS core location to be further northbound in the vCMTS architecture, owing to the capabilities of Ethernet optics in terms of reach and wavelengths. This leads to opportunities for facilities reductions and consolidations in a massive way.

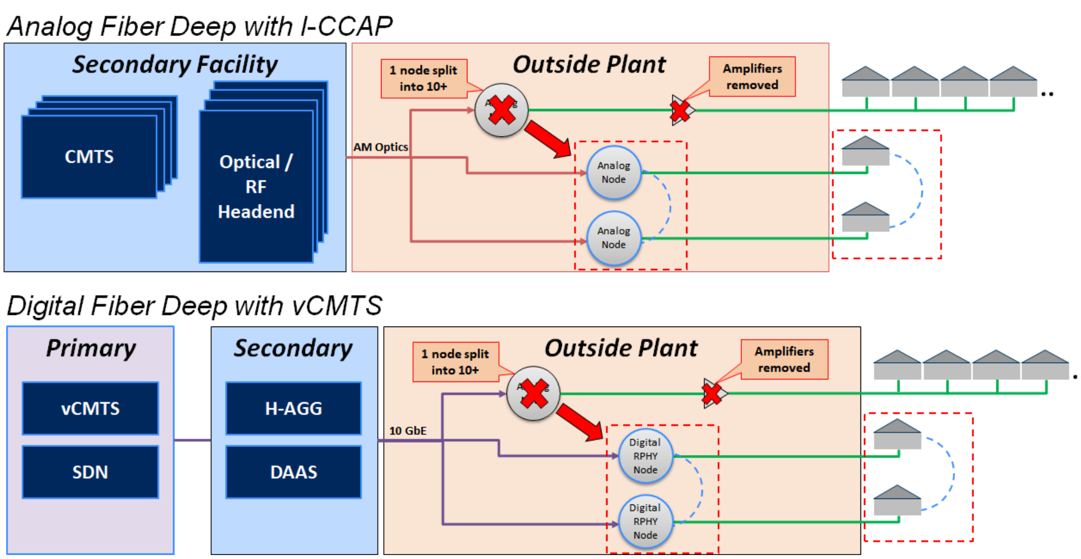

Figure 11 - DAA Aligns Cable Systems to Networking Technology that Leverages Wide Scale Data Center Solutions and Protocols

4. 10G: DOCSIS 4.0 Full Duplex Progress Update

DOCSIS 4.0 Full Duplex (FDX) enables multi-gigabit symmetrical capability over a coaxial last mile, a massive increase of upstream capacity and speeds. With the nature of consumer traffic being highly asymmetrical, cable operators optimized broadband for this reality, but are now getting ahead of requirements for new services and emerging applications that tax the upstream path – high definition (HD) home security cameras, cloud gaming, IoT, machine-to-machine applications, and those we cannot yet imagine. What FDX achieves is significantly more upstream bandwidth without a downstream bandwidth penalty, delivering new speeds while staying on top of CAGR of data traffic. Another key 10G pillar is low latency, which is also a key requirement of future applications, including many of those described above.

4.1. Key Innovations of FDX

Figure 12 shows the fundamental concept of FDX – access to massively more upstream. The FDX upstream is based on the same DOCSIS 3.1 physical layer – 96 MHz OFDMA blocks of 25 kHz or 50 kHz subcarriers, the same QAM profiles, and the same FEC. Six possible 96 MHz OFDMA blocks are added in the 108-684 MHz band as a complement to an 85 MHz "mid-split" system.

Figure 12 also makes clear the fundamental challenge of FDX – the new upstream overlaps with the downstream. This differs from the basic and, until now, exclusive Frequency Division Duplex (FDD) principle of cable TV. This dates back to the 1970s, when downstream and upstream signals began to exist in in separate regions, as in Figure 1. Coaxial cable is an inherently broadband medium, necessary originally to handle a multiplex of analog video channels, which deferred the need to be especially efficient with its use.

Meanwhile, phone companies had to work with much poorer physical medium – twisted pair copper wires -- to enable data services on a system optimized for voice bandwidth. Later, wireless networks had to manage within regulated spectrum. In both cases, spectrum efficiency was at a premium from the outset. Consequently, powerful technologies, including a technique known as Echo Cancellation (EC), were developed and matured to use available bandwidth in full-duplex mode. Years later, cable systems are now poised to leverage multi-octave, broadband EC as one of two key innovations of FDX, as efficient use of cable spectrum becomes increasingly important.

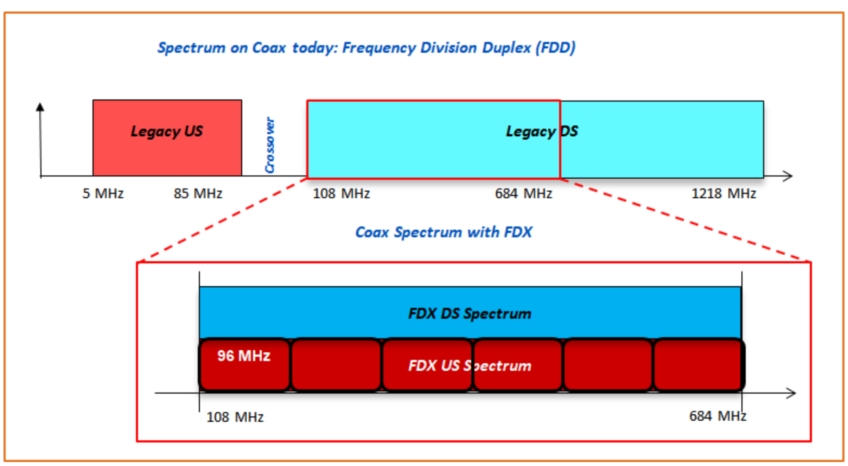

Figure 12 - Spectrum Fundamentals of DOCSIS 4.0 Full Duplex (FDX)

4.1.1. Learn to Share

The unique aspect of FDX is that the 108-684 MHz upstream can also be occupied by DOCSIS 3.1 downstream traffic. Echo Cancellation allow the simultaneous use of spectrum by removing the undesired signal from the desired signal at the receiver when they are sharing common spectrum. The term “Echo Cancellation” captures that: from a DAA node receiver perspective, it operates on undesired reflections of the downstream that bounce back from the plant. In fact, some of strongest interfering energy originates within the node itself, as either limited RF isolation or an internal reflection. Regardless of origin, it is handled by the same EC mechanism.

From the FDX RPD perspective, the EC function is conceptually straightforward – the downstream must be “subtracted” before the upstream receiver can process the upstream data. This is achieved through RF isolation of the downstream signal from the upstream receiver, and the ability to cancel residual copies of the downstream that make it to the receiver. It relies on the fact that the receiver “knows” the signal being transmitted and can therefore theoretically subtract the interference, as shown in Figure 13.

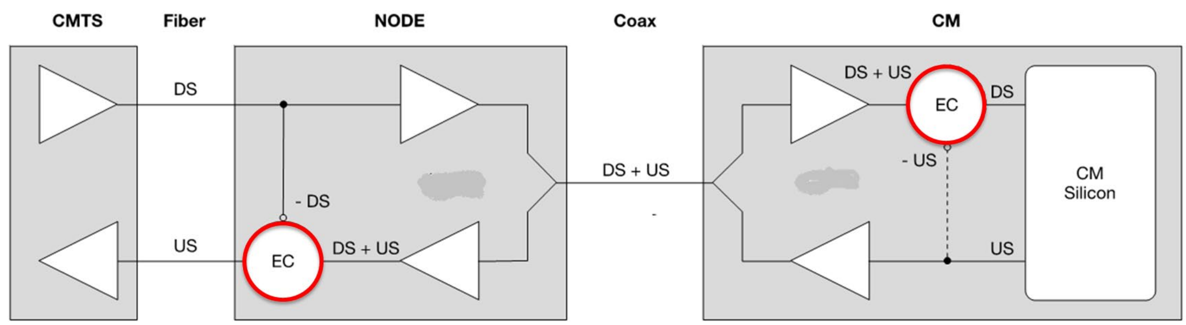

Figure 13 - Echo Cancellation Removes Undesired Transmit Energy at the Receiver

In practice, the upstream receiver must also measure the noise from the downstream transmitter and cancel it out for highest fidelity and maximum upstream bandwidth efficiency.

One important architectural difference of HFC compared to the telco last mile is that HFC is point-to-multipoint. This difference means one additional innovation is needed for FDX. That innovation involves clustering the users into groups that could inadvertently interfere with one another if any in the group were listening to the downstream while another in the group was transmitting. This is an extra scheduling exercise for the CMTS, which first determines the clustering, and then as part of its normal job of assigning transmit time slots to modems checks the upstream bandwidth allocation map (MAP) against downstream destinations, and if necessary, adjusts to avoid possible collisions.

These clusters, an important element of FDX technology definition and management, are called Interference Groups (IGs) and aggregates of IGs are called Transmission Groups (TGs). Figure 14 represents the two new techniques required to make FDX work, from the perspective of the FDX RPD.

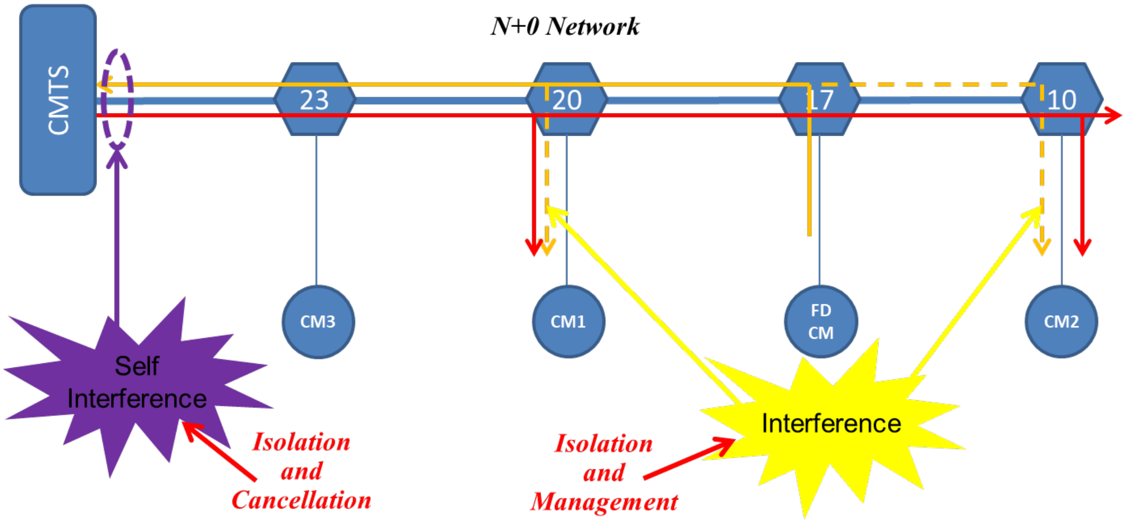

Figure 14 - Two Key Innovations for FDX: Echo Cancellation and Interference Group Determination and Scheduling

4.1.2. When You Can't Share, Be Fair

Echo Cancellation is based on having or gaining knowledge of the transmit signal and noise to be cancelled. This can be a challenge in a point-to-multipoint architecture, and from the CPE's perspective. While an FDX CM has knowledge of its own signal, it will not have knowledge of its neighbor's. Only enough neighbor-to-neighbor RF isolation will prevent an FDX band transmission from interfering with another's downstream receive signal.

Unfortunately, enough isolation within a single tap or across taps cannot always be guaranteed. Some taps, in particular mid-to-high value Taps (29 dB, 26 dB, 23 dB), will have sufficient RF isolation from tap to tap. These dB relationships will be calculated by the CMTS as part of the network "sounding" process, used to determine IGs, TGs, and QAM profiles, and update them periodically. Figure 15 is an analysis of a single RF leg in an N+0 environment of a likely IG/TG determination using typical tap specifications.

As shown in Figure 15, users on the same tap are naturally going to form an IG. However, as we will see later, the statistical likelihood of such a small group competing for use of FDX bandwidth in a way that impacts services is negligibly low until FDX penetration becomes very high (for which further features can be enabled). This leads to defining TGs – aggregating IGs simplifies the scheduler aspect by having a smaller number of groupings to manage: Divide to conquer.

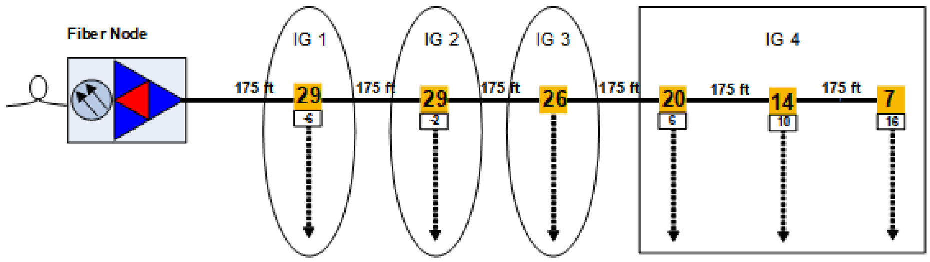

		SNR in Total Channel (108 MHz to 684 MHz)					
RECEIVE	SNR (dB)	TRANSMIT					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	-4.7	40.4	40.4	40.4	40.4	40.4
	Tap 2	40.4	-7.2	36.2	36.2	36.2	36.2
	Tap 3	40.4	36.2	-5.4	32.2	32.2	32.2
	Tap 4	40.4	36.2	32.2	-7.3	24.0	24.0
	Tap 5	40.4	36.2	32.2	24.0	-5.7	13.7
	Tap 6	40.4	36.2	32.2	24.0	13.7	-7.2

Figure 15 - FDX SNR per Interference Group Analysis using Current Tap Performance

Figure 15 also shows the FDX band SNRs for each IG. Beyond the "same tap" limitations (shown in the white cells, with negative SNR if an FDX CM is transmitting at the same time), we can see, for example, within IG4, a useable but low SNR would lead to 128-QAM maximum and as low as 8-QAM when an FDX CM on one of these taps was transmitting. By forming an IG across these taps, these low SNRs are avoided by choosing different downstream and upstream time windows for the users in the IG.

Once IGs are established, the CMTS scheduler executes its "normal" operation with the additional rules. The CMTS scheduler is the arbiter of access to the network, allocating time slots for downstream and upstream transmissions to fairly and efficiently serve the aggregate demand. The CMTS does not typically or explicitly need to pay attention to a single user's downstream and upstream access. That changes with FDX, as there is an awareness needed that accounts for IG/TG relationships. This extra scheduler "step" makes sure that a downstream receiver signal of a potentially vulnerable user is not interfered with by an adjacent FDX upstream.

4.2. Symmetric Multi-Gigamania

What happens as FDX users increase in penetration? To enable fair and efficient access, these IGs and TGs can be granted slices into sub-blocks of the very sliceable OFDMA spectrum, so that availability of FDX bandwidth is always possible for any TG.

Figure 16 shows an example whereby the entire 108-684MHz FDX band is activated. Three TGs each have 384 MHz of downstream and 192 MHz of upstream capacity in the FDX band, or about 3 Gbps downstream and 1.5 Gbps upstream (including the 85 MHz legacy upstream). Sub-band allocations are called Resource Block Assignments, or RBAs. Of course, the FDX band is not all of the available DOCSIS downstream. For example, there could be another 192 MHz of non-FDX D3.1 spectrum in the downstream, for a total of 5Gbps, and also DOCSIS 3.0 spectrum.

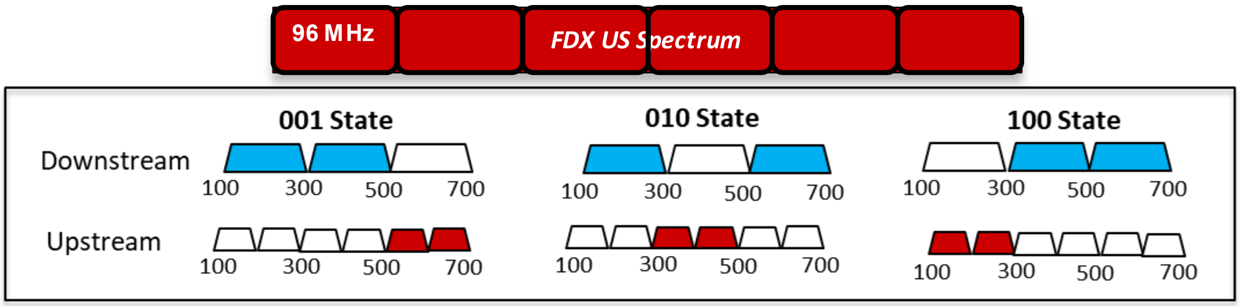

Figure 16 - OFDMA RBAs Example Fully Utilizing the FDX Band

FDX requirements have been written such that RBAs can be dynamically allocated as traffic demands shift with time. Figure 17 demonstrates this feature, with FDX band allocation on the vertical axis and time on the horizontal axis.

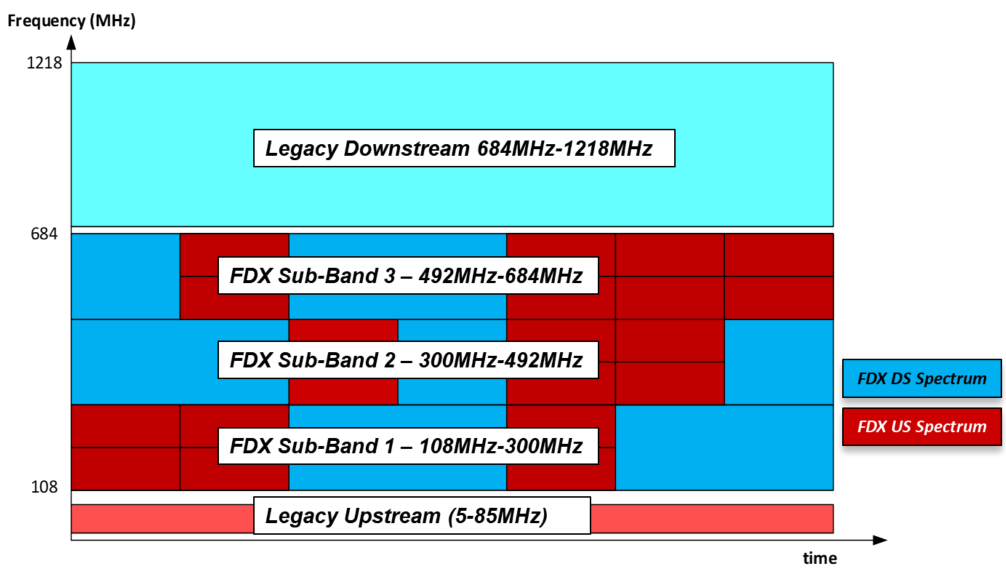

Figure 17 - Example of FDX Band Utilization vs Time Enabled by Dynamic Resource Block Assignments (RBAs)

Two last important points are in order with regard to FDX spectrum:

1) Because FDX is based on DOCSIS 3.1 technology, FDX can only be activated as FDX spectrum where downstream DOCSIS 3.1 spectrum is located.

2) Existing DOCSIS 3.1 devices using the FDX band must be aware of FDX, as this ensures they are effective participants in the IG/TG mechanism. This mode – a software (SW) upgrade to existing CPE – is known as FDX-Light (FDX-L) mode.

4.3. Network Components to Support FDX

Figure 18 identifies components of the end-to-end network effected by FDX. The red highlights identify areas of the network impacted by FDX that we will describe in the sections below.

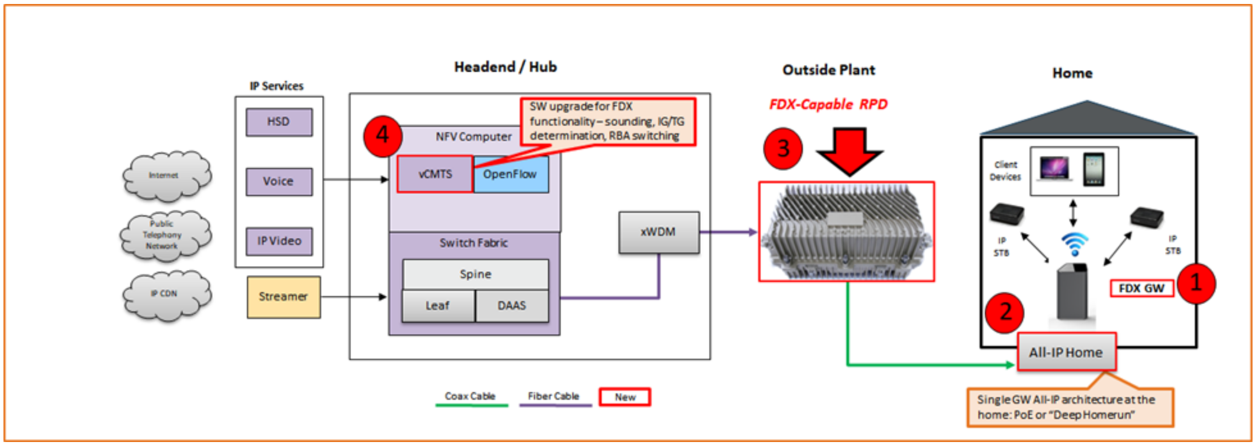

Figure 18 - Network Component Upgrades for DOCSIS 4.0 Full Duplex

4.3.1. *Consumer Premises Equipment (CPE)*

As with DOCSIS 2.0, DOCSIS 3.0, and DOCSIS 3.1, support for FDX requires new silicon development to implement the wider upstream, the additional OFDMA blocks to fill it, and to enable EC. Support for the new scheduling aspects are also necessary, although this is SW functionality and not necessarily a new hardware requirement. Existing DOCSIS 3.1 modems are expected to participate in an FDX system in "FDX-L" mode, as previously described.

The overlapping spectrum means that the CPE gateway RF design deviates from the typical diplexer-first design, relying on an arrangement based on splitters/couplers to accommodate the overlap while maintaining legacy bands.

4.3.2. *DAA Fiber Node*

Following the CPE upstream to the node, any node in an FDX system must also support the expanded bandwidth, OFDMA blocks, and Echo Cancellation. As with the home gateway, the internal RF design of the node must account for overlapping bandwidth in a way different than a traditional FDD node. This is one of the more complex parts of the node design – how to manage downstream and upstream paths to support additional couplers, while integrating EC functionality and minimizing RF losses affecting levels.

An example block diagram is shown in Figure 19, with new FDX items highlighted in red. The diagram represents a single node port or RF distribution and receive leg.

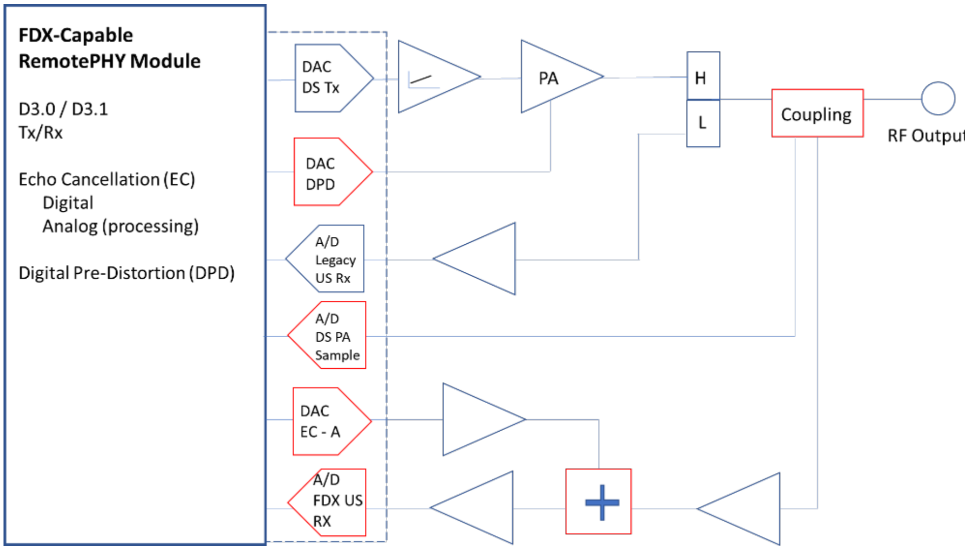

Figure 19 - Simplified FDX-Enabled Node Diagram (Single Leg)

It is possible that a DAA node can be a traditional FDD node on Day 1, and an FDX node later, when needed. Such a node would become FDX via a software upgrade and appropriately built-in hardware processing and RF capability. This would prevent future visits for FDX hardware upgrades.

4.3.3. CMTS Core

As indicated, CMTS "sounding" determines IG/TG relationships. Fundamentally, the sounding process measures relative relationships of devices to one another to determine who may need to be kept apart to avoid interference when spectrum is being shared. It uses that information to adapt the scheduler to prevent interference when FDX CMs are transmitting. It does this by scheduling them in time slots that keep them away from vulnerable receivers based on the IG/TG matrix.

Note that IG/TG relationships are not static. New users can be added and subtracted, and devices may be moved around in a home or business, affecting network characteristics. While not very dynamic, the system has to periodically re-execute "sounding" to establish the correct IG/TG matrix for the scheduler.

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 that affects channel characteristics. The EC must track these and adapt accordingly. There are not explicit transient behaviors or convergence time requirements in the specifications, so this is an area that has seen significant characterization and is ripe for potential optimization and differentiation.

Lastly, as shown Figure 17, in the feature-complete implementation, the scheduler will support optimized RBA allocations that follow the needs of the traffic demands and service tier distribution among the IGs. Today's DOCSIS 3.1 schedulers are sophisticated, and their capabilities not fully leveraged. FDX will add the additional layer of IG/TG no-go zones – at first as simple as a "look-up table," but that will eventually become more dynamic based on traffic patterns, as use of the FDX band becomes better understood.

4.3.4. Home Architecture

The FDX specifications are written from a system engineering standpoint, assuming a point-of-entry terminating home gateway. Analysis was done that considers a single gateway termination, but one that may be located deep in the home, increasing RF loss. We refer to this as the "deep homerun." The advantage of this approach would be that it would give the customer the ability to center the device for Wi-Fi optimization. Also, importantly, a self-install kit (SIK) approach can be considered as a cost-effective way to deploy new devices. SIKs are also highly preferred by customers, compared to visits from technicians, pandemic or no pandemic.

These home install options are shown in Figure 20.

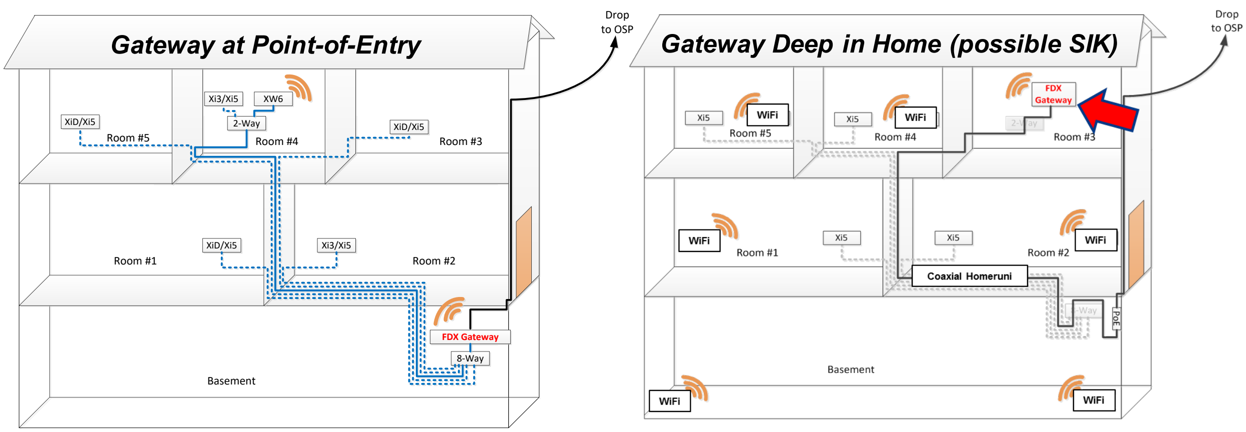

Figure 20 - FDX In-Home Architecture Options

With the practical complexities of having FDX upstream signals within the legacy downstream bandwidth of STBs, it is expected that an FDX home will become an All-IP home. This makes the "deep homerun" SIK approach more realistic if a proactive approach to determine the readiness of the home network for FDX is developed.

4.3.5. Development Progress

Full Duplex DOCSIS requirements began as an appendix to the DOCSIS 3.1 specification and was introduced as a CableLabs project in mid-2016. Recall, it uses the same OFDM and OFDMA-based physical layer for downstream and upstream communication, with the extension of the upstream band to include 108-684 MHz. The Engineering Change Release for this update was in late 2017, and the "cleaned up" spec released in early 2018. Comcast embarked on a node RFP for a DAA-based FDX node and RPD late in 2018 and launched the FDX development program mid-2019. All aspects of development above are underway. In the case of the CMTS, as noted previously, all of our DAA deployments are done with a virtual CMTS (vCMTS), and all FDX will be done via DAA. Thereby, development of FDX features such as MAC management messages, sounding, scheduling, RBA switching, etc., are features under development on the vCMTS.

Current expectations are that 10G FDX end-to-end integration will be taking place in labs across key technology partners, CableLabs, and Comcast in 2021. Operationalizing FDX via tools, training, processes, and integration of services with existing back-office system will also be taking place in order to ready FDX for scalable roll-out. There will be more news to come in 2021 as the integration comes together and we look towards early sites to prove out 10G FDX technology and multi-gigabit symmetrical coaxial last miles. Stay tuned for Gigamania!

4.4. FDX Spectrum Definition

As would be expected, a significant part of the development of the FDX specification was defining the spectrum plan. Lifting the upper spectral edge of the tried-and-true 5-42 MHz upstream/reverse signal path was no small task. The detailed outcome, published in [17], is the 108-684 MHz FDX band – so-called regardless of whether FDX channels occupy the entire band. Figure 21 shows the band plan definition as written into the DOCSIS 4.0 PHY specification [17].

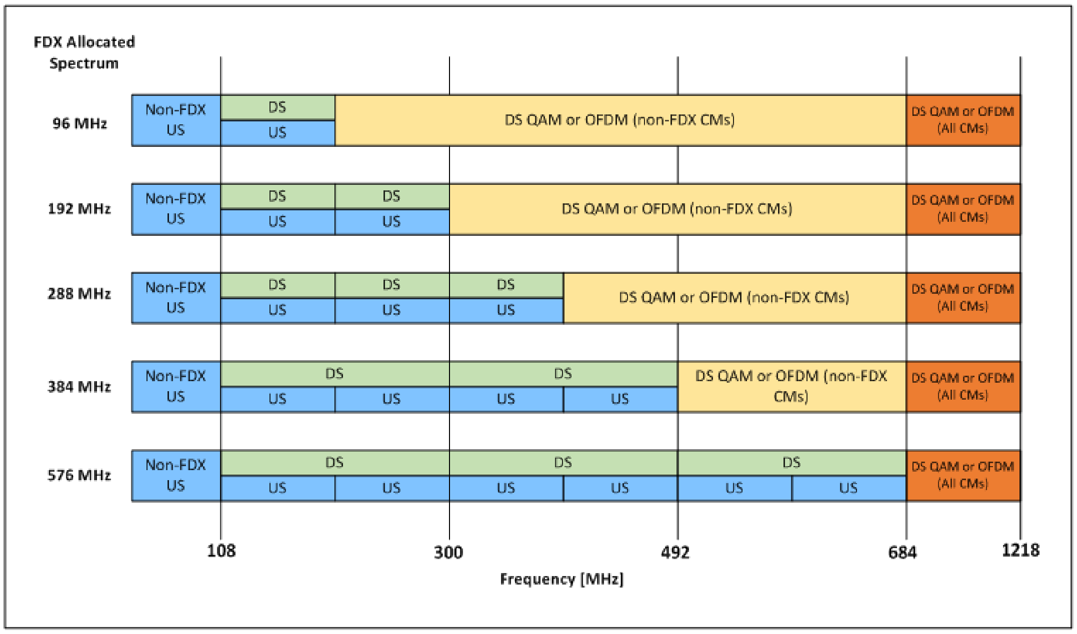

Figure 21 - DOCSIS 4.0 FDX Band Plan [17]

The MUST capabilities as articulated in the DOCSIS 4.0 specification include:

- Allocated FDX spectrum starts at 108 MHz
- 96 MHz of FDX Bandwidth: Located in a single FDX Sub-band from 108 MHz to 204 MHz
- 192 MHz of FDX Bandwidth: Located in two FDX Sub-bands from 108 MHz to 300 MHz, consisting of two 96 MHz downstream channels and two 96 MHz upstream channels
- 288 MHz of FDX Bandwidth: Located in three FDX Sub-bands from 108 MHz to 396 MHz, consisting of three 96 MHz downstream channels and three 96 MHz upstream channels
- 384 MHz of FDX Bandwidth: Located in two 192 MHz FDX Sub-bands from 108 MHz to 492 MHz, each consisting of a single 192 MHz downstream block and two 96 MHz upstream blocks
- 576 MHz of FDX Bandwidth: Located in three 192 MHz FDX Sub-bands from 108 MHz to 684 MHz, each consisting of a single 192 MHz downstream block and two 96 MHz upstream blocks
- The FDX downstream channel and upstream channels sharing the same sub-band must use the same subcarrier spacing and cyclic prefix length. The subcarrier spacing and cyclic prefix on different sub-bands are allowed to be different.

Now, because DOCSIS 3.1 enables a High Split upstream capability defined as up to 204 MHz, and because 1 Gbps is the first step to multi-gig symmetric services, the DOCSIS 4.0 FDX requirements also accommodate the high split as a starting point. With FDX spectrum beginning at 108 MHz, high-split spectrum will overlap with FDX in the first 96 MHz block, between 108 MHz – 204 MHz.

Left to their own operation, a High Split upstream transmission could interfere with an FDX downstream. This is avoided by ensuring that existing DOCSIS 3.1 modems participate in the sounding and scheduling of the FDX band, although they will not be able to use the FDX downstream band. This is accomplished via the FDX-L software upgrade previously described.

The specification also calls out what are mostly intuitive rules for the use of non-FDX spectrum.

4.4.1. Operator Spectrum Management

Figure 22 shows a typical view of the spectrum for an operator that has launched FDX. The specification basis is a 5-85 MHz Mid-Split upstream and a downstream spectrum to 1218 MHz. When activated, more upstream can begin at 108 MHz by adding 96 MHz at a time. The 85 MHz upstream bandwidth, under today's CAGRs, typical node sizes, and an average utilized bandwidth at peak-busy-hour (PBH), provides an upstream lifespan for many years (5-10 yrs.) when coupled with node splits or N+0. Thus, FDX spectrum is likely to be added to allow new product speeds above the approximate 300 Mbps limit of the 85 MHz band alone.

In Figure 22, the additional FDX bands above 300 MHz can be added as product demands deem them necessary, presuming operators have managed the rest of their DOCSIS 3.0 and video spectrum accordingly. Note that the FDX upstream spectrum can ONLY be spectrum allocated to where DOCSIS 3.1 downstream activities exist, and where that is so, the DOCSIS 3.1 device should be participating as an FDX-L modem.

Lastly, we note that product speeds over an 8-10 year time frame may move from 1 Gbps symmetric to 2 Gbps/2 Gbps, 3 Gbps /3 Gbps, etc. A 3 Gbps /3 Gbps service will extend the FDX band to the 492 MHz block edge (give or take and based on bps/Hz efficiency). This upstream allocation provides the spectrum necessary for the very rare bursts of full peak speeds. However, that spectrum is also available for use by the DOCSIS 3.1 downstream, between these rare bursts. By contrast, in an FDD system such as Extended Spectrum DOCSIS, there is still a need to allocate 492 MHz of spectrum for a 3 Gbps service. However, in this case it is a fixed upstream, set aside to be used only for the very rare upstream bursts. The vast majority of time this prime portion of coaxial spectrum sits idle. To compensate for idle spectrum no longer available for downstream, the entire plant needs to be rebuilt to extend the spectrum, including every active, tap and all in-home splitters, and in so doing also pay the penalty of another lost 100 MHz due to guard band above 492 MHz.

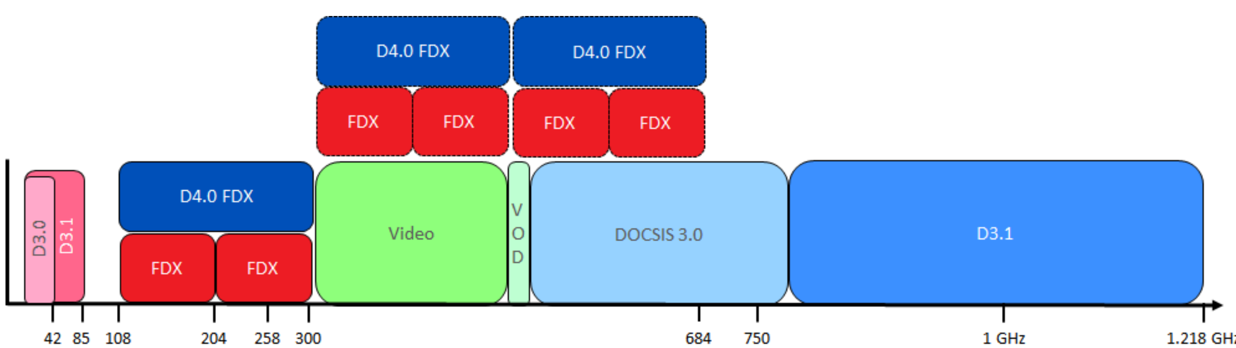

Figure 22 - "Classic" View of Mid-Split Based FDX Implementation

Figure 23 and Figure 24 show some of the practical scenarios to be accounted for as any FDX roll-out begins. In Figure 23, for example, High Split modems may be in the field alongside sub split modems and FDX modems. Often it is best to visualize the spectrum "as the node sees it" and "as the CM sees it," because the introduction of FDX and the mixing of CM splits on the plant tell a complete story. In Figure 23, we can see that as a node goes from High Split to FDX mode, the spectrum between 85-108 MHz is vacated, removing capacity from High Split only devices. This could matter to devices offering 1G upstream service, since High Split does not provide much margin above 1 Gbps of total capacity. It may, for example, force these CMs to move to an all-OFDMA upstream configuration if they are not already using that, or force the Time and Frequency Division Multiple Access (TaFDM) mode, whereby SC-QAM and OFDMA exist in the same spectrum in different time slots. Lastly, it could mean new FDX CMs for the 1G upstream users, since they are the more likely customers to require multi-gig performance sooner than the rest, and only FDX can go above 1Gbps.

The other item to identify in Figure 23 is that DOCSIS 4.0 downstream is shown in the 204 MHz to 300 MHz range, and a DOCSIS 3.1 downstream above 1 GHz. The line pointing at each is to identify this OFDM block can be in either location It is to point out that spectrum migration is a living plan, striking a balance between migration complexity and available empty bands to define the sequence of events. For example, if there is no spectrum used above 1 GHz, then adding more DOCSIS 4.0 (just DOCSIS 3.1 at that point) above 1 GHz can defer some of the complexities of overlapping FDX spectrum. Such a determination does not need to be made now, but as more DOCSIS 3.1 is demanded, the overall spectrum plan should be visited regularly against feature roadmap and device penetrations in selection of the optimal path forward.

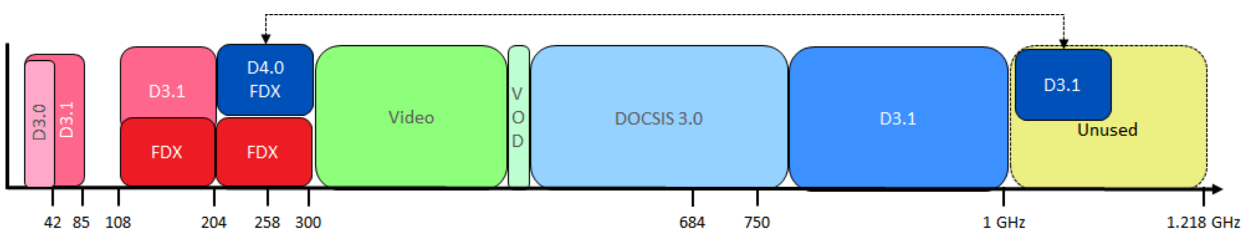

Figure 23 - DAA Node in FDX Mode with Mid-Split, High Split and FDX CMs – Mid-Split based Diplex

Lastly, consider Figure 24, which represents an FDX concept where the node in FDX mode has a High Split diplex instead of a Mid-Split diplex. New FDX spectrum is added "on top" of the High Split spectrum, as shown. There may be areas built only as High Split, such that an FDX upgrade that places 4+ FDX upstream blocks above the High Split band would enable current devices to operate as they already do, while allowing customers seeking >1G speeds upstream to obtain an FDX CM, and access the FDX channels beginning at 258 MHz.

As is common when engineers are tasked with defining what is often the far future when writing technical specifications, the specification's wording is ambiguous for anyone seeking to "comply to the letter" of its definition. However, this may be an implementable option and is currently being evaluated as a viable deployment model. As pioneering work on FDX integrations get underway, it is likely that many specification interpretations and meanings will need to be worked out among collaborating partners. This is a common component of technological advancements.

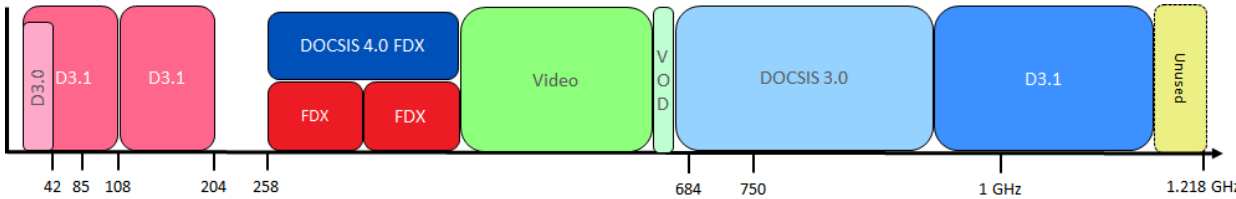

Figure 24 - FDX Node Concept as a High Split Augmentation

4.5. More DOCSIS 4.0

The 10G initiative is based on understanding the capabilities of the HFC network, and the emergence of key technologies and applications of the future. However, "10G" does not explicitly describe a "how-to." This is purposeful, because differences exist among MSO starting points and roadmaps. Different tools in the toolbox inevitably occur over time. A good example of this was different approaches taken to deal with the growth of HD television – analog spectrum recapture using Digital Terminal Adaptors, switched digital video (SDV), and/or 1 GHz expansion all were used.

Different approaches are also available for 10G. The DOCSIS 3.1 specification calls out optional use of spectrum to 1.794 GHz (which goes conversationally by 1.8 GHz). However, at the time (2012), the system engineering was deferred at that time, as this was not a 2012 priority. This has now been completed as part of the DOCSIS 4.0 specification and is referred to as "Extended Spectrum DOCSIS" or Frequency Division Duplex DOCSIS (FDD).

FDD goes beyond the original DOCSIS 3.1 upstream also, defining an upper spectral edge for upstream connectivity beyond 204 MHz, much as FDX does. The FDX upstream grid and current 96 MHz block OFDMA definition enables synergy between FDX and FDD. This synergy is shown in Figure 25.

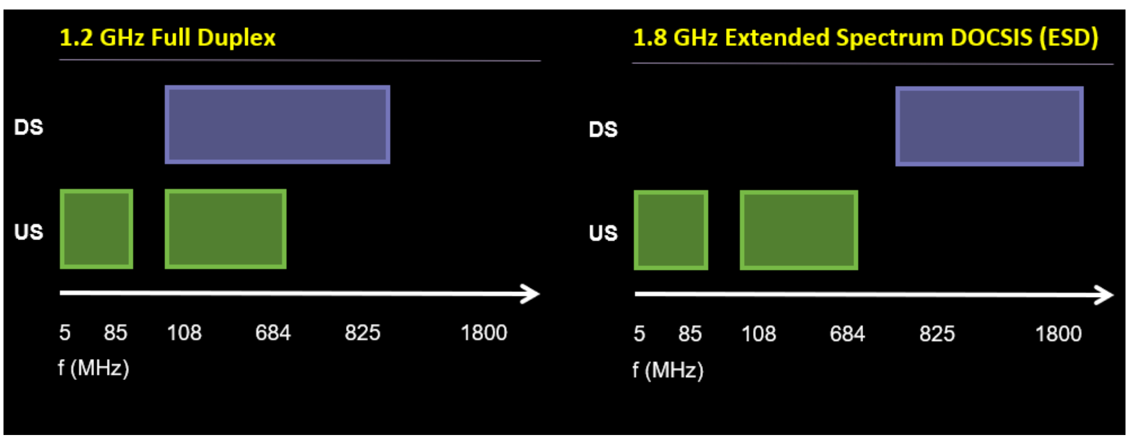

Figure 25 - DOCSIS 4.0 Full Duplex and Extended Spectrum Increase DS OFDM and US OFDMA by Different Means

As shown in Figure 25, a system design that supports both FDX functions and a 1.8 GHz downstream upper edge can have a common number of downstream OFDM and upstream OFDMA blocks. This is beneficial for silicon solution vendors when developing signal processing chains and sizing chip resources. However, they are allocated differently on the RF coaxial plant.

In a potential "DOCSIS 4.1" future, we can envision the two DOCSIS 4.0 technologies of FDX and FDD as complementary. The combined power of FDX and FDD provides yet another avenue for capacity expansion on the network. This is shown in Figure 26.

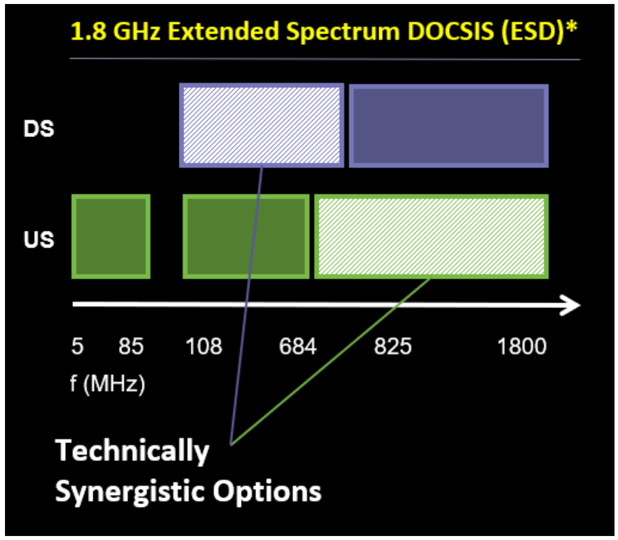

Figure 26 - DOCSIS 4.0 Full Duplex and Extended Spectrum are Complementary in Nature

4.6. FDX Amplifiers

4.6.1. Echo Cancellation-Based Amplifiers

The FDX specification was written under the assumption of an N+0 network. This simplifies the introduction of FDX by getting fixed diplexed amplifiers out of the way, and it aligns with an architecture path some operators are taking. However, it is not technically limited to an amplifier-free plant. The EC technology developed can apply at any point in the network to manage overlapping spectrum, and this can include amplifiers. Of course, these are not traditional amplifiers, but a new class of device that includes this essential signal processing.

The system engineering aspects of EC-based FDX amplifiers includes:

1) How good does the EC need to be?

2) Amplifiers can have an "amplifying" effect on IGs, depending on the HFC design. How important is this?

3) How does the amplifier specification change with respect to RF performance? Are there cascade limitations?

An EC-based amplifier concept is shown in Figure 27. The nature of overlapping spectrum and gain in both directions creates a full-circle loop gain path, by design. There is a concern even in normal HFC amplifiers that indirect loop gains occurring in certain frequency bands could cause oscillation. The EC must, as a minimum, suppress this gain loop, meaning net gain around the path must be < 0 dB.

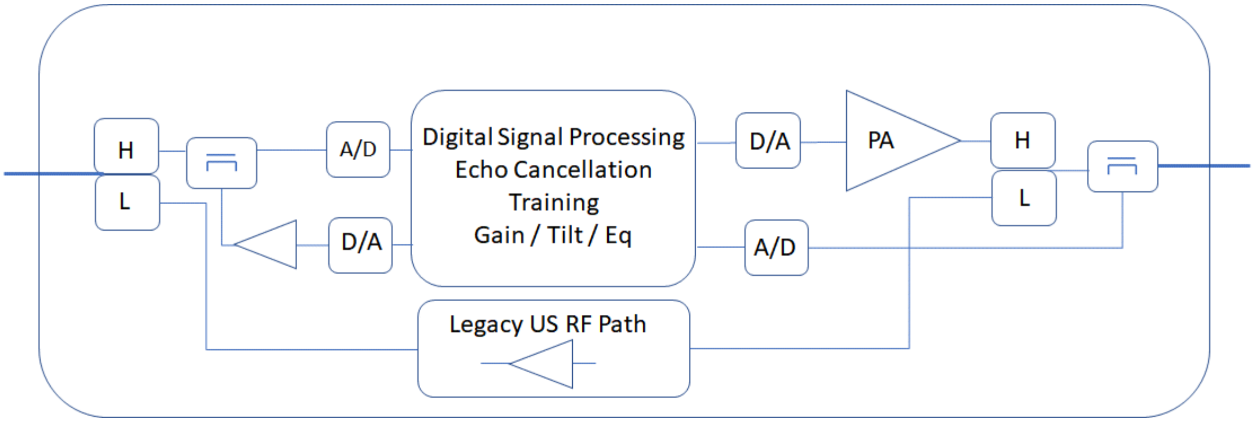

Figure 27 - An Echo Cancellation-Based Amplifier to Support FDX Beyond N+0

Echo-cancellation operates similarly to that of the node. A sample of the upstream signal is taken so that an opposite phase, equal magnitude version can be added in front of the downstream amplifier in Figure 27. Similarly, the signal from the downstream amplifier is sampled and an anti-downstream version added at the input to the upstream amplifier. As in the node, the EC adjusts for changes on the echo response.

Analysis of the EC performance necessary was calculated for two scenarios, based on common HFC amplifier RF level and noise parameters, and shown in Figure 28 [16]. The left figure shows the use of overlapping spectrum for the entire FDX band without constraints. In this case, the objective is to maintain high enough SNR (minimal residual EC degradation) to support the OFDM and OFDMA formats defined in DOCSIS 3.1, with the assumption being up to 2048-QAM in the downstream and 1024-QAM in the upstream, and assuming a reflection magnitude of -15 dB exists at each port. The EC performance requirement, shown along the X-axis, is about 75 dB before residual echo interference begins to impact the SNR required for 1024-QAM. Using the dB relationships between formats (3 dB per bit), these charts can be extrapolated in either direction (for more or less bandwidth efficiency) to arrive at EC requirements for other reflection or SNR assumptions under existing noise assumptions. For example, for a less-efficient 256-QAM upstream, which has a 6 dB lower SNR threshold, the EC requirement reduces to about 65 dB. These are challenging requirements, particularly for a single-stage echo-cancellation design.

For the second case, the amplifier is not required to support the simultaneous use of shared spectrum. In this case, the goal of the EC is to assure that the amplifiers are able to operate sufficiently in a linear mode, thus not degrading noise and distortion. This case, shown on the right, indicates that 25-30 dB is sufficient, which is a significant relaxation compared to the first case.

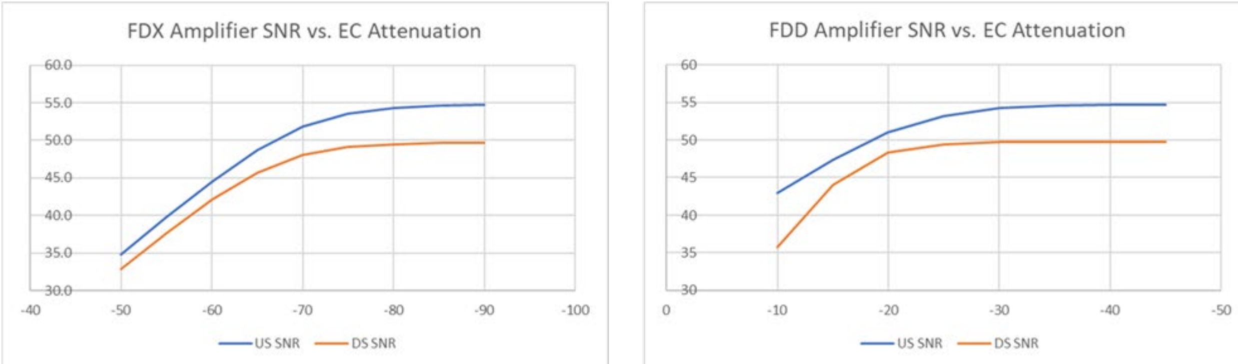

Figure 28 - Worst Case SNR vs. Echo Cancellation [16]

What this analysis points out is that FDX amplifiers are likely to be a tradeoff of complexity, cost, efficiency, HFC design, and also SW complexities associated with the additional EC to be trained and to manage. As well, possible refinement of IG/TG algorithms associated with RBA assignments to include the impact of FDX amplifiers.

4.6.2. RF Isolation-Based Amplifiers (No EC)

Because HFC has always been an FDD system, diplex filters have always been used. Their purpose is to isolate downstream and upstream signal paths from each other. Diplex splits have evolved over time from 42/54 MHz to 65/85 MHz, 85/102 MHz, and 204/258 MHz. In the future, with Extended Spectrum DOCSIS, higher splits are expected.

The downside of diplexers are that when the split needs to be changed, it is labor intensive to change all of them (2 per amplifier) or swap active components. In addition, valuable spectrum – 20-25% of the band edge – is lost in the transition band to make filters of practical cost, size, and repeatability. The amount of MHz lost to diplexers increases with absolute band edge for the split. For a 492 MHz FDD split, it's about 100 MHz of lost spectrum!

An ideal solution would be an amplifier with NO diplex filters. This would make it feasible to build a transparent, flexible network with zero-touch for upstream/downstream spectrum re-allocation, and no guard bands.

There are two significant challenges to a diplexer-less amplifier:

1. How do you create an amplifier with loop gain below 1 to prevent oscillation?

2. How do you create a system able to work with real world return losses?

Consider Figure 29, which looks at the input port of a diplexer-less amplifier, where in their place are specialized multi-port directional couplers (MC x).

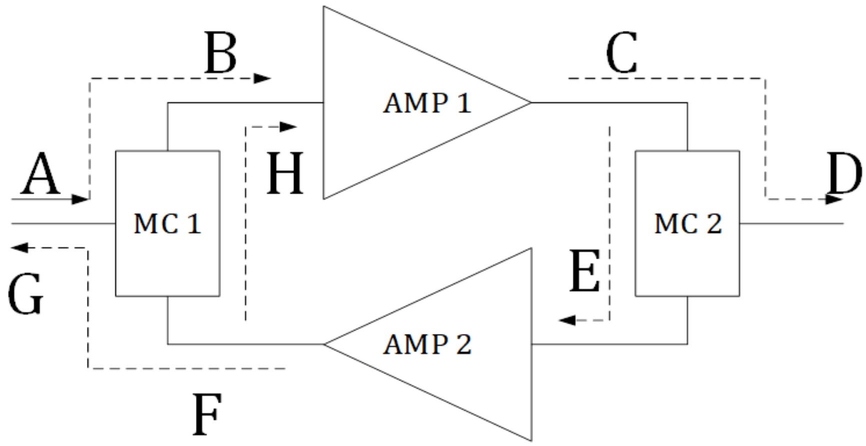

Figure 29 - Evaluating the Signal Relationship of a Diplexer-Less Amplifier (courtesy Technetix)

We define the following (all in dBmV):

A = input signal

B= (input signal A) – (insertion loss of the coupler MC1)

C = B + (gain Amp1)

D = C – (insertion loss of the coupler MC2)

E = (leakage of C in MC2)

F = E + (Gain Amp2)

G = F – (insertion loss of the coupler MC1)

H = (leakage of F in MC1)

Then, looking at the amplifier output to the downstream (point D):

D = (input A) – (Isolation MC1) + (Gain Amp1) – (Isolation MC2.)

Thus, the amplifier works only if:

(1) H < B :

(Gain of Amp1) + (Gain of Amp2) – (isolation MC2) – (isolation MC 1) < 1

(2) RL = G < A – 20dB :

(Gain of Amp1) – (isolation of MC2) + (Gain of Amp2) – (2xIL MC 1) < -20dB

An amplifier that can achieve both (1) and (2), can be a transparent, diplexer-less amplifier. By combining an amplifier with these two operating principles, along with a 54 MHz high-pass filter (HPF) and a 684 MHz lowpass filter (LPF) – the edge of FDX or Extended Spectrum DOCSIS upstream – the frequency response behavior shown in Figure 30 is obtained.

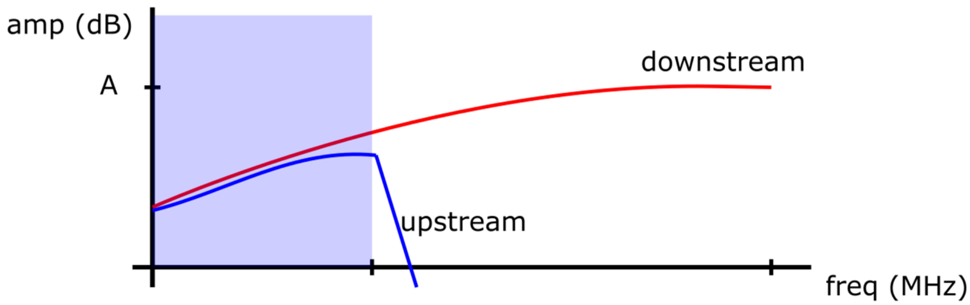

Figure 30 - Forward and Reverse Freq Response Characteristic of a Diplexer-less Amplifier (Courtesy Technetix)

In Figure 30, the purple area shows the region for which there is gain simultaneously in both the upstream and the downstream signal paths, without diplex filters in place. This creates a uniquely transparent and flexible network between the node and the customer premise, either by passing FDX signals seamlessly or by supporting any desired split that does not exceed 684 MHz in an Extended Spectrum FDD system.

From an FDX system engineering standpoint, the diplexer-less amplifier results in an extension of the echo cancellation capability in the time domain. For example, if a 5-tap string becomes a 10-tap string, reflections occur from further away points from the EC circuit. Typically, long-distance reflections also attenuate in magnitude with time. This will not be the case with gain in the return path that will forward along a reflection that is on the “south” side of the amplifier. This makes the job of the EC more challenging.

4.7. Traffic Engineering for FDX Services

Consider the case of an N+1 system based on FDX and shown in Figure 31. When an amplifier is included, using current tap RF performance, one of the effects is the expansion of Interference Group 4 (IG4) to the “south,” or home-facing, side of the amplifier. These users become part of the last IG of the tap string before the amplifier. This is because of the limited drop-to-output isolation characteristics of today’s taps.

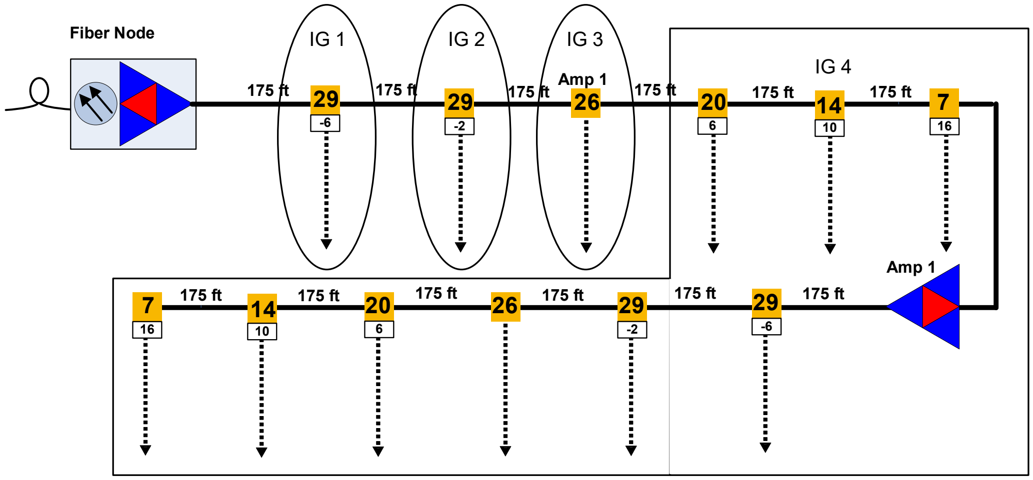

Figure 31 - Potential Interference Group Expansion due to Amplifier on an FDX Network

Taps with port-to-output isolation like those with higher tap values could eliminate this phenomenon. We have developed tap specifications with FDX requirements in mind.

Nonetheless, since current taps will be in the field for years, leading to enlarged IGs, traffic engineering analysis was performed to understand, quantifiably, the sensitivity of FDX performance to the size of an IG, the total spectrum, the service group size, speed tiers, etc. The empirically-based modeling is founded on real network traffic distributions gathered on production CommScope CMTS platforms over many years. Figure 32 shows the components of the traffic useful for helping to understand the problem and visualize the results [4].

The traffic can be broken into three components [4]:

1) Average Service Group (SG) Utilization = Number of subscribers in the service group x average utilization per subscriber at PBH.

2) Ripple = variations around the average utilization because traffic is bursty and not a constant stream equal to the average. Though statistically imprecise, the ripple is somewhat like standard deviation. However, it also encompasses a component tied to the probability of bursting to peak speed.

3) Max Burst = Just that, the moments when the traffic burst is 100% max'd out at the peak speed offering of the high speed data (HSD) product in this service group

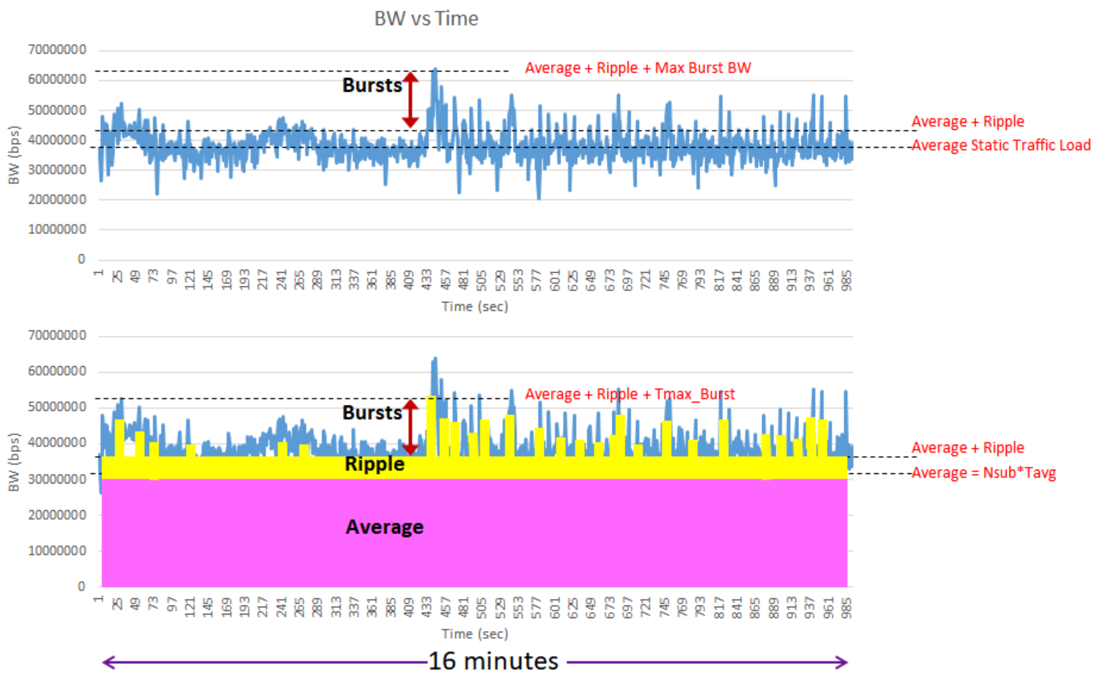

Figure 32 - Breakdown of HSD Traffic Components [4]

HSD product speeds directly translate into a minimum spectrum allocation needed to guarantee a maximum speed burst. Operators commonly allocate peak spectrum and an empirically derived additional capacity that accommodates the three components of Figure 32. On real networks, rules have been developed that link service group size, speeds, and percent capacity utilization to define the total capacity requirement. Then, with an awareness of device penetrations, this can be translated to DOCSIS 3.0 and DOCSIS 3.1 spectrum requirements. These rules were re-calibrated again with the introduction of DOCSIS 3.1 and 1 Gbps downstream speeds.

All of the above analysis and empirical rule making apply to FDX – it is still about downstream and upstream spectrum requirements to support certain product speeds. What changes for FDX is that the FDX band is allocated for both downstream and upstream. The infrequent bursts of peak speeds can therefore be called upon to service the downstream or the upstream, and they are not completely independent because of the IG phenomenon. Although from the node perspective there is full duplex spectrum operation, from the CM perspective the introduction of IGs places constraints on who can simultaneously access the spectrum. Figure 33 depicts this observation.

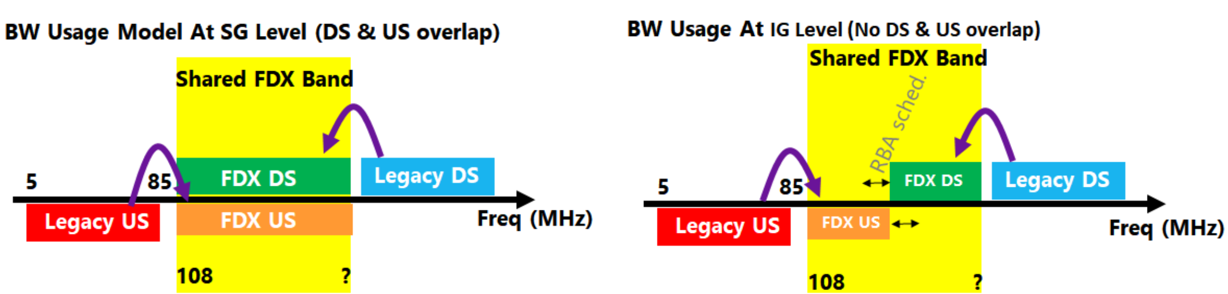

Figure 33 - Managing FDX Bandwidth to Guarantee Peak Speed Bursts [4]

Shared spectrum does not automatically mean peak bursts within an IG will compete for FDX spectrum. It depends on how much FDX spectrum is allocated, the distribution of peak speeds, the maximum speed offering, and of course the size of the IG and the number of subscribers in general. However, the FDX band carries around 5 Gbps. So, as speed tiers increase, at some point, from the perspective of an IG, there may not be enough FDX spectrum to avoid downstream and upstream burst demand exceeding available FDX band, as depicted in Figure 34. Naturally, it is much more likely to be needed for the upstream, since there is a substantially larger downstream allocation for DOCSIS 3.1.

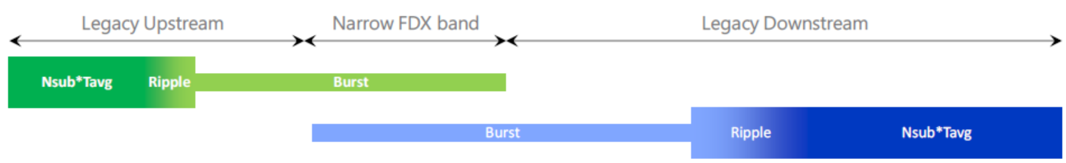

Figure 34 - Representation of Colliding Simultaneous DS and US Bursts of Shared Spectrum [4]

Some context may be helpful at this point. Development of the IG/TG concept was formulated during the early specification discussions of FDX. DOCSIS bandwidth has always been shared in the downstream and in the upstream with great efficiency by using relatively large service groups. For FDX development, instinctively there was a thought that having an IG as small as reasonably possible was a sound goal and would not significantly change the capacity utilization dynamics. A tap being its own IG, or several taps being part of the same IG, is what occurs using typical RF performance. There were not major concerns about the impacts of such small IGs; there was more concern about the process of sounding the network to define them.

The concern for IG/TG size increased when the concept of the FDX amplifier came into the picture to support FDX in non-N+0 systems. It was during system analysis with amplifiers that the IG elongation phenomenon shown in Figure 31 was uncovered. This led to questions:

1) How large can an IG be before there is an impact to the customer experience?

2) What service speed / IG size / spectrum rules exist when downstream and upstream traffic engineering become co-mingled in FDX?

We set out to find answers to these questions, working closely with experts with sophisticated HSD modeling tools at CommScope.

Figure 35 shows a model assessing the subscriber count, spectrum, QAM efficiency trade space for a set of input parameters projected for 2028 that includes a post-Covid average utilization baseline, assumed capacity CAGRs of 35% downstream and 30% upstream, and speed tiers up to 4 Gbps/4 Gbps. In these models every subscriber is given this peak service tier. Although the highest speed tiers are, in fact, typically the lowest penetrated, this obviously makes for a conservative analysis.

What can be concluded from this analysis and Figure 35 is that the subscriber group size (IG size) sharing FDX spectrum from 108-684 MHz can be as large as 64 (red circle) and stay within 1200 MHz of total spectrum, or even within 1 GHz for a 9 bps/Hz net downstream throughput average efficiency. The upstream, even for the 4 Gbps/4 Gbps (green) case, stays within the 576 MHz of maximum FDX band allocation, even for an IG size above 128.

What these studies reveal is that FDX bandwidth can be used extremely efficiently, shared across larger IGs than perhaps initially envisioned, and the statistical nature of peak bursts is better understood. What we have learned is these peak bursts are very infrequent per user, and the collision of bursts from concurrent peaking users is more infrequent still. With an 85 MHz legacy upstream, the upstream bandwidth is sufficient for PBH average utilization for a subscriber size of up to about 150-200 subscribers, shown in Figure 25 as approximately where the red arrow points to the FDX bandwidth breach. Thus, allowing the FDX upstream allocation to be a bandwidth reservoir for peaking users while still serving downstream capacity is an extremely efficient use of precious HFC spectrum, compared to setting aside hundreds of MHz of the highest quality HFC to, mostly idle.

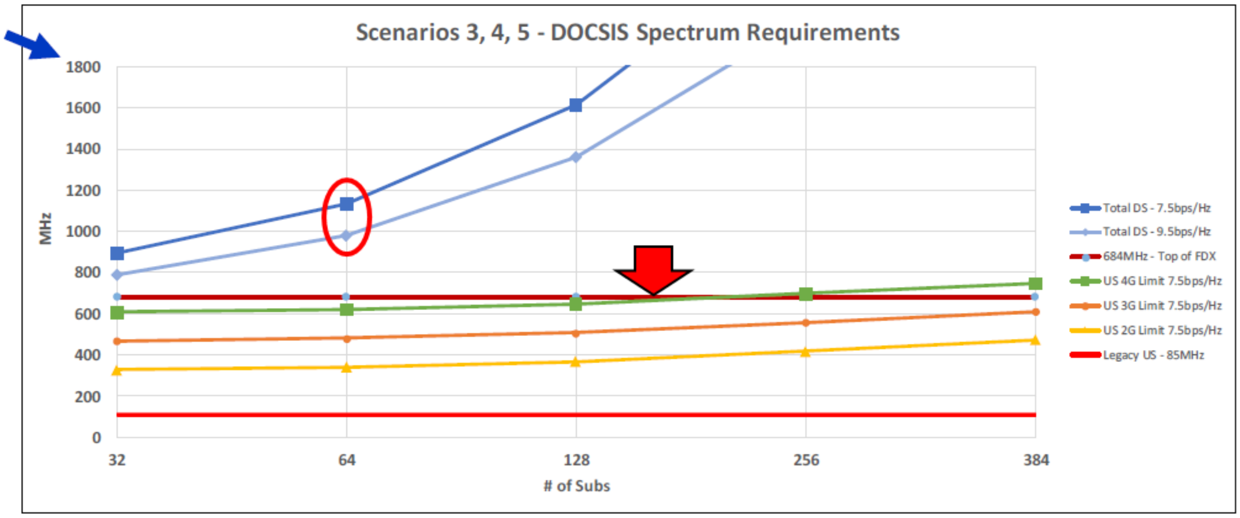

Figure 35 - Total RF Spectrum Required vs Subs Sharing an IG for 3 Gbps/3 Gbps, 4 Gbps /2 Gbps, and 4 Gbps /4 Gbps [4]

4.8. Operationalizing FDX

4.8.1. Echo Cancellation and Diagnostics - Live!

As a starting point for understanding EC math, consider first a non-FDX signal. Let y(n) represent a noiseless signal presented to a receiver, which is a summation of the received signal, and its attenuated and delayed signal copies of delay, D.

$$y(n) = \sum_{d=0}^{D} b_d x(n-d) \quad (\text{non} - \text{FDX Signals})$$

The signal copies come from plant defects including micro reflections, and other frequency response effects like filter roll-off and group delay variation (GDV). Each delayed copy will be attenuated and phase-shifted by a complex coefficient, b_d .

Non-FDX signals can be made quite robust, thanks to equalization technology available in DOCSIS today, enabling an efficient means of adapting equalizer coefficients that compensate for b_d, and deliver y(n) ≈ x(n).

By contrast, an FDX system will have both downstream and upstream signal components. At the node receiver, since signals overlap with no filters to separate them, the received signal y(n) is:

$$y(n) = y_{ds}(n) + y_{us}(n) \quad \text{(FDX Signals)}$$

The FDX echo canceller must estimate both the downstream signal, $b_0 x_{ds}(n)$ and its echo profile, $\sum_{d=1}^{D} b_d x_{ds}(n-d)$ to minimize the downstream contribution at the upstream receiver such that $y(n) \approx y_{us}(n)$, where:

$$y_{ds}(n) = \sum_{d=0}^{D} b_d x_{ds}(n-d)$$

In addition to echo-cancelled MER and the FEC, the coefficient b_d can assist the MSO in understanding the impairments at the FDX receiver at any given time, similar to the familiar linear distortion impairment analysis that can be obtained from upstream equalization [15]. EC analysis can lead to a wealth of PNM knowledge previously unavailable because the non-FDX or traditional node receiver echo profile view was blocked by a diplex filter.

A sample echo profile is shown in Figure 36 in the brown trace labeled “Vendor EC”, representing a snapshot of what an echo canceller would need to compensate for, when connected to a cable plant where the RF taps are spaced between varying lengths of cable (90-132 ft). If a time domain reflectometer (TDR) were to connect to the plant at the same location as the node, it would produce a similar echo profile, shown in the bottom blue trace labeled “VNA TDR.”

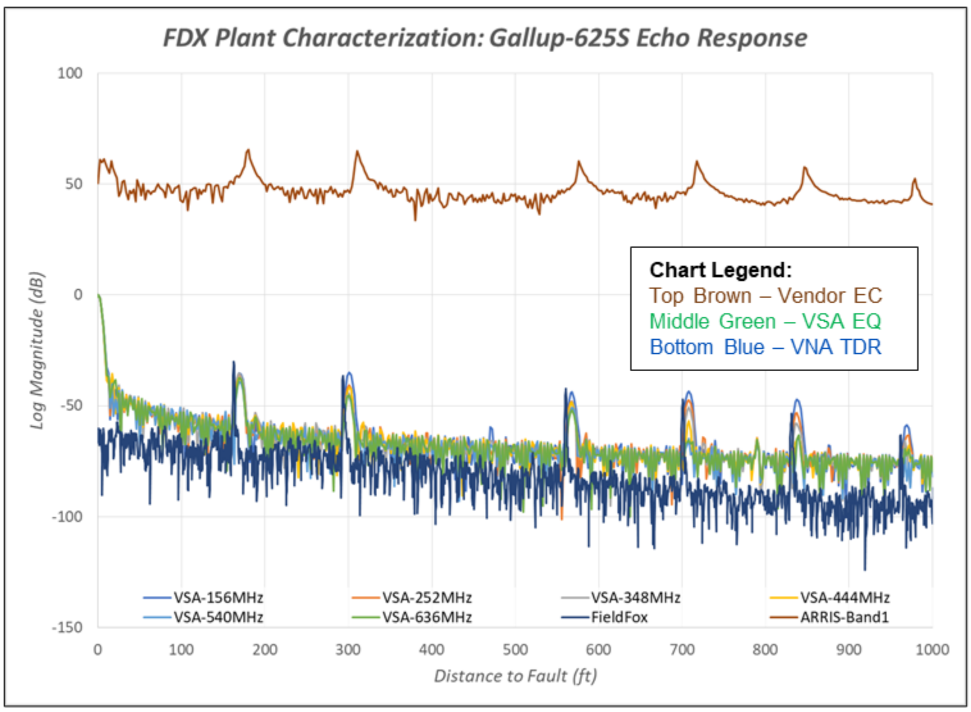

Figure 36 - Echo Profile from New London, CT FDX Field Trial

In Figure 37, the shaded block illustrates the coupling of a vector signal analyzer (VSA) and DOCSIS cable load generator (CLGD) via a resistive (low isolation) splitter at the same location of the node used in Figure 36. All of the remaining traces, labeled “VSA-156 MHz” through “VSA-636 MHz” are fully adapted VSA equalization coefficients, b_d, at all of the FDX center frequencies associated with a 96 MHz OFDM block transmitted by a CLGD and received by a VSA. All VSA traces mostly overlap with one another except for minor differences at each of the tap locations, where it is clear the overlays are visible.

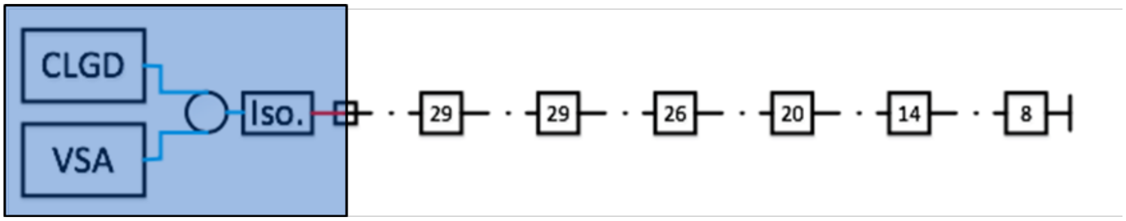

Figure 37 - CLGD/VSA Echo Profile Coefficient Capture

Diurnal and seasonal trends typically result in small and nearly continuous response variations over time. This includes changes in either temperature, or other weather-related events including wind loading. EC response data can tell us the same thing we'd learn by connecting a TDR at the node location. They can correlate to a plant map's precise location of every component, including taps, passives, splices, and cables. Each spike in the echo profile of Figure 36 corresponds to each tap within the feeder leg shown in Figure 38. The distance to the fault value of each spike directly relates to the length of cable between the node and each of the taps, specifically by a factor of 2. So, 90 ft between the node and the 2nd tap – a 2-port, 29 dB tap in Figure 38 -- is represented as approximately 180 ft from the fault measured in Figure 36.

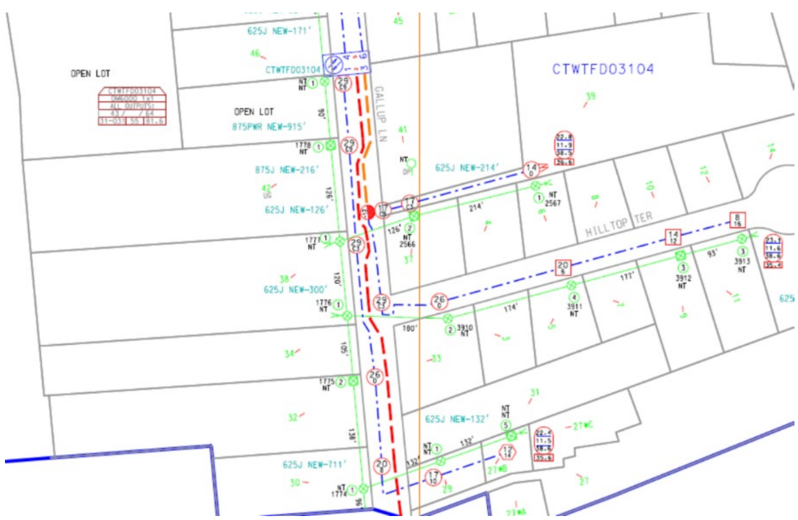

Figure 38 - CATV Plant Map Corresponding to Figure 36 Echo Profile

It is expected that it will be reasonably easy for EC technology to adapt to these static or slow-moving echo profiles. Transient behaviors, on the other hand, can represent a more stressful scenario for EC. Plant maintenance activities alone can result in rapid and potentially discontinuous response variation over time. Some common plant maintenance activities include:

a) Tap face plate removal and insertion

b) Interference testing

c) Seizure screw checks (loosening and tightening)

Tap port termination quality was evaluated as well and showed negligible impacts to the echo profile.

Figure 39 shows how the reflections behave on a millisecond scale, as a technician removes or installs a tap face plate. When the faceplate is removed, an internal continuity switch occurs, resulting in a spike representing a nearly 30 dB change. In these cases, the EC will need some time to adapt to the abrupt changes, and both MER and FEC readings will also reflect this transition. However, MER and FEC will not be able to identify the precise location of where this change occurred. This is precisely where EC may provide MSOs with new value.

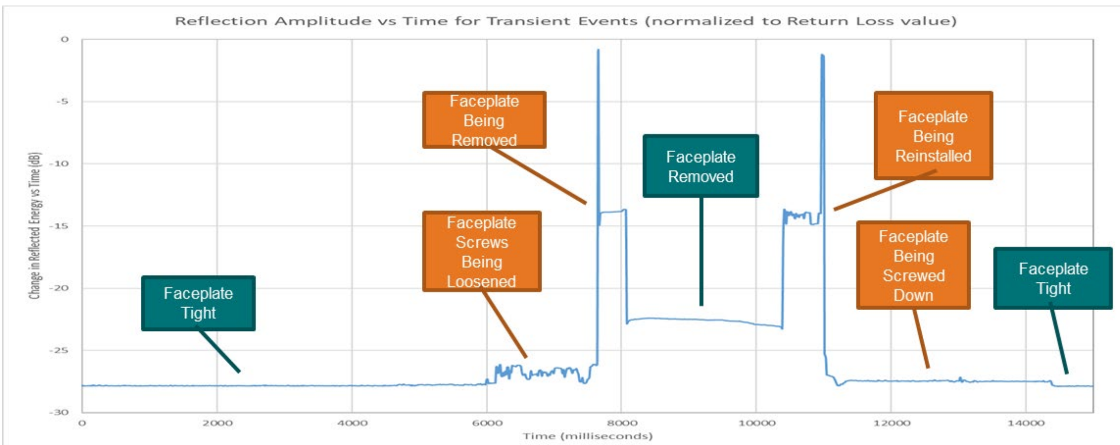

Figure 39 - Single Coefficient Variation from Faceplate Removal and Insertion

A potentially new operations dashboard may abstract away many of the gory details associated with analyzing EC echo profiles, perhaps only raising awareness when the EC has moved into a more aggressive mode of adaptation, due to a breach of a discontinuity threshold. This mode mostly provides confirmation of ongoing plant maintenance activities, but it can also assist MSOs in recognizing nearly equivalent changes occurring outside of planned plant maintenance activities. Those alerts, combined with precise location of the fault, could represent an unprecedented level of Proactive Network Maintenance.

4.8.2. Proactive Network Maintenance

Speaking of Proactive Network Maintenance (PNM), it has been an integral part of the DOCSIS specification since DOCSIS 3.0. For nearly 15 years, cable operators have been learning how to harness the vast capabilities available within the PNM suite of tools. With the addition of Machine Learning and Artificial Intelligence (ML/AI), operators are gaining unprecedented visibility into the echo profiles of their coaxial distribution networks. These tools enable operators to remotely evaluate the spectral performance at nearly every location of plant that has a cable modem installed.

The (partial) list below summarizes why operators appreciate what PNM brings to the network reliability and positive customer experience fronts. The first four help to optimize the channel performance in the face of network variations that are different from one end-to-end connection to another. The bottom four are powerful tools for monitoring, identifying, troubleshooting, and remediating problems found.

- Upstream adaptive pre-equalization
- Channel estimate coefficients
- Downstream equalization coefficients

- RxMER per-subcarrier analysis
- Full Band Capture (FBC) spectrum analysis
- Spectral Impairment Detection (SID)
- Triggered upstream spectrum capture
- Active and quite probes

What do these PNM capabilities all have in common? Most of them can play at least some role in directly or indirectly measuring return loss and reflected energy within the coaxial plant. These tools can be used together with system design maps to inform echo cancelers in FDX deployments.

Typically, the PNM tools are used to detect, analyze and troubleshoot problems within their respective RF spectrum and direction. For example, in the case of upstream adaptive pre-equalization, DOCSIS station maintenance requests are used to continuously monitor and adapt the upstream channel for micro-reflections (echoes). Our PNM systems provide a high-resolution analysis of DOCSIS channels in use. In the case of a typical 24-tap SC-QAM pre-equalizer, there are resolution limits to the equalizer. For example, a 6.4 MHz wide SC-QAM channel will have a symbol rate of 5.12 Mbps, resulting in approximately 213.33 kHz resolution bandwidth (RBW). However, using DOCSIS 3.1 OFDMA upstream pre-equalizer coefficients, RBW will be as high as 25 kHz, and typically 50 kHz.

Collecting, analyzing and grouping the pre-equalizer response signatures results in an exceptional system for locating small pockets of network impedance problems that cause low frequency echoes. Figure 40 illustrates the effectiveness and accuracy of this technique. While the current DOCSIS upstream spectrum operates at lower frequencies than FDX, this will undoubtably play an important role in troubleshooting and maintaining our physical plant. Having the echo response at these low frequencies enables operators to pinpoint virtually every impedance mismatch on the outside plant. It's also important to note that pre-equalization is available on all DOCSIS compliant devices since version 1.1.

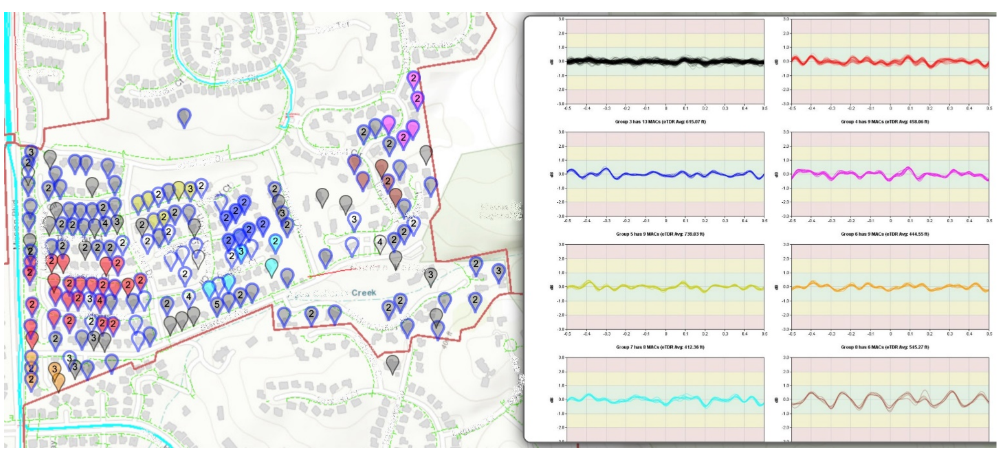

Figure 40 - Upstream Adaptive Pre-Equalization Analysis; Echo Response Groups, Mapped

Arguably the most useful PNM capability for troubleshooting RF performance is downstream Full Band Capture (FBC). The previously described technique of the pre-equalizer response grouping has been extended to downstream spectrum analysis. Figure 41 illustrates a common example of a standing wave, created by a reflection cavity on the plant. Reflection cavities are created by two or more impedance mismatches, causing echoes that become superimposed to form standing waves. These waves have predictable periodicity that can be used to calculate the exact distance of a fault, if the cable type and length is known.

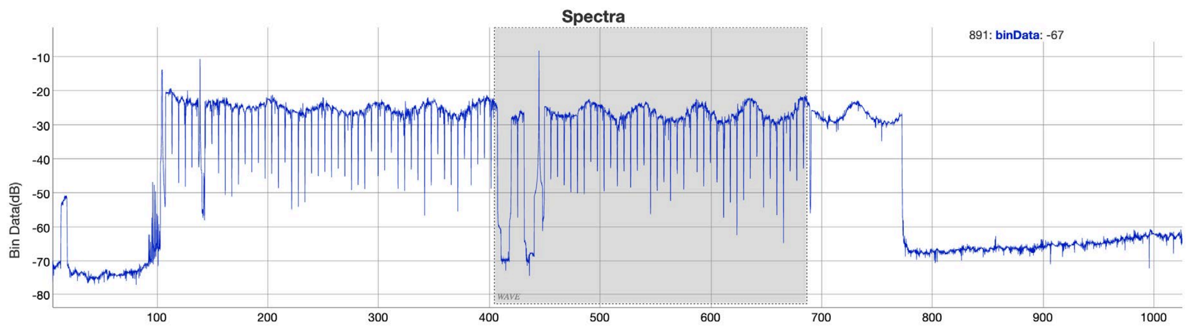

Figure 41 - Downstream Full Band Capture, Standing Wave

26 MHz to 1026 MHz at 117.18 kHz RBW

Similar techniques of signature matching based on Artificial Intelligence (AI) prove very effective at grouping common response issues. Figure 42 shows that the standing wave signature is localized to a group of 47 similar frequency responses. These 47 cable modems share a common 1842 foot .875 P3 trunk cable with a Velocity of Propagation (VoP) factor of 0.87. Having a peak-to-peak frequency separation of 47 MHz, the echo distance can be calculated using the speed of light in feet-per-second: (983,571,056 / 2) * 0.87 / 47 = 9 feet. In the case of this wave, the higher attenuation at lower frequency indicates that the primary reflector is near the launch end of the cable and the secondary reflector is most likely at the output of the node. In this example, the echo profile of this faulty cable can now be calculated in both directions, at any frequency within its operating range. Relying on well-known principals of antenna reciprocity, the important scattering parameters (S-parameters) can all be calculated for this network segment.

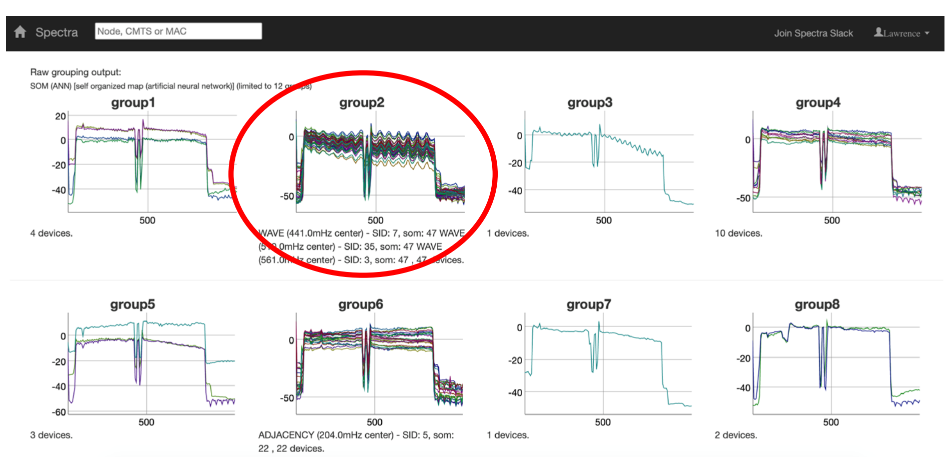

Figure 42 - Artificial Neural Network, Response Grouping; Self-Organized Map (SOM)

In summary, PNM tools have been around for a long time and have become extremely sophisticated at pinpointing plant issues, and typically before a customer would notice an impact. FDX relies by its very nature on overlapping spectrum and knowledge of the plant echo characteristic over time. The advances in PNM align perfectly for the introduction of FDX, and the introduction of FDX will further empower PNM through the availability of EC coefficients, we well as the sounding and isolation information developed for FDX operation.

4.8.3. DOCSIS 3.0 and DOCSIS 3.1 Participation in FDX Sounding

Channel "sounding" is an essential tool for FDX. Cable modems transmitting at high levels in higher frequency spectrum will create new opportunities for co-channel interference (CCI) and adjacent channel interference (ACI) where isolation is insufficient. The ability to coordinate a sounding routine across all modems will be used to assign interference groups (IG). Similar to the PMA in DOCSIS 3.1, there will be opportunities to externalize the sounding routines, allowing for more context and better-informed IG modeling.

For example, a typical FDX deployment will co-exist with a majority population of DOCSIS 3.0 and DOCSIS 3.1 modems. In the case where a customer would like symmetrical multi-gigabit speeds, requiring an upgrade from DOCSIS 3.1 to DOCSIS 4.0 FDX, they will receive a new cable modem/gateway to replace the existing device. The condition of that customer's network and its upstream impact on adjacent customer communications will not be known until after the new modem is connected and sounding has occurred.

This scenario can be addressed by extending the sounding routines and adapting them to use the external PNM infrastructure. Having full band DOCSIS 3.0 and 3.1 spectrum analysis capabilities within the cable modems, the existing equipment can participate in the sounding routine. There are limits to using the log-magnitude data reported by the full band capture, such as lacking phase information. However, in early years of FDX availability, this type of analysis will provide useful information where FDX sounding clients may be sparse or non-existent.

4.8.4. "Inverted" Plant of FDX Band

Along with FDX, and like the High Split, comes a new paradigm of high-power transmission in higher frequency spectrum from the CM. In addition to CCI and ACI, there is another potential problem from these "inverted" levels in the plant. Due to dissimilar carbon content in the metallic cable components, and in the presence of electrical current, ionic migration is constantly occurring. In other words, corrosion is commonplace, and is especially exacerbated in the presence of water.

Corrosion in the form of a molecular junction is known to behave similarly to a diode and is often referred to as a corrosion diode. These corrosion diodes can function in unpredictable ways, often having resistive, capacitive and cascaded diode properties. One of the most common side-effects is a non-linear mixing effect that causes 2nd and 3rd order RF products. In traditional coaxial networks, this is known as Common Path Distortion (CPD) and typically occurs near the output of amplifiers, where composite RF power is high. In digital cable systems, the downstream channels become down-converted and mixed with the upstream frequencies, which are much lower in RF spectrum. The summation of the 2nd and 3rd order products typically manifests as an elevated noise floor in the upstream spectrum.

It can be anticipated that a significant number of corrosion diodes may exist in our collective drop networks, which tend to be older and less well maintained than the hardline portions of the plant. However, with the relatively low power levels of downstream RF in the drop cables, there is insufficient power to forward-bias any diodes that may exist. In the case of High Split or FDX, there will be higher power in the drop cables, which could begin activating countless dormant corrosion diodes. The resulting

2nd and 3rd order products could have significant impacts on performance. This phenomenon is referred to as Passive Intermodulation (PIM).

Fortunately, corrosion diodes tend to create impedance mismatches that can be detected and located using some of the previously mentioned PNM analysis techniques. The SCTE Network Operations Subcommittee for PNM has an ongoing workstream to evaluate the effectiveness of these tools for locating potential PIM sources.

4.8.5. Field Tools

The fundamentals of cable haven't changed much since the advent of forward and return path RF, separated and protected by diplex filters. Our protocols have been designed to provide robust and reliable service, but human troubleshooting and maintenance are still required at times. FDX introduces a new level of complexity and data analytics that will become increasingly reliant on automations and algorithms to augment technicians existing tools for troubleshooting and repair. Many of the tried-and-true techniques for isolating service problems will have limited use in a world of bi-directional RF, and new processes will be developed and refined as FDX rolls out

5. Bringing Low Latency DOCSIS to Life

Without latency and jitter improvements, just delivering faster speeds will not guarantee the best customer experience. Low latency is pivotal to residential services such as gaming and video conferencing and are almost always tied into service level agreement (SLA) requirements for commercial services applications. The expectation that low-latency services will become increasingly important also revolves around AR/VR, IoT, and machine-to-machine (M2M) communications. Higher downstream and upstream utilization levels have an impact on latency, jitter and packet loss, that cannot be resolved only by increasing speed. These metrics and measurement techniques, in addition to speed and availability, must be integrated into design and operations and turned into actionable data. Although speed tests have been used as a measure of performance, other test tools tests have not yet been integrated into operations. Recent architecture changes have enabled MSOs to measure latency under load and without load. In this section, we will provide techniques and example performance that can be achieved in today's networks. We will then describe emerging features to deliver lower latency services as part of the 10G roadmap.

Current speed tests first check connection latency. Since the speed tests are performed during relatively idle utilization periods, connection latency measured between the customer router and test server within the MSO's network includes mainly basic DOCSIS media delay, as shown in Figure 43. For idle utilization levels and good network conditions, latency and jitter distributions do not have heavy tails of large values in their histogram.

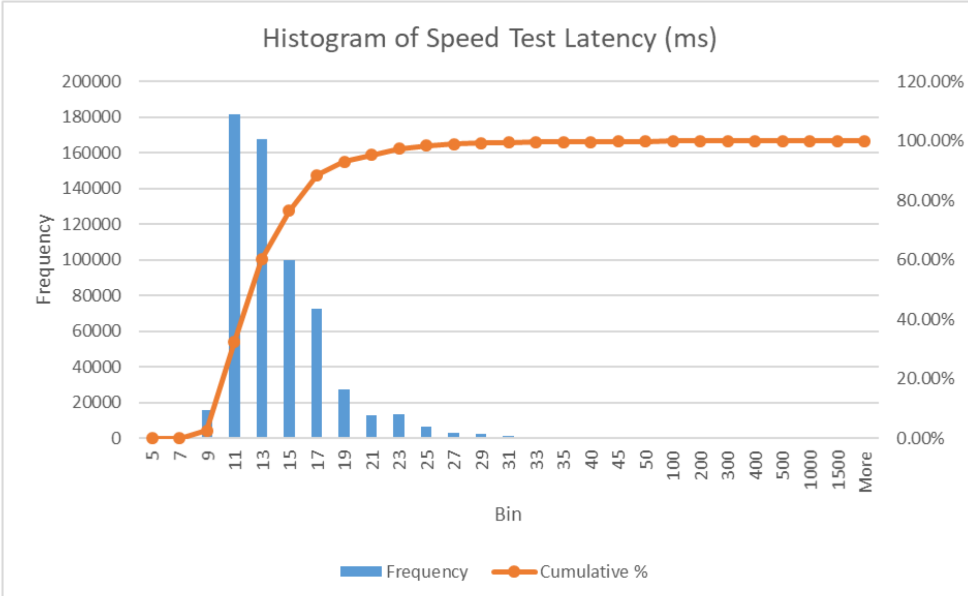

Figure 43 - CDF of RTT Measurements Under Low Utilization Levels for a Mix of Different CM Types and Speed Tiers

However, as shown below in Table 1 (https://cablela.bs/low-latency-docsis-technology-overview-february-2019), queuing is the largest source of delay when the home network is congested. Both downstream (CMTS) and upstream (CM) HSD service flow queues are single queues. Low latency services such as gaming may be delayed due to concurrent queue building (QB) applications, such as file uploads and downloads. It's the networking equivalent of "take a number." Furthermore, unmanaged large queues can lead to the "buffer bloating" effect. Finally, bonding group utilization and network conditions may delay packet transmissions, increasing queueing latency.

Table 1 - Delay Sources in DOCSIS Networks

Delay Source	Upstream	Downstream	Total	Notes
Queuing	0 – 100+ms	0 – 100+ms	0 – 200+ms	Largely caused by TCP
Media Acquisition	2ms – 8ms	0	2ms – 8ms	Request/Grant process w/ 2ms MAP interval
Serialization/Encoding	0.38 – 2.8ms	60µs – 720µs	0.44ms – 3.5ms	Based on Channel Configuration
Propagation	8µs - 300µs	8µs - 300µs	16µs - 600µs	8µs/mile one-way based on furthest CM
Switching/Forwarding	1-20 µs	1-20 µs	2-40 µs	Implementation dependent
TOTAL	**2.4 – 111+ms**	**68µs – 101+ms**	**2.5 – 222+ms**	

Queuing algorithms vary widely among current CMTS and CM models, and DOCSIS versions. Buffer Control in D3.0, and Active Queue Management (AQM) in D3.1, aim to remove buffer bloat. Latency under load (LUL) measurements reflect the maximum latency/jitter a customer may experience under heavy home network utilization. Generally, these tests are done under downstream and upstream load, sequentially. Bi-directional tests are also used to measure the impact of downstream load on the upstream latency and vice versa. The FCC's "Measuring Broadband America" program, as well as third-party ISP rating tools, use LUL as a performance indicator.

If the queue size is not managed, customers may be affected by buffer bloat. An example is shown in Figure 44, which has a large queue setting in the CMTS. Traditional CMTSs have large physical queues. In this example, higher downstream speed tiers have downstream LUL smaller than 500 ms, compared to lower downstream speed tiers that can have ~2 sec maximum downstream LUL (top left) and 1-2 sec mean downstream LUL (bottom left). After optimizing downstream queue parameters, the same network has improved latency without degrading speed test results (top right, max LUL and bottom right, mean LUL).

Note that Figure 44 includes outliers and speed test failures that may be due to test and network issues in the test environment. Figure 44 also shows that although most upstream LUL readings are < 30ms, it may reach >100 ms depending on the upstream speed tier, buffer control, and queue management algorithms in cable modems. For example, CMs with buffer control where a 250ms default target latency is used have upstream LUL of around 250-300ms, much higher than CMs with optimized buffer control or AQM settings.

Figure 44 - Maximum, Minimum and Mean DS and US LUL with (Left) High Target Latency Values for DS AQM and (Right) Optimized Target Latency Values for DS AQM

Figure 45 displays the Cumulative Distribution Function (CDF) for round-trip-time (RTT) measurements between the CPE and measurement within the MSO network under real traffic conditions for a given test case. As expected, the latency values are between the idle latency shown in Figure 43 and the latency under load reading shown in Figure 44. The aggregate values show that the 85th percentile has RTTs of less than 50 ms. Although high speed tiers help to improve packet delay, the delay is not bounded, and packet delay variation (jitter) due to fluctuations in home network utilization, with bursty queue-building traffic, can cause a degraded experience. These bursts may be short lived and may have little impact on aggregate values, but the spikes may have a significant impact on services like gaming, videoconferencing and AR/VR applications.

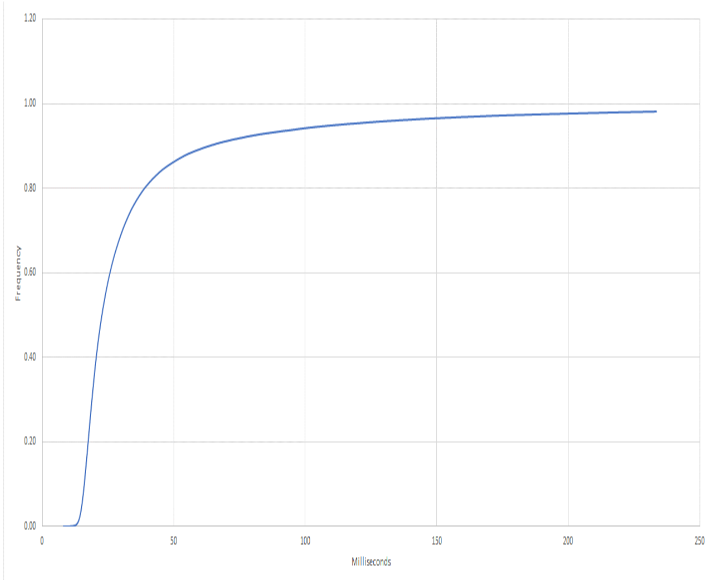

Figure 45 - CDF of Connection Latency Based on TCP Handshake RTT Measurements Under Customer Traffic Load for a Mix of Different CM DOCSIS Versions and Speed Tiers

Additional DOCSIS 3.1 LLD specifications aim to support bounded latency with three main new functionalities:

- Dual queue coupled AQM with queue protection
- Reduced MAP interval
- Proactive Grant Scheduling

Early results show that dual queue AQM can provide <10ms RTT for 95-99th percentile of packets for latency sensitive services. Proactive grant scheduling (PGS) aims to support ~1 ms RTT the 99th percentile of packets for latency sensitive services.

Deployment of a dual queue AQM approach requires other changes in the current architectures, including low latency service classification and markings integrated into operations. In addition to access networks, home (e.g. Wi-Fi) and core networks need to support these features for end-to-end support.

PGS requires more algorithmic development to optimize efficiency gain vs overhead with real traffic and to understand deployment costs and operational models.

6. 10G Security Initiatives

The 10G vision is most often associated with the capacity behind the "10 Gigabit" nomenclature. However, low latency is high among the 10G pillars, driven by gaming applications and more recently with the rapid growth of video conferencing, driven by the pandemic and work-from-home mandates/recommendations. However, low latency has a broad range of residential and commercial applications beyond gaming and Zoom/Teams/pick your most representative video conference app. The explosive growth in M2M and IoT are good examples.

The explosion of IoT devices and data, and other future applications, draw attention to another ever-more-critical area for customers – security. The smart home, security cameras, telemedicine, business data, and industrial networks present services and application use cases with a high premium on data confidentiality and on the overall security of access and communications. Customers are hyper aware of the need, but at the same time don't want to think about it, and just expect it to be there, 24/7.

To consolidate a common 10G security vision across operators, CableLabs has several parallel programs targeting essential elements that look at both new applications as well as emerging technologies for operators' networks [18].

On behalf of operators, CableLabs is participating in key IoT Security forums in the Internet Engineering Task Force (IETF), National Institute of Standards (NIST), and the Open Connectivity Foundation (OCF), among others. Security of IoT services and devices is a common ground across the device and telecommunications ecosystem, and the cable operator perspective is being brought to these forums as standards and best practices are developed.

From a network perspective, the operator edge and access networks are beginning to incorporate DAA, which uses Ethernet-based optical transport, 10GbE today, and distributed compute and networking power in unsecured outside plant. Most operators' 10G vision involves DAA. This introduces a different type of security risk level compared to cable-specific analog optics and "dumb" HFC nodes. In addition, proprietary CCAP platforms – not invulnerable but purpose-built and highly customized systems – are giving way to virtual platforms, consisting of off-the-shelf servers and compute for packet processing, again creating a different category of security threat type. CableLabs' Distributed Virtualization Security is focused on threat assessment and recommendations for these NFV-based DAA systems, including understanding how similar problems are being tacked in other sectors moving in the same direction, such as the wireless and telco industries.

Another innovative program underway at CableLabs in support of 10G is the MicroNets project [19]. As the name may imply, MicroNets apply to the home, and aim to harden the home network against the increasing number of devices with a range of integrity. With applications racing ahead and devices multiplying, bad actors are at an advantage as standards for security mature. MicroNets recognizes this range of trustworthiness or hackability, and the basic premise is to create smaller networks, "micronets," within the traditionally single home LAN, with the walls of division associated with the security level of the devices and service. Issues that arise are isolated to a specific MicroNet, ensuring that as evolving IoT security addresses devices over time, adjacent devices and service are not impacted.

There is much more to come on the security front, as both emerging networking for 10G and target applications act to increase the significance of thinking through security up-front in the design development process.

7. Conclusion

Announced merely 21 months ago, 10G has rapidly become the industry's new North Star. As can be observed herein, the extremely efficient name of the initiative describes quite a wide range of technology elements that it is building upon and building anew. There are many reasons to be looking for the reset button for the year 2020. However, with the direct and critical contributions of our very industry, people and the network are adapting. Timelier than ever, progress towards 10G has not lost a step – indeed it is accelerating, as operators are moving from technology assessments and white board sessions to investment decisions, project definition and development (the fun stuff!). A comparison of 10G material from the 2019 SCTE Expo and the 2020 SCTE Expo shows remarkable progress on many fronts. It only magnifies the excitement for what 2021 will bring, masks or no masks!

Abbreviations

ACI	Adjacent Channel Interference
AI	Artificial Intelligence
CAGR	Compound Annual Growth Rate
CCI	co-channel interference
CLGD	cable load generator
COTS	Commercial off-the-shelf
CP	Cyclic Prefix
CPE	Customer Premises Equipment
CPD	Common Path Distortion
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
EC	Echo Cancellation
FBC	Full Band Capture
FDD	Frequency Domain Duplex
FDX	Full Duplex DOCSIS
FEC	Forward Error Correction
Gbps	Gigabits per second
H-AGG	Hub Aggregation
ICCAP	Integrated Cable Modem Termination System
Igs	Interference Groups
IUC	Interval Usage Codes
LDPC	Low Density Parity Check
LUL	Latency Under Load
M2M	Machine 2 Machine
MER	Modulation Error Ratio
ML	Machine Learning
MMP	Multiple Modulation Profiles
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
PBH	Peak busy hours
PIM	Passive Intermodulation
PMA	Profile Management Application
PNM	Proactive Network Maintenance
QAM	Quadrature Amplitude Modulation
QB	Queue Building
RBAs	Resource Block Assignments
RBW	Resolution Bandwidth
RP	Roll-off Period
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SG	Service Group
SID	Spectral Impairment Detection
SIK	Self-install Kit
SLA	Service Level Agreement
TaFDM	Time and Frequency Division Multiple

TDR	Time Domain Reflectometer
TGs	Transmission Groups
UGS	Unsolicited Grant Services
VOP	Velocity of Propagation
VSA	Vector Signal Analyzer

Bibliography & References

[1] Barker, Bruce E, and Claude Bou Abboud, Erik Neeld, "Access Capacity Planning: Staying Well Ahead of Customer Demand Helped Ensure Stability During COVID-19," SCTE Cable-Tec Expo, Oct 13-16, 2020.

[2] Bienkowski, Tom, "Pandemic Life Online: What is Everybody Doing," www.netscout.com.

[3] Cloonan, Dr Tom, John Ulm, Ayham Al-Banna, Frank O'Keefe, and Ruth Cloonan, "Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years, SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.

[4] Cloonan, Dr. Tom, Ruth Cloonan, and John Ulm, "Updated FDX Traffic Engineering Analysis," CommScope presentation to Comcast, 7/13/20-8/10/20.

[5] Dolan, Andy, Maintaining Confidentiality on the 10G Network https://www.cablelabs.com/maintaining-confidentiality-10g-network

[6] Howald, Dr. Robert L, Comcast, "Aboard the Technology Wave: Surf Report," SCTE Cable-Tec Expo, Sept 26-29, 2016 Philadelphia, PA.

[7] Howald, Dr. Robert L, Comcast, "The Fiber Frontier," INTX Spring Technical Forum, May 16-18, 2016, Boston, MA.

[8] Howald, Dr. Robert L, , "Network Preparation: Maximizing Capacity ROI," SCTE Cable-Tec Expo, Oct 21-24, 2013, Atlanta, GA.

[9] Howald, Dr. Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, "What is 10G – The Technology Foundation," SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.

[10] Mutalik, Venk, D. Rice, K. Subramanya, J. Wang What Gets Measured Gets Done / What Gets Analyzed Gets Transformed, , https://www.nctatechnicalpapers.com/Paper/2018/2018-analytics-for-a-wider-deeper-network-view, SCTE Cable-Tec Expo, Oct 22-25, 2018.

[11] Nandiraju, Nagesh, "Distributed Access Architecture – Goals and Methods of Virtualizing Cable Access," SCTE Cable-Tec Expo, Sept 26-29, 2016 Philadelphia, PA.

[12] Prodan, Dr. Richard, "Optimizing the 10G Transition to Full Duplex DOCSIS 4.0," SCTE Cable-Tec Expo, Oct 13-16, 2020.

[13] Rice, Daniel, M. Harb, J. Ferreria, B. Santangelo, R. Spanbauer "A Machine Learning Pipeline for D3.1 Profile Management," SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.

[14] Rice, Daniel, M. Harb, J. Ferreria, B. Santangelo, R. Spanbauer, "Full Scale Deployment of PMA: Lessons Learned from Deploying the Profile Management Application System at Scale and Plans for Expanding the System Beyond the OFDM Scope", SCTE Cable-Tec Expo, Oct 13-16, 2020.

[15] Thompson, Robert, C. Moore, J. Moran, R. Howald, Optimizing Upstream Throughput Using Equalization Coefficient Analysis, https://www.nctatechnicalpapers.com/Paper/2009/2009-optimizing-upstream-throughput-using-equalization-coefficient-analysis, 2009.

[16] Wall, Dr. Bill, "Practical Considerations For Full Duplex Deployments In N+x Environments", SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA

[17] CableLabs DOCSIS 4.0 PHY Specification: CM-SP-PHYv4.0-I01-190815.docx

[18] https://www.cablelabs.com/10g/security

[19] https://www.cablelabs.com/technologies/micronets

