Hattery Will Get You Everywhere

A technical paper prepared for presentation at SCTE TechExpo24

Dr. Robert Howald
Fellow
Comcast
robert_howald@comcast.com

John Chrostowski
Executive Director
CONNECT Access Engineering
Comcast
john_chrostowski@comcast.com

Chris Marinangeli, Comcast

Dr. Richard Primerano, Comcast

Table of Contents

List of Figures

List of Tables

1. Introduction

The ground-breaking launch of data over cable service interface specification (DOCSIS®) 4.0 full duplex DOCSIS (FDX) technology in October 2023 has created a new set of "X-Class" customers enjoying multi-gigabit symmetrical speeds. However, as fascinating as FDX technology is, and as powerful as X-Class speeds are in enabling immersive new applications, the customer experience is defined by much more than speed. And the network operations experience is defined by much more than the DOCSIS 4.0 network upgrade.

With that in mind, Comcast's 10G team has been as focused on building tools and processes for successful deployment of DOCSIS 4.0 FDX at scale. This paper will focus on two critical automated assessment tools designed, built, and implemented that have been critical to our customer experience and scalability objectives for DOCSIS 4.0 migration.

The first, the Neighbor Home Assessment Test (nHAT), evaluates homes and businesses on an FDX systems for risk of interference from FDX upstream (US) transmissions. The nHAT "score" triggers analysis of deeper layer customer experience metrics, leading to decisions on automated spectrum and radio frequency (RF) power management strategies, or physical options to remediate, if necessary.

The 2nd tool – FDX readiness Home Assessment Test (fHAT) – determines if a customer interested in X-Class services can be activated via a self-install kit (SIK). There are powerful benefits to SIK: customers prefer to not schedule appointments for new services, and Comcast prefers to not roll trucks for device installations that can reliably be remotely activated and provisioned.

Attendees in this session will understand the theory of operation behind these tools, learn how these tools can be applied to DOCSIS 4.0, whether FDX or frequency domain duplex (FDD), see metrics from the FDX installations, and hear about lessons learned. We will discuss roadmap implications for these tools as FDX evolves to deploy over node + x actives (amplifiers) (N+x) systems, and as FDX spectrum and X-Class speeds grow over time.

2. X-Class Launch Update and Roadmap

Comcast launched its DOCSIS 4.0 FDX-based multi-gigabit symmetrical speed services under the name X-Class beginning in October 2023. The initial set of X-Class service tiers are 300/300 Mbps, 500/500 Mbps, 1000/1000 Mbps, and 2000/2000 Mbps. These are branded X-300, X-500, X-1000, and X-2000. FDX-enabled passings will be in the millions for 2024, with multiple vendors providing FDX-capable remote-physical layer (PHY) based digital nodes (RPDs) as upgrades where applicable, or swaps across the existing Digital node footprint. The launches have begun exclusively in the "Fiber Deep" node+0 actives (N+0) footprint. Comcast has a large N+0 footprint from past years of targeted construction of this architecture to build the FDX foundation. Then, as the FDX-capable amplifiers become available for use in production later this year, FDX-enabled passings will accelerate to include both N+0 and N+x.

The high-level plan for FDX enablement is to follow the large-scale virtualization architecture, or distributed access architecture (DAA) mid-split upgrade plan announced in September of 2022[1]. The FDX upgrade is a synergistic complement to these upgrades, as DAA is foundational to DOCSIS 4.0 FDX, and FDX is built upon a mid-split network architecture. Digital node locations will be upgraded with FDX-capable RPDs, and subsequently, the mid-split amplifiers, which have been strategically built on a common and widely deployed bridger and line extender platform, will be upgraded to FDX-capable amplifiers.

[1] https://corporate.comcast.com/press/releases/comcast-expand-evolve-wifi-largest-multi-gigabit-network

One of the powerful benefits of DOCSIS 4.0 FDX is that the speed tiers can be upgraded via software-based configuration, spectrum, and provisioning operations. Whereas DOCSIS 4.0 FDD is based on physical filter choices pre-positioned in nodes and amplifiers or swapped at a future date to accommodate other frequency splits, FDX offers more flexible and granular options. X-Class speeds and the associated FDX spectrum configurations are based on business decisions around X-Class growth and capability, while always maintaining the legacy 85 MHz US for DOCSIS 3.0 and DOCSIS 3.1 traffic.

The broadband roadmap with the launch of DOCSIS 4.0 FDX includes future symmetrical speed increases, such as 3 Gbps / 3 Gbps and 4 Gbps / 4 Gbps, as well as potentially asymmetrical tiers such as 3 Gbps / 1 Gbps, 4 Gbps / 2 Gbps, and 7 Gbps / 5 Gbps. As always, these will be business decisions that consider the proper balance of DOCSIS bandwidth vs quadrature amplitude modulation (QAM) video bandwidth, including the downward trajectory and pace of QAM video and DOCSIS 3.0 services. These make less efficient use of spectrum but serve millions of customers that must be considered with respect to disruption and the entirety of the customer experience. Table 1 shows the trade-off space predicted for various X-Class speeds with respect to management of services and consumer premises equipment (CPE), based on today's traffic utilizations and with a built-in compound annual growth rate (CAGR) used for long-range planning exercises.

Table 1 - Balancing X-Class Spectrum Requirements and Legacy SC-QAM Support

D4.0 Speed Tier		Video QAM	D3.0 QAM
Downstream	Upstream		
3 Gbps	3 Gbps	28	32
3 Gbps	2 Gbps	40	32
4 Gbps	2 Gbps	24	32
4 Gbps	4 Gbps	24	28
5 Gbps	2 Gbps	15	16
5 Gbps	3 Gbps	15	16
5 Gbps	1 Gbps	15	16
5 Gbps	4 Gbps	15	16

Figure 1 shows an example roadmap template of spectrum allocations to deliver particular speed tiers. Different regions have different levels of capacity utilization, QAM video line-ups, "special" channel conditions, and cable modem (CM) distributions. As such, while this template serves as rough guidance for planning purposes, automating the spectrum allocation, taking into account these regional variables, is a high priority initiative. This will eliminate the practice of spectrum management by manual spreadsheet tools in the 10G era, allowing machine learning and automation to deliver optimal spectrum lineups based on local variables.

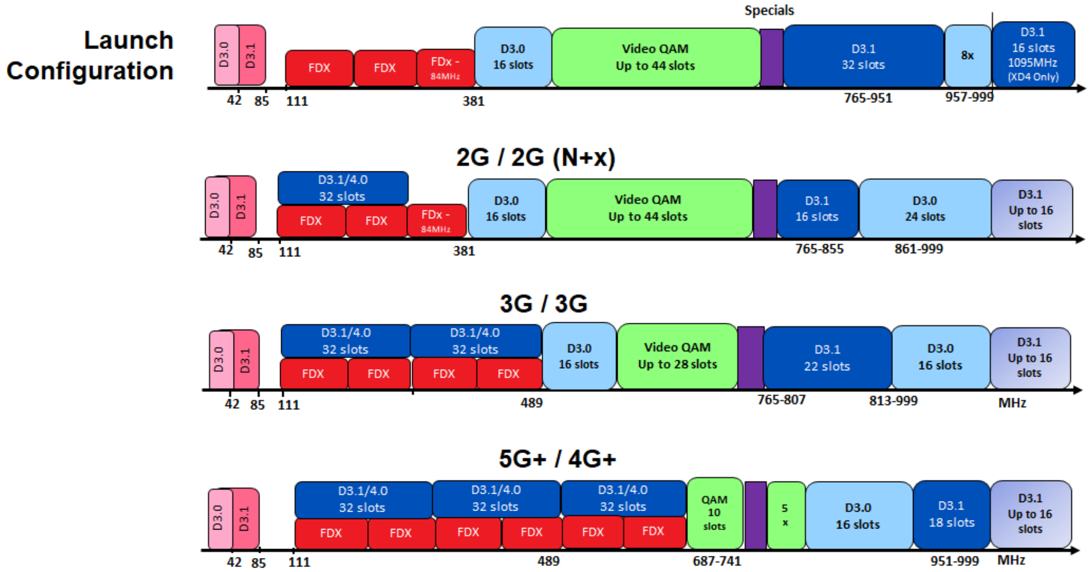

Figure 1 - Spectrum Roadmap Template and X-Class Speeds

3. Operational Challenges of DOCSIS 4.0 Technology

3.1. Neighbor Interference Risk

Adjacent channel interference (ACI) has been well documented for mid-split and high-split systems, and tools are in place to evaluate this interference to mitigate risk. ACI occurs when frequencies enter the CM or set-top box (STB), but do not overlap the downstream (DS) spectrum into the CM or STB. This additional energy is bursty and impacts the automatic gain control (AGC) of the CM or STB, which can cause degraded signal to noise ratio (SNR), modulation error ratio (MER), errors, and video tiling.

In mid-split systems, the risk of ACI is mainly within the same premises as the CM and legacy STBs are co-located. The additional power of the mid-split band is not high enough to cause issues with neighboring devices on different tap ports, or different taps. For high-split systems, the additional power in the high-split band could be noticeable at a neighboring device. For FDX and FDD systems, the additional US bandwidth and power from 108-684 MHz can cause ACI on neighboring devices, both STBs and DOCSIS 3.0 CMs. The additional power of the US with the activation of this extended band can be significantly higher than in mid-split or high-split systems. For example, for a modem transmitting with a flat power spectral density, with the single-carrier QAM (SC-QAMs) transmitting at 45 dBmV per 6.4 MHz, the additive power of the FDX band in the US is 9.6 dBmV, as shown in Table 2.

Table 2 - ACI Power Addition - OFDMA and FDX Channels

	BW(MHz)	Power per Channel Bandwidth			Power per 6.4 MHz SC-QAM
SCQAM1	6.4		45.00	dBmV	45.00
SCQAM2	6.4		45.00	dBmV	45.00
SCQAM3	6.4		45.00	dBmV	45.00
SCQAM4	6.4		45.00	dBmV	45.00
		SC QAM TOTAL POWER (dbmV) =	51.02		dBmV
OFDMA Power	39.0	dBmV Per 1.6 MHz			
OFDMA Bandwidth (MHz)	45.6	OFDMA Channel Power=	53.5		dBmV
		Total Power SC + OFDMA =	55.5		dBmV
FDX 1 Power	39.0	dBmV Per 1.6 Mhz			
FDX 1 Bandwidth (MHz)	96.0	FDX1 Channel Power=	56.8		dBmV
		Total Power SC + OFDMA + FDX1 =	59.2		dBmV
FDX 2 Power	39.0	dBmV Per 1.6 MHz			
FDX2 Bandwidth (MHHz)	96.0	FDX1 Channel Power=	56.8		dBmV
		Total Power SC +OFDMA + FDX1+FDX2 =	61.2		dBmV
FDX 3 Power	39.0	dBmV Per 1.6 MHz			
FDX 3 Bandwidth (MHz)	96.0	FDX3 Channel Power=	56.8		dBmV
		Total Power SC +OFDMA + FDX1+FDX2+FDX3 =	62.5		dBmV
FDX 4 Power	39.0	dBmV Per 1.6 MHz			
FDX 4 Bandwidth (MHz)	96.0	FDX3 Channel Power=	56.8		dBmV
		Total Power SC +OFDMA + FDX1+FDX2+FDX3+FDX4 =	63.5		dBmV
FDX 5 Power	39.0	dBmV Per 1.6 MHz			
FDX 5 Bandwidth (MHz)	96.0	FDX5 Channel Power=	56.8		dBmV
		Total Power SC +OFDMA + FDX1+FDX2+FDX3+ FDX4 + FDX 5 =	64.4		dBmV
FDX 6 Power	39.0	dBmV Per 1.6 MHz			
FDX 6 Bandwidth (MHz)	96.0	FDX6 Channel Power=	56.8		dBmV
		Total Power SC +OFDMA + FDX1+FDX2+FDX3 +FDX4+FDX5+FDX6	65.1		dBmV
		Total Power Increase With FDX Channels	9.6		dBmV

In a hybrid-fiber coax (HFC) network, there are signal paths from port-to-port on the same tap and from port-to-port on adjacent taps. The isolation across these ports is high, but with an FDX or FDD CM transmitting significant additional power in the downstream band of legacy devices, the signal at neighboring ports and taps may be high enough to interfere with these devices. Referring to Figure 2, potential interference paths exist across the tap and to neighboring taps. When FDX or FDD systems are deployed and this US spectrum is activated, these interference level needs to be evaluated. Using the transmit power of the DOCSIS 4.0 CM and the topology of the plant, the interfering signal level hitting the neighbor home can be calculated.

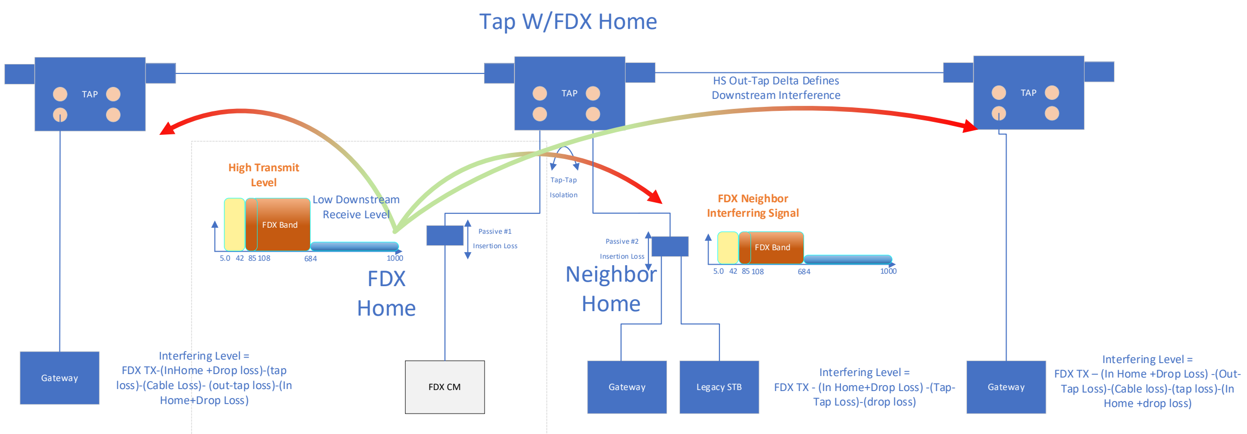

Figure 2 - FDX (or FDD) Neighbor Interference Paths

Table 1 showed the additive transmit power of the FDX band with a flat transmit power spectral density. The actual transmit profile needs to overcome the path loss of the plant and will have uptilt. The DOCSIS 4.0 spec for the CM allows for up to 12 dB of uptilt in the FDX band. Table 3 shows an example with the FDX modem transmitting with a more realistic 8 dB of uptilt. The signal level at the neighbor home, neighbor interfering receive (Rx) power, can then be calculated as follows.

Neighbor Interfering RX power =

FDX transmit *(TX) power-FDX customer drop loss – tap to tap isolation – neighbor drop loss – neighbor in-home loss*

Table 3 - FDX Neighbor Interference Example Calculation

FDX Band (MHz)	FDX Transmit Power (dBmV)	FDX Customer Drop Loss (50 ft RG-6) dB	Tap-tap isolation (dB)	Neighbor Drop Loss (50 ft RG-6) dB	Neighbor in-home loss (dB)	Neighbor RX power (dBmV)
108-204	51.3	1.2	25	1.2	4	19.9
204-300	53.7	1.6	25	1.6	4	21.5
300-396	55.5	1.9	25	1.9	4	22.7
300-492	56.9	2.2	25	2.2	4	23.5
492-588	57.9	2.4	25	2.4	4	24.1
588-684	59.0	2.6	25	2.6	4	24.8
TOTAL	64.2					30.8

With a DS band from 804-1218 MHz and a DS input level of 0 dBmV / 6 MHz, the total composite power (TCP) of the DS band equals 18.38 dBmV. The TCP of the FDX interfering signal of 30.8 dB is shown in Table 3.

The TCP interfering signal delta equals FDX interference TCP (30.8 dBmV) at the neighbor minus the DS TCP (18.4 dBmV). This gives a TCP interference delta of 12.4 dB. As will be seen in the test data presented, this TCP delta has the potential to cause interference and thus may need to be mitigated.

3.2. Home Environment

The term HFC network tends to refer to the outdoor fiber and coaxial network that connects hubs and headends to homes, including the drop network. Unfortunately, the complete HFC network extends into the home coax, until it lands on a CPE device such as a STB or a CM. From a customer perspective, their experience extends to, and is typically dominated by, the Wi-Fi network. For both the in-home coaxial and wireless connections, these parts of the network are mostly outside the operator's control. There are, of course, opportunities to exert some control over it during visits to residential and business locations. However, though good RF practices may be put in place or restored during such times, this can change since customers have access to cables, outlets, and devices connected to them. Operators know through experience that home network practices of a customer often are a root cause of service issues.

The use of extended US frequencies for both DOCSIS 4.0 FDX and DOCSIS 4.0 FDD, combined with a fixed available CM total available US transmit power, increases the challenge in the US for FDX or FDD. For FDX, and like the node, the FDX CM includes a powerful echo canceller (EC). The EC operation at the CM differs from the node in one important way. At the node, there is simultaneous DS and US in both time and frequency. At the CM, there is DS and US spectrum sharing, but never simultaneously in time. Figure 1 showed the nature of the spectrum overlap feature from the node perspective, not from the CM perspective. Figure 5, as denoted by the complementary colored FDX CMs in different transmission groups (TGs), shows a view from a CM's perspective of FDX frequency allocation.

Because of this, the CM EC only needs to provide cancellation of co-channel interfering (CCI) noise floor from the US power amplifier (PA), and any adjacent channel spectrum leakage from one resource block assignment (RBA) – one US Tx orthogonal frequency-division multiple access (OFDMA) block – into another DS Rx block. With this capability, the EC at the CM can isolate DS and US transmissions to ensure a high-performance DS even as a DOCSIS 4.0 FDX CM transmits US without a filter between the two. How effective the EC is in delivering a high-fidelity DS is related to:

- CM US Tx level and fidelity,
- CM DS Rx level,
- EC capability, and
- the echo environment of the home

As part of FDX system optimization, a range of home environments are built in the lab and characterized by the nature of their reflection environment. And a set of CM EC test cases are defined for characterizing the performance of different combinations of RBA settings for these different home environments. Figure 3 and Figure 4 show the test case scenarios and test system respectively.

FDX Band Allocation	1) 108-396 MHz – 3 FDX Sub-Bands of 96 MHz each 2) 108-492 MHz – 2 FDX Sub-Bands of 192 MHz each 3) 108-684 MHz – 3 FDX Sub-Bands of 192 MHz each		
RBA Settings 288 MHz (108-396 MHz) 384 MHz (108-492 MHz) 576 MHz (108-684 MHz)	00 (DD) 01 (DU) 10 (UD)	000 (DDD) 011 (DUU) 101 (UDU) 110 (UUD)	000 (DDD) 011 (DUU) 101 (UDU) 110 (UUD)
Reflection at CM Port	24 dB 18 dB 12 dB		

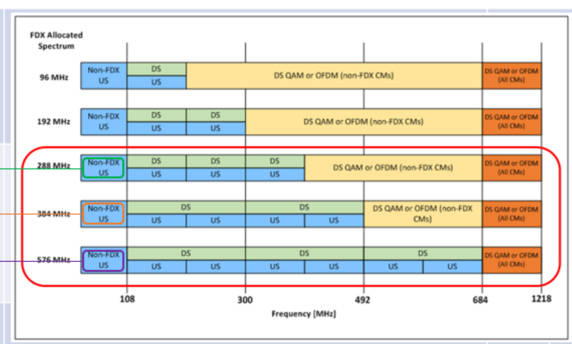

Figure 3 - FDX CM Echo Cancellation Test Cases

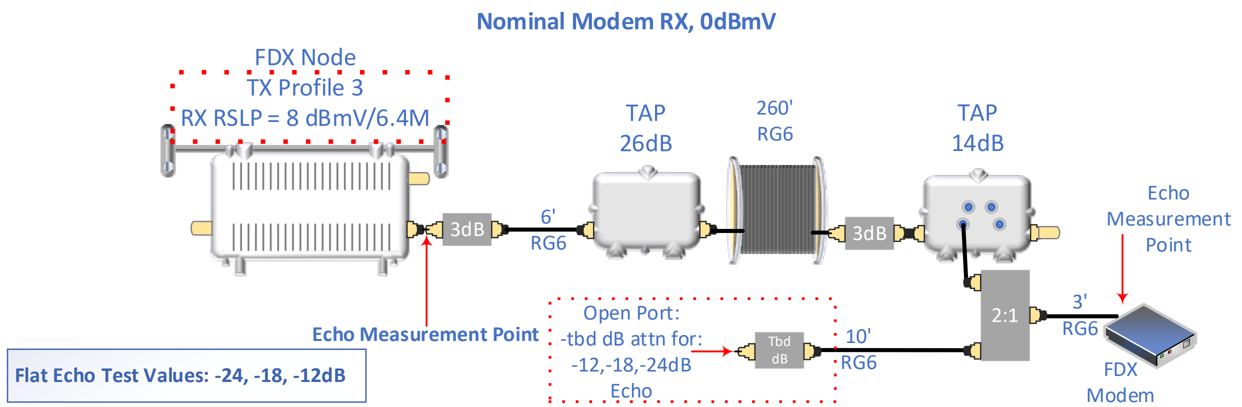

Figure 4 - FDX CM Echo Cancellation Test System (showing DS Rx level of 0 dBmv/6 MHz – also tested @ -5 dBmv, +5 dBmv, +10 dBmv)

Note that in all fidelity cases, the CM US Tx is set to the maximum total transmission power possible. Although most devices with not transmit at their peak, some will, justifying the need to validate this scenario and performance. Lastly, expectation for future versions of the profile management application (PMA) is to allow US CM Tx and US RPD Rx to be variables to be optimized. This will result in more CMs transmitting close or at their maximum TCP to achieve the highest FDX band US MER.

With the aforementioned challenges of higher US losses of the extended FDX (or FDD) upstream, and the significance of the home environment to achieving high performance, the common use of SIKs, also known as get started kits (GSK), must be considered. Customers of Comcast enjoy over 80% success rate with the self-install process today, avoiding the need to schedule technician visits. The DOCSIS 4.0 standard, as a reference architecture, assumes a point-of-entry (PoE), single-device termination at the residence or business. To facilitate this, installations by trained technicians are required. The cost of this comes in the form of both customer inconvenience, as well as the direct cost associated with a truck roll (TR) being required for each customer desiring DOCSIS 4.0-based services.

Rather than accept a "pro-install" via trained techs as an inevitability, Comcast has developed a tool that enables a pre-assessment of a home for FDX readiness, and an accompanying back-end check after installation to validate performance when the activation was via self-install. Of course, in a "pro-install," a technician will validate the installation onsite. This tool, fHAT, will be described in a subsequent section.

Lastly, note that DOCSIS 4.0 FDD using CMs that receive in the DS up to 1.8 GHz *cannot* benefit from fHAT or SIK. This is due to the extended DS frequencies being beyond the bandwidth of deployed in-home splitters, whereas FDX-based systems utilize the same bandwidth of DOCSIS 3.1 systems. The maximum DS frequency is defined as 1218 MHz, and capacity achieved by using this spectrum more efficiently via the bi-directionality feature of FDX. FDD systems, in addition to the need to replace or upgrade all actives and passives to achieve 1.8 GHz, have a significant operational and cost penalty associated with premise installation. This is due to the requirement for a technician to visit every DOCSIS 4.0 FDD premise when the premise has cable band splitters and the services delivered require use of the spectrum in the extended downstream band.

3.3. FDX Transmission Groups and Relationship to nHAT

As described previously and shown in Figure 5, the CM perspective on RBA allocation is different than at the node. An individual CM has a spectrum block defined as either DS or US, but never at once.

However, the node, as also shown in Figure 1, has simultaneous DS and US in the same spectrum blocks, or RBAs. To complete the description of FDX that makes clear the operation, consider multiple CMs, which can have alternative and complementary RBA assignments, such that the DS and US in each is fully utilized from the perspective of the CM endpoints as well. This is shown in Figure 5.

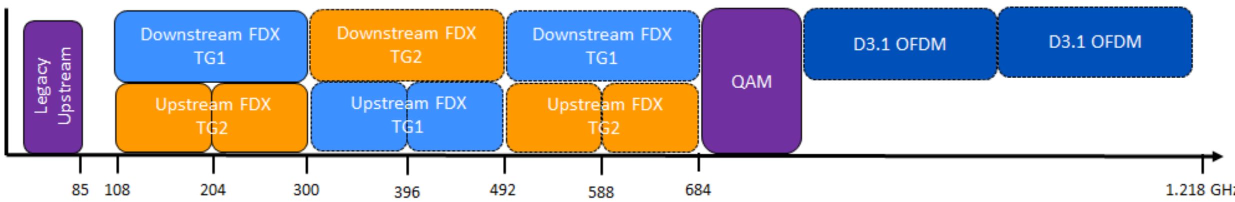

Figure 5 - Complementary Use of RBAs by FDX CMs Assigned Transmission Groups (TG1, TG2)

The terminology of FDX that describes the phenomenon whereby one set of CMs use one RBA assignment, while another uses its complement (or extending this to additional sets of complementary use > 2) is known as interference groups (IGs) and transmission groups (TGs). Just as a single CM never transits and receives simultaneously in the same bandwidth, this also applies to a set of modems that are "close" in RF distance to one another with respect to isolation. An FDX system determines which CMs might interfere with one another if they were to transmit and receive in the same band, and identifies them as an IG, for which they are managed together with respect to RBA assignment. A TG is a logical aggregation of IGs, since every IG defined by RF does not necessarily need its own scheduled RBA domain; this is a function of traffic and utilization.

We mention the concept of IG and TG here because it is often inaccurately viewed as a serious limitation of FDX because of how devices may interfere with one another. Substantial analysis and empirical evidence now prove that these concerns were unfounded, and that FDX band traffic management relies on oversubscription and capacity analytics virtually identical in concept and result as non-FDX traffic management [4]. What the analysis shows instead is that the FDD inefficiently uses dedicated US spectrum, which sits idle most of the time.

Comcast has built a tool known as the virtual service gateway (vSG) which granularly monitors traffic, providing views of short-term peak bursts as well as long term aggregate views. Figure 6 shows data observations for both DS and US for peak burst extremes, which aligns well with the behaviors as predicted in [**4**].

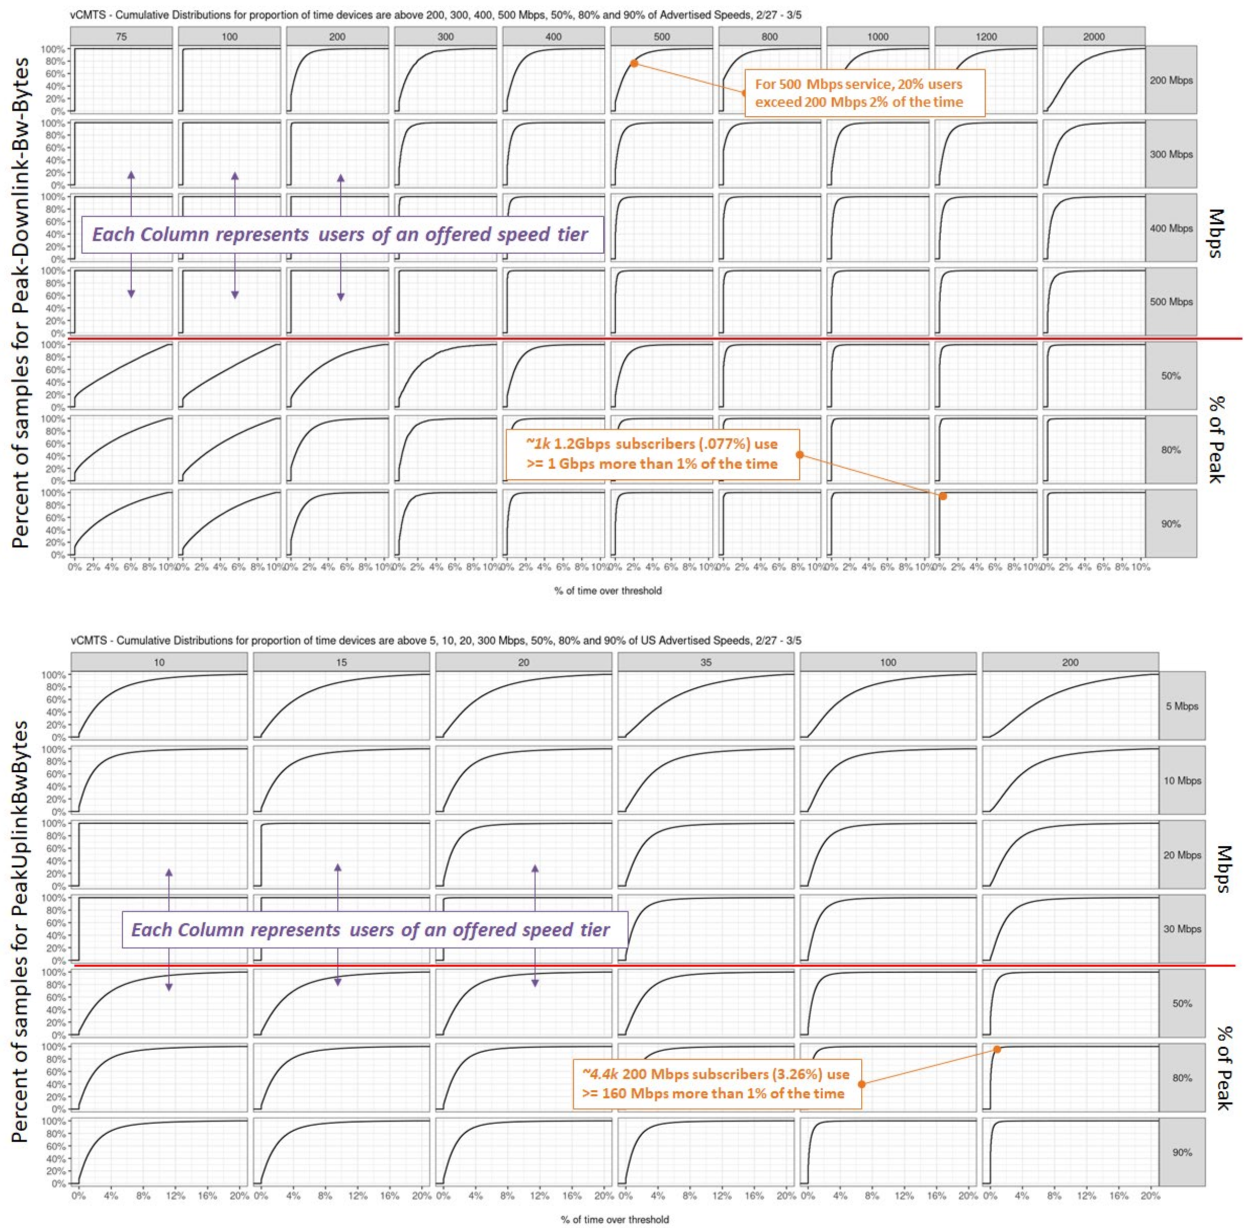

Figure 6 - Actual Traffic Utilization Confirms the Rarity of Peak Speed Bursts

A wide spread of speed tiers (columns) is shown in Figure 6, and for each the percentage of the devices' proportional time spent above a given threshold of absolute Mbps, or as a % of their speed tier in the lower three rows. The latter is done because the raw Mbps range used can exceed the speed tiers in some cases. Therefore, those users can never exceed the Mbps, and the cumulative distribution function has a "step function" look to it (0 users exceed the max speed) accordingly. This pattern can be seen, for example, in columns one and two, rows three and four. The % rows at the bottom remove this artifact.

Some stateable conclusions can be drawn from this analysis that clearly tell the story:

For a sample size of >1 million subscribers with 1.2 Gbps DS speed

- 23 (.002%) use ≥ 1 Gbps more than 10% of the time

- 1000 (.077%) use ≥ 1 Gbps or more 1% of the time

With respect to US speeds, based on the 200 Mbps peak tier (135k subscribers at time of measurements)

- 88 (.065%) use ≥ 160 Mbps more than 10% of the time
- 4400 (3.26%) use ≥ 160 Mbps more than 1% of the time

And of course, whether DS or US, the probability of *two users* accessing peak speeds is even more rare; and the impact, should that occur, at the application level will almost always be non-customer impacting.

Perhaps equally valuable, the IG/TG mechanism in FDX is derived by the process of "sounding" as described in the DOCSIS 4.0 FDX standard [7], paragraph 7.6. The nHAT operational tool that will be described further is, effectively, operating as a network sounding tool. By implementing nHAT with the introduction of FDX subscribers, nHAT is, in effect, building the IG/TG groups outside of the complex sounding process described in the standard, with each FDX customer installation.

Through the implementation of the nHAT tool, and the retention, updating, and integration with FDX back-office operations, IG/TG relationships become defined and can be made available for DOCSIS 4.0 schedulers, without requiring new virtual cable modem termination (vCMTS) features to accomplish this task. In fact, nHAT goes one step further than sounding, and is an improvement on the process. FDX sounding is about identifying IG/TGs for other connected DOCSIS 3.1 modems only. The nHAT tool recognizes the interference risk of video STBs and DOCSIS 3.0 devices, which is something that the FDX sounding operation cannot accomplish.

4. New Tools for FDX Operationalization

4.1. Neighbor Home Assessment Test (nHAT)

4.1.1. Prior Analysis and Testing Applied to FDX

To understand the potential for interference, the ACI total interference power thresholds for legacy devices need to be determined for the extended US bandwidth. Interference testing was previously completed for various STB and CMs to determine the TCP delta that would cause tiling on neighbor STB or cause forward error correction (FEC) errors on legacy neighbor cable modems. This testing was completed with a signal generator simulating an FDX US transmission for various channel maps with various duty cycles and periods.

The channel maps tested are based on Figure 1**,** shown in Figure 7. The modifications in rows 2 and 3 versus Figure 1 present a more conservative test case than those same rows in Figure 1**,** and the incremental changes allow a coarse sensitivity analysis and rules-based extrapolation to further bandwidth extension.

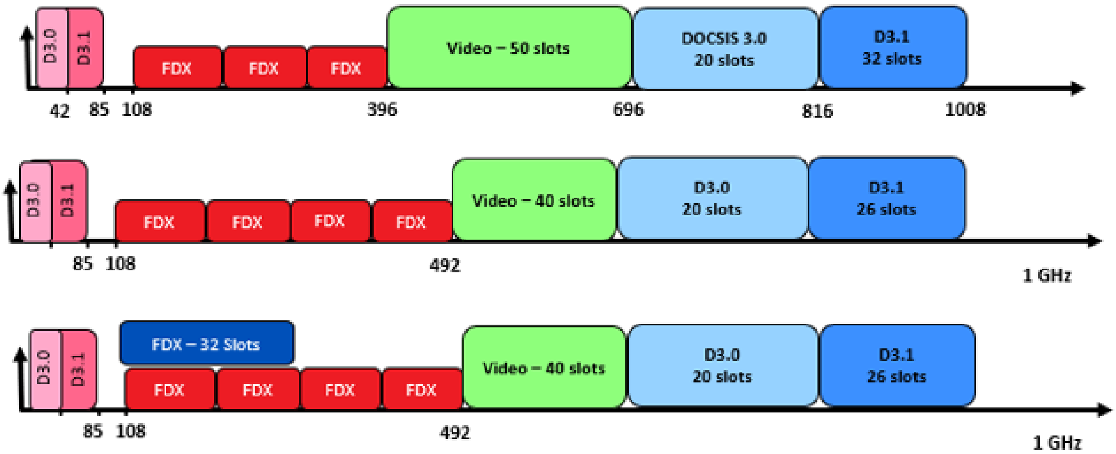

Figure 7 - Channel Map Configurations Used for Neighbor Interference Testing

Video interference testing was performed with a simulated FDX US transmission with varying duty cycles and periods to simulate a “bursty” US signal. The signal level of the interference was varied while physically watching the video content for tiling. At the point of tiling the TCP delta threshold was recorded. A sample of the data for the thresholds of interfering power which caused tiling for STBs for video at 495 MHz is shown in Table 4. The minimum ACI TCP delta is 4.5 dB across all devices and frequencies.

During neighbor interference testing, multiple video channels at various frequencies were measured. Table 5 shows the summary data for TCP delta threshold for video tiling for all the video channels tested for channel map 2. The worst-case performance is not always at the lowest frequency (closest to the ACI signals)! This is an indication that, in addition to the AGC of the device not being able to track bursty interference, distortion from the STB front end is also a contributor to ACI. The lowest TCP delta threshold across all video channels tested is 4.3 dB.

Table 4 - Video Tiling TCP Delta Threshold for Map 2 and DS Video at 495 MHz

Period	10 ms			70 ms			200 ms		
Duty Cycle	10%	50%	90%	10%	50%	90%	10%	50%	90%
STB Model	TCP Delta (dB)								
Model A	14.9	14.1	13.2	9.5	10.1	8.4	10.3	9.7	11
Model B	14.9	14.1	13.2	6.6	8.2	8.4	8.3	10.7	11
Model C	15.7	14.1	14.1	8.3	14.1	14.3	8.3	8.6	12
Model D	13.5	13	13.2	4.5	5.2	8.4	5.4	5.9	6.1
Model E	4.4	7.1	12	4.5	5.2	12.5	5.4	5.9	8.1

Table 5 - Video Tiling TCP Delta Thresholds for Map 2 for Various Video Channels

Video Ch Freq	495 MHz	549 MHz	651 MHz	729 MHz
STB Model	TCP Delta (dB)			
Model A	9.5	7.4	9.8	9.6
Model B	6.6	7.4	8.7	8.4
Model C	8.3	8.3	10.8	10.9
Model D	4.5	5	5.7	6.4
Model E	4.5	4.3	5.9	5.5

Similarly, interference data was taken with DOCSIS 3.0 cable modems being the “victim” device. Like the case for STBs, this was done with a signal generator simulating FDX signals with varying duty cycles and periods. The TCP delta threshold was recorded at the point where uncorrectable errors were noted on the target device. A summary of this data for all duty cycles and periods is show in Table 6. The worst-case TCP delta is -1.23 dB, and a value of -0.23 dB is the lower threshold for the majority of the cable modems tested.

Table 6 - TCP Delta Threshold – DOCSIS 3.0 CMs

Calculated FDX TCP (dBmV)	Calculated Legacy TCP (dBmV)	TCP Diff. (dB)	CM 1	CM 2	CM 3	CM 4	CM 5	CM 6	CM 7	CM 8
23.81	18.04	5.77	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL	FAIL
22.81	18.04	4.77	FAIL	FAIL	FAIL	FAIL	FAIL	PASS	FAIL	FAIL
21.81	18.04	3.77	FAIL	FAIL	FAIL	FAIL	FAIL	PASS	FAIL	FAIL
20.81	18.04	2.77	FAIL	FAIL	FAIL	FAIL	FAIL	PASS	FAIL	FAIL
19.81	18.04	1.77	FAIL	FAIL	PASS	FAIL	FAIL	PASS	FAIL	PASS
18.81	18.04	0.77	FAIL	FAIL	PASS	PASS	FAIL	PASS	FAIL	PASS
17.81	18.04	-0.23	FAIL	FAIL	PASS	PASS	FAIL	PASS	FAIL	PASS
16.81	18.04	-1.23	PASS	PASS	PASS	PASS	FAIL	PASS	PASS	PASS
15.81	18.04	-2.23	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS
14.81	18.04	-3.23	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS

With the continued development of the FDX network and devices, interference testing has continued with a complete FDX system, including FDX cable modems under FDX-enabled vCMTS control. Interference testing is being conducted on STBs, DOCSIS 3.0 cable modems and DOCSIS 3.1 cable modems without FDX-limited (FDX-L). FDX-L is a software upgrade to DOCSIS 3.1 CMs that make them aware that they are in a DOCSIS 3.1 system, ensuring they will never be receiving data when there is a potential interferer, and providing them with the awareness to protect themselves from ACI overload.

Previous testing was performed with a signal generator with varying duty cycles and periods to simulate US traffic dynamics. Interference testing with a complete FDX system is performed with different US traffic from 10 Mbps to 2000 Mbps. As the US traffic on the FDX device is varied, the US channel utilization varies, in both time and frequency, which will change the TCP level at which interference occurs.

On an FDX configured RPD, FDX, DOCSIS 3.0 and DOCSIS 3.1 cable modems are connected and provisioned. A traffic generator is connected to the FDX cable modem for the US interfering signal, and the target DOCSIS 3.0 and DOCSIS 3.1 cable modems monitored for interference and FEC errors. In addition to the normal splitter or tap port-port interference path, an interference path has been added from an FDX cable modem using splitters at the FDX modem, and target devices to inject the US of the FDX cable modem into DOCSIS 3.0 and DOCSIS 3.1 devices. This interference level can be varied by adjusting programmable attenuators 1, 2, and 3, as shown in Figure 8.

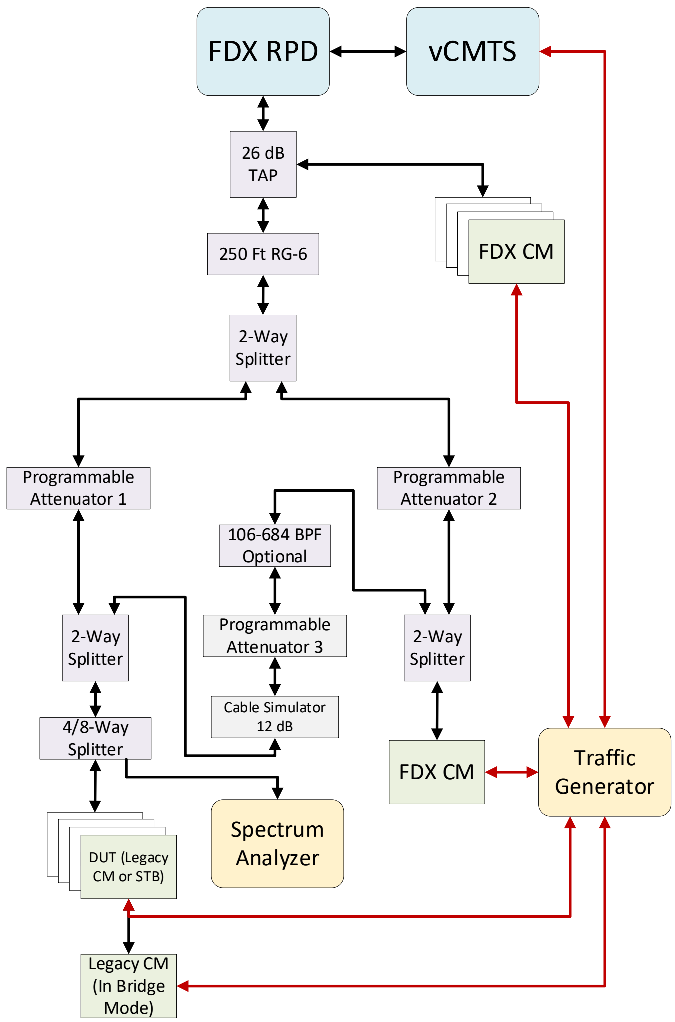

Figure 8 - FDX ACI TCP Delta Threshold Test Configuration

Speed tests are run at varying traffic rates, which in effect varies the duty cycle and occupied US bandwidth of the FDX cable modem transmission. At each traffic rate, the US signal is captured, and a histogram is created showing the counts of TCP power distribution during the capture. At low data rates, the counts of high transmit power are low; and as the data rate increases, the counts of high transmit power increase, as expected. Figure 9**,** Figure 10, and Figure 11shows histograms with 10, 100, 500, 1000, and 2000 Mbps throughput.

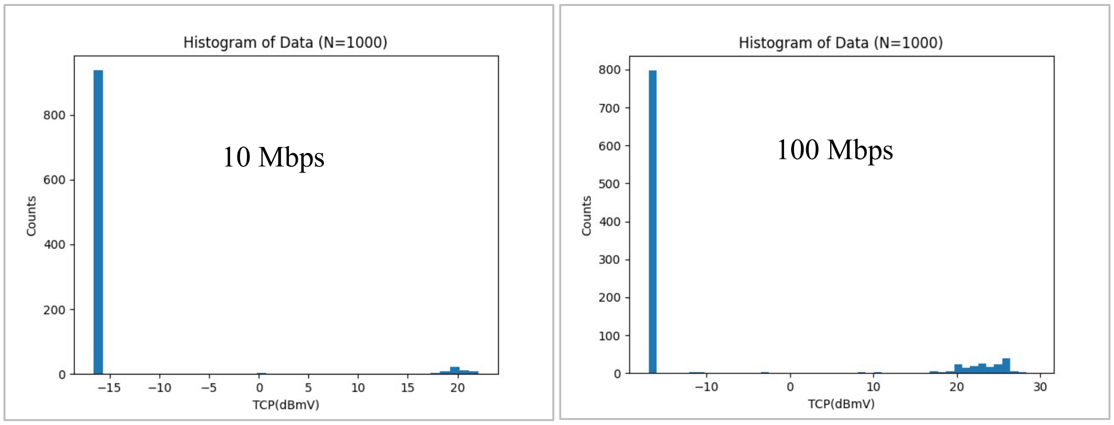

Figure 9 - FDX Cable Modem Transmit Histogram for 10 Mbps and 100 Mbps

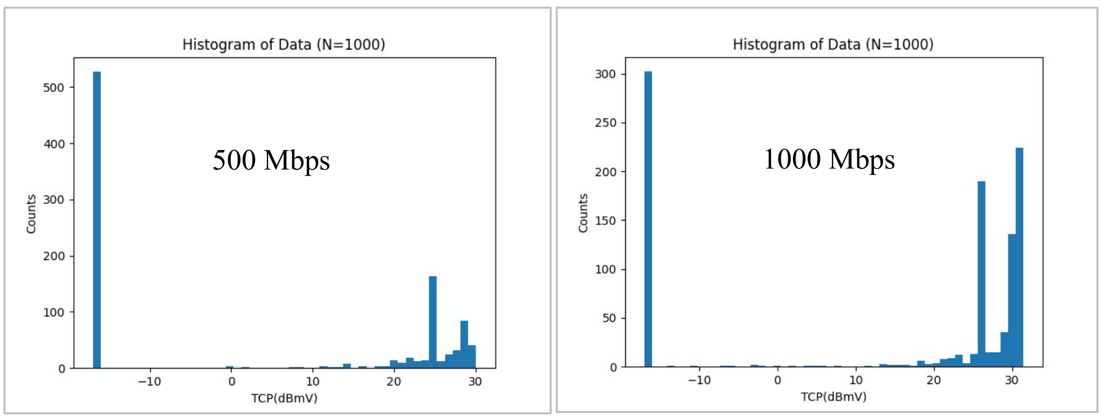

Figure 10 - FDX Cable Modem Transmit Histogram for 500 Mbps and 1000 Mbps

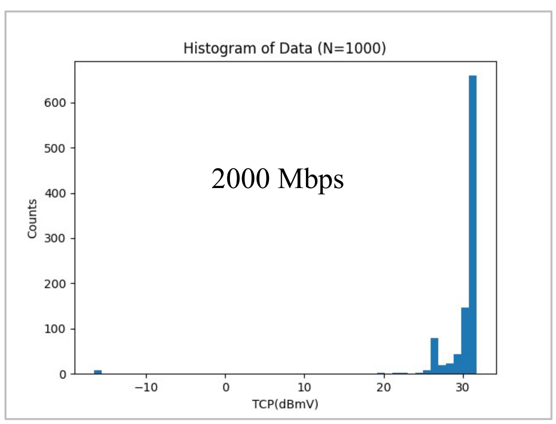

Figure 11 - FDX Cable Modem Transmit Histogram for 2000 Mbps

For the DOCSIS 3.0 and DOCSIS 3.1 cable modem case, the interfering signal from the FDX cable modem was increased until post-FEC errors occurred. Data showing the neighbor ACI TCP interference thresholds for channel map 1 is shown in Table 7. Data shows the lowest ACI TCP delta which causes interference to be around 0 dB at both 10 Mbps and 1200 Mbps. At 10 Mbps, the histogram shows very low counts with high TCP. Achieving the TCP to cause errors with these low counts is not feasible. Data rates of 1000 Mbps and 1200 Mbps also show 0 dB TCP delta. The histograms show significant counts at high TCP with this threshold. 0 dB TCP delta correlates with the previous testing completed with the signal generator simulating FDX traffic which had a worst-case TCP delta of 0.23 dB.

Table 7 - TCP Threshold delta for DOCSIS 3.0 and DOCSIS 3.1 Cable Modems (Post FEC Errors)

	Interference Threshold TCP Delta (dB)		
US Throughput (Mbps)	**DOCSIS 3.0 Cable Modem**	**DOCSIS 3.1 Cable Modem #1**	**DOCSIS 3.1 Cable Modem #2**
10	-0.5	-2.80	-2.80
50	1.5	-3.7	-0.56
100	2.7	-0.7	1.32
300	4.0	1.23	0.23
500	3.2	-0.71	1.29
800	1.1	-0.14	0.14
1000	.07	-1.92	-1.92
1200	-.03	0.17	-0.86
1500	3.7	1.17	0.16
2000	2.7	-0.51	-0.01

Video impairment testing was performed at two video frequencies: 483 MHz, and 687 MHz. Results show a worst case TCP delta (minimum value) of 5.19 dB at 10 Mbps data throughput. At 10 Mbps, the histogram shows very low counts with high TCP. To achieve this total power at 10 Mbps would require a very high cable modem transmit power, which is not realistic. Ignoring this low data rate, the next values, which most predominantly cause interference, are 5 to 6 dB TCP delta. This is close to the minimum value of 4.5 dB from the previous testing with the signal generator simulating the US traffic. See Table 8 and Table 9.

Table 8 - TCP Threshold delta for Legacy Video Set-top boxes at 483 MHz

	STB1		STB2		STB3		STB4		STB5		STB6	
US FDX bitrate (Mbps)	TCP Delta US to DS (dBmV)	Status	TCP Delta US to DS (dBmV)	Status	TCP Delta US to DS (dBmV)	Status	TCP Delta US to DS (dBmV)	Status	TCP Delta US to DS (dBmV)	Status	TCP Delta US to DS (dBmV)	Status
10	8.12	Not Tiling	8.12	Tiling	4.7	Tiling	4.66	Tiling	8.12	Not Tiling	8.12	Not Tiling
50	14.11	Not Tiling	8.86	Tiling	5.0	Tiling	5.04	Tiling	14.11	Not Tiling	14.11	Tiling
100	15.28	Tiling	9.33	Tiling	6.3	Tiling	6.32	Tiling	14.23	Tiling	14.23	Tiling
300	15.03	Tiling	9.04	Tiling	6.0	Tiling	6.04	Tiling	13.98	Tiling	13.98	Tiling
500	16.07	Tiling	10.66	Tiling	6.0	Tiling	5.98	Tiling	14.67	Tiling	15.63	Tiling
800	16.99	Tiling	13.41	Tiling	6.4	Tiling	6.35	Tiling	15.00	Tiling	16.51	Tiling
1000	17.46	Tiling	13.87	Tiling	7.0	Tiling	6.96	Tiling	14.97	Tiling	16.46	Tiling
1200	16.90	Tiling	15.92	Tiling	8.0	Tiling	8.01	Tiling	14.93	Tiling	15.92	Tiling
1500	17.43	Tiling	15.46	Tiling	9.5	Tiling	10.93	Tiling	14.96	Tiling	16.49	Tiling
2000	17.75	Not Tiling	17.31	Tiling	14.2	Tiling	15.74	Tiling	14.71	Tiling	17.31	Tiling
2500	17.92	Not Tiling	17.48	Tiling	10.0	Tiling	9.96	Tiling	14.99	Tiling	17.48	Tiling

Table 9 - TCP Threshold delta for Legacy Video Set-top boxes at 687 MHz

	STB1		STB2		STB3		STB4		STB5		STB6	
US FDX bitrate (Mbps)	TCP Delta US to DS (dBc)	Status	TCP Delta US to DS (dBc)	Status	TCP Delta US to DS (dBc)	Status	TCP Delta US to DS (dBc)	Status	TCP Delta US to DS (dBc)	Status	TCP Delta US to DS (dBc)	Status
10	7.77	Not Tiling	6.53	Tiling	5.19	Tiling	5.19	Tiling	7.77	Not Tiling	7.77	Not Tiling
50	14.87	Not Tiling	9.06	Tiling	5.77	Tiling	7.22	Tiling	14.87	Not Tiling	14.87	Not Tiling
100	16.53	Not Tiling	9.70	Tiling	6.74	Tiling	6.74	Tiling	16.53	Not Tiling	16.53	Not Tiling
300	18.27	Not Tiling	9.35	Tiling	6.36	Tiling	7.01	Tiling	18.27	Not Tiling	18.27	Tiling
500	17.95	Not Tiling	10.55	Tiling	7.07	Tiling	7.07	Tiling	17.95	Not Tiling	17.95	Not Tiling
800	17.88	Not Tiling	13.91	Tiling	7.95	Tiling	7.95	Tiling	17.88	Not Tiling	17.88	Not Tiling
1000	17.91	Not Tiling	15.93	Tiling	7.99	Tiling	7.99	Tiling	17.91	Not Tiling	17.91	Not Tiling
1200	17.87	Not Tiling	17.87	Not Tiling	10.99	Tiling	10.99	Tiling	17.87	Not Tiling	17.87	Not Tiling
1500	17.78	Not Tiling	17.78	Not Tiling	12.78	Tiling	12.78	Tiling	17.78	Not Tiling	17.78	Not Tiling
2000	17.71	Not Tiling	17.71	Not Tiling	15.75	Tiling	17.71	Not Tiling	17.71	Not Tiling	17.71	Not Tiling
2500	17.65	Not Tiling	17.65	Not Tiling	12.26	Tiling	12.26	Tiling	17.65	Not Tiling	17.65	Not Tiling

4.1.2. nHAT Field Results of X-Class Launches

The nHAT tool has been available and used for X-Class launches since Day 1, albeit not scalable operationally. A manually triggered test was built, with an eyes-on-glass engineer looking at results on a dashboard, synchronized with scheduled dates and times of installations in the sales funnel. Subsequently, automation and field processes are being put in place (writing as of May 2024) for scalability to support activation of hundreds of nodes per week.

For the first months of activation, which were initially low-scale deployments, FDX operations were supported by key engineering leads and the Comcast "incubation" resources, who are critical to transitioning new technology solutions out of engineering, through trials and deployment. The engineers who were focused on nHAT relied on running manual scripts and reviewing results on the dashboard shown in Table 10 to assess the risk profile.

Table 10 - nHAT Dashboard View: Results Summary, Tested Devices, RF Isolation Measurement Distribution

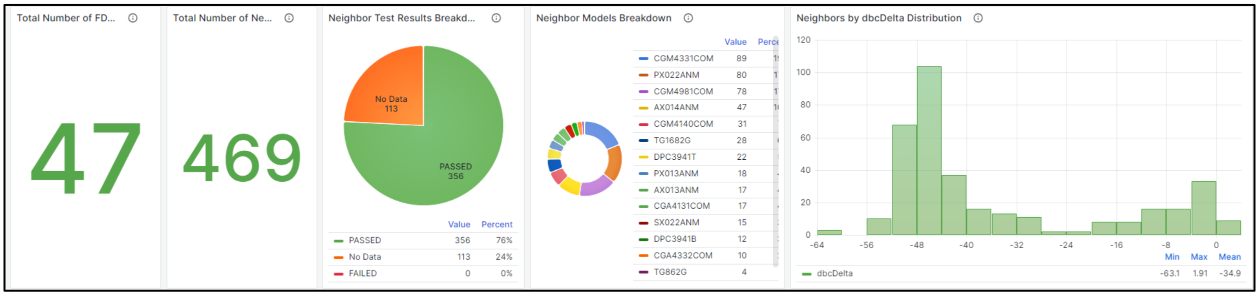

Notably, at this juncture there were 47 deployed FDX cable modems bounced against 469 neighbor CMs or STBs considered at potential risk. There are **zero cases** of nHAT alerts of at risk for ACI-related interference. 356 devices passed, which is also good news. But not good news was that the responsivity to the test was providing no data 24% of the time. In practice, this meant engineers would need to manually re-trigger the test until 100% of potential "at risk" devices were assessed, or at least everyone that was verified as online. These types of data or accessibility issues are addressed in production code error handling.

Of course, zero issues is highly encouraging. The sample size is small but not trivial. It would, in fact, imply that no special nHAT tool is required at all! Or, at least not something that needs to be built into a production process – a triage tool, perhaps. However, there are several reasons to maintain a healthy skepticism and remain vigilant developing the nHAT tool and processes:

1) N+x Systems – Current deployments of FDX are 100% in N+0 footprint. These are all recently upgraded (2016+), re-built plants. As FDX moves into the N+6 footprint, older and more challenging FDX environments are anticipated, possibly more prone to neighbor interference (NI).
2) MDUs – High rise MDU environments are likely to be the most challenging physical architecture and RF environment for NI. The current footprint has not landed on many high rise MDUs. Most have been of the garden style or townhome multiplex variety, which are of lesser concern.
3) PMA 2.0 – PMA has been in production at Comcast for years. It is also active in FDX deployments, for both DS and US and in the FDX band. However, there is new feature development focused on the FDX US band. The additional capabilities being added are to support independent US optimization on the much wider (than mid-split) band, accessing more knobs and levers. One of these knobs is the CM US Tx level. Additional flexibility to direct higher US Tx can help optimize the FDX band US MER. However, higher US Tx comes at the expense of NI risk, and thus anticipate employing algorithms to jointly optimize.
4) Invaluable as an Operational Tool – Should we still find that NI is negligibly small, it is nonetheless clear at this point that as a triage tool, application on a tech meter, or for MDU screening (loop wiring, splitters on taps), there will be value in the development of a scalable version even if access to the tool is more limited, and its use not automated within a production workflow.

4.1.3. nHAT Field Results of X-Class Launches

In addition to the *active* probing component of nHAT, Comcast has developed a *passive* approach to FDX interference detection that relies on already collected telemetry data. In so doing, we can monitor devices long after the active testing phase of initial deployment without the overhead associated with initial testing.

Passive monitoring looks for correlation between US traffic from the FDX customer and error counts in the legacy customer's devices. The theory here is that if the FDX customer is causing interference, it will be most evident during periods of high FDX US traffic. Figure 12 shows the codeword rate from an FDX device (top) and error rate from a neighboring legacy device (bottom) that were detected by the algorithm and appear highly correlated. The US traffic in this case was induced by a speed test. In a real deployment, speed testing could be used to force traffic if there is not enough customer traffic on the channel to assess risk, although then the test becomes less completely "passive."

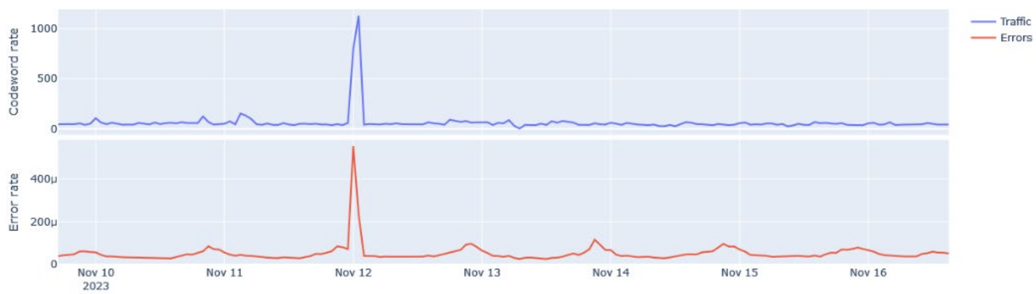

Figure 12 - A High-Correlation Pair Detected by the Algorithm

For each FDX device on a node, a correlation coefficient can be calculated against each legacy device. To reduce computational burden, the tool can leverage digitized network topology information so that the calculation is only performed between the FDX customer and its physical neighbors.

Figure 13shows the dashboard that was developed to view FDX correlation data within each node. The upper-left pane shows a scatter plot of all FDX-legacy device pairs on the node. The x-axis is the correlation coefficient calculated between the two devices, and the y-axis is the error rate of the legacy device. The dashed horizontal "Customer Impact" line denotes a threshold above which the error rate results in customer impact. The vertical dashed "Correlated Threshold" represents a point where we determine the correlation to be statistically significant. These lines divide the plot into four regions with the upper-right containing samples that are likely customer impacting and due to FDX interference.

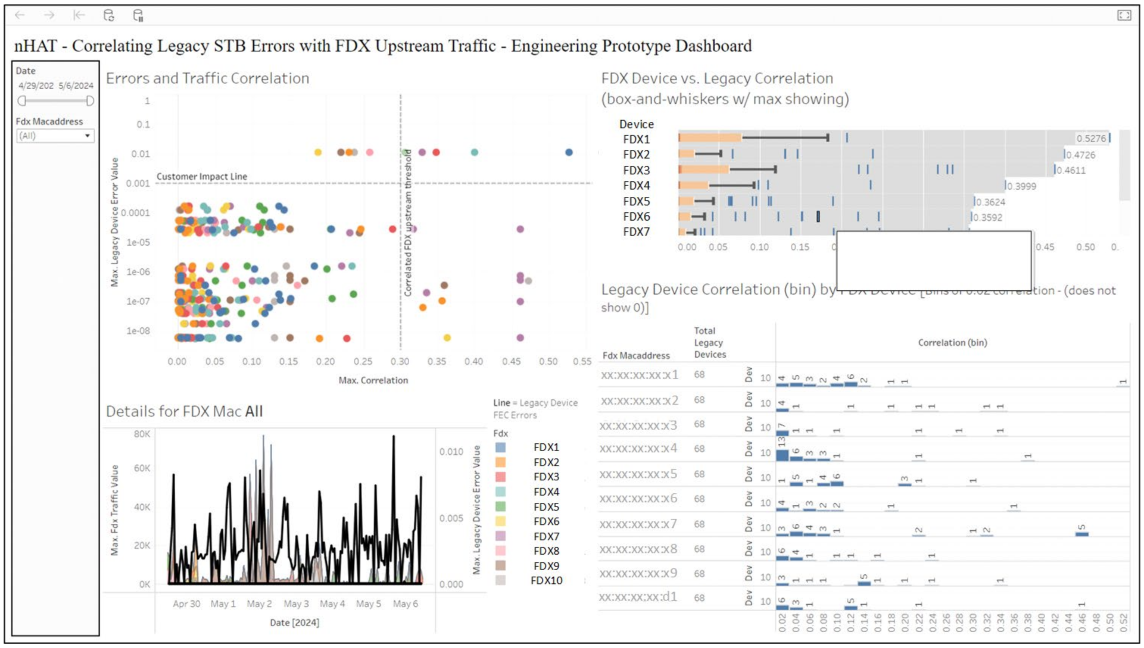

Figure 13 - nHAT Passive Monitoring Dashboard Prototype

In the final deployment of nHAT, the passive monitoring dashboard will contain logic to generate DS alerts like those generated by active testing.

The implementation of this technique is viewed as an effective complement to the "active" test probe-based nHAT tool. At installation, an active test is run, to immediately identify any at-risk customers and determine if they are to be acted on. Of course, since the network is not static, an option would be to repeat this test periodically. However, this is an intrusive test, and as FDX scales overhead will increase it may require coordination to avoid interfering with live services and disrupting vCMTS operation to manage it. The passive test is an ideal substitute that will just listen for potential issues following this first active RF assessment, and not create more network "noise" that only grows as FDX scales.

Remediation Options

Detecting potential interference, and being able to identify the violation, location, and alert at operational scale is the first step towards managing the risk. The next step is making findings actionable and driving automated ticketing as necessary. There are multiple remediation options available, and all have relevant use cases that can be applied depending on different scenarios.

However, as mentioned, nHAT has been operating live in production where X-Class services have launched. An assessment of every QAM STB and CM device on an FDX node leg is triggered when an FDX CM is discovered. X-Class has been live for 7 months as of this writing, and there has yet to be an at-risk home identified, which is a testament to solid design and field practices, and robust device front end designs. In addition, there have been no trouble calls that have been attributable to FDX interference on a neighbor.

Note also that a future nHAT software upgrade will restrict the "blast radius" examined, since we know a-priori the physical proximity to devices sharing a Tap, or one Tap removed, are devices most at risk, not an entire node leg's worth of devices. This scale reduction of nHAT tests will be based on our digitized

HFC network database and become more important as FDX moves into N+x footprints, with many more devices per node than in N+0 footprint.

Based on the ongoing findings of no interference risk detected and no trouble calls, the current approach is to do *no proactive remediation* based on an nHAT flag. Proactive remediation would prevent empirical correlation of nHAT results to customer experience impacts, which are often different than proxy measurements of technical parameters made in a lab to set thresholds. Instead, if risk has been identified, the potential "victim" device and associated customer account will undergo continued observation and any remediation applied will be on a case-by-case basis of findings.

Multiple possible causes of video impairment exist of course, and all of these common explanations should first be evaluated through the care process prior to executing a remediation associated to possible FDX interference. The remediation options and use cases are listed in Table 11below:

Table 11 - Possible Approaches for Remediating Neighbor Interference

Tactic	Explanation	Consideration
Disable FDX OFDMA	Removes interfering energy completely	Most X-Class class speeds not available – temporary fix until diagnosis at victim home
Reduce FDX OFDMA Spectrum	Removes some interfering bandwidth and thereby energy	X-Class customer's service could be impacted – speed dependency
Reduce FDX OFDMA US Tx	Reduces interference level	X-Class customer's service could be impacted – speed dependency
Blocking Filter on Victim Device	Block interfering energy at impacted device	Operational challenges and future impacts of filters in homes
Tap Change	Reduces interference level	Operational challenge of identifying specific Taps with best isolation characteristics

4.2. FDX Home Assessment Test (fHAT)

4.2.1. fHAT Theory of Operation

The foundational theory of the fHAT tool is straightforward and is shown in Figure 14 from a graphical perspective. Assumptions are as follows:

- N+0 network (not a gating functional dependence)
- Future X-Class customer has existing DOCSIS broadband service.
- The location of the DOCSIS 4.0 FDX CM at the premise, when activated via the SIK process, will be where the current CM is located.

The algorithm basics are:

1) Obtain the Node RF DS transmit power spectrum profile and the Node's US receive level. The former is a fixed, known setting and accessible via network design data or node configuration.
2) A full-band capture (FBC) of the DS spectrum at the existing cable modem is taken. This is available as normal CM telemetry.

3) The CM US Tx power is obtained – again, typical available telemetry. US Tx is reported per 6.4 MHz for DOCSIS 3.0 SC-QAM, and per 1.6 MHz for DOCSIS 3.1 OFDMA.

4) The new FDX US band is the same path as the FDX DS band, and prior to FDX is the same as the existing DS allocated to either QAM video (typically) or DOCSIS signals. The next step is to arithmetically subtract (see Figure 13 labels) the DS receive level determined by the FBC (2) from the node DS transmit power (1) to obtain the path loss and the frequency response of the path loss. This will be the response that an FDX CM will have to contend with when using FDX US.

5) Calculate the required FDX band CM US Tx from the path loss in (4) and the receive level set point (RLSP) (assumed – typical – flat over the US Rx bandwidth).

6) Calculate the CM TCP of legacy US + FDX US.

7) Determine if the CM is within its maximum TCP limitations. If it is, then this is a suitable candidate for SIK.

8) If not, determine the US Rx level for each FDX OFDMA block in the FDX band that will be available at the Node. Calculate the loss in US capacity attributable to lower levels leading to lower MER for each block.

    a. PMA can operate at lower granularity than 96 MHz, so this is a conservative assumption; with PMA active, the US capacity can only be equal to or better than the average over 96 MHz.

    b. There is a "too low for the channel to range" boundary condition, below which the capacity of an FDX channel that is beneath this is assumed to be zero which must be accounted for in the calculation.

    c. Another conservative assumption is that all the insufficient US Tx power penalizes the FDX band. With optimization of ranging and PMA algorithms, and mid-split maximum TCP settings, the available TCP can be balanced for better bandwidth efficiency across FDX and legacy, so can only be equal to or better than the assumption herein.

9) Determine if the X-Class speed tier to serve the home can be met with the available capacity calculated in step 8). If YES, then the premise is a candidate for SIK. If NO, then the SIK option is unlikely to be successful, and a professional installation should be scheduled and the customer notified accordingly.

10) Upon delivery and activation, the FDX CM initiates a speed test via the client installed on the device and determines if the X-Class speed for that customer is being met within acceptable limits. If not, then a notification is sent both to operations to identify a failed SIK has been detected, and to the customer with an opportunity to schedule a technician visit. Back-end process of operations with respect to their notification are to be determined – whether to pro-actively reach out or allow the customer to make this determination. Note that the customer will not be offline or unable to use the service. Rather, for example, they might be only able to achieve a speed of 1.4 Gbps instead of 2 Gbps.

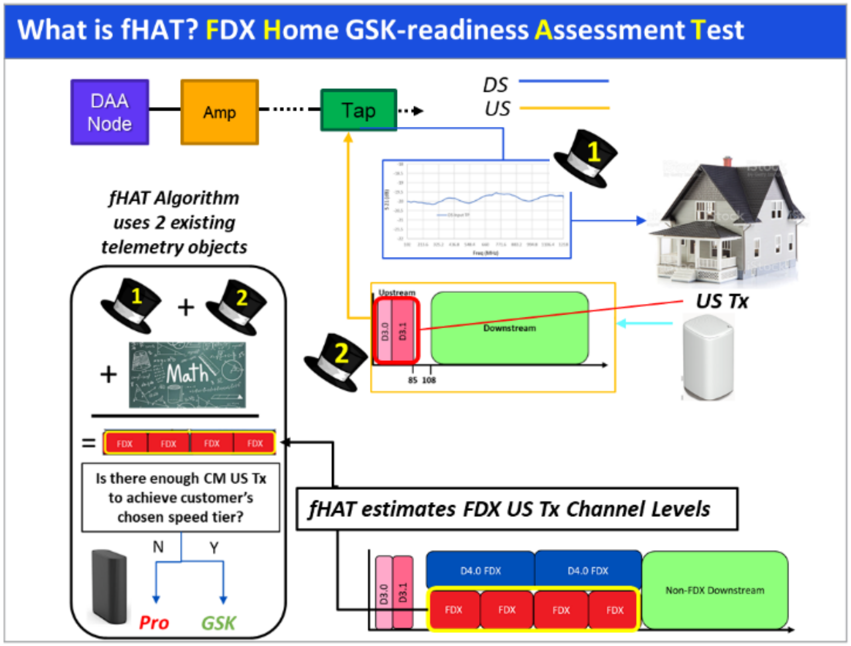

Figure 14 - Pictorial Representation of the fHAT Algorithm

Note that the fHAT algorithm does not require any DOCSIS 4.0 CMs, nodes, or active FDX spectrum. It ideally runs on a DOCSIS 3.1 system aligned to the approximate timing of a network upgrade to DOCSIS 4.0 technology, either just before or just after a DOCSIS 4.0 node is installed, and before FDX service is enabled. In this way, every resident or business on that node will have an fHAT "score" that becomes an attribute in the serviceability database for that account. A real-time service request can be automated and ticketed for a self-install or technician install at time of sale.

Note that the above applies to the case of an existing customer who is getting an upgrade to X-Class services that requires a DOCSIS 4.0 CM. Also note that X-300 is DOCSIS 3.1-based and requires no fHAT assessment. The fHAT screening relies on telemetry from the existing modem as inputs to the algorithm. In the case of a new customer, there are no metrics to rely on. In this case, deployments will be 100% professional installs, at least at first, but this likely pivots in the future.

Over time, as fHAT data accumulates and trends derived, the option to pivot to a "neighborhood average" approach may become reasonable. The idea here is that homes within a neighborhood are often built based on a similar plant and drop RF characteristics. Existing customers in the area of a new customer may represent a reasonable proxy. There is added risk to this approach, so this must be a careful balance of acceptable SIK failures versus the added customer disruption to support technician installation. The thresholds for such criteria would be derived from localized observations and SIK success results. It may be determined that there is not sufficient localized correlation to make this approach reliable enough. This will be discussed further in a subsequent section describing business policies.

Finally, note that the above describes the N+0 network for simplicity. The fHAT concept applies to an N+x network, but with differences in the RF profiles used at the node for the DS launch and for the US receive level (receive level set point, or RLSP). The RF actives separate DS and US paths within their clamshells, but the compensation of the US FDX path required to be implemented in an FDX amplifier can still be estimated from the measurement of the FBC at the existing CM.

4.2.2. fHAT Guided Self-Install (SIK) Predictions

Available information to predict SIK success are:

- Distribution of US Tx on existing DOCSIS 3.1 CMs, in particular those with OFDMA enabled, shown in Figure 15below. The working assumption for self-install is that the FDX CM will be installed at the same location, with the same path loss, as the DOCSIS 3.1 CM that is being used for the fHAT assessment.

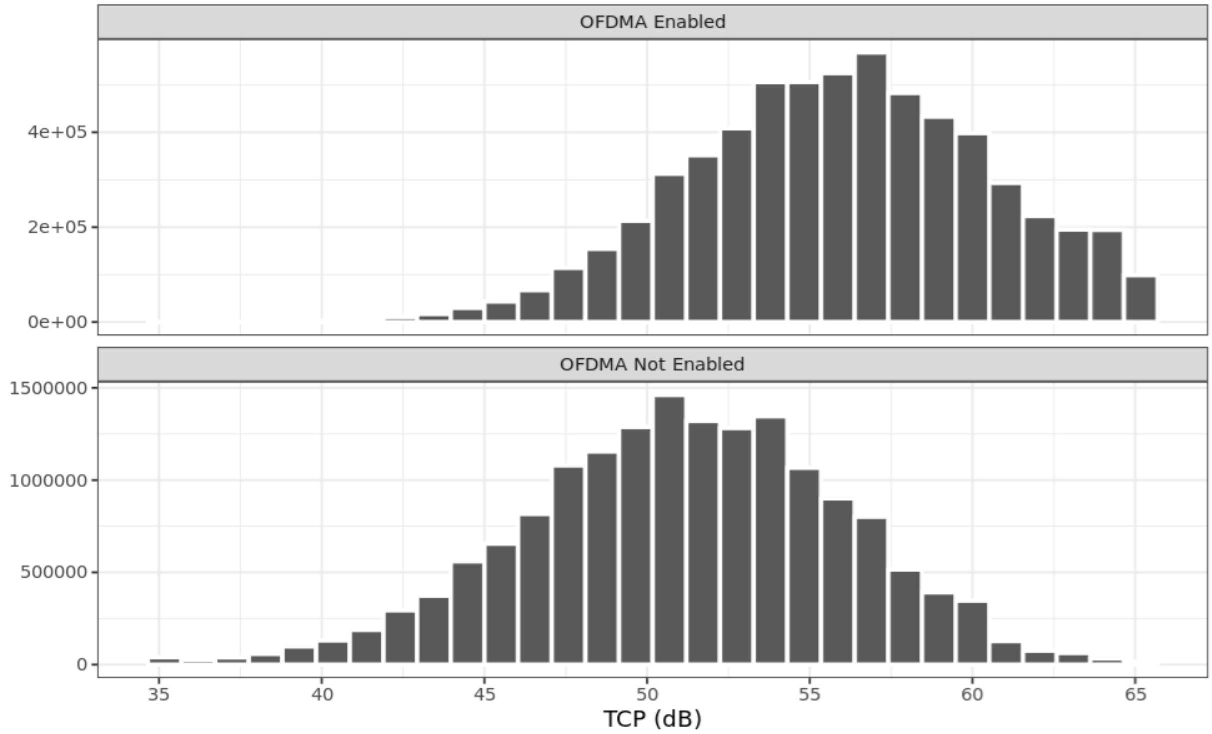

Figure 15 - OFDMA US Tx TCP Distribution of Mid-Split Enabled DOCSIS 3.1 CMs

- Assumed FDX CM transmit tilt profile from the DOCSIS 4.0 FDX Standard (for pre-fHAT SIK analysis purposes). The calculation fHAT specifically does is to determine the projected US Tx profile, based on the FBC at the measuring CM, and the known DS RF power spectral density (PSD).
- Baseline MER performance of an FDX system operating at US Rx of nominal return level set point (RLSP) and corresponding QAM profile with PMA active (2k-QAM), based on system modeling and analysis, as well as what is measured in the field.
- RLSP range of operation – this is the window around the nominal US Rx level at the RPD receiver. The RPD has a separate US Rx for the legacy US and the FDX upstream. The RLSPs can be independently set and configured. Because of the additional 276 MHz of spectrum added for X-2000 (and more in the future), the FDX receiver is configured with the generous ability to range channels much lower than in the mid-split band, and telemetry and tools behaviors adapted accordingly – i.e. CMs will not be put into partial service (some capacity is better than zero capacity), and alarms will not be triggered, for example, if a CM arrives at the RPD, for example, 6 dB below the RLSP. This means additional US PMA profiles to define and assign in the FDX band, but this is anticipated in the standard and well within the capability of the technology.

The analysis and the tool in practice will use this information to predict the capacity achievable. Adjustments are made (margin added) for empirically derived speed test deltas to theory and acceptable "pass" thresholds.

Today, the success rate of the SIK process is between 80-90%, which is baked into resource planning and budget exercises. There are natural variations of device types, age, scope of swapping (data, video), etc. We will use 85%. SIK failures are usually associated with issues at the home – Wi-Fi related, wiring errors, user execution errors during activation, low-split home amplifiers, etc.

The use of SIK is enormously powerful, providing significant cost savings and customer experience benefits. Every truck roll needed to get the customer their desired service is a cost and disruption to be avoided whenever possible. So, getting a handle on how these "baked in" benefits are effected by X-Class is of critical importance. For X-Class tiers of (all units of Mbps) 300/300, 500/500, 1000/1000, and 2000/2000, below are the expected implications on SIK:

300/300

This is a tier delivered by mid-split capable DOCSIS 3.1 devices, so remains at 85%.

500/500 and 1000/1000

This tier requires the FDX band, but not all of it, and it is not necessary that it be used at nominal efficiency. Even partial services within the FDX band support these tiers (i.e. only one or two FDX channels operating). There is a borderline case for 1000/1000 where, due to poor channel conditions of whatever cause, and the ranging algorithm trying its best to deal with this, that only one FDX channel successfully ranges, and the RF fidelity of that channel limits the modulation to 256-QAM. This scenario is unlikely, difficult to predict, and yet would still technically meet speed test minimum acceptable for X-1000, so we do not account for it as a failed case. The analysis then decomposes to the case of all FDX channels failing to range. For this we will use a maximum ranging lower limit of – 9 dB. Also, very conservatively, assume *no* ranging behavior takes place, for example, to enable one FDX channel by re-allocating power away from the others to get one channel enough level to get online. In this scenario 11% of the devices will miss the lower limit target based on today's US Tx distributions, making 89% eligible for SIK. At an implementation success rate of 85%, the "budget" associated with SIK cost savings would be an assumption of 76%.

2000/2000

This peak tier is the lone one that requires a nominal level of OFDMA efficiency for the amount of FDX US spectrum allocated. This takes the eligible SIK percentage to 65% and the budget number for the cost benefits of SIK therefore becomes 55%.

Aggregate of All Tiers Relative to Current SIK Success Rate

When accounting for weighting of tier penetrations observed to date – roughly 10%/40%/40%/10% across X-300/500/1000/2000 - the net of SIK for X-Class becomes **74.8%,** or 10.2% lower than the 85% baseline.

4.2.3. Future Expectations for SIK

A few items are favorable going forward for SIK success:

1) Because FDX is, in general, following the mid-split upgrade footprint, the hygiene related to activating OFDMA mid-split is taken care of before FDX arrives. In particular, home amplifiers are being minimized and trap filters removed. While there are mid-split drop amps available and in inventory, guidance to the field in mid-split activations when found is to solve without the crutch of an amplifier if at all possible. Because of this, 15-20% of homes with a drop amp – which would prevent FDX and gets in the way of mid-split OFDMA and contribute to truck rolls when the SIK fails – will be lower where FDX is going.

2) Current installations of FDX are N+0 networks, where the plant is the most physically stretched, increasing RF losses, to maximize homes passed coverage. N+x systems have more favorable path loss characteristics to the first active and node, which we expect will improve SIK success rate over time.

3) Self-install kits today come with simple instructions. An instruction for an FDX installation, where all coaxial STBs will be removed in favor of Wi-Fi only internet protocol (IP) STBs, will be the elimination of any splitters on the coaxial path to the FDX CM that the customer can identify. No assumptions are made around this guidance being taken or done correctly. However, it seems reasonable to assume many customers will be able to execute this simple task and decrease the path loss that fHAT assumes is there in its calculation of eligibility.

4) PMA 2.0 – PMA as used today in the FDX US is very similar to the mid-split band, using MER-based readings periodically to adjust QAM profiles to optimize use narrower segments of available spectrum. A prelude to PMA 2.0 is the ranging adjustment mentioned above, which limits partial channels in favor of the available capacity that can be obtained despite lower US Rx levels. Other parameters in the RPDs dynamic range window (DRW) have been made configurable for the FDX US Rx. With PMA 2.0 the other end of the link – CM US Tx – will also be provided with more flexibility. The updated algorithm will allow for "water filling" type optimization using knobs on both ends of the link for each CM and allow for more US profile options. There will be less adherence to the static, fixed settings around which today's legacy US is managed, benefitting capacity, speeds achievable, and thereby SIK success.

Working *against* SIK success rate relative to day 1 are:

1) An X-Class customer becomes an all-IP customer, so every X-Class installation for a customer that has QAM video will involve video STB swaps to Wi-Fi IP STBs. While this process is in place today, and Wi-Fi STBs simplify the activity, the 85% baseline is a blend that includes some SIKs that are CM-only upgrades, which are simpler than CM + STB.

2) Increases in speeds will add more FDX channels in parts of the band with more loss and require more CM transmit power. There may be a penalty to SIK eligibility at each new peak speed relative to X-2000. The good news is that for future X-Class speeds, there will be a mountain of X-2000 installation data to learn from and adjust algorithms and practices accordingly.

4.2.4. Business Requirements

As described above, with all launches of speed products, financial considerations are made concerning the ratio of pro-installation and guided SIK. As noted, Comcast's network supports a higher percentage of SIK today. The introduction of our X-Class speeds poses a new challenge. With this new challenge, the business unit provided the following requirement:

Comcast must have the ability to approximate the success of a new, or existing, customer achieving their purchased speeds, and that mechanism will be incorporated into the guided self-install kit versus pro-installation logic for all X-Class products.

From an engineering point of view, fHAT came into existence to perform the speed approximation calculation. After further decomposing the requirement, the network technical product team needed to consider how to make the test results consumable for new or existing customers.

Part 1: fHAT Test Result

When the fHAT calculation was described to the product team, and compared to the requirement, two core design tenants were established:

1. How to limit the number future iterations needs, and
2. how to limit dependencies between the network software teams producing results, and DS software teams consuming the results.

One option was to run fHAT on-demand, when a customer was moving through the purchase process. The customer facing system would log the desired speed product, initiate fHAT to run the calculation against the requested speed, consume and use the result to offer, or deny, SIK. Unfortunately, that tactic would end up being quite costly over time, as X-Class overtook legacy products and the product portfolio is not static, which in turn would mean continuous iteration and management.

As the requirement was re-examined, the mindset shifted such that on-demand test results were not required. Consideration for this was principally because large fluctuations in healthy RF plant from day to day, the main factor in the fHAT calculation, are minimal. The approach determined then was that the first iteration of fHAT would be calculated and logged once daily. From there, design conversations could begin.

Comcast's speed product structure of defined US and DS speeds was the main concept to contemplate. The need was a way to represent the data that was easily consumable, and not tied to specific speed product offering(s). With that, the established test results will be translated into Mbps and round the results down to the closest hundredth. By representing the test result in Mbps, it removes a number of future iterations and software team dependencies when the business adds, alters, or removes specific products from their portfolio, as all products at their core are Mbps.

The decision to round down was out of an abundance of caution until a larger data set can be gathered to improve the overall logic and influence future improvements of the calculation itself.

Part 2: Test State and Customer State

To fulfill the second portion of the business requirement, the ability to use this test result for existing and new customers, the first need was to identify a data point that persists continuously, whether or not there is an active account and device. The data point chosen was location – every individual location where Comcast can provide service.

With the source data set for fHAT results identified, the next step is defining the lifecycle of a location to the account/device. The lifecycle states were determined as:

1. Existing customer-device of a location
2. Recently disconnected customer-device of a location
3. No customers-devices ever connected to location.

From these three lifecycle states, the next consideration was how to translate these into fHAT test states to fulfill the requirement. Existing device fHAT calculations were quick to be defined; however, recently disconnected or never connected are more difficult as there is no device on premise to measure.

The logical progression of the location to account/device lifecycle would roughly be as follows. First, there will be no device on location. Second, a device would be added to a location, and finally a device could come and go from a location.

There is a possibility the location may fall into never connected, as there is a time dimension delineating inactive and never connected. Locations left inactive for a period become never connected. Obviously, the true lifecycle is not linear and could likely change between the three states.

Nonetheless, merging the lifecycles determined the fHAT test states were:

1. ***Active***: active device on location, with an fHAT result
2. ***Inactive***: active device was on location, and produced an fHAT result, which will be persisted for a period of time for possible future use
3. ***Average***: active device was never on location; therefore, average the fHAT results across the node associated to the location for our best approximation

Active test results have the highest confidence level for success because it is using an active device for calculation. The necessity of "new" customers having this test result was discussed. However, it was concluded that using data from neighboring locations is more advantageous than making a business policy not using any data.

To minimize translation of these test states, they are represented using 3 attributes, as shown in Table 12: Test Date, Test Result, and Active Device.

Table 12 - fHAT Scoring - Test Results and Attributes

Date	Test Result	Active Device	Score Type
Today	Mbps	Yes	Active
Yesterday – 1 year	Mbps	No	Inactive
Today	Mbps	No	Average

4.2.5. Business Operations

The operational flow is described as follows:

- When an RPD configuration is changed to contain FDX spectrum, a notification is sent to fHAT. Then fHAT will ingest the RPD name, and pull all devices connected to an RPD.
- All active devices connected are identified, an fHAT calculation is executed, and the results translated into Mbps.
- Once a day, the active test results will be batched to a centralized data store, separate from fHAT, that has all locations and the device media access control (MAC) address associated.
- Each test result is merged based on the device MAC address, and the test date and test result are updated.

- After the centralized data store has ingested and updated the active test results, the second logic set initiates. All test results on an RPD are queried, averaged and rounded down to the closet test result.
- The average is inserted for every location that does not have an active device, with one exception. On subsequent daily batches, if an active device is removed from an RPD, when the data is synchronized, the date and test result for the now absent device does not update, thus representing the inactive status.
- When customers enter the purchasing flows, the customer's address is paired to a location, a query is sent to retrieve the fHAT result and ingested.
- Logic will use that score to determine if a pro-installation is required, or if SIK is suitable. Note that an fHAT score is one of several inputs to the pro-installation/SIK logic, not the sole determining factor.

As Comcast continues to expand the X-Class product offering, by adding, altering, and removing from the portfolio, fHAT will continue to operate with minimal iteration. The core development for fHAT for the network teams lies with ensuring that as we expand the FDX spectrum on a node configuration, the number of fHAT results being translated into Mbps will need to be managed and documented. On the consumer side, as designed, the ingestion of Mbps will remain as is, and the only development work is ensuring the translation of Mbps to the new products.

4.2.6. fHAT Proof of Concept

As described previously, fHAT uses the actual transmit and receive levels from the RPD and customer's existing CM to estimate RF path loss. This path loss, along with the return level setpoint of the RPD, is used to estimate the transmit power of the FDX cable modem in the FDX bands. The RPD output level and tilt settings can be obtained from the RPD configuration file or design data, and the CM receive and transmit levels can be polled directly. Testing shows that using the SC-QAM receive levels and the SC-QAM transmit levels in the US provides the best baseline to calculate this link loss. The mid-split OFDMA channel typically has tilt, due in part to uncorrected loss characteristic of the CM. Due to this uncertainty of calibration accuracy, current testing shows better precision using the SC-QAM telemetry. This may change over time.

Figure 16 shows the path loss estimation of the SC-QAM frequencies from measurements taken on several cable modems in both the forward and return bands. Note that XB6, XB7, and XB8 are Comcast DOCSIS 3.1 Gateways, and XD4 is an FDX CM. This is compared to a network analyzer measurement of the path loss. A virtual CM loss was also calculated, which matches the network analyzer measurement.

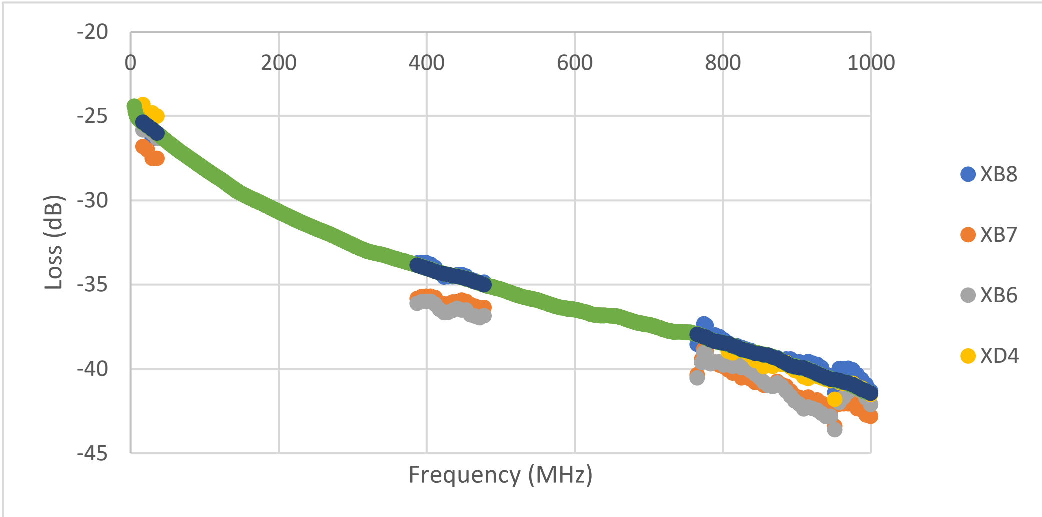

Figure 16 - Path Loss Estimation from Polled Cable Modem and RPD Receive and Transmit Values

Once the loss is estimated at the SC-QAM frequencies, the loss across the FDX spectrum needs to be calculated. This was completed using both a linear regression and a square root of frequency estimation.

Figure 17 shows the calculated path loss using both methods, linear regression, and square root of frequency. The square root of frequency method is within 1 dB accuracy of the actual path loss. Comparisons were made to the actual transmit levels of an FDX modem to the calculated transmit levels. This is shown in Table 13for the linear and square root of frequency estimated transmit levels. The difference between the estimated and actual transmit levels of the FDX cable modem are within 2 dB.

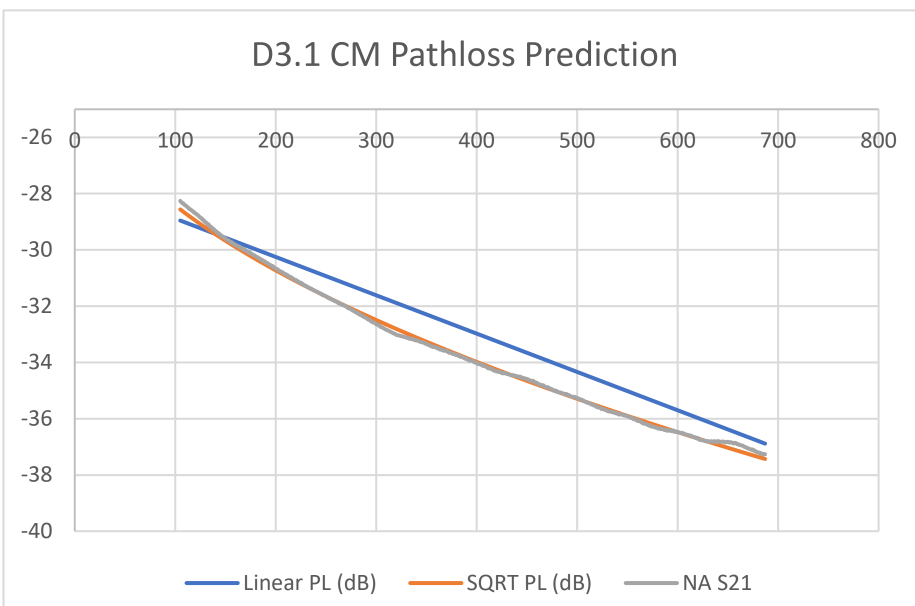

Figure 17 - Path Loss Estimation Using Linear Regression and Square Root Frequency

Table 13 - Accuracy Comparison: Linear Regression versus Square Root Frequency

Linear Regression Frequency (MHz)	Predicted Transmit Power (dBmV)	FDX Reported Transmit Power (dBmV)	Difference (dB)
156	49.3	47.9	-1.3
252	50.6	49.7	-0.9
348	51.9	51.7	-0.2
444	53.2	52.7	-0.5
540	54.5	54.7	0.2
636	55.8	55.7	-0.1
TCP (dBmV)	**60.9**	**60.6**	**-.03**
Square Root Frequency (MHz)	**Predicted Transmit Power (dBmV)**	**FDX Reported Transmit Power (dBmV)**	**Difference (dB)**
156	49.3	47.9	-1.4
252	51.3	49.7	-1.6
348	52.9	51.7	-1.2
444	54.2	52.7	-1.5
540	55.4	54.7	-0.7
636	56.5	55.7	-0.7
TCP (dBmV)	**61.7**	**60.6**	**-1.1**

A proof-of-concept tool that represents the core of the fHAT algorithm has been developed to calculate the path loss and estimated FDX transmit power for devices in the field. A sample output of the tool is shown in Figure 18 and Figure 19. Note that since the tool is doing a mathematical extrapolation of a legacy mid-split band CM, it is able to extend to the case of full FDX bandwidth to 684 MHz, as well as showing the case for half of the FDX band used as US – which is very close to the X-2000 configuration. It is a necessary capability for fHAT, of course, to be able to calculate any FDX configuration.

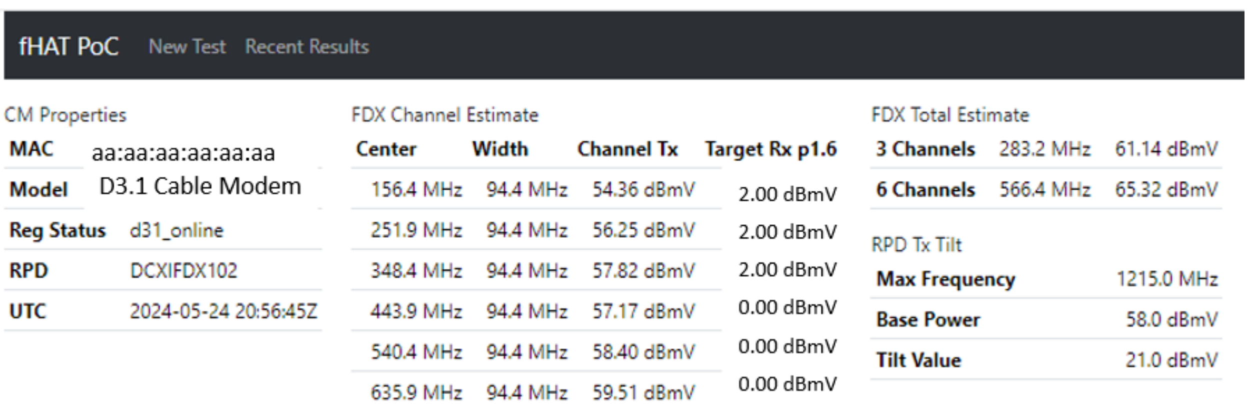

Figure 18 - fHAT Proof of Concept FDX Channel Transmit Power Estimation

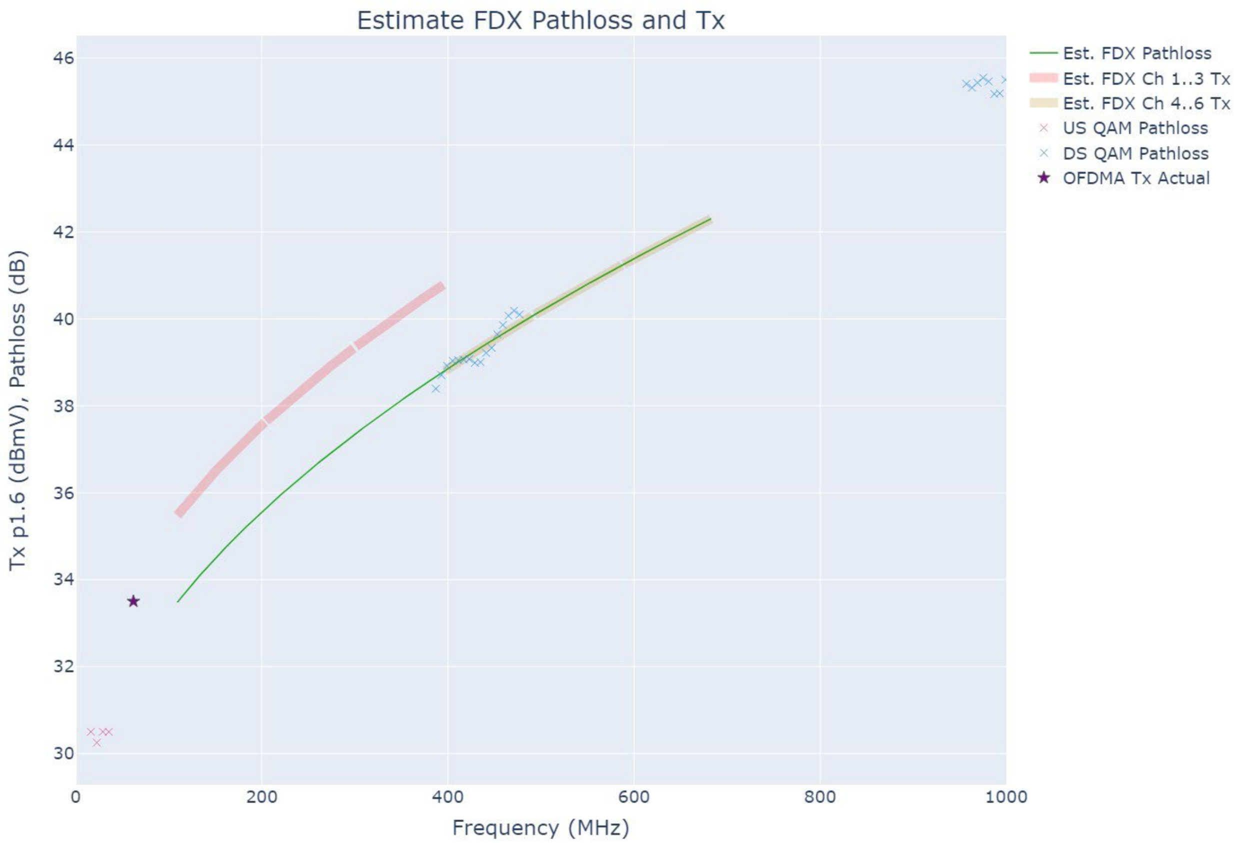

Figure 19 - fHAT Proof of Concept FDX Channel Transmit Power Estimation Graph

This tool shows the estimated TX power for each FDX channel and the total power of the FDX channels. This is broken into two groups, the lower 3 FDX channels and all 6 FDX channels. The total power of 65.32 dBmV for all six FDX channels is just above the TX limit of 64.5 dBmV.

As we have seen in earlier analysis and considered within the SIK predictions, the RLSP of one or more of the FDX channels would need to be reduced to utilize the full FDX band in this example case. Alternatively, if this customer is a lower speed tier customer, just the lower three FDX bands could be adequate. Again, this is a key observation and important element implemented as part of the SIK-or-not decision process.

4.2.7. fHAT Future Roadmap

The fHAT algorithm was founded based on how FDX is configured and operated today. As part of the learnings of launching FDX and deploying X-Class services in scale, optimizations are being implemented now and planned to go forward that will impact how fHAT behaves, and thereby the roadmap for fHAT. We will touch briefly on three areas that will play into fHAT's roadmap, changes to the algorithm and processes, and potential software updates: speed plans, DS inferences, and additional commentary on PMA 2.0.

<u>Speed Plan and fHAT Impacts</u>

The launch of X-Class has peak US speeds of 2 Gbps. Accordingly, the FDX spectrum allocation is aligned to ensure that the existing legacy mid-split band, when augmented by the additional FDX upstream, can deliver 2 Gbps. The resulting FDX US allocation to do this is 108 MHz to 384 MHz, or just shy of three 96 MHz OFDMA US blocks. The way the tools works, as previously described, is to estimate the CMs US transmission profile, assuming it was an FDX-capable CM transmitting from that same location in the home and looking to transmit in the US up to 384 MHz.

Of course, with speed increases over time, the FDX allocation will increase. The current business plan calls for two more expected FDX allocations – use of four 96 MHz OFDMA US blocks, and the use of the full FDX band allocation, which is six 96 MHz OFDMA US blocks. For these additional allocations, the fHAT estimate will then be based observing the DS by this same additional spectrum amount to assess readiness for these new speeds. The fundamentals of the algorithm do not change.

Note that devices that have already been installed with X-Class services can be reassessed, this time based on their actual US Tx of the existing service. In this way, it can be discovered whether an existing customer who desires a speed upgrade will be able to get that speed without a visit by a technician.

Profile Management Application (PMA) 2.0

The use of PMA allows operators to ensure they are running as bandwidth-efficiently as possible on their DOCSIS 3.1 channels. PMA takes advantage of the DOCSIS 3.1 multiple modulation profiles (MMP) feature, existing telemetry, and a cloud-based application to periodically monitor channel fidelity and FEC statistics and adjust them to maximize bandwidth efficiency. Because the network is no longer static, there is not a “one-and-done” QAM profile setting that has in the past forced operators to incorporate significant margin above MER thresholds for a particular QAM profile to cater to the lower fidelity set of users.

With PMA in production for both DS orthogonal frequency-division multiplexing (OFDM) and US OFDMA over the mid-split band, attention has turned to the optimization of metrics, dynamics, number of profiles, and thresholds for the FDX band. The FDX upstream, because of the expanded frequency range over which the CM must transmit, and with fixed TCP, is inherently more challenging and dependent on robust PMA.

Part of the feature development of PMA is to allow for flexibility in the ranging process – enabling a wider range of acceptable US receive levels at the node, and to be able to manipulate the CM US transmit levels. A configurable (but fixed) US ranging window has been already incorporated, as previously noted. Generally, it is expected that PMA 2.0 will bias the CMs to transmit closer to their maximum US TCP. This will maximize the node Rx MER in the FDX band, countering some of the RF loss challenges alluded to above that result in reduced fidelity.

However, in the context of the nHAT discussion, what is “optimal” US Tx may be a balance between maximizing modem capacity enabled, versus ensuring sufficient modem capacity while minimizing the risk of interference created. The beauty of a virtual network function operating a “PMA 2.0” outside the vCMTS function itself, and the use of machine learning (ML) algorithms, is that we can apply any set of variables of interest, weight them accordingly, and use them as input to an algorithm with a more global view of optimization.

By having PMA and nHAT virtual services integrated through the back-office and based on a common data architecture, this capability, along with active FDX bonding group management, opens options for SW-based nHAT remediation if it is determined that remediation is required.

Consideration of DS Metrics

We have discussed the nature and new challenges of the FDX US. However, the FDX OFDM DS also provides metrics than can be insightful for fHAT discovery. In particular, the DS Rx level is a reasonable proxy for whether a device is deep within the home. Low DS Rx may reduce efficiency of the OFDM, which is one of the reasons the DOCSIS 4.0 standard identified a point of entry CM as the reference architecture – to minimize losses that occur in the home.

The DS Rx of an existing DOCSIS 3.1 device in a home, if below some selected threshold that corresponds to an unacceptable low FDX DS MER, can also be used as an input to screen for success probability of a pro-install candidates. Such thresholds will be learned as RBA switching and CM EC performance discussed previously becomes characterized.

At installation itself, an SIK bring-up will include a check on US Tx and speed capability enabled. This data will be used to compare against the fHAT prediction and used for additional optimization and algorithm tweaks.

In addition, with an FDX CM now in place, the device can also take a snapshot of its echo environment. For homes that are mis-wired, poorly wired, bad connectors, etc., a high echo would result in a high level of US reflection back into the DS Rx that must be handled by the EC. If poor enough, DS MER may be impacted enough to affect the maximum DS QAM profile that the CM can receive on and reduce its capacity. For peak speed tiers, this could impact their ability to reach their peak. As the devices boots up and comes online, it can be programmed to execute a self-test, and in doing so discover issues in the home, send a notification to operations about the findings, and prompt for follow-up by a technician.

In addition, a notification can be sent to the customer to alert them to a likely speed issue, and prompt them to schedule a call to repair, and/or make recommendations that may resolve the issue themselves. This notification and business process would mirror the similar case discussed for an FDX US that was found to fall short of its speed objective after SIK was executed, and speed test performed.

5. Conclusion

DOCSIS 4.0 FDX provides new technology that enables significant new speed options, and in particular symmetrical multi-gigabit speeds. While the technology has matured and service launched, operationalizing FDX and X-Class services at scale is also emerging. Prioritizing automation and software-centric tooling and process management, two new and effective tools have been developed that will bring efficiency, performance, and scalability to the DOCSIS 4.0 migration and enablement process.

nHAT previews potential service risk in areas where FDX has launched, and provides a mechanism to assess the level of risk, automate mechanisms to notify and create actionable alerts, and remediation options should it become necessary to protect existing service above and beyond the capabilities of the typical RF properties of an HFC network today.

fHAT provides a way to extend SIK practices into X-Class services, by determining in advance whether an X-Class installation when implemented via SIK has a high likelihood of success in meeting the speed tier desired. SIK provides substantial economic benefit for the company, reducing truck rolls substantially, and in turn drives higher customer net promoter scores (NPS), as most customers prefer not to require an appointment with Comcast to have their services installed if the alternative is simple and effective.

With these two tools, Comcast is poised to deliver X-Class service effectively, efficiently, and with an eye towards the customer experience – for X-Class customers and their neighbors – as DOCSIS 4.0 expands rapidly across the footprint.

Abbreviations

ACI	adjacent channel interference
AGC	automatic gain control
CAGR	compounded annual growth rate
CCI	co-channel interfering
CM	cable modem
CMTS	cable modem termination system
CPE	consumer premises equipment
DAA	distributed access architecture
DOCSIS	data over cable service interface specification
DRW	dynamic range window
DS	downstream
EC	echo cancellation
FBC	full-band capture
FDD	frequency domain duplex
FDX	full duplex docsis
FEC	forward error correction
fHAT	FDX readiness Home Assessment Test
HFC	hybrid fiber-coax
IG	interference group
IP	internet protocol
MAC	media access control
Mbps	megabit per second
MHz	megahertz
ML	machine learning
MMP	multiple modulation profiles
N+0	node+0 actives
N+x	node + x actives (amplifiers)
nHAT	neighbor home assessment test
NI	neighbor interference
NPS	net promotor scores
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
PA	power amplifier
PHY	physical layer
PMA	profile management application
PoE	point-of-entry
PSD	power spectral density
QAM	quadrature amplitude modulation
RBA	resource block assignment
RF	radio frequency
RLSP	receive level set point
RPD	remote phy device
SC-QAM	single-carrier QAM
SIK	self-install kit

STB	set-top box
TCP	total composite power
TG	transmission group
TR	truck roll
TX	transmit
UHS	ultra-high split
US	upstream
US Rx	upstream receive level
US Tx	upstream transmit level
vCMTS	virtual cable modem termination system
vSG	virtual service gateway

Bibliography & References

[1] Howald, Dr. Robert and Jon Cave (Comcast), Olakunle Ekundare (Comcast), John Williams (Charter), Matt Petersen (Charter), Developing the DOCSIS 4.0 Playbook for the Season of 10G, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[2] Howald, Dr. Robert, Roaring into the 20's with 10G, 2020 SCTE Expo, Oct 13-16 (Virtual Event).

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

[4] Howald, Dr. Robert L, and John Ulm, Saif Rahman, and Dr. Zoran Maricevic, Collision-Free Hyper-Speeds on the Bi-Directional FDX Highway, 2022 SCTE Expo, Sept 19-22, Philadelphia, PA.

[5] Prodan, Dr. Richard, 10G Full Duplex DOCSIS Implementation Exceeds Expectations, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[6] Prodan, Dr. Richard, Optimizing the 10G Transition to Full Duplex DOCSIS 4.0, SCTE Expo, Oct 13-16, 2020 (Virtual Event).

[7] *DOCSIS 4.0 Physical Layer Specification*, CM-SP-PHYv4.0-I01-190815, CableLabs 2019