Full Duplex DOCSIS

John T. Chapman, Cisco Fellow & CTO Cable Access, jchapman@cisco.com
Hang Jin, Distinguished Engineer, hangjin@cisco.com
Cisco

Abstract

A method is proposed to use the same spectrum for simultaneous downstream and upstream transmission. This technique could allow a DOCSIS 3.1 plant to run at 10 Gbps downstream and 10 Gbps upstream in the same spectrum.

This technology would boost the upstream throughput by 50x what it is today without having to extend the maximum frequency of the plant. It will provide the ability to deploy true symmetric services.

This white paper will discuss the hardware and software technology of FDX. It will also discuss different deployment plans and options.

Cisco started working on a full duplex concept for DOCSIS in 2014. In early 2015, the FDX DOCSIS project received internal research funding from the Cisco Technology Fund. This white paper is a result of that work and is checkpoint half way though our project.

TABLE OF CONTENTS

INTRODUCTION

Looking Back

For the first 30 years, say the mid-60's to the mid-90's, the cable plant was just a cable plant. There was no fiber and generally there was no upstream return path. When VOD services where introduced in the 80s and 90s, there was a need for upstream signaling for STB and an upstream OOB carrier was established (there was already a broadcast downstream OOB carrier).

This large coax system would not scale for data so in the 1990's optical nodes where deployed to create an HFC plant that had a limited number of households passed in a coax segment so that the noise funneling from the homes was tolerable enough that data could be passed over the upstream path.

Data started off with a shared 2.5 Mbps upstream channel. That later grew to 10 Mbps with DOCSIS 1.0 and to 30 Mbps with DOCSIS 2.0. The number of channels that the silicon could support also grew from one to four. Today with bonding and DOCSIS 3.0, the return path data capacity is around 90 Mbps in a 42 MHz or 65 MHz return path.

In practice, the spectrum below 42 MHz is full of legacy traffic and will likely stay as ATDMA traffic. Although the OFDMA traffic from DOCSIS 3.1 could theoretically be mixed with the ATDMA traffic, the bandwidth gain provided by OFDMA and LDPC would be quickly offset by the time and frequency guard time requirements that must exist between DOCSIS 3.0 and DOCSIS 3.1 spectrum.

With DOCSIS 3.0 and DOCSIS 3.1, it is possible to run a higher frequency upstream return path. The two favorite options are 85 MHz and 204 MHz. Both these return paths push the start of the downstream spectrum to a higher frequency. In todays FDD (Frequency Division Duplex) system, new upstream spectrum is taken from existing downstream spectrum. As an added penalty, more bandwidth is lost to the cross-over band as the split increases in frequency.

Enter Full Duplex.

Today's plant is run as FDD (frequency division duplex). The proposal is to run the plant in FDX (full duplex). This comparison is shown in Figure 1.

The physical path of a N+0 FDX DOCSIS system is shown in Figure 2. The

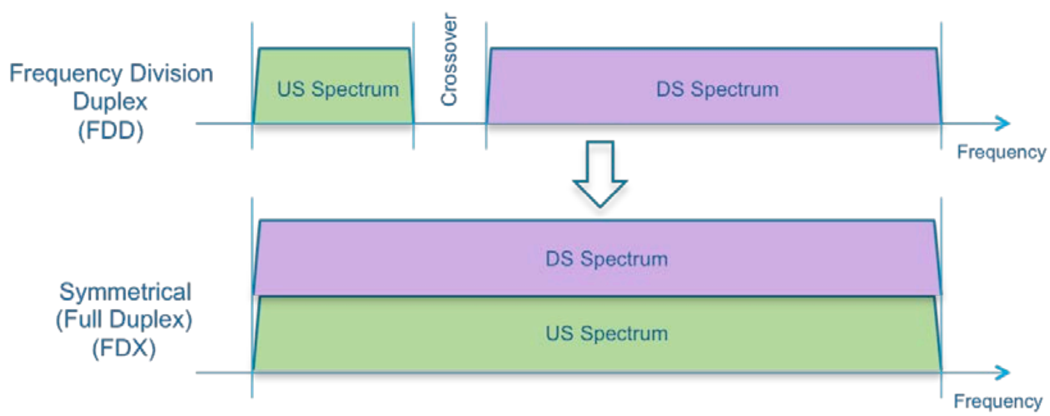

Figure 1 - FDX Spectrum

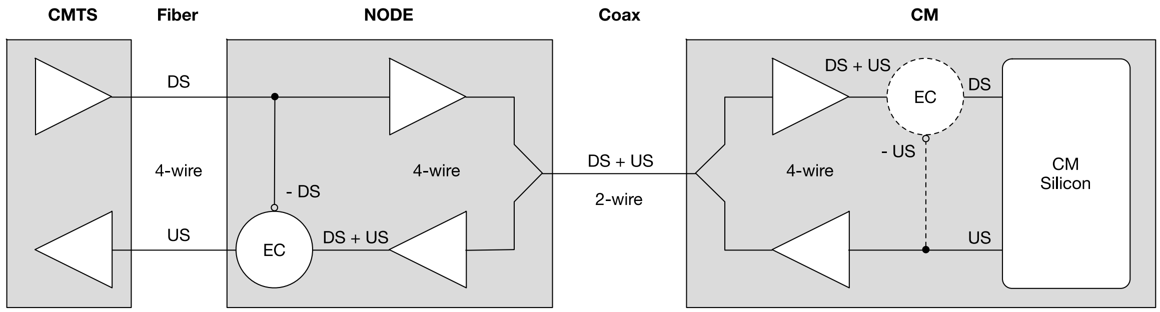

Figure 2 - FDX Path

CMTS is a 4-wire path. A 4-wire path is a path with separate transmit/ground pair and a receive/ground pair. The CMTS signal travels to the optical node where is converted to a 2-wire path. A 2-wire path is where transmit and receive signals are on the same wire and there is a common ground.

Note that energy can travel simultaneously in opposite directions on a passive coax. It is only when amplification is needed that the transmit and receive signals need to be separated into different paths. For the node and the CM, this can be done with a power splitter/combiner.

The combined DS and US signals are combined and arrive that the CM receiver. As we will see, this is not a problem that is solved by an echo canceller, as there can also be interference from adjacent CMs.

In the reverse direction, the upstream path from the CM leaves the 4-wire domain of the CM and becomes 2-wire. The combined DS + US energy from the coax enters the node upstream. Here there is an echo canceller. The echo canceller may be a combination of an analog echo canceller followed by a DSP based digital echo canceller. If the echo canceller can calculate the proper level of the DS signal to cancel out from the DS + US signal, only the US signal will be left. In practice, the remaining attenuated DS signal becomes the new noise floor.

Before diving into the technical details of how the system works, lets look at how the system might be deployed. Lets define the problem first, and then look at the solution.

SPECTRUM PLANS & USE CASES

What is Possible?

Full Duplex DOCSIS has two fundamental impacts on the DOCSIS implementation. The first impact is the full duplex operation itself. This includes the echo cancellers instead of diplexers, interference mitigation, TDD or FDD CM operation, and a series of RF challenges. The second impact is just more upstream bandwidth. Maybe 10 to 50 times more bandwidth. That's a lot.

This section looks at the impact of that available bandwidth and how it would be used. To define a product that is deployable, we need the following things to be true:

1. Working CM and CMTS products with the right silicon and software.

2. Working Optical Nodes at the frequencies of interest. This implies a higher frequency and higher power return path.

3. A spectrum plan that makes sense.

Frequency Templates

To illustrate the FDX spectrum plan, it is convenient to use a log 2 frequency base. This allows observation of the many frequencies of interest at lower frequencies. This is shown in Figure 3. Note that as you look towards the right of the figure, the amount of spectrum per inch is doubling. So on the left, an inch might be 25 MHz of bandwidth but on the right it is 400 MHz of bandwidth.

For the sake of calculating throughput and providing comparison between all approaches, the DOCSIS 3.1 OFDM spectrum will be evaluated at 4096 QAM in the downstream and 1024 QAM in the upstream. In practice, some deployments may require lower modulations at the higher frequencies due to lower CNR.

The spectrum graphs will also include basic services such as OOB (Out-of-Band), FM, and Video.

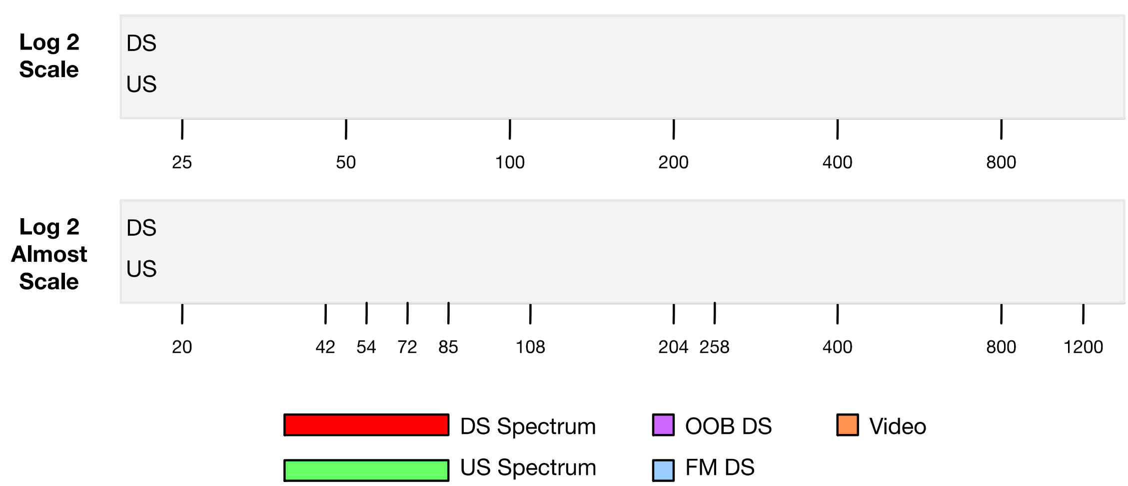

Figure 3 - Frequency Template

OOB Channel

The OOB is very significant as almost all North American upstream expansion plans are limited by the location of the OOB channel. The OOB channel is well over 20 years old and was designed prior to DOCSIS as a two-way channel for managing set-top boxes. The OOB downstream was specified as being located between 72 MHz and 130 MHz. The typical deployment for OOB was the 4 MHz gap at 72 to 76 MHz that is between channels 4 and channel 5.

As we will see, the existence of the OOB channel interferes with the growth of the return path in North America.

FM Band

The FM band in North America is more about interference. Caution is taken that the over-the-air FM band does not interfere with video signals on the cable plant, and that cable signals on the plant do not interfere with over the air FM signals. FM used to be service over cable in North America, but has long been removed.

In Europe, the FM band is an actual service over the HFC plant. Subscribers can plug in their FM radios to the cable plant and receive FM radio. Since it is a one-way service, the cable operators cannot tell how many people actually use the service and thus how many subscribers would be disrupted by the removal of that service.

As we will see, the existence of the FM band interferes with the growth of the upstream band in Europe.

LTE Band

Another band worth noting is the LTE band located between 698 MHz and 806 MHz. LTE shares the same concern that the FM band does. Good HFC plant design with low leakage must be done so that LTE does not interfere with cable signals and cable signals do not interfere with LTE.

Standard FDD Spectrum Plans

In Figure 4, three common spectrum plans that were defined with DOCSSI 3.0 and DOCSIS 3.1 are shown.

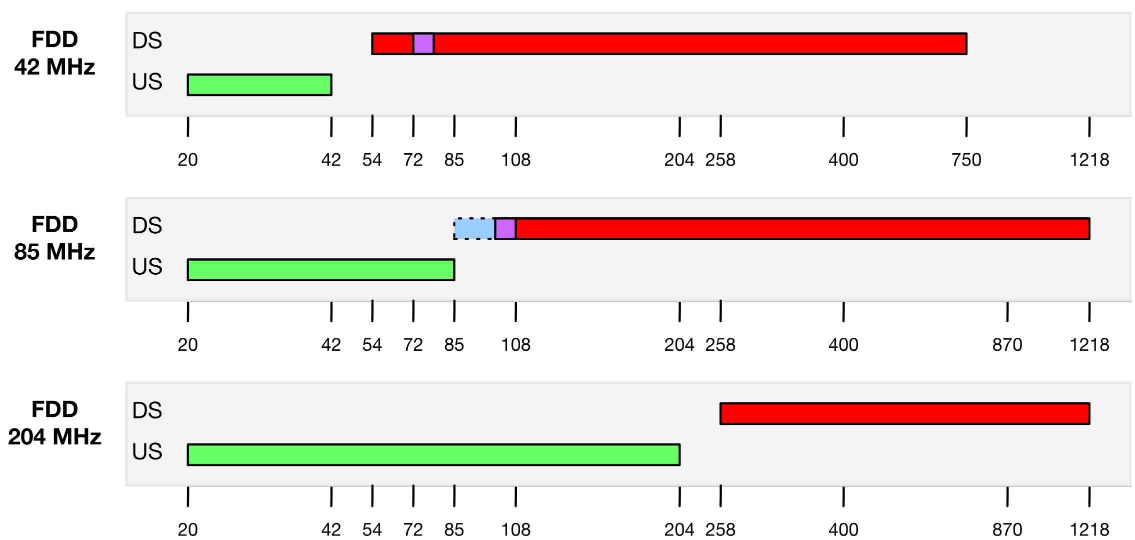

Figure 4 - Standard FDD Spectrums

42 MHz Return

The first plan is the 42 MHz return path. This is the default return path deployed in North America. Europe uses a 65 MHz return path. In the 42 MHz return path, frequencies below 20 MHz are generally not usable. Frequencies close to 42 MHz, say above 40, are often not usable due to group delay. The remaining spectrum is generally good enough to support three 6.4 MHz and one 3.2 MHz D3.0 upstream channels with a throughput of about 90 Mbps.

The maximum downstream frequency depends upon the age of the HFC plant. The example provided in Figure 4 is 750 MHz. Almost all plant is at least 750 MHz today. Newly upgraded plant is 870 MHz and 1002 MHz.

85 MHz Return

This is a new return path option that is just now being deployed. The return path is extended up to 85 MHz and the forward path begins at 108 MHz. Since it is a new install, it can also take advantage of the new optical nodes that can support 1218 MHz.

To accommodate 85 MHz, the OOB channel is moved up in frequency toward the upper end of the FM band, just below the new 108 MHz start of the downstream band. The analog or digital video channels from 54 MHz to 108 MHz must also be removed.

For the North American plant, the additional bandwidth is about 40 to 45 MHz, depending upon group delay limitations. Compared to the previous 20 MHz, an 85 MHz return path has about 3.5x the bandwidth as before. In data capacity, this is about 300 Mbps.

204 MHz return

This alternative return path is 5 to 204 MHz with the downstream starting at 254 MHz. This upstream has 10x the effective spectrum of a 42 MHz return path system with an effective shared throughput on the order of 1.4 Gbps (this is based upon 1024 QAM, 40 MHz to 204 MHz, adding ~1.3 Gbps to the existing 90 Mbps from D3.0 located at 20 MHz to 40 MHz).

1.4 Gbps in the upstream is *5 times more bandwidth that an 85 MHz return path* (1.4 Gbps / 300 Mbps) and *16 times the bandwidth of a 42 MHz return path* (1.4 Gbps / 90 Mbps).

The thing to notice about the 204 MHz return path solution is that there is no support for OOB or for FM. For FDD, both OOB and FM are downstream services and must be removed for the upstream return path to extend to 204 MHz.

Conversely, it is the existence of OOB and/or FM that is preventing many operators from using a 204 MHz return path. This is one of the features that FDX could potentially address.

Video Restrictions

In the beginning, analog video occupied the spectrum from 54 MHz to 552 MHz. Digital video was then added from 552 MHz to 750 MHz. DOCSIS was added either into the 550 to 750 spectrum or above 750 MHz.

Today most, if not all, analog video has been removed from today’s HFC spectrum. The digital video has been moved into the lower part of the spectrum. There are several concerns that limit the location of video.

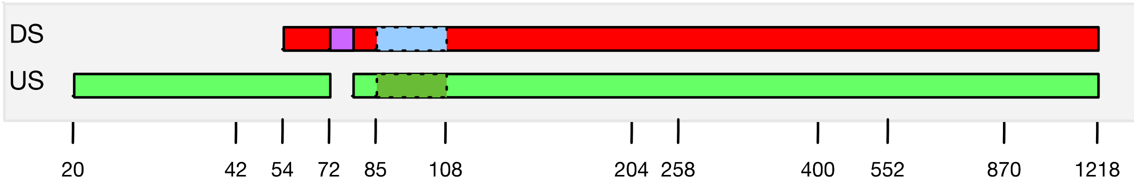

Figure 5 - FDX 1218 MHz with no Video

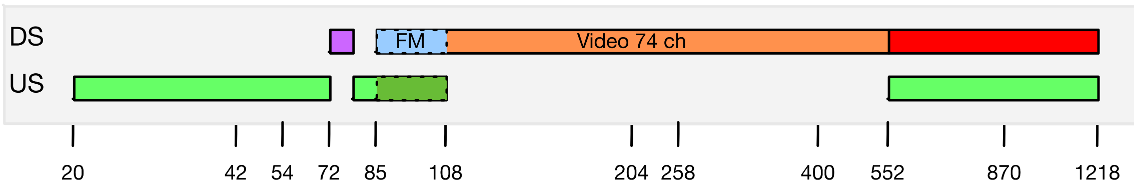

Figure 6 - FDX 1218 MHz with Video

The first is legacy STBs and DTAs (Digital Terminal Adaptors). The upper end of these devices are either 750 MHz or 870 MHz. Some even got to 1002 MHz. However, 750 MHz is the comfortable upper limit.

The next impact is the LTE band. The LTE band starts at 698 MHz and extends to 806 MHz with even more bands beyond that. DOCSIS is more tolerant to a lower CNR (Carrier to Noise ratio) than video, so the general preference is to use the spectrum above 696 (6 MHz band edge) for DOCSIS only.

This means that video – broadcast, SDV, and VOD – generally is between 54/108/258 MHz and 696 MHz.

FDX High Band

Lets start with the maximum FDX use case shown in Figure 5 where the entire upstream would support FDX. One of the rules that will be discussed is that the FDX upstream is not allowed under any broadcast service such as MPEG-TS video. Thus, this solution would not support any MPEG-TS video. All video would have to be video over IP over DOCSIS.

The DOCSIS throughput of this solution could be 11 Gbps x 9.6 Gbps.

This solution would require high power return path amps in the CM that would be able to transmit at 1218 MHz and have enough power to cover the increased attenuation at these higher frequencies.

In Figure 6, a MPEG video band from 108 to 552 MHz (74 channels) has been added. This pushed the FDX action to a high band above 552 MHz. The DOCSIS throughput for this scenario is 6.4 Gbps x 5.8 Gbps. As the MPEG video is decreased, this scenario will approach the previous scenario.

If a video guard band is required (see Section “Video SC-QAM”), there will either be less video channels or less DOCSIS FDX upstream throughput.

FDX Mid Band

Maybe a 10 Gbps upstream is not needed in the near term, or the cost of the amplification and silicon is too much or the power of the return amp is too much. There may be an interim step partially way through the band. As a somewhat arbitrary mid-point, the spectrum in Figure 7 uses 552 MHz. Years ago, 552 MHz was a top frequency in the downstream. 552 MHz is also a good number that matches current analog-to-digital converter performance.

In this scenario, the video is pushed above 552 MHz. To stay below 696, there is only room for 24 QAM channels (at 6 MHz each). This may actually be enough for a plant with small node sizes – say 80 HHP – and where the main video services are over IP and 24 channels are for legacy.

If a video guard band is required, the upstream may be limited to 435 MHz (552 / 1.27).

FDX Low Band

The FDX low spectrum scenario is targeted at existing silicon and existing return path amplifiers, both of which can operate at 204 MHz.

In Figure 8, the upstream spans from 20 MHz to 204 MHz. 2 MHz is skipped to allow the OOB channel at 72 MHz. Below 20 MHz is omitted just to allow proper comparison of spectrum. The FM band may also be skipped for European operation.

In this scenario, it is also illustrated how the 20 to 42 MHz bandwidth may be used by legacy DOCSSI 3.0 and earlier ATDMA services. 42 MHz to 72 MHz are left for non-FDX DOCSIS 3.1 services. OOB is 72 MHz to 74 MHz. Above 74 MHz is FDX. This example also leaves a video guard band between 204 and 258 MHz. There are many variations of these themes.

In this particular scenario, there is room for up to 73 video channels. If all those video channels are used, the data capacity of the system is 7 Gbps x 1.5 Gbps. This is sufficient bandwidth to support a 1 Gbps

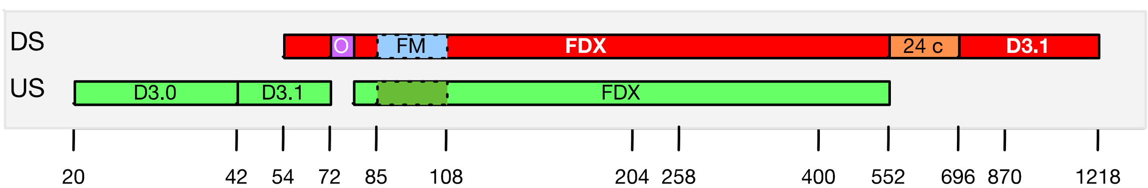

Figure 7 - FDX 552 MHz

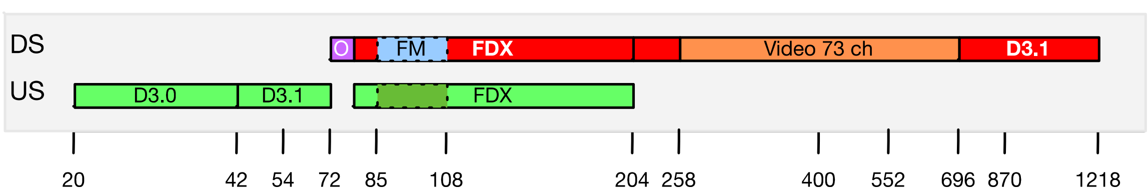

Figure 8 - FDX 204 MHz with Overlap

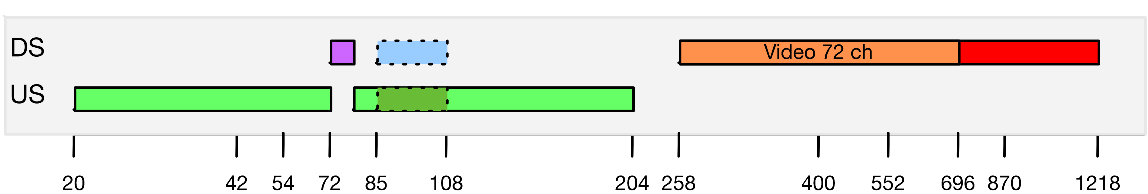

Figure 9 - FDX 204 MHz without Overlap

upstream service tier.

A variant of this channel plan is shown in Figure 9 where these is actually no overlap in spectrum. However, it is still a FDX system as there is no diplexer in the system. As a result, the upstream spectrum can be multiplexed with the OOB and/or the FM services in the downstream.

The throughput of this system, assuming up to 72 channels of video, is 5.2 Gbps x 1.4 Gbps.

This example illustrates that with sufficient downstream bandwidth, it may not be necessary to add complexity to recover downstream bandwidth from 54 MHz to 258 MHz.

In this simplest of deployment scenarios, with no over lapping spectrum, FDX allows an operator to retain OOB and/or FM and also deploy a 204 MHz upstream path.

FDX Spectrum Summary

Table 1 - FDX Spectrum Choices

Option	Low Band	Mid Band	High Band
US Max Frequency	204 MHz	552 MHz	1218 MHz
DS Max Frequency	1218 MHz		
Max Data Capacity	11 x 1.5 Gbps	11 x 4.2 Gbps	11 x 10 Gbps
Cost	$	$$	$$$
TTM	Sooner	Later	Maybe Even Later

In this section, we have discussed three fundamental deployment scenarios. These are a low band, mid band, and high band. These solutions are summarized in Table 1.

There are many variations of these themes. The common consideration fro design, though, it the maximum upper frequency. So a system that extends to 1218 MHz in the upstream is still a high band system, regardless if it uses a 100 MHz or 1000 MHz of spectrum for the upstream FDX path.

In the grand scheme of things, the choice will be driven by the intersection of:

1. CM silicon availability and price
2. Return path bandwidth/power availability and price, for CM, amps, and node.
3. CMTS silicon and throughput capability
4. The needs of the market for higher bandwidth services.

The low band solution is very interesting if existing silicon works, but that depends if higher bandwidth silicon become available before the FDX market is established.

The mid band solution is very interesting in that it delivers more than enough bandwidth for the next few years without stressing out the design elements and without having to deal with the extreme attenuation at higher frequencies.

The high band solution is interesting if technology permits it to come into existence at the right price point. Then we are done.

DEPLOYMENT CONSIDERATIONS

Simplifying Assumptions

N+0 (Deep Fiber)

FDX requires careful use of echo cancellers. Every time there is a 4-wire to 2-wire conversion, the signal from one direction become noise in the other. The easiest scenario to design for would be a passive coax plant where the only 4-wire to 2-wire conversion occurs on the optical node and the CM.

This is not to say that an FDX amp could not be inserted, but success would be declared if the system worked with N+0. (Optical Node plus 0 amps)

By definition, since there is no amplification, the length of the HFC plant is on the order of 500 to 1000 feet.

R-PHY Only

Since the preferred architecture for deep fiber is Remote PHY, a good starting assumption is that the FDX node will be an R-PHY Node. That allows the RPD (Remote PHY Device) silicon to potentially host FDX circuits and allows the FDX to re-use the ADC and DAC components in the RPD silicon.

Same Home Works

This assumption recognizes that when an FDX CM is installed in a home, whatever changes are required to make FDX for that home to work can be assumed. This might extend to a swap out of all the video gear in that home.

This reduces over specifying the FDX solution and focuses the FDX technology on being compatible with equipment in neighbor's homes that are separated by the tap.

Silicon

It is recognized that the CM and CMTS silicon will have to be upgraded to include any protocols changes but also to accommodate the increased upstream speeds.

Passive Plant Model

A passive plant is referred to as N+0 for optical node followed by 0 amplifiers. A node typically has 4 inputs/outputs. Each I/O feeds an independent segment of plant.

Each plant segment is composed of a series of 5 to 7 taps. Each tap has 1 to 8 (typically 4) drop cables to the home.

So, on a maximum dense node, there would be 7 taps x 4 HHP/tap = 28 HHP per I/O . This would be approximately 120 HHP for a node with 4 I/O. Given large neighborhoods, a minimum might be 25 HHP per node with an average being 60 HHP per node.

For FDX, each node I/O is isolated, so and CM to CM interference is isolated to a single segment.

Since the first four to five taps have a higher downstream signal and a lower upstream signal, the interference groups discussed later will tend to be per tap. The last few taps may get combined into a single interference group.

Adjacent Device Interference Analysis

Table 2 - Adjacent Device Interference Analysis

<table>
<tr><th>Scenario</th><th>Different Frequency</th><th>Same Frequency</th></tr>
<tr><td>Same Home</td><td colspan="2">• Impact
• Solve at install time within the home with new equipment
• Allow for QAM video in the spectrum</td></tr>
<tr><td>Adjacent Home</td><td>• Should be no impact
• Evaluate each scenario</td><td>• Impact likely
• CMTS Interference Mitigation Required</td></tr>
<tr><td>Remote Home</td><td colspan="2">• No impact</td></tr>
</table>

When the FDX CM is installed, it may interfere with legacy equipment. That equipment can include other DOCSIS CMs and other STBs. Any interference has be either acceptable or mitigated somehow.

Based upon location, there are three scenarios as show in Table 2. The type of interference can be within the same piece of spectrum (co-channel) or an adjacent piece of spectrum (adjacent channel).

ADI – Same Home, Same Frequency

This only applies to FDX and will be managed using interference mitigation (see section “Interference Mitigation”).

ADI – Same Home, Different Frequency

Devices within the same home as the FDX CM would be the most subject to interference. The default plan is to replace those components with components that work with an FDX CM. Some allowance may still have to be provided for video.

Regardless, there may be a need for a video guard band for MPEG QAM video if the video receiver is relying on the attenuation of a diplexer for it receiver isolation. This would only impact spectrum allocation. This item is under study.

ADI – Adjacent Home, Same Frequency

This is only going to happen with other FDX devices. FDX scheduling will be used to prevent CMs that can interfere with each other with actually interfering with each other. See Section FDX Scheduling” for more information.

If the downstream service is broadcast in nature, the upstream spectrum will be notched. See the section “Downstream Considerations” for more details.

ADI – Adjacent Home, Different Frequency

The testing results in the SCTE 2012 White Paper: “Do You Have Legacy Issues in Your New Upstream Plant?” by John T. Chapman [1] showed that there should be enough attenuation between homes combined with enough out-of-channel rejection and that STB in adjacent homes where not impacted by high power upstream carrier. The reference diagram used in [1] is shown in Figure 10.

This should also be true for CMs, although a guard band may have to be observed. More study is needed.

ADI – Remote Home

For devices located far from the home, either outside of the interference group or even on the other side of the optical node, there is not interference. This case is worth mentioning, but no work needs to be done.

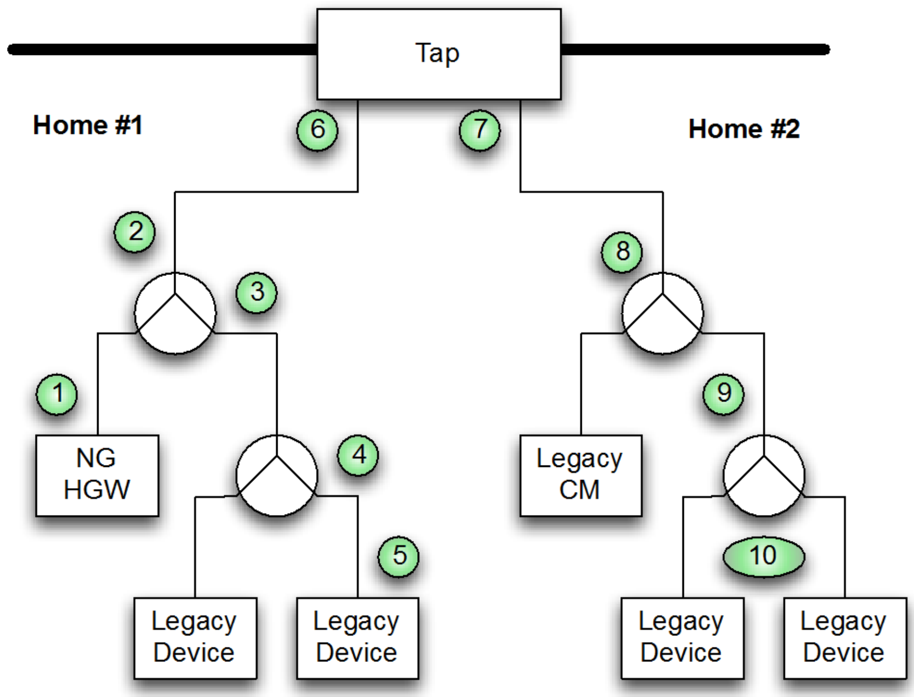

Figure 10 - Adjacent Home Splitter Model

OPERATIONAL DETAILS – PHY

Here is where we get into the nitty gritty of how the PHY layer works. If you are the RF type, this section is for you. If you are not the RF type, feel free to jump to the next section and tell your buddy to read this section.

FDX Operation Mode

A hybrid operation mode is adopted for HFC full duplex network. Same spectrum will be used twice at the same time, one for DS, and one for US. From the RPD point of view, it is a true full duplex operation, it transmits and receives signal on the same spectrum at the same time, but the same spectrum is not assigned to the same CM. So, CM still runs in the simplex mode. This hybrid mode requires minimum changes to CM.

Since HFC runs in a hybrid mode (i.e. full duplex on RPD, simplex on CM), the interference resulting from full duplex operation will be different in the RPD and in the CM. The RPD receiver sees the co-channel interference and adjacent channel interference coupled from its own transmitter.

The CM receiver sees the adjacent channel interferences from its own transmitter and sees both co-channel and adjacent channel interferences from neighboring CM transmitters.

Echo Cancellation at RPD

To suppress the co-channel and adjacent channel interference at RPD, echo cancellation must be implemented at the RPD for supporting full duplex operation. The echo cancellation will be implemented in two stages: the RF stage and digital stage.

Echo cancellation at RF stage

Echo cancellation at RF stage will cancel out the echo before it hits receiver ADC. RF echo cancellation is readily done through port-to-port isolation. Typically one could have a 2:1 combiner that offers 30dB ~ 35dB port-to-port isolation for the whole bandwidth (10 MHz to 1218 MHz).

Higher port-to-port isolation can be achieved. Nevertheless, one needs to consider all the possible paths through which the transmitter signal is coupled/ or leaked to receiver.

Direct coupling

Transmitter signal can be coupled to receiver locally at RPD due to limited port-to-port isolation. Improving port-to-port isolation will help reduce the echo that is coupled directly through the 2:1 combiner.

Reflections

The transmitter signal can go down the coaxial cable and get reflected back at any HFC dis-continuities. In full duplex operation, any reflections of transmitter signal will appear as co-channel interferences on receiver. For example, say there is a tap 100 feet away from the RPD and its input port has 20dB return loss.

A transmitter signal will be then be reflected back with 20dB attenuation. Consider the loss of 100 foot cable (both ways) and a 4dB loss of the 2:1 combiner, when it hits the receiver, the signal level of the reflection will be -25 dB to -30 dB down compared to the transmitted signal level. So, effectively, you will have 25 dB to 30 dB port-to-port isolation, even there is no direct coupling.

One could build a RF echo cancellation to cancel out not only the echo through the direct coupling, but also the reflection. To achieve this, you need an adjustable attenuator with a resolution of a fractional dB, and a tunable phasor with a resolution of <1ns, and they all need to work from 100 MHz to 1218 MHz.

You probably need to cascade a few pairs of them as there may be multiple reflections resulting from multiple discontinuities. As we speak, a tunable phasor with <1 ns resolution and 100 MHz to 1218 MHz bandwidth does not exit.

In summary, echo cancellation at RF stage has its limitation: it can cancel out the direct coupling, but will hardly cancel out all the reflections. In reality, most of echoes seen by the receiver result from reflections, which is in range of 25dB to 30dB down from the transmitter signal level.

A 2:1 combiner with 35dB port-to-port isolation is adequate, as it has suppressed the direct coupling echo under the reflections.

Echo cancellation at digital stage (DSP)

Any echo residue that does not get cancelled out at the RF stage will pass through the ADC and be further reduced through digital echo cancellation. Digital echo cancellation utilizes modern DSP algorithm and can achieve up to 55 dB echo cancellation.

The actual echo cancellation required depends on the echo and input signal levels that in turn depends on transmit and receive signal levels and HFC network condition.

Interference Mitigation at CM

In a full duplex HFC network, CM still runs in simplex mode, which means CM will not transmit and receive on the same frequency/channel at the same time. The FDX is realized through pairing CMs in the network: one CM transmit at certain frequency/channel, and other CM may receive at the same frequency/channel if there are enough isolation among them.

Which CMs can be paired for FDX operation depends on the isolation among them: sufficient RF isolation is required to ensure the receiving CM will not be impaired by the interference of the transmitting CM if they are on the same frequency/channel. The isolations among CMs depend on the HFC topology and will vary case by case.

To sort out specifically the isolation among CMs, a sounding scheme for supporting FDX is proposed. In the sounding scheme, a CM transmits with a known signal level, and all other CMs listen and report the received signal levels.

Based on the reported received signal levels, CMTS scheduler will know exactly the path losses (isolation) among CMs, and which CMs may or may not be paired for FDX operation.

The DS and US scheduler will schedule the DS and US transmissions accordingly based on these pairing information.

Interference suppression at CM

To further suppress the interferences among CM, one could leverage the subcarrier orthogonality of OFDM. OFDM subcarriers remain orthogonal if their frequencies are perfectly synchronized and times are synchronized within cycle-prefix (CP) window. As a result, there are no adjacent

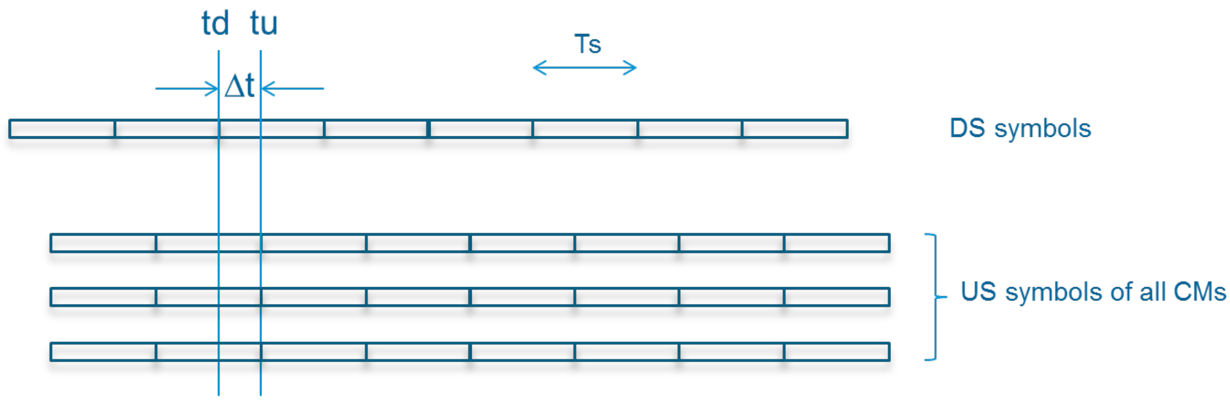

Figure 11 - DS and US Symbol Alignment

carrier leakages among neighboring subcarriers.

Therefore, if there are active subcarriers on certain spectrum and null subcarriers (excluded) on its neighboring spectrum, there will be no energy leakage from the active subcarriers onto the null subcarriers. DOCSIS 3.1 requires subcarrier orthogonality among all the DS subcarriers, and subcarrier orthogonality among all the US subcarriers.

Please note: DOCSIS3.1 requires DS and US are synchronized on frequency but does not require them to synchronize on timing. In reality, DOCSIS 3.1 DS and US timing may have arbitrary offset. As a result, DOCSIS 3.1 DS and US subcarriers are not orthogonal, and DS subcarriers may interfere with US subcarriers if they locate on their neighboring subcarrier, and vise versus.

In the case of FDX operation, a CM receives composited signals that are a superposition of its desired DS traffic and the US interferences from its neighboring CMs.

The subcarriers of the US interference need be aligned with the subcarriers of the DS in frequency and time to make them orthogonal if one likes to utilize the subcarrier orthogonality to suppress the US interferences to DS.

To align DS and US timing, DS and US must use the same cyclic prefix (CP) and symbol lengths and have timing alignment on symbol boundaries. This is shown in Figure 11.

Where:

- Δt is the timing offset that can be set:
- Δt (ns) = 0 to align DS and US symbols at RPD
- Δt (ns) = 2*L(ft)/0.87 will align DS and US symbols at a distance L from RPD

One could set L equal to the max coverage of the RPD, which effectively to make the DS and US subcarriers synchronized at the node edge.

With a single Δt, there will be some offset between US and DS timing at CM, depending on the CM distance to RPD. Given the N+0 deployment case, the offset will be within or close to CP length, thus maintaining ~100% orthogonal (quasi-orthogonal) among DS and US subcarriers throughout the whole node coverage .

This will ensure there are no or little US interference to DS on the adjacent subcarriers.

FDX US to DS pilots and PLC

The subcarrier orthogonality between the DS and the US through the DS and the US timing alignment is required to prevent the DS pilots and the PLC from any US interference. CMs need to lock to the DS pilots and the PLC to remain frequency and time synced with the RPD.

There are three types of DS pilots:

1. Preamble in PLC – 8 symbols x 8 or 16 subcarriers.

2. Continuous pilots – pilots on every symbols and on some fixed subcarriers

3. Scattered pilots – pilots scattered over time and frequency

To avoid the DS pilots from being interfered from any CM US transmissions, the CMTS must designate all the subcarriers that are occupied by the PLC and all continuous pilots as excluded subcarriers for US transmissions. As PLC and continuous pilots take <1% spectrum, leave them un-used for US wouldn't cause much overhead.

As explained earlier, due to the subcarrier orthogonality through DS and US timing alignment, as long as those subcarriers that are occupied by the pilots and PLC are excluded from US transmissions, there will be no or little interferences from US adjacent subcarriers on the continuous pilots and PLC.

CM uses scattered pilots for channel estimation, thus it needs scattered pilots only on the channels where it has DS traffics.

Full duplex DCOSIS enables full duplex operation for the network (same frequency used for both DS and US at the same time), but each individual CM is still on simplex, i.e., its DS and US frequencies do not overlap. Thus CM US does not interfere with its own DS, including the scattered pilots in its DS traffics.

Moreover, CMTS has the knowledge on whose US will interfere with whose DS through the sounding process and the scheduler will avoid assigning any US transmissions that may impair the DS of any CM in its coverage.

The scheduler will assign the US resources in the 2-dimensional time and frequency space and ensure no US/DS overlaps among interfering CMs.

Pilots are just one type of DS signal. In a big picture, full duplex DOCSIS needs to solve the interferences of:

a) RPD transmission (DS) to its reception (US) at RPD

b) CM US to its own DS at CM

c) Adjacent CM's US to adjacent CM's DS at CM

d) CM US to DS broadcast messages.

a) is resolved through echo cancellation, b) is resolved by limiting CM to simplex, c) is resolved through sounding and intelligent scheduling, and d) is resolved through null (excluded) US transmissions on the spectrum where DS broadcasts locate on.

FDX US to OOB

The OOB occupies 2MHz spectrum and centers on the spectrum 72MHz to 76MHz. With FDX operation, one needs to protect the OOB from CM US transmissions.

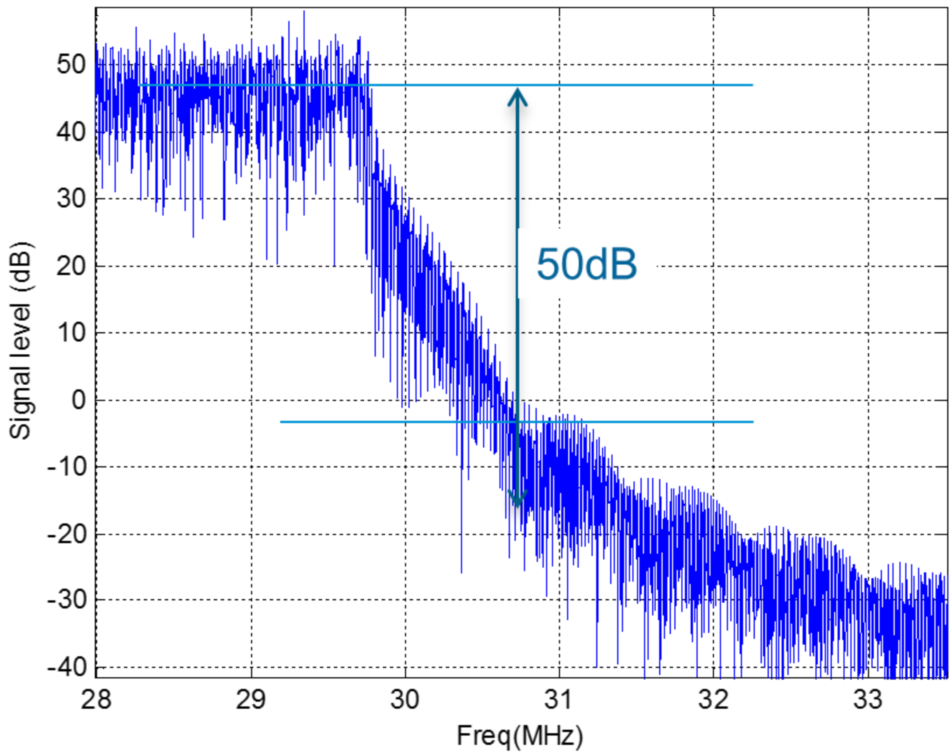

Figure 12 - OFDM Adjacent Channel Leakage

To protect the OOB, the CMTS needs to null the US transmissions on 72MHz to 76MHz, this effectively leaves 1 MHz guard band on both sides. With 1 MHz guard band, OFDM signal leakage into OOB will attenuate by ~50dB.

This is shown in Figure 12. For a CM, the max signal level of US interferences from the neighboring CMs will be 17dBmV/ch (6MHz) (max 57dBmV/ch – 40dB tap-to-tap isolation).

With 50dB attenuation on adjacent channel, OFDM leakage into OOB bandwidth will be -33dBmV/ch. OOB employs low QAM orders with input signal level close to ~0dBmV, it is not expected there will be any issues on OOB reception if the interference level is as low as -33dBmv/ch.

FDX US to FM

The same scheme that is used to protect OOB from FDX US interferences can be used here to protect FM from FDX US interferences.

FM spectrum occupies total 20.5MHz in 87.5MHz to 108MHz. CMTS will exclude US transmissions on 86.5MHz to 109MHz, leaving 1MHz guard band on each side to protect DS FM radio.

FDX US to Video

Similar to OOB and FM, CMTS needs to carve the DS spectrum that is used by video and do not run FDX on that spectrum. As a general rule, leave 1MHz guard band on each side to ensure no or little leakage from OFDM channel into video spectrum.

FDX US to FDX DS – Same CM

In a FDX network, the CM still runs in simplex mode. For each individual CM, the US and the DS will be on different frequency or channel. As a result, CM does not have co-channel interference from its own transmitter.

Nevertheless, this doesn't mean that CM receiver is 100% immune to its own transmissions. For supporting FDX operation, CM needs to have full frequency agility, that is, it's US and DS channels can be assigned on any FDX spectrum, as long as they do not overlap.

For that purpose, the diplexer currently used in CM to separate DS and US spectrums needs be replaced with 2:1 combiner. The common port connect to the HFC coaxial cable, one individual port connects to TX and the other individual port connects to RX.

The US signal will be sent to the HFC through the common port of the 2:1 combiner and DS will be received through the same common port of this 2:1 combiner.

Due to limited port-to-port isolation, the US signal will leak over to the RX. There is a concern that even if the US and the DS are on different channels, and there are no co-channel interferences, the level of the US leakage into the RX may be so high that RX front end may get saturated.

With nominal received signal level 0 dBmV/ch for DOCSIS and 6dB power boosting for video, the total received DS power at CM is 26 dBmV.

The max interferences seen by CM from its own US will be 64 – 35 = 29dBmV, 3dB higher than the nominal DS signal at the CM; where 64 dBmV/ch is the max CM output power, 35dB (average) is the port to port isolation at the CM.

The CM may require some tweaks to its front end to ensure its receiver will not be saturated and can maintain proper operation with the presence of intermittent interference which are 3dB higher in power than its nominal DS received power.

FDX US to FDX DS – Neighboring CM

In an FDX network, the intelligent scheduler will co-schedule DS and US to ensure there are no co-channel interferences to DS. The intelligent scheduler utilizes the sounding scheme and interference groups to achieve this purpose.

Nevertheless, CM can still see US interference on adjacent channels. This could occur with CMs on the same tap or on different taps. A brief calculation shows this adjacent channel interference would not be an issue.

The max interferences seen by CM from its neighboring CM's US will be 64 – 4 – 39 = 21 dBmV, 5 dB lower than the nominal DS signal at the CM; where 4 dB is the 2:1 combiner loss and 39 dB is the minimum isolation among CMs (tap-to-tap isolation).

CM should be able to maintain its normal operation, given the interference level is a few dB below its nominal DS receiver level. In the case where new CM (FDX CM) and legacy CM (non-FDX CM) co-exist in the same network, one could keep the legacy CM on FDD and only operate FDX on new CM.

A new CM may be assigned to transmit on the spectrum that is a part of the DS spectrum for legacy CM as long as the legacy CM does not receive any DS traffic on that spectrum.

Downstream Considerations

When the same spectrum is used for both upstream and downstream transmission, coordination is required. The general scheme used in this white paper is that any one CM does not transmit at the same time on the same frequency.

This poses some challenges for any downstream broadcast/multicast or common upstream spectrum usages. Lets go through each use case and determine the best course of action. Most of these cases will result in upstream exclusion bands and are illustrated in Figure 13.

OOB

The out-of-band channel, as defined by [2] and [3], typically is located at 72 to 74 MHz, which is a gap between channels 4 and 5 [4]. The OOB channel is used to control legacy MPEG set-top boxes. There is also an upstream OOB channel typically located at 8 to 12 MHz.

In a conventional FDD system such as DOCSIS 3.1, the existence of the OOB downstream channel often limits the upstream spectrum to an 85 MHz return path.

The plan would be to notch out the upstream spectrum under the downstream OOB channel. There should not be a need to have guard band.

Since the upstream OOB is at a sufficiently low frequency, it is not impacted by the FDX operation other than having to be included in any upstream digitization of the spectrum.

FM

FM is a service on the downstream plant. Years ago, FM was a service in North America. However, that is generally not the case now. Europe still has FM as a service. Since FM is a one-way service, there is not easy way to determine how many users exist. So it is hard for European operators to judge the impact of removing FM as a service.

In a conventional FDD system such as DOCSIS 3.1, the existence of the FM band often limits the upstream spectrum to an 85 MHz return path.

The plan would be to notch out the upstream spectrum under the FM band. There should not be a need to have guard band.

DOCSIS 3.0 SC-QAM

DOCSIS 3.0 deployments today have

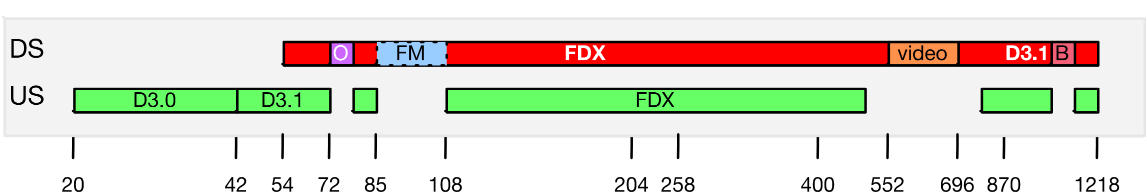

Figure 13 - Downstream Exclusion Bands

between 8 and 32 QAM carriers on the plant. If the FDX protocols, specifically the sounding and the scheduling cannot be retrofitted back into DOCSIS 3.0 and earlier CMs, then the upstream spectrum below DOCSIS 3.0 cannot be used.

If the FDX protocols can be extended to DOCSIS 3.0, then the DOCSIS 3.0 traffic would be broken up into unicast QAMs and multicast/broadcast QAMs.

The unicast QAMs would participate in the FDX protocols, while the broadcast and multicast QAMs, such as primary QAM channels that contain MAPs and other MMM would have spectrum notched out in the upstream.

2 MHz guard band is required between DOCSIS 3.0 and DOCSIS 3.1 per DOCSIS specifications. This 2MHz guard band rule shall be applied here as well to separate DOCSIS 3.0 spectrum and FDX spectrum.

DOCSIS 3.1 Non-FDX

DOCSIS 3.1 in a non-FDX mode is just being deployed. If the FDX protocols – namely sounding and scheduling – can be applied to non-FDX DOCSIS 3.1 CMs and there are no RF issues, then it should be possible to use the upstream bandwidth below the unicast DOCSIS 3.1 spectrum. Otherwise, non-FDX DOCSIS 3.1 spectrum would have to be excluded.

DOCSIS 3.1 PLC

The PHY Link Channel (PLC) is a broadcast channel that is used for CMs to register. It has a preamble for syncing, a timestamp, and some basic signaling. The upstream spectrum must be notched out below the PLC channel.

For supporting FDX operation, DS and US subcarriers need to maintain orthogonality by synchronizing DS and US timing. With DS and US subcarriers orthogonality, no additional guard band is required around PLC subcarriers.

DOCSIS 3.1 NCP

The Next Code Word Pointer (NCP) is a challenge. The NCP exists initially in every symbol and points to the beginning of the next code word. The NCP is then scrambled across time and frequency due to the interleaver process. This makes it hard to omit upstream bandwidth for the NCP field.

The easiest solution is to allow the CM to ignore the NCP when the upstream transmission trashes the downstream transmission. This requires some research to see if this is acceptable. There is one aspect of the NCP where there is a countdown when the NCP modulation is going to change. Missing that count down could be problematic.

Both these problems could be addressed

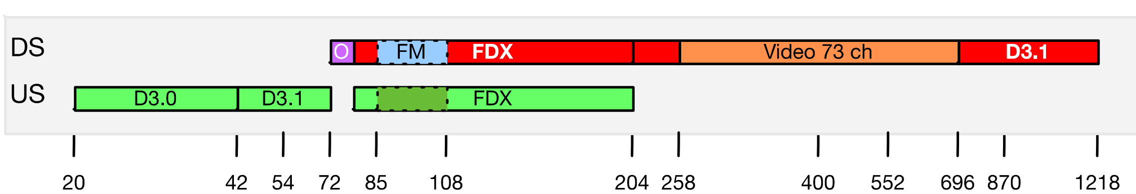

Figure 14 - Video Guard Band

by introducing additional signaling in the PLC channel. This is an issue that requires further study.

Video SC-QAM

Video has three forms – linear which is broadcast in nature, Switched Digital Video (SDV) which is multicast in nature, and Video on Demand (VOD) which is unicast in nature. All these three forms are uncorrelated with DOCSIS and hence with potential FDX protocols. As such, all upstream bandwidth at the same frequency of video must be excluded.

There may be a necessity to have a video guard band as explained in section "Video SC-QAM". If this is the case, here is how to calculate the guard band. The amount of guard band is dependent upon the order of the filter used in the diplexer. This means that it takes a certain amount of frequency band to create enough isolation.

The classic diplexer ration is 1.27. For example:

54 / 42 = 1,285
108 / 85 = 1.27
258 / 204 = 1.265

These ratios vary slightly as the frequency cut-offs are rounded to the nearest 6 MHz channel boundary. So, if the upstream band extended to 204 MHz, video could not start until 258 MHz (204 * 1.27). If video stopped at 696 MHz, and an upper guard band is required, the upstream spectrum could not restart until 882 MHz. (696 * 1.27 rounded down to a channel boundary).

It is not clear at this time that an upper video guard band or even a lower video guard band is required. This is under study.

Another factor that comes into play is the migration of video services to all IP. If a home that upgrades to FDX DOCSIS can also upgrade to all video over IP, then no video guard bands may be required.

Downstream Data Path Traffic

In the earlier section, we covered layer 1 and lower layer 2 messaging such as PLC and NCP. In this section, we will look at the layer 2 and layer 3 traffic that are in the data path. The focus will be on DOCSIS 3.1 traffic.

Downstream Broadcast

This would include the MAP message that is sent to all CMs. It would also include certain MMM (MAC Management Messages). All these are contained in profile A. However, we cannot cut out the spectrum below Profile A as Profile A is multiplexed with the other profiles, and the result is then interleaved across time and frequency.

The only reasonable solution seems to be to isolate the traffic from profile A onto its own channel with its own time and frequency assignment, with its own interleaver. Another solution would be to put the MAP and MMM on the PLC channel, although this might increase latency and exceed that bandwidth of that channel.

Both of these approaches require a change to the usage of the DOCSIS 3.1 protocol. For backwards compatibility, it would probably make sense to make this part of the standard DOCSIS 3.1 operation.

Downstream Multicast

Downstream multicast traffic would refer to any IP multicast traffic over DOCSIS. This typically might be a video stream. It also might be some sort of IP signaling such as an IPv6 signaling multicast. There are at least two solutions.

The first solution is to put all multicast traffic in with the broadcast traffic in a dedicated DOCSIS 3.1 channel. This is simple. It is also useful if you believe that there is not much multicast traffic to begin with, so optimization is not needed.

CMs are grouped into interference groups (IGs). Those IGs could be further sorted to line up with multicast groups. This is a bit more complex, but not impossible.

Upstream Considerations

Upstream Contention Requests

During a contention interval, any CM may send a request. To avoid interference from contention requests, CMTS needs to assign sub-slots for contention requests and exclude the corresponding DS subcarriers from DS usage to avoid interferences from contention requests to DS.

To avoid overhead, the sub-slots will be assigned on fixed mini-slot locations over time.

This would be a *downstream* exclusion band.

Upstream Initial Ranging

During initial ranging, any CM may send a request. To avoid interference from initial ranging, CMTS needs to designate initial ranging zone and exclude the corresponding DS subcarriers from DS usage to avoid interferences from initial ranging signals to DS.

Initial ranging zone does not have a frequency permutation with the rest of US mini-slots. This helps to locate the corresponding DS subcarriers that need be excluded to avoid interferences

To avoid overhead, the initial ranging zone should take a narrow width on subcarriers (take one mini-slot)

This would be a *downstream* exclusion band.

Interference Groups

Sometimes there are just noisy neighbors and you are stuck with them. When they party, you can't hear yourself think. The same is true in the tiny world of FDX.

If one CM is too close to another CM, its upstream transmission may interfere with the other CMs downstream reception. And since the interference is coming from another CM, a classic echo canceller cannot cancel the signal out.

FDX deals with situation using discovery followed by mitigation. Basically, FDX finds the noisy neighbors and forces them to get along.

Interface Group (IG) Discovery

The CMTS is generally not aware of the physical topology of the HFC plant. It may know the name of a particular node group and can attach CMs to that node group, but it does not have observably on which tap a CM is on or how long the cables are.

There is only one real way to determine if one CM will interfere with another CM and that is to test for interference. To do this, one particular CM will send out an interference signal and the other CMs will listen for it. If the transmit power is known and the receive power is measured, an actual CM to CM path attenuation can be calculated. If signals are sent at different frequencies, tilt can also be calculated. Multiple tests will probably be needed that correspond to each channel boundary.

The interference group testing will either re-use existing signaling within the DOCSIS 3.1 complex – such a PNM – or new signaling can be created.. It will be important for the CM to be able to accurately measure the power of the received signal. One concept would be to reserve specific

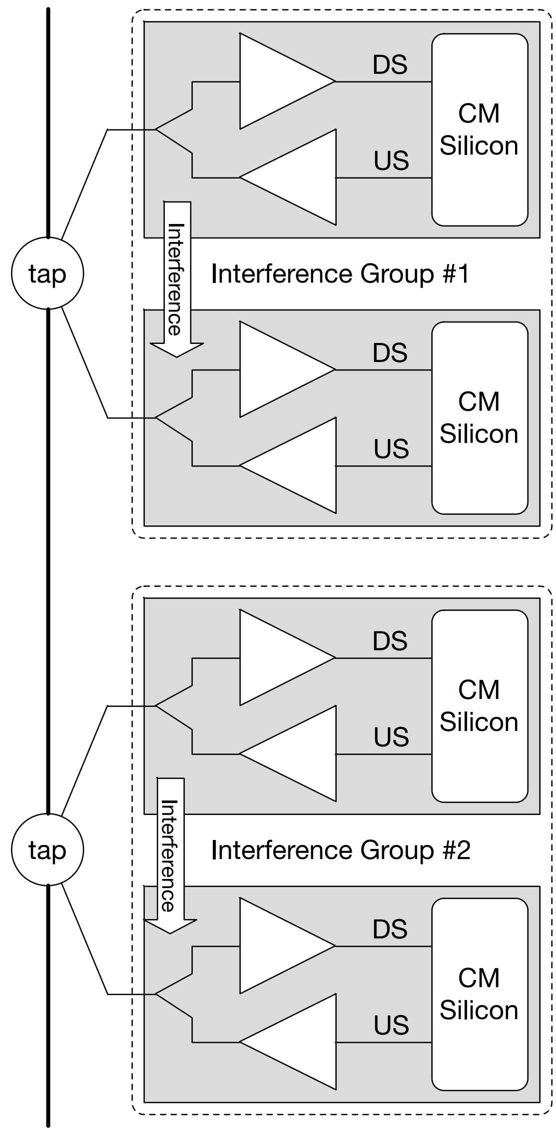

Figure 15 - Interference Groups

frequencies that are only used for IG Discovery. That way there will not be any artificial background noise.

When interference is discovered between two CMs, those CMs are placed into an interference group. Typically, it is expected that each tap group will also be an interference group. The last few tap groups at the end of a line may also group together into an interference group.

Transmission Groups (TG)

IGs are based on actual measured interference between neighboring CMs. If there is good isolation between taps, there may be quite a few independent IG. For scheduling convenience, it may be convenient to have a minimum number of groups.

This goal is achieved by assigning interference groups (IG) into transmission groups (TG). So, whereas IG are groups based upon a physical property, TGs are based upon a logical need.

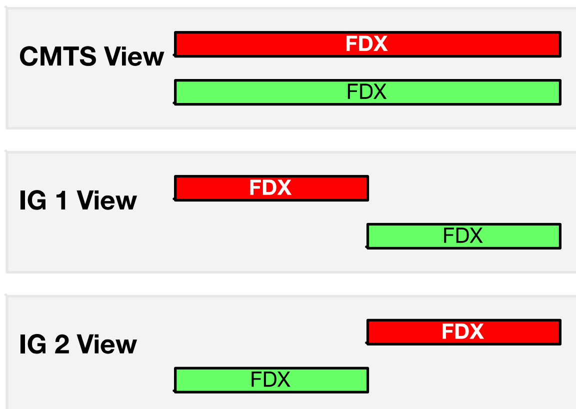

Figure 16 - Duplex - Simplex Scheme

For example, a 80 HHP node may have 20 taps. If each tap were a separate IG, then there would be 20 IGs. For scheduling purposes, a simple scheduler may only need two TGs. IG are assigned to TGs based on traffic density, multicast groups, or any number of criteria determined by the CMTS.

For nomenclature, it is proposed that IGs are numbered where as TG are assigned capital letters. So IG-1 and IG-2 may belong to TG-A.

Interference Mitigation

The rule is that *within an interference group one CM is not allowed to transmit at the same time and frequency that another CM is receiving.*

This rule is enforced by the CMTS through scheduling. To schedule properly, there needs to be a frame work to schedule within and rules that can be enforced. This is what we will explain next.

Duplex-Simplex Scheme

What this really implies is that *a CM IG operates in a simplex mode while the overall plant operates in a duplex mode.* This is illustrated in a simple way in Figure 16. Figure 16 shows frequency allocation within a spectrum at a specific point in time.

So, within an IG (or a TG), the CMs cannot transmit and receive at the same time on the same frequency. They may transmit and receive on different frequencies (FDD) or at different times (TDD). Meanwhile, while one IG is receiving only, another IG can transit only. This is represented in Figure 16 with IG View #1 and IG View #2. This is what is implied with simplex within an IG.

When scheduled right, the entire downstream spectrum, at each point in time, is used for CM reception and the entire upstream spectrum can be used for CM transmission. This is duplex. This is represented in Figure 16 with the CMTS view.

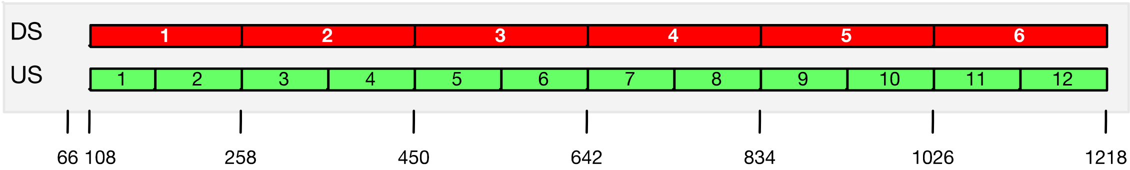

Figure 17 - High Band Channel Plan

There are three fundamental ways to create these simplex partitions: FDD, TDD, and a mix of FDD/TDD. Note that Figure 16 is equally applicable to FDD where the horizontal access is frequency, or TDD, where the horizontal axis is time.

FDX – FDD

This would be spoken as "Full Duplex with FDD". FDD is Frequency Division Duplex and would the technique for enforcing transmission rules for TGs and hence IGs. FDD is how DOCSIS and the cable plant works today. For the sake of illustration in the following scenarios, a full spectrum system is assumed.

FDX Channel Plan

So, how many channels are required for an FDX full spectrum system?

In the downstream, the answer is six. Six 192 MHz OFDM channels will span from 1218 MHz down to 66 MHz. If there is video to be skipped over, the number might be less.

If a dedicated broadcast channel of say

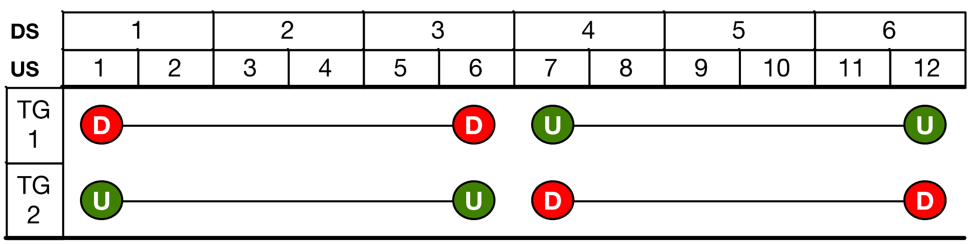

Figure 18 - FDX-FDD With 2 TG

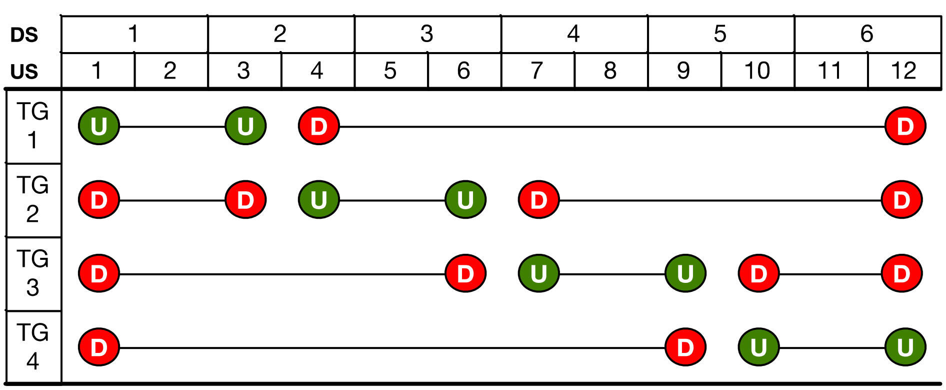

Figure 19 - FDX-FDD with 4 TG

24 MHz is used, then one more channel would be needed (so seven). Note that 66 – 24 = 42 MHz, which conveniently is the upper part of the existing upstream spectrum.

In practice, FDX would not go below 42 MHz so as to not disturb legacy DOCSIS 3.0/2.0 traffic.

In the upstream, the answer is twelve since each OFDMA channel is 96 MHz, which is half the width of a downstream OFDM channel.

Other combinations or partial channel maps are possible. However, Figure 17 will be used for illustrative purposes.

Scheduling a FDX-FDD System

In Figure 18, the top row shows 6 downstream channels and 12 upstream channels. On the left are transmission groups (TG) one and two. The barbell style diagrams are bonding groups.

In this example, each TG is given 50% of the bandwidth for downstream and 50% of the bandwidth for upstream.

Within a TG, different frequencies are used for isolating the downstream and upstream transmission. Between TGs, it does not matter. Thus, upstream transmissions will never interfere with downstream reception.

From the CMTS viewpoint, 100% of the spectrum is used for both downstream and upstream transmission.

What if the SLA that you are trying to create is not symmetrical? Figure 19 shows a system with four TGs where the upstream is 25% of the bandwidth and the downstream is 75% of the bandwidth within each TG. Overall, though, 100% of the spectrum is used for downstream and upstream.

FDX – TDD

TDD – Time Domain Duplex – is another scheme all together. It is not compatible with the current DOCSIS architecture. Let analyze what it would look like and then see if it is needed.

TDD at the CM would mean that a CM, or rather the IG or TG associated with that CM, is either transmitting across the full OFDM channel or receiving, but not both. That decision to transmit or receive would come from scheduling at the CMTS.

To do that, the CMTS has to know that there is downstream traffic on the wire for the TG that cannot be interrupted. That is not possible today. Today, the CMTS does not really know what is on the downstream wire and when. Take a look at the DOCSIS 3.1 system today in Figure 20 (source: [5]).

When the downstream packets hit the DOCSIS QoS engine, they are queued and potentially delayed. This is the first delay. The profile buffers and the code word builder are the second delay. Then there is the time and frequency interleaver which really slices and dices things. Even if you could predict where a packet was, its bits are now scrambled all over the place.

The DOCSIS upstream is better. You actually know exactly what packet from which CM will be sent at what exact time. You just have to calculate it 2 ms or so in advance.

The only way to have a TDD system is to construct one.

First, the downstream and upstream symbols have to be aligned. This means that the downstream and upstream channels that are occupying the same spectrum will use the same cyclic prefix.

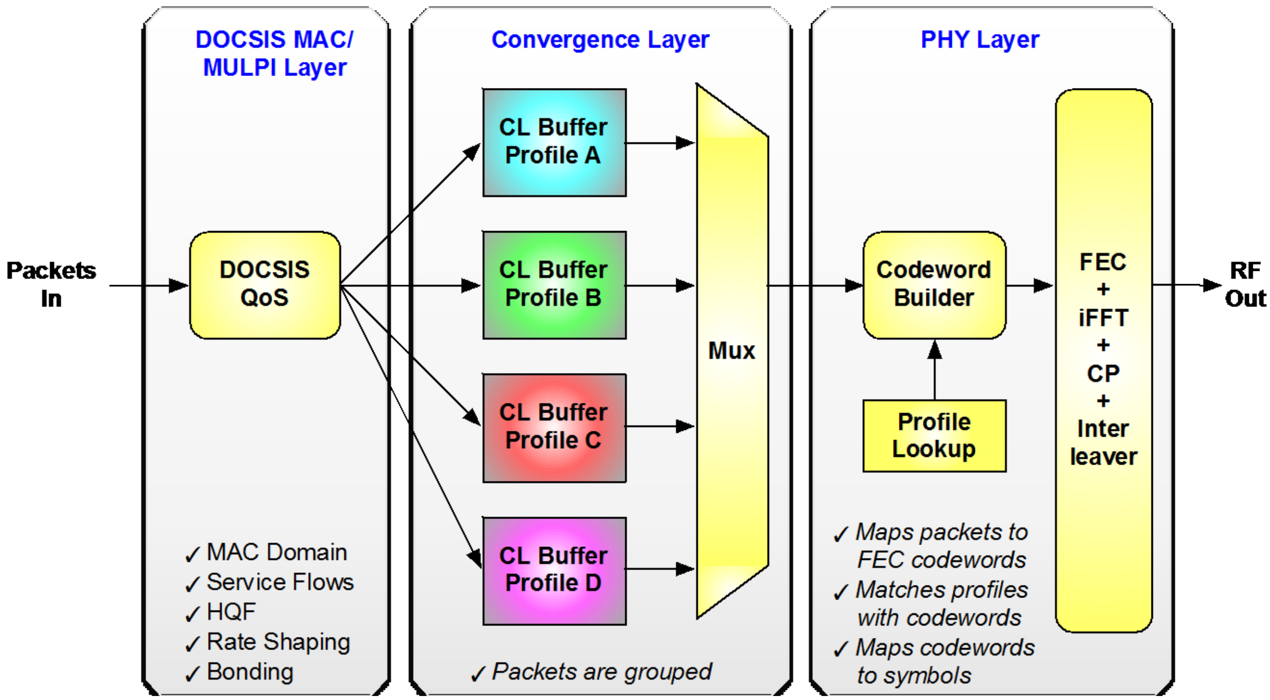

Figure 20 - DOCSSI 3.1 MAC-PHY

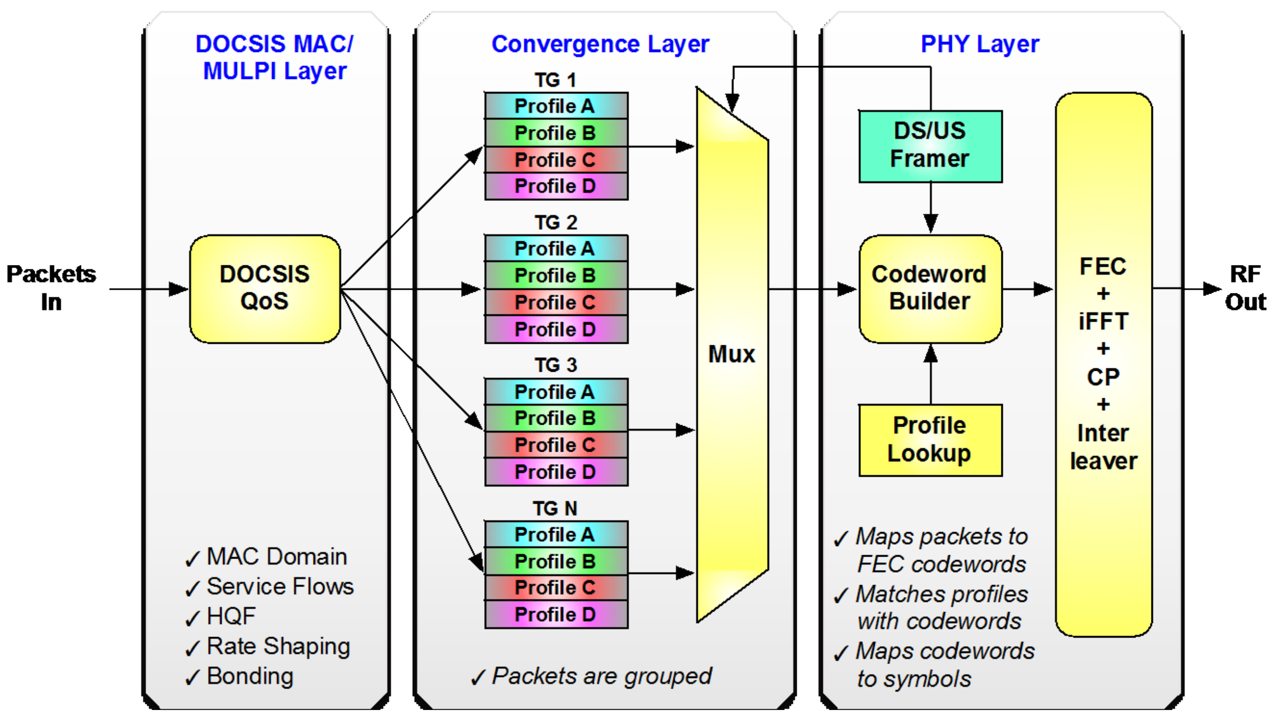

Figure 21 - DOCSIS 3.1 MAC-PHY with TDD

A frame structure has to be introduced. Since the downstream already has inter-leaver boundaries and the upstream already has minislot boundaries, a frame structure that recognizes the superset of both of these is required.

Next, the queuing and profile hardware needs to support these frames boundaries. A proposed system is shown in Figure 21.

When the packets are released by the DOCSIS QoS stage, they are directed towards a downstream TG. Each TG has its own profile buffers. There is than a framer that pulls an entire frame contents from one set of profile buffers.

The assignment of DS and US frames to TGs is dynamic and can be done based upon the relative bandwidth that each TG requires.

FDD vs TDD

So, is a TDD system for TGs needed if there is already a FDD system available? It probably boils down to granularity. An FDD system has the granularity of channels and might break the available bandwidth into two or for frequency chunks.

A TDD system would break the bandwidth into multiples of frame boundaries. It is quite possible that a TDD system would have better granularity than an FDD system. However, in a large system, the granularity of a FDD system may be enough.

Ironically, a TDD system must support more channels than an FDD system. For a given bandwidth resource block, an FDD system can use less channels but is continuous in time.

A TDD system uses less time and therefore requires more channels. For example, for the same throughput, an FDD system may use one channel 100% of the time while an equivalent TDD system might use two channels 50% of the time.

FDX – FDD/TDD

It is worth noting that FDD and TDD systems could be combined to provide a three dimensional matrix of time, frequency and DS/US space, all under full scheduler control. The theoretical ability to do this would have to be compared to the practicality of doing this. But that discussion is for another time and place.

CONCLUSIONS

10 Gbps x 10 Gbps theoretically possible for DOCSIS. That put DOCSSI on equal footing with 10 Gbps Ethernet and 10 Gbps PON.

The optical node silicon (which is probably the CCAP silicon for an RPD) have to support an echo canceller. However, the biggest impact of FDX is more upstream throughput. This means more OFDMA channels and more CPU cores for scheduling.

CMTS and HDC plant are FDX. The entire spectrum is used for both downstream and upstream transmission. Each CM, however, is simplex using FDD or TDD. (partial spectrum or partial time is possibly a more accurate way of describing the concept.)

The design is being optimized for R-PHY Node and N+0 deep fiber. This is a realistic deployment plan going forwards and represents the minimum viable product.

Getting some proof of concepts, writing the specifications, creating the silicon, building the product, and establishing the market will determine availability of FDX. The next step then is to create a plan to match silicon and product availability with market requirements.

REFERENCES

Related Papers by the Same Authors

[1] John T. Chapman, Do You Have Legacy Issues in Your Upstream Plant? Dealing with ADI and OOB on a High-Split Plant, SCTE Cable-Tec Expo 2012, Oct 17, 2012, http://www.johntchapman.com/Chapman-LMA-(2012-08-30).pdf

Related Industry Specifications

[2] SCTE 55-1 2009, Digital Broadband Delivery System: Out of Band Transport Part 1: Mode A, SCTE, http://www.scte.org/documents/pdf/standards/SCTE%2055-1%202009.pdf

[3] SCTE 55-2 2008, Digital Broadband Delivery System: Out of Band Transport Part 2: Mode B, SCTE, http://www.scte.org/documents/pdf/standards/ANSI_SCTE%2055-2%202008.pdf

[4] North American Television Frequencies, Wikipedia, https://en.wikipedia.org/wiki/North_American_television_frequencies

[5] DOCSIS 3.1 Physical Layer Specification, CM-SP-PHYv3.1-I08-151210, CableLabs, Dec 12, 2015, http://www.cablelabs.com/wp-content/uploads/specdocs/CM-SP-PHYv3.1-I08-151210.pdf