Network Preparation: Maximizing Capacity ROI

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Dr. Robert L Howald
(Formerly) Fellow of the Technical Staff
ARRIS
robert_howald@cable.comcast.com

Jack Moran
cablejack.moran@gmail.com

Robert Thompson

Ken Couch
Director of Marketing
ComSonics
kcouch@comsonics.com

Daniel Howard
SVP and CTO
SCTE
dhoward@scte.org

Abstract

Cable's hybrid fiber coax (HFC) downstream channel represents perhaps the world's most capable radio frequency (RF) channel in terms of capacity. Theoretical capacity boils down to two simple ingredients: Signal-to-Noise ratio and bandwidth. The cable downstream has relatively large amounts of both. The cable upstream is less so enabled. However, this is largely by choice of spectrum allocation and architecture, and is more historical artifact than inherent limitation.

Nonetheless, HFC remains far from optimally exploited. Given the persistence of compounding traffic, more efficient use of precious HFC resources must be realized. Initiatives in the cable industry (DOCSIS®) and telecommunications community at large (IEEE 802.3bn) are addressing this head-on. Modern digital communication tools obtain performance quite close to theoretical limits, and cable will take advantage of these tools – multi-carrier techniques, more bandwidth efficient quadrature amplitude modulation (QAM), and modern forward error correction – as well as continued architecture evolution to maximize HFC capacity. This paper will quantify downstream and upstream capacity as a function of architecture, spectrum scenarios, HFC performance, and emerging HFC technologies. We will also associate this capacity to projected user bandwidth over time.

While capacity is very insightful, reality can be much more eye-opening. In practice, any channel's ability to maximize capacity is also tied to its ability to conquer network imperfections and non-additive white Gaussian noise (AWGN) noise sources. The range of cable's exposure to each is quite large, and resulting implications to actual capacity is effected accordingly. We will quantify these effects based on field measurements of various forms of interference and observed network operating characteristics, consider how these new tools are impacted, and compare to capacity theory. From these insights, we provide HFC Readiness recommendations. Our guidelines will enable operators to get the best capacity bang for their investment buck on the road to optimizing their valuable HFC assets.

CAGR versus Capacity

Bandwidth: How Much is Enough?

The persistent growth of downstream HSD bandwidth has raced along at a compounding rate of about 50% per year (referred to as 50% compound annual growth rate or CAGR). This presents tremendous challenges for operators with capacity limitations in the HFC plant and a downstream spectrum that is already fully utilized.

Figure 1 shows a commonly used reference [2] that identifies how Internet bandwidths to the home have charged ahead for basically the entire history of "broadband" access – where obviously the definition of "broadband" has dramatically changed over the years. Figure 1 represents the problem statement for multiple system operators (MSOs) who subscribe to the inevitability of continued bandwidth growth.

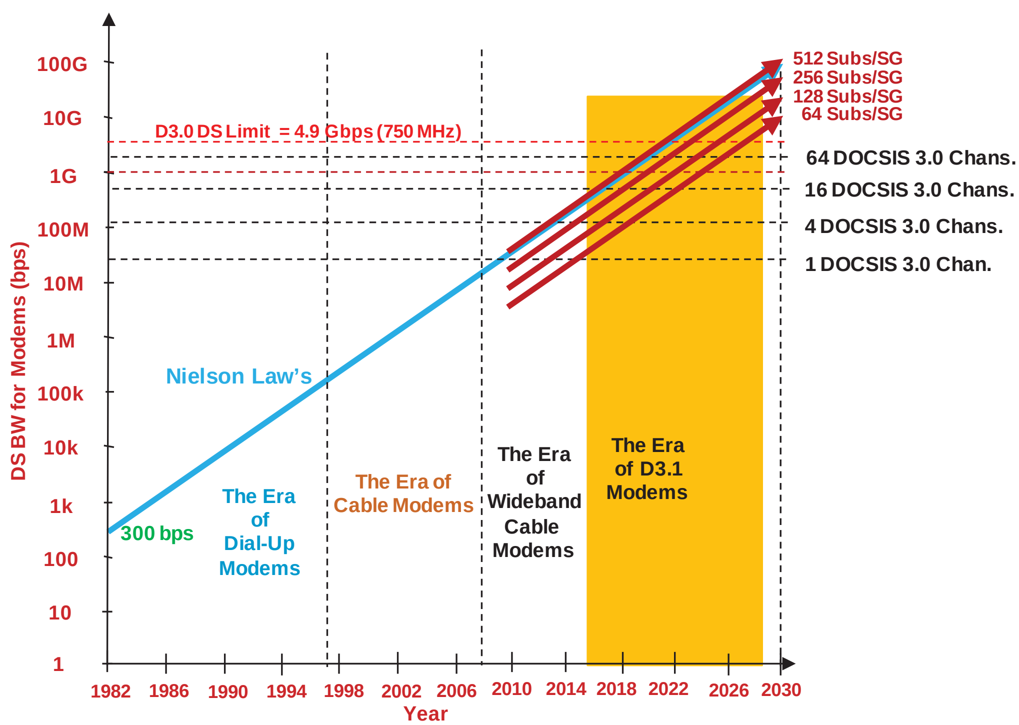

Figure 1 – Bandwidth Growth to the Home Has Been Aggressive and Consistent

So, how much bandwidth is enough? Well, there is no winning a battle against CAGR if growth persists without end – you can only hope to offset the growth in a way that makes network investment manageable as the growth continues. And, unfortunately, Figure 1 does not even show the full magnitude of the problem. While it captures DOCSIS® 3.0 expansion, and we can estimate what DOCSIS® 3.1 and additional spectrum will add, the fact today is that the vast majority of HFC spectrum is *not available* for DOCSIS® growth due to legacy services consuming that spectrum.

So, no matter how attractive we can make the end game lifespan of HFC bandwidth through the combination of spectrum, technology, and architecture, there is a very complex capacity management problem to navigate while preserving revenue and investing prudently. This multiple variable problem can be quantified, and many detailed analysis and use cases for capacity management have been discussed in prior work. These analysis cover examples that consider network lifespan as far out as the 2030 time frame [5][8][11][12][14]. Some MSOs even argue that the lifespan of HFC could extend for another 30 years which means we must consider CAGR effects as far out as 2040.

Lastly, there is the philosophical question of "Forever" CAGR versus Tapering CAGR. One possible consumer-centric reason for a tapering CAGR are the limitations of human sensory experiences and the limited ability of humans to consume media content and multi-task effectively..An operator-centricpossibility is the gradual removal ofhigh data rate users such as business customers off of the HFC network and surgically moving their service onto fiber to the premises (FTTP) solutions such as radio frequency over glass (RFoG) or metro Ethernet service. There are enormous architectural implications to how such variables actually play out. The referenced papers include scenarios that cover both angles, and with differing assumptions on the service evolution stimuli, using a Capacity Management Timeline approach. Below we summarize the Capacity Management Timeline, and briefly discuss some of the key concepts, conclusions and relevance to DOCSIS® 3.1.

<u>The Capacity Management Timeline: The Intersection of BW Growth and HFC Capacity</u>

Figure 2 is an example of a tool that takes the growth referenced in Figure 1 and maps it against the capabilities and architecture evolutions of the HFC network.

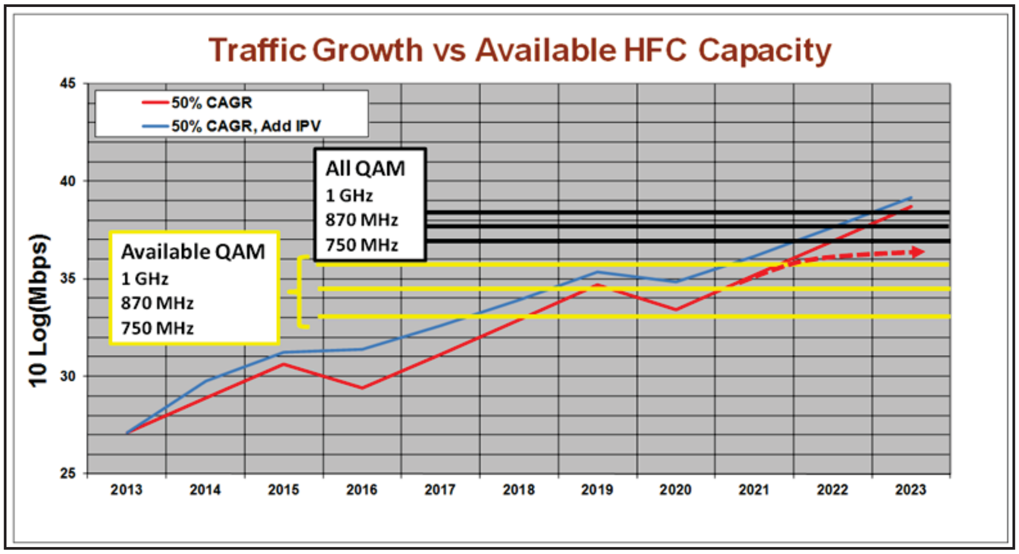

Figure 2 – Persistent Downstream Bandwidth Growth and its Relationship to HFC Capacity Thresholds

The red trajectory represents a 50% CAGR from a starting point of 10 downstream DOCSIS® channels, while the blue trajectory represents a 50% CAGR but while simultaneously setting aside additional DOCSIS® channels, for example, for managed IP Video (IPV) services (if such an approach is preferred).

Breakpoints in the trajectory for both red and blue represent ideal node splits. Note that the blue trajectory does not drop by the same amount as the red at a node split year. This is because of an assumption of an incrementally growing set aside of IP Video bandwidth as the segmentation is taking place. The black thresholds represent the available bandwidth for growth if *all* of the spectrum were allocated to DOCSIS®, for three different examples of forward path bandwidth, and assuming all of these channels were carrying 256-QAM.

The yellow thresholds represent a scenario where it is recognized that not all of the spectrum can be set aside for DOCSIS® – some must remain to support legacy services. Many different possibilities exist here – is there analog video or not and how much over time, is there switched digital video (SDV) or not and how much, how much video on demand (VOD) is consuming spectrum, etc. The main point is simply that MSOs do not have free spectrum – they must find a way to create it with various downstream tools and levers. These variables in play can be quantified and the actual capacity available for growth determined.

<u>The Role of DOCSIS 3.1</u>

On its very best day, DOCSIS® 3.1 will enable 50% more bandwidth efficiency than prior DOCSIS® versions or of current MPEG-TS digital television (DTV). Both, of course, use 256-QAM. Furthermore, by going to 8k or 16k QAM, DOCSIS® 3.1 could increase to 75% more bandwidth efficiency over time.

As with other knobs and levers for capacity analysis, the effect of DOCSIS® 3.1 is easily represented on a Capacity Management Timeline, and its impact on lifespan assessed. Along with other forward-looking possibilities, an example of new thresholds for DOCSIS® 3.1 is shown in Figure 3 [11].

Figure 3 plays out both the "Forever CAGR" scenario as well as the "Tapered CAGR" scenario. With tapered CAGR, an asymptote of subscriber consumption occurs over time. In this example, the asymptote is associated with the limits of simultaneous media consumption and associated practical limits on Quality of Experience (QoE). This particular comparison and the logic for (and against) a CAGR asymptote are specifically described in [11]. The figure also includes a move to N+0, or the equivalent service group size by other means, in 2024 with the third breakpoint in the trajectories.

Key conclusions from Figure 3 are three-fold. First, "forever CAGR" shows a lifespan carrying into the late 2020's. Second, with an assumption of two additional node segmentations, taking the households passed per node into the 100-125 range, the

DOCSIS® 3.1 10 Gbps objective can be seen to be sufficient bandwidth under the assumptions of tapering used [11]. Third, with N+0, we can achieve the tapered threshold with about 3 dB (2x) of error margin for systems limited to 750 MHz. This is important since such systems may not just be bandwidth limited (obviously), but also limited in their cost effective upgradeability due to the implication of age where 750 MHz actives are still in place.

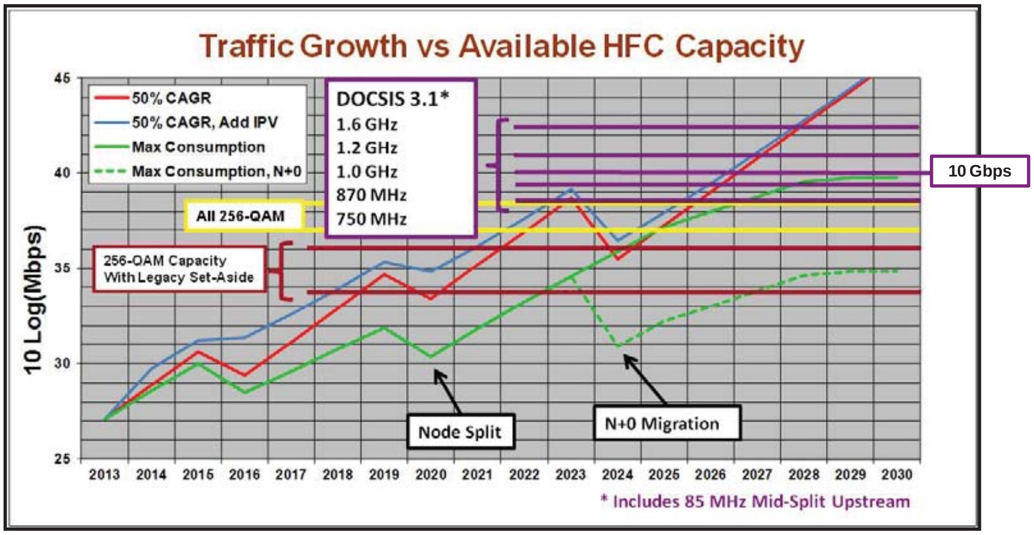

Figure 3 – Forever and Tapered CAGR vs Legacy and D3.1 Thresholds

Figure 3 represents a lot of what MSOs lose sleep over when pondering their migration alternatives in the face of persistent CAGR. While tapering shows tremendous potential benefit, *counting* on it is obviously a risky bet. There will be quite a bit of trend observation of CAGR in the coming years to gauge where HSD bandwidth growth might be headed. And, the other nice conclusion that can be drawn from analysis like Figure 3 is that a runway of years are possible to do exactly that in order to guide reasoned decision-making.

The tapering logic and model used in Figure 3 is described in detail in [11]. Summarizing it briefly, it basically involves assumptions about the next generation of High Definition taking over as the dominant primary screen service over time, all-IP delivery and its inherent traffic engineering variables becoming the dominant delivery format, service group size shrinking manifesting as increasing unicast, and other practical demographic and behavioral assumptions about human media consumption.

Lastly, we have seen how Figure 3 puts into perspective what DOCSIS® 3.1 does for the network from a capacity management and lifespan extension standpoint when combined with other architectural assumptions. Not obvious from Figure 3 but very important to point out is the DOCSIS® 3.1's 10 Gbps objective, which puts cable on a level playing field with 10 Gbps EPON. DOCSIS 3.1 allows this threshold with just a modest extension of bandwidth over 1 GHz, which is anticipated to be achievable without major implications to architecture.

Upstream with an Extra Paddle

As mentioned, it is persistently aggressive downstream growth that keeps cable operators up at night. And, it is well-understood in operations that downstream segmentation to manage new growth is often wisely accompanied by simultaneously addressing the upstream while visiting the plant, if not necessarily at the cable modem termination system (CMTS) receive port. As downstream data growth goes, so also goes the upstream – but generally not at the same pace due to the nature of the difference in the dominant traffic types.

In fact, CAGR in the upstream has been uite slow compared to the downstreamover the past several years – a far cry from the chaos wrought by “Napster” that dealt a severe blow to “phone line” Internet access over a decade ago. However, operators are wary about turning their attention too far away from upstream, mainly because the possible “fixes” are few and far between. While the downstream problem is complex, there are at least many knobs and levers under the operator’s control, including the service mix itself, to help manage the problem.

In the upstream, there is not much under the operator’s control from a service standpoint, the spectrum constraints present a hard stop on new channels to add, and the spectrum itself is nearly full – at least where operators have drawn the line today with advanced time division multiple access (A-TDMA) only. Nonetheless, there is not much panic about the upstream, because it is relatively straightforward to paint a strong capacity and lifespan management picture for the upstream, as shown in Figure 4, especially when including DOCSIS® 3.1.

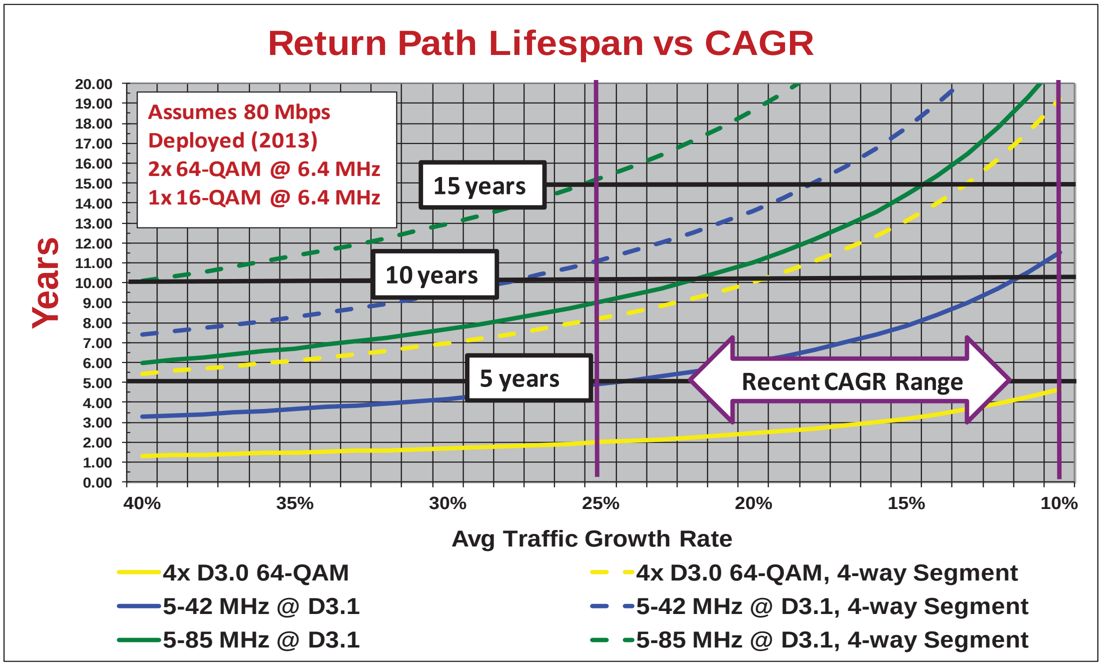

Figure 4 – Upstream CAGR has Lagged Downstream in Recent Years, Leading to a Healthy View of Lifespan Management

Some key take-aways from Figure 4 are the extended lifespan DOCSIS® 3.1 provides, in part due to the larger percentage gains of new modulation profiles compared to the downstream (since the upstream is starting lower), the range of ways to achieve > 10 years of lifespan with modest assumptions of CAGR similar to those observed today, and what segmentation combined with DOCSIS® 3.1 can mean to the congested 42 MHz return band.

New Tools and Techniques

Capacity Optimization Simplified

Theoretical capacity is a based on two simple variables – Bandwidth (BW) and Signal-to-Noise Ratio (SNR). Shannon Capacity, the well-known limit, represents the maximum error-free rate in additive white Gaussian noise (AWGN), and is expressed as

$$C = [BW] \, Log_2 \, [1+SNR \text{ (linear)}] \quad (1)$$

This can be yet further simplified for cable networks, especially downstream. If the SNR is high, we can instead write, with SNR *expressed in decibels* (dB):

$$C \approx [BW] \, [SNR \text{ (dB)}] / 3 \quad (2)$$

This simplification of Shannon Capacity is accurate to within less than 0.5% above about 15 dB SNR.

According to (2), more capacity is available with higher SNR, but with logarithmic proportionality. For example, 50% more spectrum yields 50% more capacity, but so does 50% more SNR – in . However, turning a 30 dB SNR into a 45 dB SNR is a significant network performance leap. Nonetheless the SNR proportionality is the premise around migration to architectures that create higher SNR, such as deeper fiber, digital optics, and home gateways.

In practice, SNR has two key practical components:

1) Improving the link end-to-end SNR itself, which translates to higher order modulation formats, and higher spectral efficiency in bps/Hz.

2) Use the best Forward Error Correction's (FEC) available to enable a given SNR to deliver the most bandwidth efficient QAM formats.

Efficient Use of Spectrum: Advanced M-QAM Formats

Today's cable systems implement a maximum M-QAM format of 256-QAM (8 bits/symbol) downstream and 64-QAM (6 bits/symbol) upstream. Digital communications technology and information theory have progressed significantly in the years that these have been available. New tools and formats are now possible.

Figure 5 shows current modulation profiles (upstream and downstream) and two of the higher order constellations anticipated in DOCSIS® 3.1– 1024-QAM and 4096-QAM (512-QAM and 2048-QAM are not shown). All of the M-QAM formats shown in Figure 5 are at an *equivalent uncoded bit error rate (BER)* of 1e-8. Since they are 6 dB apart for each 2 bits/Hz step upward in density, the SNRs are therefore 28 dB, 34 dB, 40 dB, and 46 dB, respectively.

A common end-of-line HFC cascade performance requirement for digital channels is a composite carrier-to-noise (CCN) of 42 dB. Given that 256-QAM requires 34 dB (1e-8) without coding, and J.83 Annex B error mitigation is provided, it is apparent why today's networks successfully implement 256-QAM. In fact, some are likely able to support 1024-QAM robustly already using similar "J.83B" tools [13]. Previous challenges with 1024-QAM were primarily due to excessive phase noise in customer premises equipment (CPE) and that is no longer the case with modern CPE

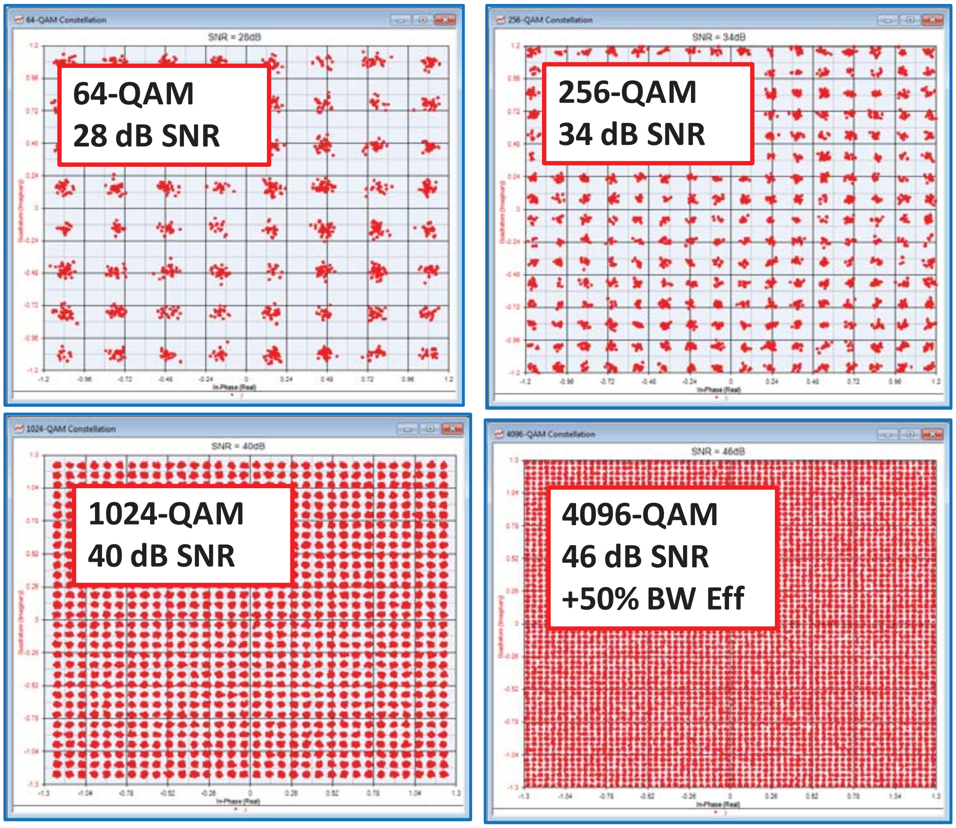

Figure 5 – Increasingly Spectrally Efficient M-QAM Formats

However, it is also apparent given the relationship between HFC performance and 4096-QAM uncoded requirements that there is necessary work to be done to hope to achieve this with robustness. It can come in the form of FEC (a must), architecture modifications, technology improvements, or all of the above.

<u>A View from the Field</u>

Figure 6 shows pioneering field analysis performed by a major North American MSO that statistically quantifies the reported SNR for a large sample of cable modems (CM) [20].

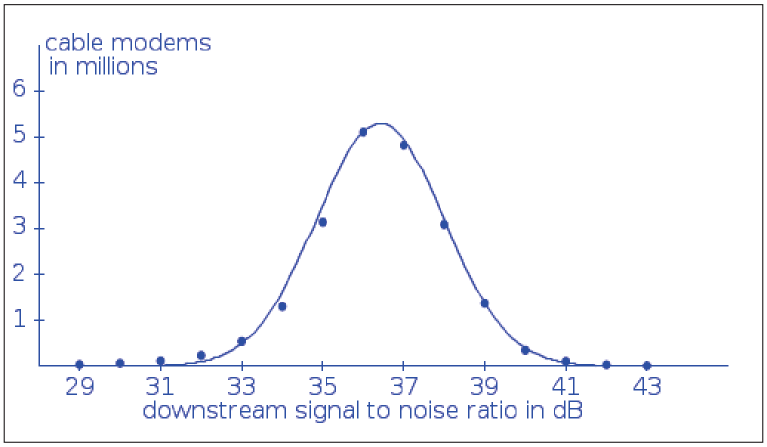

Figure 6 – Major MSO Cable Modem SNR Distribution [20]

There are important differences between a CM's reported SNR and HFC delivered SNR. The most important ones are:

1) The CM actually is reporting modulation error ratio (MER), which includes *all* impairments on the channel, including the CM contribution itself.
2) The CM's contribution to SNR is strongly dependent on the location of the CPE in the home.
3) The CM was implemented with 256-QAM in mind, so the MER measurement limitation is perfectly sufficient if it is in the low-to-mid-40's.

Despite these limitations, Figure 6 serves as an excellent baseline for the distribution of the channel quality, and therefore is an excellent, conservative source for assessing the use of new QAM profiles, FEC, and architecture to maximize bandwidth efficiency.

FEC, 15 Years Later

As powerful as J.83B's Reed-Solomon-based FEC is, 15 years of technology evolution means more sophisticated FEC can now be processed in real-time. Better FEC reduces the SNR required for a particular QAM format, increasing bandwidth efficiency. The family of codes offering cost-effective performance that is closest to Shannon capacity today is Low Density Parity Check Codes (LDPC). LDPC codes have been around since the 60's, but the computing requirements were beyond that possible in consumer electronics or CPE until recent years.

In Figure 7, we show the BER vs. SNR performance for the DVB-C2 family of LDPC codes [19] for all square constellations from 16-QAM to 4096-QAM. Observe the encircled SNR requirements enabled by LDPC under the "Highest Code Rates" label in Figure 7 (90%). These are the nearest comparisons to the code rate used by J.83B today.

The true power of LDPC can be seen in the SNRs required to deliver vanishingly low error rates. For a 90% code rate, the FEC comes with a 10% efficiency penalty. However, it delivers *9-11 dB* of SNR gain, which means that *one-tenth* of the SNR is tolerated for just 10% lost efficiency. The 4096-QAM case, for example, reduces to 35 dB, compared to the 46 dB number without FEC. The 9-11 dB range of SNR advantage is a testament to the power of LDPC codes.

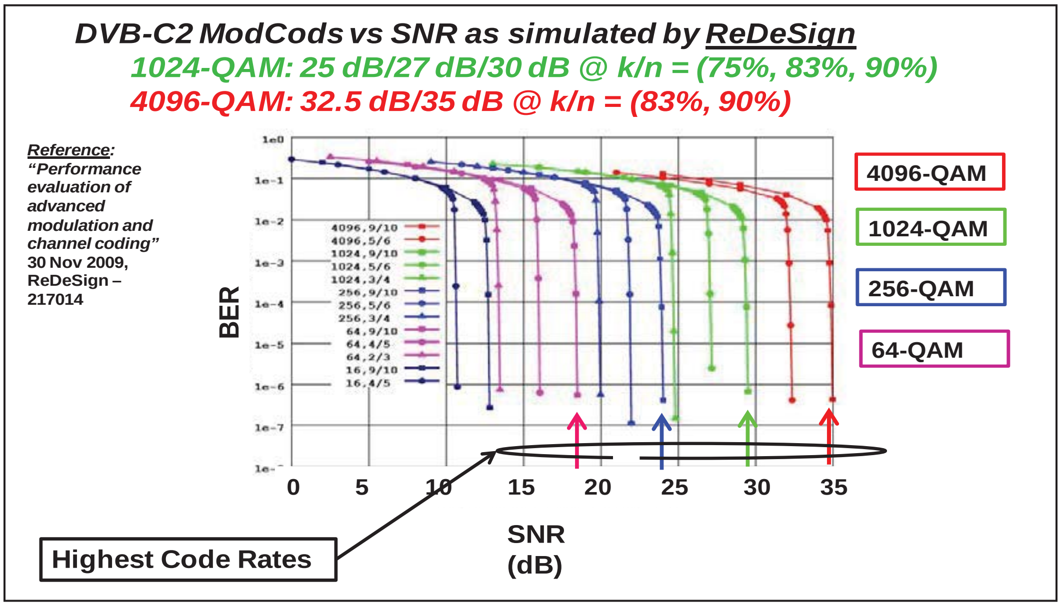

Figure 7 – Bandwidth Efficient M-QAM Enabled by LDPC FEC [19]

Figure 8 illustrates how closely the DVB-C2 based LDPC code is to Shannon theory.

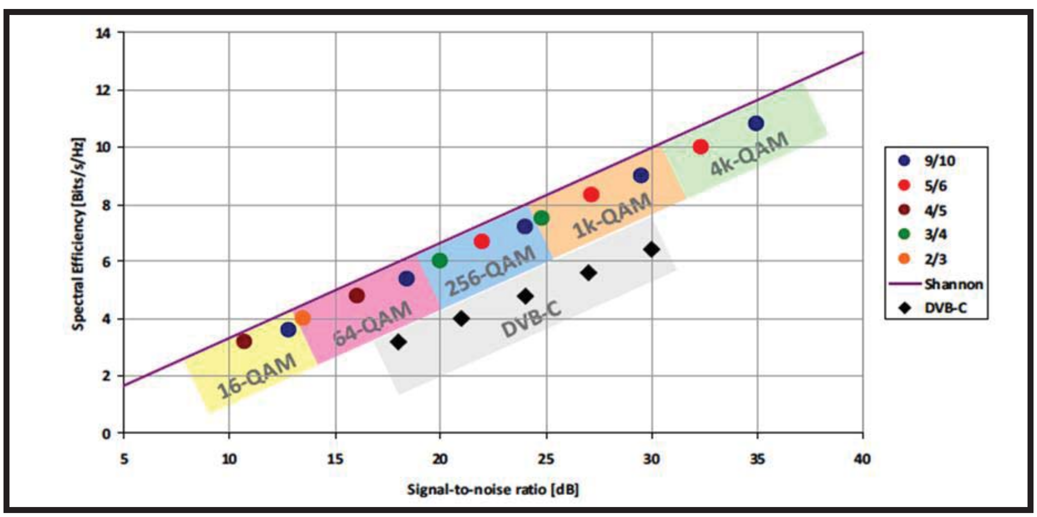

Figure 8 – Advanced FEC Enables More Bandwidth-Efficient QAM and Performance Closer to the Shannon Capacity Limit [19]

One final impression of the power of modern FEC can be seen by looking at the poor quality of the received constellations that are still able to be repaired by the LDPC FEC in the AWGN noise case. This is shown in Figure 9.

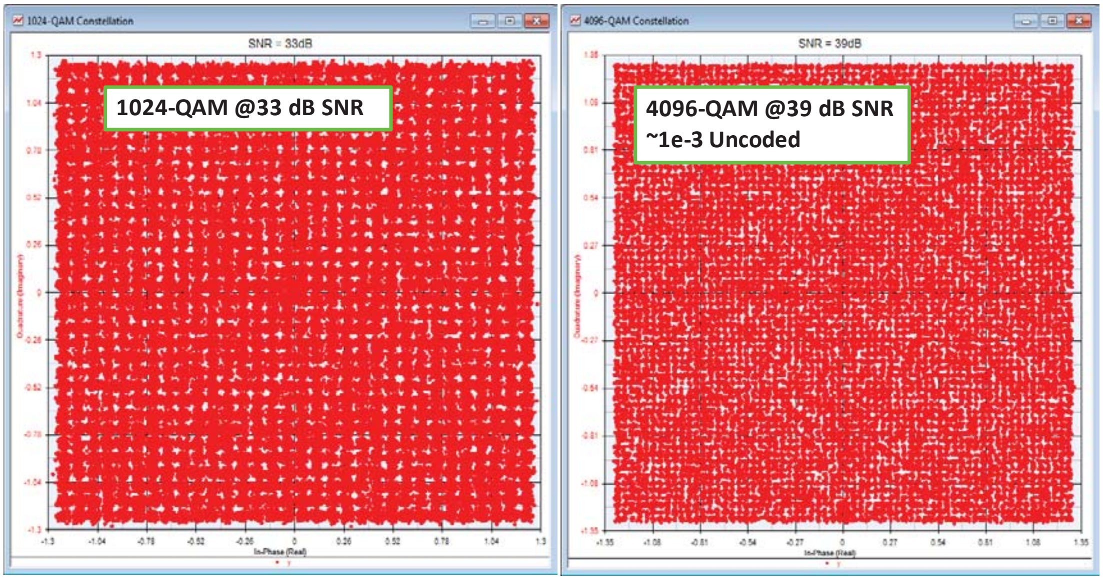

Figure 9 – Can't Tell Which Symbol Was Sent? An LDPC-based Decoder Can!

<u>QAM-FEC Meets HFC</u>

We can put the QAM/FEC relationship shown in Figure 7 with the field reported SNR distribution in Figure 6 to understand what we might hope to achieve with DOCSIS® 3.1 given today's measured-in-the-field end-to-end performance. Again, we keep in mind the caveat that this performance is guided by what was necessary to achieve 256-QAM. This is shown in Figure 10 [21].

From the distribution itself, which closely approximates Normal statistics, we can determine that over 98% of the CMs are reporting an SNR higher than 32 dB. This 32 dB SNR is approximately the slicer threshold for 2048-QAM, indicating that the majority of CMs may be capable of increasing their modulation order from 256-QAM to 2048-QAM. However, accounting for practical margin, it might instead be prudent to suggest that a 32 dB SNR instead assures robust 1024-QAM for over 98% of the CMs. In this case, a 25% bandwidth efficiency improvement appears to be a lock for the majority of users, and thus also as a minimum increase of aggregate capacity.

Even applying this practical margin, however, we can state that well over 50% of the CMs can support 2048-QAM with 3 dB of margin (35 dB SNR) and indeed a sizeable population should comfortably be able to support 4096-QAM (38 dB SNR). Based on the Normal distribution fit statistics applied [20], the net percentage of capacity increase is about 37.5%, or the equivalent of 2048-QAM on average, as can be eyeballed from Figure 6 and a Gaussian assumption.

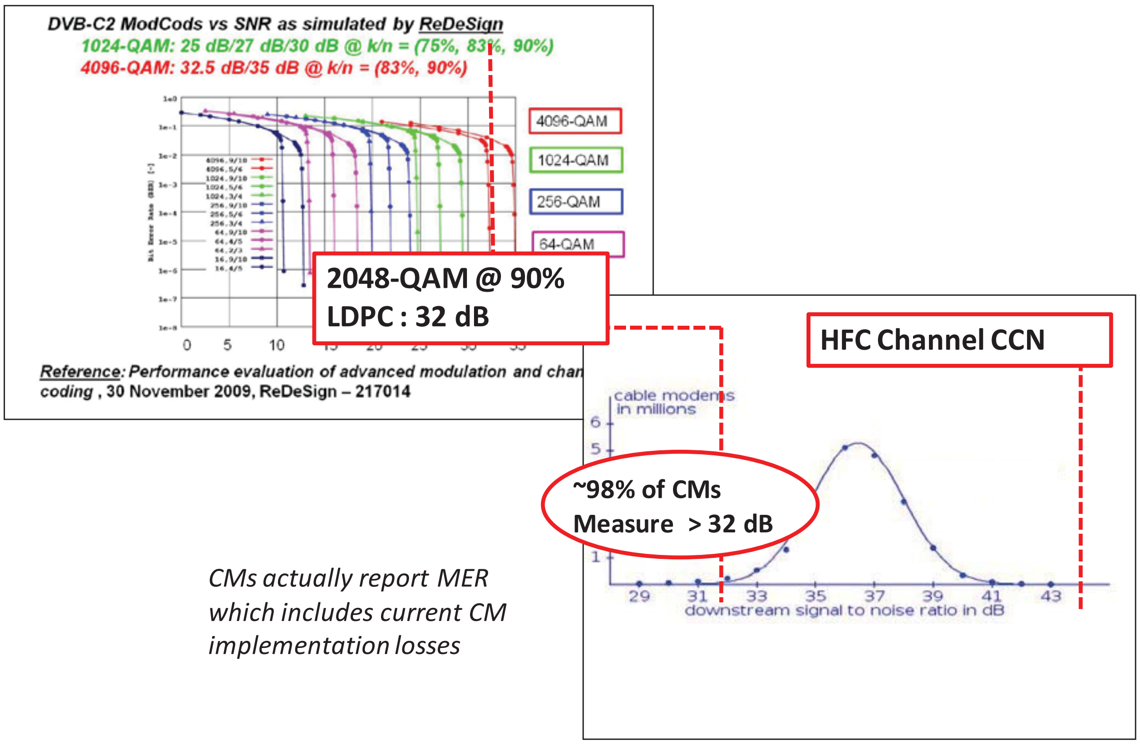

Reference: http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf

Figure 10 – Bandwidth Efficient M-QAM Enabled by LDPC FEC

As encouraging as Figure 10 appears, the distribution of CM SNR, as we have pointed out, is conservative. Both the transmitter and receiver ends of the link (the CMTS or edge-QAM (EQAM) and the CM or the multimedia terminal adapter or MTA) were designed with a requirement only to robustly support 256-QAM. Transmit fidelity and receive sensitivity are implemented to achieve this level of performance cost-effectively over an HFC plant – a plant with inherent noise and linearity capability to support analog video. How much better might we do without the constraint imposed by equipment targeting 256-QAM?

We can get a sense of “can we do better” by extracting from Figure 6 the endpoints that are likely to get in the way based on today’s design criteria. Instead, we can look specifically at the capability of the outdoor plant equipment that creates the RF channel conditions (which, of course, includes the effects of the HFC optics). We can also recognize that this channel performance is not completely fixed. As operators further segment the plant with fiber deep, as shown in Figure 11, channel conditions can also improve. To maximize what DOCSIS® 3.1 can achieve, new targets for plant performance might make sense for new investment.

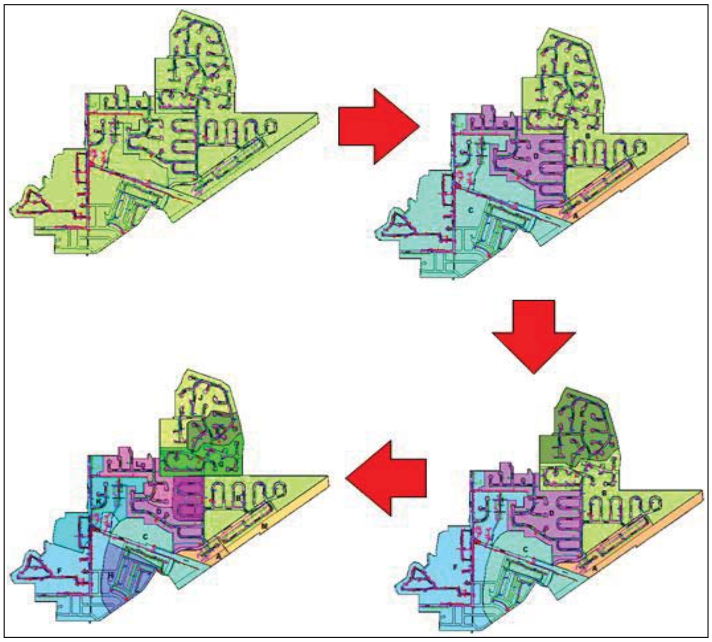

Figure 11 – The Fiber Deeper Double Bonus: More Average Bandwidth per Subscriber and an Improved HFC Channel

An example of this improvement for a typical 1310 nm point-to-point HFC link with fiber pulled deeper and shortening of the RF cascade is shown in Figure 12. In Figure 12, the QAM format possible as a function of the input level at the CPE (y-axis) for various drop-home architectures for a given tap port design level (15 dBmV) and HFC performance is shown. When the colored trajectory drops below an input level threshold in black, that format is enabled by the combination of variables on the x and y-axis.

Note also in Figure 12 that 8192-QAM begins to appear as a possibility from an SNR standpoint for N+0 with a friendly drop-home architecture. Indeed, 8k-QAM and 16k-QAM are part of long-term DOCSIS® 3.1 M-QAM considerations. The increasing architectural constraints in Figure 12 for the highest order QAM formats, including 4096-QAM, is one of the primary reasons that digital fiber architectures with RF only in the last coaxial mile are drawing attention. These would contribute robustness and perhaps higher M-QAM capability downstream.

Note that the x-axis is labeled “HFC CCN” where CCN stands for Composite Carrier-to-Noise ratio. This term recognizes that the noise floor we observe, for example, in a spectrum analyzer measurement, is composed of both thermal noise components such as laser and amplifier noise, but also of digital distortions which inevitably splatter across the band as intermodulation (IMD) products are formed. HFC equipment design and level setting recommendations are about achieving the highest, robust, CCN (with some margin of headroom), and is a careful balance of increasing levels to rise above the noise but not so much as to move too close to regions of nonlinear distortion. This careful balance will be at a premium for increasingly high QAM formats.

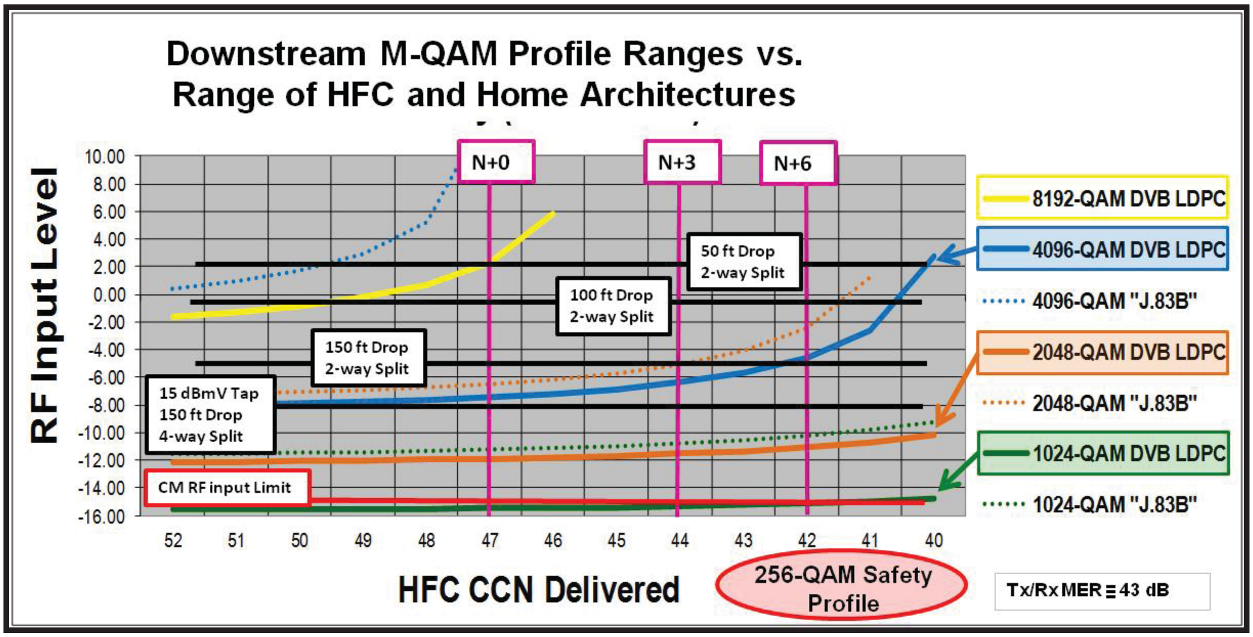

Figure 12 – The HFC Network's Performance Range Can be Used to Identify the Range of QAM Profiles the Downstream Should Support

When the HFC network is properly level aligned, the noise floor "looks" like a typical noise impairment, but in fact has digital intermodulation components among its contributors. From the standpoint of predicting error rate performance, central limit arguments allow us to still consider the noise floor contribution to digital detection as Gaussian additive noise. In fact, however, in measuring and predicting CCN we are capturing both the noise and nonlinear distortion characteristics of the downstream channel. When no analog carriage exists, CCN includes *all* of the intermodulation distortion energy.

Upstream: A Similar Story

The above discussion focused entirely on detailing the use of advanced M-QAM and new LDPC-based FEC on the downstream. However, the same toolkit of new QAM formats and FEC applies to the upstream. There are choices of FEC for the upstream, (as there are in today) because of the range of possible channel conditions and packet sizes. Also, the QAM formats are less aggressive in general because of the nature of the channel performance by comparison to the downstream. These are detailed in [6][10]. The end result, however, is still that more bps/Hz will be available.

The combination of QAM, FEC, modern distributed feedback (DFB) optics (or 10-bit digital return), and expanding the split to 85 MHz are summarized by Noise Power Ratio

(NPR) performance shown in Figure 13. Note, as with downstream, composite upstream "SNR" is also a combination of noise and distortion. Proper set up of the return path levels ensures operating in and around the sweet spot of NPR performance that delivers sufficiently high signal-to-composite noise and distortion, yet has enough headroom from compression and clipping that these effects are minimized. Level alignment has been an upstream issue since the first days of DOCSIS®, and has resurfaced with the growth in number of upstream channels. Getting this right will be critical to maximizing DOCSIS® 3.1's capacity potential.

An important part of the upstream is that the channel performance described by the Noise Power Ratio (NPR) curve shown in Figure 13 is much more likely to be compromised by external sources of interference, in particular at the lower end of the band. As the upstream split expands, the spectrum is expected to become more efficiently used as these types of interference and impulse noise tend to be minimally disruptive above about 20 MHz, at least in North America. We will discuss interference and impulse further in a subsequent section.

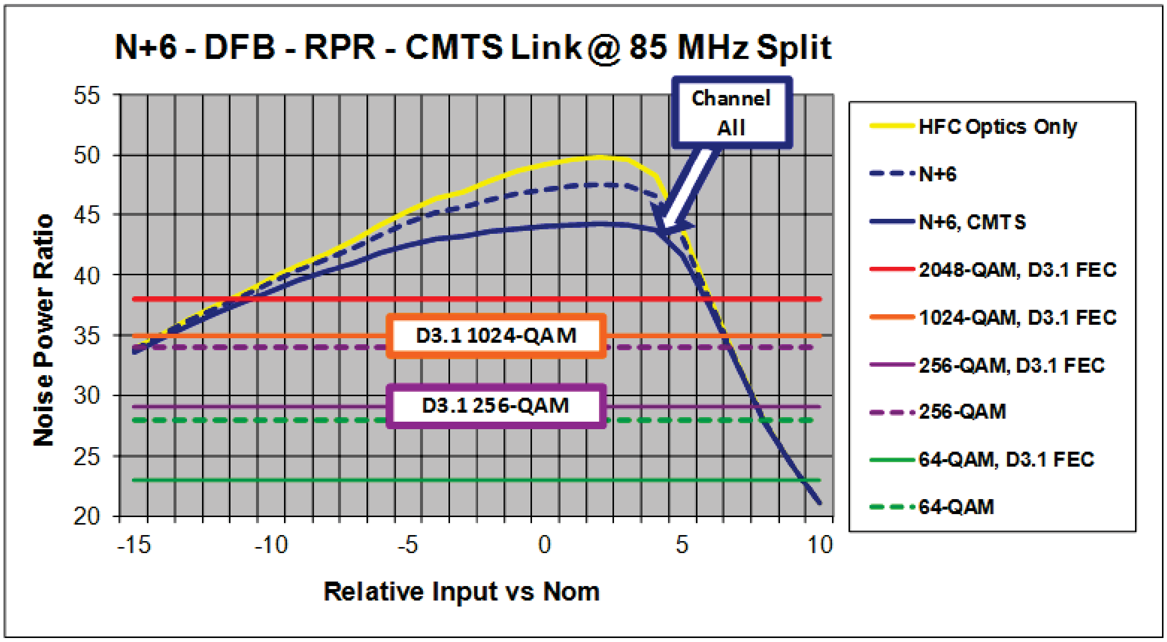

Figure 13 – The Upstream is also Enabled with New Capacity Potential if DFB or Digitized Returns are in Place to Exploit it

Lastly, Figure 13 assumes a DFB laser (although baseband digital solutions are sufficient as well) and assumes no combining of returns paths post-optical receiver in the Headend. A remote demodulation architecture would act from a performance standpoint like baseband digital return. However, the "NPR" performance in this case is hidden to the operators since the A/D converter directly feeds the receiver, and it would have the advantage of front end AGC to drive the A/D at the proper input level.

While late model Isolated FP lasers could support multiple channel DOCSIS® 3.0 to some degree, they did so with little margin. These must be replaced by DFB or digital return solutions to squeeze the most out of the return using DOCSIS® 3.1.

Headend return path receiver combining, of course, incurs an SNR penalty – 3 dB on equivalent links. This 3 dB is a half-modulation order (1024-QAM down to 512-QAM, for example). So, while it is understandable that plant segmentation in the upstream may occur before the ports are needed in the CMTS because of synergistic segmentation with the downstream, there is a price to be paid in terms of bandwidth efficiency. However, if the upstream is uncongested to begin with (the segmentation is being done for the D/S), then this 3 dB loss and accompanying efficiency loss may not be significant at the time of the split.

CCN, NPR, MER – What Else?

While CCN and NPR represent commonly used access network terminology, DOCSIS® 3.1 CMs will also be reporting MER as they do today. And, this is important because this singular value does capture all of the impairments, not those simply contributed by the access network. However, as modulation orders get higher, a single parameter as a reference for signal quality becomes less sufficiently descriptive on its own. The reason for this can be correlated to the SNRs that we saw rise with the “M” in M-QAM in Figure 5 to achieve the same uncorrected BER. Not only are these signals more sensitive to AWGN, they are sensitive to all other non-AWGN impairments as well, some of which perhaps would not impose any degradation to 256-QAM.

For example, some forms of optical non-linearity in wavelength division multiplexing (WDM) systems may cause a crosstalk among optical wavelengths that measures 45 dB below the desired RF spectrum in the coaxial domain. This may be insignificant to 256-QAM, but for a profile such as 4096-QAM that might be relying on 40 dB of SNR, a degradation of over 1 dB is incurred that could be impactful to robustness.

In addition, many forms of non-AWGN impairments are frequency dependent (2^{nd} and 3^{rd} order distortions), so how these are expressed in a wideband system – as opposed to 6 MHz channel slots – to derive QAM profiles will be meaningful. Figure 14 shows how third 2^{nd} and 3^{rd} order discrete (analog) and digital distortion products can be distributed across a forward path bandwidth.

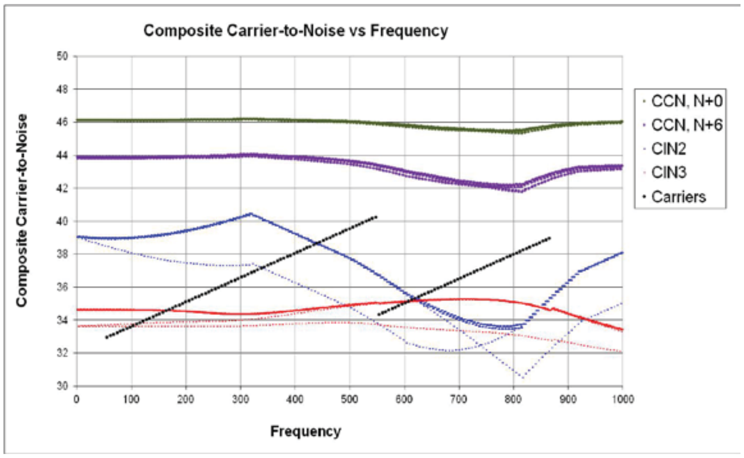

Figure 14 – Frequency-Dependent Impairments Have Different Connotations for Wideband Signals Compared to 6/8 MHz Channelized Systems

Finally, some forms of non-AWGN impairments create non-uniform effects across the symbol points. The example of phase noise is shown in Figure 15.

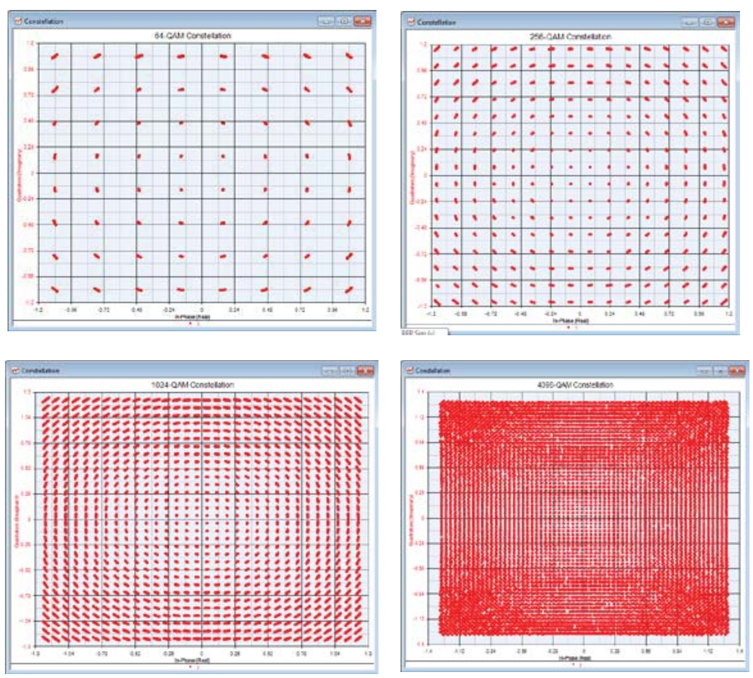

Figure 15 – Phase Noise is an Example Distortion Mechanism that Can Have a Non-Uniform Impact Across Symbols

For DOCSIS® 3.1, then, to ensure added robustness given the increased sensitivities of higher order modulation, we will need to "see" things we perhaps did not see or even try to see before. This means a deeper set of diagnostic metrics and proactive maintenance tools for gathering data on operating systems. The implementation of full-band capture A/D front end receivers and wide band signals aligns well with the goal of additional waveform information, such as the capture of burst noise and waveform histograms to complement FEC data and reveal information about network alignment (is the laser clipping, is there a transient noise issue). Within the receiver, MER will be augmented by deeper information about the constellation itself, as opposed to just the average noise cloud across all symbol point that MER describes.

Examples of potential metrics that may be made available, and which operators will need to invest in from a system assessment, alarm threshold, and personnel training standpoint , include:

- Full band spectrum capture
- Per-subcarrier MER
- Per-symbol MER
- Constellation I,Q values
- Equalizer coefficients
- Waveform histogram
- Transient capture
- NPR notch depth
- Spectral efficiency in bps/Hz
- FEC statistics in particular the corrected and uncorrected codeword errors:
    - All three metrics for the concatenated Bose, Ray-Chaudhuri, Hocquenghem (BCH)-LDPC approach in DOCSIS® 3.1: the error rate before any correction ("Pre-FEC") the error rate after LDPC ("Pre-BCH") and the final error rate after all correction ("Post-FEC").

In some cases, the metrics will only be available through the use of a pre-defined "off" or "quiet" transmit mode built into the specification to allow an observation window (time and/or frequency) that might otherwise be unavailable in normal operations.

The DOCSIS® 3.0 Proactive Maintenance team made excellent strides through the use of the Pre-Equalization tool in helping to better diagnose DOCSIS® networks. DOCSIS® 3.1 will build on this with a new set of metrics consistent with the need to support advanced QAM formats more likely to suffer from more subtle non-AWGN effects.

Spectrum Evolution

Figures 3 and 4 each imply an evolution of HFC spectrum to enable anticipated growth and achieve the DOCSIS® 3.1 objectives. In principle, the 10 Gbps downstream objective can be met by modestly increasing the downstream beyond 1 GHz – modest enough that many taps and 1 GHz actives singularly may have excess bandwidth above 1 GHz that can be exploited.

Simple algebra can lead to the 10 Gbps conclusion: 1 GHz of spectrum used at an efficiency of 10 bps/Hz is 10 Gbps. 1 GHz of downstream spectrum would push today's 54 MHz-1 GHz forward bands just over the 1 GHz point. Furthermore, if we anticipate an upstream expansion to 85 MHz, then the downstream will begin at about 108 MHz. 1 GHz of spectrum will then push the downstream just beyond 1100 MHz.

Now, 1024-QAM modulation is 10 bits/symbol. Ideally then, 1 GHz of 1024-QAM is 10 Gbps. However, factoring in spectrum shaping and OFDM and FEC overhead, 1024-QAM alone does not quite get there. But, recognizing that we have 11 bits/symbol (2048-QAM) and 12 bits/symbol (4096-QAM) at our disposal provides the extra oomph to overcome the efficiency losses and achieve 10 Gbps when channel conditions permit.

Of course, the quality of bandwidth above "1 GHz" components it is not guaranteed to be flat. It will roll-off with cascading, and perhaps quite severely, making 10 Gbps accessible only if architecture evolution and cascade shortening enables it. And, where there is roll-off there is also group delay dispersion that creates intersymbol interference (ISI) in single carrier systems, and potentially effects efficiency in OFDM systems as well. Nonetheless, 1100 MHz is meaningful because what has been found in 1 GHz components is excess spectrum ranging up to 20% in some cases (i.e. 1200 MHz). Figure 16 shows the excess bandwidth concept on a single 1 GHz tap, and the range of variation such devices may have in practice.

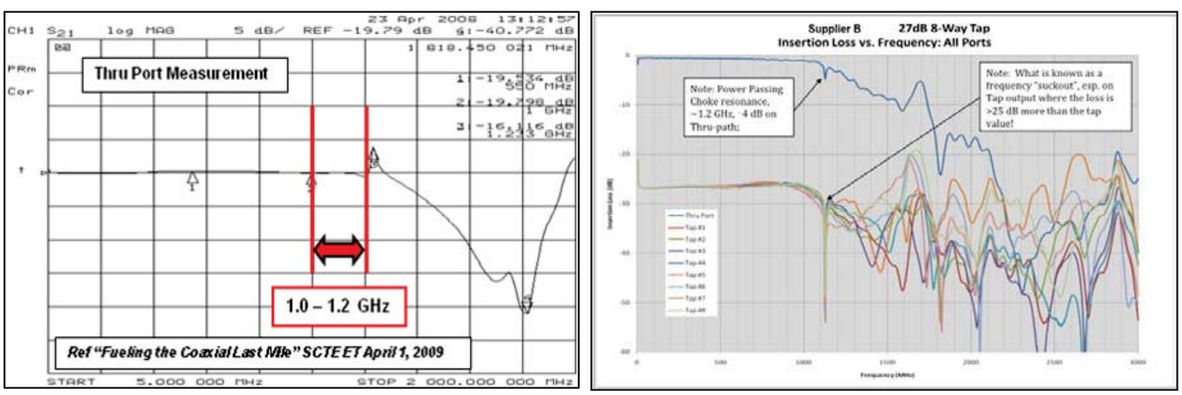

Figure 16 – "Excess" Bandwidth Potential of 1 GHz Tap Products [1]

While 10 Gbps is potentially achievable with a little extra downstream spectrum, should we be able to access it, the other DOCSIS® 3.1 objective is for 1 Gbps in the upstream. Again, some simple math from several angles can show why we will *not* achieve this using DOCSIS 3.1 in 85 MHz.

First, using (1), if we can assume a robust 40 dB of SNR over 80 MHz of bandwidth, we will find 1066 Mbps of *theoretical* capacity. However, we are not at 94% efficiency relative to Shannon limits. Another way to view this is that the spectral efficiency needed is (C/B) = 40/3 or 13.3 bps/Hz. This is better than ideal 8192-QAM in the upstream, which is not even an optional upstream profile in DOCSIS® 3.1. Lastly, consider that the highest defined upstream profile (which is optional) is 4096-QAM, which is 12 bits/symbol. After all overheads, if we could manage to achieve an aggressive 11 bps/Hz upstream using 4096-QAM, then we'd only reach 880 Mbps. Note that a spectral efficiency of about 11 bps/Hz was recently demonstrated at the 2013 Cable Show by the NCTA.

Because of these limitations, DOCSIS® 3.1 has defined both the 85 MHz Mid-Split and the 200 MHz High Split, with the idea that the 200 MHz upstream case will indeed enable 1 Gbps of upstream capacity, and even a 1 Gbps upstream speed if necessary at some future point. However, if the upstream is going to extend to 200 MHz, then the downstream is going to have to be shifted accordingly to maintain its 10 Gbps of capacity. This then requires plant components that go beyond 1 GHz, by design. This is what is behind the 1.6 GHz threshold in Figure 3, and also is behind DOCSIS® 3.1 specifications identifying an extended spectrum option.

Tap products already exist that achieve an extended bandwidth, such as is shown in Figure 17.

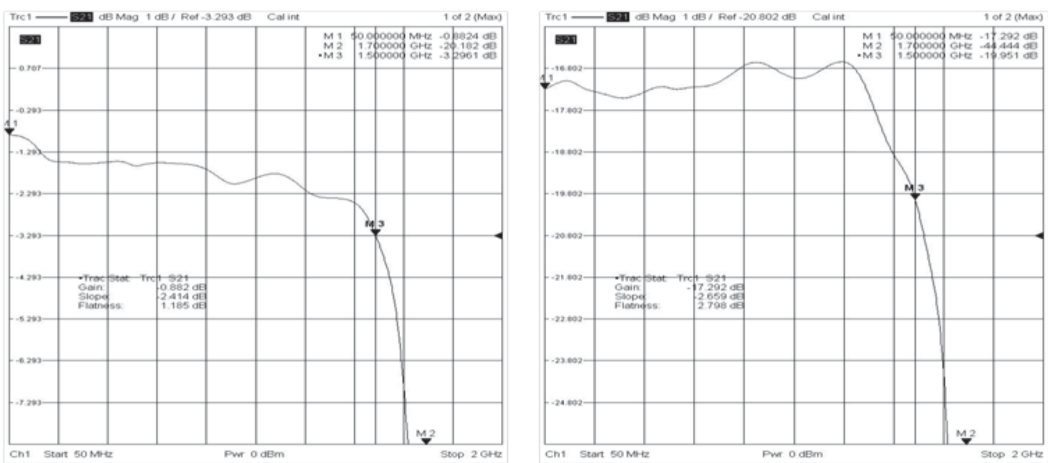

Figure 17 – A 1.5 GHz, 17 dB Tap Response (Through Port, Coupled Port)

It is not clear that amplifier products will follow tap solutions with extended band operation, as important implications to housings due to microwave effects with high gain become significant. However, at a minimum, fiber optic nodes that can drive this bandwidth will be required if it is to be exploited, and this would entail many of the same issues RF amplifiers have.

Other potential issues arise as well with such an extended bandwidth. In particular power loading of the RF spectrum becomes more challenging. Figure 18 shows how an extended band might impact plant actives from the power loading perspective. As shown in Figure 18, the 1.2 GHz case (same level, extended tilt) may be manageable with technology evolution, incurring a 2.7 dB of increased (tilted) load. However, a fully loaded spectrum to 1.7 GHz over linear optics looks very stressful to the design of plant actives that require high linearity.

It is this reason and the housing limitations of RF amplifiers themselves, among others, that drive architectural solutions for the 1.7 GHz band towards digital optics and "narrowband" overlay, modular, solutions.

Lastly, though we display the classic "uptilt" use of RF loading in Figure 18, the logic to this is very clear when using equal capacity digital channels. When some channels (or spectrum) can be used more efficiently than other channel (or spectrum), it makes sense to consider using a disproportionate share of the RF load on the more capable channels, as opposed to burning it away on excess margin for 256-QAM channels, for example. This is something likely to be explored in DOCSIS® 3.1 optimization of shared HFC spectrum.

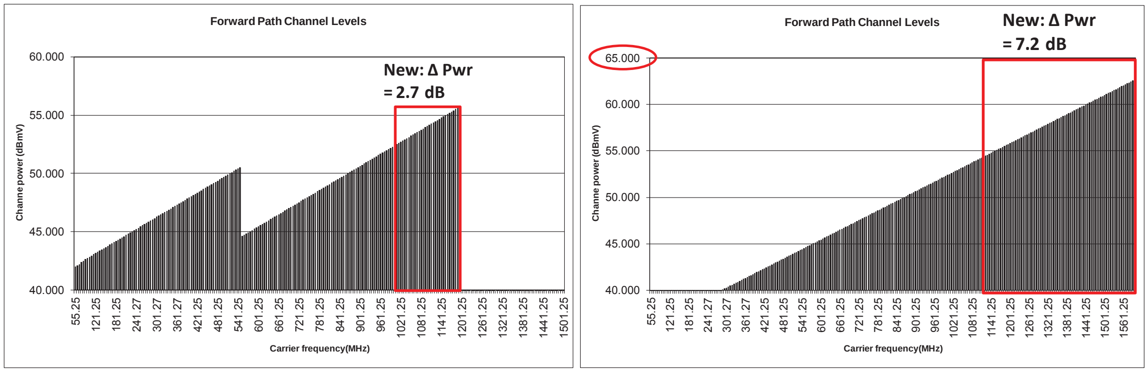

Figure 18 – Adding Spectrum Increases the RF Power Load Disproportionately on the Tilted Coaxial Output

Putting all of the above pieces together in a long term spectrum evolution plan – with an assumption that 10 Gbps/1 Gbps is an operator's end state – leads to Figure 19.

It is quite possible that 1.2 GHz/85 MHz becomes a reasonable "end state" from a business planning perspective given the capacity management analysis and lifespan previously discussed, coupled with the likely speed tiers MSOs will consider important to ensure by allocating sufficient spectrum.

For any extension of forward spectrum, and acute awareness of MoCA™ and satellite TV downconverted bands will be necessary, and the possibility that in-home filters become a regular part of the home solution will be likely unless a true HFC-terminating gateway architecture is deployed.

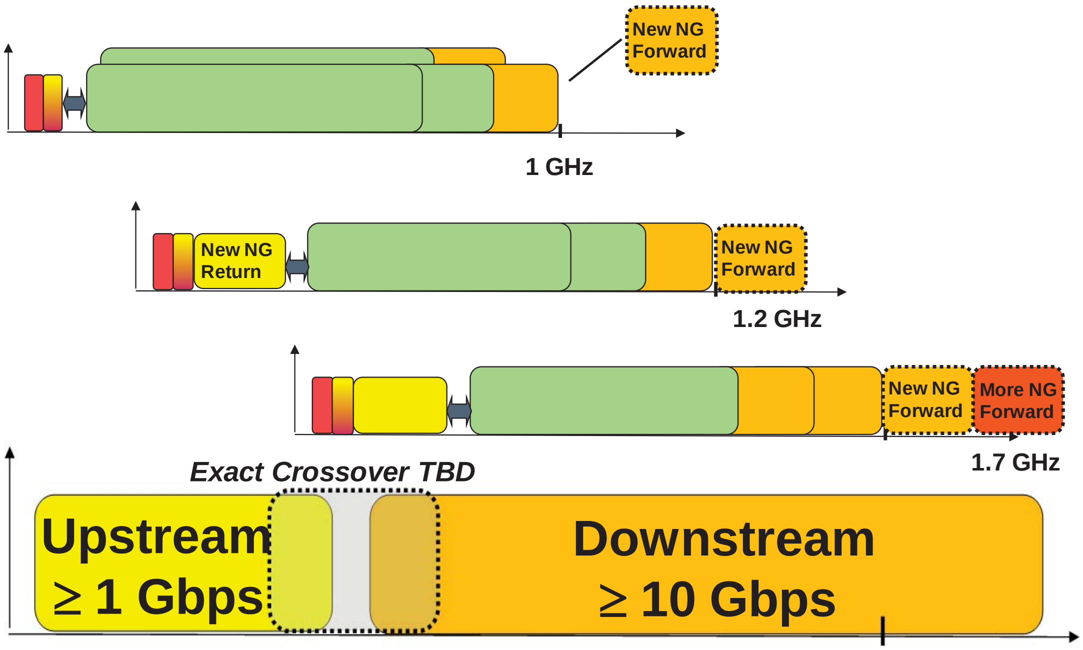

Figure 19 – Possible Long Term Evolution of the Coaxial Spectrum

<u>What About this OFDM Thing?</u>

Capacity Implications

At the beginning of this section in equation (1), we wrote this:

$$C = [B] \ Log_2 \ [1+SNR \ (linear)]$$

Then, we simplified this to this (2):

$$C \approx [B] \ [SNR \ (dB)] / 3$$

Well, we are going back to un-simplification the capacity equation. This is because a good way to interpret the OFDM approach in terms of its capacity-maximizing effect is to write the simplified expression for capacity in "long" form:

$$C \approx (1/3)\sum_{\Delta f} [\Delta f] \ [P(\Delta f) \ H(\Delta f) / N(\Delta f)]_{dB}$$

Here, instead of bandwidth, we have used a summation of spectrum chunks using a set of small frequency increments, Δf. The sum of all Δf increments is the bandwidth available, B. Instead of SNR, we have broken it down into its components: signal power (P), noise power (N), and channel response (H) – over small Δf increments. In practice, Δf is a parameter that represents how OFDM operates, which is shown in Figure 20.

We can consider Δf the width of one OFDM subcarrier. And, while Figure 20 shows these subcarriers overlapping in frequency, which they do, the fundamental premise of OFDM is that they do not interfere with one another because the specific uniform spacing chosen between them ensures orthogonality when they are generated by one transmitter or synchronized among several transmitters. The principle of orthogonality is not unique to OFDM – only the dimension to which it applies. Orthogonal *codes* are behind upstream synchronous code division multiple access (S-CDMA) to allow parallel, non-interfering transmission. The lower picture of Figure 20 shows how subcarrier frequencies individually remain orthogonal from one another over a single OFDM symbol period, T_{OFDM}.

Getting back to the long form capacity expression above, the total capacity is then simply the summation of the individual capacities of chunks of spectrum. An advantage of this form of the expression is to recognize that channels may not have a fixed SNR characteristic, such as in the case of uncharacterized spectrum above today's 1 GHz forward band as we saw in Figure 16.

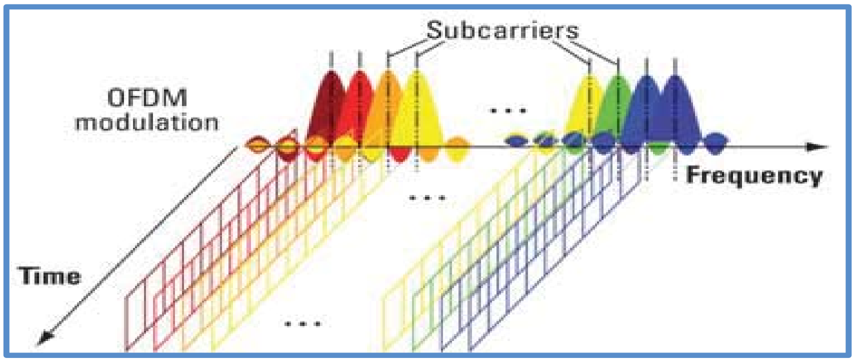

http://www.lightwaveonline.com/articles/ofdm-to-power-high-bit-rates-in-next-gen-optical-access-networks.html

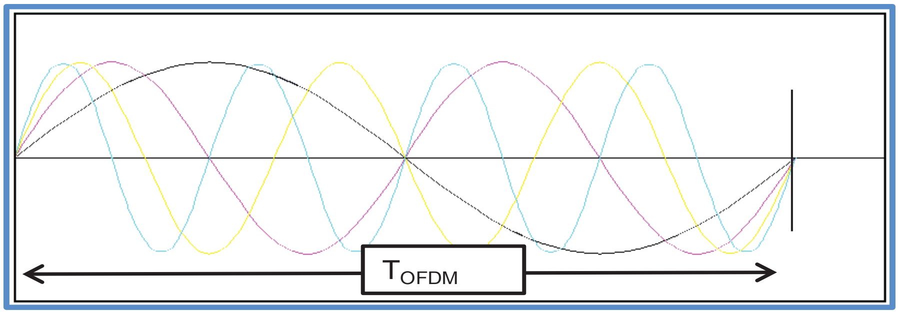

http://www.dipol.ie/newsletter2/inf_dipo_2009_32.html

Figure 20 – The Fundamental Frequency and Time Domain Principles of OFDM

When such variation exists, the capacity of a not-flat SNR region can be calculated by looking at it in small chunks that, because of their narrow width, themselves approximate flat channels.

A similar approach to net capacity applies when there is, for example, interference. The effected OFDM subcarriers will have a lower SNR (in this case S/(N+I). This flexibility is a key advantage of OFDM – very narrow channels, each of which can be individually optimized. Alternatively, whether talking about frequency response variation or interference, a specific implementation may include processing that effectively averages the SNR over the subcarriers, and an optimum QAM and FEC chosen from the output of this estimate.

An important difference between frequency response variation and interference, however, is that in the case of frequency response variation, the signal and major noise contributions (the linear optics) may be varying together, so subcarrier SNR may not be affected, except as far as CM receive levels impact the link SNR. For interference, the signal to interference ratio (SIR) is a direct hit to baseline SNR.

Note that we have used the acronym "OFDM" exclusively in the above discussion. In the upstream, however, the technology will in fact be OFDMA – Orthogonal Frequency Domain Multiple Access, whereby different transmitting users may have a subset of the subcarrier set of a given OFDM symbol. Nonetheless, the governing principles and advantages of OFDM apply in the case of OFDMA, but with the added complexity of scheduling, RF power management, and maintaining orthogonality across users (as in S-CDMA). Figure 21 shows a comparison of OFDM and OFDMA.

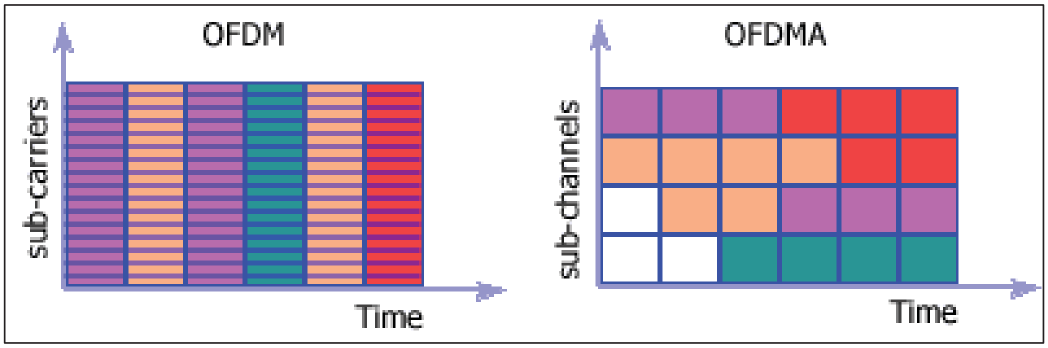

Figure 21 – OFDM vs. OFDMA [17]

Linear Distortion Implications

We described the frequency response roll-off above the specified bandwidth of 1 GHz components above. This is an example of linear distortion. However, significant variation of frequency response is not limited to operating HFC outside of the guaranteed component bandwidths. The forward and reverse path can take on droop and ripple over time, sometimes quite significant. Ripples in the broadband response are often associated with aging cables and connectors that spawn poor terminations and ingress/egress points. Poor terminations exist in the plant, and it is almost a guarantee there will be open terminations in the drop and home network contribute to frequency domain ripple. Cable CPE in the home itself has a minimum in-band return loss of just 6 dB. Lastly, multiple dwelling units (MDUs) e.g., apartments are often fertile ground for the more extreme linear distortions due to the stretching of coaxial design guidelines, different rules of engagement, and the dynamic nature of the terminating RF environment – turnover and occupancy may vary.

Typical HFC linear distortion components tend to be static or very slowly varying, and their effect on transmission is generally handled by OFDM's channel-optimizing qualities. In wireless systems, for example, the environment to deal with is the much more rapidly varying, fading environments as users transit cells. Frequency response due to signal reflection (multi-path) is common in HFC.

Periodicity (frequency ripple) in the frequency domain is correlated to the time delay between the two reflections points through a very simple relationship:

$$\text{Delay } (\mu\text{sec.}) = 1/(\text{Frequency response peak-to-peak, MHz})$$

Of course, because distance and time are related, we can determine the distance, which helps figure out the contributing components through (distance) = (c')/(delay/2), where c' is the speed of propagation in the coax, which is approximately 1 nsec/ft (this value is highly dielectric dependent – check the constants for your own cable types).

An understanding of the micro-reflection delay profile guides the design requirements for the equalization subsystem in traditional single-carrier systems. However, for OFDM it instead drives the selection of a critical OFDM system design parameter known as the Cyclic Prefix (CP), which from a reflection energy perspective acts asa guard interval. Table 1 is drawn from channel model derivations applicable to the HFC downstream and upstream recently determined through available DOCSIS® 3.0 metrics and contributed to IEEE 802.3bn working group. [22].

Table 1 – Downstream and Upstream Micro-reflection Assumptions

Downstream Echo	**Downstream Echo Profile, dBc**	**99%**	**Majority**
	.5 usec	-20	-30
	1 usec	-25	-35
	1.5 usec	-30	-35
	2 usec	-35	-40
	3 usec	-40	-45
	4.5 usec	-45	-50
	5 usec	-50	-55
Upstream Echo	**Upstream Echo Profile, dBc**	**97%**	**Majority**
	.5 usec	-15	-10
	1 usec	-20	-20
	1.5 usec	-25	-30
	2 usec	-30	-30
	2.5 usec	-35	-30
	3 usec	-40	-30

Single carrier transmissions manage micro-reflection energy through tapped-delay line equalizers – a Linear Equalizer (LE) in the downstream, and a Decision Feedback Equalizer (DFE) in the upstream. The total time span of the taps is constructed to mitigate the range of dispersion that the signal may endure in traversing the channel. Figure 22 shows an impulse response estimate available using DOCSIS® 3.0 tools in the upstream that captures a micro-reflection about 27 dB below the main signal with a delay of about 1.3 usec (actually beyond worst case DOCSIS® channel assumptions shown in red).

As frequency response degrades and the data speeds get higher (the signal bandwidth gets wider), shorter per-tap time spans are required and thereby more taps are needed, and the complexity rises. It is perhaps more easily intuitive to understand the complexity in the frequency domain – it is harder to equalize, or flatten, a wide channel than a narrow one. But, the single reflection time-domain concept is simple enough – if a delayed version of the signal arrives at some amplitude and phase 1 usec after the main signal (which according to the table will be attenuated by 25 dB, 99% of the time), a tap at 1 usec could be adjusted to cancel this out. Otherwise, if the SNR is 35 dB for the signal, then it would be unacceptable to add a 25 dB interference source and expect to maintain performance.

In a more general sense, since perfect cancellation and inverted frequency response leads to noise enhancement, the equalizer taps are adjusted to minimize the mean-square error of noise plus intersymbol interference (ISI).

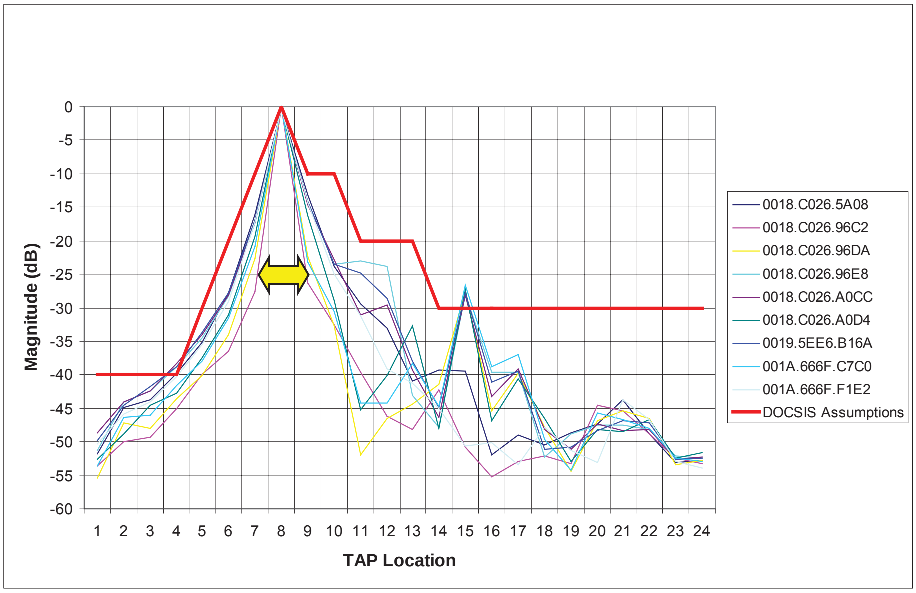

Figure 22 – Classic Impulse Response View of Single Carrier Equalization

OFDM, by contrast, works on a simpler, and less complex, premise. The symbol rate is much slower, and the subcarriers therefore very narrow. Over any independent, narrow, subcarrier, the channel looks flat. Different channels may require individual adjustment to equalize the band – a single tap multiplication – but any specific subcarrier sees a reasonably well-behaved channel because the subcarriers are so narrow.

In the time domain, the important relationship to maintain is that the reflection energy (multipath is the common wireless terminology term) cannot arrive in a way that interferes with a subsequent OFDM symbol. This is avoided by essentially "waiting out" the delay and sending the next symbol after the reflection delay – or at least at a point whereby the reflection energy is significantly attenuated so as not to contribute significant ISI energy to the link SNR. This phenomenon is depicted in Figure 23.

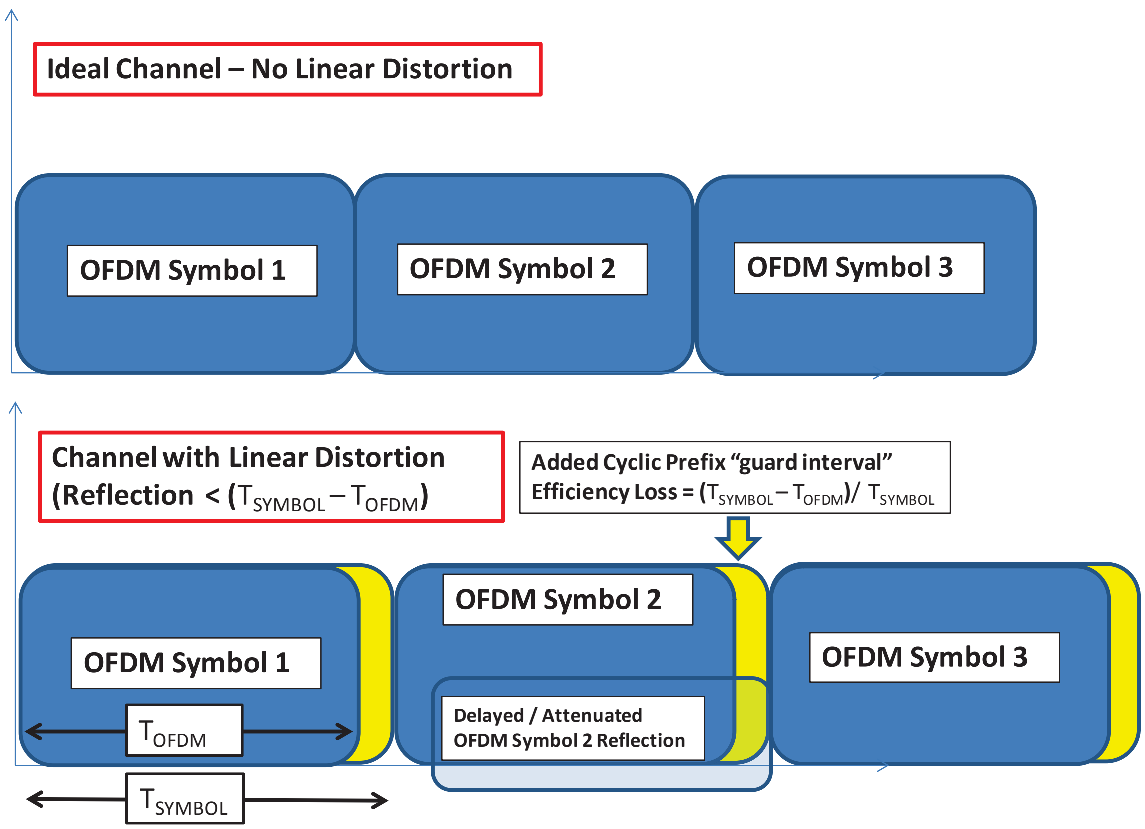

Figure 23 – OFDM's Handling of Micro-reflections (Multi-path)

This "guard time" or "cyclic prefix" (CP, referring to how the delay is implemented in the IFFT versus simply a "guard time") means that no data is being transmitted for a period of time, and thus link efficiency decreases. Therefore, an important OFDM system parameter is to make sure the CP is small relative to the symbol duration. Table 1 can therefore be used in OFDM to select a suitable range of CP options, as well as being a guideline to the subcarrier spacing which defines an OFDM symbol period.

The CP is how OFDM handles linear distortion in the form of micro-reflections. However, while we have described amplitude variation in prior discussion, we have not significantly addressed the fact that wherever there is amplitude variation there will also be phase variation. The roll-off area in Figure 16 is a good example, because where there is steep roll-off, a steep curve of group delay variation (GDV) will ensue, and this is dispersive.

GDV can be viewed as kind of a "continuous" form of micro-reflection that is a granular function of frequency. The use of CP as it applies to micro-reflection applies also to the case of GDV. If the GDV from one end of the OFDM band to the other is high, then one subcarrier will arrive delayed from the other. If the delay is longer than the discrete micro-reflection source, then this dispersion will guide the CP choice.

External Interference

This section summarizes the primary sources of interference that have the potential to limit network capacity and some recommendations on what can be done about them [3]. The problem in fact has two reciprocal aspects, both of which are important to operators: Ingress and Egress.

First, we recognize the two basic mechanisms for injecting interference onto the cable network.

- Radiated: Over-the-air signaling penetrates the cable shielding, generally due to cable integrity flaws resulting in insufficient shielding. Another ingress portal is associated with consumer electronic (CE) devices in the home connected to the cable network. If these devices lack sufficient shielding to prevent ingress, they also become potential ingress conduits. LTE, TV and radio stations, shortwave transmissions, and HomePlug® are common examples for radiated mechanisms.

- Conducted: The sources of conducted ingress are typically found within the consumer home. Signals are directly injected into the cable network, usually via the home wiring. Grounding integrity also plays a role in conductive interference. Examples: satellite TV, MoCA™, electro-mechanical systems in the home.

The prevention of ingress is dependent on whether the ingress source is of the radiated or conducted type. We now examine potential interference sources below, and, in particular, note the band of operation and mechanisms for each.

<u>Fixed Frequency Interference</u>

4G Wireless / Long Term Evolution (LTE)

LTE wireless service is aggressively being deployed in the US and Europe and will become the primary wireless transmission type over the next 5 years. LTE differs from previous cellular technologies such as 3G in that it uses OFDM for the downlink (as does DOCSIS® 3.1) and Single Carrier – Frequency Division Multiple Access (SC-FDMA) for the uplink. SC-FDMA can be viewed as a peak-to-average optimized form of OFDMA, so is also a multi-carrier approach. LTE currently operates from 700 MHz to 800 MHz in the USA, and 800 MHz to 862 MHz in Europe. However, there is already talk that future LTE services may also occupy the 600 MHz to 700 MHz spectrum. The transmit power at the base station is typically 62 dBm (Effective isotropic radiated power or EIRP) to achieve a 1 Mbps data rate.

In the case of LTE, there are two primary scenarios – transmission from the base stations (cell towers), and from the User Equipment (cell phones). These scenarios are shown in Figure 24. The distance from these sources will have a direct correlation to the field strength that can potentially enter the cable network. Homes very near LTE cellular towers will be most susceptible to high field strengths and users operating cell phones furthest from the towers will tend to transmit the highest powers. Further field studies need to be done to fully understand the relationship between performance, cell design parameters, and ingress/egress levels.

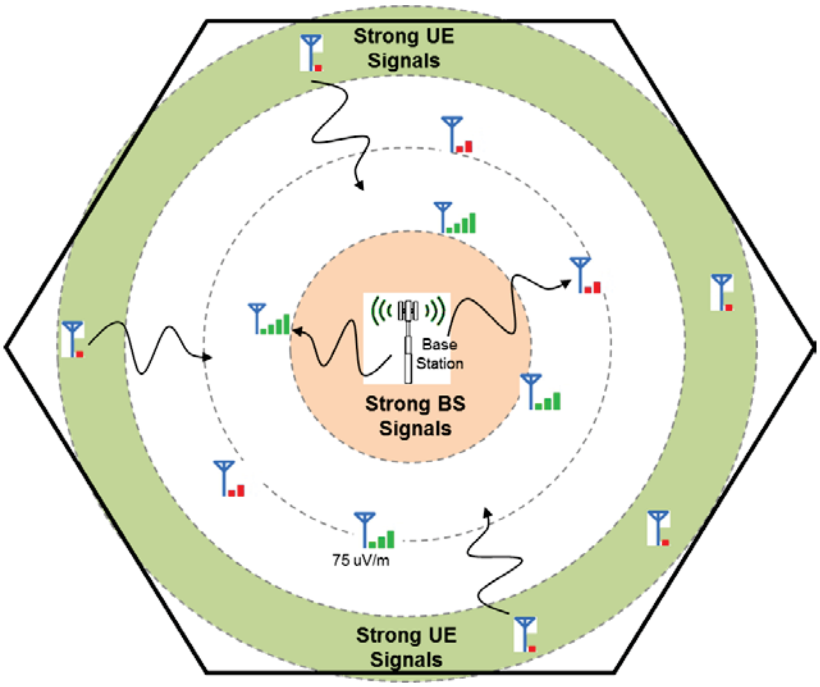

Figure 24 – Potential LTE Interference Scenarios [3]

Cable egress in the LTE frequency range can interfere with LTE services. Egress is RF signaling that escapes from the cable infrastructure, normally due to a cable shielding integrity flaw. Digital QAM leakage in the 700 MHz to 800 MHz range (USA) has been known to raise the noise floor in affected locations, and this can degrade the performance of the LTE service, as shown in Figure 25 [3].

And, of course, LTE ingress into the cable network can cause performance issues for cable services near LTE frequencies, often associated with the same causes as egress.

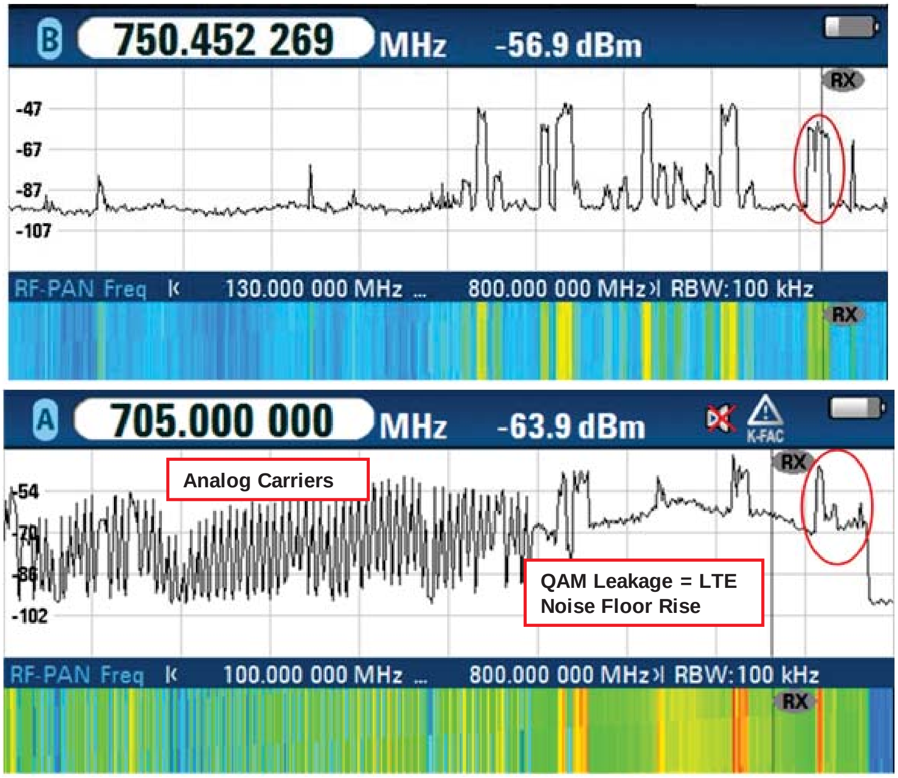

Figure 25 – OTA Spectral Scan Before and After a Large Cable Leak [3]

In order to prevent LTE ingress from entering the cable network via cable integrity flaws, the prevention methodology is to find and eliminate egress sources. Based on early data findings, the frequency used for the egress detection should be less than ½ of an octave from the LTE frequency band.

In order to prevent ingress from entering the cable network via the attached end-user CE devices, operators and equipment vendors will need to ensure the shielding effectiveness of these products is sufficient to mitigate the ingress. For example, an LTE cellular phone transmitting at full power generates a 23 dBm signal inside the home. This is sufficient power to affect the overlapping QAM channels if the cable shielding is inadequate. New shielding effectiveness standards for in-home cable equipment should have minimum recommended ratings.

In general for cable integrity issues leading to ingress and egress, a systemic approach of detection and cable repair must be taken. An effective methodology for detection is based on a fleet vehicle approach that covers approximately 90% of the cable plant every 3 months. The vehicles are outfitted with cable leakage detection equipment that is able to locate and log leakage found into a mapping platform. The number of vehicles outfitted with such equipment is dependent on their ability to cover a high percentage of the cable plant in a given period of time. Then, a methodical approach to repairing the leaks should follow.

While cable integrity has always been a major part of cable operations, and it is always a sound recommendation to ensure cable and connector integrity, emergent LTE services and the introduction of higher order QAM formats in DOCSIS® 3.1 have upped the ante.

Satellite TV

Satellite TV downlink signals operate well above the cable network band. However, the signals are down converted at the dish to operate in the 250 MHz to 750 MHz band for Ka-Band systems, and 950 MHz to 1450 MHz for Ku-Band systems, and enter the house on coaxial cable. When these down-converted signals are improperly connected to cable network, interference issues can result from the overlapping spectrum, and the impact could potentially be felt in neighboring households as well.

Diligent install practices must be in place to avoid this mixing of services over a common coaxial home network, for example, when a customer with broadband Internet from the MSO has Satellite TV services. The best prevention is to ensure that consumers using a satellite TV service are using a network isolated from the cable company's network. A filtering solution may be possible and in some cases required to avoid Satellite signals from exiting the home and interfering with a neighbor's service.

MoCA™

Multimedia over Coax Alliance (MoCA™) uses the existing in-home cable infrastructure to create an IP home network, eliminating the need to run new Ethernet cabling or rely on WiFi™ coverage. MoCA™ also uses OFDM technology to ensure robustness in the unpredictable, uncontrolled subscriber home network.

MoCA™ specifications allow operation from 500 MHz to 1500 MHz. For coexistence with cable services, MoCA™ is used in the frequency bands that are unused by the cable system. For 1 GHz systems, for example, MoCA™ channels would be confined to the 1GHz to 1.5GHz portion of the spectrum.

The problem that MoCA™ can cause to a cable system, even though it may be operated outside of the cable band, is the overloading of existing CPE that are receiving cable signals. Strong MoCA™ signals that enter a tuner or analog front end can cause intermodulation components that end up in the cable band, and are high enough to interfere with low-level cable signals and sensitive analog video. Care must be taken to properly isolate MoCA™ signaling, including signals that could escape the home and travel to surrounding homes.

A MoCA™ node has a maximum output power between -4 dBm to +8 dBm when operating in band A, B, C, or D (875-1450 MHz), and between -1 dBm to +7 dBm when operating in band E (500-600 MHz), at every supported MoCA™ channel frequency. By comparison, a nominal CPE received power on a given 6 MHz channel may be 0 dBmV, or -48.75 dBm.

Isolation filtering in the home and at the cable point-of-entry (POE) to avoid disturbing the neighbors is often required. These filters are also generally installed on STB's that are not MoCA™ compliant. Many operators, having deployed CPE with this capability, have established recommended practices for MoCA™ home LANs. These practices may need to be continually updated as MoCA bands expand and new specification revisions become standardized.

HomePlug®

HomePlug® is the name given to the technology that uses the existing in-home AC electrical wiring as the copper media for the IP home network. The idea of HomePlug® is to allow any outlet in the house to be a Home LAN access point, mitigating access constraints of other wired technologies such as MoCA™, and of course providing a solution for homes that do not have coaxial infrastructure at all.

HomePlug® is also based on OFDM and operates from 2 MHz to 30 MHz – sitting right on top of a large swath of the already limited HFC return band, making it a risk to the upstream if it makes its way onto the home coax network.

Of course, since every cable CPE plugs into a wall outlet, it is potentially exposed to HomePlug®. HomePlug® ingress is a result of radiated emissions in the 2 to 30 MHz from the in-home electrical wiring onto the cable network. These emissions enter the cable network via integrity flaws within the home such as poor connections, bad connectors, splitters, consumer devices, etc. Prevention calls for use of a cable leakage detection device during any installs or service calls to ensure the cable shielding integrity is intact. While this does not prevent consumers from altering the cabling inside their homes, it will at least reduce ingress caused by improper installations.

TV "White Space"

TV White space is traditionally defined as the unused spectrum between broadcast analog TV stations. Full power analog television broadcasts ceased operating in 2009 in the Unites States per the digital switchover mandate. In 2008, the FCC had voted 5-0 to approve the unlicensed use of white space. The use of this spectrum in an unlicensed fashion still mandates that the service must NOT interfere with the existing over-the-air licensed services, such as digital TV.

An example of a proposed "White Space Device" (WSD) is extended WiFi devices that would use the white space to provide outdoor coverage. These devices concern cable operators, particularly in the areas of channels 2 through 4 because of direct pickup sensitivities. As has long been the case, cable operators in rural areas are faced with intercepting very weak television signals from remote locations for local residents' entertainment. High-powered WSDs for data access have the potential to overpower weak signals occupying nearby spectrum.

White space remains a wild card for the industry. Proposals exist recommending that the channel 2-4 spectrum be mandated free of WSDs.

Transient Interference Sources

Cable systems are also exposed to a wide range of environments, and the operator's network itself includes the inside of every home it is connected to, at least until the drop/home architecture is evolved to an HFC-terminating point-of-entry (POE) gateway. As such, there are many potential sources of transient noise. The majority of the noise with the energy to do great damage exists at the low end of the return band.

While the upstream is more troublesome, data and analysis of field captures show that transient, or "burst" noise can also take place in the downstream. Characterizing time-domain interference parameters is important for optimization of the OFDM system design. The symbol energy and time-domain interleaving aspects of SC-QAM and OFDM are quite different, and the proper FEC design must account for the likelihood and statistics of burst events.

Figure 26 provides a snapshot of a cable upstream. The two charts in this snapshot are of both frequency (top yellow) and time (bottom green).

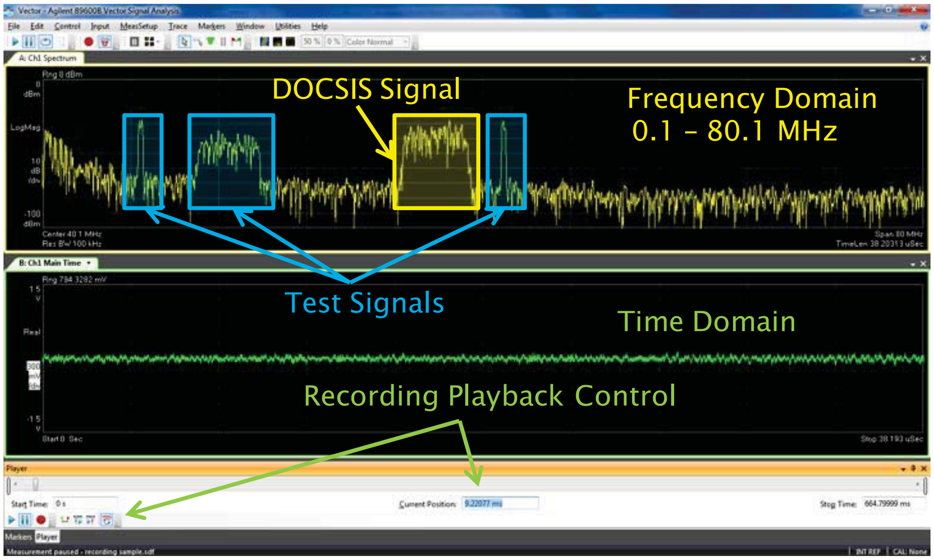

Figure 26 – Snapshot of a Return Path in the Frequency and Time Domain

The operating band of the cable upstream in Figure 26 is 5-42 MHz. Measurements typically include 80 MHz of spectrum to differentiate laser clipping from burst noise, because they behave similarly and can be easily confused with one another. The test signals shown are constant carriers for reference to assess the impact of burst noise. The DOCSIS® signal is an actual customer's transmission.

The instrument used to provide this snapshot is an Agilent MXA-Series Vector Signal Analyzer (VSA), model N9020A. Because burst noise is a transient impairment it is often advantageous to record the impairment in real-time and then playback (via the playback control shown in the bottom of the Figure 26) in slow-motion in order to observe the true nature and impact of the impairment. Currently, accurate characterization is limited to expensive test instrumentation and expertise often times not readily available to HFC technicians.

Sources of Burst Noise

Burst noise impairments can originate from a variety of sources including hair dryers, auto-ignition circuits, dimmer switches, and fluorescent lights. Under normal operating conditions, emissions from these devices should be benign because networks are maintained and properly isolate these potential sources. However, problems arise most often when there are defects in the craftsmanship or maintenance of the coaxial network, aggravated by subscriber-"engineered" coaxial service distribution.

Defects can occur in a variety of ways including loose or corroded F-connectors, damaged cables, and unterminated ports or drop cables. These defects result in degraded or poor system matching, and open the door for the above electronic emissions to ingress into the network.

Many of the burst noise impairments observed are short in duration, lasting 10s to 100s of microseconds. The RF energy associated with these impairments is typically very high and has very broad spectral properties, limited only by the filters in the coaxial network itself used that separate the upstream from the downstream. Often times, burst impairments occur in a repetitious manner.

The impact that burst noise impairments will have on an upstream transmission is temporarily degraded SNR. This degradation could impact a series of SC-QAM symbols or a series of OFDM-QAM symbols (but many fewer). The universal best way to fix a burst noise issue, regardless of signaling scheme, is to diagnose, characterize, and ultimately isolate the defect that enables the impairment to enter the network. This is not always immediately practical, especially given the nature of many of the sources that are within the subscriber's home.

Another way to combat burst noise impairments is to mitigate its effects using common digital signal processing (DSP) tools. With typically vastly different symbol durations, the impact to SC-QAM and OFDM-QAM and the strategies to mitigate the burst noise will vary. However, they generally involve some combination of FEC, interleaving, and design of the waveform itself. Tradeoffs can be made in overhead, latency, and complexity among these tools.

Characteristics of Burst Noise

Figure 27 shows a burst impairment observation. There are three key characteristics of burst noise: duration, relative amplitude, and periodicity.

From the time-domain chart of Figure 27, we can see the duration of the burst noise event. Each division is approximately 3.8 μsec long and the energy of the burst noise impairment spans less than that.

The impact to DOCSIS® signals may be assessed by noting the SNR associated with the duration of the burst noise. SNR is severely degraded as observed in the frequency domain chart. In Figure 28, we can see that the increase of RF energy throughout the RF spectrum, and, in the lower time domain plot converted to dB units, the degradation can be quantified by measuring during the burst and outside of the burst.

In this specific case, the SNR will degrade by approximately 15.7 dB over the duration of the burst. Thus, if the SNR of the DOCSIS® signal is 40 dB with no burst impairment, then the SNR will degrade to approximately 24.3 dB during the burst event and likely corrupt all the symbols during that time. For 64-QAM, 24.3 dB SNR represents an uncoded BER of approximately 1e-4, which is an uncomfortable operating point likely to consume significant FEC margin.

Figure 27 – Burst Noise Observation

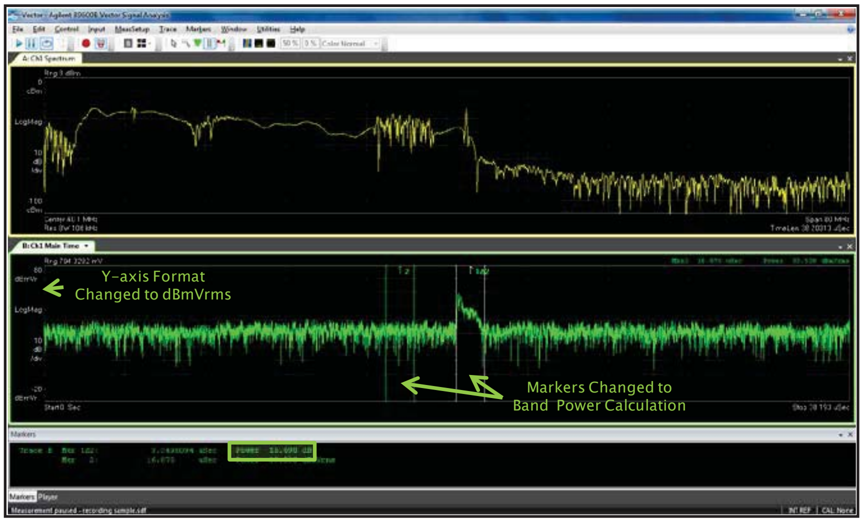

Figure 28 – Characterizing Relative Amplitude

Lastly, the periodicity of the impairment can be easily evaluated as shown in Figure 29 when viewed on the proper time scale.

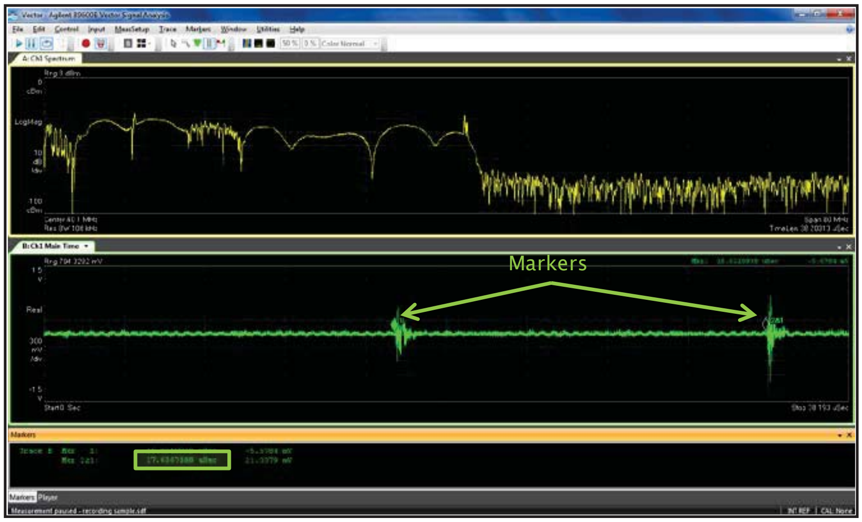

Figure 29 – Characterizing Periodicity

In summary, burst noise typically has short duration, significant RF energy, and often also has periodicity. Ideally, these parameters are characterized statistically so that system designs can be optimized, and operations can recognize when conditions deteriorate beyond "normal" transient noise behaviors. Figure 30 shows an example of a large sample of *downstream* bursts at a particular location on a cable plant revealing key burst duration vs. amplitude relationships that can drive statistical fit characterizations. Significant recent progress has been made in understanding downstream burst impairments, and this will allow for better optimization of OFDM system parameters, and possible lead to future variations of OFDM "numerology."

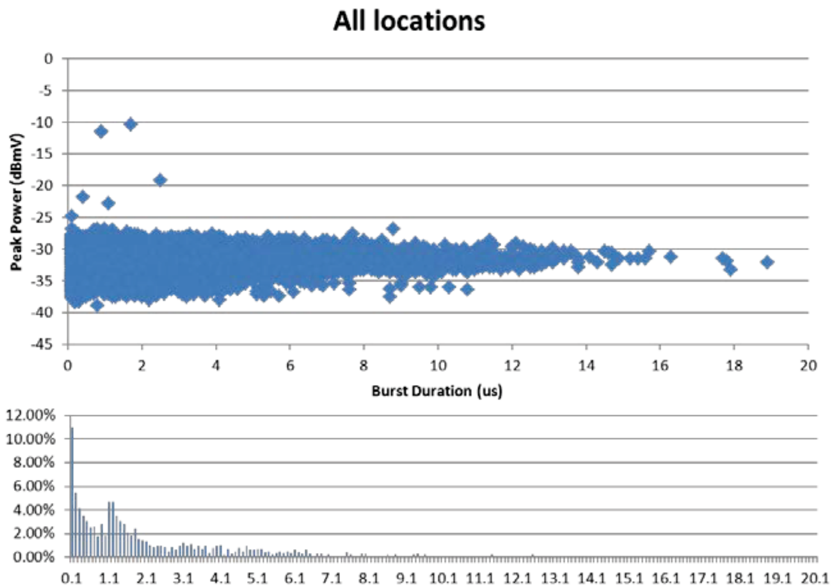

Figure 30 – Robust Downstream Burst Characterization Is Emerging in New Generations of CPE (Courtesy Maxlinear)

Accurate characterization combined with spatial relationships of the CMs with degraded SNR and the understanding of which burst impairment characteristics are associated with which emissions of electronic devices can lead to isolation and minimization of their contribution to cable network performance degradation.

What Will OFDM Do?

Fixed Frequency Interference

As has been described, key benefits of OFDM are its handling of difficult channels and adaptability to conditions. Choosing OFDM/OFDMA for DOCSIS® 3.1 was to take advantage of its ability to maximize capacity under any conditions – for example, allowing 4096-QAM where possible, but also increasing robustness so that typical performance or graceful degradation can be maintained under deteriorated conditions as opposed to outages. In this latter sense, the benefits of OFDM's channel-optimizing qualities extend to the interference scenarios described above. We briefly describe how DOCSIS® 3.1 OFDM is outfitted to deal with these impairment types when they are present.

Figure 31 points out the primary conceptual difference of "narrowband" interference when we move from a traditional SC-QAM system to OFDM. Namely, what is narrowband for 6 MHz SC-QAM slots may not be narrowband for OFDM subcarrier bandwidths. Also, because the power per subcarrier is a small fraction of the total power, an interferer that may be insignificant to a full band QAM carrier may seriously degrade the MER of a subcarrier or several that it falls near (nearby because of the slow roll-off of OFDM sidebands).

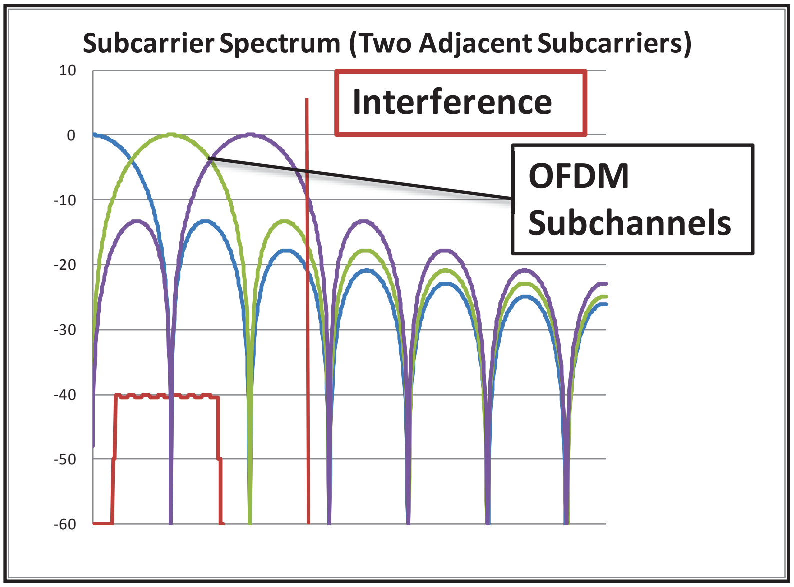

Figure 31 – Narrowband is Often Not-So-Narrowband for OFDM

In Figure 32, a practical cable example shows how this plays out for guaranteed "interference" sources composite second order (CSO) and composite triple beat (CTB) for systems that have analog loading – though the mechanism is that of a nonlinear distortion product. For the SC-QAM signal at the top of Figure 32, the CSO and CTB beats have bandwidths on the order of 10-20 kHz. These are narrow, modulated piles of spectrum, at levels of at least 47 dBc down each (to digital QAM, via 53 dBc FCC analog requirement) and typically much better.

The lower figure of Figure 32 compares the relationship a CSO/CTB distortion beat pile may have relative to a 25 kHz OFDM subcarrier bandwidth if it falls completely beneath it. It is no longer a "narrow" interference but a wider noise block. Furthermore, what was 53 dBc for a 6 MHz reference QAM power is now 24 dB worse (29 dBc) because of the fraction of power that a single subcarrier will have on the average. Since these particular distortions have a varying amplitude, the peak of the distortion interference can actually be significantly higher.

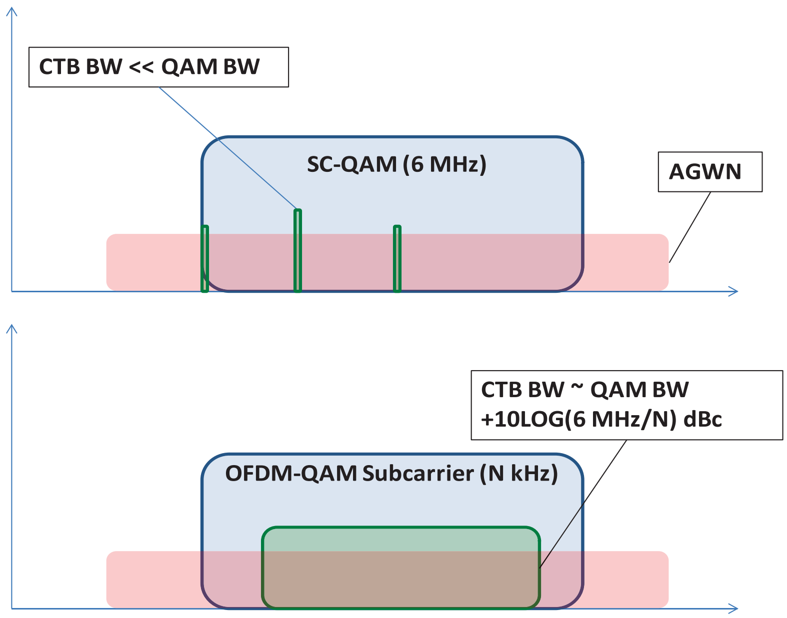

Figure 32 – SC-QAM and OFDM-QAM with Analog Distortion Beat Interference

For SC-QAM, it is important that the interference stay below a certain threshold, beyond which a BER "crash" can ensue when AWGN is imposed on top of a signal-to-interference ratio (SIR) offset. For OFDM, a graceful degradation approach could be managed by bit loading to the subcarrier quality. For example, an SNR of 29 dBc with a 3 dB margin of operation would still enable 512-QAM according to Figure 7. If the peaking of interference led to much worse degradation, the subcarrier might be nulled altogether. In either case, only a small number of subcarriers are affected of the whole, so the effect is a minor loss of capacity. This is in contrast to a thresholding effect that single carrier systems have with narrowband interference levels and bandwidths beyond what ingress cancellation can handle.

Alternatively, with soft decision decoding, the low SNR subcarrier is part of a large codeword of many subcarriers, and decoder performance operates on the average SNR of the set. The end result may be a lower average SNR, but if they number very few (such as CSO/CTB beats), not so much lower as to impact QAM profiles enough to degrade capacity – just lost margin. These are implementation specifics of QAM-LDPC decoding algorithms.

In either case above, distortion interference and potential capacity or margin losses can be eliminated with the removal of bandwidth-inefficient analog video services.

Figure 31 and 32 described cases of discrete narrowband interference. An emergent interference concern for cable operators is the growing LTE services described earlier and occupying the 700-800 MHz band in the USA. Figure 33 pictorially tells the impairment and mitigation strategy of LTE using modern OFDM processing tools.

In the top and next lower figure of Figure 33 we see what 10 MHz of LTE ingress beneath SC-QAM slots in the downstream may look like. "Typical" interference, where it exists, may be 30 dB when present on the plant [21][22]. It can (and has) wiped out specific 6 MHz QAM carriers, which may result in the loss of a dozen digital TV programs or make a data carrier unusable.

In the third figure from the top in Figure 33, we can see how LTE interference might look to an OFDM spectrum. Because all of the interference falls in a contiguous band, there is a reasonable probability it exists within one FEC codeword, as shown, although this is not a necessary condition to manage the interference (and in fact, spread across is better). As with do with time domain interleaving and burst noise using Reed-Solomon J.83B today, it is better to spread out errors across codewords to make the FEC work most effectively. We can manage the same effect using the principle of duality with OFDM and subcarrier interleaving when there is frequency domain interference.

The bottom figure of Figure 33 shows how the degraded subcarriers are spread amongst codewords following de-interleaving at the receiver. Spreading the errors around allows the FEC to be most effective in correcting channel errors, yielding the smallest degradation from ideal AWGN and the best probability for achieving the highest order QAM profiles. Frequency domain interleaving is a tool unique to OFDM because of its use of narrow subcarriers.

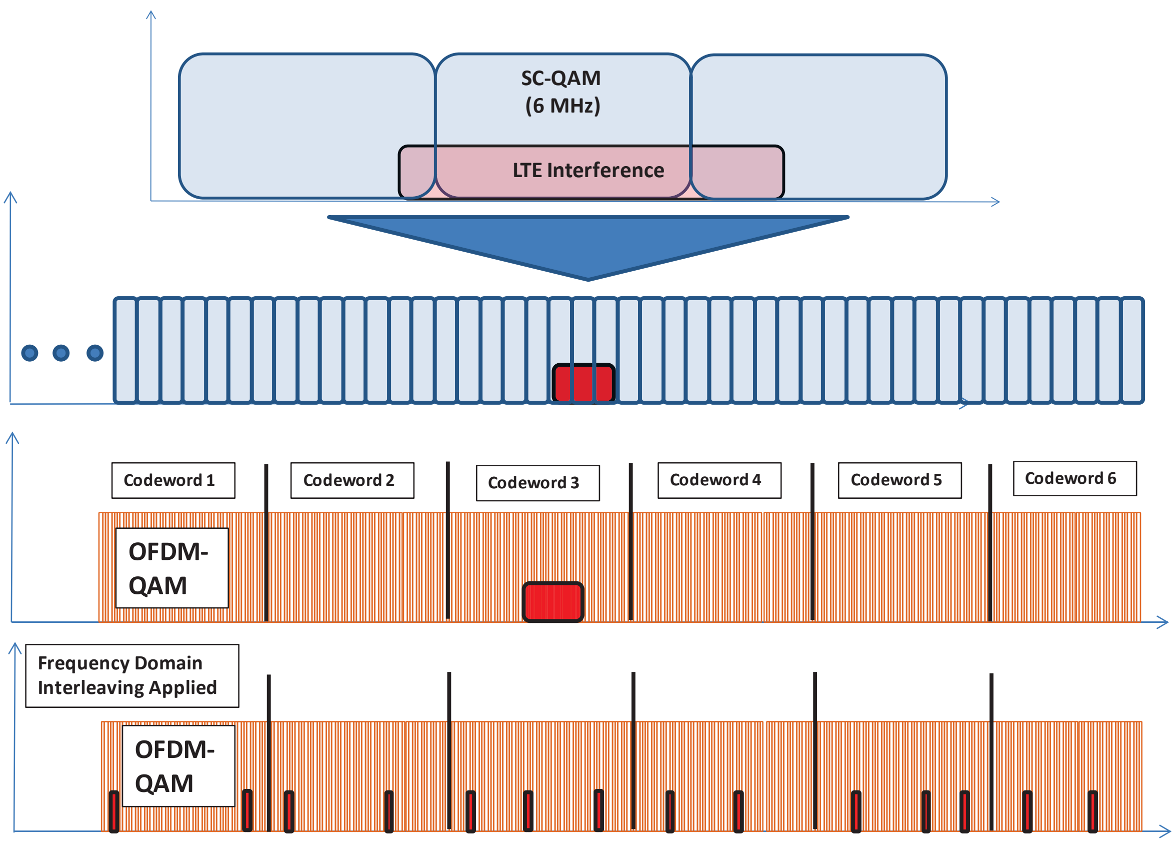

Figure 33 – SC-QAM, OFDM-QAM, and Frequency-Interleaved OFDM-QAM with an Interference Band (e.g. LTE)

Transient Interference

In many cases burst impairments could wind up being the most significant problem MSOs face when attempting to fully load the HFC upstream. Intermittent poor performance is hard to pinpoint. Transient noise events capable of wiping out substantial numbers of consecutive SC-QAM symbols have been observed, and similarly a full band of OFDM subcarriers could be wiped out. Temperamental upstream performance and uncorrectable downstream codeword errors – indicating transients that outlast the J.83B time-domain interleaver – suggest this can happen in both directions, though the upstream's frequency band and funneling topology makes it inherently more prone. Figure 30 is further, new, evidence that the downstream is not immune to burst noise.

Note that damage can occur in both legacy DOCSIS® because symbols are quite short relative to the duration of some transient impulses, and in DOCSIS® 3.1 because of its wideband nature. Whether a transient is classified as impulse or burst is a matter of whether the noise profile in the frequency domain is colored or white (flat) in nature.

It is important to note the likelihood that this impairment type will become even a larger concern since there is an increasing list of potential contributors, including HomePlug, HPNA, and LTE (at the home, LTE upstream has burst qualities to poorly shielded cabling and CPE). Adding to the difficulty is that the primary symptom of burst noise with today's tools is an error rate problem when there does not seem by other traditional metrics to be any problem at all.

What can DOCSIS® 3.1 do about burst noise? Two techniques already deployed in DOCSIS® are effectively combined and re-deployed for DOCSIS® 3.1. First, OFDM symbols are much longer in time than SC-QAM symbols. In many cases they may be longer than the burst. The OFDM symbol duration advantage is analogous to S-CDMA "symbol spreading," which is done instead by codes for that technology. Thus, an inherent symbol energy relationship exists that may be favorable to longer duration symbols for some burst characteristics. With SC-QAM, the only hope for bursts exceeding a few μs in duration is the de-interleaver's ability to spread the affected symbols out in time at the receiver so that that the FEC can operate closer to a random noise scenario and not exceed its maximum number of errors-per-codeword that can be corrected.

DOCSIS® 3.1 also will employ interleaving. An example of OFDM interleaving is shown in Figure 34 [15].

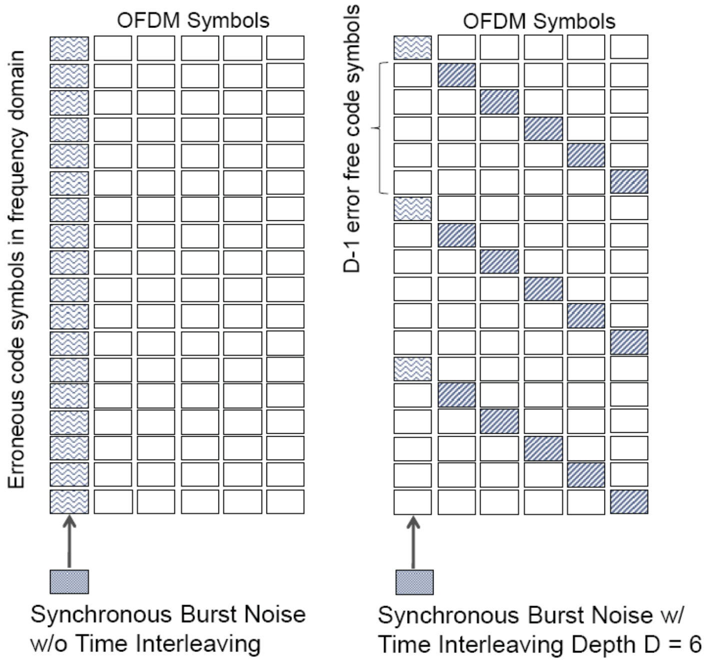

Figure 34 – Interleaving of OFDM Symbols to Manage Burst Noise [15]
(Diagram courtesy of Qualcomm)

With the long symbol times of OFDM, a challenge becomes the poor time granularity of interleaving and the therefore the impact that this can have on latency. Each symbol of interleaver depth is a relatively long delay, and a depth necessary to mitigate bursts can threaten latency budgets. There is a necessary trade-off between burst correction capability and latency. However, the interleaving concept is fundamentally the same. In OFDM, a large burst harms a set of parallel transmitted subcarrier QAM symbols, and the interleave process is to spread these degraded QAM symbols that are from the same time slot into adjacent OFDM symbol time slots uniformly to help the FEC decoder do its job.

Figure 35 is an example of how OFDM interleaving can be effective for a given QAM, symbol duration, burst, and latency parameters consistent with commonly assumed cable parameters and levels [15]. While burst mitigation is powerful, there is nonetheless up to a 3 dB loss, which translates to a half-modulation loss of capacity loss due to burst. For an SIR of just 20 dB during the burst duration, at least the result is not an outage.

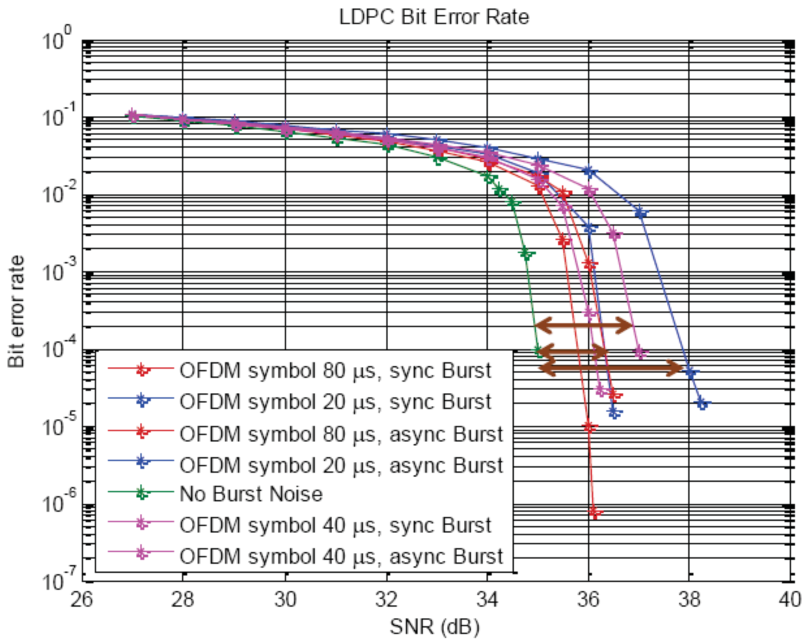

Figure 35 – BER vs SNR and OFDM Symbol Time for 4096-QAM with DVB-C2 FEC (16,200, R = 8/9) in Burst Noise (20 usec, 20 dB SIR) at Fixed Interleaver Depth [15] *(Simulations courtesy of Qualcomm)*

Capacity Optimization

Some useful definitions were contributed as part of the IEEE 802.3bn initiative in order for the members of the standards committee, ad-hoc working groups, and adjacent functions could be speaking the same language. We discussed the optimization feature of OFDM for difficult channel previously, and formalize terms here [4].

Bit Loading – QAM order for each subcarrier can be selected based on its individual narrowband channel quality. This ability to optimize subcarrier QAM orders across the total bandwidth of the communications channel is known as *Bit Loading*. The term bit loading commonly assumes that FEC does not vary across subcarriers.

Modulation and Coding Scheme (MCS) – Another way to match the OFDM subcarriers to channel quality is referred to as the *Modulation and Coding Scheme (MCS).* This technique is realized by applying the same modulation and coding to all the subcarrier carrying the code word, matching the average channel quality. In general, different MCSs may differ both in terms of the QAM order of modulation and code rate (k/n) and typically, for each QAM order, there may be multiple code rates available to select.

The above descriptions express the allocation of subcarriers as a function of frequency. When bit loading or modulation and coding can adjust over *time*, they become *Adaptive* Bit Loading or *Adaptive* Modulation and Coding. Link adaptation can vary from a few times a day or hourly, for example, in a relatively static channel such as cable systems, to being session-based on even frame-to-frame based in more dynamic environments such as wireless.

We can also factor in a user dimensions to adaption and optimization. In the case of a “single profile” approach, the selected scheme is applied in the downstream direction for all users. To be effective, this approach is based on targeting the worse user channel conditions in the plant. DOCSIS® 3.1 is adopting an approach with *Multiple Modulation Profiles*, or MMP. In MMP, the selected scheme is applied for a subset of the users in the plant, where users are grouped into bins according to their channel conditions.

MMP is a *major* deviation for DOCSIS® 3.1 that does not exist in today’s broadcast downstream, where all users receive 256-QAM (or 64-QAM). As such, the plant and drop-home architecture are maintained to ensure that the worst conditions in the plant support 256-QAM. MMP is an optimization that allows operators to scale up their capacity when a subset of users on the network can support it, and dial down the modulation efficiency when necessary according to channel estimation. Note that the case of a single profile is merely a subset of MMP.

The expectation of MMP for DOCSIS® 3.1 is based on system analysis and field performance is that a small number of profiles (4-6) should sufficiently span the range of channel conditions, as shown in Figure 36 based on the data originally shown in Figure 6. The use of MMP in the time domain is shown in Figure 37.

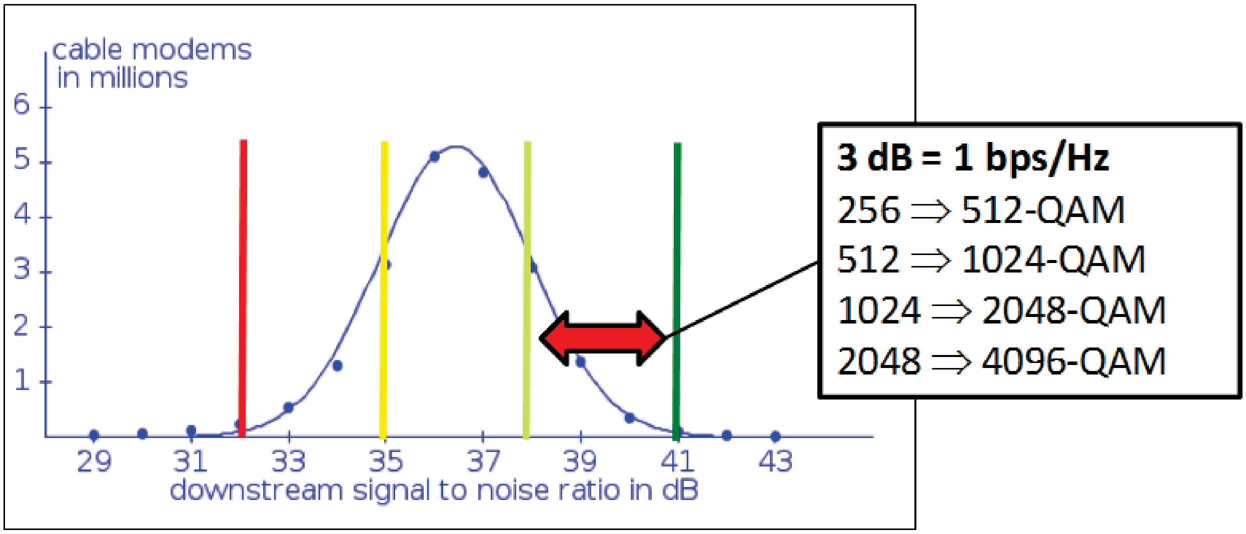

Figure 36 – Fielded Cable Modems Also Tell a Story Favorable to the use of Multiple Profiles in the Downstream [6][20]

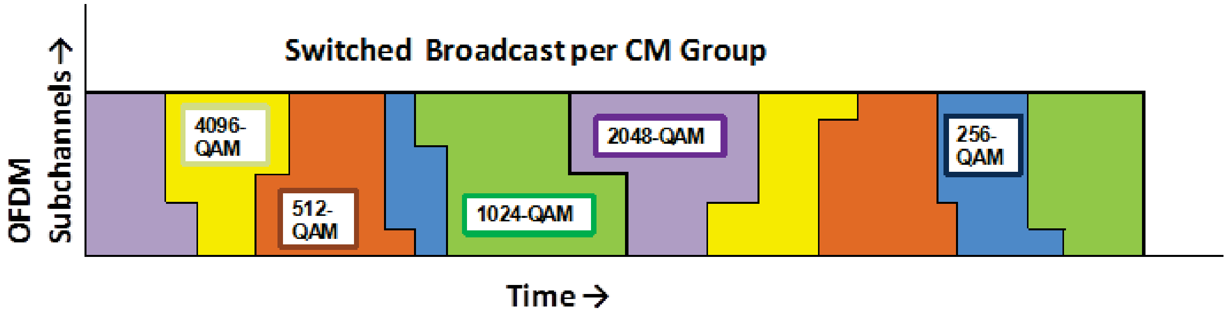

Figure 37 – "Switched" Broadcast Downstream Modulation Profiles Exploit the Range of Expected CM SNRs [6]

MMP functionality presents vendors with an opportunity to differentiate in the area of algorithms for profile selection, efficiency of the MMP mode, capacity optimization, channel estimation metrics, and balancing performance and configuration. It also creates a new set of configuration variables for operations teams to understand the value of this form of capacity optimization is to be fully leveraged.

HFC Readiness Recommendations

Many optimizations, suggestions, and recommendations were discussed across several disciplines in this paper. We try to capture them here along with a brief summary so that they are all listed in a single place, bucketed by the nature of the recommendation.

Many are already general practice recommendations of today, whereby the use of DOCSIS® 3.1 simply provides stronger reasoning to adhere to them. We are otherwise imposing more painful speed limits on the technology when compared to today's top speeds and capacity. While features of DOCSIS® 3.1 enable it to be used as an effective tool to overcome imperfections in the plant and maintain or improve service, the main motivator for DOCSIS® 3.1's new PHY tools and technologies was to increase the capacity of the network via the spectral efficiency in bps/Hz with little to no heavy lifting required on the plant itself. But since the maximum spectral efficiency is achieved on a nearly perfect plant, our goal is thus to maintain and minimally upgrade the plant so that the full capacity of DOCSIS® 3.1® can be achieved over time.

We use four categories, though clearly overlap exists and some concepts can fit in more than one bucket.

<u>Investment</u>

DFB Lasers or Digital Return – The return path will simply run into unnecessary speed limits with Fabry-Perot (FP) technology, even modern isolated FPs. Modern DFBs or better baseband digital return solutions are recommended even for multi-channel DOCSIS® 3.0, and will be a requirement to maximize upstream capacity in DOCSIS® 3.1

Remote Architectures – Robust performance of the highest order QAM formats should be possible with modern, high quality, well-maintained and aligned linear optics. However, a digital optical link potentially offers more dB of SNR, and thus possible capacity exploitation by removing the noise and distortion contribution of the linear optics. DOCSIS® 3.1 may go beyond 4096-QAM in the future, and digital optics may very well be a key enabler for this order of QAM.

Migration with a (Frequency) Plan – Business-As-Usual node splitting is effective but moving forward it must be done mindfully if it is to be efficient in preparing for DOCSIS® 3.1. These opportunities become the time to adapt technology to new frequency bands and architecture shifts. It is also opportunities for vendors to differentiate themselves by architecting maximum flexibility into the modular node architecture. Operators and vendors must collaborate to ensure the right optical platforms and module flexibility is in place for future directions.

Analog Loading – Operators are already well aware of the bandwidth reclamation benefits of removing analog video carriers. However, analog loading also generates discrete distortion components that will extract a price in capacity or margin. Removing analog also potentially frees up a small amount of RF loading into the digital band that can translate to positive dB on the QAM link budgets, and paves a much cleaner way for architecture convergence and use of WDM for deeper fiber.

Reducing the Number of Actives – As we strive to achieve the highest orders of modulation, we'll want to use the highest possible power levels in plant equipment and CPE subject to the amount of composite intermodulation noise (CIN) and the ingress/impulse margins required. Fewer actives means less CIN at the same total power level, so reducing the number of actives in the plant enables the higher order modulations we seek to obtain. Alternately, if new amplifier technology becomes available that permits higher power levels at the same or lower levels of intermodulation distortion, these components should be migrated into the network where needed most. OFDM technology in fact allows for peak-to-average power reduction schemes that may let actives to run at a higher operating point for the same total RF load. As OFDM takes over the HFC spectrum, these benefits could amount to significant new dB of available link performance. They could also be leveraged for power savings.

And, of course, fewer actives means less maintenance and smaller system MTBFs.

CCN and MER Targets – We saw evidence of essentially location and architecture dependent link performance in Figures 6 and 12. The use of MMP is an attempt to exploit this range of performance. However, using the "rising tide lifts all ships" philosophy, there is no substitute for improving HFC performance itself. This may involve incrementally improving performance, as we have seen occur with WDM over time, major shifts in architecture per the aforementioned digital optics approach, and/or more attention to response sweep (see metrics), maintenance periodicity, tap port levels, and drop/home quality.

Plant Equipment Above 1 GHz? – We saw the potential for manageable use of above 1 GHz under certain assumptions of 1 GHz equipment (and similar assumptions would apply for above 750 MHz or above 870 MHz for equipment specified to those limits). However, the extension to 1.7 GHz or thereabout is a new generation of outdoor plant gear. Vendors would need convincing to invest R&D in new housings for amplifiers and nodes without compelling evidence that this spectrum will be utilized. Even with the microwave physics in place, other obstacles (loading, new modules, CPE) exist that would require substantial new investment to exploit the band. Such bands may lend themselves to more complex craft issues as well (in the category of “Practices”).

Operations

Laser Alignment Levels – Better upstream technology is essential. However, it is wasted without proper attention to alignment of levels at the input to the laser or A/D converter. Under-loading robs the link of SNR and capacity, and overloading introduces accelerated compression and laser clipping, and leaves the link more vulnerable to external interference overload.

Limit/Eliminate HE Combining of Returns – Combining return paths haves the port costs, but incurs upstream SNR loss and multiplies funneling domains. If the port is not yet needed, then perhaps the lost capacity is not very significant.

Forward Alignment Levels – Traditional practices have run all 256-QAM or 64-QAM at the same relative level. When higher order QAM formats are introduced alongside legacy services, a different balance of levels might make more sense from a system standpoint, at the expense of increased operational complexity.

New Metrics – More information is being made available about network frequency response through the granularity of OFDM itself. Other important new information will come in the form of QAM constellation quality and signature, equalization coefficients, new and more detailed FEC statistics, waveform histograms, buffers for full-band capture of transients, full-band RF spectrum capture and NPR notch analysis. Given the adaptability of OFDM, it is also possible that an operational approach of periodically testing individual subcarriers with the highest order of modulation for brief intervals to determine the number of errors that occurs may also be possible. Quiet period configurations in real time can also provide a deeper view of what is going on in the network, such as nulling a single or group of subcarriers to perform an NPR type measurement. Tools to track, extract, and efficiently alert and inform will be necessary to avoid analysis paralysis.

Leakage Detection (plant & home) – Cable signal egress requirements already exists, so tools and processes to manage this are in place. The reciprocal problem of ingress has escalated for the downstream (it has been persistent in the upstream already), and cable and connector integrity in the plant and RF isolated headends in some cases will be required to manage its impacts. In the home, the problem is more complicated for two reasons – cabling is generally out of 24/7 access reach of the operator, and the shielding of current CE equipment is simply inadequate to provide the kind of isolation that may be required in situations, as a function of proximity (better shielded CE alone would actually fall under the "investment" category). One way in particular that OFDM can be used for leakage detection is to take advantage of the modulation-free carriers used as pilot tones, and ensure they are regularly placed at specific frequencies that are relatively free of existing off-air signals and detect them in field sweeps similar to how leakage detection is done today.

Fancy Test Gear Has Value – Modern MXA-series Vector Signal Analyzer (VSA) as shown herein are expensive but easy to operate, run remotely, and provide a tremendous depth of capture and analysis that can ensure getting to the bottom of the most perplexing performance issue.

Next Generation CMTS Capabilities – With next generation CMTS or converged cable access platform (CCAP) silicon having dedicated processing for the capture and analysis of the entire upstream with no impact on normal data signal processing, they will be able to capture and perform FFTs on the upstream with performance that could approach expensive VSA systems and be built into all networks, with centralized repositories of data for analysis by headend and network operations center (NOC) techs throughout the MSO's footprints.

Practices

Home Architecture – The reference architecture for DOCSIS® 3.1 is the point-of-entry home gateway, one splitter-deep ideally. This preferred install has benefits to receive levels, transmit power, ingress sensitivity, and eventually (with no splitter) access network abstraction. Over time, if the home network becomes primarily wireless and/or Ethernet based, a nearly complete isolation of the home network from the cable operator's network could eventually be achieved for the ultimate network integrity.

Grounding/Shielding Integrity – The universally best solution to reducing the ingress and impulse noise that tends to burden the upstream much more than the downstream is reducing it at the home via proper grounding practices and adequate shielding of CPE and home wiring. Various estimates put the ingress received upstream as 80% generated by RF and electro-mechanical sources inside of homes that push large current around sensitive RF equipment.

Of course, the home network is virtually uncontrollable to operators, although it is becoming more visible via TR-069 and other CPE and home network management solutions being deployed. Install practices might standardize around some basic shielding/leakage and grounding integrity tests. This offers the best opportunity to minimize radiated ingress entering the cable from the home (and escaping out) and conductive sources taking advantage of insufficient grounding. Techs of the future should be driving towards performing more spectrum analysis, either with field instrumentation or via new InGeNeOs-based methods to certify new home installs are not impaired by impulse/burst noise. Lots of information may be gleaned from just a minute or two or more of maximum hold spectrum traces.

For future truck rolls and new install practices, an area that is more accessible and important for ingress mitigation is the drop/home cabling and connection. Drop cable is of lower quality than trunk but exposed to the same outdoor or underground environment, so is prone to degradation and is at the mercy of the homeowner. The integrity of the connection into the home, which in the case of the home run may be several connectors at the side of the home, and the use of 75 Ω terminations where possible will help to gradually minimize interference sources getting onto the HFC plant if practiced methodically over time on house calls.

MoCA™ overload – In general, practices for the coexistence of MoCA™ and cable service are established, which typically includes filtering of the home and of CPE in the home that might have their RF front end overloaded.

MDU Care and Feeding – Apartment homes frequently bend the coaxial design rules, are prone to low taps and long cable runs, are difficult to ensure wiring integrity and the apartments themselves represent a dynamic termination environment depending on turnover and occupancy.

Multiple home networks – New home networking technologies have overlapping spectra to the cable upstream, such as HomePlug®, and as such could ingress into the cable network and disturb the return path (for everyone). Satellite signals have and will continue to pose a concern for the possibility that different coaxial networks are unwittingly tied together.

Training

The list of training topics for field staff should be quite obvious by now. It includes, but is not limited to:

OFDM and OFDMA; differences and similarities to SC-QAM
Higher-Order Modulation
Digital Optics and Remote Architectures
Multiple Modulation Profiles vs. Broadcast and Configuration of Downstream Users
Upstream Laser Loading (existing but worth a re-visit)
Impairment Locating through Diagnostics

CPE Architecture/Install for Home (POE) and MDU (Rules & Guidelines)
New Pro-Active Metrics, Diagnostic Tools, Analysis Algorithms

MER per-QAM symbol
Constellation signatures
NPR notches
Burst characterization
Use of FEC statistics (especially the new ones from the BCH-LDPC coding used in DOCSIS® 3.1)
Equalization, Pre-EQ, coefficient and channel estimation analysis
Full band capture analysis
LTE signatures
Burst noise signatures

Conclusion

Cable's HFC downstream channel represents perhaps the world's most capable RF channel in terms of capacity. With DOCSIS® 3.1, the industry will take on exploiting it to its maximum potential. To do so, new technology tools are being leveraged, and different architecture and spectrum possibilities are being considered. These were described herein and in some examples we quantified the benefits they could provide.

However, technology and system design can only take us so far down the path of optimization. Operations and best practices must be aligned with the new objectives, which place more pressure on the network to be robust and well-behaved. Ensuring this is the case itself is a daunting problem statement that has to do with rethinking various aspects of personnel, training, metrics, infrastructure, and maintenance and installation guidelines, rules, and processes. In this paper, we touched on some of the key components for operations to think about in order to make sure that DOCSIS® 3.1 delivers all it is meant to.

Acknowledgements

The authors would like to thank John Ulm, Ayham Al-Banna, and Doc Cummings of ARRIS for their valuable contributions to this paper. Key contributors also include Bruce Currivan (Broadcom), Leo Montreuil (Broadcom), Alberto Campos (CableLabs), Tom Williams (CableLabs), John Chapman (Cisco), Tim Brophy (Cisco), Saif Rahman (Comcast), Joe Solomon (Comcast), and Zan Zhang (Maxlinear).

Bibliography

[1] Chapman, John, Mike Emmendorfer, Dr. Robert Howald, Shaul Shulman, *Mission Is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS*, 2012 Cable Show, Boston, MA, May 23-25.

[2] Cloonan, Tom and Sam Chernak, "HFC Spectrum Planning in the Era of the HSD Bandwidth Explosion," 2012 Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[3] Couch, Ken, "Ingress and Egress Impairments," Draft Contribution to SCTE Special Working Group on HFC Readiness," June 28, 2013.

[4] Garavaglia, Andrea, Robert Howald, and Hesham El-Bakoury, *Adaptive Bt Loading Definition,* Contribution to EPoC Special Interest Working Group, Nov 2012.

[5] Howald, Dr. Robert L, *Boundaries of Consumption for the Infinite Content World*, SCTE Cable-Tec Expo, New Orleans, LA, October 20-22, 2010.

[6] Howald, Dr. Robert, *Breathing New Lifespan into HFC: Tools, Techniques, and Optimizations,* 2013 Cable Show Spring Technical Forum, Washington, DC, June 10-12.

[7] Howald, Dr. Robert L, *Fueling the Coaxial Last Mile*, SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[8] Howald, Dr. Robert L, *Looking to the Future: Service Growth, HFC Capacity, and Network Migration*, 2011 Cable-Tec Expo Capacity Management Seminar, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Atlanta, Ga, November 14, 2011.

[9] Howald, Dr. Robert, *The Grown-Up Potential of a Teenage PHY,* 2012 Cable Show Spring Technical Forum, Boston, MA, May 21-23.

[10] Howald, Dr. Robert L, *Time on Our Side: Using the Past to Pave the Way Forward*, 2012 Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[11] Howald, Dr. Robert L, Dr. Sean McCarthy, *Accounting for Techies: Taking it to the Ultra*, The Cable Show Spring Technical Forum, June 10-12, Washington, DC.

[12] Howald, Dr. Robert L, Dr. Sean McCarthy, *Bits, Big Screens and Biology*, The Cable Show Spring Technical Forum, May 20-22, Boston, MA.

[13] Howald, Dr. Robert, Michael Aviles, and Amarildo Vieira, *New Megabits, Same Megahertz: Plant Evolution Dividends*, 2009 Cable Show, Washington, DC, March 30-April 1.

[14] Howald, Dr. Robert L. and John Ulm, *Delivering Media Mania: HFC Evolution Planning,* 2012 SCTE Canadian Summitt, March 27-28, Toronto, ON, Canada.

[15] Pietsch, Christian, OFDM Numerology, Contribution to IEEE 802.3bn (EPoC) Working Group, January 23-25, 2013, Phoenix, Az. http://www.ieee802.org/3/bn/public/jan13/pietsch_01_0113.pdf

[16] Ulm, John and Gerry White, *Architecture & Migration Strategies for Multi-screen IP Video Delivery*, 2012 SCTE Canadian Summit, March 27-28, Toronto, CA.

[17] http://maninair.in/orthogonal-frequency-division-multiple-access-ofdma

[18] http://www.dipol.ie/newsletter2/inf_dipo_2009_32.html

[19] http://www.ict-redesign.eu/fileadmin/documents/ReDeSign-D08_Report_on_Cable_Channel_Model.pdf

[20] http://www.ieee802.org/3/bn/public/adhoc_mmp/adhoc_mmp_urban_01.pdf

[21] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf

[22] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_02_0313.pdf

[23] http://www.ieee802.org/3/bn/public/nov12/schmitt_01a_1112.pdf

[24] http://www.lightwaveonline.com/articles/print/volume-26/issue-9/industry/ofdm-to-power-high-bit-rates-in-next-gen-optical-access-networks.html.

Abbreviations and Acronyms

A-TDMA	Advanced Time Division Multiple Access
AWGN	Additive White Gaussian Noise
CAGR	Compound Annual Growth Rate
CCN	Carrier to Composite Noise
CE	Consumer Electronics
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premises Equipment
DFB	Distributed Feedback Laser Technology
DFE	Decision Feedback Equalizer
EPoC	Ethernet PON over Coax
EPON	Ethernet Passive Optical Network
EQAM	Edge QAM
FEC	Forward Error Correction
FP	Fabry-Perot Laser Technology
GDV	Group Delay Variation
HSD	High Speed Data
IFFT	Inverse Fast Fourier Transform
IMD	Intermodulation Distortion
IPV	Internet Protocol Video
ISI	Intersymbol Interference
LDPC	Low Density Parity Check Codes
LE	Linear Equalizer
LTE	Long term Evolution
MoCA™	Multimedia Over Coax Alliance
NPR	Noise Power Ratio
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
POE	Point of Entry
QoE	Quality of Experience
S-CDMA	Synchronous Code Division Multiple Access
SC-FDMA	Single Carrier Frequency Division Multiple Access
SC-QAM	Single Carrier QAM
SIR	Signal-to-Interference Ratio
SDV	Switched Digital Video
SNR	Signal-to-Noise Ratio
VOD	Video on Demand