Characterizing and Aligning the HFC Return Path for Successful DOCSIS 3.0 Rollouts

SCTE Cable-Tec Expo®
Denver, Colorado
October 28-30, 2009

Dr. Robert L. Howald
CTO Office
Rob Thompson
Systems Engineer
Dean Stoneback
Senior Manager, Access Systems Engineering
Vipul Rathod
Manager, Access Systems Integration & Test
Charles Moore
HFC Services

Motorola Home & Networks Mobility

Phillip Chang
Principle Engineer

Comcast Network Engineering & Technical Operations

I. Introduction - Upstream Capacity Evolution

HFC has evolved since inception, with periodic heavy investment in access network "upgrade" cycles taking place roughly every five years. These upgrades have generally been associated with creating new downstream spectrum to add more channels, digital channels, data services, voice services, VOD, high-definition services…..the list goes on and will continue to. Along the way, in addition to adding physical spectrum in the form of coaxial plant bandwidth, virtual bandwidth was also created by splitting nodes to create smaller service groups. Doing so provides more average bandwidth per home, which, combined with investment at the Headend to add more narrowcast QAMs to take advantage of the increasing number of service groups, increases the amount of services each subscriber can simultaneously consume.

Unfortunately, the act of adding coaxial bandwidth by increasing the useable physical spectrum has no benefit to upstream service bandwidth. All prior and current upgrade cycles – 550 MHz, 750 MHz, 870 MHz, 1 GHz, etc., add downstream bandwidth only, because the diplexer split of upstream and downstream is at 42 MHz (North America) – higher in other parts of the world, but still fixed at the low end of the coaxial spectrum. However, the act of *virtual* bandwidth indeed has a positive effect on the upstream, with added benefits when accomplished through physical node splits (as opposed to node segmentation within the node). While a secondary effect of downstream segmentation, this has been the primary mechanism to increasing upstream bandwidth per home to date.

The tools being used for downstream service expansion offer tremendous new potential capacity [1]. There is not a similar available, convenient toolset for the upstream to achieving capacity improvement on the scale of the downstream. As a result, the upstream, limited to its roughly 4-5% of the coaxial bandwidth, is vulnerable to Moore's law-type growth rates without relatively major steps being taken to combat this asymmetry. The 4-5 % is, in fact, a positive spin on the situation. Most MSOs avoid the spectrum below 15 MHz for HSD services, because the channel properties in this region are highly variable and can be problematic. As a result, around *3% or less* characterizes many plants upstream/downstream ratio. Furthermore, the downstream has the advantage in channel quality, so the capacity asymmetry is more staggering. As we will discuss, one very important upstream, S-CDMA provides the incremental tool to utilize the full spectrum, including the low end of the band. This will allow MSOs to buy time to plan a next move, which is an inevitable one, in order to enable the throughput growth that will be necessary as more upstream bandwidth is consumed. As will be seen, use of S-CDMA offers nearly 50% more return capacity by operating where A-TDMA simply is unable to, or, where A-TDMA can operate, offering the ability to deliver a higher bit rate than A-TDMA on the same spectrum, using a higher modulation profile with less overhead and fewer errors.

There is another important upstream tool available to deploy with relative convenience – namely DOCSIS 2.0's (and 3.0's) 64-QAM modulation mode. Use of 64-QAM allows, for example, a 10 Mbps 16-QAM upstream to become a 15 Mbps upstream. Or, for DOCSIS 3.0, it allows a 16-QAM, 20 Mbps upstream channel to become a 30 Mbps channel. Obviously, in both cases, this represents a 50% throughput improvement. Considering the combination of S-CDMA and DOCSIS 3.0, about twice the capacity of the typical upstream of today can be achieved within the given allocated spectrum. However, DOCSIS 3.0 using 64-QAM requires some insight into the upstream channel to be deployed successfully. Note that, while 64-QAM mode delivers new capacity, DOCSIS 3.0 upstream channel bonding does not, in principle, deliver any new *capacity*, and thus does nothing to address the data rate growth issue. The use of more *spectrum* certainly adds capacity. As such, if bonding is deployed by putting a new DOCSIS carrier adjacent to an existing one, that system has added capacity. However, the added capacity is a result of deploying the additional channel, which could occur with or without channel bonding. Upstream bonding provides an approach to delivering higher *tier* service rates, although the performance value of such tiers given other network throughput constraints is probably less important than the marketing value of a higher Mbps number.

While peer-to-peer traffic first drew attention to upstream beyond basic web-browsing, the evolving trend toward video social networking and user generated content – real-time service – will provide the next spotlight on the upstream. Today, the expectation of poor quality, low-resolution video (e.g. YouTube) offers cover for unpredictable upstream streaming capability. However, this too will change and drive a need for enhanced upstream services and speeds. The question is when this occurs, so that operators can determine at what point it is a priority to invest in upstream solutions. Table 1 offers some guidance on this important question.

Table 1 – Lifespan of the HFC Upstream

When Do You Run Out of Upstream @ 2x per 18 months Increase						
Homes Passed *HHP/Node*	**Penetration** *Subs/HHP*	**Concurrency** *Simultaneous Subs*	**Start pt Mbps** *2009*	**Upstream of 100 Mbps** *yrs supported*	**Add S-CDMA - 150 Mbps** *yrs supported*	**5-65 MHz - 300 Mbps** *yrs supported*
500	60%	1%	2/5/20	6.0/4.0/1.0	7.0/5.0/2.0	8.0/6.5/3.5
		2%		4.5/2.5/will not	5.5/3.5/0.5	7.0/5.0/2.0
300	60%	1%	2/5/20	7.0/5.0/2.0	8.0/6.0/3.0	9.5/7.5/4.5
		2%		5.5/3.5/0.5	6.5/4.5/1.5	8.0/6.0/3.0
180	60%	1%	2/5/20	8.0/6.0/3.0	9.0/7.0/4.0	10.0/8.5/5.5
		2%		6.5/4.5/1.5	7.5/5.5/2.5	9.0/7.0/4.0

Table 1 offers some keen insight into when operators can expect an upstream bottleneck that requires action be taken. Operators face a balancing act of upstream initiatives between major capital intensive capacity additions, incremental capacity additions such as DOCSIS 3.0 64-QAM and S-CDMA, and DOCSIS 3.0 upstream channel bonding to deliver higher service tiers.

The way to interpret Table 1 is as follows:

1) HHP/Node – Household passed per node – the commonly used plant segmentation metric
2) Penetration – Also the commonly used metric referring to the number of data (DOCSIS) services as a percentage of HHP (60% is aggressive, but it is a desirable growth objective and a system design boundary should such a growth "problem" occur)
3) Concurrency – Simultaneous use percentage or reciprocal of oversubscription; how many users can actually use the tier rate offered at the same time without blocking
4) Start pt Mbps – Key swing point of chart: What is the currently assumed required tier that is or must be offered today to be competitive? This will serve as the starting point of growth calculations that follow to determine lifespan. It is an operators individual decision based on the competitive market and positioning whether to lean towards a demand-pull based number here, or a market-push based number (such as in FTTH

markets). Regardless, what is assumed as necessary today for tier offerings forms the basis for accumulating average growth in subsequent years using a Moore's/Nielsen's Law basis – in this case an aggressive 2x every 18 months. The column references three data points to cover the range of considerations: 2 Mbps/5 Mbps/20 Mbps.

5) What years of life remain under the prior column's plant and data rate assumptions when the upstream can support 100 Mbps of total throughput using DOCSIS 3.0 ATDMA
6) What years of life remain under the prior column's plant and data rate assumptions when the upstream can support 150 Mbps of total throughput using DOCSIS 3.0 ATDMA and including use of the added capacity using S-CDMA
7) What years of life remain under the prior column's plant and data rate assumptions with the European split of 65 MHz

In the table, red numbers highlight conditions that say action must be taken within the next two years or less, as this essentially means planning steps should be ongoing or taken immediately.

A few highlights from the table:

- The power of the node split is clear, as would be expected. It is important to recognize that node splits do more than just add per-home BW through service group splitting. Node splits also create cleaner channels with every step, offering the opportunity for more effective use of higher order modulations both upstream and downstream [3].

- S-CDMA provides a valuable time-buyer in planning and deferring node splits, at essentially no new costs other than the learning curve of turning on this upstream mode.

- Similar to S-CDMA, the higher split "buys time," but it does so in a much more intrusive and expensive way where 65 MHz or greater is not already the standard. An effort to increase the split should consider a more potent long-term step given the nature of that investment and the overall asymmetry.

- The concurrency factor is important, and will be even more important going forward. As social networking video (real-time) services become prevalent, the ability to continue to overbook the channel as the data rate increases may be lessened, particularly if these services and other user-generated content (UGC) begin to take on higher quality video characteristics.

- The decidedly linear expansion possible via node splits, S-CDMA, and higher return splits, still give way in time to the overall power law growth rate of high-speed data services. Relative to the downstream, this represents a significant bottleneck down the road, and within the planning horizon in many cases.

Clearly, this table can be sliced and diced in many ways under varying assumptions to create the larger picture of what to do with the upstream, longer term. For the requirements right around the corner, DOCSIS 3.0 provides the means to further empower the upstream via 64-QAM, upstream bonding to increase tier rates, and enabling unused spectrum via S-CDMA. The DOCSIS 3.0 suite of enhancements provides a powerful set of tools for cable operators to move their high-speed data business forward. Among the most notable features is the ability to bond channels together to offer higher rate service tiers in both the upstream and the downstream directions. The potential for a higher upstream tier comes on top of DOCSIS 2.0 modifications that enabled wider band, higher-order upstream modulation (64-QAM @ 5.12 Msps). DOCSIS 2.0 also introduced S-CDMA, which saw additional enhancements in DOCSIS 3.0. With years of successful DOCSIS services behind them, most operators are either beginning to deploy, planning deployments, or considering the next steps to turning on the new DOCSIS features that empower the HSD service to greater heights of customer satisfaction.

The rest of this paper focuses on aspects of optimizing use of the upstream to take advantage of these new features. In particular, we will focus on the most advanced modulation profiles, and the ability of the upstream to support them. Given the limited spectrum available, particularly in North America, the ability to most effectively mine the available bandwidth and maximize throughput will be critical as the natural rate of speed requirements continue and market pressures continue to increase. Unlike the downstream, the upstream has a wide array of potential obstacles and impairments, creating one of the most challenging digital communication channels to manage and fully exploit. However, it nonetheless has a relatively high average SNR in most circumstances, although the SNR and its related, non-equivalent cousin Modulation Error Ratio (MER) itself can be a tricky variable. The challenge is how to ensure that the capacity associated with the upstream channel, or as much of it as possible, gets realized for services and revenue. To do so requires a thorough understanding of a diverse set of HFC and digital communications variables. More importantly, variables that did not matter very much for 16-QAM operation now become not just relevant, but critical to understand for successful deployment of 64-QAM, and to a lesser extent, 32-QAM.

II. Upgrading the Upstream

A. The 64-QAM Effect

In the Beginning

As a short history lesson for context, DOCSIS 1.0 provided a robust means to get cable modem service up and running in the dynamic environment that is the return path. It incorporated very rugged modulation profiles over a range of potential data rates, in an effort to accommodate poor upstream channels, whether due to RF limitations or older optical technologies, and to enable more cable spectrum with the tools to deliver higher speeds. While pre-DOCSIS upstream cable services could in some cases be considered near real-time, aside from early telephony systems, significant latency was tolerable. In addition, the kind of rates required to execute transaction-oriented and polling-type of services was very low pre-DOCSIS, harkening to dial-up speeds before they got “fast.” The only requirement was to close the link reliably, with little regard to efficiency or speed. As such, proprietary upstream modems were based on very robust, but not bandwidth efficient modulations, such as QPSK or FSK. And, the transmitter and receivers spoke using very simplistic and inefficient protocols, such as Aloha – essentially a single-wire free-for-all that relied on the lack of persistent traffic to provide adequate performance and very low cost. Because little attention needed to be paid to the return path characteristics for these early services, and its future was far from obvious when these services were first brought online, the most rugged modulations running at very low rates – and thus narrowband from a noise receiving standpoint – provided a reliable, practical, and low-cost platform for the enhance video services the upstream enabled. The use of frequencies in some cases that would later be observed to be not typically very healthy return path spectrum for high-speed data perhaps further emphasizes how little attention was paid to the upstream and its dynamic characteristics.

DOCSIS 1.0 provided a much more sophisticated TDMA multiple-access protocol to service the real-time browsing experience and efficiently use spectrum, with an eye towards where cable modem service might go in the future – which clearly was faster and faster, with only how fast it would accelerate in question. In addition to QSPK symbol rates jumping in octaves from 160 ksps (320 kbps) to 2560 ksps (5.12 Mbps), the DOCSIS 1.0 profiles included a 16-QAM upstream, doubling the rate possible at each symbol rate increment over what QPSK offered, with the highest symbol rate at 16-QAM offering a 10 Mbps service. To further harden these already robust modulations, DOCSIS incorporated powerful Reed-Solomon (R-S) forward error correction (FEC), chosen in part also because of its burst error correction capability, such as that associated with the prevalent impulsive noise on the upstream. The combination of frequency agility, modulation choice, symbol rate choice, and burst error correction provided a sound approach to delivering an upstream service robust enough for high-speed (relatively speaking) web surfing interactivity.

The implementation of 16-QAM compared to QPSK comes at the price of about 7 dB of SNR, ignoring coding gain aspects. This is, of course, not an insignificant number, but in the context of what is possible on a cable upstream with modest performance, can still leave quite comfortable link margin. This is a very important point to recognize. ***Link margin and its implications is a critical concept to understand well in order to successfully deploy 64-QAM***. As we will see below, operators will not be able to look at link margin the same way when deploying 64-QAM, simply because its inherent requirements in terms of channel SNR have crept into the ballpark of what some optical links are delivering today. Pile other RF impairments on top of this reality, and once comforting dBs of margin has shrunk considerably. In cases where margin was not so comfortable, it may have evaporated entirely. There are high performance reverse links and improvements possible – higher power and lower noise lasers, digital return links, for example. However, we will focus on what is common and in the field today as DOCSIS 3.0 is being rolled out.

Upstream Link SNR

Let's put some of the above discussion in the context of some quantified parameters. Refer to Figure 1 below. It shows the calculated optical link characteristics of various laser technologies at a link length of 25 km. These curves represent the "worst case" scenario – which generally means that the most challenging environmental conditions are imposed on each end of the link, and specifications are at the limits of their respective requirements set on both ends. In short, it is the type of information that system designers need to assess their worst margin conditions.

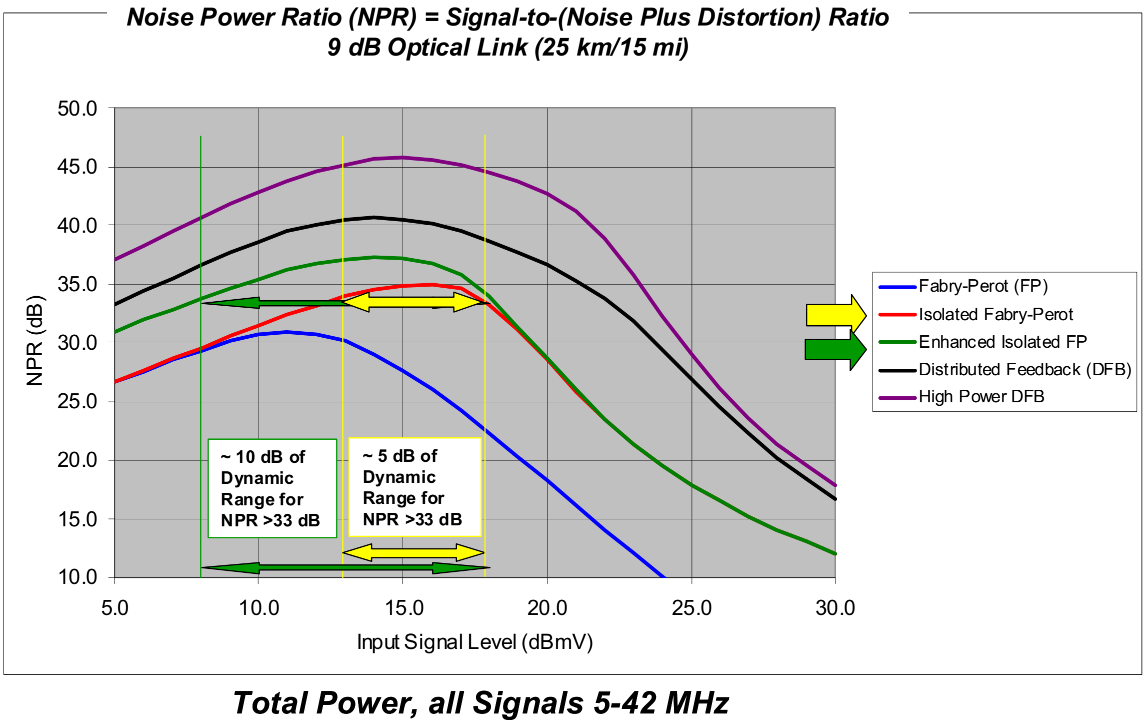

Figure 1 – Noise Power Ratio (NPR) Performance for Various Laser Types

The "Noise Power Ratio" technique offers a glimpse at the SNR provided by the optical link only across the ***entire*** return path spectrum, were that spectrum filled with QAM carriers with uniform loading. It represents the range of per-channel SNRs as a function of laser drive, using a "power-per-Hertz" loading methodology, which will be discussed in a subsequent section.

Refer to Figure 2. The left side of the NPR curve represents the linear operating range, while the right side represents distortion-induced degradation. For a deeper discussion of the NPR curve, see [2] and [7]. Figure 2 is taken from reference [7] and shows the various regions of the NPR curve. It is important to understand that the left hand side of the curve is truly *signal-to-noise ratio* (SNR), where the noise is the multiple optical contributions that deliver an additive Gaussian White Noise (AWGN) characteristic. It is also very important to recognize that w***hat a CMTS reports as SNR is, in fact, Modulation Error Ratio, or MER***. The difference between SNR and MER is that MER is determined by comparing the demodulated constellation to the ideal version, thus capturing all impairments along the way – AWGN, distortions, impulsive noise, ingress, etc. Thus, the noise or SNR region of Figure 2 and the value reported by the CMTS are not one in the same.

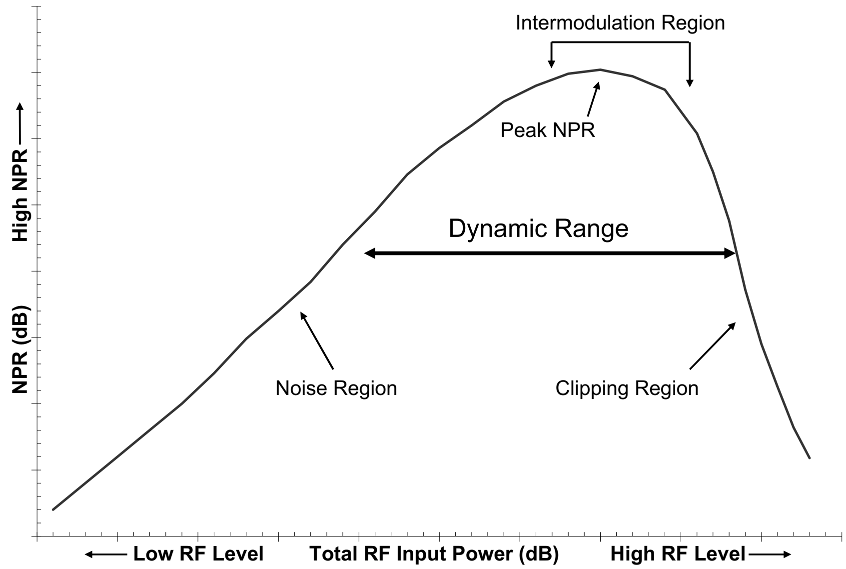

Figure 2 – Noise Power Ratio (NPR) Dynamic Range Example

It is also very important to understand as we transition to 64-QAM, that NPR (and SNR from it) and MER are *averaging* measurements – not insignificant considering the bursty-noise nature of the upstream. MER is always lower than the actual SNR, because it includes the AWGM *plus* the rest of the imperfections. Various techniques are employed to mitigate these additional impairments and allow as close to ideal symbol sequence detection as possible. These include that R-S FEC, narrowband ingress incision techniques, and the powerful tapped delay line equalizer to correct for frequency response distortions and repair reflections that deliver multiple symbol copies to the receiver. The CMTS ***Equalized*** MER (EQ-MER) means it is calculating MER after it has done its work to fix these frequency response distortions. It's what the symbol detection has to work with to take the information off of the channel. By reporting EQ-MER, the CMTS is also in effect describing how well it handled the combination of impairments. What errors created that it cannot correct for is left to the R-S decoder to try and sort out.

Moving from south to north on the NPR plot covers the range of lasers typically deployed today, from older, non-isolated, low power Fabry-Perot (FP) lasers, through to deployed Distributed Feedback (DFB) lasers. Ten-bit digital return links look very much like the best DFBs, with the added advantage, of course, of having link-length independence, in contrast to analog lasers. The parallel to DFB performance is not the case for all digital return links, however, as the performance is dependent on the number of bits used to capture and reconstruct the spectrum, and possibly also a function of digital signal processing applied prior to reconstruction.

As can be seen in Figure 1, the NPR peak varies nearly 15 dB across the set of lasers shown. This is not to infer that the peak is the proper operating point, it is simply the most easily referenced point. More important than the peak is the range over which a "good" NPR can be held, as plant variations require a "window" over which a link performance can be assured.

QAM Thresholds, NPR, and Link Margin

What is evident from Figure 1 is that the lowest performing laser delivers an NPR peak just above 30 dB. With some margin of signal level back-off from the peak in order to operate in the linear region and have some protection against the elements, there is still, for an example of 3 dB headroom margin, and about 28 dB of SNR. While older FP lasers themselves create spurious emissions which can look impulsive in nature, this 28 dB reference alone explains the ruggedness with which DOCSIS 1.0 was able to be rolled out and subsequently allowed to continually increase its rate to keep pace with consumer demand for speed. Without even considering coding, in place to mitigate harmful burst noise more as much as to provide AWGN advantage, QPSK requires about a 15 dB SNR to deliver 1e-8 BER, meaning there is in excess of 13 dB of link margin with the lowest performing optical link shown in Figure 1. It also means that, for 16-QAM, which delivers 1e-8 without FEC at about 22 dB SNR, there is in excess of 13 dB – 7 dB = 6 dB of link margin. This is not a particularly high margin, but this is again referenced to the lowest performing link. This example also readily shows the major impact modulation profile has on the available margin. If we consider 4-way node combining, we have chewed through nearly all available margin for 16-QAM on this lowest performing laser, and we have allocated no link budget for RF noise or other possible disturbances. It is clear to see why the older FP lasers are generally not well-suited to DOCSIS moving forward.

From Figure 1, considering the standard DFB curve, we have added 10 dB of *new* margin, all else the same, and suddenly we have a situation with 16-QAM nearly analogous to the prior QPSK case – excess link margin of greater than 10 dB. A way to spend that excess margin is by combining nodes to save upstream ports if the average bandwidth per user remains adequate. Again, a 4-way combine would then degrade this margin to 4 dB plus coding gain. In that scenario, once again the battle is on for wiggle room, variations in the plant, alignment errors, and new growth. Because changing lasers is a big deal, an existing standard DFB is probably not very likely to become a high power DFB soon, whereas any older non-isolated or low power FPs remaining may have the opportunity to be upgraded to this better tier of DFBs or a digital solution. In the case where the FP is a newer model, however, performance margin can be adequate, as indicated by the green arrows in Figure 1, where 10 dB of margin to a 33 dB MER exists, assuming no subsequent RF combining.

In the case described above, it is apparent that effective link gain is achieved by removing node combining, which is a natural consequence anyway of increasing the average bandwidth per concurrent user to support the continued need to increase the service tiers. Removing node combining for the same modulation profile is a logical step to take to increasing capacity prior to turning up a new, more advanced modulation profile. There is a drawback from a marketing perspective, which is that 16-QAM limits the peak rate to the 10 Mbps, or for DOCSIS 2.0 and 3.0, 20 Mbps. This may not be as competitive as required across all footprints.

Let's assume that we have removed node combining, and once again have more than 10 dB of margin using 16-QAM and a standard DFB, and we now want to increase this channel to support the maximum DOCSIS 2.0 profile of 64-QAM @ 5.12 Msps, or about 30 Mbps. Each increment of 2^n-based M-QAM yielding a square constellation incurs a 6 dB penalty. Thus, 16-QAM to

64-QAM requires yet 6 dB *more* SNR, while 64-QAM to 256-QAM would require still another 6 dB, as it did in the downstream when the conversion was made.

For 1e-8 BER, and not including coding gain, we then have the following reference points for theoretical SNR, and margin to a hypothetically (but practical) 33 dB:

Table 2 – Theoretical SNRs @ 1e-8 BER, No FEC, Margin @ 33 dB

Modulation	SNR	Margin
QPSK	15 dB	18 dB
16-QAM	22 dB	11 dB
32-QAM	25 dB	8 dB
64-QAM	28 dB	5 dB
256-QAM	34 dB	------

The 64-QAM profile is a DOCSIS 2.0 and DOCSIS 3.0 requirement. Focusing on this, it is easy to see in a more profound way by observing the margin impact of the instant loss of 6 dB moving from 16-QAM to 64-QAM. It is, in fact, a *minimum* of 6 dB, because that is the difference for *AWGN-only* given by *SNR*, and there is increased sensitivity of 64-QAM to other impairments compared to 16-QAM as well. The actual SNR-only can only be better than 33 dB, and the degradations that move the observed MER downward to 33 dB will be more deleterious to higher order modulations than the 6 dB AWGN relationship. We count on the mitigation tools in place to minimize their effect in practice, such that this 6 dB difference holds. We will see later that the advanced equalization in DOCSIS 2.0 and 3.0 allows the two modulations (16-QAM and 64-QAM) to be dealt with within 2 dB of performance of one another against micro-reflection impairments, for example. Another way to look at Table 1 is to recognize that a 33 dB (reported by the CMTS as SNR) is an MER representative of a slightly higher actual link SNR. This does not, however, mean it is "better" than it looks – in fact, the opposite. The combined impairment scenario leading to an MER of 33 dB is more complex to deal with than a 33 dB AWGN-only environment.

A commonly asked question is "can I run 64-QAM upstream over my FP link." Of course, the answer is "it depends." More specifically, as Figure 1 shows, the SNR required to deliver a 28 dB is available, actually for all cases. It becomes a question of ***margin,*** which we wil raise repeatedly. As mentioned, operators will likely have to become accustomed to looking at link margin differently with the deployment of 64-QAM. Moving to 64-QAM simply means that the theoretical requirements for the modulation profile bump up against the SNR that can be delivered over the optical link, especially when imposed upon by the set of other upstream impairments. This is especially true for FP links.

Now consider the yellow arrows of Figure 1. The yellow arrow shows that, on a link using an Isolated FP at a typical link length, a 33 dB NPR value, which is 5 dB of margin above the 28 dB reference in Table 1, has about 5 dB of dynamic range - not a particularly large range. The enhanced FP in this case adds only a couple dB more *peak* SNR, but more importantly roughly doubles dynamic range to about 10 dB above this 33 dB threshold, which is a more comfortable situation.

Summarizing the QAM-NPR situation, then, whereas DOCSIS 1.0 modulation profiles demanded just 15 dB or 22 dB or SNR, with 64-QAM we are taking steps to truly use the channel closer to its capability. The result is to eat into what was essentially wasted link margin in the early DOCSIS days, and spend some of that margin on new capacity. It becomes clear when considering the above why DFB lasers become valuable for successful 64-QAM deployments, in particular when there is the potential that node combining may occur. As previously described, such combining can consume significant margin, lost through the noise addition. Also, node combining brings another factor into play that can often be ignored – contributions of upstream RF amplifiers. As nodes have been split and cascades shortened, these contributions continue to shrink. However, when cascades are long and nodes get combined, the number of upstream amplifiers accumulates, and can be enough to create a non-negligible noise contribution. The effect would be noticed first on Figure 1 by the pushing down NPR peaks through "soft" distortion mechanism of RF amplifiers (the "Intermodulation region" of Figure 2), but not to effect the dynamic range over which a given SNR threshold is set. For the vast majority of practical cases and in particular in places where DOCSIS 2.0 and DOCSIS 3.0 will be implemented, this is the only effect noticeable on the channel's baseline noise performance. However, when the situation also translates into more homes on a link, the likelihood of link degradation due to ingress and impulse generated by the homes served increases, and thus so does the probability of a difficult channel.

Now, it is important to consider a further system design element, in particular with respect to the Isolated Fabry-Perot. While there is a clear observable impact as various laser types are compared to one another, link quality degradation as a function of link *length* is more subtle, at least for small variations from typical. Consider Figure 3. The Isolated Fabry-Perot transmitter (IFPT) was deliberately chosen, as we have noted how these higher quality FPs basically sit on the border of "acceptable" for 64-QAM, making the secondary conditions more important – RF plant quality, laser power loading, optical receiver settings, and link length, shown here. This particular case is chosen because we note that by going from a 10 dB link to a 15 dB link, about 6 dB is lost on the left hand (SNR) side of the NPR curve. As discussed, 6 dB represents the difference between one square M-QAM constellation and the next, for $M = 2^n$. Thus, link length could contribute to the difference between having comfortable margin or not for 64-QAM links. Based on this simple comparison, note that alternative methods to power-per-Hz loading may be useful to consider on lower performing optical links. This comes, of course, at the expense of complexity. Power loading algorithms that are "optimized" in a raw BER sense have been developed in anticipation of such cases [4].

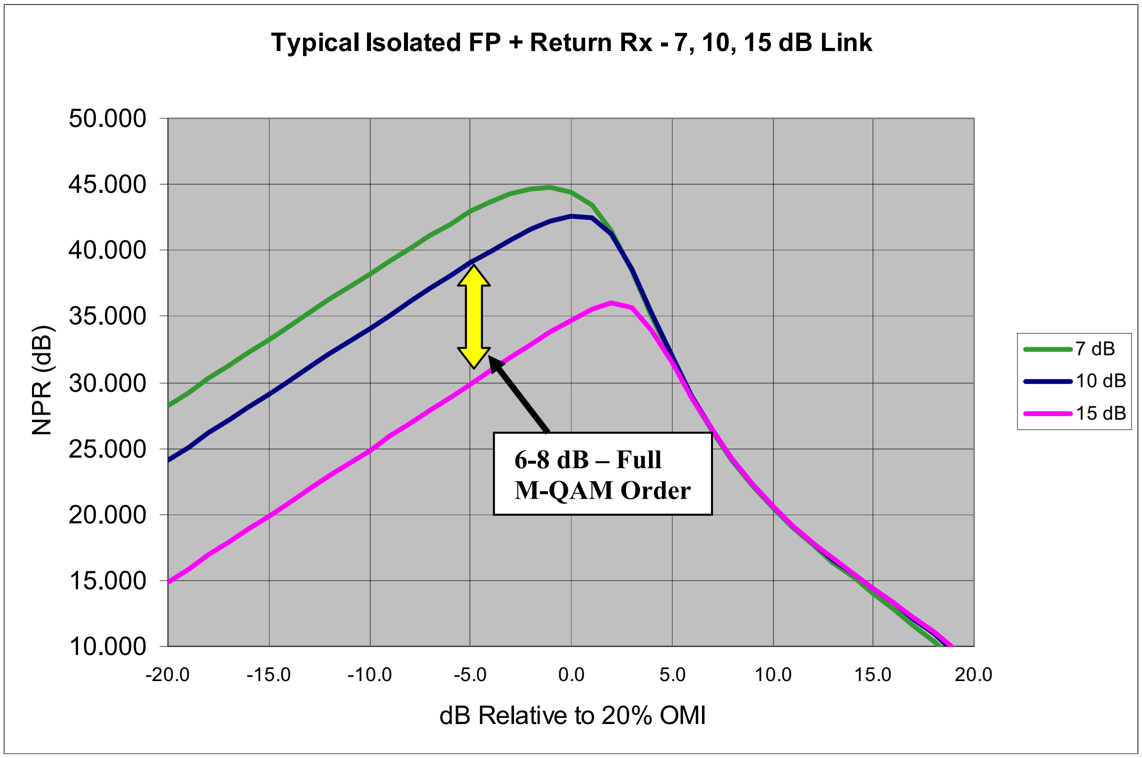

Figure 3 – Link Length Dependence of NPR and QAM Margin Sensitivity

While the DFB adds the additional overhead to the optical link to provide a comfortable level of margin for 64-QAM from an SNR perspective, as previously described, 64-QAM is increasingly sensitive to ingress, impulsive noise, and distortion as well. To compensate, DOCSIS has also evolved to handle ingress and frequency response distortion better as well. Notably absent from that list is changes to provide impulse noise enhancement. In fact, there are two further elements of DOCSIS meant to deal with this inevitable component of the return path – interleaving and S-CDMA. These will be discussed in subsequent sections as they come into play when discussing R-S functionality and FEC statistics, and when we address spectral planning and usage.

It is worthwhile to point out during this continued dive into the needs of 64-QAM that an intermediate step using 32-QAM is a logical approach. QAM relationships being non-linear as they are, the difference in stepping to 32-QAM from 16-QAM is actually less than 3 dB, but at 5.12 Msps, 32-QAM delivers a 25 Mbps service (vs. 20 Mbps for 16-QAM and 30 Mbps for 64-QAM). In Table 1, for significant digit purposes, we rounded the SNR requirement, and for margin in the proper direction we rounded it upwards. An intermediate step based on 32-QAM is a simpler and more controlled upgrade with higher probability of success, from which a measure of the additional pain of moving all the way to 64-QAM can be estimated.

While stressing a gradual migration, it is also worthwhile to point out that 64-QAM does not necessarily present the upper limit to the achievable QAM profile on the upstream from an SNR perspective. This is again evident by considering the data shown in Figure 1 and Table 1. It is clear that DFBs offer the potential for high SNR links, in particular in the absence of node combining. Digital links offer the same potential, as has been pointed out. NPR curves show 40-45 dB is available, which is clearly above the 1e-8 / 34 dB SNR requirement of 256-QAM, and with good margin at the high end. It is this relationship that brings about discussions about potentially implementing 256-QAM on the upstream. Of course, at these high SNRs especially, the dominant impairment mechanisms will be the key survival factors, and plant fidelity and maintenance is taken up another notch. However, the evolution of HFC towards deeper fiber and smaller serving groups are all favorable to channel conditions improving over time. Also, noteworthy in the digital return solution with potentially very high NPR performance is that factoring in margin associated with optical link length no longer applies.

B. The 6.4 MHz Effect – Microreflections & Linear Distortions

Spectrum Considerations

In addition to accommodating higher order modulation, DOCSIS 2.0 and 3.0 also increased the maximum symbol rate from 2.56 Msps (3.2 MHz bandwidth) to 5.12 Msps (6.4 MHz bandwidth). DOCSIS 3.0 further enables bonding of upstream channels for a higher possible peak burst rate service. However, since this technique is logical bonding – single channels demodulated individually and glued together at the logical layer – there is not a difference in how the physical layer behaves due to the action of bonding itself. There are variables around using bonding that can affect the link that operators should be aware of. In particular, there is the proper distribution of upstream channel power to consider if not already following a per-Hz philosophy. And, there is the impact of what the use of more spectrum means to the probability of encompassing more impairments, and to increasing distortion. But, there is nothing inherent about the bonding process itself that changes the way the physical layer demodulation behaves.

We will consider channel loading and the channel bonding aspect in a subsequent section. Touching briefly on these secondary spectrum considerations, there is a need to be wary of channel bonding from a frequency domain perspective, assuming it is through the addition of new spectrum that the channel bonding is being turned on – as opposed, for example, to already functioning upstream channels that are now being converted to be part of a bonding group. In general, the concern with any new spectrum is simply to be wary when moving away from the return path "sweet spot" that exists roughly between 20-35 MHz. As spectrum below 20 MHz gets turned on for services, it is recommended that a characterization of the behavior in that part of the band take place pro-actively, as this region becomes susceptible to combined ingress and impulse noise. Note that this refers to any part of the signal bandwidth being deployed below 20 MHz, not just the center frequency.

The effects are magnified as the spectrum dips further below 15 MHz, and continues to degrade to the 5 MHz band edge. To extract capacity from this low end of the return band, use of DOCSIS 3.0 S-CDMA is recommended. A discussion of S-CDMA will be given in a subsequent section.

Time Domain Perspective

As spectrum becomes deployed above 35 MHz, it is important to recognize the contribution that frequency response distortion may have. When the roll-off of cascaded diplex filters occurs over deep cascades, it effectively eats into the return band from the top down, gradually increasing attenuation and adding group delay distortion in the channel, causing intersymbol interference (ISI). Expressed in the time domain, this is the distortion evident in an eye diagram that sees excessive closure because the adjacent symbol are spreading into one another's sampling intervals, making the work of the detection process more difficult. Fortunately, upstream equalization has evolved with DOCSIS, and includes both a longer tap delay line, as well as pre-equalization functionality to help mitigate these effects. However, when combined with inherent microreflection energy in the plant, the equalizer can be overwhelmed and unable to fully correct for the distortion, as observed in a poor EQ-MER value.

The frequency response distortion phenomenon above can occur for any DOCSIS bandwidth, in principle. However, it is aggravated in particular by the 6.4 MHz bandwidth mode, and further aggravated by the increased sensitivity of the higher order modulation. From the time domain perspective, the symbol rate for this mode is now half of that of the 3.2 MHz mode – the previous maximum symbol rate – at just below 200 nsec. While the equalizer has evolved to a powerful 24-tap structure, the fact that it is a *fixed* 24-tap structure means that each symbol rate increase corresponds to a reduction in the total time span of this T-spaced (symbol period-spaced) tapped delay-line equalizer. Whereas for 2.56 Msps, the time span was (~500 nsec) x (24) ~12 usec, it is now halved to about 6 usec. Furthermore, as will be seen in subsequent analysis, only about 2/3 of this span is associated with equalizing the ISI distortion associated with delay, or about 4 usec. The result is a time span that creeps into the neighborhood of possible cable run lengths that, if it corresponds to a cable run that induces reflections, can miss the equalizer capability altogether and contribute directly at its interference level. This situation is atypical, but possible. Particular segments of plant geographies are more prone than others.

Frequency Domain Perspective

From a frequency domain perspective, the wider bandwidth brings into play more potential frequency-selective multi-path reflection potential. More importantly, however, it means that more of the channel's passband could fall into an area of spectrum suffering attenuation or group delay distortion. Put another way, at the upper end of the return band, it may be necessary to deploy the wider band channel at a lower center frequency in order to avoid the impact of band edge roll-off for the 6.4 MHz wide bandwidth, compared to the 3.2 MHz bandwidth. Because the roll-off of the channel is a function of the cascaded diplexers in the HFC plant, it turns out that the center frequency selection is a function of the cascade depth. This is most evident for the 6.4 MHz carrier, because the channel's width and cascade depth, in combination, may demand that the center frequency drop into a spectral region generally considered the return "sweet spot" previous described.

While both reflections and frequency response distortion individually create degradation, they are also individually of limited threat under normal, practical conditions. A more probable scenario associated with performance degradation is the *combined* effect of strong micro-reflections, within the equalizer's range, and frequency response distortion. Both of these require the equalizer to work to mitigate the impairment energy, and in combination can overstress the equalizer itself, especially in the case where it is being asked to handle 64-QAM.

We will now take a more detailed look at the above distortion mechanisms individually, and quantify what they mean to wideband QAM channels.

Frequency Response Distortion Modeling

Attenuation Distortion (AD)

Amplitude Distortion is undesirable variation in the channel's amplitude response. Common forms of AD include tilt, ripple, and roll-off. A common cause of AD is the upper return band-edge carriers, aggravated by long reaches of coaxial plant. Long reaches of coax accumulate diplex filters from devices including amplifiers and in-line equalizers. While individually contributing small attenuation versus frequency, when part of a deep RF amplifier cascade, the combination can create appreciable response variation. An example of amplitude roll-off is shown in Figure 4.

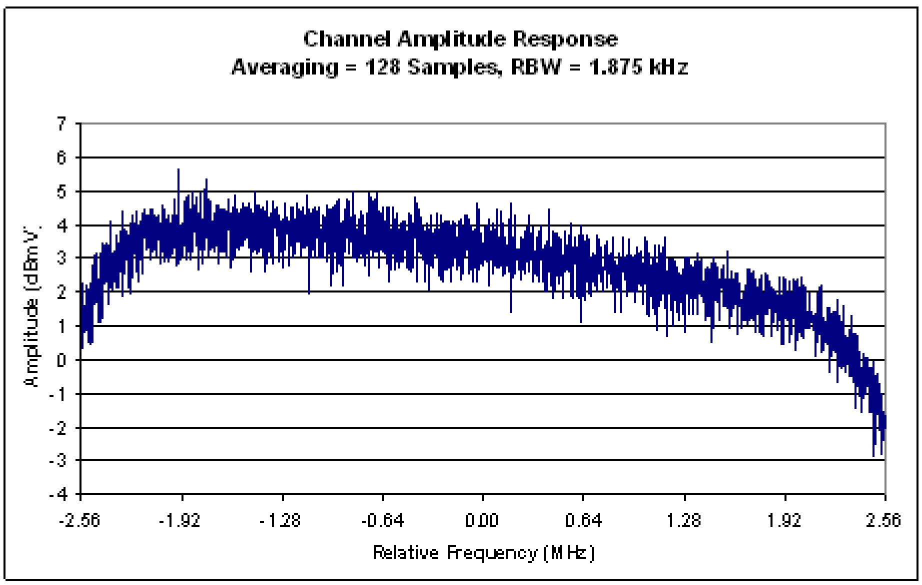

Figure 4 - Amplitude Roll-Off Impairment – Channel Amplitude Response

In a QAM constellation, the amplitude roll-off causes the symbols to spread in a pattern similar in *appearance* to Additive White Gaussian Noise (AWGN). Figure 5 and Figure 7 show constellations for 16-QAM and 64-QAM respectively, which have been impaired by "equivalent" levels of AWGN to the attenuation response shown in Figure 4, whose effect is shown beside the AWGN cases in Figures 6 and 8.

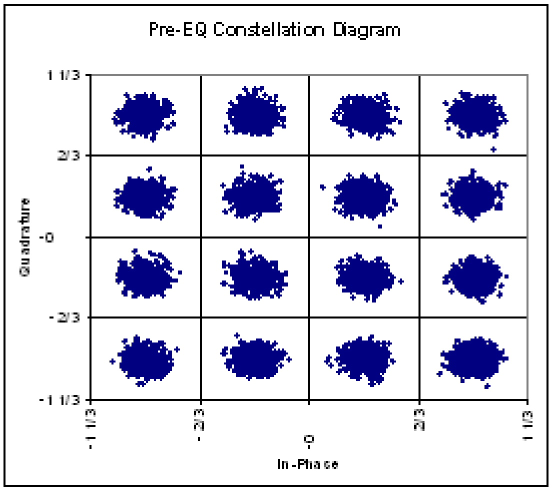

Figure 5 - AWGN Impaired 16-QAM

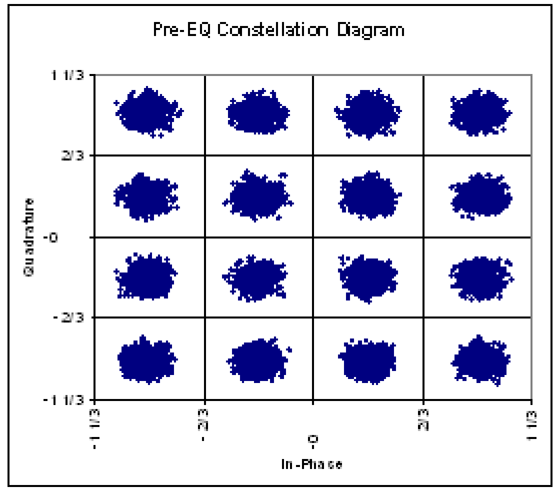

Figure 6 - AD Impaired 64-QAM

Per our earlier discussion of QAM relationships, increasing the AWGN contribution by 6 dB would show that 16-QAM as close to the decision boundaries as 64-QAM is in Figure 7. Conversely, reducing AWGN contribution by 6 dB would show 64-QAM is now just as far away from the decision boundaries as 16-QAM is in Figure 5. In fact, the 64-QAM case is bordering on a catastrophic link result without some intervention.

Note that similar assumptions must not be made regarding the impairments discussed in this section. For our case here, simulation and test is crucial for characterizing the true nature of the relationship which exists between these more complex impairments and modulation complexity, and in particular for multiple simultaneous impairments.

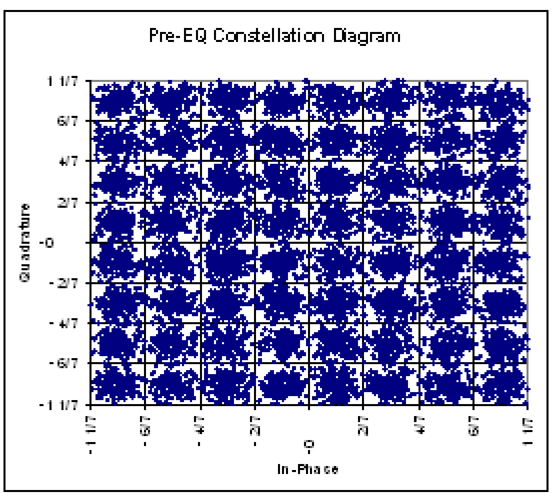

Figure 7 - AWGN Impaired 64-QAM

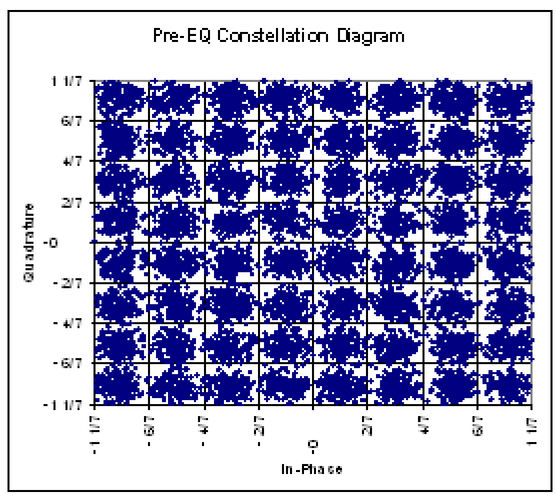

Figure 8 - AD Impaired 64-QAM

Group Delay Variation (GDV)

GDV is undesirable variation in the communication channel's phase response, resulting in distortion of the digital signal phase, or a variation in the propagation of frequency components of the signal across the channel.

As is the case for AD, one major cause of GDV in the HFC plant is upper-band-edge operation, combined with long reaches of coaxial plant. The reasoning is the same as in the AD case. Note that filtering functions typically induce nonlinear phase responses as the band edges are approached, so the combination of AD and GDV in the same band region, understanding that diplex filtering is the cause, is perfectly expected. Different filter functions induce different GDV responses, just as different filter functions induce different stop-band characteristics. It is common that the sharper the roll-off, such as would be the case for long cascades, the worse the GDV will be.

In a QAM constellation, GDV causes the symbols to spread in a pattern similar to what has already been illustrated for AWGN and AD, Figures 5 through 8. As expected, 16-QAM is less sensitive to GDV than 64-QAM because of reduced decision boundary size of 64-QAM.

Micro-Reflection, (MR)

As seen by a receiver, a MR is a copy of the transmitted signal, arriving late and with reduced amplitude. The result of the additional copy is the familiar image ghosting in analog video reception, but for digital communications the result is intersymbol interference, or ISI.

MR sources are composed of pairs of HFC components separated by a distance of cable. HFC components facilitate the propagation of signal copies in a variety of ways including return loss, isolation, mixing, and combining.

Figure 9 illustrates one of many possible MR source configurations. Two devices with poor return loss, acting as signal reflectors, are separated by a length of cable. The CM is acting as the second reflector, but any HFC component has the potential to achieve a similar result. The reflector return loss and the loss between the reflectors determine the amplitude of the MR. The delay encountered as a signal copy traverses the red path of Figure 9 will determine which equalizer tap is responsible for correction.

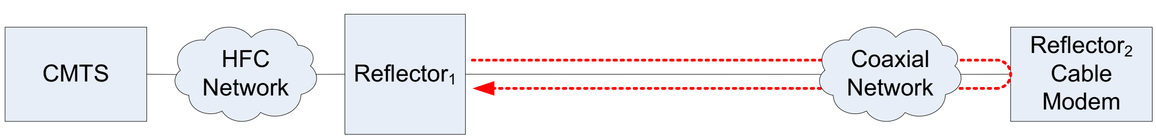

Figure 9 – Micro-Reflection (MR) Source

Note that the CM has as a design limit of 6 dB return loss, meaning it may reflect up to 25% of its incident power. In the plant, design limits are typically significantly better, but over time will degrade as the plant ages and elements that contribute to good RF matching – connectors, cable, splitters, interfaces on PCBs – degrade.

On a spectrum analyzer, an MR appears as amplitude ripple. The peak-to-peak amplitude and frequency of the ripple, a sample of which is shown in Figure 10, are directly related to the MR's amplitude and delay. In this case, the signal is impaired by a MR whose relative amplitude is -20 dBc and whose delay is 4 symbol periods.

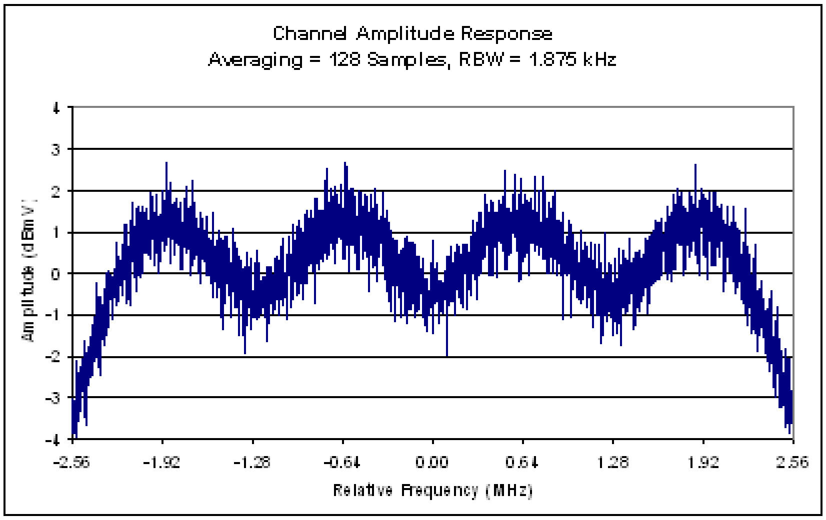

Figure 10 – MR Impairment – Channel Amplitude Response

In a QAM constellation, a micro-reflection causes the symbols to spread in a miniaturized pattern similar to the full QAM constellation itself. Additionally, phase distortion may cause the spread symbols to appear rotated. Figure 11 shows the effect of a MR on a 16-QAM signal's constellation diagram, where it becomes evident why the reflection is in fact a "signal copy." The MR's characteristics are those depicted in Figure 10. Note the spread throughout the symbol region on each 16-QAM point, and subsequently how now less additive noise has more likelihood of causing a symbol to jump a boundary and create a hard decision error.

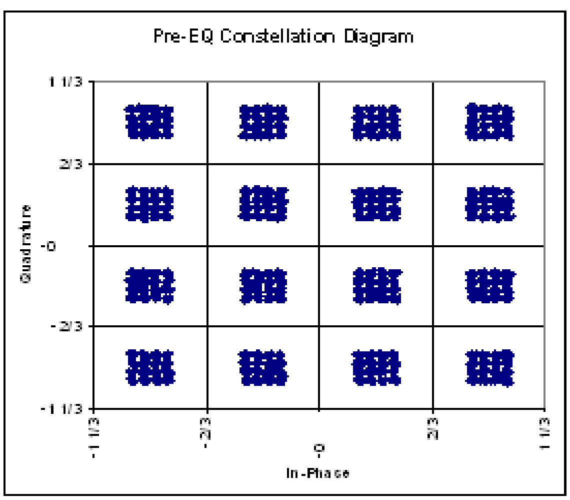

Figure 11 - MR Impaired 16-QAM

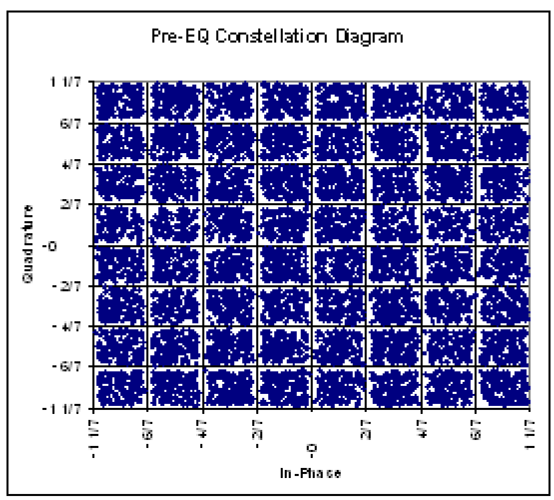

Figure 12 - MR Impaired 64-QAM

Figure 12 repeats the same MR scenario for 64-QAM. In comparing Figure 11 and Figure 12, the same level of MR impairment has spread the symbols of the 64-QAM constellation appreciably closer to the symbol boundaries than in 16-QAM constellation. The 64-QAM situation is clearly catastrophic situation without equalization applied to undo the MR.

Combined Impacts

Simulation and tests were performed of increasing single dominant MR impairment. The results of these tests reveal the highest MR impairment level that could be corrected by DOCSIS 2.0/3.0 Pre-Eq. 16-QAM and 64-QAM have both been evaluated in [6].

Simulation and measurement for both 16-QAM and 64-QAM, illustrated in Figure 13 and Figure 14, reveal DOCSIS 2.0 mitigation of impairments is appreciably higher than what is assumed by DOCSIS to be present in the HFC environment. Additionally, there is a reduction of correction capability as 16-QAM signals are migrated to 64-QAM. Note that this reduction is, approximately, only 2 dB on average.

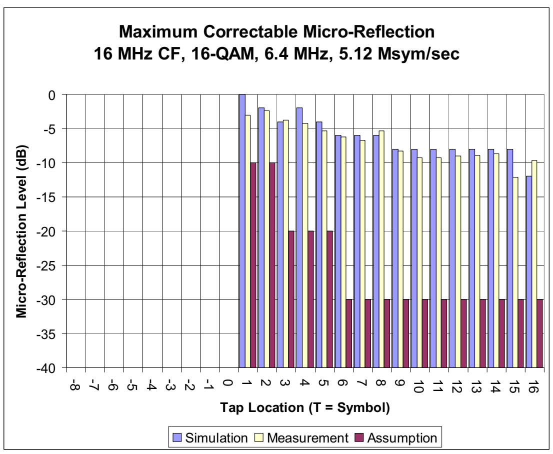

Figure 13 - 16-QAM MR Correction

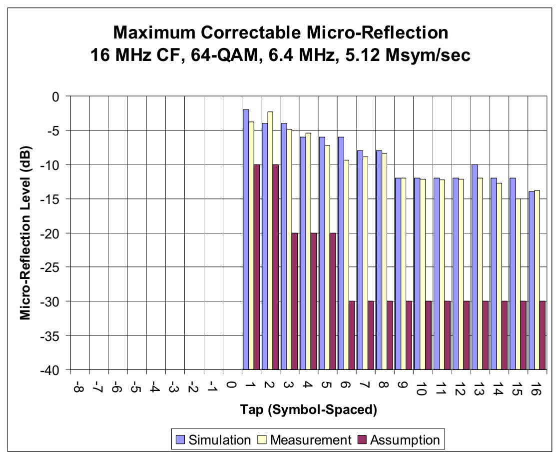

Figure 14 - 64-QAM MR Correction

Simulation and test of increasing MR impairment were repeated adding AD and GDV impairments by locating the signal near the upper band edge of the return path [6]. Figure 15 and Figure 16 are the AD and GDV present at the band edge.

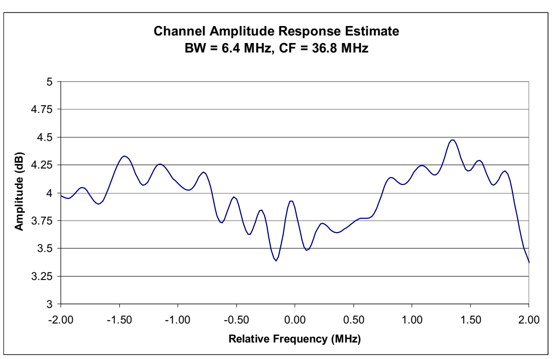

Figure 15 – Channel AD

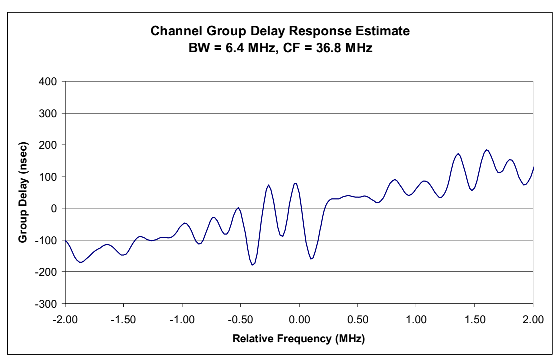

Figure 16 – Channel GDV

The results, illustrated in Figure 17 and Figure 18 for 16-QAM and 64-QAM respectively, reveal a slight decrease in correction capability of the DOCSIS 2.0 Pre-Eq with the additional increase in AD and GDV impairments.

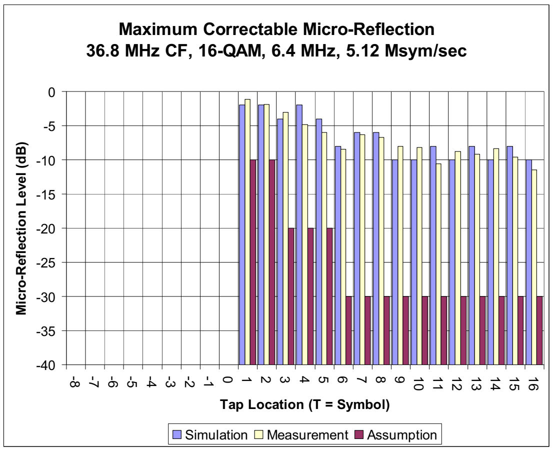

Figure 17 - 16-QAM MR, AD, GDV Correction

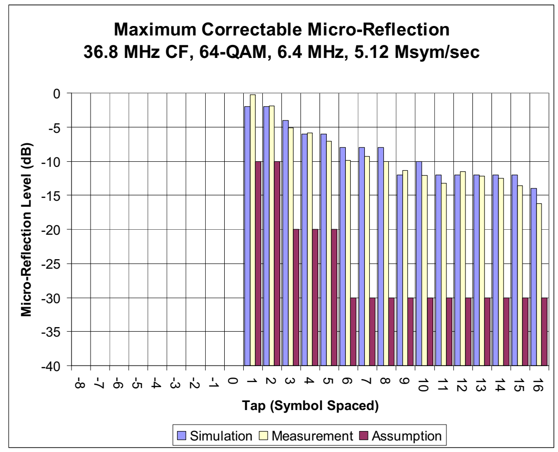

Figure 18 - 64-QAM MR, AD, GDV Correction

Summarizing, the results shown in Figure 13 through Figure 18 are significant for the following reasons:

1. Discussed impairment levels can significantly exceed DOCSIS channel assumptions and still be corrected by DOCSIS 2.0/3.0 Pre-Eq
2. Simulation results closely agree with laboratory measurement
3. MR impairment impact (2 dB) on modulation complexity is different from AWGN impairment impact (6 dB)

Continued investigation of impairments and combinations thereof can complete the acceptable performance limit requirements of DOCSIS 2.0/3.0 Pre-Eq. An understanding of the impairment limits and relationships versus modulation complexity will help cable operators define maintenance requirements and transition toward optimal upstream throughput.

Frequency Response Distortion – Measurement Results

To get a comprehensive understanding of the frequency vs. cascade depth relationship, a thorough lab characterization is required to support the modeling. Analysis used four A-TDMA 6.4 MHz bandwidth carriers as follows:

				Start Freq.	**Stop Freq.**
• Carrier 1	=	16.4 MHz	-	13.2 MHz	19.6 MHz
• Carrier 2	=	23.2 MHz	-	20.0 MHz	26.4 MHz
• Carrier 3	=	30.0 MHz	-	26.8 MHz	33.2 MHz
• Carrier 4	=	36.8 MHz	-	33.6 MHz	40.0 MHz

Note that the upper carrier frequency is the highest 6.4 MHz wide center frequency that can be used for a 40 MHz return split without violating the upper limit. The N+9 cascade tested included a structural MR as might be encountered in any HFC plant. In addition, the optical link was based on an FP Laser – whereas recommended practices call for a DFB link for 64-QAM on the return path. The FP, of course, limits the attainable MER. The FP return laser coupled with a built-in "plant" MR makes the laboratory N+9 cascade quite similar to what is often observed in real plant environments in the field.

Cascade Depth Distortion Results

Tables 3, 4 and 5 summarize measured GDV, AD, and MER performance. GDV was averaged across 4 MHz bandwidth. Overall, GDV and AD conditions degrade with increasing cascade and increasing center frequency as the carrier approaches the upstream band edge of 40 MHz.

Table 3 - Cascaded GDV vs. Center Frequency

GDV (ns/MHz) vs. Cascade Depth vs. Center Frequency						
	N+4	**N+5**	**N+6**	**N+7**	**N+8**	**N+9**
36.8 MHz	44.42	49.04	60.20	89.03	93.44	97.41
30.0 MHz	20.15	22.70	24.63	25.62	26.98	26.48
23.2 MHz	16.13	16.87	17.71	13.61	15.13	13.84
16.4 MHz	10.32	9.82	9.23	7.69	7.53	6.88

Table 4 - Cascaded AD vs. Center Frequency

AD (dB/Channel) vs. Cascade Depth vs. Center Frequency						
	N+4	**N+5**	**N+6**	**N+7**	**N+8**	**N+9**
36.8 MHz	0.48	0.58	0.62	1.28	1.72	1.46
30.0 MHz	1.52	1.43	1.41	1.40	1.41	1.51
23.2 MHz	1.09	1.11	1.12	1.09	0.86	0.98
16.4 MHz	0.55	0.73	0.62	0.39	0.65	0.43

Table 5 - Cascaded MER vs. Center Frequency

MER (dB) vs. Cascade Depth vs. Center Frequency						
	N+4	**N+5**	**N+6**	**N+7**	**N+8**	**N+9**
36.8 MHz	40.32	37.06	39.31	19.46	*	*
30.0 MHz	41.69	37.67	42.60	34.40	35.98	32.88
23.2 MHz	43.13	40.03	42.46	37.98	34.72	34.81
16.4 MHz	42.41	41.77	41.94	39.75	38.29	35.74

Main tap compression (CMP) and main tap spreading (SPR), shown in Table 6 below, can be used to assess the stress caused to an equalizer by impairments present in the channel. With no impairment present and modest noise levels, the main tap level is 0 dB. As the equalizer approaches instability, the main tap level will drop about 2 dB or more. This has been observed as a point that signals the equalizer's effectiveness against channel impairments is being significantly reduced, as it indicates that appreciable energy in the symbol is actually moved into adjacent signaling interval. This is observed by MER summary of Table 5 as the carrier frequency increases along with the cascade depth, variables each leading to the presence of significant frequency response distortions.

N+9 Failure Thresholds

36.8 MHz is the highest center frequency available for a 6.4 MHz upstream channel used which supports a 5-40 MHz return path spectrum. The amount of GDV and AD present at 36.8 MHz appreciably stresses the equalization function, resulting in very little margin available to combat other channel impairments that may arise, like MRs. It is this combined set of impairments that places deep cascades at most risk – the inherent frequency response distortion of the lowpass cascade of diplexers, coupled with always-present MRs, which only have more opportunities to be present as the cascade lengthens.

An assessment was performed to determine how well the pre-equalization function would be able to sustain the DOCSIS link as the center frequency and cascade depth were incrementally increased to create more imposing frequency response distortion. The results are shown in Table 6. The conclusion for this N+9 cascade was that ***37.8 MHz*** was the last observable center frequency at which the DOCSIS link was still active (cable modem registered). Based on this observation, the link conditions – FP return and a single MR, and the quantified frequency response measurements at this center frequency shown in the table below, a high-speed data customer would likely be on the verge of experiencing data connectivity issues when the combination of GDV and AD reach 103.78 ns/MHz and 1.46 dB, respectively (or ***approximately 100 nsec/MHz and 1.5 dB*** for rounding purposes). Thus, this represents a ***zero-margin*** threshold, beyond which link connectivity is at risk.

Table 6 - N+9 MR, AD, GDV vs. Center Frequency

DOCSIS Link Measurement Approach for N+9 vs. Center Frequency (200 kHz Step)				
	GDV (ns/MHz)	**AD (dB)**	**SPR (dB)**	**CMP (dB)**
36.4 MHz	65.11	0.65	-0.56	-1.15
36.6 MHz	71.25	0.63	-1.65	-1.40
36.8 MHz	75.39	0.54	-2.24	-1.56
37.0 MHz	78.91	0.52	-2.79	-1.71
37.2 MHz	84.38	0.61	-3.35	-1.88
37.4 MHz	87.99	0.91	-4.00	-2.10
37.6 MHz	99.12	1.06	-4.74	-2.37
37.8 MHz	103.78	1.46	-5.48	-2.67

Summarizing, the table above indicates that 37.8 MHz represents the zero-margin threshold under the following conditions:

- 64-QAM modulation
- 5.12 Msps (6.4 MHz BW)
- Pre-Eq = On
- N+9 Motorola cascade (diplexer group delay performance is vendor-dependent)
- Single MR
- FP Return

Because zero margin is not a practical operating condition, the table above suggest about 1 MHz of margin to this zero-margin threshold. This added margin recognizes that the onset of equalizer degradation has begun, but is enough removed from major reduction of capability. The recommended maximum operating point for N+9 occurs at a GDV of approximately 75 nsec/MHz.

As the cascade is shortened below N+9, there is approximately 200 kHz/amplifier of bandwidth that can be freed up to accommodate DOCSIS signals (i.e. 36.8 MHz can become 37.0 MHz at N+8, etc) down to about N+5 or N+6. Below this cascade depth, concerns tend to dissipate with Pre-Eq turned on, as the frequency response distortion is benign enough for the Pre-Eq system to support any reasonable final center frequency for upstream paths of typical link quality.

Acceptable margins between DOCSIS link operation and DOCSIS link failure is not easily and uniformly defined, because of the multitude of practical impairment combinations, of which the comprehensive set of possibilities is nearly infinite in the return path, and link implementation variables. For example, the presence of inline equalizers and in home amplifiers implies the need to have additional GDV margin. Also some older actives have less than ideal frequency response at the diplex filter than newer diplex filters. Operators are advised that the center frequencies listed here are for guidance only, and test the network limitation ultimately answers these types of questions.

Proactive Channel Assessment Using Pre-Eq Channel Information

While we note the function of the equalizer to correct for channel distortions above, there is also a wealth of information available by closely looking at the equalizer tap coefficient values settled into for a given channel to help determine what ails a particular return path. DOCSIS 2.0 Pre-Eq coefficients are a list of the 24 complex values (for the 24 Taps) and may be interpreted in multiple ways. Figures 19 through Figure 23 represent a digital communications channel with negligible levels of impairment as a reference baseline for impairment identification.

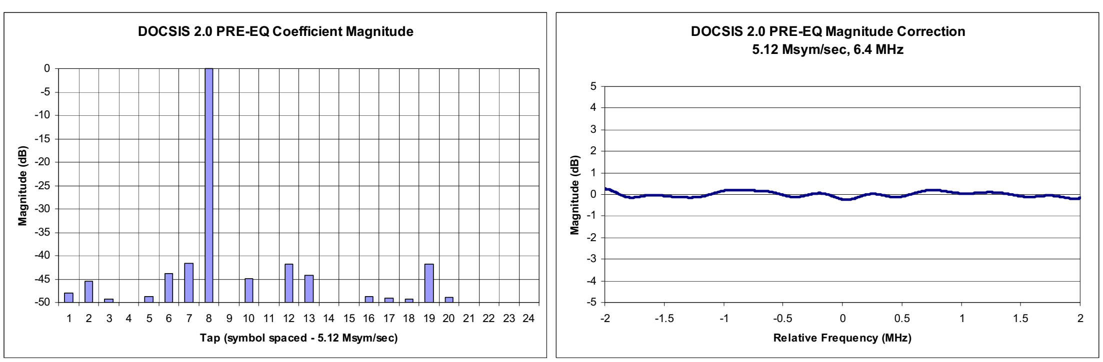

Figure 19 – Well-Behaved Frequency Response – Time and Frequency Domain

Figure 19 is the magnitude of the equalizer's amplitude response in the time and frequency domains. On the left-hand side is the time domain. A single line at Tap = 0, known as the main tap, on the x-axis would be the ideal response, indicating if an impulse is applied to the channel, and impulse comes out the other side. The main tap represents the desired symbol energy, while the remaining taps represent negligible correction magnitudes of < -35 dBc. The very small random magnitudes of the non-zero taps are primarily due to simulated system noise.

The right hand side is the frequency response. The amplitude response would be ideally constant (flat) across the channel's bandwidth. However, a well-behaved response illustrated in Figure 18 may include negligible amounts of amplitude correction, which in this case is about 0.5 dB.

Three important regions of the impulse response to focus on are the *pre-tap, post-tap* and *main tap* regions. Dominant micro-reflections typically impact the post-tap region, which consists of tap 9 through tap 24. Dominant AD and GDV typically impact the main-tap region, which consists of taps adjacent to the main tap, numbers 5, 6, 7, 9, 10, and 11. Dominant GDV impairments have resulted in higher post-tap level than pre-tap levels, where the pre-tap region includes tap numbers 1 through 7. Note that pre-tap energy in the non-ideal channel requiring attention results in the main tap shifting to T =8.

For coefficient analysis of multiple CMs such as in a real plant, it is helpful to break-down impulse response measurements into regions in which dominant impairments will have the greatest impact. Numerical sorts based upon these impaired regions could facilitate efficient organization of similarly impaired CM groups, and this can help in diagnosing issues.

Figure 20 is the phase of the equalizer's impulse response, $\theta_e(t)$. The impulse response phase appears randomized between $-\pi$ and π, except for the main tap whose phase correction is 0 radians. While this plot looks "noisy," recognize from the magnitude response that the amplitude contribution of the phased imperfections is negligible.

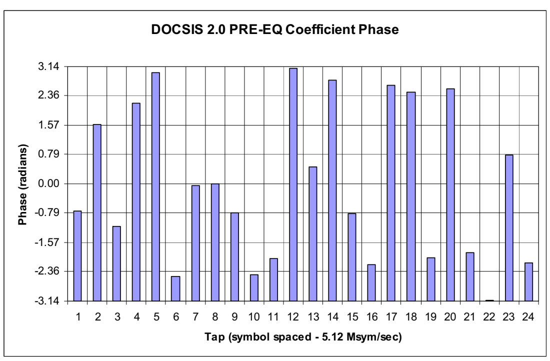

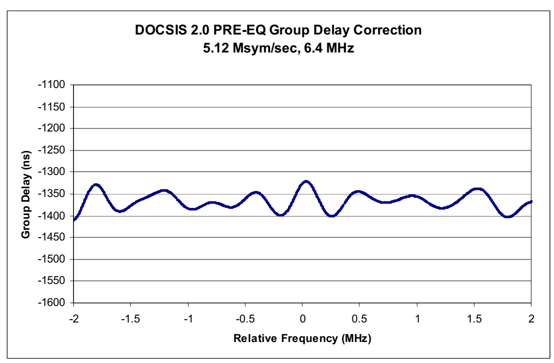

Figure 20 – Well-Behaved Phase Response – Time and Frequency Domain (GDV)

Figure 20 is the equalizer's phase response, $\theta_e(f)$ shown in the time and frequency domains. The equalizer's phase response would be ideally linear throughout the channel's bandwidth. The more commn parameter used for phase response is the Group Delay (GD) response, $GD_e(f)$, shown on the right-hand side. The relationship between group delay and phase response is

$$GD_e(f) = -\left(\frac{\delta\theta_e(f)}{\delta f}\right)$$

Note that the equalizer's GDV is approximately 17 nsec across the channel's bandwidth. Group delay is another way of describing the phase characteristics, but in a way more intuitively descriptive. It represents the absolute time delay each frequency component across the band will endure. As such, it is the variation of this delay that matters most, as components of frequency arriving at different times result in distortion and ISI.

With this baseline, we will catalogue the Pre-Eq tap responses under some common impairments, in order that it's diagnostic capabilities can be illuminated.

Micro-Reflection

A single dominant micro-reflection impacts the DOCSIS Pre-Eq coefficients in the post-tap region as illustrated in Figure 21. This response is identical to that of the Figure 10 MR: amplitude -20 dBc relative to the main tap, and delay 4 symbol periods later than the main tap, as indicated by the frequency response on the right-hand side.

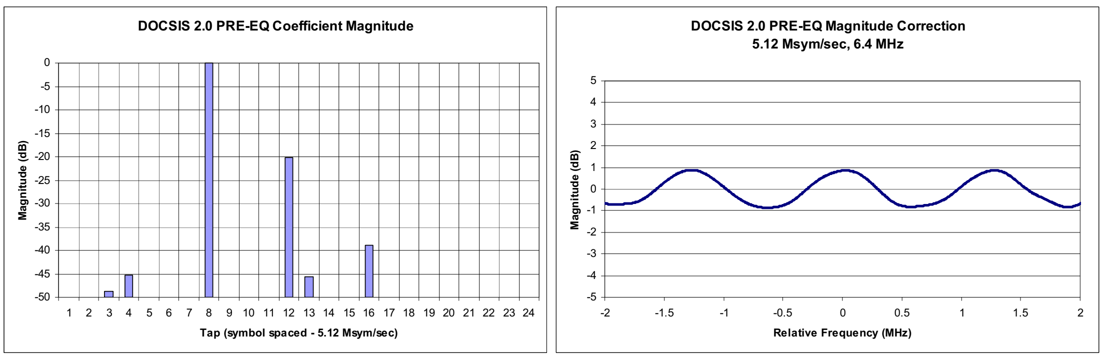

Figure 21 – MR-Impaired Impulse Response Magnitude

The impulse response phase reveals negligible phase distortion of both the desired symbol and the micro-reflection impairment.

Note that, relative to the comparison in Figure 10, the equalization process estimates the *inverse* of the digital communication channel response, $H_c(f)$, and applies it to the incoming signal per the equation below.

$$H_e(f) = \frac{1}{H_c(f)} = \frac{1}{|H_c(f)|} e^{-j\vartheta_c(f)}$$

The phase response shows some nonlinearity across the channel's bandwidth, especially when compared with the previous well-behaved response. The GD response of Figure 22 clarifies this phase distortion with appreciably higher GDV than for the well-behaved case. Note that the equalizer's GDV is approximately 37 nsec across the channel's bandwidth, while the symbol duration itself is less than 200 nsec.

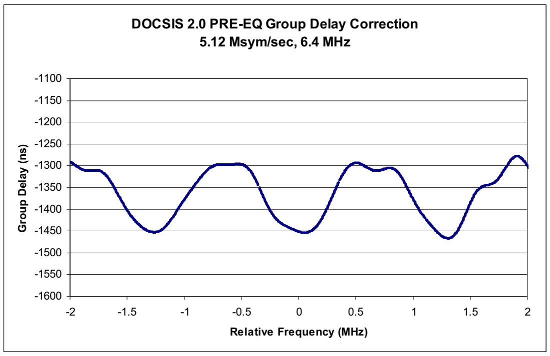

Figure 22 – MR-Impaired FR GD

Attenuation Distortion (AD)

Amplitude roll-off effects the near main-tap region of the impulse response magnitude, illustrated in Figure 23, reveals main-tap spreading in the region of main tap +/- 3 taps. The amplitude response of Figure 24 reveals appreciable amplitude correction, consistent with band-edge roll-off such as that observed in Figure 4.

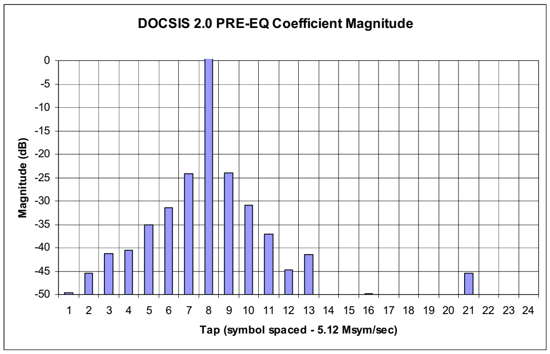

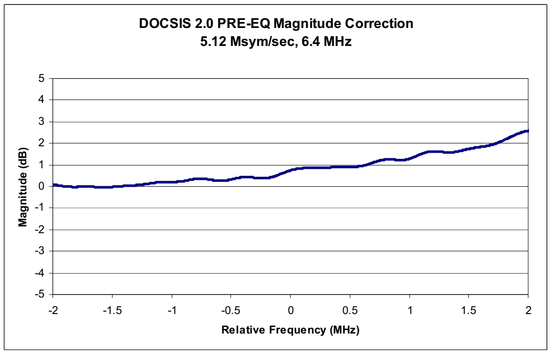

Figure 23 – AD-Impaired IR Mag **Figure 24 – AD-Impaired FR Mag**

Group Delay Variation (GDV)

The main-tap region of the impulse response magnitude, shown in Figure 25, reveals main-tap spreading similar to what was illustrated for the amplitude roll-off impairment in the equalizer's taps leading up to the main tap, but the post tap region has higher tap levels than the pre tap region. The amplitude response, shown in Figure 26, reveals negligible amounts of amplitude correction. Since the induced impairment is specifically phase related, this makes sense.

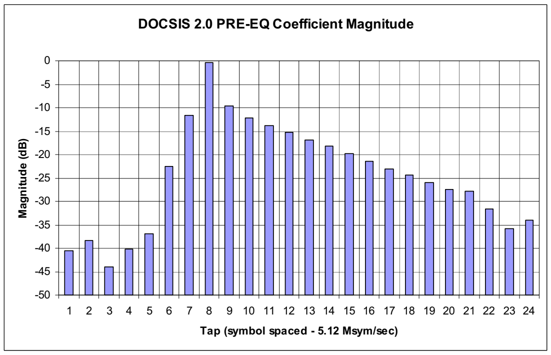

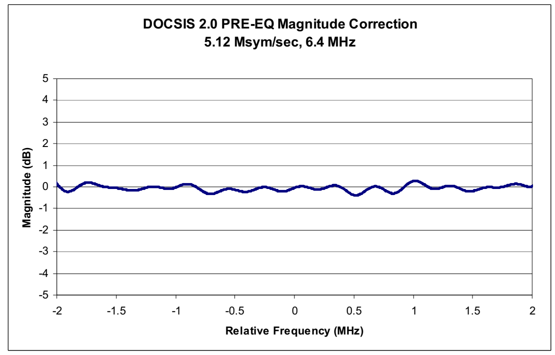

Figure 25 – GDV-Impaired IR Mag

Figure 26 – GDV-Impaired FR Mag

As expected, Figure 27 reveals appreciable GD correction over the GD correction associated with the phase distortion that co-exists with roll-off impairment scenario exhibited in Figure 4. Note equalizer's GDV was approximately 54 nsec across the channel's bandwidth.

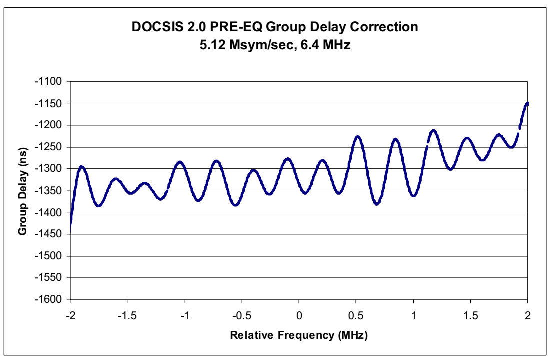

Figure 27 – GDV-Impaired FR GD

A Diagnostic Process

Clearly there are ways to take advantage of the wealth of information provided by Pre-Eq to isolate DOCSIS Pre-Eq related impairments. The following step-by-step process is suggested to isolate impairments using equalization coefficient analysis.

Step 1

Ensure that the majority of DOCSIS links are supporting at least DOCSIS 2.0 with Pre-Eq enabled. The resolution of the 24-tap equalizer of DOCSIS 2.0 is better suited to identify impairments, compared to the 8-tap equalizer of DOCSIS 1.1.

Step 2

Query the DOCSIS 2.0 CM population using an SNMP query tool similar to Modem Pre-Eq Response Tool, illustrated in Figure 28. The Modem Pre-Eq Response Tool, custom software developed at Motorola, can query multiple DOCSIS terminal devices based on an IP address list. Periodic polls of coefficient values and other relevant physical layer (PHY) metrics are displayed and/or stored into a log file for post processing. This tool also provides users with a graphical view of either the impulse response or amplitude response for each CM poll. This tool can help cable operators establish a baseline of performance, and identify problem DOCSIS links, based on CM IP addresses.

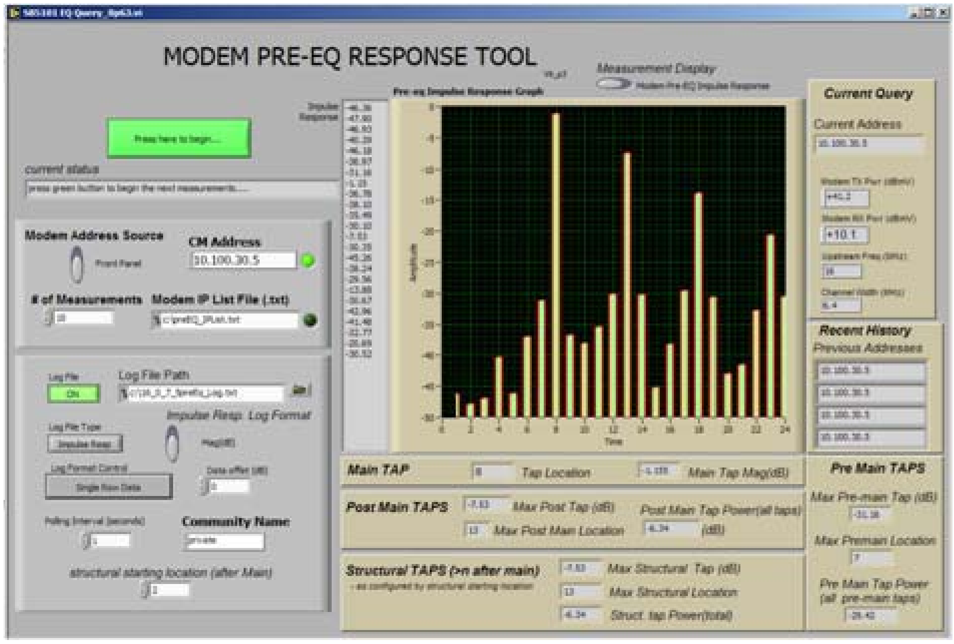

Figure 28 - Transmit Pre Equalization Query Tool

Step 3

Identify impairment problems by sorting, on increasing levels that sum the DOCSIS Pre-Eq regions previously defined. For example, determine which CMs experience the greatest amount of micro-reflection impairment by sorting on the levels which result from summing the taps located in the post-tap region.

Step 4

Understand the above described problems and how they originate in HFC plant. For example, one micro-reflection source has been discussed in the micro-reflection fundamentals section, but many possible permutations of micro-reflection sources must be understood for successful isolation.   Use results to define what impairment levels will likely result in service calls, consider potential thresholds to address proactively, and what margin is necessary to prevent calls.

Step 5

Leverage the CM population to differentiate between CMs experiencing an impairment problem and those that are not. For example, a query of the CM population of the HFC coaxial feeder segment illustrated in Figure 29 (an actual case study) reveals that CMs located off of amplifier 1 are reporting a MR problem, while CMs located off of amplifiers 2, 3, and T1 are not reporting a problem.

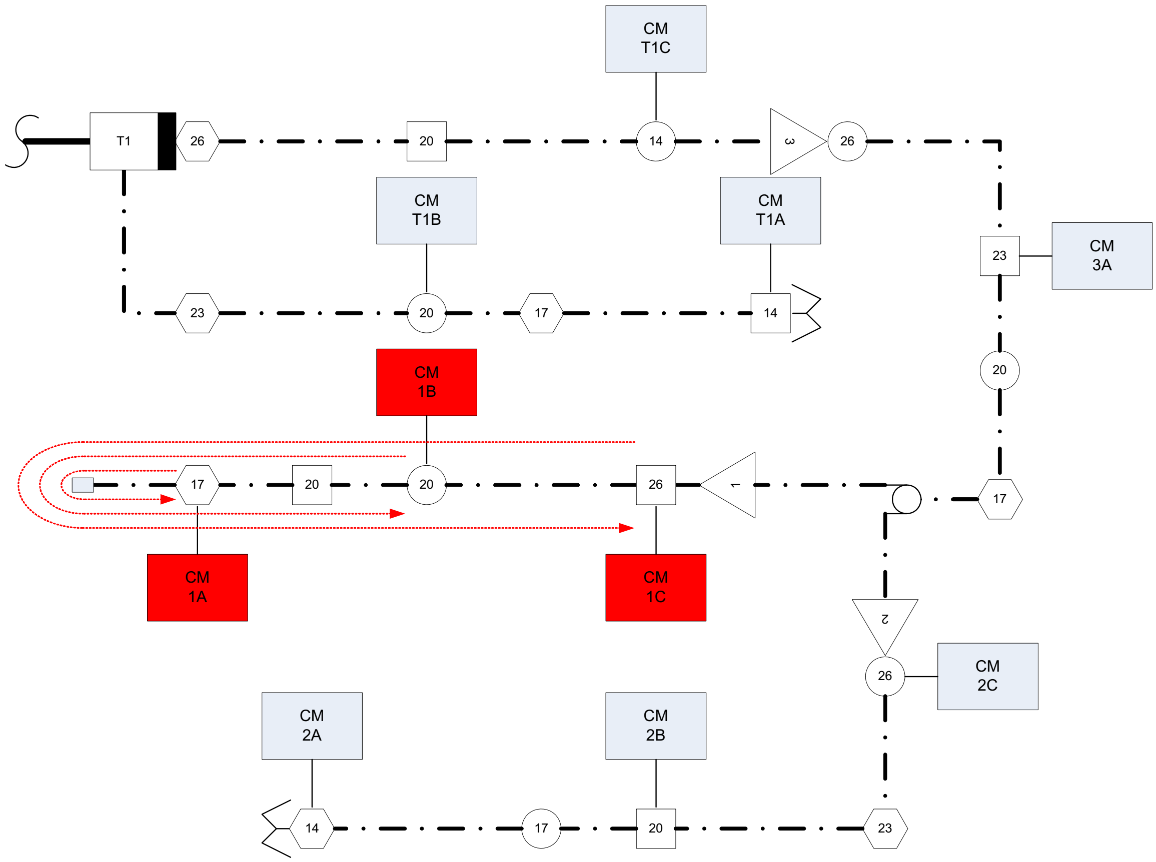

Figure 29 - Micro-Reflection Impairment Isolation Example

Step 6

Identify suspected HFC components using results from Step 4 and Step 5. In the example above, therefore, all of the HFC components fed from amplifier 1, and amplifier 1 itself are suspect.

Step 7

Inspect and repair as necessary all suspected HFC plant components resulting from Step 6. Again, referring back to the example provided, inspection of the suspect components showed that the micro-reflection source to be a combination of tap-to-output port isolation loss and an unterminated cable splice at the end of the amplifier 1 feeder run. Properly terminating the splice reduces the micro-reflections to negligible amplitudes.

Step 8

Repeat CM population query and compare to baseline captured in Step 1 to ensure that the impairment problem has been eliminated and the improvements are sustainable.

As this section has shown, fully understanding the Pre-Eq function and deploying some simple tools to perform equalization coefficient analysis on the data gathered by this function makes it possible to identify the dominant impairments for which the DOCSIS 2.0 Pre-Eq is compensating. Further simulation and measurement are required to determine all of the points at which DOCSIS 2.0 Pre-Eq will work under various impairment combinations and levels. Understanding these limits will help cable operators establish when proactively maintaining the HFC plant will be most beneficial, leading to a refined process and help cable operators leverage the benefits of DOCSIS Pre-Eq coefficient analysis.

C. Spectrum Usage

Return Alignment Theory

The concepts of properly selecting operating levels and aligning the return path have not changed since the mid-1990's, when carrying a full spectrum of digital services on the return path was first discussed. The cable industry was at a crossroad during that time, uncertain as to whether the return path could offer a reliable connection for data and telephony services. Motorola (then General Instrument) presented a paper at SCTE's Emerging Technologies Conference in San Francisco in 1996 [B]. That paper introduced the concepts of tap/drop equalization in the RF return plant to decrease in-home level variability and reduce ingress and the "Constant Power per Hz" methodology of selecting laser drive levels.

Many people in the cable industry thought that the concept of Constant Power per Hz was too simplistic and proposed other more complex methods of optimizing the available laser drive level for the services being deployed. Over the years, some MSOs adopted these alternative methods and aligned their lasers to be fully loaded with the limited signals in service at that time. Now, as DOCSIS 3.0 is being deployed and MSOs are preparing to fully load the return path spectrum, they find themselves needing to readdress the issue of laser loading. As MSOs are preparing to offer bonded upstream channels, they are now adopting a Constant Power per Hz methodology, even if that's not what they call it. Similarly, the concept of using equalization in the tap to reduce the variance of return path levels was viewed to be too cumbersome or expensive. Yet, today, as MSOs work to utilize the entire spectrum and desire to be able to have high bandwidth signals originate deep in the subscriber's home, many are adopting the use of "window taps," which allow 35, 32, 29, and 26 dB taps to be replaced with 23 dB taps that allow return signals into the network with less loss, yet include a "cable simulator" in the tap to equalize forward path levels to the proper value. Some are also using a different configuration of taps that allows loss to be added to the return path in low value taps.

During the initial years after that paper was presented, many papers and workshops were presented at various SCTE meetings to teach the industry how to properly design, align, and maintain the return path of an HFC network. Through much hard work, the cable industry showed that the return path can, in fact, be operated and maintained in a manner sufficient to provide an excellent user experience. Now, almost 15 years later, virtually every cable plant in North America and majority of the plants in the world are running two way services. However, many of the people that designed the initial networks have moved on to other jobs or positions.

Much of the knowledge of how to properly select plant, laser and headend levels has been lost, and systems are simply following the rules established in the past. As the industry moves towards fully loading the entire return path spectrum with DOCSIS 3.0 services, it is necessary to be certain that today's engineers and technicians understand what levels should be selected and how to design the network to achieve those levels. Additionally, as the upstream becomes congested, operators need to segment nodes, and have started to pull fiber deeper into the network, converting amplifiers into nodes. The current generation of technicians needs to understand how to align the network.

To assure that the HFC return path is optimized for DOCSIS 3.0 upstream signals, we will discuss the critical aspects of selecting levels for each portion of the network and aligning the gain between each portion so that each segment operates at its optimal level.

Choosing Operating Levels

Any discussion of levels in an HFC network should be divided into several distinct areas, including levels in the home and RF plant, levels in the optical link (laser transmitter and receiver), and levels in the headend. The levels in each of these three areas are determined using different criteria. Proper alignment of the return path cannot be accomplished until the operating levels are selected for each area. All portions need to be operated at the correct signal level simultaneously. Once the levels for the plant, the laser transmitter and the headend are chosen, you will know how much gain or loss is required inside the optical node, which is where the RF plant feeds the optical link, and in the headend between the optical receiver and the CMTS or other demodulator.

Figure 30 shows the major components of the HFC return path network. The major components are:

1. RF transmitter in the home
2. Loss between the transmitter and the RF amplifier station port
3. RF amplifier station port
4. Node station port
5. Return path laser module RF input
6. Headend optical receiver RF output
7. CMTS or other upstream RF demodulator input

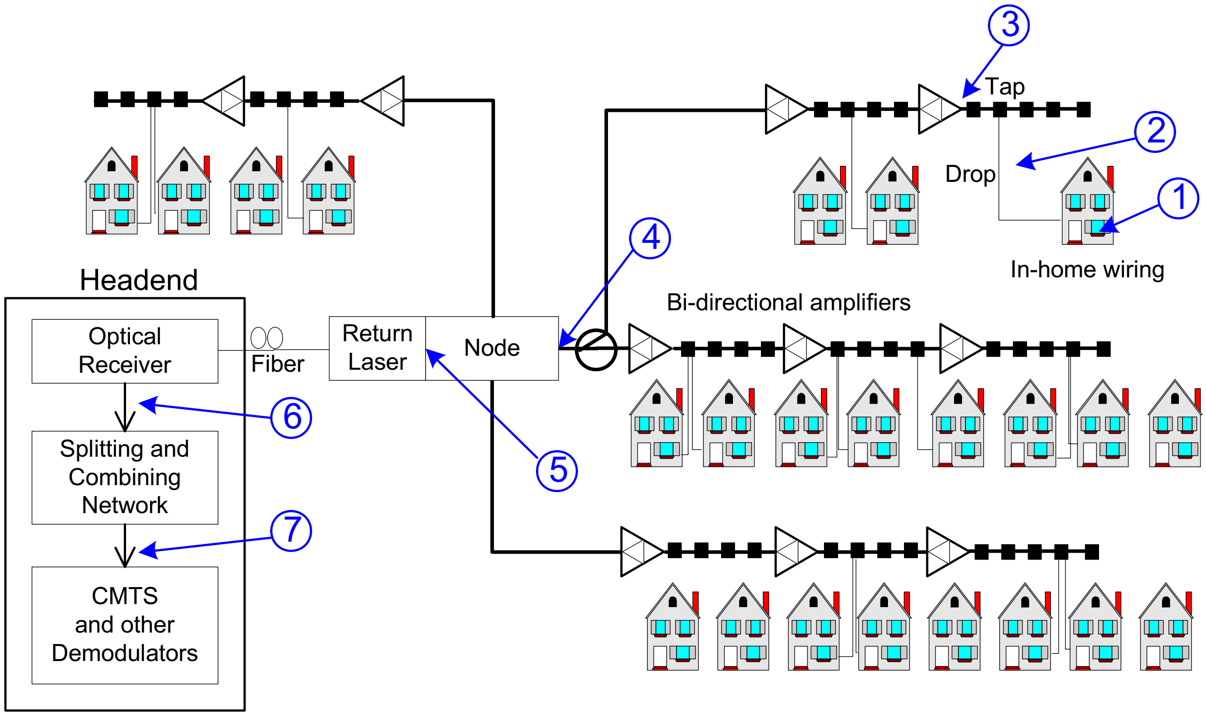

Figure 30 – Return Path Network

Long Loop AGC

Long loop Automatic Gain Control (AGC) refers to the gain control that must occur within every service that operates on the return path. You need to understand how long loop AGC works in order to properly set levels for each portion of the network. The long loop AGC is the mechanism through which upstream levels are adjusted. "Long loop" refers to the fact that the commands are issued all the way from one end of the plant (the headend) to the other end of the plant (inside the home) to affect a signal originating inside the home and destined for the headend. Thus, a long loop is formed all the way across the plant and then back again. "Gain control" refers to the process of automatically adjusting gain. Strictly speaking, it is the level and not the gain that is controlled, but the term "long loop AGC" is widely used anyway.

All levels are controlled from the headend. In a DOCSIS network, this controller is the CMTS. The CMTS commands the modems to raise or lower their transmit levels until the level received from that modem in the headend is at the proper level. Because the amount of loss between the subscriber's equipment and the cable plant (in house cable, in house splitters, drop loss, feeder loss, and other losses) is not known, there is no way to predict what transmitter level is required to ensure that the return path signal enters the plant at the correct level. Nevertheless, the network must be designed so that the modem is capable of producing the required level and that, when the received level at the CMTS is optimized, the levels in all portions of the network are also optimal.

Referring to Figure 30 and working from the headend out into the plant, we see the following relationships that must be adjusted:

1. **Headend Splitting/Combining Network Gain:** The gain of the splitting/combining network must be adjusted such that the output of the optical receiver (6) is at an optimal level when the CMTS (7) is receiving it's target upstream level.
2. **Optical Receiver Gain:** The gain of the optical receiver must be adjusted so that the input to the laser (5) is at an optimal level when the output of the optical receiver (6) is at it's optimal level.
3. **Node Return Gain:** The gain inside the node must be adjusted so that the RF level at the node station port is at the optimal plant level (4) when the input to the laser (5) is at it's optimal level.
4. **RF Amplifier Unity Gain:** The gain of each RF amplifier must be adjusted such that the RF level at each RF amplifier station port (3) is the same as the RF level at the node station port (4).

The operator must always keep Long Loop AGC in mind when making changes to the gain structure. For instance, decreasing padding in the headend will lower levels everywhere (the home, the plant, and the laser). But, decreasing padding in the node will only lower levels in the plant and in the home, without affecting the level into the laser.

Determining the Relationship Between the Levels of the Various Channels

Since the return path laser is the device with the lowest dynamic range, any discussion of selecting the proper relationship between channel powers focuses on the fiber optic link. As mentioned previously, Motorola has been suggesting for over 13 years that all channels operate at the same power density at the laser. This is commonly referred to as "Constant Power per Hz." We do not preclude the use of other more complicated optimizations - and have even suggested some [4] - but believe that constant power density is the best way to achieve acceptable results without overly complicating the network.

Laser Operating Point

Before allocating the total available power to the various channels, the optimal total power must first be determined. As discussed previously, Noise Power Ratio (NPR) is a good way to determine an optimal operating point. Additionally, the performance of the laser can be verified with a BER test to be certain that the selected operating point has the expected dynamic range with real traffic. Since BER is generally error-free over a very large range of laser drive levels, NPR is normally used to determine the exact desired operating point and BER is used to verify proper operation across the dynamic range. The SCTE standard for measuring BER has very good instructions for performing a BER test of a laser [8].

The NPR test is performed by injecting broadband noise with a notch in the middle of the spectrum into the laser at various power levels. As the power level is varied, the depth of the notch is measured. Figure 31 shows an example of the input signal and Figure 32 shows the output signal when there is significant intermodulation noise and clipping. The results are plotted as shown previously in Figure 2.

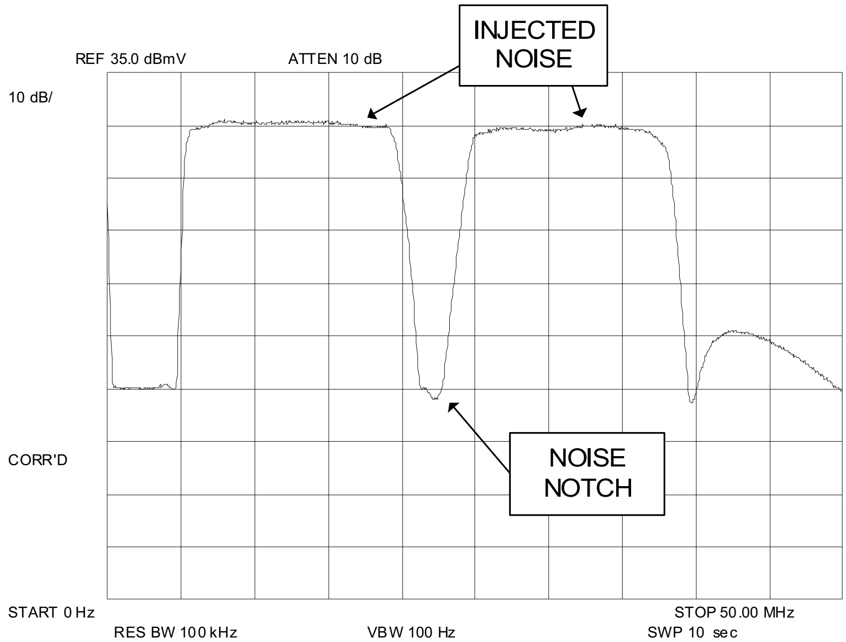

Figure 31 – Spectrum of Injected Signal

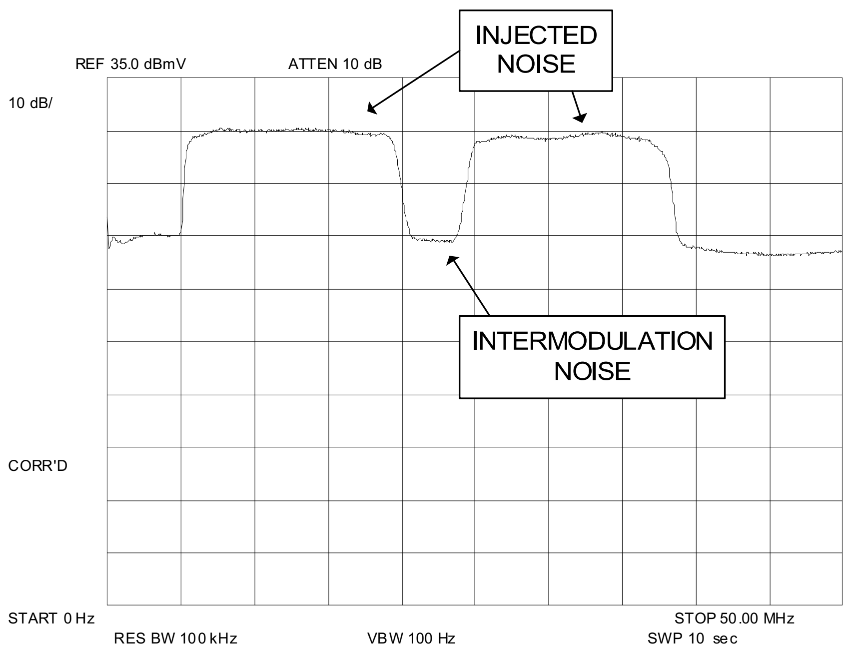

Figure 32 – Signal with Intermodulation and Clipping Distortion

The objective of the NPR test is to determine the best operating point for the laser. If the RF drive level into the laser is too low, the SNR will be low. If the RF drive level is too high, the laser will be very close to clipping and any additional power from ingress or improperly aligned devices will push the laser into clipping. As can be seen in Figure 2, the clipping side of the NPR curve is very sharp. Thus, operating too close to clipping must be avoided. The trade-off is determining how far to the left one must operate, knowing that the SNR continues to get worse as input levels are lowered, until being comfortable that the network is safe from clipping. The range of RF input levels between the lowest input level that gives acceptable SNR and the highest input level that gives the same SNR (but is operating in clipping) is called "Dynamic Range." All networks require dynamic range because, in the real world, the levels of each channel from each subscriber's home are never at the "perfect" optimal level. In general, the optimal operating point is close to the middle of the dynamic range, since the signal level variation in live plant can be both above and below the desired level.

Motorola has been characterizing return lasers for over 15 years and has determined that the optimal operating point for DFB lasers is 20% OMI (optical modulation index) total power. These lasers are aligned such that they are at 20% OMI when driven by a total RF power that is equal to the "Recommended Input Level." This operating point is the best trade-off between SNR and headroom. For FP lasers, Motorola aligns for 35% OMI at the Recommended Input Level. This is to provide for better SNR (since the FP lasers have poorer noise performance) at the expense of clipping headroom. MSOs that use other vendor's lasers should find out how those lasers are aligned so that the proper operating point can be determined.

Operators should be careful not to be so afraid of clipping that they select operating points that are so low that SNR becomes a problem. As mentioned previously, level errors in real plants go both up and down, so it is best to operate in the middle of the dynamic range. Additionally, ingress is not likely to significantly add to the total RF power at the laser. Consider Figure 33, which shows significant ingress. The purple (lower) curve illustrates ingress that is below the signals. At frequencies above 20 MHz, the carrier to ingress is about 30 dB, and below 20 MHz, the carrier to ingress varies from 20 to 5 dB. The point here is that most of the spectrum below 20 MHz is unusable for TDMA and the MSO would be working to fix the problems if the plant looked this bad. Nevertheless, all that ingress only adds 0.2 dB to the total power at the laser. Ingress with inhibit communication long before it gets bad enough to add significant total power to the laser.

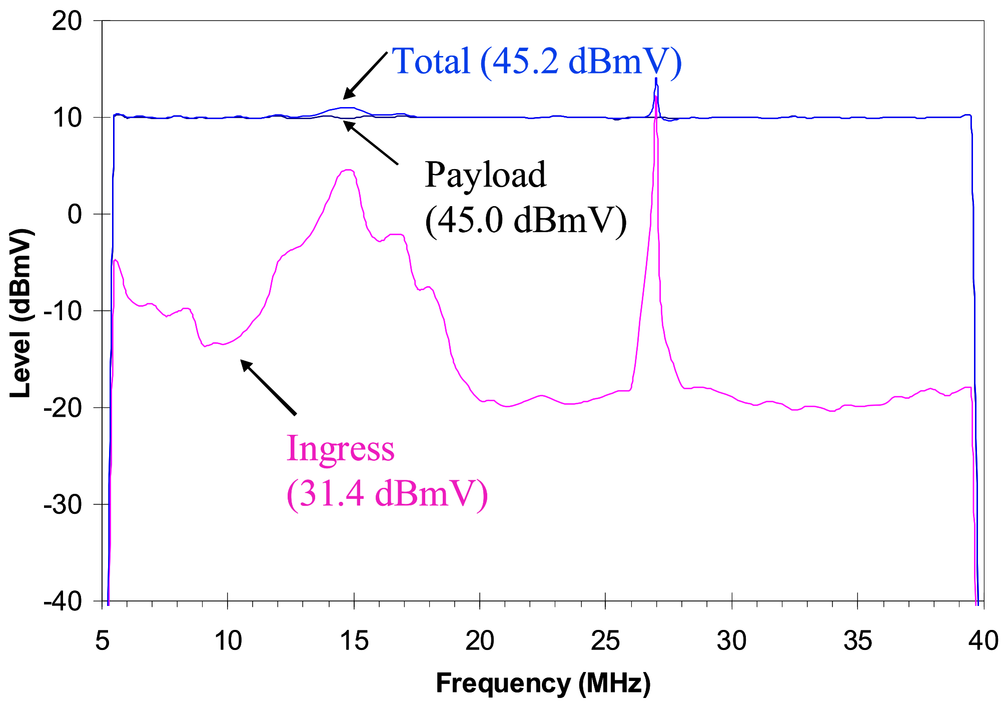

Figure 33 – Large Ingress Does Not Add to Total Power

Constant Power per Hz

Once the optimal laser operating point is determined, it's time to allocate that total power to the various services (channels) that will exist in the return path spectrum. To use the Constant Power per Hz method, you simply follow these three steps:

1. Start with the total RF power at the laser's input when it's at the optimal operating point
2. Distribute power over the entire bandwidth to be used (calculate the power per Hz)
3. Assign power to each service (channel) based on that channel's width

Consider the following example:

1. The laser's optimal RF input is +45 dBmV total power
2. Calculate the Power per Hz:

$$\text{Power per Hz} = \frac{\text{Total Power}}{\text{Total Bandwidth}}$$

   In log terms this is:

   Power per Hz = total power –10* log (total bandwidth)

   Power per Hz = 45 –10* log (35 MHz) = -30 dBmV / Hz

3. Assign power to a 3.2 MHz wide DOCSIS channel:

   Channel power = power per Hz +10*log (channel bandwidth)

   Channel power = –30 dBmV/Hz +10*log (3.2 MHz)

   Channel power = 35 dBmV

When power is assigned to a channel, the channel spacing, rather than the symbol rate or noise bandwidth, should be used to calculate the channel power.

Figure 34 shows what the spectrum might look like when fully loaded with channels. This figure is from the 1996 Emerging Technologies paper [5] and is used here to illustrate that the fundamental concepts of operating the return path have not changed. As you can see from the figure, in 1996 the vision was to use a lot of narrow band channels. Since then, most operators are using fewer wider channels. Additionally, the low end of the spectrum is best served with S-CDMA channels, rather than many narrow TDMA channels.

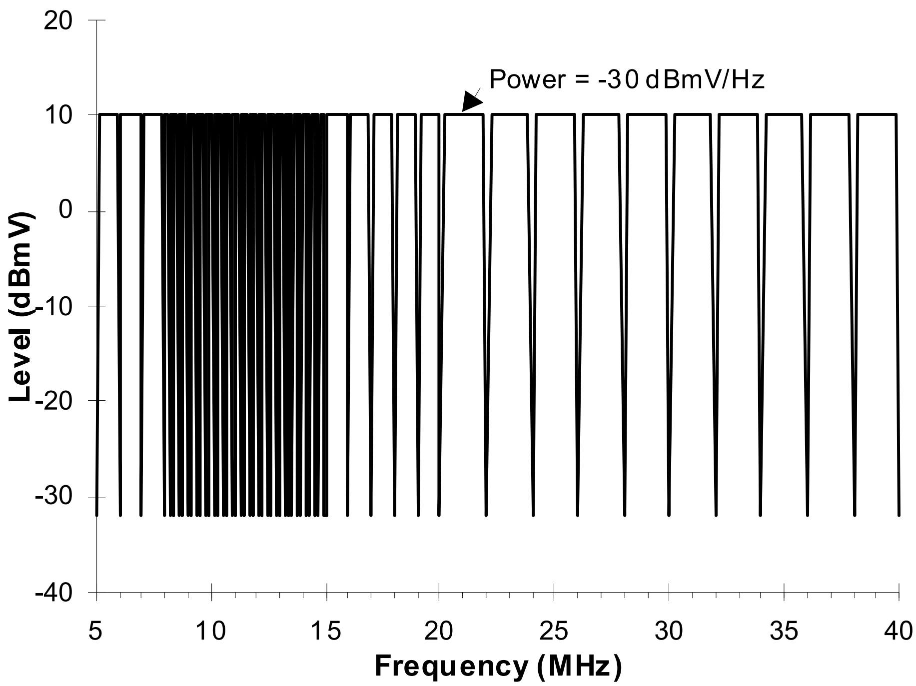

Figure 34 – Example of Return Path Spectrum

Optical Receiver Levels and Headend Combining / Splitting

Designing a headend combining and splitting network is a huge task, beyond the scope of this paper. Nevertheless, it is critical that the operator understand the net gain between the optical receiver and the input to the demodulator (generally the CMTS). The CMTS commands levels up and down based upon the level it receives. If something is changed or incorrectly aligned in the headend, the levels in the rest of the network will be incorrect. The headend gain must be set such that the CMTS receives the target level when the laser is driven by the optimal level for that channel. Details on how to accomplish this are available in several references, including [9], [10] and [11],

One final comment on headend levels: There is great confusion in the industry regarding how to adjust the optical receiver in the headend for optimal operation. Several vendors produced receivers over the past decade or two that have a poorly designed attenuator circuit. Those attenuators significantly reduce the carrier-to-noise of the signal when used. Many of these vendors are teaching operators to never use the gain adjustment control of the receivers to adjust levels. They say to always run the receiver at full output and then pad after the receiver's RF output. While this is a good recommendation for the operation of those receivers, many MSOs have been confused to believe that all receivers should be adjusted this way. ***This practice of always running at full gain should not be implemented with all vendor's products***. For instance, some receivers have very high internal gain and should not be run at full output on very short links. Additionally, the gain control does not significantly reduce carrier-to-noise and *should* be used to adjust levels. Operators must be careful to always follow manufacturer's instructions and not apply instructions from one manufacturer to another's equipment.

Plant and Home Levels

A full discussion of how to determine proper plant levels is beyond the scope of this paper. The basic concepts will be covered here. Those desiring more information should consult the very detailed discussions in [9] and [11]. The goal is to select an operating level for the plant which is as high as possible, thus forcing all upstream transmitters in the home to transmit at high levels. This provides for the maximum possible carrier-to-ingress. However, if the target plant level is too high, devices in the home will not be able to comply, resulting in alarms and, perhaps, unreliable communication.

As mentioned previously, the HFC feeder and tap network has a large range of upstream losses between the various tap ports and the amplifier housing port. Consider the following example, which is a real feeder system in operation on which DOCSIS 3.0 will be rolled-out:

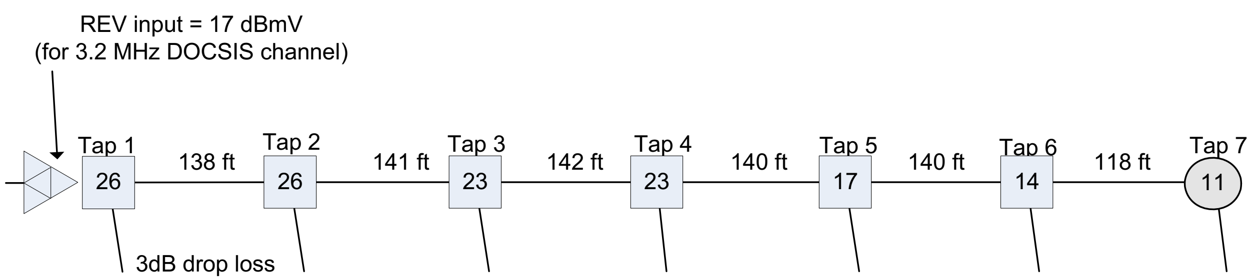

- Feeder cable = P3-750
- Maximum drop cable = 200 feet of RG-6
- Upstream frequency used for calculations = 40 MHz
- Input level to the reverse input of the fiber node:
    - 17 dBmV, referenced to 3.2 MHz DOCSIS channel
    - 20 dBmV, referenced to 6 MHz channel
    - 28 dBmV total power
- In house cable loss is not included.

Consider a typical home shown below connected to one of the taps. The levels at the taps and in the home are shown in Table 7.

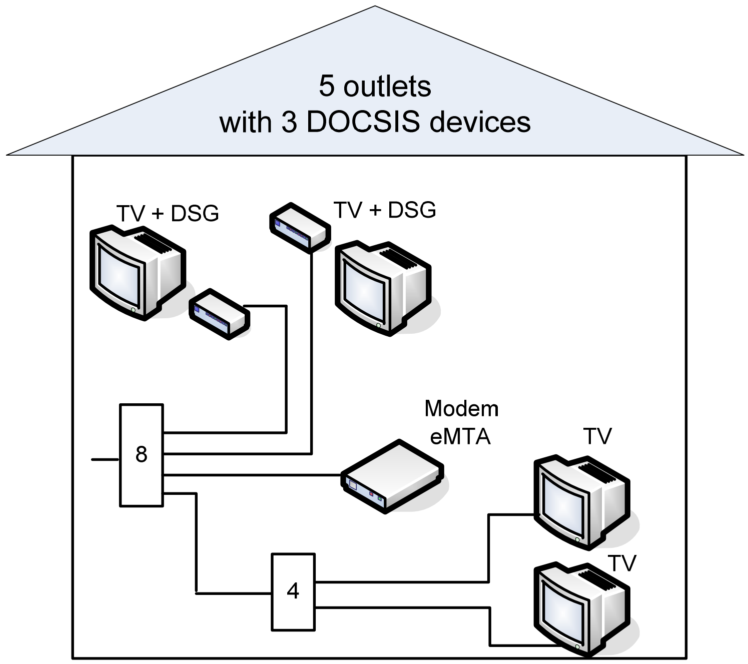

Table 7 – Upstream Transmit Levels (Maximum Tap Value = 26)

Location	1st Tap	4th Tap	7th Tap
Tap Port Level	43	42.5	**34.6**
Modem or DSG STB Level	54	53.5	**45.6**
TV or STB Deep in House Level	**58**	**57.5**	**49.6**

If we compare this to the required maximum transmit power of a cable modem (shown in Table 8), we see that the modem or DSG STB behind a 4-way splitter will work, but only in single channel mode and with very little margin for changes in plant gain over time and temperature. To work deep in the house, the modem would need to be a DOCSIS 3.0 modem running in single channel QPSK mode, which is hardly practical.

Table 8 – Minimum Pmax for DOCSIS Cable Modems

Modulation	DOCSIS 2.0	DOCSIS 3.0, Single Channel	DOCSIS 3.0, Four Channels
QPSK	58	61	**55**
16-QAM	55	58	**52**
64-QAM	**54**	**57**	**51**

To lower the required levels into the first tap, the 26 dB taps could be replaced by 23 dB taps as shown below. The results are in Table 9. The required levels at the first tap have come down 3 dB as expected, but the levels required at the 4th tap are slightly higher. This is due to the extra insertion loss of the 23 dB taps. This could be corrected by replacing the third and forth 23 dB taps with 20 dB taps, but there would be even more insertion loss with such a change. Eventually, there will not be enough excess gain in the forward path to accommodate further reductions in tap value.

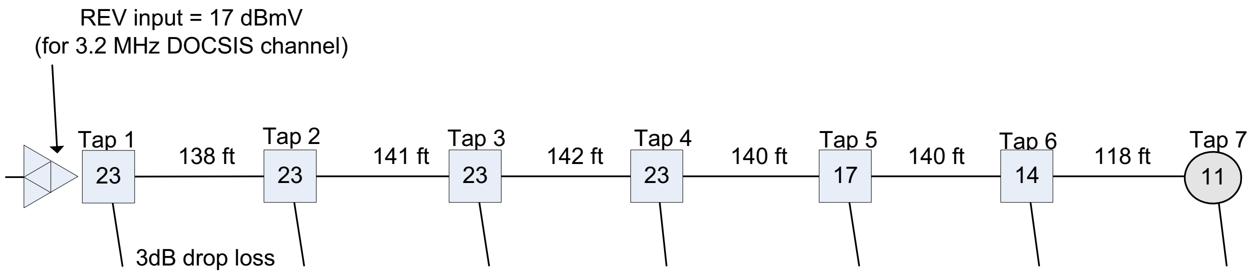

Table 9 – Upstream Transmit Levels (Maximum Tap Value = 23)

Location	1st Tap	4th Tap	7th Tap
Tap Port Level	40	42.9	**35**
Modem or DSG STB Level	51	53.9	**46**
TV or STB Deep in House Level	**55**	**57.9**	**50**

The above examples illustrate that it is very difficult to transmit wide bandwidth channels from deep in the house at high orders of modulation. Several potential solutions exist, but none are desirable:

1. Use lower power density (lower channel power) for the channels that originate deep in the house. Unless there's excess margin in the SNR design, this also necessitates using a lower order of modulation.
2. Lower the target plant levels by increasing the gain in the node. This would be done by reducing the pad in front of the laser and would lower the levels in the plant for all services. This is a nice solution, except that it will lower the carrier-to-ingress level of the entire plant proportionately to the reduction in plant levels.
3. Lower the plant and laser levels by decreasing padding in the headend. This is the least preferred approach because the laser SNR will go down in addition to the carrier-to-ingress ratio going down. Some MSOs take this approach because the change can be made without sending a tech to the node, but the resulting decrease in SNR could cause errors due to decreased operational margin.
4. Use drop amps with return path gain in the homes connected to high value taps.
5. Replace high value taps with lower value taps. This will lower the level to the problem home, but will increase the insertion loss in the plant.
6. Use dual-drops to homes with a lot of cable outlets. This allows splitting loss to be decreased in the home without requiring an active-return drop amp.

In most homes, three DOCSIS devices can be used (off the ports of a 4-way splitter). To enable more devices, it is likely that MSOs will choose some combination of 1, 2 and 4. As node sizes get smaller and ingress decreases, it should be possible to lower plant levels. There will always be some problem homes, and drop amps with active returns could be used in those homes.

Implications of Filling the Return Spectrum

As MSOs take steps to fill the RF spectrum with channels, it will become increasingly difficult to troubleshoot the return path and to detect ingress. Once there is no more free space in the return spectrum, technicians will need to rely on advanced spectrum monitoring techniques that use digital analysis to determine the types of impairments that are under the channel. Other advanced techniques can look at the channel between TDMA bursts. None of these techniques will be possible with a standard spectrum analyzer. The technicians will need to be trained to use new tools.

As the downstream signals are converted from analog video to digital channels, it will also become more difficult to detect and troubleshoot common path distortion. The characteristic 6 MHz separation between distortions in the return path band will no longer exist, because the digital channels are spectrally flat. Once the downstream is 100% digital, it will be nearly impossible to detect whether the impairments in the return path are caused by ingress or common path. This difficultly is compounded by the fact that the return path will eventually be filled, so the only way to discern what's happening will be with advanced DSP techniques.

Operators will need to be more diligent than ever at making sure there is no downstream leakage. That is still the best way to keep the plant tight and minimize any ingress.

Implications of Using Upstream Channel Bonding

If the return system was designed to handle a full load of digital channels, upstream channel bonding has no impact on the return path plant or lasers. The RF amplifiers or lasers are not impacted in any way when multiple upstream channels come from one modem vs. coming from multiple different modems. However, if the return path has not been designed and aligned to carry a full band of digital channels, the operator must follow the guidelines in the previous sections to properly design and align the system so that the spectrum can be filled with channels. If you've followed the principles that were first outlined in 1996, you're all ready to go!

Bonding does, however, impact levels inside the house. This is because the power available per channel decreases when a modem is transmitting in bonded mode. As discussed in the previous section, bonded modems will only work if they are connected to a 4-way splitter or less. If more level is needed, a drop amp with return path gain must be used, or target plant levels must be lowered.

III. MER, CER, PER Performance Evaluation

A. Upstream Samples – Pre-Assessment Prior to Launch

At this stage of the evolution of the upstream, there is a common understanding that the upstream channel can present a diverse set of highly variable characteristics, both spatially (across different parts of the same plant or different plants) and in time on any given return path. While the variation is wide, it is also the case that most returns have their own signature. It may have diurnal variations, but will tend towards some predictable ranges and types of disturbance. These can change slowly over time, as many of the disturbances come from variable sources – subscriber home wiring "engineering," and gradual plant aging resulting in corrosion of connectors, loose seals in actives and passives, and coaxial drop cuts, bends, and general environmental damage. The many-to-one nature of the upstream, the lack of control at the endpoints, the overall funneling mechanism, and the band of operation are the primary reasons for the innate struggle to maintain a clean spectrum. Nonetheless, the DOCSIS upstream was designed from a system engineering perspective with these problems in mind. Robust modulations, wide dynamic ranges, powerful FEC, and a sophisticated F/TDMA protocol allowed essentially immediate success. The robust modulation, modest data rates (lower bandwidths), flexible choice of empty spectrum, and access to virtually as much laser drive as desired all played key roles in this success. For DOCSIS 3.0, however, ***all of these advantages are disappearing or are already gone***. As such, it is imperative now to more fully understand the RF nature of the return channel.

This section identifies the types of things commonly observed on returns, the vast majority of them "working" returns. However, they fall into the category of likely needing help if they are to take on a robust DOCSIS 3.0 and be suitable for additional channels. In a subsequent section, we methodically introduce impairments such as these displayed in a controlled fashion to quantify upstream performance, at least for a small subset of profile variables that allows a reasonable set of permutations.

This short refresher of upstream headache-makers includes a short description of the troubling characteristics, and some of their resulting impacts. In the case of spectrum plots, an average and a peak hold curve are shown, representing coarsely both the typical (static) and worst case (dynamic) channel conditions. A CMTS that reports the spectrum quality on a granular time interval basis for further post-processing is a valuable tool, as impulsive noise can best be captured in this way and, as we will see, can be the dominant performance-limiting impairment. Note that on a spectrum analyzer in classical averaging mode, the "max hold" function highlighted below can help to identify its presence, but little else about it other than its bandwidth and overall peak spectral density during the course of the measurement. Periodic burst noise can appear as static in an averaging frequency domain measurement.

In the captures below we note the following:

70 MHz Span (other SA parameters shown)
Blue = Max Hold, Yellow = Average
DOCSIS signal mid-to-high band (varying) 3.2 MHz wide

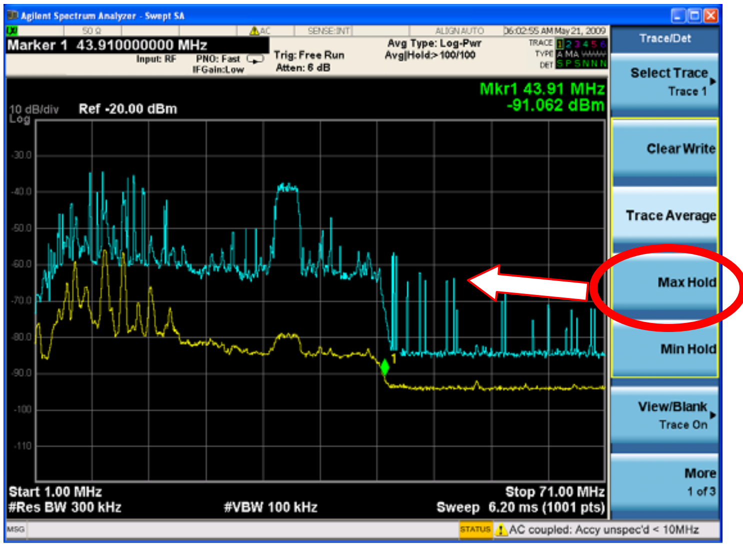

- Narrowband interference
- High pk-avg, impulsive
- Possible laser overload (wideband components)
- Likely supports 16-QAM
- Insufficient for 64-QAM

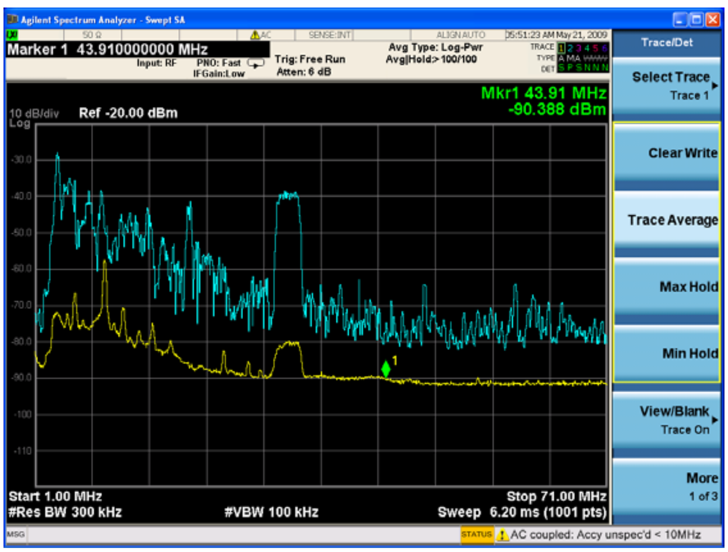

- High level low frequency impulse
- Otherwise good spectrum & SNR in carrier region
- ATDMA not well-suited to low end spectrum
- Wideband channel near 20 MHz would enter poor 15-20 MHz spectrum

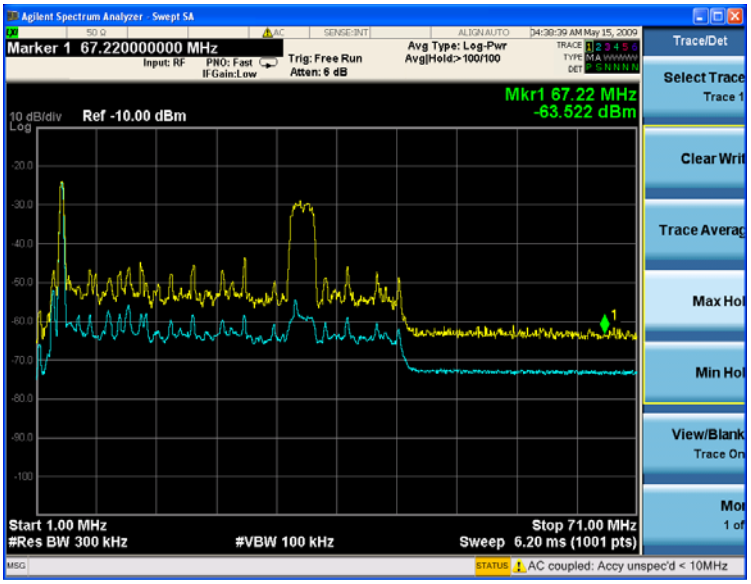

- Shortwave Interferer (high level)
- Evidence of CPD tones – tolerable but can degrade and associate with a microreflection
- Noise floor of FP return, stable
- Adequate for 16-QAM, poor conditions for 64-QAM

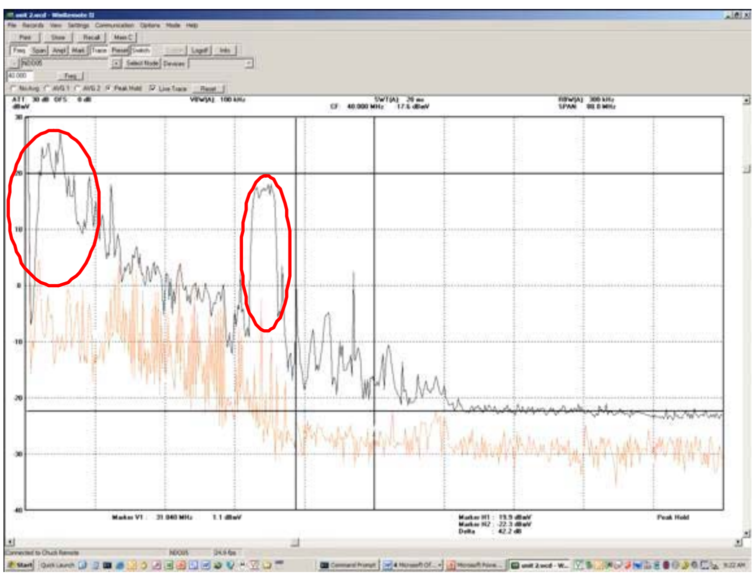

- Impulsive interference power > signal power
- Power Loading Issue – headroom needed for new carriers
- ATDMA not well-suited to low end spectrum

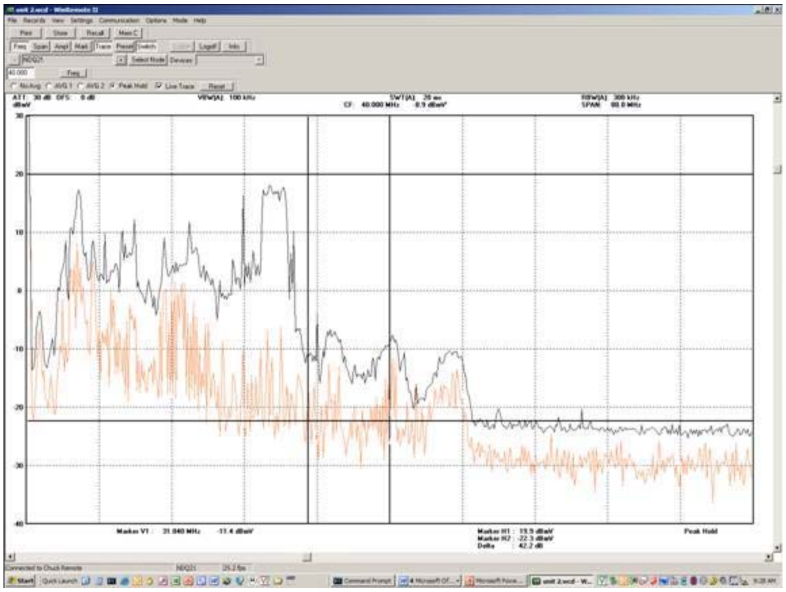

- Impulsive noise into DOCSIS Band at ~ -15 dBc
- 16-QAM supported but likely with FEC working hard
- Insufficient for 64-QAM

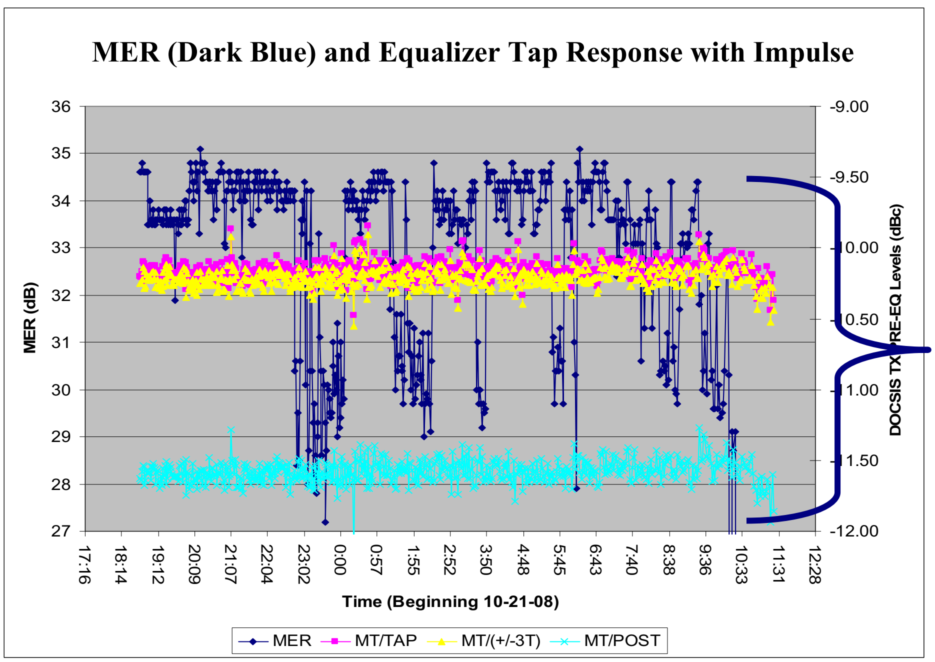

- Impulse noise – impact on MER
- 8 dB pk-pk
- More than an order of modulation

B. Burst Correction Mechanisms – R-S FEC and Upstream Interleaving

It may be hard to fathom, but Reed-Solomon forward error correction turns 50 years old next year. It may seem unusual, but if so it is only because the linear algebraic roots preceded its cost-effective implementation, in particular for real-time processing. But, in fact, given that Shannon's information theory bible was published in 1948, perhaps it starts to make sense as that touched off a race to the capacity holy grail. Shannon sent communications engineers on a hunt for the grand coding scheme that his mathematics said had to exist and deliver the predicted "capacity" at arbitrarily low error rates. Reed and Solomon's step forward in the search for better codes was this subset of what are known as BCH codes. Only the ability to implement the processing stood in the way of taking advantage of high rate (i.e. efficient) codes with excellent *distance* properties (high gain). Code rate is simply the ratio of data symbols to data+parity (overhead) symbols in a code, and the higher the better from a use of the channel standpoint. Distance is a linear algebraic measure describing how "far" in code space different valid codewords are, and thus how difficult it is to make mistakes between them. The space is multi-dimensional, so picturing it is difficult, but good distance properties lead to better performance. R-S codes enjoy a property known as "minimum distance separable," which is another way of saying they can be made as far apart as algebraically possible given the amount of parity check overhead. That is all of the linear algebra we will introduce into this discussion, but there are (literally) books devoted solely to R-S codes.

The properties above essentially translate to meaning that lots of dBs of link budget can be gained using R-S coding as part of the system solution. For example, the 28 dB of 64-QAM previously noted as a threshold can drop to, say 24 dB, depending on the R-S code selected. So, in exchange for some added overhead, or minor bandwidth inefficiency, and a few added chips or processing in a receive modem chip, the system designer gets to ease the burden on his transmitters, receivers, etc, in his link budget. Often, a few extra pennies or dollars in a receiver's IC is a much more cost-effective solution than a couple dB of additional transmit power or noise figure. It is the processing breakthroughs that allowed implementation and spurred widespread use. No convincing that they were excellent codes was required. Figure 35 [12] shows a sample of the gain employing a RS(255,x) code – the maximum size of the DOCSIS codeword – across a range of value of "t," which represents how many symbols in the codeword can be received in error and still be able to be corrected (t = 1 through 16 are all valid DOCSIS modes). Note that the channel symbol error is not one in the same as a QAM symbol, but represents a R-S byte in the codeword. However, clearly there is a correlation between 4-bit or 8-bit QAM symbol errors (16-QAM and 64-QAM) and an 8-bit R-S symbol being in error.

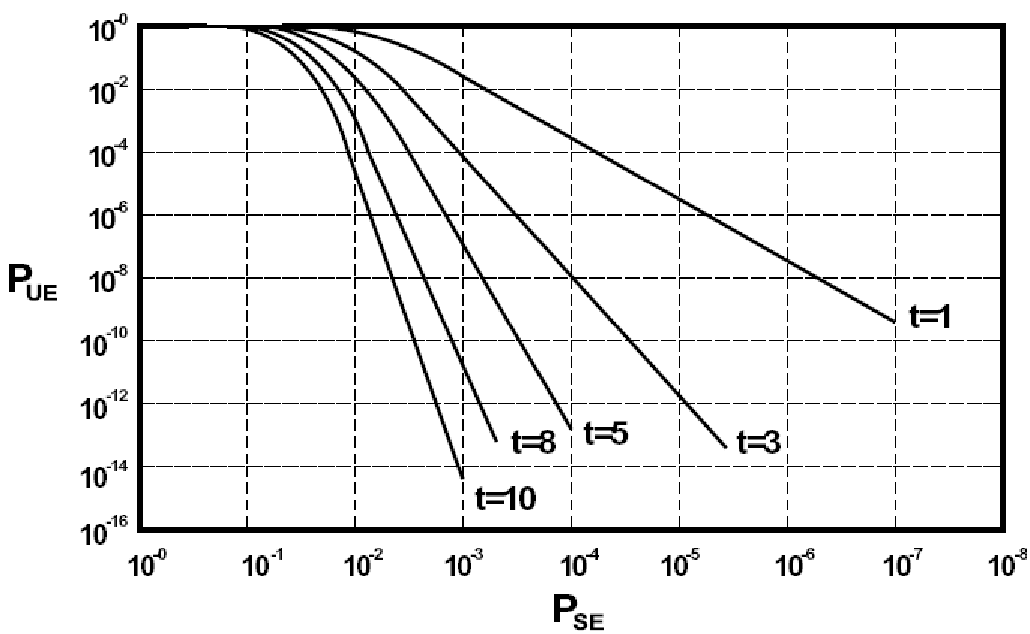

Figure 35 – Probability of Uncorrectable Codeword vs Channel Symbol Error

Note the steepness of the slope for the lowest rate, highest-correcting, codes (t = 8, 10). The figure is in some ways misleading in this regard. When plotted against SNR, as symbol error rate (SER) and BER curves often are, the P_{ue} slope would appear even steeper. Indeed, a P_{se} range of 1e-3 through 1e-7 on Figure 35 covers just about 4.5 dB of SNR range. The above curve is important to keep in mind as we test impairments in an actual HFC environment described later, including impulse noise capable of wiping out many consecutive symbols. As the above chart shows, symbol errors rates (again – not exactly the same QAM errors but strongly related) as low as 1e-2 through 1e-4 see orders of magnitude improvement. Furthermore, the above R-S code in terms of codeword size is very close to what is tested, and the code rate is even lower in the test case (t is larger).

In addition to the properties that make it an excellent code from a coding gain standpoint, Reed-Solomon FEC has a couple of very valuable properties for the HFC upstream. They can be based on non-binary algebra (contributing to the processing intensity), they are systemic block codes – meaning the actual message stays intact as is in creation of the codeword, and they are powerful against burst errors, a consequence of the non-binary property. By using non-binary algebra, one R-S symbol in the codeword instead of being one bit can be a digital byte as it is for DOCSIS, obviously a convenient unit of digital processing. The code corrects, depending on the code rate, up to 16 symbol errors in a codeword size of up to 255 bytes. That means 16 *bytes* may be in error – a bit of the byte or the whole byte itself in error is simply one R-S symbol error. Thus, this creates an inherent burst error capability, in this case by a factor of eight or more. Consecutive bytes can also be in error if the total remains below the “t” value. Most codes deliver high performance in “random” error environments, such as AWGN, but can get overwhelmed by bursts that demolish one codeword beyond recognition. Reed-Solomon codes have excellent gain properties in a random error environment, but have the added advantage of increased ruggedness to burst errors, with the burst tolerance required being an element of the code design that can be matched to the channel characteristics.

While we have lauded the burst-correcting strength of the R-S FEC, it is still the case that the burst robustness, as we simultaneously increase the symbol rate and deal with different swaths of return spectrum, still may be overrun by a determined garage door opener or hair dryer. Thus, a byte interleave option is available to increase this further – “byte” because the R-S symbols are bytes. Burst robustness is scaled by a factor of the interleaver depth, the maximum of which is a function of the codeword size, and bounded by the size of the storage block of 2048 bytes. The concept is quite straightforward – the transmission may suffer a burst event but once rearranged (de-interleaved) at the receiver prior to the R-S decoder, the burst of errors is spread out across multiple codewords, increasing the probability that they will individually be correctable and no single codeword will be at risk for the disturbance taking out more symbols in the codeword than can be recovered. The process attempts to create a random error correcting environment from a burst-generating one. The interleaver function includes a dynamic mode, whereby the interleaving block size is structured to for uniformity of depth by choosing depth based on the packet size, so the burst robustness is preserved across the packet.

In the data to come, link statistics available from the DOCSIS upstream receiver as a function of a set of channel impairments will be analyzed. These are tabulated along with the output packet error rate that they correspond to using long and short packets, using *fixed* time (as opposed to fixed packet count or codeword count). In all cases, the upstream interleaver is turned off, and maximum R-S correction turned on. Having powerful FEC shapes the transmission error rate curves drastically, such that when combined with the many impairment permutations of the upstream, numerical characterization is the best reality check. In general, as discussed, long block codes generate tremendously steep error rate curves compared to the classic symbol or bit error rate slopes. Typically, at operating points of interest, the error rates drop orders of magnitude, creating substantial extra margin in exchange for that sensitivity. Symbol and bit error curves, valuable for the physical channel characterization, are no longer individually as meaningful when symbol-based block coding is used when measuring output packets delivered. Instead, internal decoding metrics include the codewords (blocks) with no errors, with errors that are corrected, and blocks with errors that cannot be corrected. These can be converted to percentages and probabilities.

C. Impairment Performance Catalog

The noise contributions of the optics can be modeled as an AWGN source. While useful, it is often the combination of other impairments that dominate the noise performance, and particularly as the lower end of the spectrum is used. The contributions of these other disturbances on the DOCSIS channel are captured by the CMTS, where they are displayed as SNR, which as indicated is actually reporting MER, a term that includes all of the contributors measured at the receiver that create non-ideal constellation reception. It is by calculating error from this constellation that MER is derived.

There are four points to emphasize about the MER term:

1) MER is always worse than actual link SNR because other impairments are added to the AWGN contribution.

2) It is ***not possible*** to extract from an MER measurement alone the breakdown of the contributing impairments. However, with an observable constellation there is information available through the shape of the distorted pattern that can help identify the main causes.

3) MER is an averaging measurement over some defined period or observation interval. That period of time may or may not contain an impulse, so ***an impulse noise environment is not easily captured*** in an MER measurement.

4) Finally, and most importantly, unlike SNR, ***there is NOT a unique, one-to-one relationship between MER (reported as SNR by the CMTS) and BER***. It can be close, and it may negligibly differ, but it is not a unique relationship. The same MER can result in two different BERs; the same BER can result from two different MERs.

The last fact is simply the result of the MER being an aggregate of different impairments. As an averaging error magnitude measurement, the same MER can be reported for two different combinations of AWGN + ingress + impulse. However, each of these effect error rate degradation differently. If the dominant impairment is AWGN, it will have a different BER result than if the dominant impairment is ingress or impulse. Another way of looking at it is that an SNR of 28 dB for 64-QAM may yield the 1e-8 BER, but a 28 dB S/I, for a narrowband interferer, will yield something different (actually zero errors in the absence of other additive impairments for a static CW carrier). All of the above is true whether or not we are considering uncoded or coded transmissions.

Laboratory Measurements

Because of this non-unique relationship between MER and BER, a set of combined measurements was created to offer insight into how much variability can occur over a fixed MER range, and thus provide a level of awareness and guidance in the diagnosis of possible link error conditions. A Motorola CMTS was connected to over 40 DOCSIS 2.0 modems across common vendors, with ten modems having input stimuli from a traffic generating test source, using a common and uniform profile. The testing was done with the following conditions and parameters:

Plant & Signal

- 20 km DFB link
- N+6 cascade
- 36.5 DOCSIS carrier frequency
- 64-QAM @ 5.12 Msps
    - o 16-QAM subset @ 5.12 Msps to draw comparisons
- 1518-byte packets
- FEC: K=219, T=16
- No Interleaving
- Ingress canceller ON
- Pre-Equalization ON

Note that the chosen frequency does not violate the prior criteria for cascade depth (36.8 MHz @ N+9). However, it does place the signal into an area where it will experience significant amplitude and group delay roll-off, and thus exercise the equalizer in the face of the impairments below. Including the added rule of thumb of 200 kHz/active, the threshold for this cascade would be 37.4 MHz.

Also, a set of data for short packets (128 bytes) was also taken, but not included here for brevity. The short packets clearly create some advantages with respect to the amount of data is lost when errors occur, and this reflected in those results. Long packets thus create a conservative example.

Impairments

AWGN Noise

- SNR = 35 dB
- SNR = 27 dB

Static Ingress – Baseline Levels

- Single CW Carrier @ -10 dBc
- Three CW Carriers uniformly spread @ -15 dBc, -20 dBc, -25 dBc
- Single FM Carrier @ -10 dBc
- Three FM Carriers uniformly spread @ -15 dBc, , -20 dBc, -25 dBc
- Mid-band and +/- 1.5 MHz used for three; -1.5 MHz used for single

Impulse – Baseline Levels

- 4 usec duration AWGN pulse @ 100 Hz @ -10 dBc, -5 dBc
- 10 usec duration AWGN pulse @ 1 kHz @ -10 dBc, -15 dBc

Note the following in describing the impairment levels and types chosen:

- The three primary contributors to return path impairment are used: AWGN, impulsive noise, and narrowband interference. Use of the HFC link allows the other secondary contributors (frequency response distortion, clipping) to be part of the results at a nominal level.

- The 27 dB AWGN case begins in an uncomfortable area for 64-QAM symbols. It represents a 3e-7 hard decision, uncoded, region of operation, and is also just 3 dB from a 1e-4 QAM symbol detection range. The 35 dB SNR case supports both comfortably.

- Both ingress and impulse levels were identified through baseline testing to align the range across good performance to poor performance, looking for the break point thresholds individually and eventually in combination.

- The baseline levels above were extended in an effort to reach threshold points, again to offer a glimpse of when particular impairment combinations cause link closure issues. This included narrowband interference with no impulse, and fixing interference while decreasing impulse *amplitudes* (periods and duty were not varied beyond these two conditions).

- The purpose of FM ingress was to provide a broader bandwidth, continuous-looking spectrum interferer at the same amplitude to challenge the ingress filter. The BW was approximately 20 kHz.

- The impulse "dBc" levels are as measured by the AWGN floor when gating of the noise is turned off. In other words, they are the static SNR if the noise were on all of the time.

- A shorter duration of impulse was preferred, but due to equipment limitations, 4 usec was used. It represents about 20 QAM symbols at this symbol rate, and thus 80 bits/10 Bytes in a 16-QAM transmission and 120 bits/12 Bytes in a 64-QAM transmission – smaller than T=16 bytes of correction, but a substantial percentage of the maximum correctable words. It represents a 4e-4 duration of average symbol count should all effected symbols be wrong and all others be correct. The long duration bursts were meant to consume the R-S correction capability from a time domain perspective, and represents a 1e-2 duration of average symbol count with all symbols during the burst being incorrect and all of the rest correct.

- As previously discussed, poor symbol error rates can become acceptable, or at least vastly improved, based on the R-S curves shown in Figure 35. However, the use of impulsive effects brings into play the burst-versus-random correcting aspect of R-S codes – they have advantages, but are not bulletproof.

- While the support S-CDMA brings to solving the impulse and combined impairment problem is not tested here, we will summarize some recent, encouraging results with S-CDMA in the section that follows.

These are significant impairments, and in particular for the impulse dynamics – the situation most dynamic on the return path and most unpredictable, link-to-link. Nonetheless, the combination represents the types of things that are observed in real plants, and in particular as frequencies move below 20 MHz, and more so even when moving below 15 MHz. They are more unlikely to all be encountered in the "sweet spot" between 20-35 MHz, but, again, as the prior figures of plant conditions show, such characteristics can be found to extend into higher frequency bands on a non-negligible number of nodes in an average Headend, perhaps 10-20% as observed in the field. In this case, they will allow us to see where system breakdowns can be quantified, and point out where the 16-QAM margin and robustness is evident. In combination, it is difficult in principle for a receiver to mitigate simultaneously both narrowband and wideband (impulse) interference.

In terms of measurements taken, while QAM symbol error rate (SER) or BER are excellent physical layer representations of what is happening on the wire, the fact is the QAM modulations are supplemented with FEC. What is available to measure are FEC statistics, and these are calculated by the receiver and made available to the CMTS for reporting. Conversion to BER may offer insights into what is happening on the coaxial line, but more valuable for operators is the delivery of packets, which of course take advantage of the FEC. Since the service is delivered out of the CM or CMTS as packets, we use a packet-error test with the DOCSIS link in between to ensure we are capturing values that are representative of reasonable packet error rates (PER). In this case, the PERs will be on the high side because we are looking to watch links degrade and break. An advantage of high error counts for these relatively short test sequences is that they increase the statistical validity of the results. Further, more granular, testing will ensue, but the data here gives some snapshot guidelines of combined impairment thresholds for 64-QAM.

Since PER is not available as a measurable statistic on an operating system, we also compare PERs to measured correctable and uncorrectable codewords – something operators do have access to on a per-modem basis and per-CMTS port basis, as well as MER. As we have emphasized, and as will be seen, MER is not a unique value relative to the error rates. As an average value, it is likely to not represent impulsive noise environments well, and this will be demonstrated. The data provides the ability to draw some correlations between FEC errors and packet errors across different impairment combinations.

Note that the FEC statistics are gathered over a given measurement interval using before(time = t1) and after (time = t2) absolute counts of uncorrectable codewords, errored but corrected codewaords, and unerrored codewords tabulated. The corrected and uncorrected codeword percentages are found simply as follows:

Total Codewords (TCW) = Unerrored + Errored Corrected + Errored Uncorrectable

Then,

Correctable % = [Errored Corrected(t2) – Errored Corrected(t1)] / [TCW(t2) – TCW(t1)]

Uncorrectable % = [Errored Uncorrectable(t2) – Errored Uncorrectable(t1)] / [TCW(t2) – TCW(t1)]

Tables 10-13 summarize some of the key data taken as impairments were swept across ranges in different combinations. The data represents a subset of measurements, and further results are available by contacting the authors. Also, obviously, the measurements represent a subset of the possible permutations of impairments. For each table, the metrics shown for which the impairments are parametrically varied are MER, FEC statistics UCER and CCER – uncorrectable codeword error rate and corrected codeword error rate – and packet error rate, PER. The MER reported is a port average as reported by the CMTS, while the FEC statistics are from queries of the received data from ten modems actually ingesting and transmitting the traffic payload. The PER is calculated using the closed loop test equipment external to the DOCSIS link. Along with each table are some key take-aways from reviewing the tabular results.

Table 10 – Ingress-only Thresholds for 64-QAM

1518-Byte Packets			
Noise Floor = 27 dB	MER	CCER/UCER %	PER
None	26.90	0 / 0	0.00%
CW Interference			
1x @ -5 dBc	26.00	8.6 / 0.018	0.10%
1x @ -10 dBc	26.20	7.02 / 0.00176	0.00%
3x @ -10 dBc/tone	26.00	9.5 / 0.08	0.50%
3x @ -15 dBc/tone	26.10	9.5 / 0.0099	0.06%
3x @ -20 dBc/tone	26.10	8.2 / 0.00137	0.00%
FM Modulated (20 kHz BW)			
1x @ -10 dBc	25.80	15.66 / 0.33166	1.00%
1x @ -15 dBc	26.40	6.2 / 0.0008	0.04%
3x @ -15 dBc/tone	25.50	19.48 / 0.639	2.00%
3x @ -20 dBc/tone	26.00	10.68 / 0.00855	0.03%
Noise Floor = 35 dB	MER	CCER/UCER	PER
None	32.60	0 / 0	0.00%
CW Interference			
1x @ +5 dBc	28.50	0.24 / 0.09	0.50%
1x @ 0 dBc	30.00	0.006 / 0.013	0.00%
1x @ -10 dBc	31.40	0 / 0.0065	0.00%
3x @ -10 dBc/tone	31.20	0.002 / 0	0.00%
3x @ -15 dBc/tone	31.50	0 / 0	0.00%
FM Modulated (20 kHz BW)			
1x @ -5 dBc	30.60	0.004 / 0	0.04%
1x @ -10 dBc	31.10	0.003 / 0	0.00%
3x @ -10 dBc/tone	30.00	0.01 / 0.0009	0.08%
3x @ -15 dBc/tone	30.80	0 / 0	0.00%

Some items of note:

- Note baseline MERs recorded with no impairment (i.e. "None") can be referenced for all subsequent discussion.
- For the AWGN floor of 27 dB, the FEC is clearly working much harder, as the CCER % at 5-15% of the receive codewords getting cleaned up indicates. This is consistent with 64-QAM being imposed on by the 27 dB SNR limit, which creates countable 64-QAM symbol errors that translate to R-S symbols in error within a codeword.
- Without impulse noise, the only real pain in this set of data from a PER standpoint is for the 27 dB case, and in particular with FM ingress. The combination of modulated ingress and the high noise environment challenges the ability of the ingress filtering to deeply suppress the interference, as a result of the two-fold randomness that must be adapted to.
- It is not as easily apparent here, but we will see a steady consistency between 1x@-10 dBc interference and 3x@-15 dBc interference. Thus total interference power, or signal-to-interference, for a small number of interferers, with relatively wide band spread, appears to be a reasonable way to consider interference power.
- Note the threshold for 35 dB SNR and a single CW interferer to beginning to count packet errors of ***+5 dBc***. This is impressive and a natural method of evaluating ingress protection, but also insufficient to rely on exclusively. Note also that the CW and FM cases differ by 5-10 dB in this area, the modulated case as expected being more difficult to handle.

Table 11 – Ingress Thresholds with Fixed Impulse for 64-QAM

	Impulse (dBc-pk) - Gated AWGN @ 10 MHz BW								
1518-Byte Packets	4 usec @ 100 Hz						10 usec @ 1 kHz		
Noise Floor = 27 dB	-10			-5			-15		
Ingress (dBc)	MER	CCER/UCER	PER	MER	CCER/UCER	PER	MER	CCER/UCER	PER
None	26.70	3.29 / 0.001	0.01%	26.70	3.59 / 0.0866	0.50%	26.60	9.86 / 0.0046	0.04%
CW Interference									
3x @ -15 dBc/tone	25.90	11.84 / 0.286	0.90%	26.10	8.88 / 0.59	2.50%	25.70	16.9 / 1.21	8.00%
3x @ -20 dBc/tone	26.10	10.71 / 0.084	0.50%	26.10	10.33 / 0.42	2.50%	25.90	16.14 / 0.51	3.00%
3x @ -25 dBc/tone	26.20	7.92 / 0.02	0.10%	26.20	8.06 / 0.227	1.50%	25.90	16.23 / 0.14	0.50%
3x @ -35 dBc/tone				26.30	5.8 / 0.13	0.70%	26.00	14.4 / 0.016	0.20%
FM (20 kHz BW)									
3x @ -15 dBc/tone	25.40	19.4 / 1.29	5.00%	25.40	19 / 1.53	6.00%	25.30	23.1 / 5.56	18.00%
3x @ -20 dBc/tone	25.90	14.15 / 0.16	1.00%	26.20	8.12 / 0.27	1.70%	25.70	13.8 / 8.52	5.00%
3x @ -25 dBc/tone	26.20	8.46 / 0.021	0.20%	25.80	12.8 / 0.47	0.60%	25.90	16.0 / 0.18	1.50%
3x @ -30 dBc/tone				26.40	5.2 / 0.11	0.70%	26.20	10.7 / 0.016	0.05%

	Impulse (dBc-pk) - Gated AWGN @ 10 MHz BW								
1518-Byte Packets	4 usec @ 100 Hz						10 usec @ 1 kHz		
Noise Floor = 35 dB	-10			-5			-15		
Ingress (dBc)	MER	CCER/UCER	PER	MER	CCER/UCER	PER	MER	CCER/UCER	PER
None	32.30	0.69 / 0	0.00%	32.30	0.68 / 0.04	0.30%	31.90	6.5 / 0.0007	0.00%
CW Interference									
3x @ -15 dBc/tone	31.40	0.5 / 0.24	1.50%	31.50	0.33 / 0.57	2.80%	31.00	6.36 / 0.98	5.00%
3x @ -20 dBc/tone	31.60	0.6 / 0.19	1.00%	31.40	0.38 / 0.46	2.20%	30.80	6.63 / 0.47	2.50%
3x @ -25 dBc/tone	31.70	0.66 / 0.056	0.40%	31.60	0.47 / 0.29	1.70%	31.20	6.66 / 0.14	0.80%
3x @ -30 dBc/tone							31.50	6.75 / 0.027	0.20%
3x @ -35 dBc/tone				31.60	0.6 / 0.08	0.50%			
FM (20 kHz BW)									
3x @ -15 dBc/tone	30.60	0.38 / 0.55	2.80%	29.90	0.2 / 0.85	3.70%	29.80	4.59 / 4.44	21.00%
3x @ -20 dBc/tone	31.10	0.47 / 0.35	1.70%	31.00	0.26 / 0.82	3.50%	30.30	5.66 / 2.1	12.00%
3x @ -25 dBc/tone	31.50	0.64 / 0.1	0.70%	31.40	0.43 / 0.34	2.10%	31.00	6.76 / 0.28	1.50%
3x @ -30 dBc/tone	31.50	0.002 / 0	0.00%				31.00	6.9 / 0.037	0.30%
3x @ -35 dBc/tone				31.80	0.65 / 0.11	0.60%			

Some items of note:

- As before, for the AWGN floor of 27 dB, the FEC is working hard
- Note the row "none," showing the measurement metrics for no static interference. It is evident that the level of the impulse-only events can count errors at the physical layer and packet layer. This is expected, given the prior discussion on the impulse peak levels in wideband SNR terms, relative to the requirements previously stated for the modulation.
- There are some observable benefits as the ingress levels are dropped with a given level of impulse. However:
  - While 35 dB yields generally lower error rates, there are not tremendous differences between the two AWGN levels – UCERs, which contribute to PER, are the same order of magnitude. PERs are also similar.
  - This points to the dominance of impulse in this case relative to the now secondary AWGN addition.
  - This is a common practical case - confusion as to why sound MER values give the same level of performance issues as poor MER values. As has previous been discussed, the MER is not unique to error rate – the same error rate can come of two different MERs, and the same MER can yield different error performance.
  - The impulse essentially sets a floor to the error rates, as the prior discussion on the duration and levels of this impairment imply could occur.

Table 12 – Impulse Thresholds with Fixed Ingress for 64-QAM

1518-Byte Packets	CW Ingress Characteristic					
Noise Floor = 27 dB	1x @ -10 dBc			3x @ -15 dBc		
Impulse Type	MER	CCER/UCER	PER	MER	CCER/UCER	PER
4 usec @ 100 Hz						
-10 dBc	26.30	8.88 / 0.1442	0.50%	25.90	11.84 / 0.286	0.90%
-15 dBc	26.30	5.8 / 0.01	0.05%	26.10	8.1 / 0.022	0.15%
10 usec @ 1 kHz						
-15 dBc	26.10	14.46 / 0.807	3.00%	25.90	16.14 / 0.51	3.00%
-20 dBc	26.30	7.8 / 0.006	0.04%	26.10	10.4 / 0.019	0.10%
Noise Floor = 35 dB						
Impulse Type	MER	CCER/UCER	PER	MER	CCER/UCER	PER
4 usec @ 100 Hz						
-10 dBc	31.30	0.49 / 0.48	1.40%	31.40	0.5 / 0.24	1.50%
-15 dBc	31.20	0.5 / 0.05	0.30%	31.20	0.5 / 0.05	0.30%
10 usec @ 1 kHz						
-15 dBc	30.60	6.03 / 2.12	6.00%	31.00	6.36 / 0.98	5.00%
-20 dBc	31.10	0.67 / 0.006	0.08%	31.30	0.7 / 0.004	0.08%

1518-Byte Packets	FM Ingress Characteristic					
Noise Floor = 27 dB	1x @ -10 dBc			3x @ -15 dBc		
Impulse Characteristic (dBc)	MER	CCER/UCER	PER	MER	CCER/UCER	PER
4 usec @ 100 Hz						
-10 dBc	25.60	23.64 / 2.55	6.00%	25.40	19.4 / 1.29	5.00%
-25 dBc	25.90	13.8 / 0.4	0.70%			
-30 dBc				25.70	14.3 / 0.31	1.50%
10 usec @ 1 kHz						
-15 dBc	25.10	18.96 / 8.0	20.00%	25.30	23.1 / 5.56	18.00%
-30 dBc	26.10	13.2 / 0.25	0.60%	25.60	15.4 / 0.31	1.50%
Noise Floor = 35 dB						
Impulse Characteristic (dBc)	MER	CCER/UCER	PER	MER	CCER/UCER	PER
4 usec @ 100 Hz						
-10 dBc	31.00	0.0014 / 0.0035	0.10%	30.60	0.38 / 0.55	2.80%
-25 dBc	30.90	0.06 / 0	0.08%	30.50	0.06 / 0.007	0.10%
10 usec @ 1 kHz						
-15 dBc	28.30	5.05 / 5.51	15.00%	29.80	4.59 / 4.44	21.00%
-20 dBc	30.70	0.80 / 0.06	0.20%			
-25 dBc				30.60	0.016 / 0	0.00%

Some items of note:

- FEC continues to work hard for the AWGN floor of 27 dB
- The drop in UCER and PER is precipitous with the drop in impulse level, and ultimately drops into manageable error rates with continued decrease of the peak noise floor associated with the impulse event.
  - o Again, this is consistent with practical scenarios – in this case MERs for a given AWGN are nearly the same and can yield different UCERs and PERs – MER is not unique and poorly represents impulse in a typical averaging measurement.
- For either AWGN case, if the fixed ingress was FM, in had a noticeable impact on the ultimate dBc level (noted in left hand column), and on performance achieved as impulse level was dropped to create a low error rate situation.
- The impulse amplitude low-error thresholds reached were relatively insensitive to the AWGN level for CW ingress, but this was not the case for modulated ingress, an observation previously noted.

Table 13 – Baseline Ingress + Impulse: 16-QAM vs 64-QAM

	Impulse (dBc-pk) - Gated AWGN @ 10 MHz BW							
	4 usec @ 100 Hz							
1518-Byte Packets	-10				-5			
Noise Floor = 35 dB	16-QAM		64-QAM		16-QAM		64-QAM	
Ingress (dBc)	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER
CW Interference								
1x @ -10 dBc		0.07%	0.49 / 0.48	1.40%	0.4 / 0.002	0.07%	0.33 / 0.75	2.50%
3x @ -15 dBc/tone	0.5 / 0.006	0.03%	0.5 / 0.24	1.50%	0.9 / 0.07	1.20%	0.33 / 0.57	2.80%
FM Modulated (20 kHz BW)								
1x @ -10 dBc	0.55 / 0.04	0.10%	0.0014 / 0.0035	0.10%	0.85 / 0.6	1.50%	0.27 / 1.33	3.70%
3x @ -15 dBc/tone	0.6 / 0.03	0.20%	0.38 / 0.55	2.80%	0.8 / 0.3	2.00%	0.2 / 0.85	3.70%
Noise Floor = 27 dB	16-QAM		64-QAM		16-QAM		64-QAM	
Ingress (dBc)	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER
CW Interference								
1x @ -10 dBc			8.88 / 0.1442	0.50%	1.0 / 0.02	0.10%	8.49 / 0.519	1.60%
3x @ -15 dBc/tone			11.84 / 0.286	0.90%	1.0 / 0.03	0.20%	8.88 / 0.59	2.50%
FM Modulated (20 kHz BW)								
1x @ -10 dBc			23.64 / 2.55	6.00%	0.9 / 0.3	0.90%	19.6 / 1.9	5.00%
3x @ -15 dBc/tone			19.4 / 1.29	5.00%	0.9 / 0.14	0.80%	19 / 1.53	6.00%

	Impulse (dBc-pk) - Gated AWGN @ 10 MHz BW							
	10 usec @ 1 kHz							
1518-Byte Packets	-15				-10			
Noise Floor = 35 dB	16-QAM		64-QAM		16-QAM		64-QAM	
Ingress (dBc)	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER
CW Interference								
1x @ -10 dBc	0.055 / 0	0.00%	6.03 / 2.12	6.00%	7.9 / 0.2	0.70%	2.05 / 9.4	38.00%
3x @ -15 dBc/tone	0.07 / 0.007	0.04%	6.36 / 0.98	5.00%	8.2 / 0.1	0.80%	3.5 / 4.3	27.00%
FM Modulated (20 kHz BW)								
1x @ -10 dBc	0.17 / 0.017	0.07%	5.05 / 5.51	15.00%	8.2 / 1.9	4.50%	1.4 / 15.4	42.00%
3x @ -15 dBc/tone	0.2 / 0.02	0.10%	4.59 / 4.44	21.00%	8.5 / 0.6	4.00%	1.4 / 12.5	44.00%
Noise Floor = 27 dB	16-QAM		64-QAM		16-QAM		64-QAM	
Ingress (dBc)	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER	UCER/CCER	PER
CW Interference								
1x @ -10 dBc			14.46 / 0.807	3.00%	8.2 / 0.03	0.20%	9.86 / 8.11	40.00%
3x @ -15 dBc/tone			16.9 / 1.21	8.00%	8.5 / 0.04	0.40%	13.12 / 9.09	42.00%
FM Modulated (20 kHz BW)								
1x @ -10 dBc	0.3 / 0.01	0.05%	18.96 / 8.0	20.00%	9.1 / 1.4	4.00%	18.48 / 16.9	46.00%
3x @ -15 dBc/tone	0.5 / 0	0.00%	23.1 / 5.56	18.00%	8.7 / 0.2	1.70%	18.3 / 15.1	47.00%

Some items of note:

- For the short duration impulse case @ -5 dBc, there is an order of magnitude PER advantage for 16-QAM over 64-QAM when there is CW interference. This advantage is diminished when the interference is FM modulated, again pointing out the added difficulty of wideband and narrowband interference when the narrowband process is more randomized.

- For the long duration impulse case @-10 dBc, catastrophic PERs for 64-QAM improved an order of magnitude or greater in 16-QAM mode. For the 27 dB AWGN case, a full two orders of magnitude improvement is noted for the CW interference case.

- For long duration impulse noise @ -15 dBc, 64-QAM PERs remained catastrophic, while 16-QAM performed with virtually no packet errors.

Expecting correlation between UCER and PER, we plot the two against one another in Figure 35. Clearly, for large packets such as 1518 bytes, there are going to be multiple codewords. For this profile, the codeword size is n=251, so the ratio is significant but only by a decimal point or so accounting for overhead bytes. This is reflected in the figure as well, as just one codeword of the bunch inside a packet will be counted as a PER. The data used for the plot was combined impulse and ingress, based on previously discussed measurements throughout Tables 11-13.

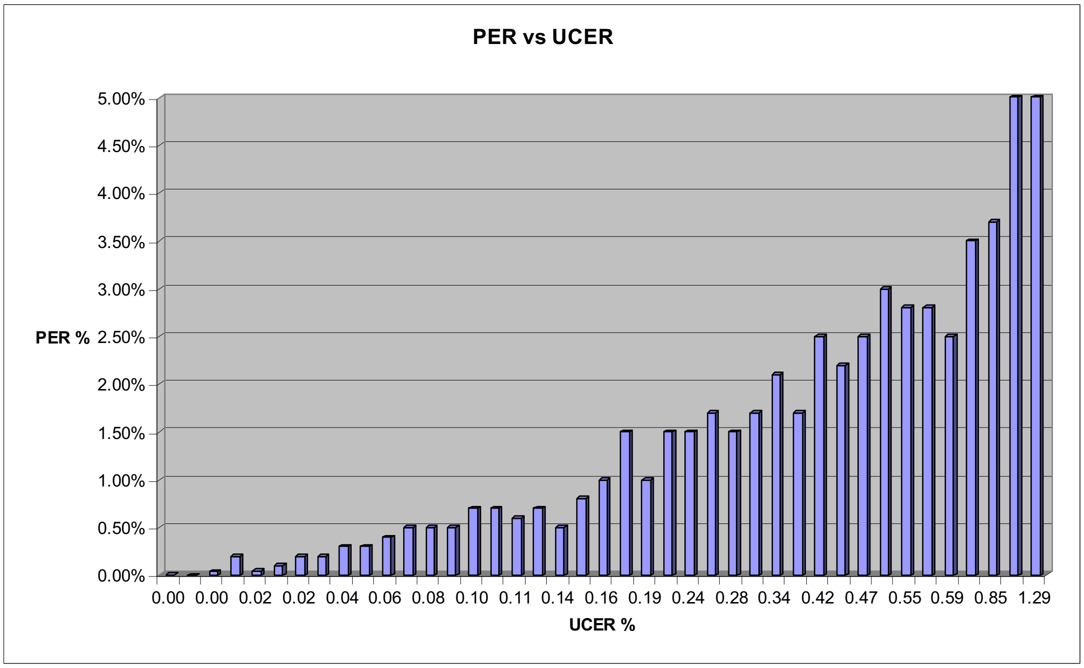

Figure 36 – PER vs Uncorrectable Codewords – Combined Impairments

What we notice is a rather steady trend up to about the 2% range, where the PER has moved to the unacceptably high side. With this much error in play, many of the assumptions and expectations break down, and quickly escalating error rates ensue. Below 2% or thereabouts looks to show a reasonable, linear average fit. However, there is not significant enough granularity or volume of data to pronounce such a curve fit, and the anomalous points need further investigation. At this point, the plot simply represents and encouraging expectation in the relationship between the two for the conditions used here. Quantifiably, this data also indicates a relatively consistent and predictable UCER-PER relationship in the .5% PER range. UCER remains below about .15% (1.5e-3) before a further increase begins to demonstrate a steadily increasing PER above this range. Given that this represents a combine impairment scenario with a dominant impulse noise case, a conservative 1518-byte packets, and a live HFC link including some frequency response distortion, this suggests that a .1% (1e-3) UCER represents a maximum value target for this range of PER.

The sets of data above, and similar sets for short packets can be used to understand the sensitivities of 64-QAM going onto the upstream plant, as well as how it can be expected to react relative to 16-QAM. Again, while the conditions here are on the worse-then-average side, they do represent conditions that can be typically found in the range of about 15 MHz. Also, below 10 MHz can be expected to be worse than these conditions for a not-insignificant percentage of nodes across a system. And, on average, there will exist a gradually improving fidelity of spectrum as frequency increases above 15 MHz.

The strong impact of impulse noise here is important – this can often be the bewildering element to a troublesome upstream, and can be in one subscribers home unknowingly and causing harm to many. Also important is how powerful the ability to cancel narrowband interference is, but also recognizing it begins to see limitations as the interfering signals take on wider spectral properties, in some case more so than varying interfering levels contribute. It also shows noticeably more struggle when the AWGN environment adds error magnitude at the receiver to deal with.

D. Role of S-CDMA

The prior section created a difficult channel environment, particularly the long burst duration case. It is an environment where A-TDMA has known performance limitations, increasing the modulation profile in DOCSIS 2.0 and 3.0, as we have seen, aggravates the situation. Synchronous Code Division Multiple Access (S-CDMA) was incorporated into the DOCSIS 2.0 standard, in recognition of its inherent superior ability to handle channels that include a *significant* amount of impulsive noise, in addition to the combined impairments of impulse and ingress. Impulsive noise is typical to find in the 5-15 MHz portion of the upstream band, with a wide range of relative levels, durations of bursts, and duty. In short, while it is nearly always present in the 5-15 MHz band of the upstream, its characteristics for a given link are quite unpredictable, but within any given link there tends to be a consistent signature. While 15 MHz and below is commonly impaired, impulse noise, in part because it is wideband by nature, will often extend out to 20 MHz, at generally declining levels.

Based on the above, it is clear why most system deployments implement upstream DOCSIS carriers above 20 MHz. Concerns over the channel quality below 20 MHz, and certainly below 15 MHz, have left this spectrum as essentially barren, at least to high speed data services. This represents an obvious waste of spectrum resources – anywhere from 25-40% – when it remains empty. S-CDMA is the tool that can mine this otherwise unused spectrum. By doing so, nearly 50% new capacity can be added on the same return spectrum, providing room for new growth and buying more time and flexibility for planning of capacity-enhancing steps that may be more intrusive and expensive. S-CDMA not only allows higher order modulations than A-TDMA in difficult spectrum, it can do so with more efficiency (less overhead), and in the most troubled spectrum can deliver throughput where A-TDMA may not operate at all. Figure 37 shows a hypothetical upstream that takes advantage of some of the features of S-CDMA to extract as much upstream capacity as can be had on the 5-42 MHz return. In this example, 45% more throughput is enabled. This could translate, based on Table 1 for example, into two years time to defer or plan new node splits for the upstream growth projections used in the table.

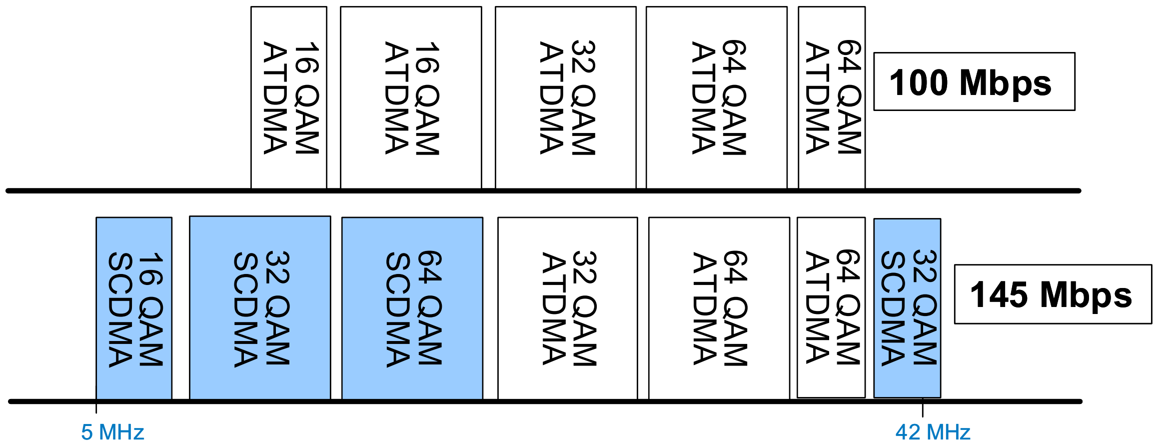

Figure 37 – Hypothetical Upstream Use – A-TDMA & S-CDMA

Features & Benefits

The key element of S-CDMA is its inherent immunity to impulse noise, and thus to enable channels that include impulse noise in addition to other impairments. The ingredient that makes the robustness to impulse noise possible is the spreading out of the symbols by as much as 128 times in the time domain, which directly translates to stronger protection against impulse noise. Noise bursts that may wipe out many QAM symbols of an A-TDMA carrier must be two orders of magnitude longer in duration to have the same effect on S-CDMA, which is very unlikely. S-CDMA has greater duration to wait out bursts and is not sensitive to packet size the way interleaving is in a burst environment. It is the characteristic of the lower part of the return path spectrum to have significant impulse noise, and thus where S-CDMA has the most benefit. It is the spread signaling approach itself, without even considering FEC settings, that enables S-CDMA to withstand much longer impulsive events. There is no reduction in throughput as a result of this spreading, of course. This is because the slower symbols, in the example above 128 of them, are transmitted simultaneously. Refer to Figure 38 below.

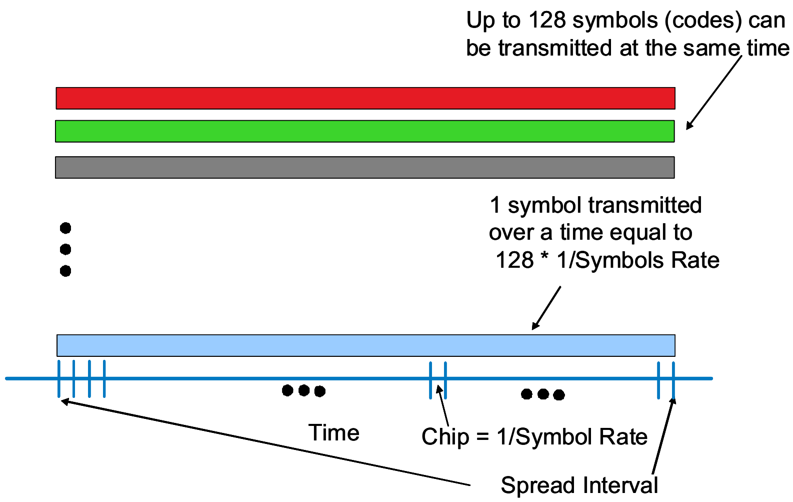

Figure 38 – S-CDMA Parallel Symbol Transmission

The way S-CDMA works is that these parallel-transmitted symbols are independently detectable at the receiver by relying on the property of orthogonality – each symbol is "multiplied" in the binary sense by a spreading code similar to what is used in wireless spread spectrum to allow simultaneous users on the same carrier frequency. For upstream HFC, multiple users are simply replaced by multiple symbols in parallel in a given time slot. The spreading codes or "chips" set the bandwidth though the rate at which they multiply the symbol. These "chip rates" are the DOCSIS bandwidths we are familiar with. At the receiver, the correlation process in the S-CDMA receiver pulls the individual symbols from the parallel stream by replicating the codes in synchronizing to them (thus the "S" in S-CDMA). This synchronization allows any other code that is not a symbol's encoded sequence to result in the correlation process netting to zero. The desired matching code thus delivers the symbol alone with zero (ideally) interference. S-CDMA is in some ways analogous to Orthogonal Frequency Division Multiplexing (OFDM), with the difference being that OFDM leverages orthogonality in the frequency domain, while S-CDMA does so in the code domain. In both cases, it is the simultaneous transmission of multiple parallel, lower rate symbols that delivers performance advantages.

Now, like OFDM, S-CDMA by nature creates a high single-channel peak-to-average ratio – one that looks noise-like. This is in contrast to QAM, whose theoretical peak-to-average is calculable in terms of its constellation format, although this is modified to something higher by the Nyquist pulse-shaping used to efficiently use the channel bandwidth. The impact of this is that, whereas for A-TDMA the concern over a high dynamic range return load on the laser may occur only under an aggregate of multiple independent A-TDMA channels sharing the upstream, for S-CDMA, just one channel can create this situation without extending the loading back-off, at the expense of SNR. This is not seen as a serious shortcoming however, as the anticipated use of S-CDMA is after other portions of the spectrum are filled, and thus aggregate signals exist. Also, with the move afoot to deploy DOCSIS 3.0, upstream lasers are generally being removed in favor of DFB lasers. Their dynamic range exceeds FP technology, and in some cases far exceeds some of the oldest FP lasers. Use of S-CDMA is one more cautionary reminder, however, of the importance of adhering to proper upstream alignment principles.

Note also that neither S-CDMA nor OFDM is a *modulation* in the purest sense of the term, which refers to information content, but are in fact multiple access or signaling format techniques. The underlying modulation in both cases is, typically, QAM, just as it is in TDMA or A-TDMA.

There are some important secondary features of S-CDMA that further enhance its potential as a return path solution. The synchronous operation reduces overhead associated with guard times and preambles for synchronization associated with burst reception. It is not an enormous savings, but contributes in particular to efficiency savings for small packets. S-CDMA includes ingress cancellation, much like A-TDMA, which is partially responsible for its improved performance under combined impairments associated versus A-TDMA. Ingress cancellation does consume some code capacity (throughput) in order to obtain the channel knowledge necessary to mitigate the interference.

Finally, S-CDMA implements Trellis-Coded Modulation (TCM), a well-known technique for maximizing coding gain without adding bandwidth overhead to do so. It does this by expanding the symbol set instead of appending parity symbols, and the combined coding-modulation approach to FEC is known to offer advantages to providing coding gain closer to theoretical channel capacity than independently performing FEC encoding and symbol mapping.

Additional S-CDMA features vary when considering DOCSIS 2.0 or DOCSIS 3.0. These will be discussed in the next section.

DOCSIS 2.0 & 3.0 S-CDMA

As mentioned previously, the addition of ingress cancellation to S-CDMA was an important step for ensuring its value. Narrowband ingress interferers are a known issue, and A-TDMA receivers have employed this mechanism to mitigate the problem. DOCSIS 2.0 includes this ingress cancellation, and DOCSIS 3.0 adds a new (optional) wrinkle called Selectable Active Code (SAC) that makes it even more powerful, in particular as DOCSIS 2.0 and DOCSIS 3.0 implement the higher order, more sensitive, QAM profiles. SAC allows only a subset of the codes to be turned on for payload use. Each code is fixed and periodic, and as such has a frequency response that can be associated with it. Because of this, some codes are more susceptible to an interferer than others. Offering the option to select the codes to use or not to use provides an extra tool to mitigating the most severe levels of ingress and support the highest order modulation profiles. This mode is called SAC Mode 2. The difference in mode 1 is that there is no selectable code usage. The ingress canceller operates through what it can learn from the number unused codes, whereby these unused codes begin at code 1 and progress sequentially. While SAC2 creates an especially powerful advantage in narrowband ingress, a very important point with it is that *all* modems on the channel must support it in order for it to be enabled.

Another DOCSIS 3.0 feature is that S-CDMA also offers flexibility in the code implementation that allows a trade-off between the power allocated per-code and the number of spreading codes turned on. This tool is known as Maximum Scheduled Codes (MSC). For example, if 128 codes are on transmitting at Pmax, each code is allocated Pmax/128. If only 64 codes are used, however, each code is allocated Pmax/64, or 3 dB more power per code. This comes at the expense of throughput, but offers some choices to the operator that may be better than an equivalent A-TDMA alternative. The added power can be used to reach deeper into highly attenuated upstream in the home, for example, or offer an SNR boost towards a higher order modulation. However, lower order A-TDMA options offer similar potential if it is a matter of simple attenuation and not also a channel quality (low return band) issue.

S-CDMA Performance

Test modems using live plant characteristics have proven out the advantage that S-CDMA creates in the poorer part of the upstream spectrum.

A sample for comparison of S-CDMA and A-TDMA on the identical return path conditions is shown in Figure 39. Apparent from the figure is that A-TDMA is taking errors in transmission at a nearly 20% clip, while S-CDMA is taking none. In this case, all of the errors were correctable errors, but clearly in the A-TDMA case the FEC is working extremely hard to maintain adequate performance, while the S-CDMA link is not struggling at all. It is not difficult to extend this to a more difficult channel, whereby the A-TDMA link would clearly begin to break before the S-CDMA link. In this trial, existing legacy STB upstream traffic in the low end of the band below 10 MHz did not provide ample unused spectrum to deploy a wideband carrier in that worst-case region. The lower band edge was thus limited by the STB carrier, pushing the center frequency well above 10 MHz to run the test creating these results. It is very important to point out in this "apples to apples" comparison that, in fact, the FEC settings are dialed up for A-TDMA, at T=16 symbol error correcting, while for S-CDMA, they are dialed down to T=2. S-CDMA inherently takes advantage of its impulse immunity properties rather than relying on FEC.

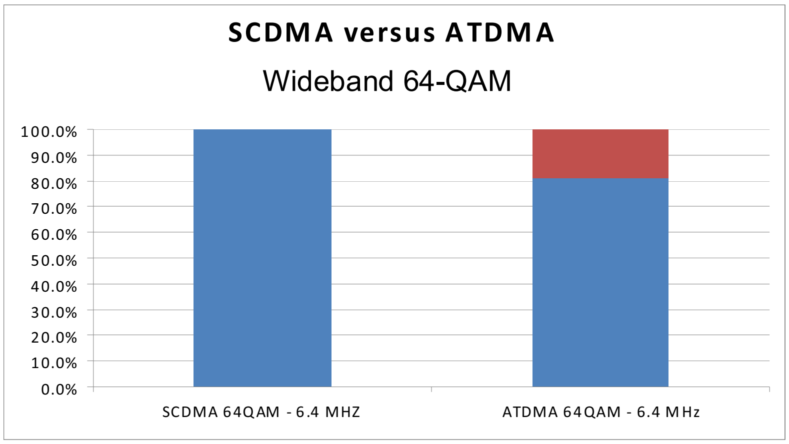

Figure 39 – Corrected Error Statistics

To attempt to quantify the improvement in link budget dBs between S-CDMA and A-TDMA, it is necessary to operate both modes head-to-head, at the same frequency, and one where both can reliably close the link – such as in the case of Figure 39. However, it is actually more desirable to have somewhat worse performance for this type of estimation, because without errors it is not possible to observe how much margin remains. In addition, a handy tool to provide some insight into the dBs of link gain is the modulation profile, because of the known relationship between the QAMs previously described. Cases where S-CDMA runs with lower error rates and fewer uncorrectable codewords were observed in the trials. Situations where the A-TDMA mode simply could not close the link reliably also were observed. Using all of test results from multiple sites, and additional lab characterization, Table 14 summarizes approximate link "gain" relationship from a modulation performance perspective using S-CDMA under the conditions tested.

Table 14 – Link Gain Estimates – S-CDMA vs A-TDMA

Frequency Band	2.56 Msps	5.12 Msps
5.8-9.0 MHz	< 6 dB	< 8 dB
9.0-13.0 MHz	< 4 dB	< 6 dB
13.0-18.0 MHz	< 2 dB	< 2 dB
18 MHz < fo <38 MHz	~Equiv	~Equiv

Note: 4-8 dB represents 1-2 orders of modulation order

– ***16-QAM → 32-QAM (~3 dB) → 25% more throughput***
– ***16-QAM → 64-QAM (~6 dB) → 50% more throughput***
– ***32-QAM → 64-QAM (~3 dB) → 20% more throughput***

The advantages of S-CDMA as lower frequencies are used in the return band are evident in this table. Further field work is planned to corroborate these early results, as well as including using S-CDMA in combined impairment testing such as is discussed in this paper.

IV. Summary - Deployment Guidelines

- Post-MER thresholds of 15/21/24/28 dB support 4/16/32/64-QAM – at a zero margin reference. ***The 64-QAM question is all about achieving and maintaining tolerable margin, noting that QPSK-to- 16-QAM is NOT a good reference guideline.***
    - Adding 5 dB to the numbers above for dynamic margin means 26 dB for 16-QAM and 33 dB for 64-QAM. Note that 26 dB would be inadequate to reliably support 64-QAM.
    - Having 30 dB of MER available means 9 dB of margin to 16-QAM, and that many ugly things can be happening on the channel and still support 16-QAM. By moving to 64-QAM its down to 2 dB and a couple of hair dryers, loose connectors, or an alignment offset and an extra cold night can set the stage for phone calls.
    - Margin will have to be looked at differently going forward with 64-QAM. There will be less of it, consistently, and because of that all the setup and maintenance dBs that once were of little significance matter more.

- MER is not uniquely associated with error rates at the bit, FEC, or packet level
    - The same MER can yield differing error rates
    - The same error rates can be observed with different MERs
    - Both of the situations above were observed in the data
    - What the CMTS reports as SNR is actually MER

- Return optics – particularly laser technology – are the dominant factor to setting a sound channel average MER
    - There are a wide range of lasers in the field and a wide range of capability. The oldest FPs have virtually no chance of supporting 64-QAM reliably, while modern DFBs can due so comfortably, generally speaking.

  - There are a variety of choices that are on the borderline of being capable, most notably isolated FP lasers. These can be made to work, but generally will do so without comfortable margin. They become sensitive to link length variables
  - Older lasers are often consistent with an older, larger plant, exposing them to more noise variables from homes and upstream amplifiers, eating further into an already smaller starting margin.

- While a sound MER foundation is a good figure of merit, as we have seen, MER does not tell the entire story, in particular to the dynamic impairments characteristic of the HFC return path
  - Static ingress only moves MER when it is of significant relative power to the channel, which can still be a small value for a relatively poor signal-to-interference ratio (S/I), such as S/I = 10 dB.
  - Impulse noise is even less noticeable, and even when very high during an event, it is hidden by the averaging nature of the measurement.
  - Thus, a problem performing node comes with these first three simple and straightforward steps:
    - Is the MER a sound foundation?
    - Is their interference in the band as a spectrum analyzer or monitoring receiver running FFTs on the CMTS see it?
    - If neither of the above is apparent, the natural candidate is impulsive events: a spectrum analyzer set to rapidly sweep, use of zero span, or use of time-based FFTs available through some CMTSs can diagnose this

- Because of the dominant effect impulse noise can have in setting the link performance floor, as has been observed, it is recommended to consider S-CDMA to utilize the low end of the band to its maximum capacity. S-CDMA is significantly more robust to impulse noise and combined impulse and ingress than ATDMA.

- Cascade depths are decreasing, but not as quickly as DOCSIS 3.0 is rolling out. Therefore, understanding the sensitivity of wideband 64-QAM to return center frequency and cascade matters. A flag should be set when the region of 37 MHz/39 MHz (40 MHz/42 MHz split) is under consideration for deep cascades per Table 6.

- Maintaining a PER of better than .5% correlated to about a .1% (1e-3) UCER for long packets. Above the .15% UCER, it is observed that PER can begin to rapidly increase. This result includes all items present in live plants and difficult channels: combined ingress and impulse, AWGN, an HFC link, and live modems sending long packets.

- There are many other physical layer parameter permutations, and layer two principles to consider further as characterization continues for a comprehensive DOCSIS 3.0 evaluation via test and measurement.

V. Conclusion

DOCSIS 3.0 offers tremendous opportunity for cable operators to deliver enhanced, compelling new services. In the upstream, it opens the door for rapidly expanding residential services with volumes of user-generated content, such in the explosion of social networking, continued peer-to-peer capability, and expanded business services offerings. However, the new technology relies on advanced modulation techniques and utilizing the full spectrum, in order that increased upstream rates go hand-in-handed with new capacity that supports the continued growth.

This paper illuminated the many challenges involved in optimizing the upstream – and make no mistake, there is still much to optimize. DOCSIS 3.0, and the growth of upstream in general, puts into the past dealing with upstream service, alignment, and maintenance with a casual attitude. However, all of the tools and techniques needed to fully utilize the HFC upstream to its maximum capacity are in place already. Now is the time to put them to use.

References

[1] Howald, Robert, "Fueling the Coaxial Last Mile," SCTE Conference on Emerging Technologies, Washington DC, April 3, 2009.

[2] Howald, Robert, and Michael Aviles, "NPR the Analytical Way," National Cable Television Association (NCTA) Show, New Orleans, La., May 7-10, 2000.

[3] Howald, Robert, Amarildo Vieira, and Michael Aviles "New Megabits, Same Megahertz: Plant Evolution Dividends," National Cable Television Association (NCTA) Show, Washington DC, April 1-3, 2009.

[4] Howald, Robert, "Enhanced Return Performance Using Modulation-Based Alignment Criteria," SCTE Cable-Tec Expo, Orlando, Fla., May 2001.

[5] Stoneback, Dean and William Beck, "Designing the Return System for Full Digital Services," SCTE Conference on Emerging Technologies, San Francisco, CA, January, 1996

[6] Thompson, Robert, Chuck Moore, Jack Moran, and Rob Howald, "Optimizing Upstream Throughput Using Equalization Coefficient Analysis" National Cable Television Association (NCTA) Show, Washington DC, April 1-3, 2009.

[7] ANSI/SCTE 119 2006, "Measurement Procedure for Noise Power Ratio," www.scte.org/standards

[8] ANSI/SCTE 132 2007, "Test Method for Reverse Path (Upstream) Bit Error Rate," www.scte.org/standards

[9] Broadband Return Systems for Hybrid Fiber / Coax Cable TV Networks, Raskin, D. and Stoneback, D; Prentice-Hall PTR, 1998, ISBN 0-13-636515-9.

[10] NCTA Recommended Practices for Measurements on Cable Television Systems, Third Edition

[11] Motorola Return Path Level Selection, Setup, and Alignment Procedure, part number 453760-001

[12] Primer: Reed-Solomon Error Correction Codes (ECC), AHA Application Note ASNRS01_0404

Acknowledgements

Many hours of support and many thanks go out to the Motorola DOCSIS team of Mike Cooper, Marc Belland, Jack Moran, Marc Morrissette, for their contributions throughout this paper.