THE GROWN-UP POTENTIAL OF A TEENAGE PHY

Dr. Robert Howald, Robert Thompson, Dr. Amarildo Vieira
Motorola Mobility

Abstract

Cable operators continue to see persistent annual increases in downstream traffic, most recently driven by aggressive growth in over-the-top video. Well-understood tools exist to manage the growth, including switched digital video (SDV), analog reclamation, service group splitting, improved encoding and transport efficiency, and RF bandwidth expansion. Beneath it all, however, the underlying RF transport approach has remained unchanged, relying on 1990's era ITU J.83 technology for PHY layer and FEC technology. Meanwhile, 15+ years of advancements in communications technology and processing power have since taken place. Many of these advances, which close the gap between Shannon theory and real-world implementation, are already being tapped in other industries. Some are now poised to enable cable to support a new generation of Gbps-class services and to mine completely the capacity of the coaxial last mile – a key element to guaranteeing an enduring HFC lifespan.

In this paper, we will present a comprehensive link analysis addressing the deployment possibilities of these communications technology advances over the HFC channel. We will focus in particular on the ability to support higher order QAM, such as 1024-QAM through 4096-QAM. We will discuss the role multi-carrier techniques (OFDM) could play and why. We will specify SNR implications to HFC, including considerations for fiber deep migration. We will describe the SNR repercussions of advanced QAM to CPE noise figure (NF), which is critical to understand as wideband, digitizing front ends replace analog STB tuners. In addition, we will dive deeper into the subtle link impairments that become potentially limiting factors as we push the boundaries of PHY technology on the cable plant. Previously less significant issues such as timing jitter and phase noise are magnified as constellations become increasingly dense. These items ultimately effect equipment requirements.

In summary, we will articulate and quantify the ability of the HFC network to support ever-increasing orders of bandwidth efficient modulation, and the impact these modern communications formats have on equipment and requirements.

INTRODUCTION

The industry is deeply engaged in long-term network planning, in recognition of the continuing growth of IP traffic and concern for the network's ability to support it. There are two key components to the problem. The first is simply determining if the infrastructure in place is physically capable of delivering on the growth, and, if so, for how long. There tends to be a consensus that the HFC architecture is capable enough, but that it has not been optimized as of today to ensure it [8, 11]. This brings us to the second part. If the answer to the first is yes, then how do initiate a transition plan from today's infrastructure to the architecture that does optimize what can be extracted from the network?

There are many spectrum and capacity management tools in the downstream. However, the operator has much less control over upstream congestion. Common to both downstream and upstream is the reliance on what is now aging, 1990's era, physical layer

(PHY) tools. In use on the downstream is the ITU J.83B Physical layer (PHY) and forward error correction (FEC) technology. The DOCSIS upstream is also QAM with a Reed-Solomon based FEC – powerful at the time but more powerful PHY techniques are available today. The result is less efficient use of cable spectrum that could be achieved with modern PHY tools.

Relying on 1990's era technology puts cable at a disadvantage. Perhaps the single most important long-term objective of the architecture transition is, in the end, to have created maximum bandwidth efficiency (and therefore maximum lifespan) cost-effectively. We will focus our attention on this one component of capacity management – more efficient use of spectrum – more bits-per-second-per-Hz (bps/Hz).

Capacity Levers

Note that theoretical capacity is a based on two variables – bandwidth (spectrum allocated in our case), and SNR. For high enough SNR, finding spectrum dominates the equation. The capacity equation can be simplified, and in so doing, it can be shown that capacity is essentially directly proportional to bandwidth, B and SNR expressed in decibels (dB):

$$C \approx [B]\,[SNR\ (dB)] / 3 \qquad (1)$$

This approximation is accurate asymptotically within 0.34% with increasing SNR. Because of the inescapable relationship of capacity to bandwidth, as cable looks to increase capacity, new spectrum is being sought after. Figure 1 is an example of a likely spectrum evolution [2], resulting in a final state of bandwidth allocations.

In the downstream, we are extending an excellent channel into an area where it will suffer more attenuation as a minimum. It will also likely have to deal with frequency response issues.

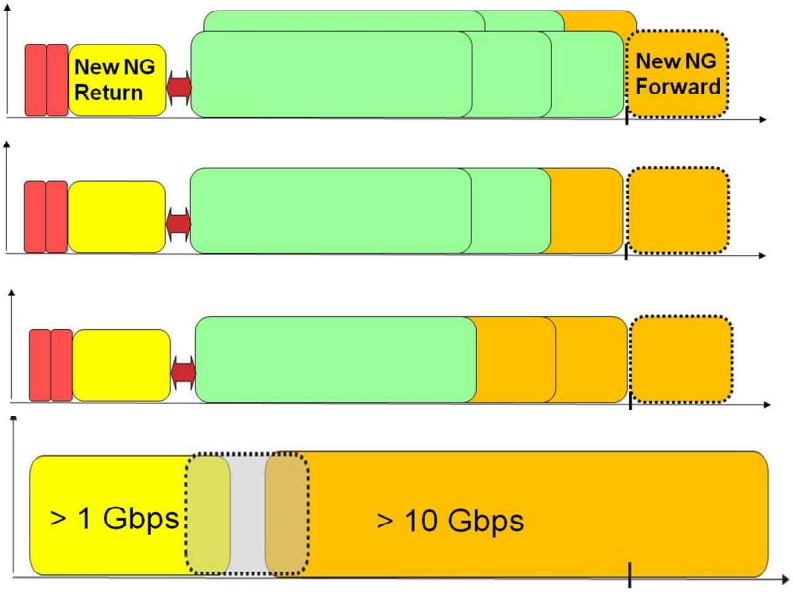

Figure 1 – Likely Cable Spectrum Evolution

In the upstream, we will be in some ways doing the opposite – extending a partially troubled channel into an area where we expect a much better behaved environment from which to extract new capacity.

We will be taking advantage of significant technology advances for enhancing the PHY. Much of what is being taken advantage of is continued advances in the real-time computing power of FPGAs and ASICs. The theoretical basis for modern PHY tools in some cases is, in fact, very old.

For example, Reed-Solomon coding itself was born in the 1959. Low Density Parity Check Codes (LDPC), the basis of today's most advanced forward error correction (FEC), is also quite old, first introduced in 1960. Information Theory is a linear algebraic discipline. However, the computing power to perform the algebraic operations required for efficient decoding of very large matrices, and using non-binary arithmetic (Reed-Solomon) came along much later.

Multi-carrier modulation (MCM, the generic name for OFDM and its variants – we will use both throughout) has a parallel history to advanced FEC in this sense. It was very difficult to implement, until the FFT version of the DFT came along, followed by computing power to calculate larger FFTs faster. IFFT/FFT algorithms form the core of the OFDM transmit and receive function. So, while we are talking about "new" technology, it is important to understand that these technologies are already very well grounded theoretically.

Of course, the single most important attribute of these advances is that they close the gap between theoretical Shannon capacity and real-world implementation – something the world has been trying to do since 1948. A simple crunching of today's HFC performance and throughput illustrates how far from this ideal we are today. Table 1 compares downstream and upstream as we use them today against the theoretical capabilities of the channel. We have accounted for code rate efficiency losses, but not framing, preamble, or other overhead unrelated to pure PHY channel transmission capacity.

Table 1 – Downstream and Upstream vs. Theory

	D/S	U/S
BW (MHz)	6	6.4
SNR (dB)	35	25
Capacity from (1)	70.00	53.33
Legacy QAM (no framing)		(t=10)
256-QAM	38.83	37.75
64-QAM	26.99	28.31
Delta Capacity		
256-QAM	**55%**	**71%**
64-QAM	**39%**	**53%**

As we can see in Table 1, today's commonly used modes – 256-QAM downstream and 64-QAM upstream – operate at 50-60% of capacity, and therefore are leaving a lot of bits on the table. With the help of new tools and supporting architecture evolution, some of the current limitations can be alleviated, putting cable in a position to deliver a new class of services and maximize the lifespan of its core architecture.

QAM LINK BUDGETS

Capacity Enhancements

Let's begin with the "SNR" part of (1). There are two elements of the SNR component of capacity. First, clearly, more capacity is available if higher SNR is available. Since it is related to SNR in dB, however, it is a compressing function, and its affect on capacity less effective in increasing C compared to spectrum. In practice 50% more spectrum yields 50% more SNR. This

is also true for 50% more SNR in dB. However, in practice, 50% new SNR in dB, such as converting a 35 dB SNR into a 52.5 dB SNR, is not reasonable in most cases. Nonetheless, it is certainly the case that more SNR translates to more capacity, and architectures that create higher SNR are architectures that open up more potential capacity. This is why, for example, when fiber deep topologies are discussed, both average bandwidth per home (because of fewer homes per node) as well as a more robust, higher SNR channel for use are both important results. We will quantify architecture effects later in this section.

Part 2 of the SNR component of capacity is using the available SNR most efficiently. This is specifically where MCM and FEC advances come into play. A good way to understand the former is to use the "long" (but not longest!) form of (1).

$$C \approx (1/3)\sum_{\Delta f} [\Delta f]\, [P(\Delta f)\, H(\Delta f) / N(\Delta f)]_{dB} \quad (2)$$

This is the same information expressed in (1), just differently. Instead of bandwidth, we have used a summation over a set of small frequency increments, Δf. The sum of all Δf increments is the bandwidth available, B. Instead of SNR, we have identified the components of SNR – signal power (P), noise (N), and channel response (H) – each also over small frequency increments.

The capacity, then, is a summation of the individual capacities of chunks of spectrum. The purpose of (2) is to recognize that channels with changing SNR – such as any "new" bands to be exploited outside the normal cable bands – that may not have a flat response. In particular, above today's forward band there will be roll-off with frequency. The capacity of this region can be calculated by looking at it in small chunks that approximate flat channels. More importantly, however, a technology that can actually *implement* small channel chunks can optimize each of those frequency increments to get the most capacity from them. This is the key advantage of MCM – very narrow channels, each of which can be loaded with the most bits possible. With a single, wide, transmission, it is difficult to achieve the same effect without very complex, and sometimes impractical equalization techniques and interference mitigation mechanisms. Thus, (2) effectively expresses why MCM is often better suited in channels with poor frequency response. MCM also accomplished this while overlapping these narrow channels. They are kept the independent through the orthogonality of frequency spacing – separated by the symbol rate of the sub-channels.

Figure 2 shows a capacity calculation of the forward band extension case described above. It plots the capacity including bandwidth above a 1 GHz network when the network is modeled as a lowpass roll-off, governed by the frequency response characteristics of 1 GHz taps [7]. The red curve shows the aggregate roll-off of five taps and a single coupled port, as well as interconnecting coaxial cable. An assumed 45 dB digital SNR at 1 GHz is used to calculate the capacity as signal power is attenuated above 1 GHz on a flat transmission profile.

The total capacity calculation if the entire forward band is taken into account (blue) is also shown, as well as the capacity over and above what is currently available in a 1 GHz network that is fully loaded with 256-QAM signals (pink). In both cases, the diminishing returns associated with the attenuation of current HFC passives – inherent implementation limitations, not barriers of technology – are obvious as the forward band goes above about 1.2 GHz. Analyses like Figure 2 point out why cable is bullish on the ability of the HFC network to support 10 Gbps data rates.

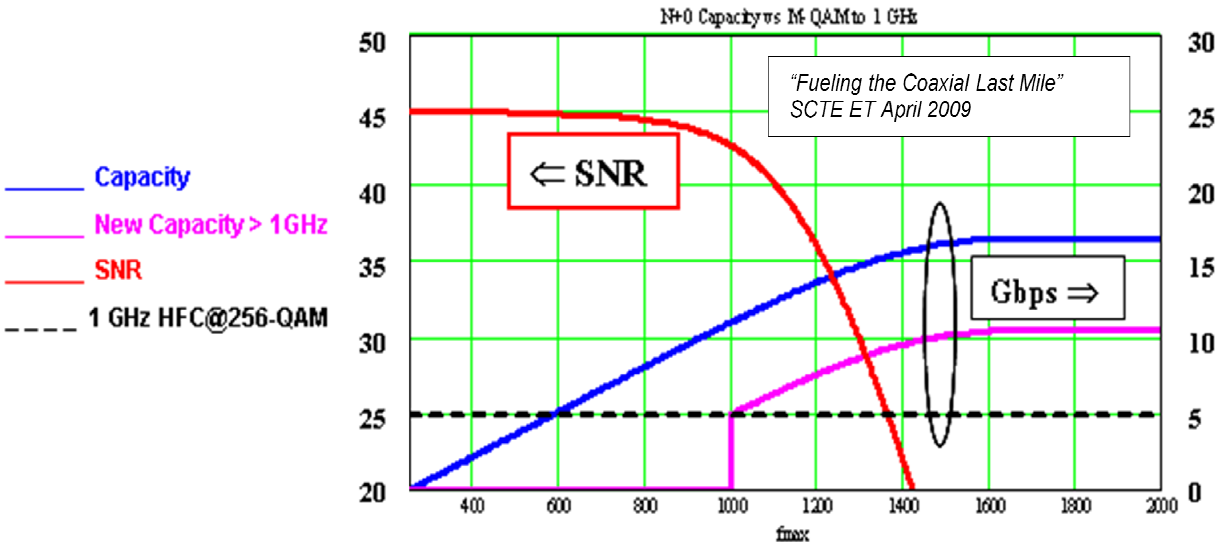

Figure 2 – Capacity Above 1 GHz on a Passive Coaxial Segment

Now let's consider the role of FEC. Using SNR most efficiently, from a coding perspective, is about finding the right codeword design. Major leaps in this capability have occurred in the past 20 years – Reed-Solomon (RS), Trellis Coded Modulation (TCM), Turbo Codes, and now LDPC. Again, the advances have mostly to do with the ability to process the complex decoding algorithms and the tools needed to design them specific to the application.

Coding theory has always been about trying to close the gap between the theory derived by Shannon and the reality on the wire or in the air. Quantifiably, this means simply getting more bps/Hz out of the same or lower SNR.

Defining SNR Thresholds

The impact on of advances in FEC is quite simple – it reduces the SNR required to achieve a particular modulation profile, increasing throughput. The SNR requirements for each QAM modulation profile without coding are theoretically well-founded. We will consider advanced FEC as we compare results of link analysis to determine what can be supported by a particular PHY and HFC architecture SNR, and to compare that to today's architecture and requirements. The thresholds that will govern the comparisons are shown in Table 2.

The three M-QAM BER columns are as follows:

1) 1e-8, No FEC
2) DOCSIS Specification (and extended estimates where QAM profile does not exist)
3) New LDPC-based FEC; assumption of 5 dB more gain

Table 2 - Downstream SNR Assumptions for M-QAM Profiles

	SNR Requirement Assumptions, D/S		
	No FEC, Theory 1.00E-08	DOCSIS Req't (J.83B)	New FEC LDPC @ 5 dB
64-QAM	28	24	19
256-QAM	34	30	25
1024-QAM	40	est. 36	31
4096-QAM	46	est. 42	37

Several important items must be noted with respect to Table 2.

- DOCSIS includes an allocation for implementation margin on top of an assumed coding gain impact. We are inherently carrying those implementation margins forward by using an LDPC gain factor and not an LDPC SNR versus QAM simulation.

- Coding gain may increase as M-QAM orders increase, but it is conversely more difficult to maintain a constant implementation loss across higher profiles. By using 5 dB, we are essentially calling these a wash. No effort has gone into infrastructure requirements to support, for example 4096-QAM, and hardware limitation can become exaggerated for these cases.

- There are no code designs selected using LDPC for North American cable. In Europe, DVB-C2, for example, has defined a range of code rates, and these are SNR reference points reported of the OFDM PHY + LDPC:

  256-QAM: 22-24 dB
  1024-QAM: 27-29 dB
  4096-QAM: 32-35 dB

  These numbers similarly require margin be applied in practice, but are nonetheless useful in understanding how efficient a network can be as other non-idealities of implementation are reduced.

- These are all AWGN-only SNR values, which is the fundamental construct of channel capacity.

- These are SNR thresholds for the downstream only. The first column, of course, is independent of downstream or upstream.

So, while the dB to use for a given profile can be debated, Table 2 gives us a ballpark starting point.

We similarly set thresholds to use for the upstream, which also includes margin for operations. Because of the variety of unknowns, the margin allotment is higher. Note that the upstream is not ITU J.83B, although it is Reed-Solomon-based with configurable error correction parameter. Thus, the RS code can be stronger or weaker, depending on configuration, although it is usually set at lower code rates (stronger).

As reference guidelines for upstream SNR, we choose what one particular operator

uses as classification for a good upstream in terms of observable metrics. These metrics are based in part on upstream SNR (actually MER) reported. A good score for the upstream includes a consistent 30 dB reported SNR. We will assume this would represent a link capable of the highest order modulation profile at all times, which is 64-QAM today. This is what is reflected in Table 3.

Note that this threshold is actually above the no-FEC threshold for 64-QAM. This is simply the nature of the margin allotted to upstream as operated today. While the RS code offers a theoretical gain similar to the RS downstream, it is configurable from none up to strong. Also, the upstream channel has, in general, been difficult to fully exploit to date because of the range of impairments and field implementations. It is certainly reasonable to expect that, as service group sizes shrink, alignment practices improve as bonding becomes prominent, and other architectural changes take place (like a point-of-entry home gateway architecture to be discussed), the quality of the upstream will improve and the margin required to support a particular modulation profile reduced. We do not make any of assumptions about any new dB associated with those "what ifs" here.

Upstream traffic typically comes in small chunks, and therefore can only be supported by smaller block-sized LDPC codes. This leads to less coding gain for a given code rate compared to downstream. We round this difference up to an even 1 dB offset compared to downstream, or 4 dB of new upstream gain from LDPC. For upstream, then we use the assumptions shown in Table 3.

We will go forth with these SNR values as we investigate the implications to key components of the HFC architecture. Note that whether we are discussing legacy single carrier QAM or MCM systems, the SNR thresholds established in these tables are the same. These are based on AWGN performance. Thus, for both legacy QAM style and MCM, the link budget analysis below is applicable. However, it can be argued, in particular for the upstream, that use of MCM offers the opportunity to eliminate some of the margin currently allotted in Table 3, since this is based on experience with today's single carrier upstream channels. The reasoning here is that multi-carrier could be more resilient to some of the things that go into setting the upstream margin.

Table 3 - Upstream SNR Assumptions for M-QAM Profiles

	SNR Requirement Assumptions, U/S		
	No FEC, Theory 1.00E-08	DOCSIS w Upstream Margin Reed-Solomon	New FEC LDPC @ 4 dB
64-QAM	28	30	26
256-QAM	34	36	32
1024-QAM	40	est. 42	38
4096-QAM	46		

DOWNSTREAM HFC MIGRATION

Table 4 quantifies the delivered performance at the end of line for an HFC network based on a classic 1310 nm linear optical link as a function a modern RF cascade, such as GaAs-based RF. It is a typical mix of bridger (multi-port) amplifiers and line extenders where a cascade ensues. Table 4 includes an assumption of partial analog reclamation – a total of 30 analog carriers remain. Most MSOs have an analog reclamation plan, though many anticipate that they will leave a basic tier in place for a long time. We will analyze a full analog reclamation case as well.

Table 4 - Downstream Performance vs Cascade

	1 GHz, 30 Analog Carriers				
	CNR	CSO	CTB	CCN	QAM CCN
N+6	51	60	65	49	43
N+3	55	62	67	52	46
N+0	57	65	69	55	49

These numbers all relate to analog levels, of course, so in reference to digital they must be lowered by the amount of the digital de-rating. Mathematically, it is straightforward to show [7] that removal of analog frees up some RF power from the total load that could be re-allocated to digital loading. This varies from approximately 1-3 dB, depending on how many analog carriers remain (zero or 30) and the tilt used from RF amplifiers on the coaxial leg. We will not account for new RF power at this stage, but come back to this point as we discuss link budget closure. QAM SNR levels – 6 dB lower than the yellow column – are listed on the far right of Table 4 using a 6 dB de-rate. An important and expected result from Table 4 is the improvement in the CNR and CCN as the cascade becomes shorter.

Note that CCN stands for Composite Carrier-to-Noise, and captures all noise floor components – AWGN and digital distortions. It is the "true" SNR, although that is an imperfect label technically because of the contributions of distortion. However, the digital distortion contributors are many and largely independent, so a Gaussian assumption is reasonable. On a 6 MHz channel, a white, Gaussian assumption is also reasonable. For wideband channels, the "white" component may deviate, but this is exactly where MCM plays a role. By its nature, it again will make the channel noise, including CCN, "look" white (flat) in a sub-channel.

Using Table 2 requirements and Table 4 performance, and assuming the lower limit of input power is assumed delivered for a QAM channel at -6 dBmV, we can derive what noise performance is needed from the CPE to meet each threshold. This is shown in Figure 3. We have extended the CCN range downward on Figure 3 compared to Table 4 to represent perhaps deeper cascaded architectures than N+6, or systems running somewhat stretched or simply below the performance of Table 4 for a variety of other design reasons.

Also shown in Figure 3 are the above calculated CCN values of 43 dB (N+6), 46 dB (N+3) and 49 dB (N+0), identified and labeled using red vertical lines. Along with the maximum noise figures plotted in Figure 3 are noise figure values (black dashed lines) representative of common CPE platforms in the field, and of today's vintage, which are lower noise designs. In the case of the "maximum noise figure" curves (color), QAM profiles are supported when the colored line identifying a modulation profile is *above* an example CPE NF threshold in black.

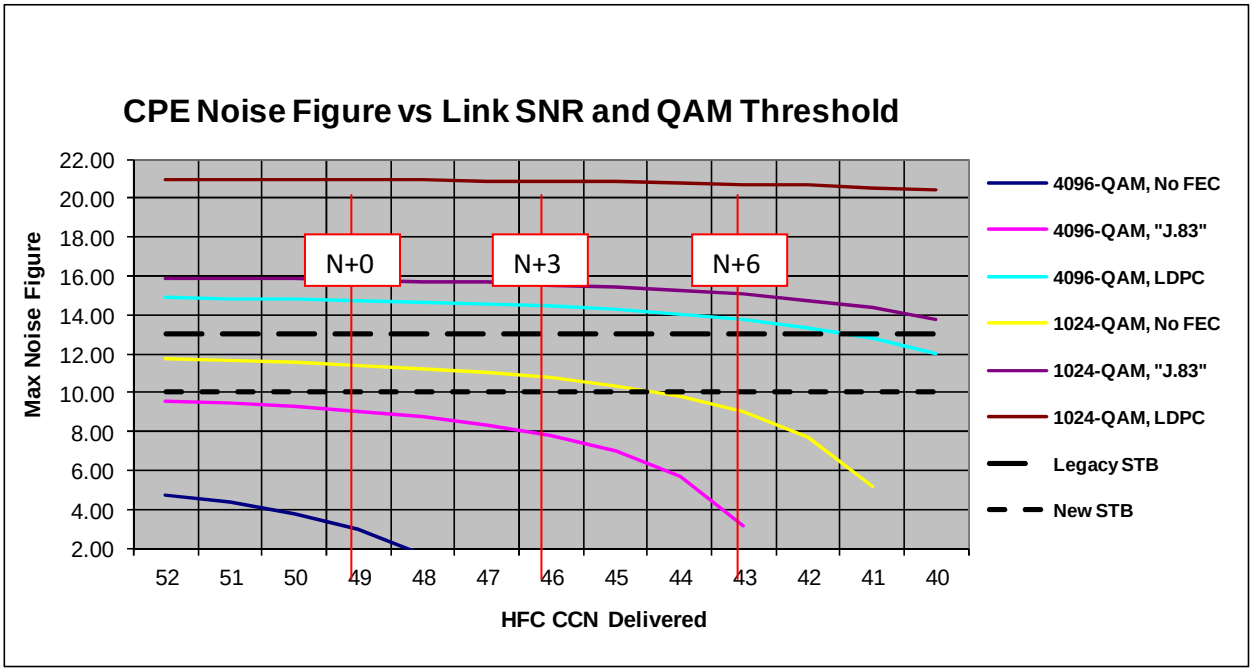

Figure 3 – STB Noise Figure Limit vs. Modulation Efficiency

Apparent from Figure 3 are three things with respect to the access network and home environment:

1) 4096-QAM is not achievable without introducing new FEC, based on today's linear optics and CPE performance. Even with new LDPC FEC, however, it is too marginal to be practical without sensitivity improvements of modern STBs. And, even with those improvements, there is just a little link budget to spare, and only if there is a fiber deep migration. Stretched architectures and high in-home losses could struggle – a QAM input below -10 dBmV to the STB instead of -6 dBmV would be insufficient for N+6, for example. Remember those possible dBs of power allocation gains of analog reclamation we identified earlier? It is cases like this where it becomes obvious that every dB of a link budget becomes critical in some cases for practical margins to be realized.

2) 1024-QAM with a J.83 flavor of FEC is achievable today with legacy STB performance, albeit there is also not much margin. For example, while 256-QAM performance requirements exist down to a -15 dBmV input in DOCSIS, this additional loss would not be able to be absorbed in the 1024-QAM link budget per Figure 3. On the other hand, 1024-QAM with LDPC is the one curve that is clearly and robustly supported – to levels as low as -13 dBmV for even existing STBs and below -15 dBmV for newer class boxes.

3) For HFC migrated to fiber deep architectures such as N+0 or N+(small) – left hand side of Figure 3 – there is little sensitivity of the NF curve to SNR variations for all cases except for a threshold using 4096-QAM without FEC, which is a non-starter.

The fact that 1024-QAM using J.83 is possible is consistent with the conclusions drawn in [9]. That analysis also pointed at

the CPE noise as a potential link limiter to 1024-QAM today.

All in all from an SNR standpoint, though, updating the FEC will be instrumental to delivering higher modulation efficiency on today's quality of HFC architectures, and newer CPE will buy important link headroom to enable robustness.

Figure 3 captures access network and home. A missing component of this analysis is that Figure 3 inherently assumes a perfect transmit fidelity. In fact, of course, the DRFI specification governs transmit fidelity today. If we assume that plant linear distortions are properly handled in the receive equalizer (a good assumption), then we can consider the DRFI Equalized MER contribution of 43 dB. There is an implicit assumption that the MER is not dominated by a few discrete spurious components when we are using an SNR analysis. We will consider discrete interference in a separate section. Figure 4 shows the result with the DRFI requirement included.

While there are major differences in Figure 3 and Figure 4, the conclusions previously drawn do not vary very significantly. 4096-QAM is just more impractical than before, and the 1024-QAM "J.83" case has lost some of its margin, but remains on the bubble of workability as long as the HFC link is very good, such as the N+0 case.

Lastly, we point out that while we have captured the DRFI MER requirement in Figure 4, that requirement is obviously a minimum. Although broadband fidelity requirements are among the most difficult to meet, suppliers compete on key parameters and therefore product performance may be better in practice.

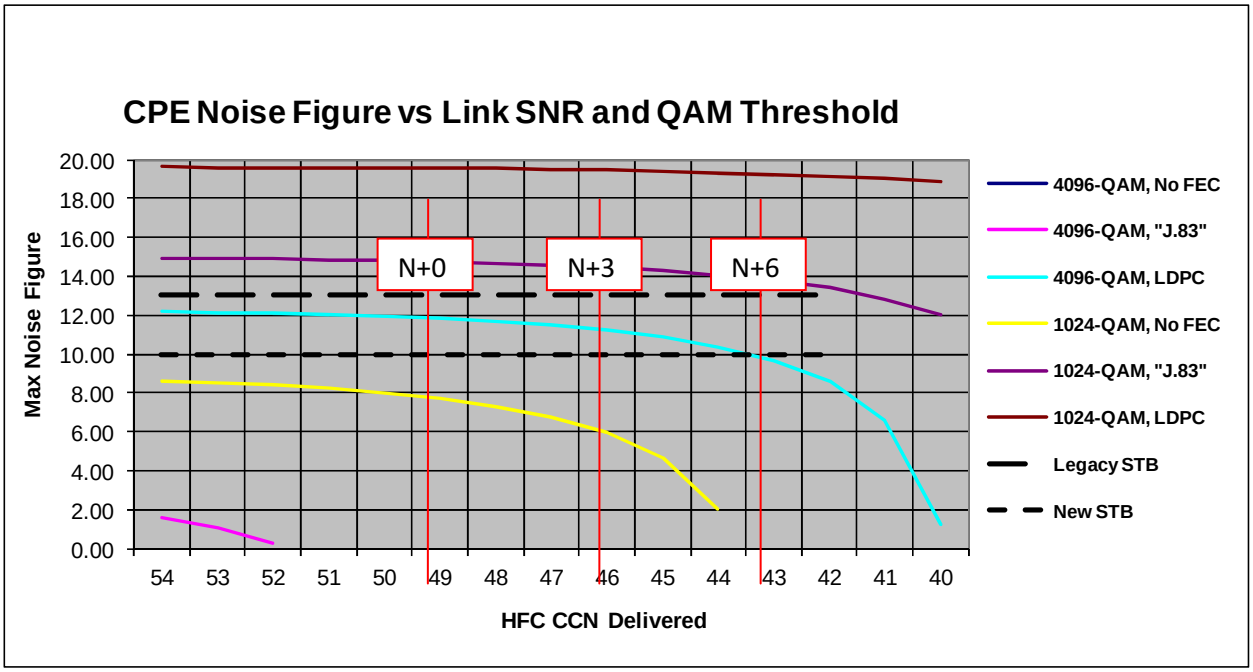

Figure 4 – STB Noise Figure Limit vs. Modulation Efficiency, DRFI MER

Multi-Wavelength Optics

Let's update the architecture now to include full analog reclamation and wavelength division multiplexed (WDM)-based linear optics commonly implemented today in multi-service fiber distribution architectures. WDM tools are becoming very valuable as operators take on Ethernet and EPON-based business services, avoid pulling new fiber where possible, and consider more consolidation of hub locations. With downstream loads moving to more QAM carriage and away from analog, the optical nonlinearities that make heavy analog loads difficult to manage over WDM become less imposing, and reaches can be extended.

Table 5(a-c) show three cases: 750 MHz, 870 MHz, and 1 GHz. In each case, a 1550 nm, ITU-grid-based, Analog/QAM transmitter is shown under nominal link conditions. Performance is also calculated with an 85 MHz upstream mid-split (slightly reduced forward load). This is a likely upstream evolution path for operators looking to exploit the full capabilities of DOCSIS 3.0 while minimizing the imposition on the downstream spectrum. The range of CCN's in Table 5 is within the ranges shown in Figure 3 and Figure 4, so these results are entirely applicable to the cases shown originally using Table 4.

Compared to classic single wavelength 1310 nm delivery, advanced optics such as these are being implemented more often, and across a variety of service scenarios, so extensive effort has gone into characterizing them across load and band variations. The extended calculations of Table 5 offer a few interesting conclusions with respect to variations in performance.

1) As the architecture shortens to N+0, the CCN going from 30 analog to zero analog improves. Accounting for the QAM-relative 6 dB de-rate with analog, for example, would leave 45 dB of QAM CCN for 750 MHz and N+0 (5-42 MHz system), instead is 47 dB. This is indicative of the effect of the analog carriers, which are higher levels than the QAM, to have a major impact on the distortion mix because of the difference in channel power. As a result, without taking advantage of any load, a couple extra dB are available. This does not included any additional dB that may be available from allocating more per-QAM power as the analog load is removed. As previously discussed, there is potential for another 2-2.5 dB on the optical link (flat loading) available that would keep the total power load the same. The tilted RF link would see less benefit.

2) The 85 MHz architecture has minimal impact on QAM CCN (0-1 dB). It removes a small chunk of forward bandwidth, but not enough to have a measurable impact. This may change at 200 MHz or more of upstream bandwidth. That case is not shown here, however, as in that case we would also expect an extended forward band. The net result should benefit CCN, since sliding the entire band results in fewer octaves of coverage, and the number octaves is important for broadband RF distortion characteristics.

3) 750 MHz systems have a 1-2 dB of SNR compared to 1 GHz systems. This is a worthwhile amount of dB gained, but perhaps not a good tradeoff relative to the capacity lost for not having the spectrum available.

Table 5 – Performance vs. Architecture a) 750 MHz, b) 870 MHz, c) 1 GHz

		750 MHz				
		30 Analog				All QAM
		CNR	CSO	CTB	CCN	CCN
Return	N+6	49	61	66	48	41
5-42 MHz	N+3	51	63	67	50	43
	N+0	51	64	67	51	47
	N+6	49	61	66	48	41
5-85 MHz	N+3	51	63	67	50	43
	N+0	52	64	67	52	48

		870 MHz				
		30 Analog				All QAM
		CNR	CSO	CTB	CCN	CCN
Return	N+6	48	61	66	47	41
5-42 MHz	N+3	49	63	67	49	43
	N+0	50	64	67	50	46
	N+6	48	61	66	47	41
5-85 MHz	N+3	50	63	67	49	43
	N+0	50	64	67	50	47

		1 GHz				
		30 Analog				All QAM
		CNR	CSO	CTB	CCN	CCN
Return	N+6	47	61	66	46	40
5-42 MHz	N+3	48	63	67	48	42
	N+0	49	64	67	49	45
	N+6	47	61	66	47	41
5-85 MHz	N+3	48	63	67	48	42
	N+0	49	64	67	49	46

Home Architecture Evolution

New CPE are taking advantage of full band capture (FBC) A/D architectures, which avoid pre-digitizing tuners that can contribute to RF degradation and simplify CPE designs. It will also make them more flexible to evolve moving forward. A low noise amplifier (LNA) precedes the A/D conversion to achieve the necessary levels to efficiently operate an A/D. This architecture is shown in Figure 5.

Analog-to-Digital converters themselves are inherently high noise figure components for typical high-speed bit resolutions because of unavoidable quantization noise. Nonetheless, this architecture does offer added flexibility for front-end sensitivity, and systems are easily optimized by choosing an external LNA, with the effects straightforward to calculate and not frequency dependent.

A quick, nominal, example using Figure 5 illustrates the simplicity:

- Assume a 20 dB gain LNA with a 5 dB NF
- Automatic Gain Control (AGC) amplification to drive A/D converter
- 12-bit A/D (at least 11 effective bits)

A well-design input cascade (LNA + AGC + A/D), can maintain a net NF of 6-7 dB for low input signal levels (where the NF comes into play the most). Thus, with input losses such as diplexers, and design focus on achieving higher modulation efficiency, NF of 8-9 dB could be the next level of sensitivity in new CPE, slightly better than the 10 dB range available today.

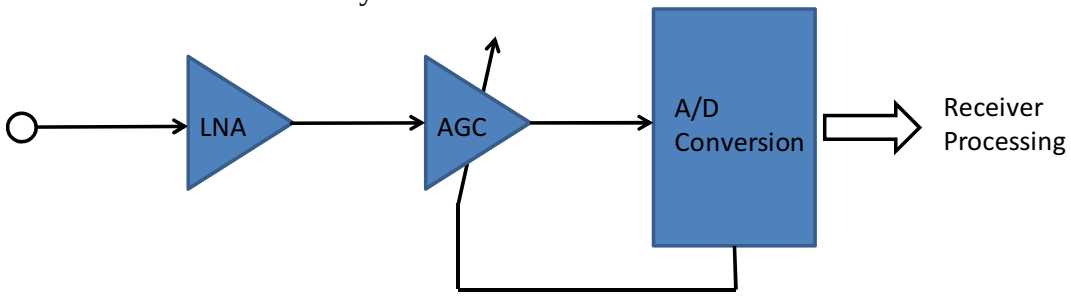

Figure 5 – Full-Band Capture Receiver Architecture

More important than the details of the receiver architecture, however, is that as part of the IP transition, operators are seriously considering investing in the next generation of home gateways based on a point-of-entry (POE) concept. Today, every subscriber is inherently *part* of the HFC access network, which makes for unpredictable results and ultimately money spent on truck rolls. The POE concept would have the cable drop go directly to an IP gateway (legacy support capabilities TBD). This IP gateway would completely abstract the inside of the home from the access network, and use only Home LAN interfaces to deliver content around the home – MoCA™, WiFi, and/or Ethernet interfaces delivering the bits over the last 100 feet. This has valuable benefits to RF losses in and out of the home for QAM receivers and upstream transmitters.

For the downstream, it amounts to benefits to the receiver SNR contribution, which can be substantial and meaningful at low input levels when advanced modulations are considered, as we have seen in Figure 3 and 4. Consider, for example, 20 dBmV tap port levels, 100 feet of drop (RG-6), 4-way splitting in the home, and 50 ft coaxial runs in the home (RG-59). At 1 GHz, we are losing RF power quickly:

STB RFin = 20 – 7 – 7 – 4 = 2 dBmV (virtual) or -4 dBmV.

If a few things break differently, the RF input will drop and the sensitivity of the receiver tested. A secondary splitter (-4 dB), extra drop length (-3.5 dB), and 15 dBmV design levels could challenge this link budget entirely, and house amplifiers may be called into play.

Now let's consider that all of the in-home loss is eliminated except for one splitter (an assumption for legacy considerations), as a POE architecture is apt to look. The loss is now the drop and one splitter, or 11 dB, a 7 dB savings. We can expect receiver NF degradation as the AGC design dials in more attenuation, but this is small relative to the improvement in signal power, so a higher SNR is obtained from the CPE.

Next, we consider the combination of the FBC architecture leading to lower noise CPE, and the POE concept, together as a "next generation" home architecture opportunity. By recalculating Figure 3 with 7 dB more of QAM power, and add a line representing the potential decrease in NF, we can recalculate NF margin to the various modulation profiles. This is shown in Figure 6.

We can quantify an example of possible NF degradation based on the Figure 5 architecture for this increased input level. Using a very simple front end cascade design, the degradation in NF calculated is less than 2 dB. It is probably much less than this, but this offers a boundary for particularly simple Figure 5 architecture. This still means at least a net SNR gain of 5 dB for the CPE contribution to the total SNR. A "degraded NF" case based on the above calculation would be identical to the 10 dB representing today's performance in Figure 6.

On Figure 6, the cascade depth marker lines used are the HFC CCN values taken from the fully loaded 1 GHz case, complete analog reclamation, and an 85 MHz Mid-Split upstream – i.e. Table 5c, orange shaded values.

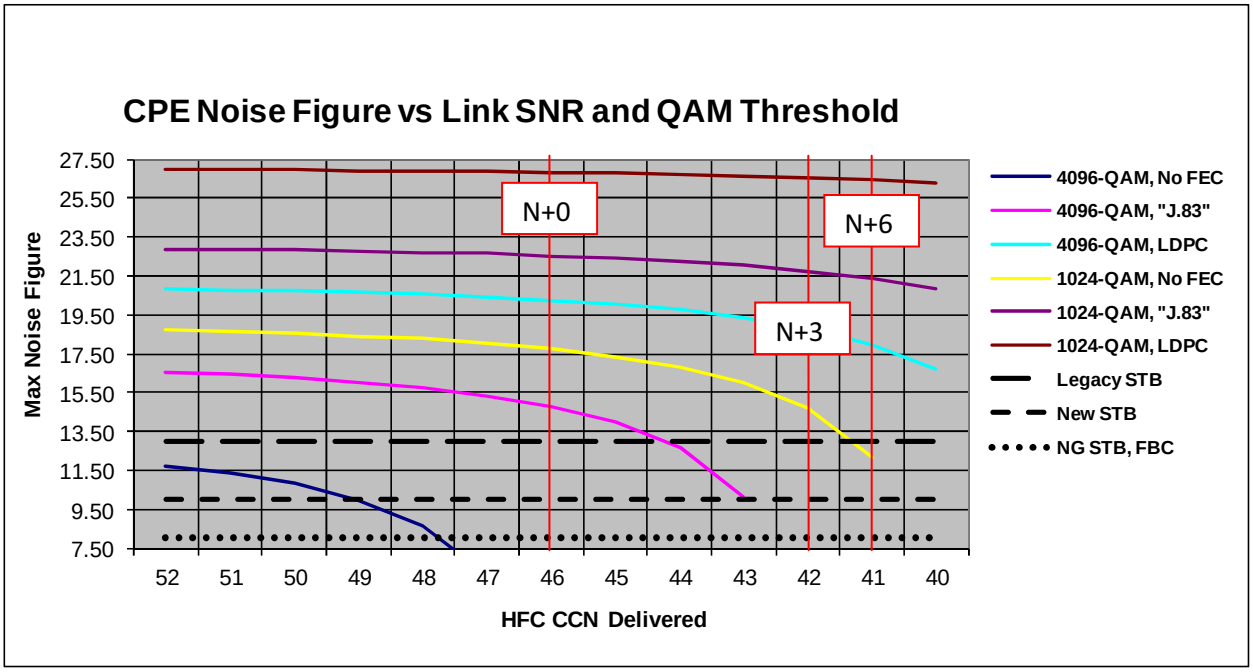

Figure 6 – STB NF Limit vs. Modulation Efficiency, POE Gateway

It would appear in Figure 6 that a tremendous amount of new margin has been created, and this certainly is the case with respect to the access network and home environment's ability to deliver the highest modulation profiles. The results suggest, for example, that "J.83" style 4096-QAM (pink) can be supported, at least on N+0 or equivalently high performing HFC links. Of course, now we are likely to see the biggest impact to the contribution of today's DRFI MER of 43 dB. The impact of this reality is shown in Figure 7.

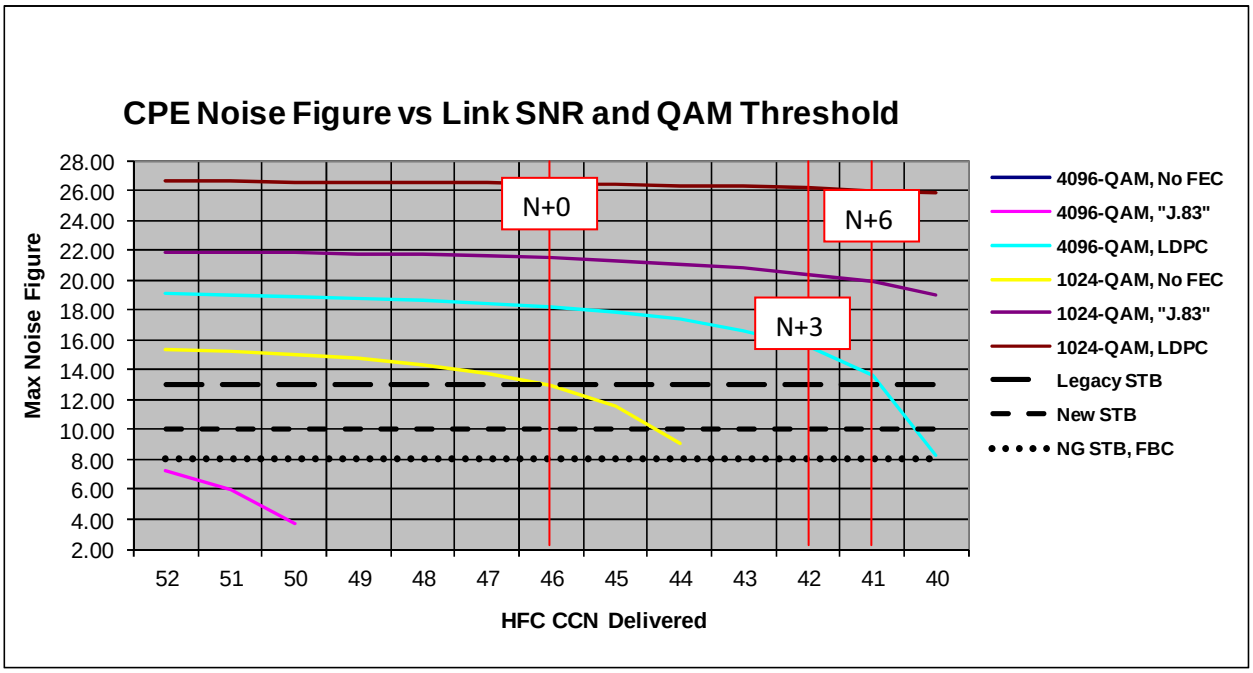

Figure 7 – STB NF Limit vs. Modulation Efficiency, POE Gateway, DRFI MER

We can draw the following conclusions when observing the full picture in Figure 7:

1) 1024-QAM is very comfortably supported from an SNR perspective with old or new FEC, cascade depth or forward band plan
2) Robust 4096-QAM would *require* advanced FEC, and a short HFC cascade – less than N+3 preferably as the curve of support is rapidly becoming sensitive to performance variation and running out of margin as the cascade lengthens.

It seems intuitive that 4096-QAM would require an updated FEC to be operational. Indeed, had it not been so, system architects would have previously considered increasing modulation profiles, as this would have indicated that HFC performance was sufficient. We can come full circle by considering the following:

1) HFC links were designed originally for analog video
2) Analog video, CNR delivered: ~45 dB
3) Digital CNR for a 45 dB analog: 39 dB
4) 4096-QAM with LDPC threshold (Table 2): 37 dB

This illustrates why we are now speaking of 4096-QAM as within range for HFC delivery, and in particular highlights the value of the advanced FEC in making this so. It also illustrates why 1024-QAM in "J.83" style is already on the verge (36 dB) [9].

Figure 8 plots one more example network architecture evolution – in this case removing the linear optical component and distributing the RF generation to the HFC node via digital optics. The assumption in Figure 8 is that this could be accomplished while maintaining DRFI-compliance (43 dB MER). Without the linear optics variation, of course, the "curves" are only a function of the RF cascade depth.

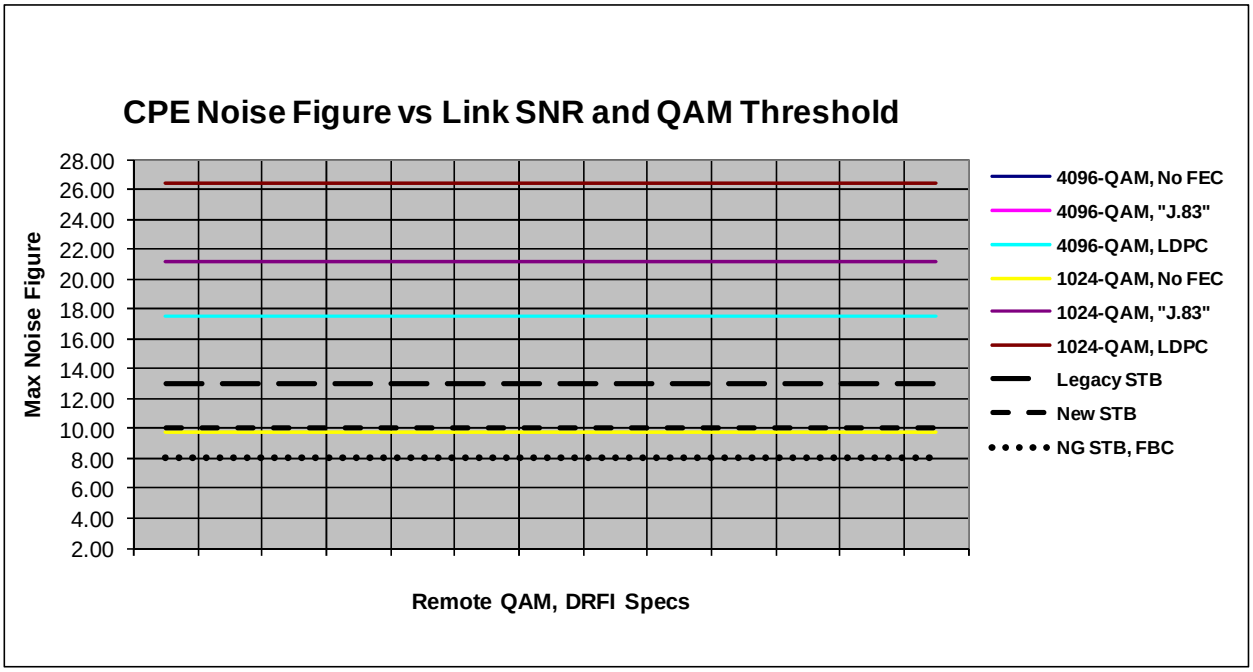

Figure 8 – STB NF Limit vs. Modulation Efficiency, POE Gateway, Remote DRFI, N+3

In Figure 8, we have assumed an N+3 cascade, using the contribution as calculated by comparing the N+0 and N+3 cases in Table 5c, and subtracting the difference. This calculation has some favorable uncertainty in it, because an "N+0" in fact includes the RF chain of the node itself, and the difference between N+0 and N+3 does not capture the effect of these gain blocks. However, Figure 8 gives a ballpark estimate of the performance of a remote QAM with DRFI performance driving an HFC cascade, and the ability this architecture has to support advanced modulation profiles. It also uses the assumptions of Figures 6 and 7 with respect to receiver and home architecture.

HFC UPSTREAM

Turning our attention to the upstream, the SNR threshold assumptions that will be used to illustrate capabilities were shown in Table 3.

In Figure 9, we show today's state of the art for a linear optical upstream. The ability to support 256-QAM upstream over 85 MHz mid-split architectures has been proven in the field, where throughputs of 400 Mbps were obtained [12, 19]. It was shown that DFB optics, coupled with higher sensitivity, higher fidelity DOCSIS 3.0 receivers, could robustly support a fully loaded 85 MHz upstream. A 12 dB dynamic range of sufficient NPR is shown. Dynamic range (DR) in the upstream is much more important than in the downstream. Network design is optimized and aligned precisely in the downstream, while upstream design and environmental variations, alignment techniques, as well as unpredictable RF channel conditions, require an SNR to be met over a range of input levels. Historically, DR on the order of 10 dB has been sought.

Measured packer error rates (PER) were taken at various input levels on the N+3 cascade used in Figure 9. These points are shown in Figure 9 where the yellow marker dots are along the blue noise power ratio (NPR) curve. These are where low PER was observed. The yellow dots on yellow trace are representative of DFB performance of transmitters such as many in the field today. The noise power ratio analysis of Figure 9 makes plain why robust performance was observed. The measured points clearly fall within an area of high NPR and SNR, where good performance would be expected, and with solid 6-7 dB of peak margin extending beyond the 256-QAM threshold identified.

Note that in all figures below, thresholds identified and not labeled as "LDPC" are the assumed "DOCSIS" thresholds of Table 3 (i.e. not the "No FEC" thresholds).

In Figure 10, we introduce the new PHY performance thresholds with advanced FEC from Table 3. In addition, we have included the net performance of the RF portion of the HFC by introducing a deep, combined amplifier cascade (orange). Because the return amplifiers contribute high SNRs individually, the effect here is minor, but we will see how this contribution may also increase as optics improves. Of course, over time, the expectation is that the RF cascade will shorten as well. Or, if the segmentation is virtual (combining in the node is removed), fewer amplifiers will "funnel" upstream. The net effect is the same – fewer amplifiers contributing noise to the return path, reducing the input noise power at the receiver and thereby increasing SNR.

It is clear from Figure 10 that "DOCSIS" coding for 1024-QAM (dashed red) would not be able to be supported from an SNR perspective. This threshold is simply above even the peak NPR. For 1024-QAM enabled with LDPC FEC, however, we can see that the threshold of operation is now exceeded by the combined HFC+CMTS link.

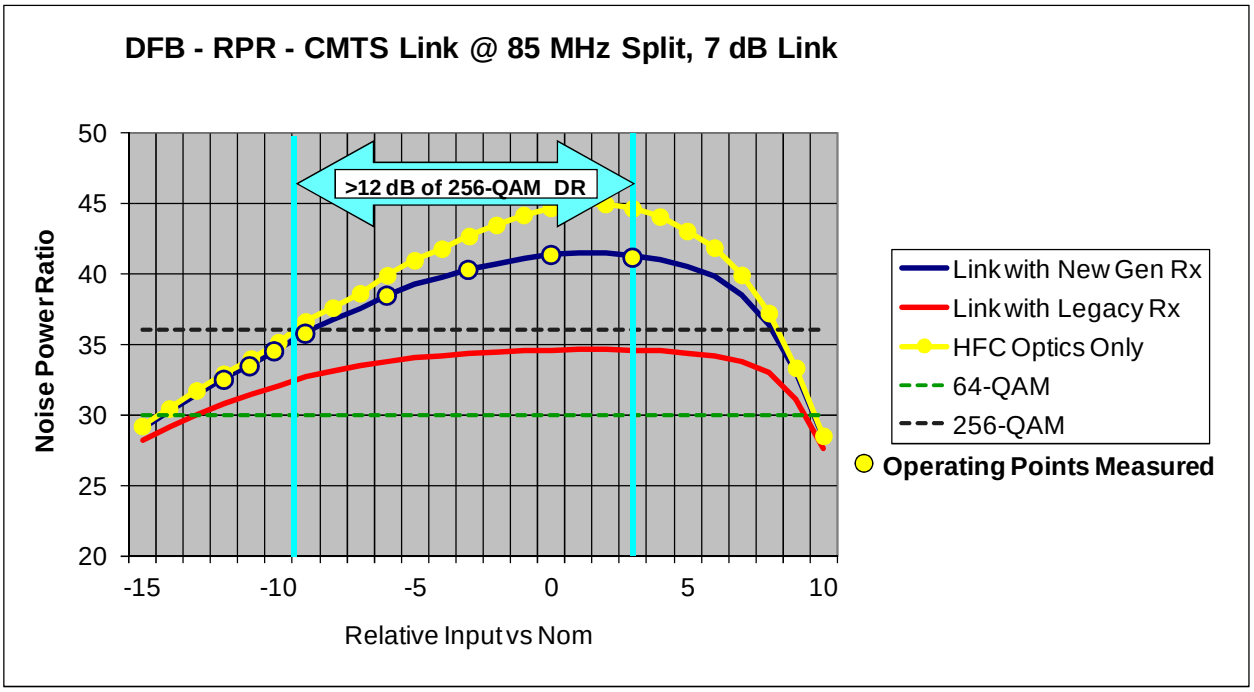

Figure 9 – 256-QAM "DOCSIS" over 85 MHz DFB Return Optics

Unfortunately, the dynamic range is reduced on the left hand side by 3 dB. It is also well understood and documented that, as QAM profiles become increasingly dense, the right hand side (soft distortion components, red arrow) also have a more significant impact, reducing dynamic range from the right. The above reduction from the right is an estimate based on prior work, which only considered up to 64-QAM [23].

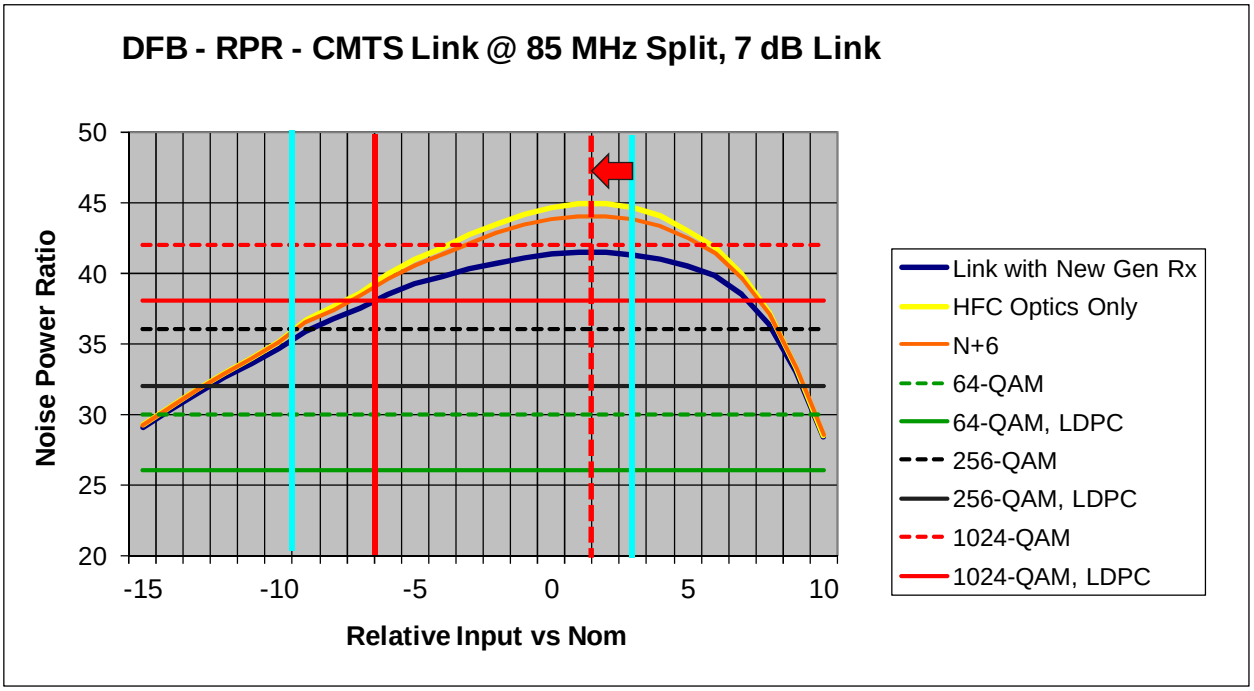

Figure 10 – 64/256/1024-QAM, "DOCSIS" and New FEC, 85 MHz DFB Optics

The net effect of the reduced range (8 dB DR) is that 1024-QAM with LDPC would be marginal in practice. It would likely work in some places, but inconsistently throughout a footprint without further network improvements. An analogous situation is the introduction of DOCSIS 3.0 64-QAM using FP upstream laser technology. FPs exhibit reduced performance compared to DFBs, and thus eat into the DR margin acceptable to run DOCSIS 3.0 64-QAM. Such is the case in Figure 10 for 1024-QAM with LDPC. The operating window is small, and this would likely be reflected in inconsistent performance.

In summary, with LDPC, it would be possible to get 1024-QAM working on well-behaved upstream channels that are aligned properly for laser input level, and only using new DOCSIS 3.0 upstream receivers with higher sensitivity. However, in practice, over a large range of plants, performance would be unreliable (and impossible at all if laser technology has not been updated and legacy receiver performance exists).

In Figure 11, we extend the Figure 10 case to a 200 MHz "high split" return – a long term approach under consideration by the industry to deliver more upstream capacity [2]. An implicit assumption is that the 200 MHz receiver could perform equivalently from a sensitivity standpoint as today's DOCSIS 3.0 receiver. Nonetheless, the pure power loading makes 1024-QAM completely impractical. The 256-QAM case without new FEC is now marginal, as the 1024-QAM case in Figure 10 was – in fact, it exhibits the same DR and signature relative to the performance threshold (highlighted in red). However, it is likely to be somewhat more robust that 1024-QAM simply because 256-QAM will be less sensitive to the types of things that the allocation of DR and margin is meant to protect against.

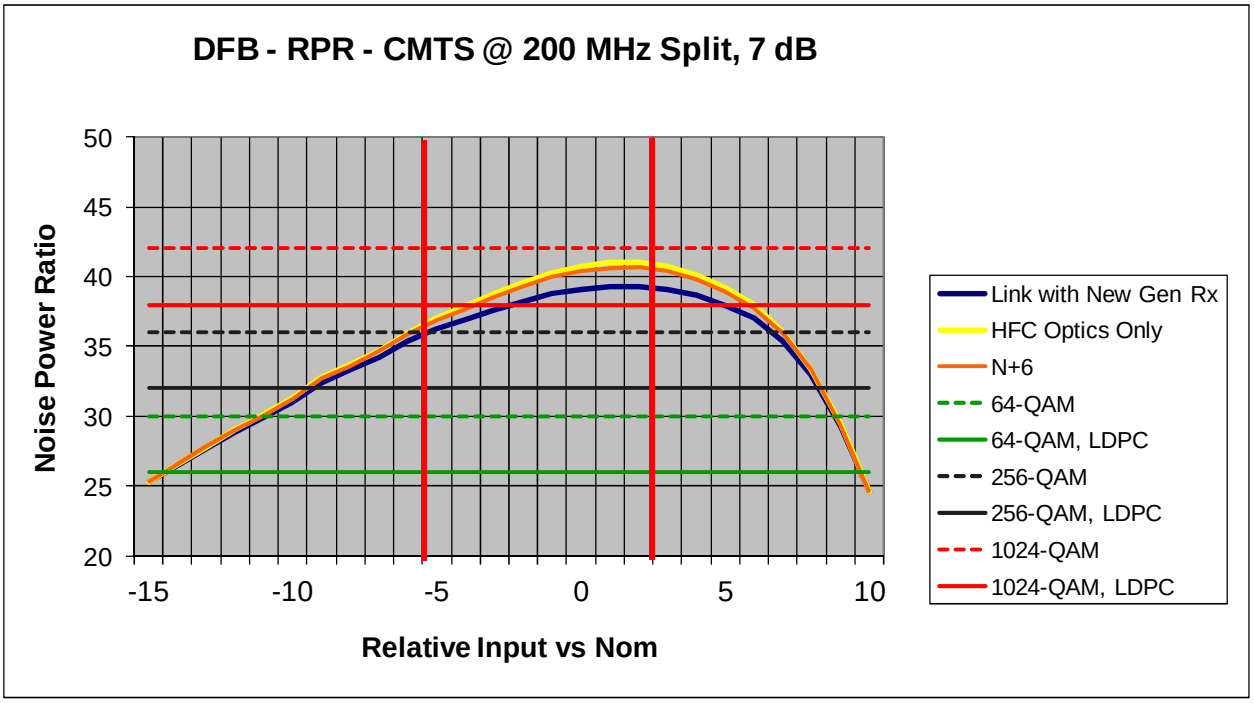

Figure 11 – 64/256/1024-QAM, "DOCSIS" and New FEC, over 200 MHz (Projected)

Enhanced Linear Optical Performance

The analysis in Figures 9-11 uses performance of DFBs like many that may be in the field, perhaps lower power (1 mw), where the upgrade from Fabry-Perot lasers (FPs) had already taken place. Today, because of the rising interest in mid-split architectures and more bandwidth efficient modulation, new development activity is focusing on analog and digital return solutions that optimize performance for extended bandwidths. Continued improvements in noise performance of both transmitters and receivers have also occurred over time. Recently measured performance of a mid-split bandwidth "DFBT3" (Motorola model, temperature compensated, 2 mw) to a return path receiver is shown in Figure 12.

In Figure 12, the performance curves from Figures 9 are also shown. A notable improvement over time in peak NPR is observable, and an associated dynamic range increase for a given threshold. Peak performance of 50-51 dB is observed – again making the point that a good DFB return path transmitter-receiver link looks very much like a 10-bit digital return link [16].

Also shown is this improved performance when combined with a DOCSIS 3.0 receiver (dashed blue). Here, it becomes clear that as the optics improves, the influence of receiver SNR contribution begins to have a larger effect. We will discuss this further later in this section.

In Figure 13, the DOCSIS and "next gen" thresholds using advanced FEC are evaluated against the improved mid-split performance identified and measured in Figure 12. Note that the performance shown is the linear optical return (yellow) only, along with the DOCSIS 3.0 receiver (blue). Thus, this is equivalent to an N+0 case, since there are no RF noise contributions from the plant in Figure 13.

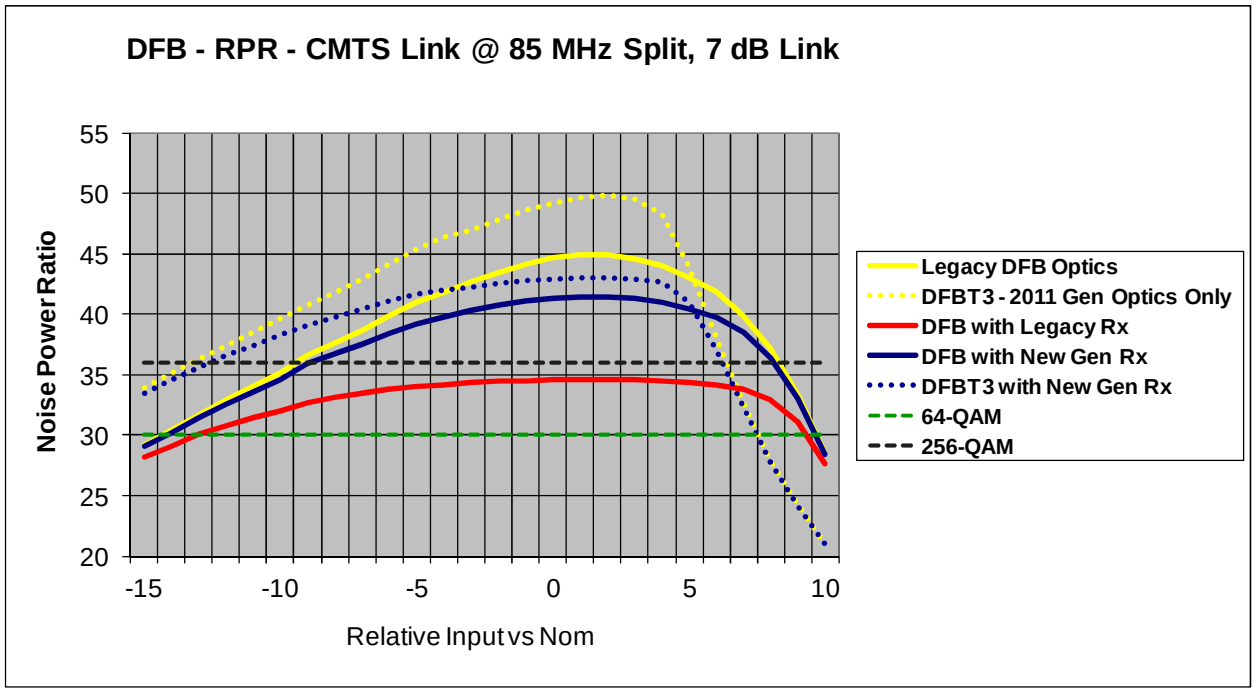

Figure 12 – Measured Mid-Split Performance, Modern DFBT3-RPR Link

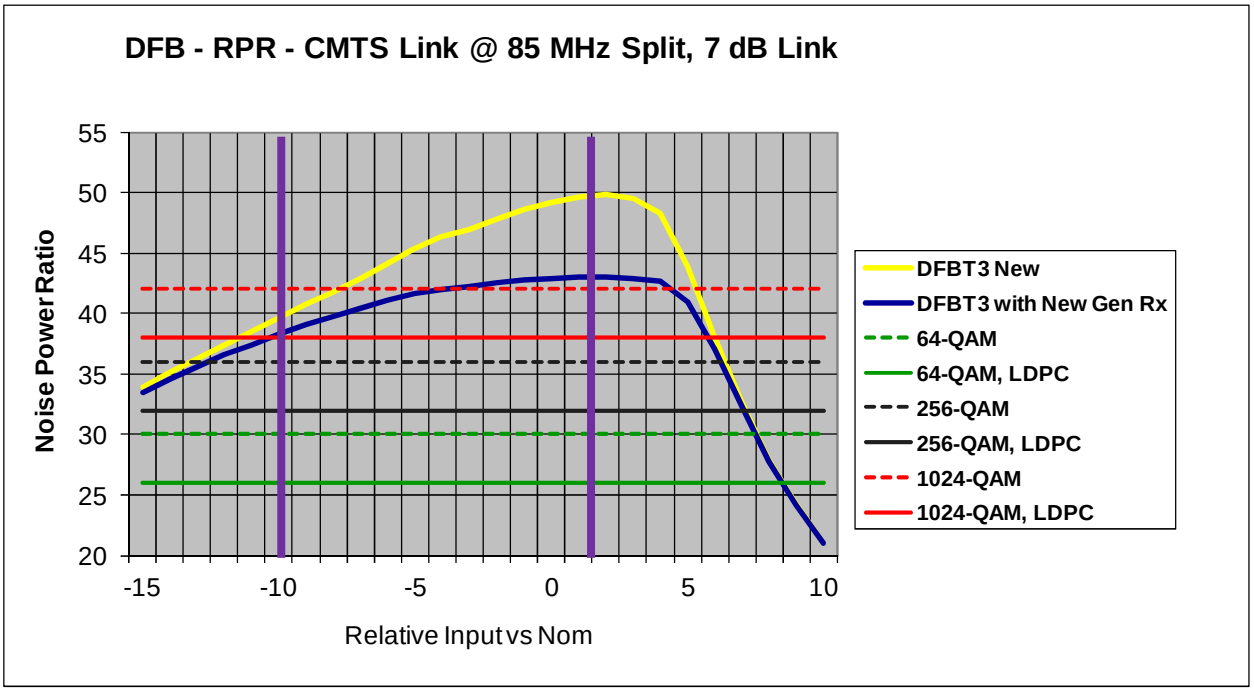

Figure 13 – High Performance DFBT3-RPR Mid-Split, All Thresholds

All of the thresholds are met in Figure 13, and the 64-QAM through 256-QAM cases comfortably supported. This is to be expected since this has been proven to be the case for 256-QAM today without any new FEC applied or improved HFC return performance.

For 1024-QAM, however, clearly "DOCSIS" PHY coding will not be sufficient using the Table 3 assumptions. The advanced FEC applied to 1024-QAM, however, does show promise. Indeed, it shows 11 dB of DR (purple markers) above the threshold identified for 1024-QAM. However, note that the margin above the threshold never exceeds 5 dB, and for half of the DR it is 4 dB or less. The receiver noise contribution, ably supporting 256-QAM (against a 64-QAM-maximum DOCSIS requirement, it should be added), comes into play here.

While the dynamic range appears robust, the low margin across the full DR range may make performance less robust than such a DR would otherwise indicate considering we are dealing with a more sensitive QAM profile. In other words, "peak" margin may need consideration as we move to more complex QAM profiles. Return path alignment is based on setting composite signal levels around a "sweet spot" close to where NPR is at its peak, but allowing for some margin of back-off to avoid the clipping and distortion region on the right hand side. At the very least, it would seem reasonable that the same peak margin available today for 256-QAM (6-7 dB, Figure 9) would be a good starting point objective for higher order scenarios like 1024-QAM.

Today's expectation of DR is built around volumes of 16-QAM and 64-QAM deployments, exclusively. Since the DR for 1024-QAM with LDPC in Figure 13 is about the same as we observed in Figure 9 for 256-QAM, we might expect 1024-QAM with LDPC to be operational under good upstream

conditions. Similarly, since the peak margin available is reduced across its DR compared to Figure 9, and because other contributions not captured by NPR will effect 1024-QAM worse than 256-QAM, its performance is likely to be less rugged than the 256-QAM case in Figure 9.

Figure 14 adds an RF noise contribution assuming a deep cascade (N+6) and combined four ways. The degradation due to the quantity of RF amplifiers is seen (orange). However, it is clear in observing the blue curves – (optics + receiver) – with and without RF noise contributions from the cascade, that the limitation about 256-QAM is in the noise contribution of the upstream receiver, at least for nominal architectures and levels as they are implemented today. And, again, we are quantifying a receiver for 1024-QAM that had a requirement to meet 64-QAM performance objectives.

In summary, from an NPR/SNR perspective, with new LDPC-based FEC, 1024-QAM is possible on high performance upstream optical links. However, it may be operationally less robust when considered across a range of possible channel environments given the decreased peak margin and exaggerated sensitivity of 1024-QAM to other link impairments not captured by NPR analysis. It is hard to be certain, but we can get a window into the performance consistency through the ruggedness with which 256-QAM becomes implemented. 1024-QAM with LDPC should be "like" this but somewhat less robust. As might be expected, we will need to ask more of next generation receivers than sensitivity supporting 64/256-QAM if implemented over linear optical returns or today's vintage of digital returns.

In Figure 15, we apply the Figure 12-14 performance across a 200 MHz upstream. The assumptions are that identical optical noise performance can be achieved (shared over a wider bandwidth) and identical receiver sensitivity can also be achieved.

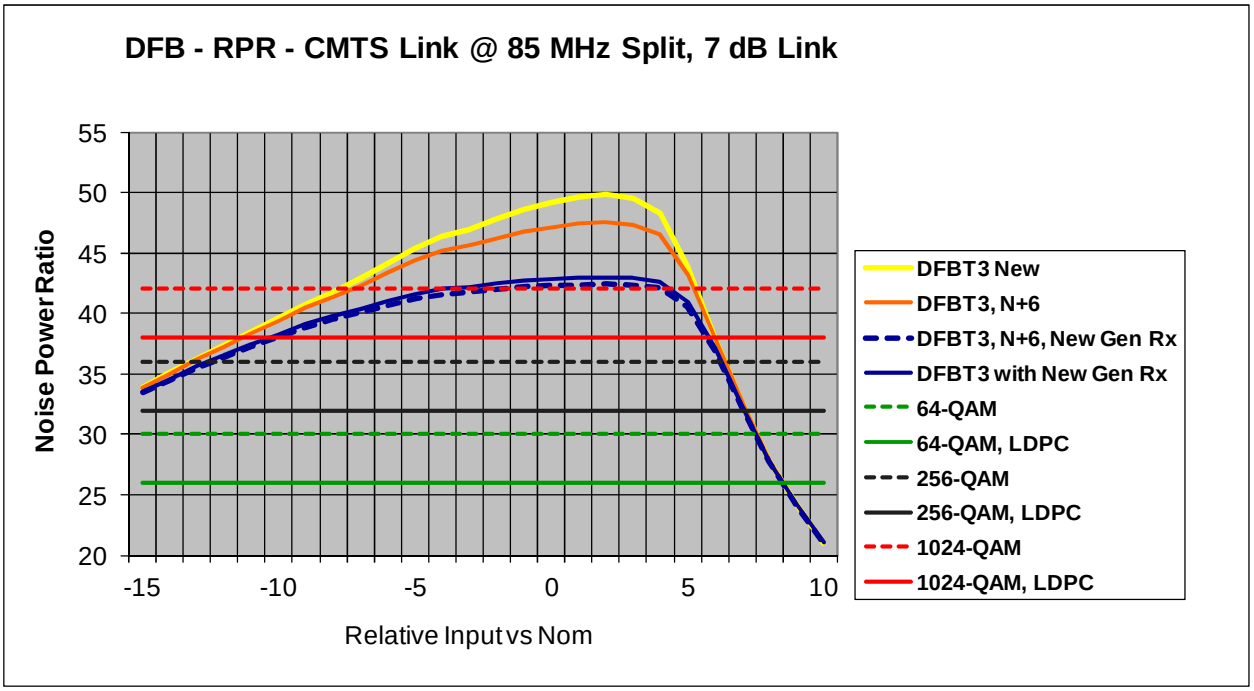

Figure 14 – High Performance DFBT3-RPR Mid-Split, N+6, All Thresholds

An encouraging conclusion is that, for this extended bandwidth case, 256-QAM of the "DOCSIS" variety appears to be operational at acceptable DR and a reasonable peak margin, at least under these assumptions of the wider band design. Since we do not have extensive field lessons for 256-QAM, the peak margin question raised for 1024-QAM could apply in this case as well. However, to be at least within reach of 256-QAM already as it exists today on an extended bandwidth return, based on today's linear optics, is an excellent side for the future of new broadband capacity for the upstream.

The 1024-QAM case now clearly has insufficient DR (purple) as well as small peak margin from the threshold – 4 dB maximum and less than that across the DR.

Figure 16 shows an enhanced return performance example, in this case based on digital return. Key advantages of digital return include ease of setup, and, from the perspective of this paper, NPR performance independent of link length. The return path performance of the digital return approach is almost entirely determined by the A/D converter resolution. Finding A/D converters that provide a high *effective* number of bits (ENOB) as bandwidths expand is the limiting performance component. The potential to require redesign for each bandwidth increasing increment is a disadvantage [16] of this approach.

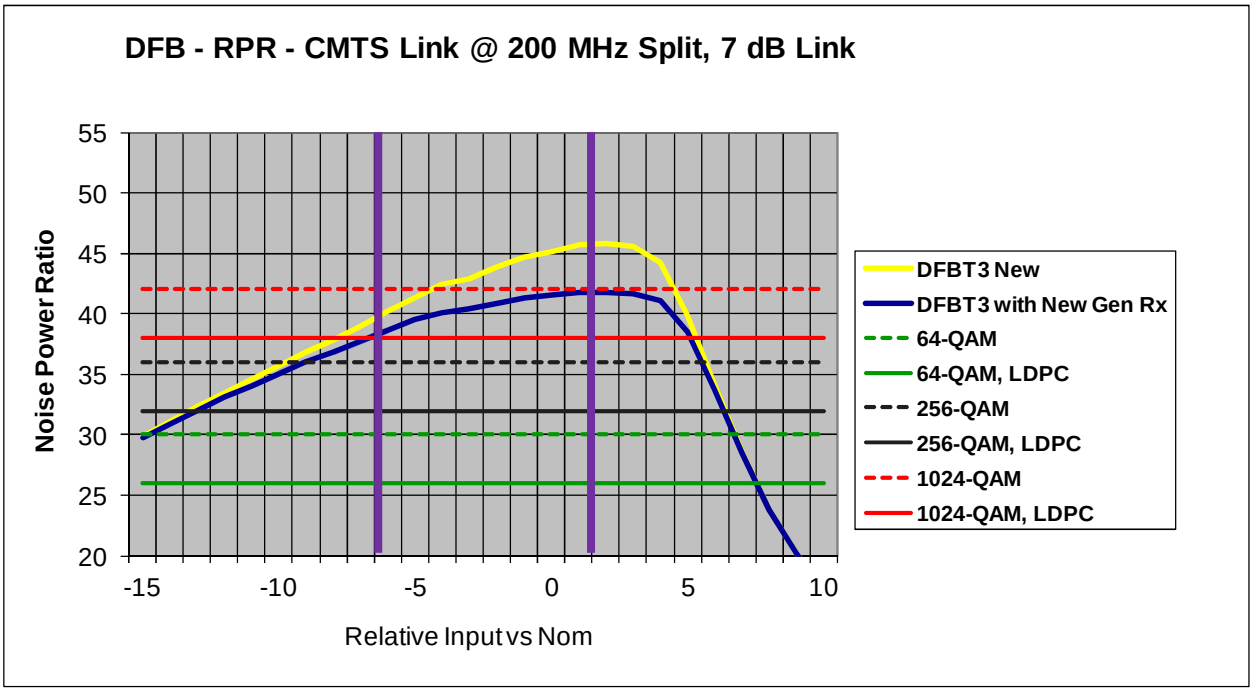

Figure 15 – High Performance DFBT3-RPR "High" Split, All Thresholds

Figure 16 uses measured performance of a solution that behaves like an ideal 10-bit A/D, or, equivalently, has an ENOB of 10 bits.

Three composite NPR curves that include RF and receiver contributions are shown with the digital return-only NPR performance:

1) Digital return, N+6 RF cascade, receiver (purple dash)
2) Low noise DFBT3-RPR link, N+6, receiver (yellow)
3) Original DFB-RPR (256-QAM analysis), N+6, receiver (blue)

The comparison indicates that the digital solution contributes a couple more dB of DR to the 1024-QAM with LDPC threshold on the left side of the NPR curve (the SNR side), which itself adds 4 dB of DR over the legacy solution. The extra couple of dB from the digital return could be the difference between robust or less robust. However, since the peak margin has not changed between the two (yellow and dashed purple), and is limited by receiver sensitivity, their performance will likely be similar – on the bubble of robust-enough margin.

A clear conclusion from these analyses is that support of 1024-QAM would require new FEC, and be aided by an improved receiver sensitivity. This is readily illustrated by observing the effect of a hypothetical receiver designed to support 1024-QAM (as opposed to DOCSIS 3.0, 64-QAM), with a receiver sensitivity improved 3 dB in so doing.

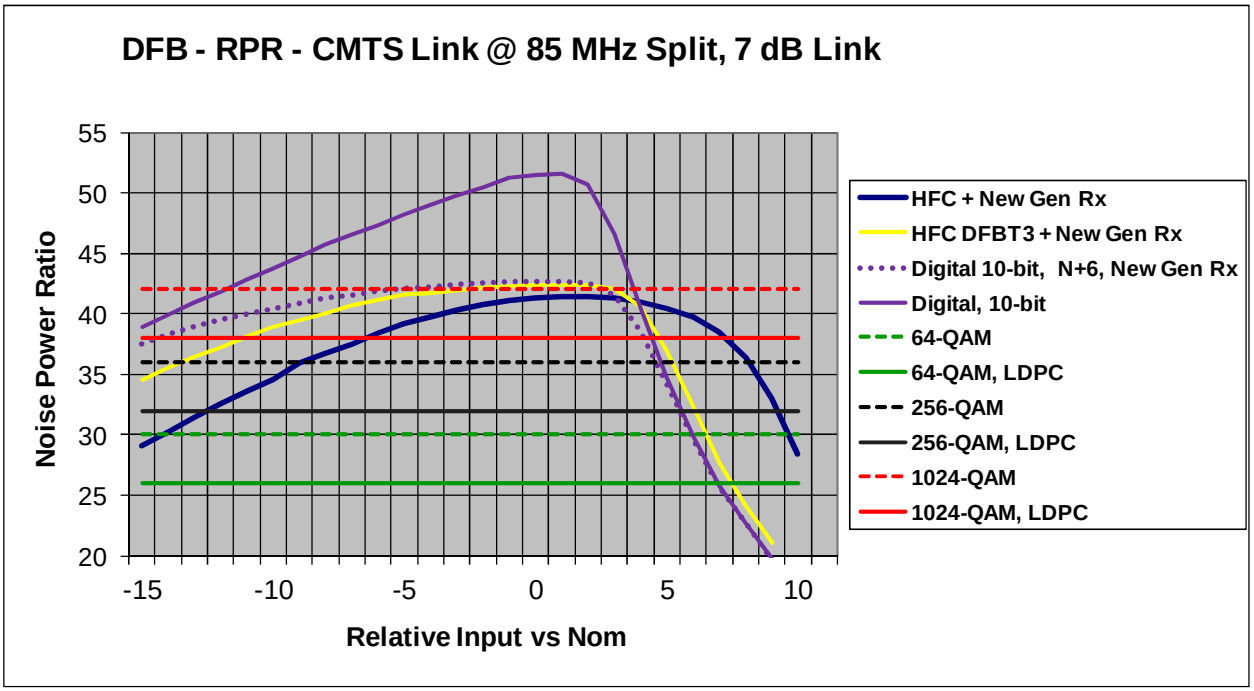

Figure 16 – 10-bit Digital Return & Linear Optics Cases, N+6, All Thresholds

Figure 17 shows the same composite NPR curves, minus the digital return-only one for this case of improved sensitivity. Figure 17 illustrates that for linear optical links or digital returns, it is clear that these would now have margin to support 1024-QAM with adequate DR. This could only be so if the 1024-QAM was accompanied by new FEC. And, there are still some open questions about whether peak margin standards should be considered, much of which may come as 256-QAM deploys under "DOCSIS" threshold conditions.

This leads to an overall conclusion that, for upstream, attention will need to be paid to upstream optical architectures, receiver performance, and possibly Headend architecture on the whole given that the optical receiver-CMTS connection today is a simple, almost forgotten pipe that can create low level inputs to CMTS ports and challenge their sensitivity. For all cases desiring to support 1024-QAM, LDPC-based FEC is a must-have, though itself not a sufficient condition.

Remote Demodulation

Future architectures may take advantage of distributed physical layers. The analogous concept in today's world is "CMTS in the Node." For a transition into new RF and IP technology or extension of DOCSIS, this may be easier to consider than it has historically been with DOCSIS. And, with Gigabit Ethernet and EPON optics available at low cost, it become very attractive to consider taking advantage of these standard interfaces, potentially eliminate linear optics from the plant, and improve performance all at the same time. Modular node platforms are now built to handle various plug-in optical and RF modules, so this approach is consistent with HFC node evolution.

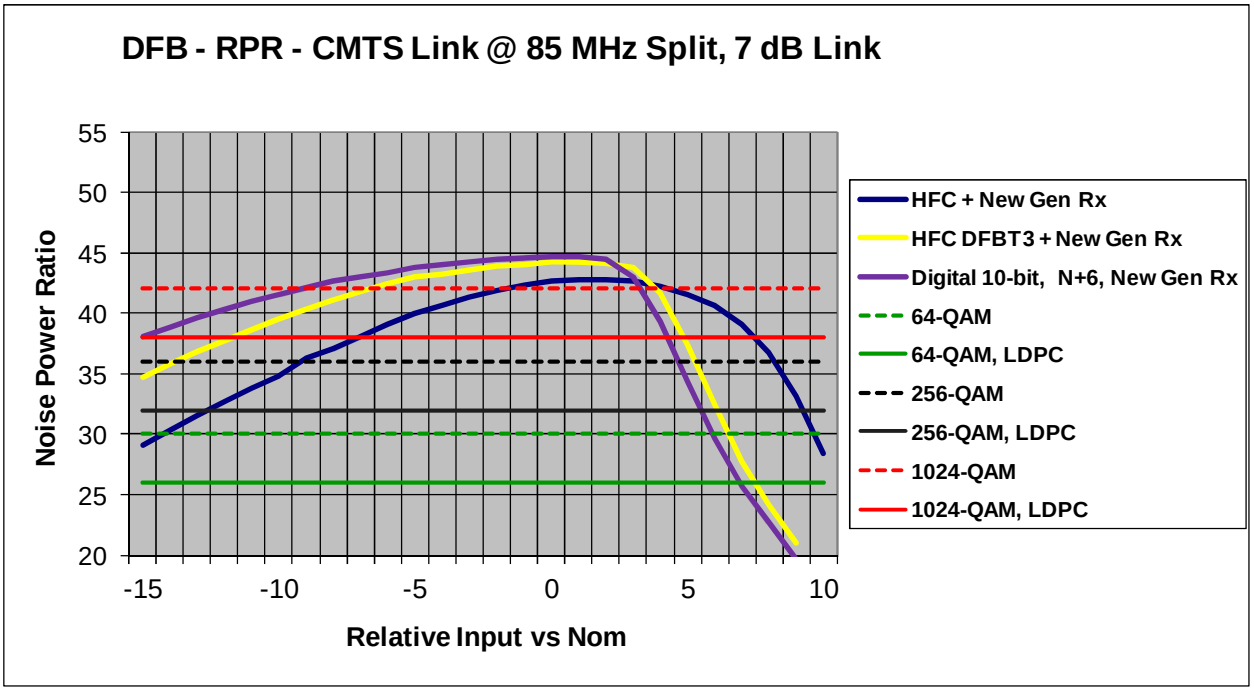

Figure 17 – All Cases of Return Optics, N+6, Improved Rx Sensitivity (3 dB)

If a remote receiver is placed in the plant, then we can simplify the analysis by removing the optical links from the equation. This can be shown using the same NPR curves as before, except now only the softly distorting amplifier limits the right hand side of the curves. While these can be characterized and specified, this is not usually done. It is nominally assumed that the upstream signal transmissions stay comfortably within the range of return amplifier linearity. The alignment of the total level of the load is critical at the upstream optical interface, but in the plant there is not the constraint of total power load to the extent that there is in the optics. Coupled with common noise figures of return amplifiers, the result is that very high SNRs are possible for a single amplifier relative to its noise contribution to the upstream cascade.

Instead of NPR, Table 6 shows a range of required Noise Figures of a module installed in node. We have made an assumption, based on our above discussion of peak margin above threshold for advanced QAM profiles, that a net SNR of 45 dB (7 dB of margin to 1024-QAM with LDPC) is the objective.

A range of port levels (low, mid, high) are shown, and the impact these levels would have on the remote receiver's NF requirement. The path loss between the coaxial input port and node module is accounted for, so this is the noise figure performance at the input to the receiver module. Provided low input levels do not reign, these are not particularly challenging. And, even at the lowest input level identified here, the NF is achievable with good design practices.

Table 6 – Calculating "Remote" Receiver Noise Requirements

Signal Level		**Signal Level**		**Signal Level**	
Ports	10.0	Ports	15.0	Ports	20.0
Node Path Loss	5.0	Node Path Loss	5.0	Node Path Loss	5.0
Node Combine	6.0	Node Combine	6.0	Node Combine	6.0
Noise		**Noise**		**Noise**	
Amplifier (NF=8)	50.2	Amplifier (NF=8)	-50.2	Amplifier (NF=8)	-50.2
Cascade	6.0	Cascade	3.0	Cascade	3.0
Combine	4.0	Combine	1.0	Combine	1.0
RF Noise	36.4	RF Noise	-45.4	RF Noise	-45.4
SNR Req'd	**45.0**	**SNR Req'd**	**45.0**	**SNR Req'd**	**45.0**
Terminating NF	**6.3**	**Terminating NF**	**16.9**	**Terminating NF**	**22.0**

SIGNAL-TO-INTERFERENCE

Single carrier techniques to combat narrowband interference amount to attempting to notch out the offender's band through an adaptive filtering mechanism, and recover the modulated carrier around it as effectively as possible. Because removing the interference involves removing signal spectrum that subsequently must be equalized and detected, the effectiveness of the process is reduced as the interference becomes wider band, or, for a fixed signal-to-interference energy (S/I), if there are multiple interferers to handle.

Test results have been observed and reported with respect to the A-TDMA DOCSIS upstream in separate studies in recent years [13,22]. Table 7 shows thresholds of uncorrectable codeword errors observed for 64-QAM.

Table 7 – Ingress Thresholds for 64-QAM A-TDMA Upstream

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%

In the study summarized in Table 8, recent results for 256-QAM for a fixed PER objective of 0.5% and 1% for a high SNR condition are derived through testing. The SNR condition applied (SNR = 36 dB) is consistent with the Table 4 assumption on a robust SNR threshold to use for 256-QAM in the upstream, and the analysis in [22] further identifies this high SNR as one that increasingly enables ingress cancellation.

Table 8 – Ingress Thresholds for 256-QAM A-TDMA Upstream

	256-QAM				
	Level (dB, dBc)	UNCORR%	CORR%	PER%	MER (dB)
Baseline - AWGN	**36**	0.000%	0.000%	0.000%	37
Single Ingressor Case					
QPSK 12kHz 0.5%	3	0.254%	0.435%	1.060%	34
QPSK 12kHz 1.0%	1	0.447%	0.944%	2.300%	34
FSK 320ksym/s 0.5%	29	0.278%	0.032%	0.110%	35
FSK 320ksym/s 1.0%	27	0.633%	0.230%	0.810%	35
FM 20kHz 0.5%	2	0.128%	0.295%	0.750%	34
FM 20kHz 1.0%	1	0.187%	0.554%	1.260%	34
Three Ingressor Case					
CPD 0.5%	28	0.297%	0.041%	0.190%	34
CPD 1.0%	27	0.698%	0.144%	0.750%	33

Tables 7 and 8 are valuable indicators of the performance of ingress cancellation (IC). However, to understand how well it is working, it helps to know how robust the individual QAM profiles are to signal-to-interference ratio (S/I) to begin with. A system simulation was performed to evaluate this sensitivity. An example of 64-QAM with a single CW interferer is shown in Figure 18, while a 256-QAM example is shown with three non-coherent interferers in Figure 19. The familiar CW "donut" pattern is clear in Figure 18.

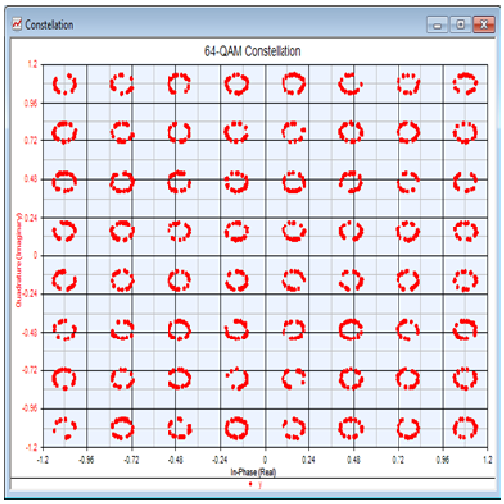

Figure 18 – 64-QAM with a Single CW Interferer

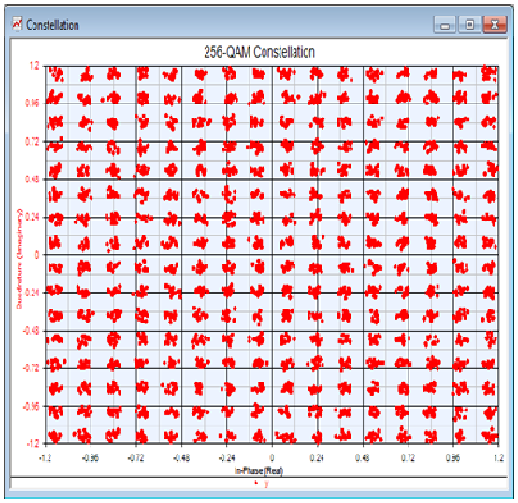

Figure 19 – 256-QAM with 3 CW Interferers

Both cases were evaluated to find thresholds of correctable low error rate. We use the more challenging three-interferer case as a reference. A subset of these simulation results are summarized in Tables 9a-d for 64/256/1024/4096-QAM. The modeling tool is described further in the Appendix.

Obviously, more dense profiles are more sensitive to S/I. The relationship of interest is to note that, for a given "DOCSIS" Table 3 SNR threshold, high enough to have robust performance and allow the IC to work, the relative S/I difference across profiles is also approximately 6 dB for when errors begin to be counted.

Now consider the IC performance based on Table 7 MER before and after IC. It can be calculated as providing roughly an effective 26-28 dB of cancellation for the case of multiple 20 kHz interferers. An estimate for the Table 8 case for 256-QAM and 0.5% PER using 256-QAM data at SNR = 36 dB, from the analysis table shown in Table 9b (est. 28 dB S/I), is that the IC is providing about 26 dB (S/I =2 to S/I = 28) of IC for a single 20 kHz interferer. Table 7 suggests that if the total power of the interference is the same, and it is narrowband, IC performance is close to the same for one interferer or three.

Table 9a-d – Ingress Thresholds for Uncoded QAM
a) 64-QAM b) 256-QAM c) 1024-QAM d) 4096-QAM

	64-QAM	S/I, N=3 Interferers										
		20	21	22	23	24	25	26	27	28	29	30
	25	1.80E-02	1.20E-02	9.10E-03	4.31E-03	3.45E-03	1.72E-03	9.35E-04	7.46E-04	5.01E-04	2.88E-04	2.09E-04
	26	1.90E-02	1.30E-02	7.08E-03	4.20E-03	1.63E-03	1.18E-03	5.63E-04	3.12E-04	2.07E-04	1.09E-04	7.90E-05
SNR	27	1.90E-02	1.00E-02	4.29E-03	2.19E-03	9.51E-04	3.26E-04	2.17E-04	1.22E-04	3.70E-05	2.00E-05	1.90E-05
	28	1.90E-02	9.70E-03	3.17E-03	1.03E-03	5.22E-04	1.36E-04	1.01E-04	2.20E-05	2.00E-05	1.10E-05	5.00E-06
	29	1.20E-02	4.47E-03	2.25E-03	1.31E-03	7.90E-05	1.30E-05	3.20E-05	1.00E-06	0	0	0
	30	6.32E-03	2.58E-03	1.17E-03	2.32E-04	1.78E-04	4.50E-05	7.00E-06	0	0	0	0
	31	1.40E-02	6.29E-03	2.10E-03	3.75E-04	7.00E-06	8.00E-06	0	0	0	0	0
	32	1.60E-02	1.13E-03	2.81E-04	3.90E-05	4.10E-05	2.00E-06	0	0	0	0	0
	33	2.69E-03	5.27E-04	1.05E-03	1.57E-04	2.00E-06	0	0	0	0	0	0
	34	2.36E-03	4.48E-03	3.10E-05	2.30E-05	1.00E-06	0	0	0	0	0	0
	35	9.62E-03	5.78E-04	6.60E-05	0	0	0	0	0	0	0	0

	256-QAM	S/I, N=3 Interferers										
		30	31	32	33	34	35	36	37	38	39	40
	30	4.02E-03	2.68E-03	1.90E-03	1.36E-03	8.94E-04	6.77E-04	5.22E-04	4.48E-04	3.53E-04	2.95E-04	2.86E-04
	31	2.86E-03	1.30E-03	9.74E-04	5.69E-04	4.25E-04	2.36E-04	2.20E-04	1.34E-04	1.11E-04	9.10E-05	8.60E-05
SNR	32	1.40E-03	9.41E-04	5.01E-04	2.03E-04	1.50E-04	7.90E-05	6.50E-05	3.50E-05	1.90E-05	1.50E-05	1.30E-05
	33	5.79E-04	3.93E-04	9.90E-05	1.18E-04	3.50E-05	2.00E-05	1.70E-05	1.20E-05	6.00E-06	4.00E-06	2.00E-06
	34	2.89E-04	2.76E-04	3.80E-05	2.00E-05	2.60E-05	2.00E-06	2.00E-06	0	0	0	0
	35	8.70E-05	1.09E-04	3.00E-05	1.10E-05	8.00E-06	0	1.00E-06	0	0	0	0
	36	3.80E-05	5.00E-06	7.00E-06	0	0	0	0	0	0	0	0
	37	3.70E-05	5.00E-06	1.00E-06	0	0	0	0	0	0	0	0
	38	1.50E-05	4.00E-06	0	0	0	0	0	0	0	0	0
	39	5.00E-06	0	0	0	0	0	0	0	0	0	0
	40	0	0	0	0	0	0	0	0	0	0	0

	1024-QAM	S/I, N=3 Interferers										
		40	41	42	43	44	45	46	47	48	49	50
	35	1.99E-03	1.60E-03	1.38E-03	1.10E-03	1.01E-03	8.36E-04	7.91E-04	8.35E-04	7.51E-04	6.92E-04	6.09E-04
	36	7.18E-04	6.75E-04	5.39E-04	4.05E-04	3.64E-04	3.26E-04	2.61E-04	2.13E-04	2.23E-04	2.04E-04	1.68E-04
SNR	37	4.34E-04	1.94E-04	2.18E-04	1.14E-04	8.40E-05	8.40E-05	4.60E-05	4.60E-05	6.20E-05	5.10E-05	5.40E-05
	38	1.33E-04	1.10E-04	5.00E-05	3.30E-05	2.20E-05	1.40E-05	1.40E-05	1.10E-05	7.00E-06	3.00E-06	7.00E-06
	39	2.60E-05	1.60E-05	1.50E-05	0	4.00E-06	0	0	4.00E-06	0	0	0
	40	5.00E-06	5.00E-06	2.00E-06	0	1.00E-06	0	0	0	0	0	0
	41	1.00E-06	0	0	0	0	0	0	0	0	0	0
	42	2.00E-06	0	0	0	0	0	0	0	0	0	0
	43	0	0	0	0	0	0	0	0	0	0	0
	44	0	0	0	0	0	0	0	0	0	0	0
	45	0	0	0	0	0	0	0	0	0	0	0

	4096-QAM	S/I, N=3 Interferers										
		40	41	42	43	44	45	46	47	48	49	50
	40	0.014	0.01	7.92E-03	6.40E-03	4.87E-03	4.24E-03	3.48E-03	2.83E-03	2.71E-03	2.28E-03	2.16E-03
	41	0.011	7.55E-03	5.34E-03	3.79E-03	2.63E-03	2.18E-03	1.56E-03	1.31E-03	1.11E-03	8.47E-04	7.26E-04
SNR	42	8.11E-03	4.64E-03	2.87E-03	1.98E-03	1.23E-03	9.98E-04	6.45E-04	5.35E-04	4.13E-04	3.17E-04	2.77E-04
	43	5.34E-03	2.49E-03	1.43E-03	1.26E-03	5.54E-04	4.43E-04	2.91E-04	1.32E-04	1.16E-04	6.60E-05	7.10E-05
	44	4.91E-03	2.18E-03	9.43E-04	3.59E-04	2.24E-04	1.02E-04	8.71E-05	3.60E-05	3.50E-05	2.50E-05	1.30E-05
	45	3.48E-03	1.53E-03	3.76E-04	2.03E-04	1.01E-04	2.50E-05	1.60E-05	8.01E-06	1.80E-05	5.00E-06	1.00E-06
	46	1.83E-03	1.04E-03	3.52E-04	1.55E-04	4.80E-05	1.20E-05	8.01E-06	1.20E-05	0	0	0
	47	2.04E-03	1.81E-04	1.68E-04	9.01E-06	1.10E-05	1.00E-06	0	1.00E-06	0	0	0
	48	9.81E-04	2.89E-04	2.70E-05	8.01E-06	1.00E-06	0	0	0	0	0	0
	49	7.19E-04	9.21E-05	1.00E-05	0	0	0	0	0	0	0	0
	50	4.51E-04	8.31E-05	6.00E-06	0	0	0	0	0	0	0	0

Based on the above, then, the IC is providing about the same amount of cancellation in these two cases tested, although there is much more SNR headroom in the 64-QAM case. However, the 64-QAM case also does not *need* additional suppression, so, while it may be available, no further IC adaption is required to deliver low error performance.

For the upstream, we are interested in extending what we have learned for 64/256-QAM to 1024-QAM as an advanced profile. Simulation results for the "DOCSIS" SNR threshold condition for 1024-QAM of Table 3 and a 27 dB S/I easily shows that uncoded error rates are horrendously high (0.1 to 0.01 range). They are so high as to possibly be unable to be corrected adequately if at all by FEC, or it would be not desired to rely so heavily on it. Better than 26 dB of IC would be required, and based on the above mentioned relationship, probably 6 dB better. It is unclear if today's IC can accomplish this for multiple interferers with bandwidth, but it likely will be necessary.

It is worthwhile to point out that for a single CW interferer, effective cancellation of about 35 dB was obtained based on Table 7. So, at least for this friendlier case, the IC function can be stronger.

Threshold Values with New FEC

We have already concluded in prior analysis that a 1024-QAM downstream will require LDPC, so let's consider how this plays out relative to S/I. Our revised threshold of 38 dB is for AWGN. For the upstream, where we allocate more margin, this is still a low error rate condition, as shown in Figure A-5 (appendix). A 40 dB SNR is the 1e-8 BER case for 1024-QAM.

The same modeling table above used for 256-QAM estimation indicates that for the 38 dB SNR case, an S/I of approximately 37 dB will leave a very correctable error rate in the 1e-3 to 1e-4 range. A 1e-5 threshold, based on Table 9c, suggests instead a 42 dB S/I. For the 37 dB objective, -2 dBc of interference successfully suppressed for 256-QAM in Table 9 then requires 35 dB of IC applied. This is *more* than 6 dB (by 3 dB) of additional IC given the 26 dB 256-QAM example. The same exercise for 4096-QAM, were it an upstream mode (using 44 dB SNR), would suggest a 42-45 (Table 9d) S/I, or 40-43 dBc of IC. This is 5-8 dB different than 35 dBc.

So, there is at this point only a range of additional IC expectations that seems to follow from these results. Intuitively, not being an AWGN impairment around which detection (and FEC) is optimized, we should expect that once the impairment is large enough to contribute to errors, the relationship would exceed 6 dB per modulation profile. Because of that and the sensitivity of FEC error rate curves to fractions of dB of SNR, it is probably a good starting point to consider a relationship such as 8 dB more IC capability until more granular modeling over a range of interference patterns and a larger sample size can be established.

Downstream Interference

In the downstream, a significant amount of study has already been performed to quantify the impact of CSO and CTB analog beat distortions on QAM. These distortions look like narrowband interferers, but the nature of their make-up (many independent distortion contributors falling close to one another) is that they have noise-like qualities, including an amplitude modulation component. This is meaningful for BER degradation. An example of a 256-QAM signal with analog distortion components from a 79-channel load is shown in Figure 20.

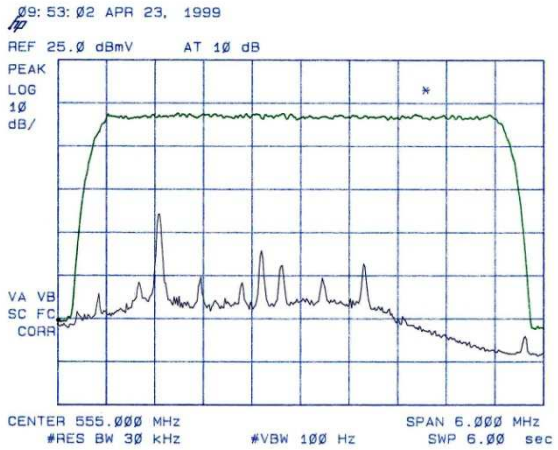

Figure 20 –256-QAM with CSO and CTB Distortions

The "donut" constellation we saw previously for 64-QAM in Figure 18 gives insight into why the beat distortion phenomenon when analog loads are prevalent matters from a QAM perspective. A CW interference example for 256-QAM at very high SNR is shown in Figure 21.

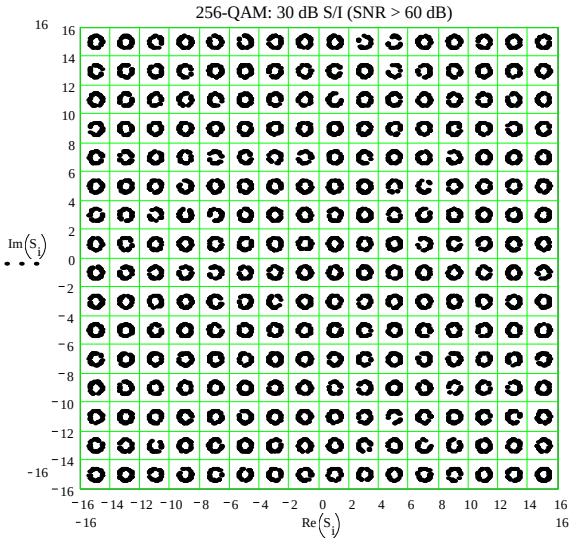

Figure 21 –256-QAM @ 30 dB S/I

In Figure 21, the SNR is 60 dB – an error free region in an AWGN-only environment, where 1e-8 occurs at SNR = 34 dB. Clearly, with only this CW carrier imposed on Figure 21, we still have an apparent error free environment. The peak-to-average ratio of a sinusoid is, of course, 3 dB – it has a constant envelope. If that envelope is not large enough to cause a decision error, then without noise there will be no decision errors even with the interferer.

The concerns with respect to CSO and CTB distortion beats are that, unlike CW or FM interference, they have a noise-like peak-to-average quality to them [20]. In fact, the envelope has a Rayleigh-like fit, which is representative of the detected envelope of a Gaussian process. This is shown in Figure 22 [20]. This amplitude modulating effect can be applied to the "donut" of Figure 21 – envision the circular symbol point breathing in and out. The nature of the distortion beat degradation is also that it is a narrowband process (10's of kHz) relative to the QAM bandwidth. Thus, a distortion beat sample will extend over many symbols in a row, and if it is high enough to induce decision errors, there is likely to be a burst of them.

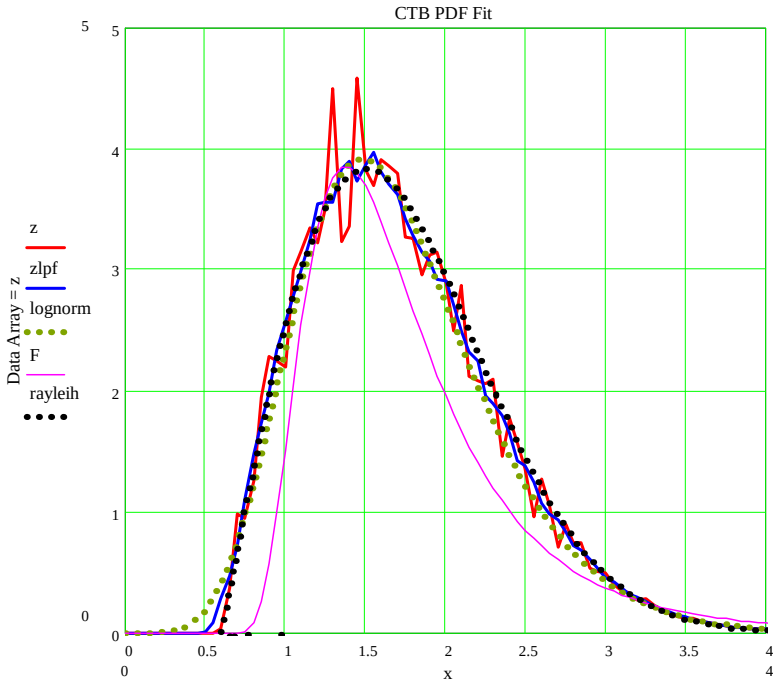

Figure 22 –Amplitude PDF of CTB

Fortunately, as we shall see and describe in more detail in the case of phase noise in the next section, in the downstream a powerful interleaver is available. It is capable of randomizing the errors, allowing the FEC to do its job better. This is quantified for phase noise, but the same dynamics apply in this case. A 20 kHz process has a 50 usec "time constant" of error generation, and a common (I=128, J=4) interleaver setting exceeds this by a factor of more than 5, easily distributing the errors into correctable codewords.

Analysis and test of the CSO/CTB impact for 1024-QAM has been performed [9]. A summary table of the results relative to these analog distortions is shown in Table 10. One of the key take-aways from that analysis – RF cascade depth as a function of analog carriers and amplifier performance – is shown in Figure 23.

For Figure 23, perhaps the single key result for purposes of this paper is that if analog video is reduced to 30 carriers, then the cascade depth that can be tolerated, under assumptions of 20 Log(N) degradation (which is overly pessimistic) is, for all practical purposes, unlimited for 1024-QAM using amplifier performance commensurate with today's plant equipment.

Table 10 – 1024-QAM Downstream Interference and Thresholds

	CW Interference		CTB Interference			
SNR	Pre-FEC Error Threshold	Post-FEC Error Threshold	Pre-FEC Error Threshold	Post-FEC Error Threshold	Post FEC > 1E-6	Post FEC Broken
50 dB	34 dB	33 dB	55 dB	55 dB	55 dB	45 dB
45 dB	35 dB	33 dB	55 dB	60 dB	55 dB	46 dB
40 dB	36 dB	34 dB	60 dB	60 dB	55 dB	49 dB
37 dB	38 dB	37 dB	60 dB	60 dB	55 dB	50 dB

Based on the CSO and CTB values we observed in Table 5, this is not surprising. The historical minimum acceptable value of 53 dBc works out to 47 dBc for 256-QAM, and this proved to be manageable with 256-QAM. With the 30-analog values of CSO and CTB ranging from 61-67 dB, 8-14 dB of better performance is occurring while a modulation profile increase of 6 dB is taken on for 1024-QAM. Also, Table 5 represents the performance at the worst case frequency. For CSO, these components tend to pile up at the low end of the band (analog). CTB tends to pile up in the middle, where QAM spectrum will be allocated – thus the focus on CTB below.

If we consider that 1024-QAM under an LDPC FEC has a threshold SNR of 31 dB based on our Table 2 downstream assumptions, then 60-something dB of CTB distortion, even with a noise like amplitude variation, would have little impact on an architecture delivering this SNR, or an SNR higher but not high enough to advance the modulation profile to 2k/4k-QAM. Basically, if 1024-QAM is workable with today's PHY, then better FEC can only make it better – it just may not achieve all of the new FEC gain of an AWGN-only noise. However, the interleaver, the relationship of the distortion levels to the CCN values in Table 5, the low pre-FEC thresholds observed in [9], suggest that it will achieve full benefits of the FEC.

We had already concluded in [9] that 1024-QAM was possible, and much more so at 30 analog carriers. What can we say about 4096-QAM?

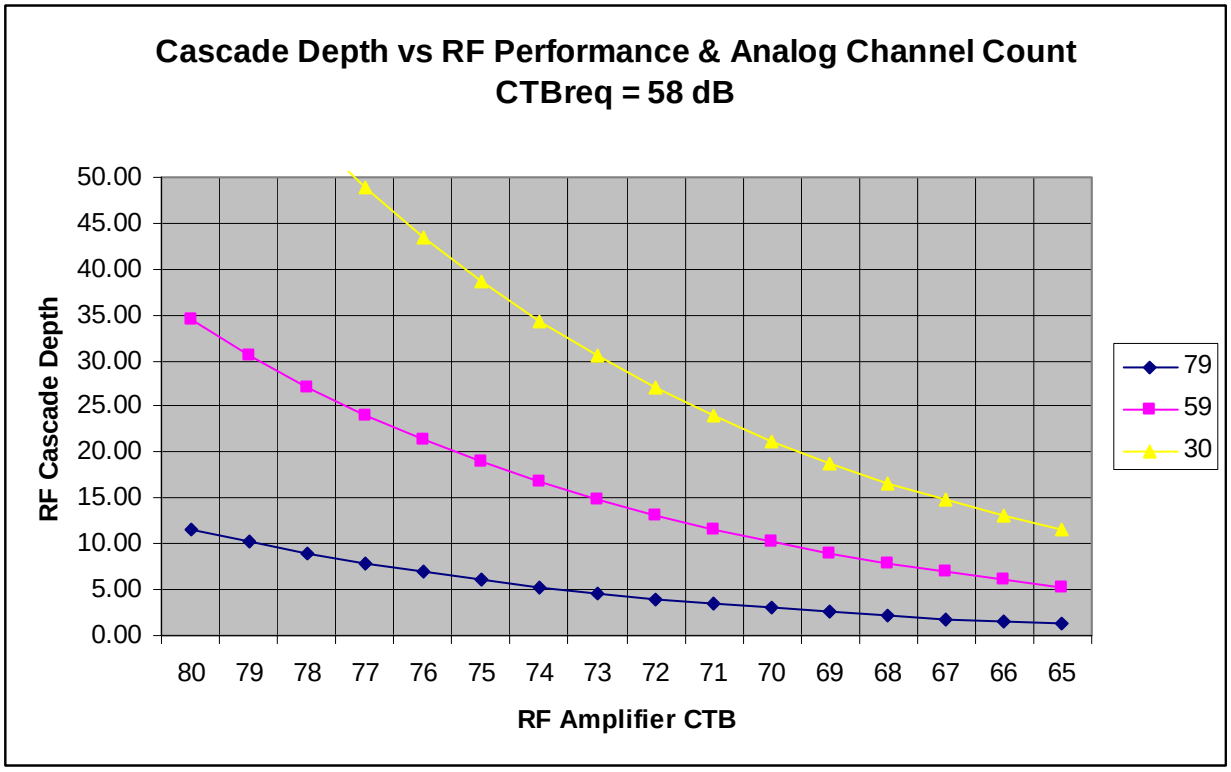

Figure 23 –RF Cascade Depth Limitations vs. Amplifier CTB for Analog Reclamations

First, a few items that may make us not care very much about that situation:

- When analog is fully removed, there are no longer any beat distortion components to worry about as interferers. We can then simply follow the CCN and the SNR analysis previously discussed.
- In the case of narrowband interference in the downstream, we could call on the upstream interference analysis above. It can only be conservative since the upstream must handle burst reception and adapt accordingly on a burst by burst basis. In the downstream, the receiver has the luxury of a constant input signal, a much simpler problem.
- Lastly, by the time we are deploying 4096-QAM, it is very likely that it is part of a multi-carrier downstream, and so narrowband interference analysis applies completely differently. We will discuss that later in the paper.

Figures 3-8 pointed out why the 4096-QAM case will require LDPC to be robust. It will provide the margin necessary to maintain performance against the 37 dB threshold established in Table 2. A high pre-FEC error count would ensue for 79-channel analog system at the 53 dBc minimum: The resulting narrowband interference would be 47 dBc on average, and peak to 33-35 dB, causing pre-FEC error rates worse than .01 (Figure A-5). They might be fixable, but this would not the ideal way to consume the FEC budget.

Therefore, 4096-QAM and full analog loading, we would suggest, is not a good combinations. 4096-QAM ought to be reserved for reduced analog loading or no analog loading. A caveat is that, as part of a transition plan of spectrum, such as Figure 1 indicates, may include new PHY above today's forward band, and beat mapping analysis would likely treat this region favorably in dBc of distortion reduction.

For reduced analog loading, CTB values of 66 dB (60 dB to QAM), peaking amplitudes would instead be 46-48 dB. The 1e-8 value for 4096-QAM without coding is 46 dB. So, it is likely errors will be counted, and corrected. They will be bursty, but the interleaving will arrange them nicely for the FEC. The situation is analogous to 256-QAM measurements with distortion in 2002 [20]. We now have 12 dB more of modulation profile sensitivity, and 13-14 dB better distortion values due to analog reclamation assumptions. Because they are noise like in effect on the constellation (MER is clouded), the similar relative relationship should yield similar results. This is also consistent with the expectation that 1024-QAM will perform very well in a reduced analog system from a distortion perspective.

PHASE NOISE

When QAM signals are put into the RF domain, they inherently have a phase noise mask applied to them. Phase noise is a measure of the spectral purity of the carrier signal itself. It is commonly measured by turning off the modulation on a waveform, and observing the level of noise surrounding the carrier at very close offset frequencies. At its most simplest, a perfect CW tone would be a single line in the frequency domain. In practice it cannot be perfect, and the amount that it is not perfect is quantified by this random phase modulation imposed at frequency offsets from the carrier frequency itself. The shape of the noise around a carrier is well understood, following many classic behaviors of semiconductor-based oscillators and the circuits that perform frequency synthesis to put modulated signals somewhere in the RF band. While there are many variables, in general, the higher the frequency, the worse the phase noise will be,

the broader the tuning range, the higher the phase noise will be, and the finer the increment of tuning, the higher the phase noise will be. The shape is also known as a phase noise "mask."

An example phase noise mask is shown in Figure 24 [10]. It illustrates some common characteristics – a close to carrier flat region and small peaking, and a region of 20-30 dB/decade roll-off. The flat region is often much wider than the example shown here. The "0" of the x-axis represents the carrier itself, and the data points plotted indicate offsets from the carrier where noise density is measured. The values at offset frequencies are given in dBc/Hz, and the total noise in a bandwidth is then the area under the curve of the mask, usually recorded as a dBc value or degrees rms.

Signal-to-Phase Noise Relationships

In this section we introduce some simple-to-understand M-QAM-phase noise relationships based on the qualitative descriptions above, some nomenclature, and a deeper understanding of the processes involved in determining its effects.

For converting dBc values of phase noise, or signal-to-phase noise ratios, to degrees rms, we have:

deg rms of phase noise = (180/π) sqrt[10^(-dBc of phase noise)]

The use of degrees rms is often very illustrative when we think about QAM constellations, as we shall see. There is a simple rule of thumb that keeps us from having to rely on the above equation. It is based on the recognition that a 35 dB signal-to-phase noise ratio is the same as 1° rms. Also, rms is a linear quantity, so doubling it is 6 dB.

Figure 24 – Example Phase Noise Mask – RF Upconverter @ 601.25 MHz

For example, if -35 dBc = 1° rms, then

-23 dBc = 4° rms
-29 dBc = 2° rms
-35 dBc = 1° rms
-41 dBc = 0.5° rms
-47 dBc = 0.25° rms, etc.

Because of its rotational effect, phase noise affects QAM constellation points non-uniformly. Figure 25 [6] shows an example of essentially noise-free 64-QAM with 1° rms phase noise, using a mathematical tool on the left and a simulation environment (for error rate analysis) on the right. While the angle of rotation is the same for every symbol point, it is apparent and geometrically expected from polar coordinate mathematics how this impacts the outermost symbol points the most relative to breaching decision boundaries.

Compare Figure 25 with Figure 26, which shows the same 64-QAM symbol with only AWGN impairment, set at a 1e-8 BER. Note how the degradation due to additive noise is randomly distributed in I and Q dimensions, whereas the phase noise impact is exclusively angular.

Furthermore, the sensitivity of M-QAM gets worse with increasing M because of this non-uniform rotational effect. The shrinking of the distance to the decision boundaries for increasing M for a fixed average power puts makes the same amount of rotation more deleterious for higher M-QAM profiles. Figures 27 and 28 show a noise-free 256-QAM constellation with just 0.5° rms phase noise imposed, and a 1024-QAM constellation with a .25° rms phase noise imposed, respectively. The similarity in Figures 25, 27 and 28 of the relative rotation to decision boundaries, for the outer symbols in particular, is clear.

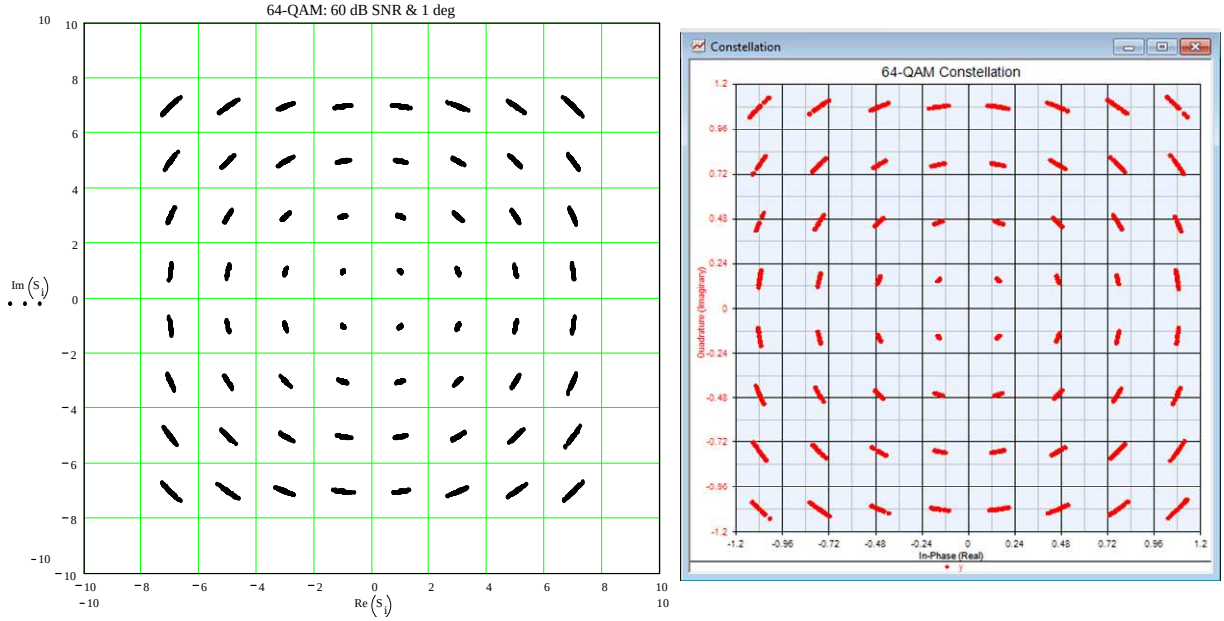

Figure 25 – 64-QAM, 1° rms Phase Noise (Analysis Tool, Simulation Tool)

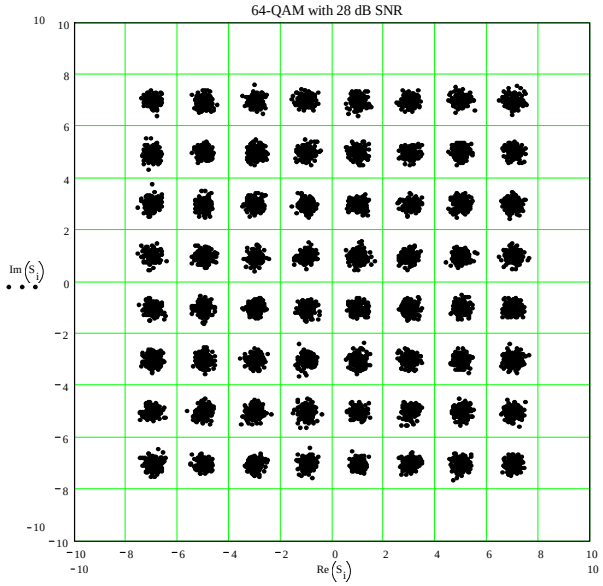

Figure 26 – 64-QAM @ 1e-8 Noise (AWGN) Level

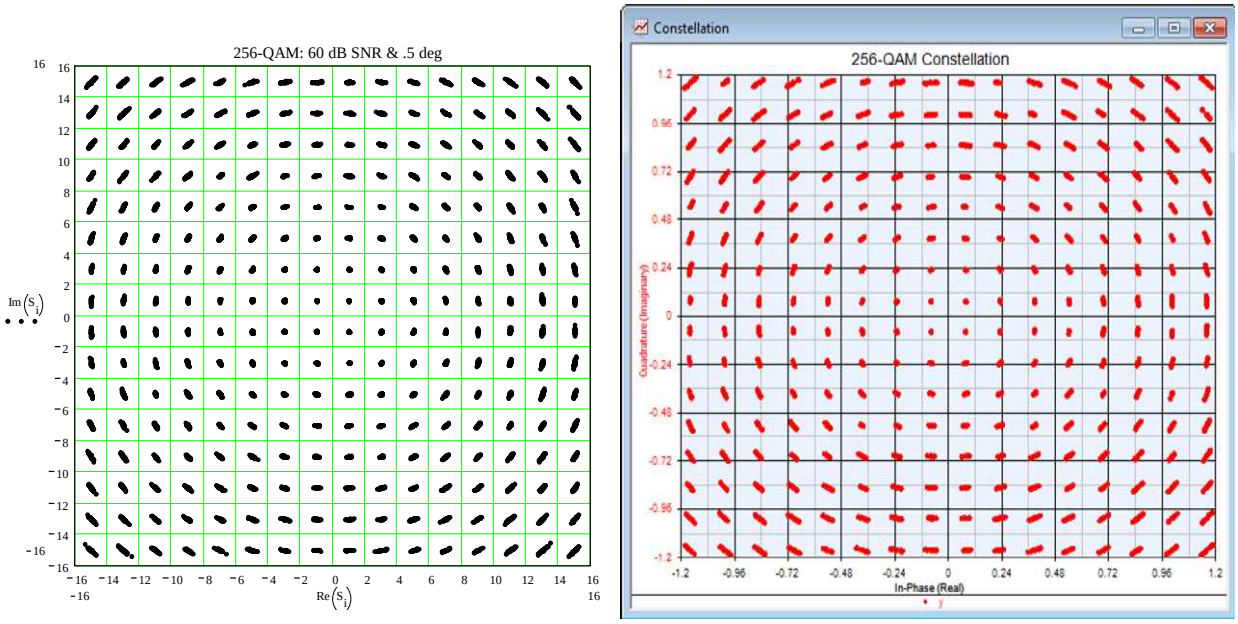

Figure 27 – 256-QAM, 0.5° rms Phase Noise (SNRφ = 41 dB), (Analysis, Simulation Tool)

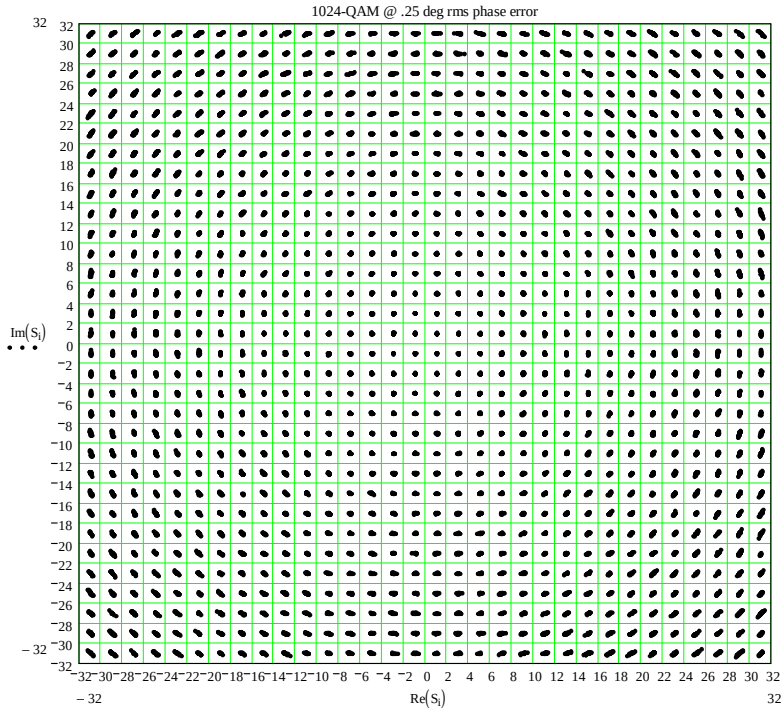

Figure 28 – 1024-QAM, 0.25° rms Phase Noise (SNRφ = 47 dB)

The nature of untracked phase noise is that it can lead to error rate floors at detection, because even without AWGN, if there is enough untracked phase noise after carrier recovery, it alone can cause symbols to cross boundaries. This is most likely for the outermost symbols, and these points can thus be used to determine limitations of phase noise necessary to eliminate flooring. More complex expressions are required to set thresholds associated with minimizing BER degradation [5, 6, 10].

It is a simple trigonometric matter to determine the rotational distance to a decision boundary as a function of M for M-QAM:

$$\varphi \text{ (decision boundary)} = \arcsin[(\sqrt{M} - 1) / M\sqrt{2}\,]$$

SNRφ Thresholds, M-QAM, and BER

Table 11 summarizes the phase error analysis across the modulation profiles of interest. It also identifies recommended levels of phase noise for minimal BER degradation, and levels beyond which the degradation curve shifts from a simple offset from theory to a more severe break from the normal steepness of descent of the BER waterfall curve.

Table 11 – Untracked Phase Noise Limits vs. M in M-QAM

	φ	dBc φ thresh	BER < 0.5 dB	SNR φ	BER on the Brink	SNR φ
16-QAM	16.8°	-10.7	1°	35.0	2°	29.0
64-QAM	7.7°	-17.4	.5°	41.0	1°	35.0
256-QAM	3.7°	-23.8	.25°	47.0	.5°	41.0
1024-QAM	1.8°	-30.1	.125°	53.0	.25°	47.0
4096-QAM	0.9°	-36.1	.0625°	59.0	.125°	53.0

The relationships shown can be deduced in part by recognizing that, since we are using a Gaussian statistical model for the jitter, the boundary merely represents a threshold on a normal curve that we can scale the rms (σ) to calculate its probability of threshold crossing. For example, a 16-QAM floor of about 1e-6 occurs for 4° rms, while for 64-QAM, a similar floor exists for 2° rms.

To exactly quantify allowable degradation with phase noise, AWGN and now phase noise can be combined together to create a composite "σ." However, they do not impact BER in a uniform fashion, as Figure 27 and 28 make apparent. Nonetheless, it is common approximation for lower order modulation formats to sum the AWGN noise and phase noise come up with a composite SNR. This simplification tends to understate the impact for high M, however.

Figure 29 shows a BER analysis that includes both phase noise (.25°, .35°, .5° rms) and AWGN contributions for 256-QAM [6]. The thresholds identified for 256-QAM in Table 11 are shown clearly on this chart by referencing the 1e-8 error rate threshold. The 0.25° rms value represents < 0.5 dB of degradation, while the 0.5° rms value has clearly is losing the characteristic waterfall shape, and on the verge of an error rate disaster.

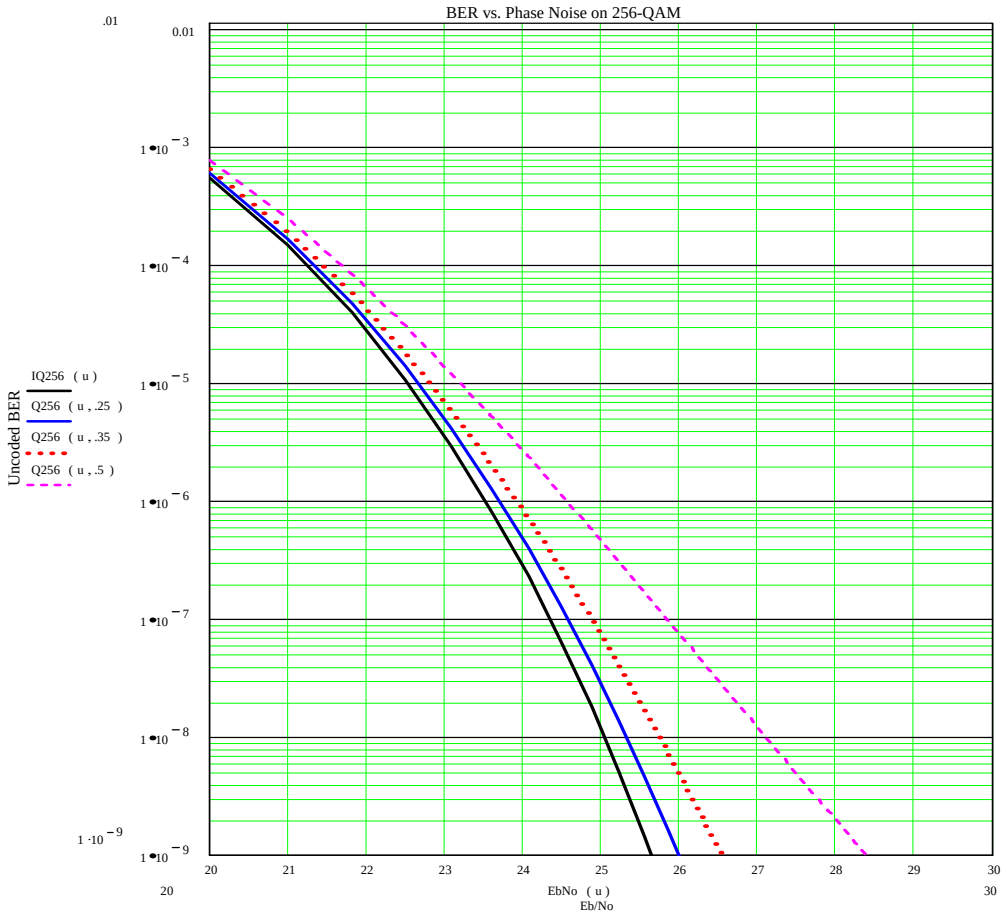

Figure 29 - 256-QAM BER with Phase Noise: .25°, .35°, and .5° rms

Note in Table 11 how the SNRφ required increases with increasing M. This relationship is similar to the relative relationship to AWGN. However, while a network architecture and new FEC may enable an AWGN performance improvement, the RF portion of the architecture that contributes to phase noise tends to remain in place and can be affected mostly in smaller ways by the tracking process. Redesign of RF equipment to achieve the same frequency agility objectives with improved phase noise is no minor proposition.

Offset M-QAM modulations and adjustments to decision boundaries in the face of a dominant phase noise impairment have been explored and this is addressed in [4] for the interested reader.

HFC Equipment Calculations

Phase noise is important because QAM, of course, encodes information in the phase of the symbol. A QAM signal contains "I" and "Q" orthogonal components, and the amplitude and phase applied to these identifies a point on a QAM constellation. This is why phase *noise* matters to M-QAM transport – noise in the phase domain translates to constellation position error or MER degradation in the signal space as we have seen in Figures 27 and 28.

The Downstream RF Interface Specification (DRFI), part of the DOCSIS portfolio of requirements, recognizes this and has a phase noise requirement, shown in Table 12.

All RF frequency synthesis or frequency conversion functions along the way contribute to the phase noise mask. The other typical major contributor in cable is the tuning function in the CPE. Though this function has been replaced in the RF circuitry sense by FBC technology discussed previously (wideband A/D conversion front ends), the clocking function of the A/D instead imparts the phase noise.

A modern wideband tuner built for digital cable and designed for compliance with ITU J.83A-C, is the Microtune MT2084. It has a specified phase noise requirement that serves as an excellent reference. These are shown in Figure 30.

Table 12 – Example Phase Noise Mask – RF Upconverter @ 601.25 MHz

Phase Noise Single Channel Active, $N-1$ Channels Suppressed (see Section 6.3.5.1.2, item 6) 64-QAM and 256-QAM	1 kHz - 10 kHz: -33dBc double sided noise power 10 kHz - 50 kHz: -51dBc double sided noise power 50 kHz - 3 MHz: -51dBc double sided noise power
All N Channels Active, (see Section 6.3.5.1.2, item 7) 64-QAM and 256-QAM	1 kHz - 10 kHz: -33dBc double sided noise power 10 kHz - 50 kHz: -51dBc double sided noise power

Phase noise (SSB)
1 kHz offset -91 dB/Hz
10 kHz offset -92 dB/Hz
20 kHz offset -93 dB/Hz
100 kHz offset -105 dB/Hz
1 MHz offset -125 dB/Hz

Figure 30 – Sample Phase Noise Mask for RF Tuner in CPE

Since coherent QAM is used – meaning the carrier frequency and phase are recovered at the receiver in order to demodulate the signal and select which of the constellation points was transmitted, the final stage of "processing" of the phase noise mask occurs in the receiver. Carrier synchronization is performed by the carrier tracking subsystem. Modern designs use a decision-directed approach, which has been shown of the alternatives to have better noise performance, at least under low error rate conditions [18].

By tracking the carrier, a carrier recovery function is inherently also tracking the phase noise imposed on the carrier up to that point. However, it is a closed loop feedback system, and cannot track all of it without risk of other noise contributors, thermal and self-generated, from disturbing the stability of the recovery process. A feedback loop is in place which creates an error signal that is constantly adjusting the tuning oscillator to keep it aligned to the incoming signal. The feedback loop has a response time set by its loop bandwidth.

Without going into great detail about specific receiver architectures, the tracking occurs roughly up to the point of the loop bandwidth, and any RF-imposed phase noise beyond that is not tracked. There is an optimum bandwidth selection that considers these factors, input noise, and self-noise, among others. It is the total of untracked phase noise that contributes to MER degradation and possible symbol error. Note that DOCSIS specifications, in order to encourage innovation and competitive advantage among suppliers, allow flexibility on the receiver functions, where most of the complexity and sophisticated processing lie. As such, things such as loop tracking architectures or requirements are not defined, only end performance objectives under assumptions on channel conditions and other system assumption.

For the receiver designers, it is important to understand the associated transmit phase noise as defined in DRFI (CMTS or EQAM transmitter) and, for the upstream, in the DOCSIS PHY specification for cable modems:

-46 dBc, summed over the spectral regions spanning 200 Hz to 400 kHz
-44 dBc, summed over the spectral regions spanning 8 kHz to 3.2 MHz

Recall, the upstream has a range of symbol rates, beginning with 160 ksps (200 kHz) and increasing to 320 ksps (400 kHz) and so on in octaves up to 5.12 Msps (6.4 MHz). This range of symbol rates is reflected in the two requirements. We will assume wider symbol rates and thus the value of -44 dBc applies. Because of receiver design variations and the phase noise contributions from the receivers themselves, it is difficult to further quantify contributors to the phase noise process. We can say more will be added, some will likely be tracked out, and that the untracked jitter will have a lowpass structure to it.

We can at least, however, estimate what the specified requirements would mean to

our M-QAM constellations, and estimate implications to receiver architectures using the requirements that are in place. Let's examine the effect of the combined DRFI specification and above tuner mask on the QAM profiles of interest. We have used the DRFI requirement in Table 12, and the tuner mask shown in Figure 30 to create a composite mask. This is shown in Table 13.

Table 13 – Composite Mask: DRFI + Tuner

Composite Mask (dBc/Hz)	
50 kHz	-90
100 kHz	-110
1 MHz	-130
3 MHz	-139
Total, dBc	-47

Assume that all of the mask beyond 50 kHz is untracked, and assume it is the dominant contributor to untracked phase noise after carrier recovery. It extends out to the symbol rate edge of 3 MHz (6 MHz double sided). This suggests a tracking bandwidth in the 50 kHz range, tied to other parameters [14, 15], and high SNR conditions in the carrier recovery architecture from self noise and input SNR.

As shown in Table 12 and 13, the composite mask is a 47 dB SNRφ, or .25° rms. The mask in Table 13 is shown on 64-QAM, 256-QAM, 1024-QAM, and 4096-QAM in Figure 31a-d.

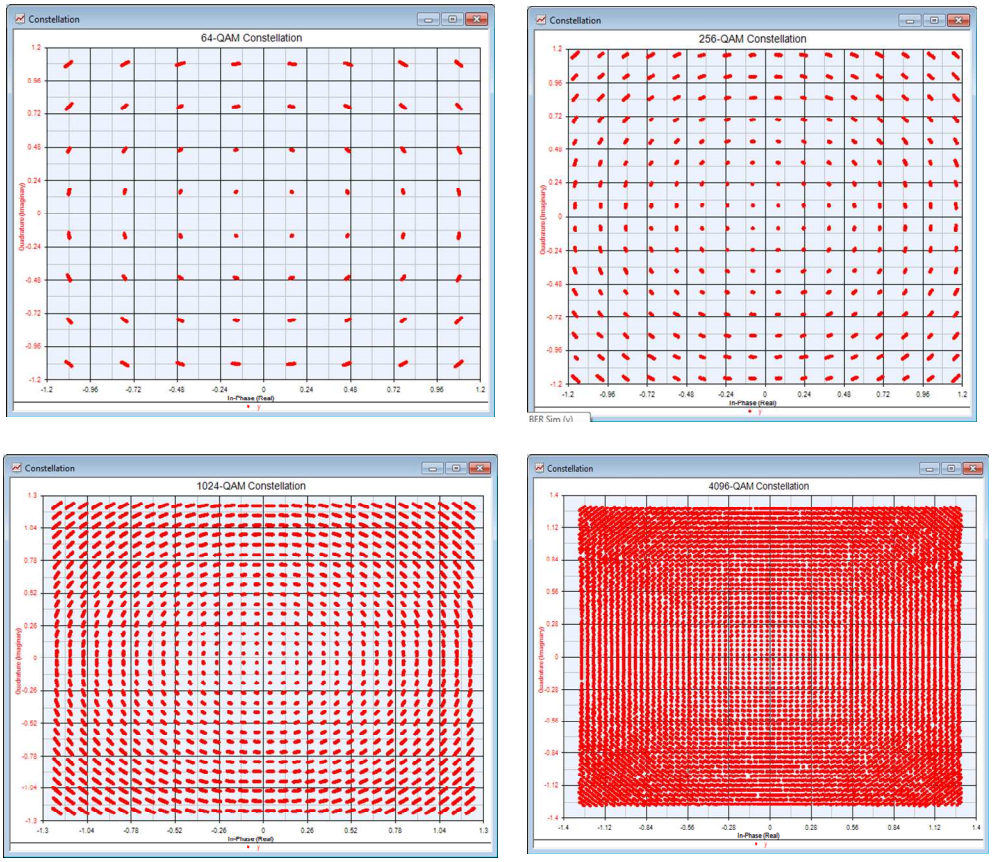

Figure 31(a-d) – Table 13 Mask Applied, Clockwise from Upper Left: a) 64-QAM b) 256-QAM c) 1024-QAM and d) 4096-QAM

In Figure 32, we have evaluated the uncoded BER for M-QAM profiles for M=64, 256, 1024 and 4096-QAM under the 47 dBc SNRφ conditions of Figure 31.

As Table 11 indicates, SNRφ = 47 dBc is the breakpoint between small degradation for 256-QAM, and the BER being on the brink of large degradation for 1024-QAM. The 4096-QAM case is untenable with this amount of untracked phase noise. This is all verified in Figure 32.

For 4096-QAM, which is clearly suffering and in practice would not be able to effectively hold the receiver locked for demodulation, it is difficult to see much in Figure 31 other than clouds of impossible-to-discriminate symbols. This case is shown again by itself in Figure 33, where you can begin to see some daylight between symbol points, mostly inner points. But, the symbol clouds at the edges are still massively intruding on each other's space to the point that they are becoming indistinguishable. This yields the very high symbol error rates, likely to overwhelm an error correction mechanism or carrier recovery subsystem.

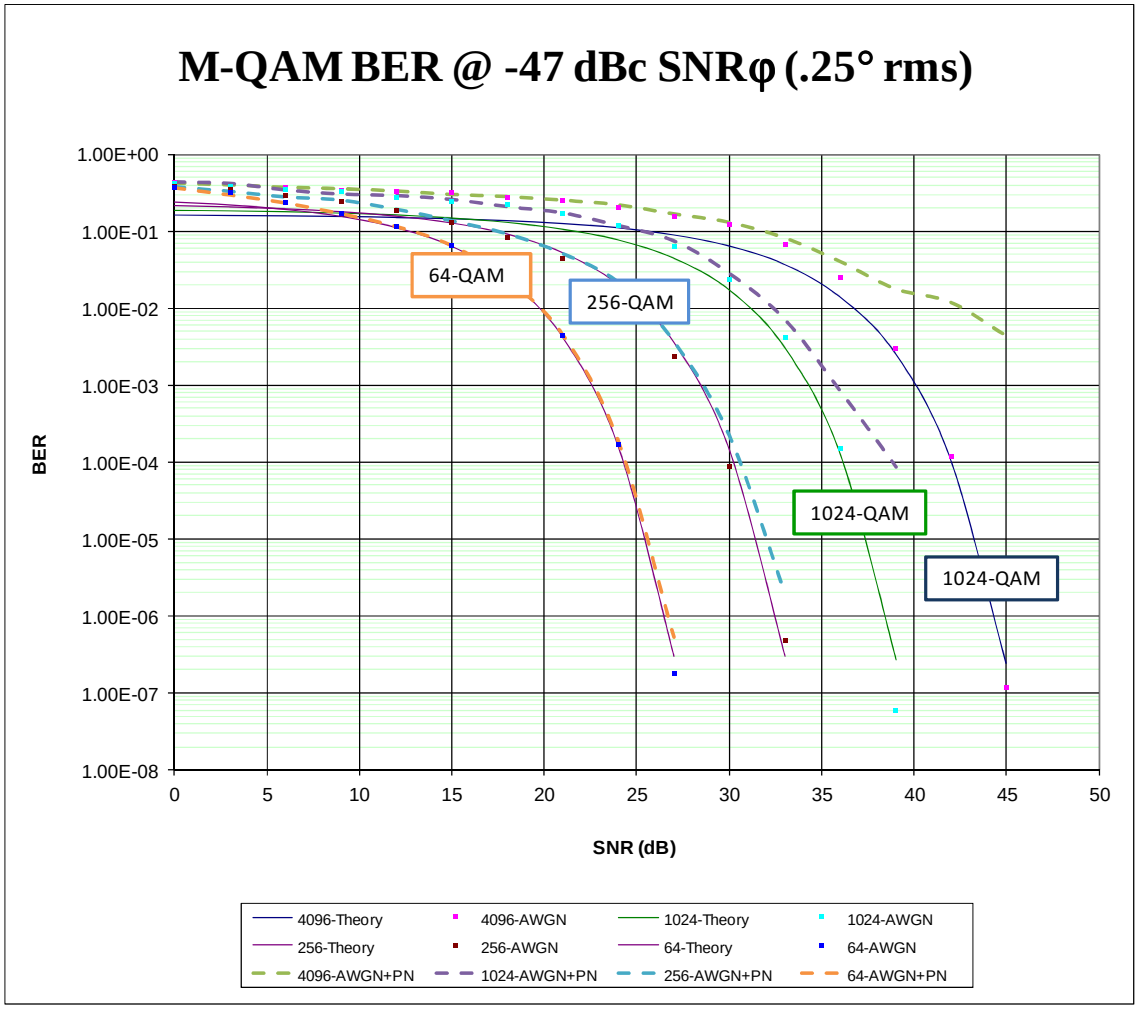

Figure 32 – M-QAM BER for SNRφ = 47 dBc (DRFI + Tuner)

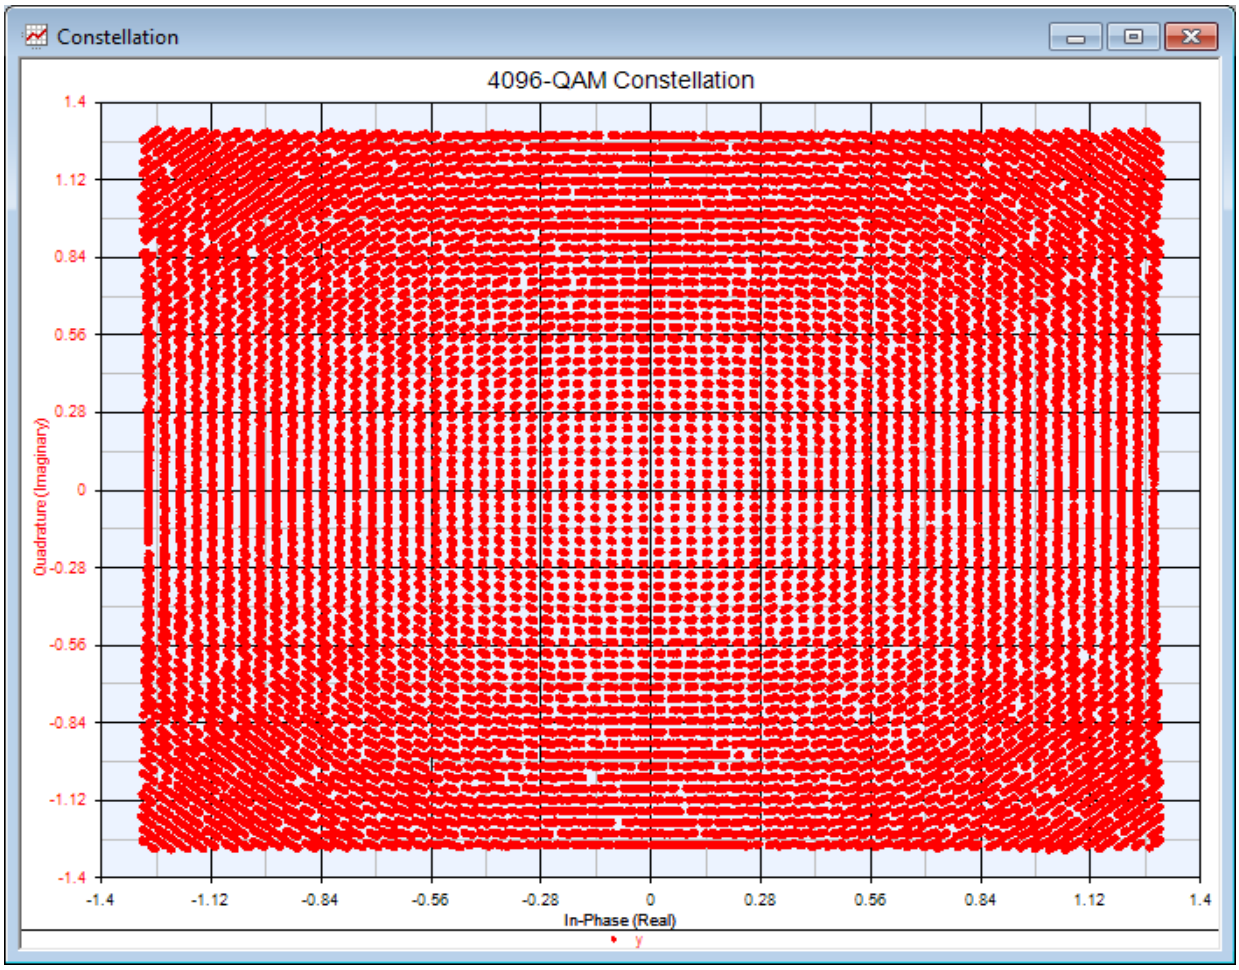

Figure 33 – 4096-QAM@ .25° rms (SNRφ = 47 dB)

According to the guidelines of Table 11, SNRφ = 53 dB is the brink of trouble for 4096-QAM, and a reasonable guideline for 1024-QAM that limits the degradation.

This 1024-QAM case, with SNRφ = 53 dB, is shown in Figure 34. The similar, relative MER characteristic compared to Figure 31b (256-QAM @ 47 dBc) is apparent. The 4096-QAM case for SNRφ = 53 dB is shown in Figure 35.

The BER evaluation for SNRφ = 53 dB is shown in Figure 36.

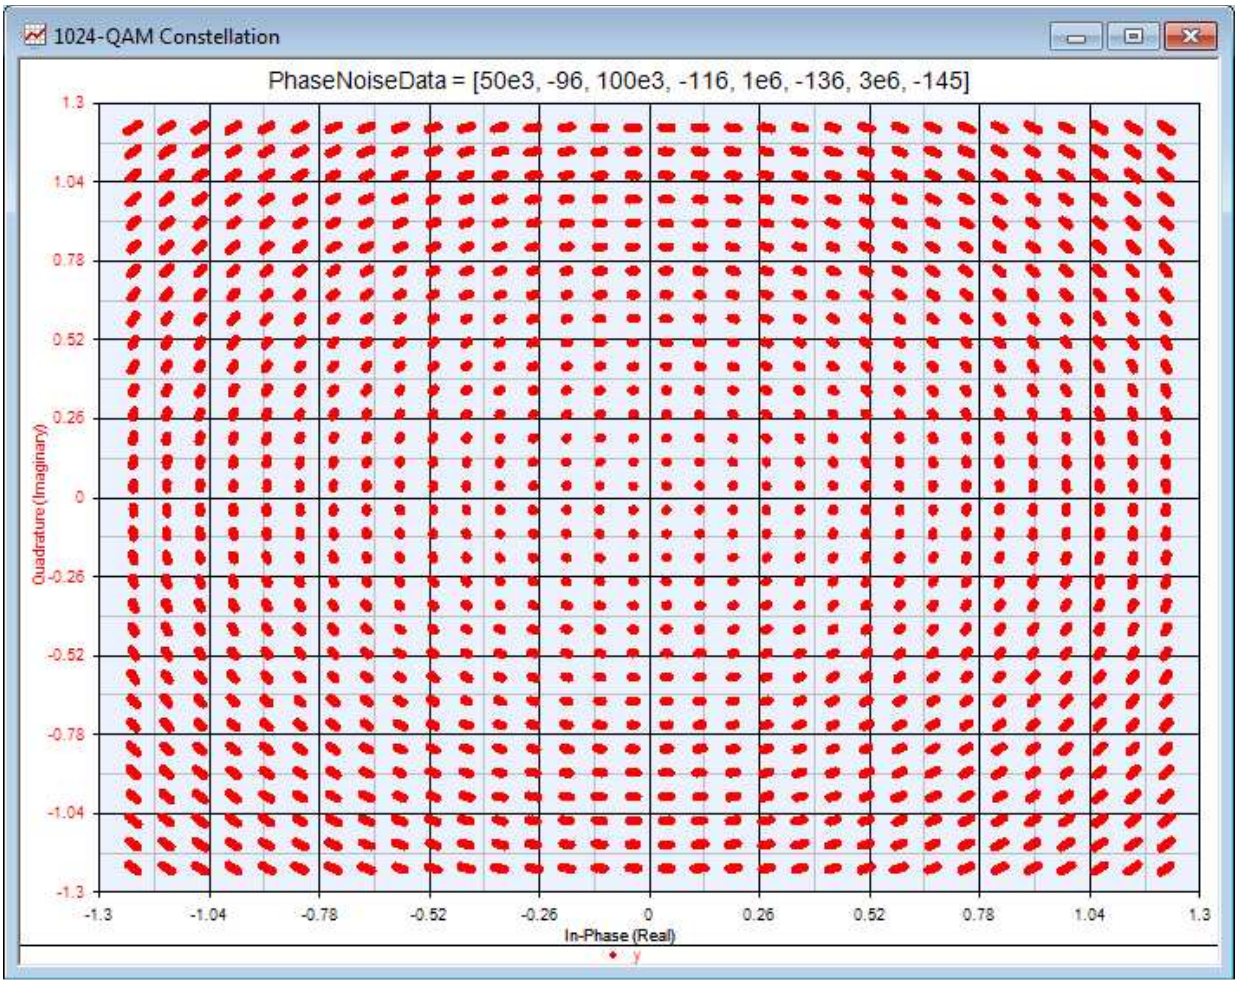

Figure 34 – 1024-QAM@ .125° rms (SNRφ = 53 dB)
(Recommended)

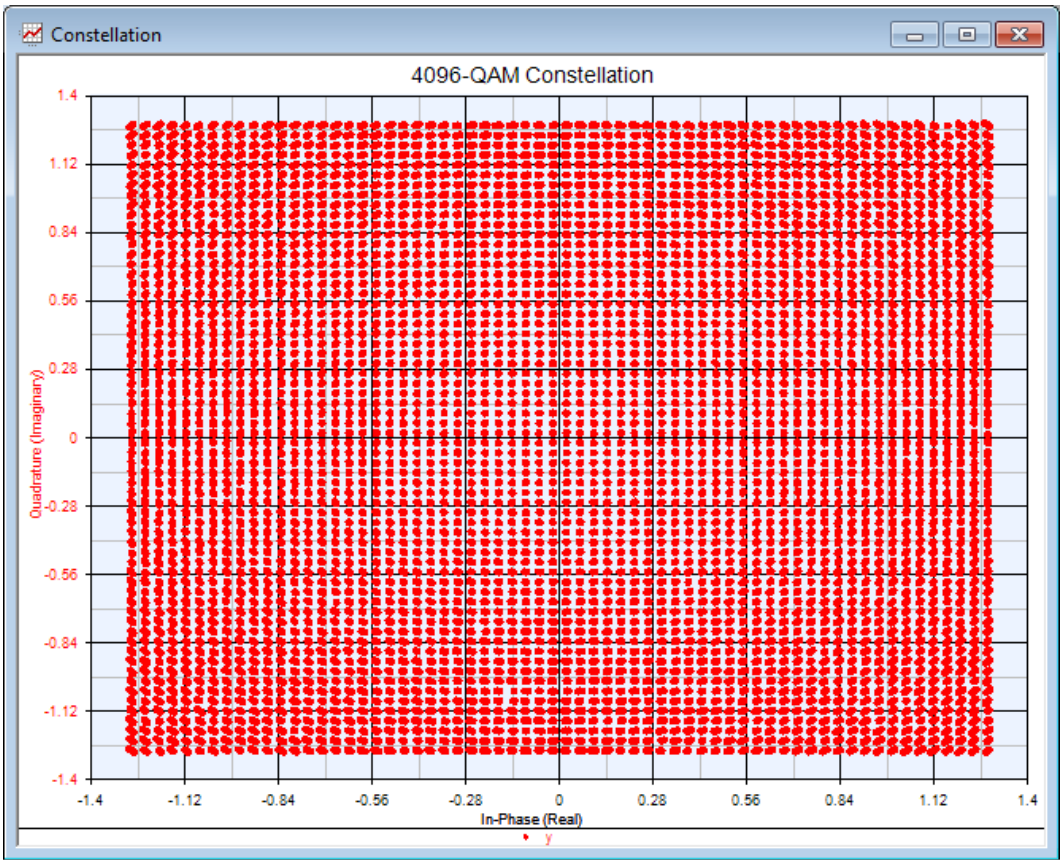

Figure 35 – 4096-QAM@ .125° rms (SNRφ = 53 dB)

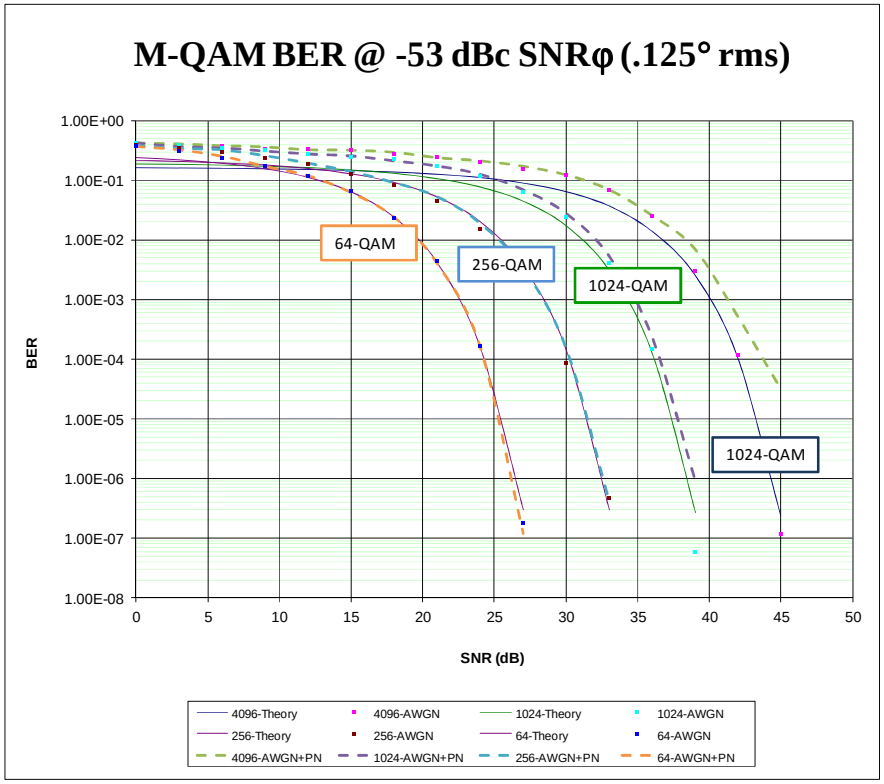

Figure 36 – M-QAM BER for SNRφ = 53 dBc

In Figure 36, we can see that 1024-QAM is now under control with modest degradation, and that 4096-QAM is on the edge of major BER performance degradation. This again is consistent with the recommendations in Table 11.

Finally, Figure 37 shows the constellation impact to 4096-QAM with the recommended maximum phase noise of SNRφ = 59 dBc, or .0625° rms. Of course, link phase noise is not going to adjust for the modulation profile, so the RF and tracking subsystem must be architected for the most sensitive modulation anticipated. The improved fidelity in Figure 37, in particular of the outer symbol points, illustrates why SNRφ = 59 dB is recommended for minimizing degradation against the theoretical performance curve.

Figure 38 shows the BER evaluation for the SNRφ = 59 dBc case, where it becomes clear that the 4096-QAM untracked rms phase noise recommendation of Table 11 is sufficient.

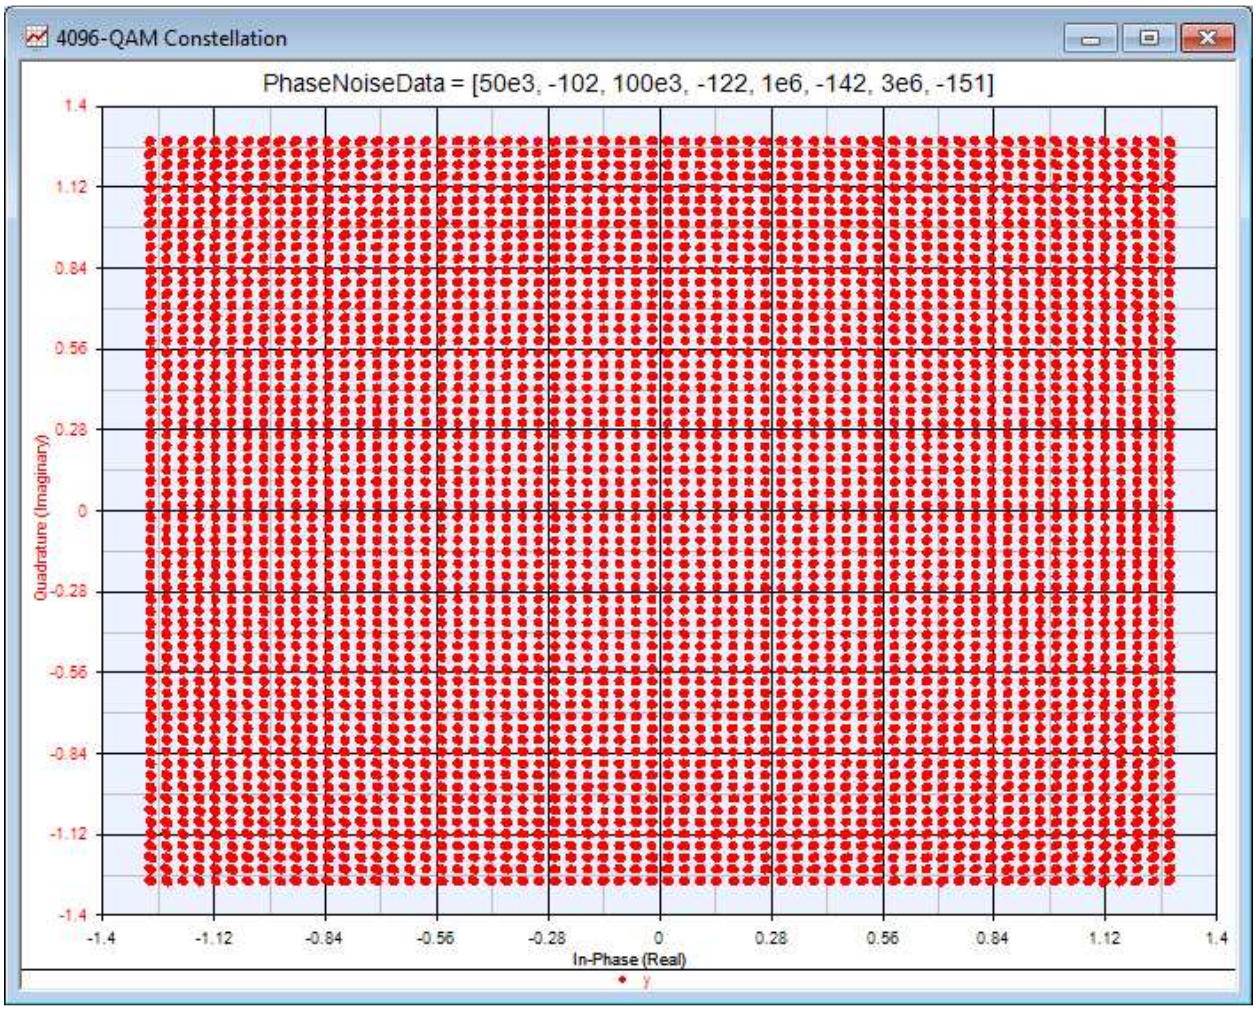

Figure 37 – 4096-QAM@ .0625° rms (SNRφ = 59 dB) (Recommended)

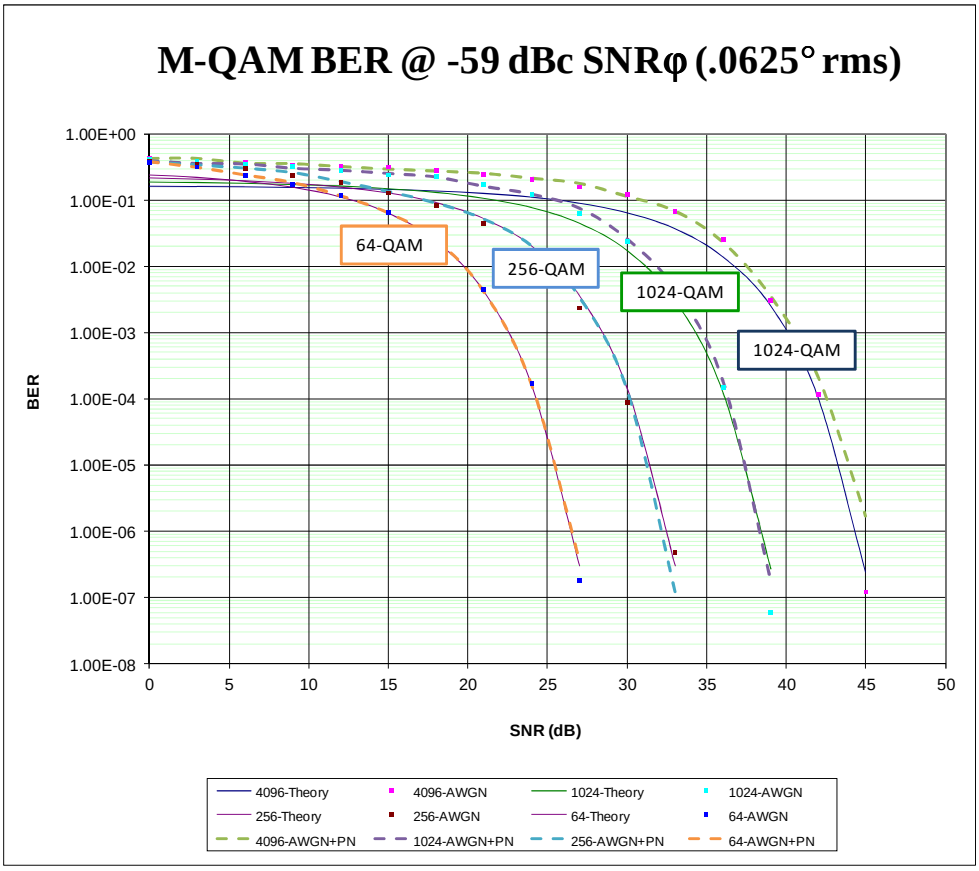

Figure 38 – M-QAM BER for SNRφ = 59 dBc

Post-Detection Processing

For high SNR systems, the loop bandwidth can be, relatively speaking, quite wide. However, it is nonetheless narrow compared to the symbol rates of single carrier QAM signals used in cable. This is important because it means that if there is enough phase noise to contribute to misplacing a symbol in the constellation, it will misplace potentially a large consecutive set of them for single-carrier systems. Because the loop bandwidth is much lower than the symbol rate, a sample of phase noise will be in about the same relative phase location for many symbols in a row – including when the sample is near a decision boundary or across one altogether. This is often referred to as the "slow" phase noise assumption, and is a common characteristic of single carrier QAM systems. The result is that phase noise, as an error mechanism itself, is bursty in nature. This puts pressure on the receiver to have burst correction either via FEC and/or interleaving. Reed-Solomon encoding is burst correcting, but the encoder in the J.83B downstream is only a t=3 symbol correcting design. It instead relies on the interleaver to provide a randomization of the symbol errors to make the RS decoding more effective, spreading out a burst of errors across codewords.

Fortunately, at least in the downstream, J.83B defines a very powerful, configurable, interleaver. It can configure burst protection from 66 usec to 528 usec (Level 2 mode with I = 128) at the expense of introducing latency. The lowest latency value is (I = 128, J = 1), where I and J describe the register structure used to feed Reed-Solomon codeword bits in and out. This setting provides 66 usec of burst protection at the cost of 2.8 msec of latency. Real time voice

is the service that is typically most carefully watched for the latency budget, and 2.8 msec can be accommodated easily in a budget that targets around 50 msec typically one-way. I-128, J=4 is a recommended setting, contributing 11 msec of latency in exchange for 264 usec of burst protection.

The symbol rate of 5.36 Msps (256-QAM) works out to 187 nsec symbol periods. Using 50 kHz to represent the rate of the phase noise process, its "period" (it's a noise process, so period is loosely used) is about 20 usec, or 107 QAM symbols for 256-QAM. The interleaver spreading exceeds 20 usec even for the lowest latency setting. Therefore, the interleaver is a very powerful helper against phase noise impairment - provided the native error rate is low to begin with.

The right-hand side column of Table 11 identifies rms phase noise thresholds that are at the edge of the native BER curve remaining stable. Because of the interleaving downstream, this column could be considered a target objective for the maximum allowable phase noise if it is within the budget of the FEC to support the error contributions from phase noise in addition to other channel impairments it may have been designed to protect against.

Measured performance is available for 1024-QAM in a pseudo "J.83" mode [9]. As shown in Table 14, pre-FEC errors are measured, and these are associated primarily with clipping and phase noise. In each case, however, post-FEC error rate is zero – meaning that the combination of the interleaver and RS FEC was able to completely eradicate any burst errors that may have been caused by the introduction of phase noise.

Unfortunately, in the upstream, we have potentially higher phase noise contributions specified, although it is specified over different ranges that may allow more of the transmit contribution to be tracked. This cable modem requirement was for SNRφ = 44 dBc. Upstream is likely to rely on lower orders of modulation, however, such as being limited to 1024-QAM. However, we have identified the 1024-QAM SNRφ threshold as 53 dBc, or 9 dB better than the cable modem requirement. This has important implications to the carrier recovery requirements in the burst receiver.

The upstream Reed-Solomon FEC is more powerful than the downstream, but still would not be capable of spanning a phase noise induced degradation of 50 kHz of noise bandwidth, much less as low as 8 kHz. This is more than five times the span, so represents about 5 times the number of symbols in error in a row at the highest upstream symbol rate – this likely outlasts the average burst size upstream entirely.

Table 14 – Pre-FEC 1024-QAM Error Rates with Zero Uncorrected Codewords

		1024-QAM Carrier Frequency		
		603 MHz	747 MHz	855 MHz
QAM @ -4 dB to Analog	MER	39.6	39.2	38.9
	BER	6.1E-08	1.12E-07	3.76E-07
QAM @ -6 dB to Analog	MER	39.0	38.9	38.6
	BER	1.5E-07	2.6E-07	2.5E-07
QAM @ -8 dB to Analog	MER	38.3	38.2	37.7
	BER	4.30E-07	2.02E-06	3.48E-06

As such, if impaired by phase noise, post-FEC results should register some low level of uncorrectable codewords. Since FEC is not a source of burst protection from a phase noise perspective, nothing is lost in moving from a RS-based FEC scheme to LDPC.

Note that SNRφ = 44 dBc is about .35° rms, which is plotted in Figure 29 for a 256-QAM – the state-of-the-art throughput available today [12]. In theory, this amount of untracked noise would lead to slightly less than 1 dB of degradation in an uncorrected BER curve. Based on the burst dynamics above, a post-FEC result would register some low level of uncorrectable codewords if there was a phase noise-induced BER contribution measureable. However, in [22], it is shown that there is error-free pre-FEC and post-FEC performance of 256-QAM upstream with SNR = 36 dB. However, this is consistent with the curve for .35° rms even if the SNRφ = 44 dBc is *all* untracked. Thus, it is not possible to learn whether or how much of an rms error reduction takes place in the carrier tracking process.

For purposes of upstream evolution, then, such as beyond 256-QAM, it is impossible to tell from 256-QAM performance, without additional measurements, whether there is adequate margin in the untracked rms phase noise to support 1024-QAM. However, without question, the current CM specification of -44 dBc over the specified bandwidth would be wholly inadequate without the ability to remove substantial induced phase noise in the carrier recovery process. This suggest these requirements may need to be updated to go beyond 256-QAM. The BER curve for 0.5° rms for 256-QAM in Figure 29 would be a reasonable approximation to the trajectory that the 1024-QAM BER would take for 0.25° rms, and this would of course get worse for 0.35°, meaning it would induce more than 2 dB of degradation at low error rates. Without interleaving, there would not be an opportunity to correct for this degradation in the upstream, so this bears consideration. Upstream phase noise for 1024-QAM may create the need for updated requirements.

In summary, to advance the modulation profiles, the phase noise requirements identified in DOCSIS and DRFI may need to be reconsidered to provide the spectral fidelity necessary to support very bandwidth efficient QAM transmissions such as 1024-QAM upstream and 4096-QAM downstream. In the upstream, 256-QAM has been shown to be supported today. It is inconclusive whether or not the phase noise margin contribution that is today adequate for 256-QAM is sufficient also for 1024-QAM. There is no mechanism in place to handle the burst noise environment phase noise-induced errors can create.

In the downstream, a measured post-interleaver, pre-FEC error floor suggests that there is some residual phase noise impact that is being handled well enough by these burst correcting mechanisms. Additionally, as shown in Figure 32, the error rate performance against the combined DRFI and tuner mask would be very poor for 4096-QAM – so badly so that the ability to manage an effective decision-directed tracking loop and successful decoding process is likely to be compromised. While the interleaver is very effective, there is an inherent assumption of low error rate to avoid overwhelming the interleaver-dispersed errors and the carrier recovery subsystem. Performance recommendations for total untracked rms phase noise for 1024-QAM and 4096-QAM are shown in Figure 34 and 37. Under these conditions, there would not be a heavy reliance on the interleaver, FEC

budget, or concern about the sensitivity of decision-aided tracking robustness.

MULTI-CARRIER MODULATION

OFDM Applications to HFC

The industry is considering, as part of the IP transition, adopting a new RF waveform and fundamentally changing the access method away from a line-up of 6 MHz frequency domain multiplexed (FDM) slots. Wideband, scalable MCM or OFDM is being considered for the next generation of RF over HFC, for many of the reasons discussed previously about expanded bandwidths and RF channel uncertainties. There are many acronyms in use that describe an implementation of the same fundamental core concept: lots of narrowband carriers instead of one wideband carrier. Figure 39 illustrates the OFDM concept.

Historically, OFDM applications have been linked by a common thread – unknown or poor RF channels. Virtually all modern RF systems implement some form of MCM – 4G Wireless, MoCA, G.hn, HomePlug AV, 802.11n, and VDSL. The differences are based on the medium and channel conditions expected affecting the band of operation, subcarrier spacing, modulation & FEC profiles, bit loading dynamics, and whether the system is multiple access in the sub-channel domain (OFDMA). Table 15 lists some common Pros and Cons of OFDM.

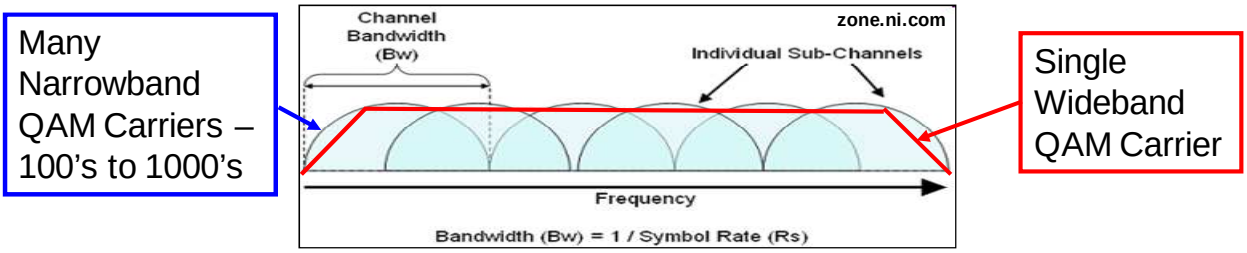

Figure 39 – Fundamental Concept of Multicarrier vs. Single Carrier

Table 15 – Pros and Cons of Orthogonal Frequency Division Multiplexing

Pro	Con (*or Comment*)
Optimizes Capacity of Difficult Channels	High Peak-to-Avg (CPE issue); (*PAR reduction schemes exist - adds OH*)
Simplified Equalization against Frequency Response or Multipath	(*Cyclic Prefix = Guard time OH*)
Robust to Narrowband Interference	Avoidance Approach – Throughput Penalty by Deletion or Mod Profile
Robust to Impulse Noise	*(Similar Principles as S-CDMA - Time Spreading and Parallel Transport)*
Modern Ease of Implementation – IFFT/FFT DSP functionality	Complexity Increase for Shaping and Wavelet schemes – trade-off C/I vs. ISI
Simple Co-Existence via Flexible Subcarrier Allocation (and Power)	Backward Compatibility with DOCSIS
More Spectrally Efficient Wideband Channel than FDM Can be Multiple Access (OFDMA)	Potentially More Sensitive to Synchronization Noise Such as Carrier Phase Jitter (loss of orthogonality)

The most powerful advantage of OFDM has been that it shines in difficult or unpredictable channel environments. With the increasing ability to do computationally complex operations in real time, OFDM implementation – once an obstacle – has become a strength through simple IFFT/FFT functionality that forms the core of the transmit and receive operations.

For HFC, of course, this primary advantage is worth a closer look. The HFC downstream is one of the highest quality digital RF channels available – it is very low noise, and very high linearity. Such channels benefit very little in performance from OFDM, and probably not enough to justify introducing a new waveform if modulation efficiency of today's forward band was the only thing at stake.

However, as discussed, operators are looking for places to exploit more spectrum, and the channel quality of extended coaxial spectrum will be less predictable. This makes OFDM well-suited to be introduced in this part of the downstream band above 1 GHz as shown in Figure 1. Then, as the IP transition moves ahead and legacy 6 MHz slots are eliminated, the spectral flexibility of OFDM through allocation of its subcarriers becomes an especially valuable transition tool.

The HFC upstream, of course, does have a troublesome part of the band at the low end of the spectrum to which OFDM is a good fit. Today, the solution available to exploit capacity here is S-CDMA, which is just now seeing growth in interest and field deployment as upstream spectrum become congested and there is nowhere else to go, but down (in frequency). Like S-CDMA, OFDM should be robust at the low end of the band if properly designed for the impulse and narrowband ingress environments in that region of the return.

Unlike the downstream, above the low end of the band, as the upstream is extended above 42 MHz, there is likely to be steadily *improving* channel conditions, at least up to the FM radio band of 88-108 MHz. As in the downstream, this part of the upstream band – the extended upstream – may have a less obvious need for the primary poor-channel performance value of OFDM. But, there are some unknowns, and the FM band looms. And, since DOCSIS carriers will exist for many, many years, the flexibility of spectrum allocation once again makes OFDM worthy of consideration in this changing environment.

A second well-earned "pro" for OFDM is the simplicity with which poor frequency response can be combated. We discussed this as part of the capacity discussion in the beginning of the paper. However, OFDM also makes difficult multi-path channels more manageable. The HFC network is prone to "multi-path" in the form of micro-reflections associated with impedance mismatches that occur naturally over time and unnaturally through the fact that, as discussed in the section on the POE home gateway, every home in the plant is also part of the access network. Unlike the mobile application, the "multi-path" is static or nearly so. Nonetheless, because the upstream is burst mode from a randomly located source, it has dynamic characteristics associated with the allocation of time slots to modems that are basically on a single frequency but have individually dependent, but unique channel characteristics. OFDM enables the simplification of the equalizer function in these cases.

Perhaps the most talked about disadvantage of OFDM is its inherently high peak-to-average-power ration (PAPR). An OFDM signal is a collection of independently modulated carriers, all sent at once. As such, the composite waveform has

noise-like qualities. This is a potential RF concern, as it requires more linearity, or higher P1dB, in the transmit power amplifier stages compared to single carrier signals to ensure the waveform does not get clipped and distorted. PAPR is primarily an issue for CPE – more transmit power headroom translates to more hardware cost. The same is true in principle (the need for more headroom) for the OFDM receiver, but the receive side is rarely tested from a distortion standpoint, processing very low level signals such that the dB differences have much less impact to design and cost. Schemes that encode subcarriers in a way that reduce PAPR have been developed.

OSI Layer 1 Standard?

An emerging analogy for OFDM is to liken it for OSI Layer 1 what Ethernet and IP are for OSI Layers 2 and 3. A complete, modern Layer 1 PHY is emerging as Multi-Carrier QAM with LDPC-based Block Code. The combination of the two drives implementation very close to theoretical capacity, so there is little else to optimize. There is a natural convergence of solutions towards this combination to yield the highest throughput efficiencies for a given channel. System parameters around the OFDM implementation would vary by application as a function of channel characteristics, as do the block sizes used for the LDPC code. The large number of modern systems based on OFDM is another important factor driving towards a PHY layer "standard" approach.

HFC is no different in this regard – looking for optimal ways to extract capacity on channels that are ill defined or know to be potentially troublesome, while adapting around legacy signals. OFDM or MCM is an alternative suited to these objectives, and some variant is a likely final evolution phase of the coaxial last mile, as introduced in [7].

Channel Impairments and OFDM

As discussed, for AWGN channels, the results relative to SNR and architectures above applies directly. While the HFC channel is never "only" AWGN, this assumption applies well to the HFC spectrum that generally represents the "good" part of the spectrum. There are often modest linear distortions comfortably handled by straightforward time domain equalizer structures. OFDM may achieve high performance with less implementation complexity in these cases, but single and multi-carrier systems would otherwise perform very similarly. The same can be said for the upstream, although the adaptive equalizer complexity in the upstream is much greater, so the weight of a simplifying architecture may be of more value.

However, it is interesting to point out that the maximum attainable bit rate expression on a channel for a given SNR is approximately the same for multi-carrier and single carrier when equalized by a DFE – the approach used today for the DOCSIS upstream [1]. The key difference, again, is that as the channel gets more difficult in terms of frequency response, the theory holds up well, but scale of implementation favors multi-carrier, and more so the worse the channel conditions become. In both cases, feedback from receiver to transmitter fully optimized the bit rate attainable.

Let's take a look at some of the impairment scenarios we quantified for single carrier and discuss how they relate to multi-carrier.

Signal-to-Interference

Single carrier techniques combat narrowband interference by notching the band through adaptive filtering mechanism, as previously discussed. OFDM, on the other hand, deals with narrowband

interference by avoidance. Subcarriers that are imposed upon by an interferer are notched out our dialed down to a more robust (less bandwidth efficient) modulation profile that can be supported, all as part of an adaptive bit loading algorithm. The effect is a capacity loss, but generally a modest one. Note that single carrier ingress cancellation requires overhead itself (lost capacity) in order to operate.

CW Interference

A simple example of the capacity loss for OFDM can be calculated by recognizing the sub-channel spectrum for OFDM when implemented in a pure FFT-based architecture, the simplicity of which being one of the reasons it has become so attractive. The spectrum of a single FFT-based sub-channel is shown in Figure 40. The roll-off of the Sin(x)/x response is slow – 6 dB per octave – so the "in-band" rejection can as a result be low when an interferer falls onto a sidelobe of a particular sub-channel. The first sidelobe has the commonly referenced 13 dB relationship to the main lobe response

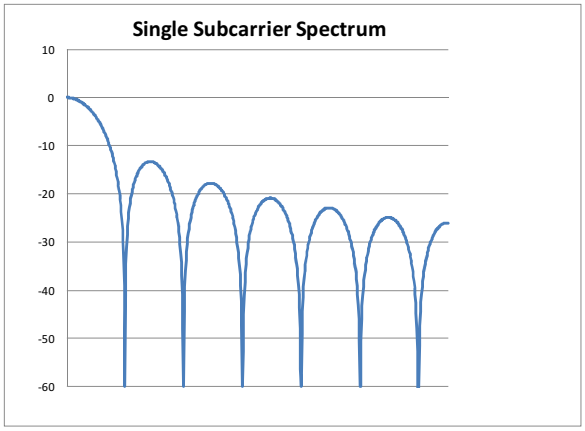

Figure 40 – FFT-Based OFDM Subchannel Spectrum

In Table 7, we noted that the A-TDMA ingress cancellation function had a limit of S/I = 10 dBc for zero corrected codeword errors for single CW ingress signal at an SNR = 35 dB. Three FM interferers could be as high as -15 dBc each (total S/I is still about 10 dBc). How would FM interference at -10 dBc affect the OFDM capacity? Figure 41 shows how several adjacent subcarriers appear in the midst of a CW interfering tone (not to scale of the numerical example).

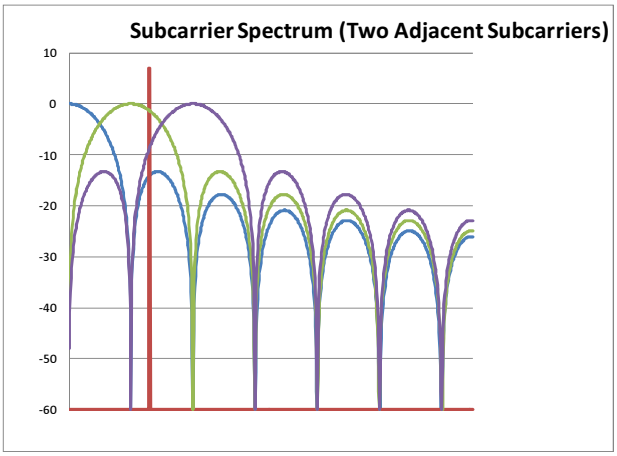

Figure 41 –CW Interference in an OFDM Channel

Let's assume channel SNR conditions are high, such that 64-QAM can be deployed. Going back to our 8k point FFT, each sub-channel is (1/8192) of the total, so the S/I on a per-sub-channel basis is (10-39) = -29 dBc on a subcarrier. Simulations performed such as were shown in prior results for 64-QAM indicate that a required 25 dB S/I for error free BER in this high SNR condition (35 dB). This would require 54 dB of rejection. If the interference coincides with a sub-channel frequency, then it does not interfere with adjacent sub-channels because of the same orthogonality properties that ensure that the sub-channels do not interfere with one another. However, this is unlikely. The worst case is it is just off center of a sub-channel so exposed to the envelope of Sin(x)/x roll-off of the spectrum in Figure 41. For rejection of 54 dB, this occurs at about 160 subcarrier indices away (320 total). If these sub-channels are all nulled, and all FFT

sub-channels are used for payload, then the lost capacity is about 3.9%.

If instead of muting, for example, the adaptive bit loading tries to implement 16-QAM where possible, requiring only 20 dB S/I, then only about 160 sub-carriers *total* are lost, or 1.95% of capacity is lost to muting. There are then (320-160) 160 new subcarriers carrying 16-QAM, which works out to 0.65% of lost capacity, for a total of 2.6%. This is an improvement over muting all of them. Another subset could use 8-QAM, QPSK, etc. This is precisely how OFDM is handy for optimizing under varying channel conditions.

Alternatively, all of the above analysis was performed without considering new error correction. Our "new" SNR requirement is 26/32/38 dB for 64/256/1024-QAM, respectively. Simulations like those already discussed show that for these SNRs, we can arrive at S/I conditions that leave codeword error rates that are easily correctable (1e-4 or lower), for delivering low PERs. This would be a different set of dBc values to meet, but the same approach to the calculation of the carrier indices effected. This approach allots FEC "budget," built around AWGN performance, to correcting for interference induced errors, which would come at the expense of SNR to some degree. Error rate curves are very, very steep compared to classic uncoded waterfall curves, so this type of analysis trade-off would require careful simulation and test.

Modulated Interference

Let's assume the interference is FM modulated. Again referring to Table 7, zero correctable codewords required a 15 dB S/I for 64-QAM with high SNR. Now, however, at 20 kHz wide, its bandwidth is roughly that of an entire sub-channel, so looks like a noise floor increase.

Figure 42 shows the implications of an additive "narrowband" modulated interferer when applied to OFDM (not to relative dBc scale of S/I = 15 dB). On a per-sub-channel level, it represents -24 dBc. The main difference in the analysis approach is that we no longer need to refer to S/I behaviors to quantify how much rejection is needed.

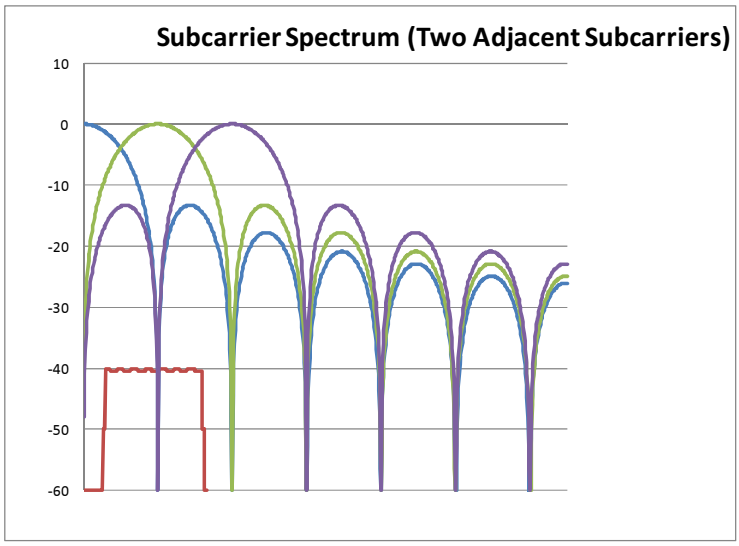

Figure 42 – Modulated "Narrowband" Interference for OFDM

We can treat it instead as a noise floor addition and refer to QAM profile performance against SNR. Modulation that creates a broad noise floor (relative to the sub-channel) and AWGN would not have the same precise effect, but that model is more relevant than a CW S/I model. Based on the thresholds used in Table 3 for upstream and 64-QAM (26 dB), the lost capacity would be close to the CW case. The differences would be in the S/I of a carrier index needing to reach 26 dB vs. 25 dB S/I, and the weighting that would apply for a broad spectrum applied versus a narrow carrier when passed through an FFT receiver.

As previously discussed, a reasonable argument can be made that the margin allotted to the QAM profiles in Tables 2 and 3, which are based on today's single carrier upstream channels, can be decreased *because* of a multi-carrier technique, as OFDM would be naturally more resilient to some of the items that contribute to the margin allotted.

Downstream Distortion Beats

We have noted that CSO/CTB interference has been a cause for concern for downstream QAM. We have also noted that, under the assumptions of analog reclamation, the CSO/CTB levels decrease dramatically. And it was noted that the bandwidth of these distortions is on the order of tens of kHz [20]. So, like the modulated interference previously described, this type of distortion is on the order of a sub-channel bandwidth for OFDM.

For full analog reclamation, the only number that matters becomes CCN, as all the distortions themselves become digital and spread across the spectrum in a noise-like fashion. Unlike the case of the FM interference above, beat distortion have an amplitude modulation component. Under the assumptions in Table 5, the worst case CTB indentified is 66 dBc. On a per-sub-channel basis, this becomes 27 dBc. Considering the noise-like peak-to-average in analysis makes sense for single carrier QAM because the distortion is "slow" in relation to the bandwidth, so peak samples exist for symbol after symbol. In the case of OFDM, the distortion bandwidth is on the same order of the sub-channel width (in this example), so noise averaging takes place just as symbol detection averaging does. As in the modulated interference case, we now can compare this 27 dBc to the modulated thresholds in noise environments to arrive at the impact to OFDM subcarriers.

Referring to Table 2, then, 256-QAM on a sub-channel interfered by this level of CTB would be supported in high SNR conditions (AWGN + CSO/CTB do not exceed the 25 dB shown). No capacity is lost in this case due to CTB for 30 analog carriers and 256-QAM. For 1024-QAM and 4096-QAM, however, threshold SNRs were identified as 31 dB and 37 dB. In both cases, only the sub-channel or two where the CTB falls will be impacted, since the spectral roll-off provides enough rejection (13 dB minimum) to meet these two SNR requirements.

This number of effected channels, too, can be calculated, as distortion noise "lumps" occur at periodic increments – two CSOs and two CTBs every 6 MHz (see Figure 20) – so there will be over 100 sub-channels imposed upon in the 192 MHz example discussed. This would be a maximum of 3.1% of the channels (128 = (192*4/6)). However, since we have shown that these channels could support 256-QAM, the capacity impact is only 0.62% for 1024-QAM and < 1.1% for 4096-QAM (less than 1.1% because some of the sub-channels could likely use 1024-QAM).

The same argument previously made about the FEC budget can be made here,

which applies in particular for 4096-QAM. The SNR thresholds are based on AWGN, so it is the combination of AWGN and new noise contributors like CTB that should meet these thresholds. The "high SNR" assumption would then assume that this beat distortion is then the dominant effect in the sub-channels where it appears. When this is not the case, the offset for the addition of noise power must be made to guide the modulation profiles that can be supported.

For example, adjacent channel rejection of 13 dB from the 27 dBc example is 40 dBc. If our AWGN performance is 40 dB, supporting 4096-QAM with new FEC, then the two combine to 37 dBc, the 4096-QAM threshold identified in Table 2, and no capacity is lost in the adjacent channel. If instead we were already maximized at 4096-QAM with a 37 dB SNR, then the 40 dBc pushes the composite SNR closer to 35 dBc. In this case, a 2048-QAM profile may be required to have a robust channel performance.

Lastly, again, this combination of impairments, when coupled with sharp error rate functions that swing orders of magnitude on a dB of SNR difference, requires robust simulation to quantify precisely.

We nonetheless conclude that, in the case of a partial analog reclamation, OFDM should support the advanced modulation formats with little capacity degradation to do distortion interference.

Other Multi-carrier Approaches

Various shaping techniques and use of wavelets for orthogonality have been studied to reduce the effect of narrowband interference. However, these add complexity, and waveforms that provide narrower frequency response are inherently creating longer symbols in the time domain, so negatively affect performance on dispersive channels. By defining expected channel conditions, an optimum balance of time domain and frequency domain robustness can be implemented.

For an OFDM system for HFC, there is still some homework to be done on the system optimization side to determine the sub-channel spacing, shaping, and channel conditions anticipated and specified as new HFC bands become part of the RF channel definitions. We can count on improvements in SNR and downstream distortions, but updated frequency response and impairment models need a careful examination to ensure the HFC flavor of OFDM is optimized for its channel the way other OFDM-based systems in the wireless and wireline world have been optimized in their applications.

Phase Noise

OFDM creates an interesting scenario with respect to phase noise degradation. A core component of the analysis for the single carrier case discussed previously was built around recognizing that symbol rates for single carrier QAM generally exceed the frequency offsets where phase noise is prevalent. We used the example of 50 kHz of phase noise "bandwidth" to point out that the assumption for degradation is "slow" phase noise. A phase noise sample that rotates a constellation tends to be in the same place over many symbols in a row. If that phase error is close to a decision boundary or across it, there is likely to be a consecutive burst of errors.

Of course, with OFDM, we are using many, many narrow sub-carriers. For example let's use 192 MHz as a maximum OFDM bandwidth – consistent with coexisting with 6 MHz and 8 MHz forward path channel line-ups. An 8k FFT implementation for OFDM would mean that subcarriers are approximately 23.4 kHz apart. For a 16 k FFT, it would be 11.7 kHz.

Compare these to the 50 kHz phase noise "bandwidth" we were using earlier. This phase noise mask then extends *beyond* the sub-channel QAM symbol rate and beyond the main lobe of the OFDM spectrum implemented via FFT. Clearly, this is no longer a case of a phase noise process that is "slow" compared to the QAM bandwidth.

Let's look at an example phase noise spectrum after carrier recovery. We saw a sample spectrum in Figure 24 of phase noise that would be imposed on a QAM carrier by an RF frequency conversion. For slow phase noise, the exact shape is not as important – it is the total rms noise that matters, because the phase noise power is dominated by "slow" or low frequency energy. As such, we referenced "dbc" values of total phase noise based on requirements currently in place. Some of the may be tracked out at the receiver, but in all cases for single carrier it can be characterized as "slow" except perhaps for the lowest upstream symbol rates, which are rarely implemented today.

For OFDM, however, the spectral content of the phase noise mask matters. Consider an example post-carrier recovery mask shown in Figure 43. This is the characteristic lowpass shape of untracked phase noise. It has components associated with RF phase noise imposition, additive noise, and self-noise of the carrier recovery process itself.

Now let's take a look at how this type (two examples) of mask might look against an OFDM sub-channel spectrum. This is shown in Figure 44.

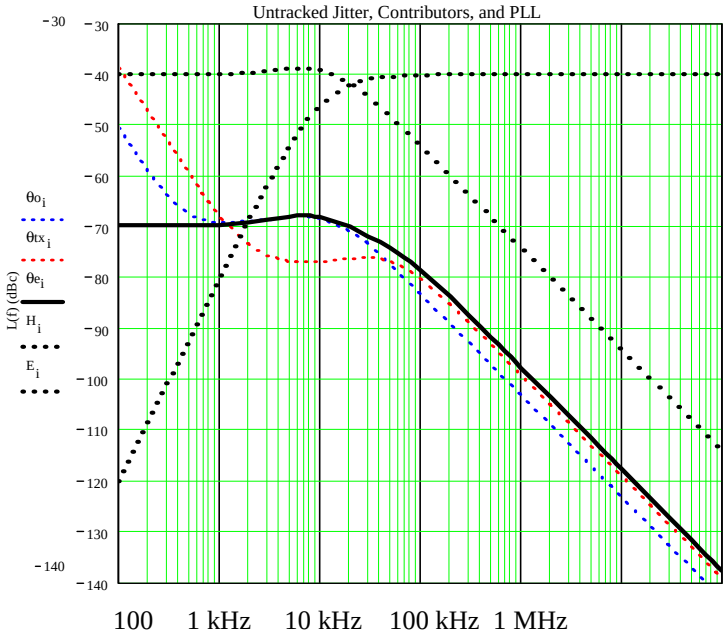

Figure 43 – Example Untracked Phase Noise Spectrum

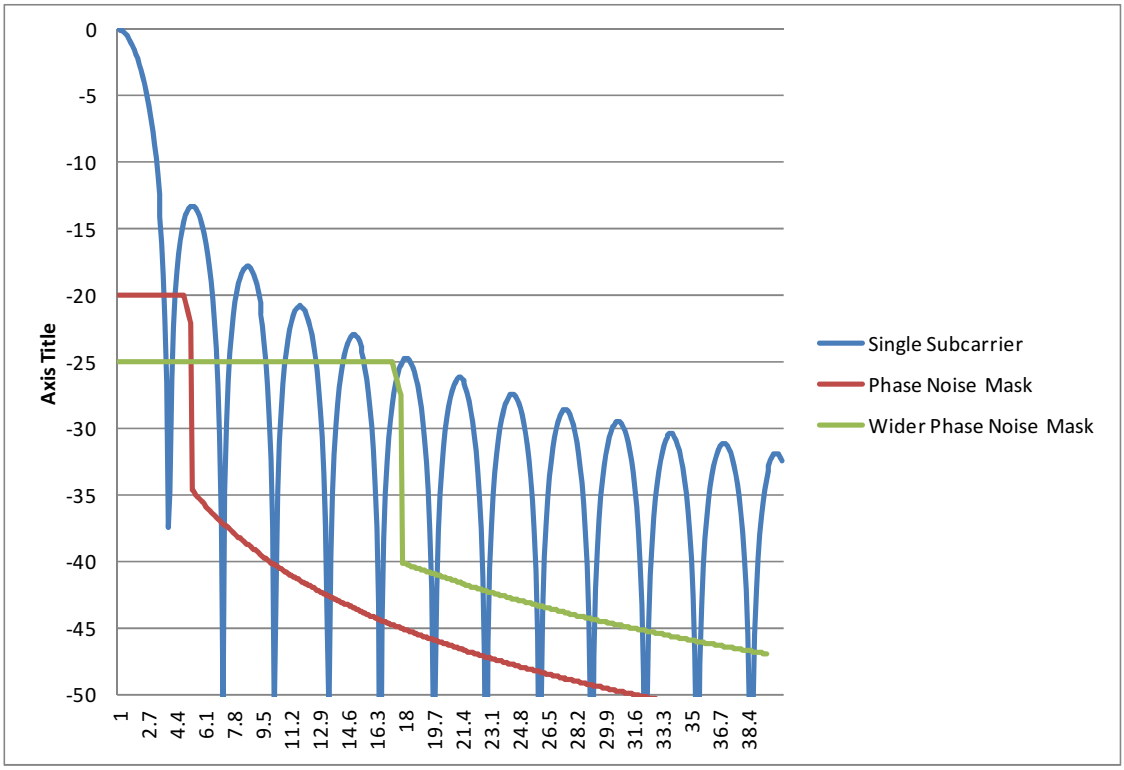

Figure 44 – Untracked Phase Noise vs. OFDM Sub-Channel

Two examples are shown. The red mask, for example, represents how the spectrum of Figure 43 relates to the 16 FFT discussed previously, with its roughly 12 kHz sub-channel spacing. Under this scenario, the 50 kHz of phase noise bandwidth we used earlier to discuss single carrier degradation is shown in green. Every OFDM sub-channel is effectively demodulated with the noise imposed by the phase noise mask.

However, as Figure 44 reveals, phase noise contributes two kinds of degradation to OFDM. There is an error common to all subcarriers related to the "in-band" effects – what for single carrier is the "slow" phase noise. Only, for OFDM, there is a good chance that the slow assumption is no longer valid, advantageously so in fact. It depends on the untracked mask – the shape or spectral occupancy of the phase noise is now important. Moderately varying or "rapid" phase noise allows some averaging over the bandwidth that the symbol is integrated over, and an average of a zero mean process is a better scenario than a single amplitude phase error sample.

However, there is also a component of phase noise that contributes to Interchannel Interference (ICI) as the masks cross into other sub-channel bands *because* of the relative relationship of phase noise mask to subcarrier spacing. This phase noise effect is additive looks nature, just as AWGN (mathematically easy to demonstrate sing the small angle assumption: $\exp(j\varphi) \approx 1+j\varphi$). In Figure 44 it is also clear that the noise in an adjacent band is a function of the phase noise spectrum itself (the shape) weighted by the Sin(x)/x response it leaks into. Not obvious from Figure 44 is that all of the subcarrier phase noise spectra combined create the full ICI effect. Because of this, and because the noise level is monotonically decreasing, the middle sub-carriers are the most effected by ICI due to phase noise.

The SNR degradation due to phase noise for OFDM has been calculated in many

papers, and is simplified in [17] for a basic coherent receiver architecture as

$$SNR(penalty, \varphi) = 1 + SNR * \varphi_{rms}$$

For less than 0.5 dB of degradation, this simply reflects that the rms phase noise would be about 10 dB better than the SNR itself – a common relationship when analyzing additive impairments, again verifying the ICI component of phase noise degradation for OFDM.

Comparing this to the assumptions in Table 11, we see that this is about 3 dB better per modulation when compared to the single carrier, no error correction, slow phase noise case. As discussed, slow phase noise and its angular rotation effect is more painful than an averaging of that noise or an additive effect such as in OFDM. The common phase error on all channels is less degrading when some of the energy is outside of the symbol bandwidth, and the energy that contributes to angular rotation is now no longer slow by definition. This appears to overcome the effect of additive noise contributions that leaks across and create ICI.

The study of these two effects has led to substantial research on the sensitivity of OFDM and studies of the proper carrier recovery approach for OFDM frequency and phase synchronization and manipulation of phase noise processes by transmission and tracking systems. In fact, you can find literature that indicates OFDM is *less* sensitive to phase jitter, or *more* sensitive to phase jitter. And, in fact they can both be correct because the nature of the relationship of phase noise to the QAM carrier has changed, and new variables come into play. The assumptions about those relationships affects the results, and a comprehensive analysis for an HFC version against phase noise mask requirements such DRFI would be necessary for 1024-QAM upstream and 4096-QAM downstream to be effectively deployed.

AND THAT'S NOT ALL FOLKS

We have offered some guidance here on requirements and impairments for new modulation profiles and access techniques, but also recognize that more information is required to make solid requirements and recommendations in many cases. Even so, we have not covered all of the potential angles of the analysis. As more work goes into defining advanced PHY profiles for HFC, we will consider yet the next level of details. This includes items such as new isolation requirements for new service on old and old service on new. And, discussion of the equalizer complexity issues of single carrier, or the pro-con trade-offs of different multi-carrier approaches. We have discussed carrier synchronization in depth, but not timing synchronization. Symbol degradation is a quantifiable problem by quantifying timing jitter relationship relative to the eye diagram and pulse shaping used. We also have not discussed timing requirements that become complex in OFDMA. All of these are important topics for future discussion, along with new depth and insights on what we have discussed here as more variables become known and information complete.

SUMMARY

In this paper, we have discussed HFC architectures and key variables for downstream and upstream in order to allow an increase in spectral efficiency and maintain robust performance. We have provided guidelines for system parameters and discussed specifications of equipment today and the implications of the requirements to support for long term bandwidth efficiency objectives. We have investigated the component parts from optical links to RF links, to CPE, and into the home itself. We have explored options for

network architectures that extend beyond today's bandwidth and carrier access methods, and quantified how such shifts in network design may affect these choice. We have analyzed how these modern multi-carrier methods may be affected by HFC conditions today and moving forward. We have broken them down into downstream scenarios and upstream

All in all, the outlook is hopeful for cable operators to be able to exploit modern tools and enable more spectral efficient use of the network, prolonging its already healthy lifespan. However, indications are that some important changes to business-as-usual may be in store to ensure the required robustness on the most advanced modulation profiles – the silicon itself of course, but also outdoor plant architectures, potential requirements changes to create the fidelity conditions that support the modulation efficiencies of interest, changes in service delivery at the interface to the home, and comprehensively defining RF channels that heretofore have not been used by cable, but would necessarily be so to deliver on the capacity needs of the future. Indeed, there is much to do, and based on lifespan projections of service mixes ahead, now is proper time to be game-planning the transition.

APPENDIX – SIMULATION TOOL

Modeling Environment

Agilent's SystemVue was used to conduct system-level analysis of M-QAM with combinations of AWGN, static narrowband interference, and phase noise. The SystemVue model was comprised of random data source, transmitter, channel, receiver, and a data sink. Data streams at the source and sink were compared bit-for-bit to approximate system impact in terms of Bit-Error-Ratio (BER). Parameter sweeps were conducted for the relative RF levels for both AWGN and interference contributions within the channel.

Modeling QAM

Construction of the baseband signal first required a random sequence generator, such as a PN15 data pattern, operating at the system bit rate. As an example, the system bit rate for 4096-QAM is 12 bits/symbol multiplied by the symbol rate, in this case 5.360537 symbols/second, resulting in a system sample rate of approximately 64.33 Mbps. The random bit sequence was then fed into a symbol mapper, whose states were organized such that errors associated with misinterpreting a received symbol as an adjacent symbol would result in only 1 bit error in the symbol. After mapping, both the In-Phase and Quadrature (I and Q) baseband signals were filtered using root-raised-cosine (RRC) filters. Below are the impulse and frequency response plots associated with the RRC filters with an alpha of 0.12 (downstream excess bandwidth).

Figures A-1 through A-4 show the time domain pulse, the baseband pulse spectrum, and the RF spectrum (SNR = 28 dB) and constellation (SNR = 28 dB), respectively.

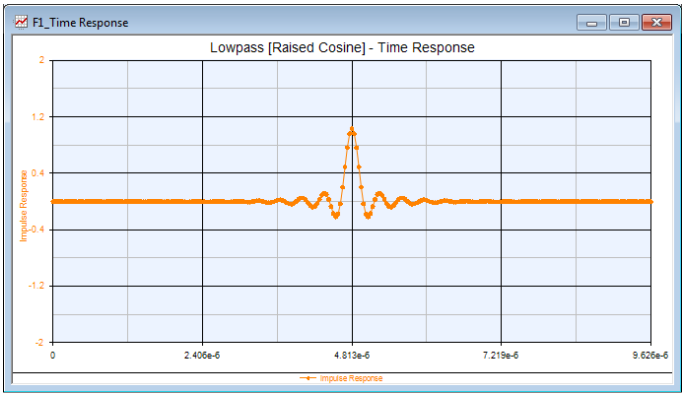

Figure A-1 – Time Domain Root Raised Cosine (RRC) Pulse Shape

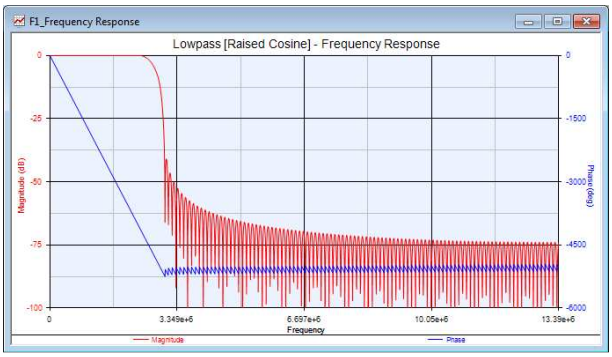

Figure A-2 – Root Raised Cosine (RRC) Spectrum

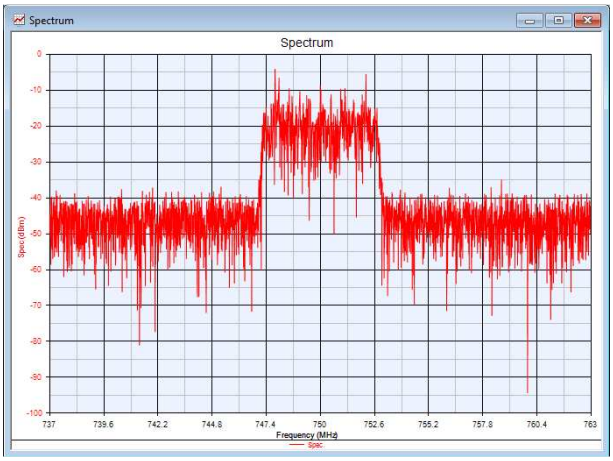

Figure A-3 – RF Spectrum with AWGN – SNR = 28 dB

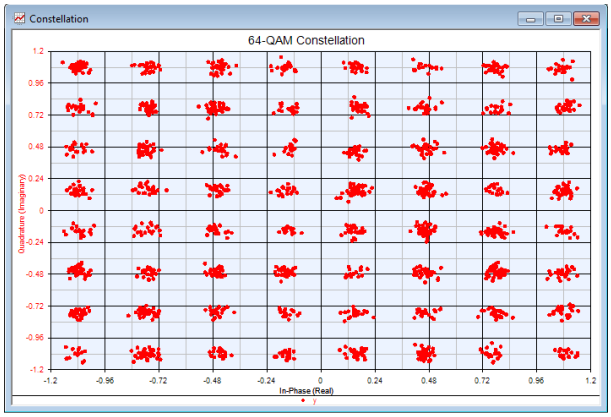

Figure A-4 – 64-QAM Constellation with AWGN – SNR = 28 dB

At the receiver, the signal is demodulated, RRC filtered and de-mapped using the same structures described above in reverse order. The response of 64, 256, 1024, and 4096-QAM to varying SNR measured against theory for uncoded transmissions is verified in Figure A-5. It can be seen that the simulated results track closely with theoretical expectations. This exercise provides the model basis for now extending channel impairments to items such as phase noise and narrowband interference.

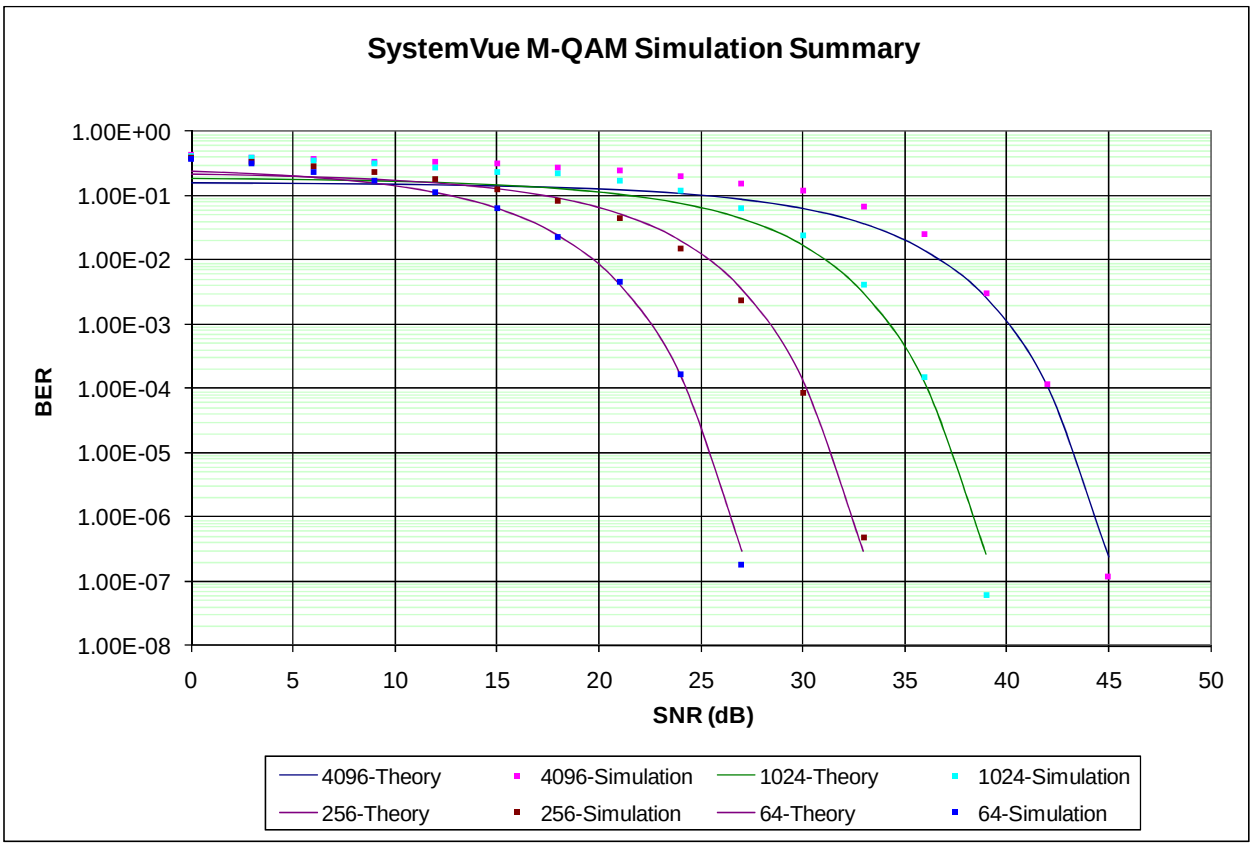

Figure A-5 – Simulated BER vs. Theoretical

REFERENCES

[1] Bingham, John C, *Multicarrier Modulation for Data Transmission: An Idea Whose Time Has Come*, IEEE Communications Magazine, May 1990.

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

[3] Howald, Dr. Robert, *Boundaries of Consumption for the Infinite Content World*, 2010 Cable-Tec Expo, sponsored by the Society for Cable Telecommunications Engineers (SCTE), New Orleans, LA, October 20-22, 2010.

[4] Howald, Dr. Robert, The Communications Performance of Single-Carrier and Multi-Carrier Quadrature Amplitude Modulation in RF Carrier Phase Noise, UMI Dissertation Services, 1998.

[5] Howald, Dr. Robert, *The Exact BER Performance of 256-QAM with RF Carrier Phase Noise*, 50th Annual NCTA Convention, Chicago, IL, June 10-13, 2001.

[6]Howald, Dr. Robert, *Fueling the Coaxial Last Mile*, 2009 Society for Cable Telecommunications Engineers (SCTE) Emerging Technologies Conference, Washington, DC, April 3, 2009.

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

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

[9] Howald, Dr. Robert, *QAM Bulks Up Once Again: Modulation to the Power of Ten*, SCTE Cable-Tec Expo, June 5-7, 2002, San Antonio, TX.

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

[11] Howald, Dr. Robert and Phil Miguelez, *Upstream 3.0: Cable's Response to Web 2.0*, The Cable Show Spring Technical Forum, June 14-16, 2011, Chicago, IL.

[12] Howald, Dr. Robert L., Phillip Chang, Robert Thompson, Charles Moore, Dean Stoneback, and Vipul Rathod, *Characterizing and Aligning the HFC Return Path for Successful DOCSIS 3.0 Rollouts*, 2009 SCTE Cable-Tec Expo, Denver, CO, Oct 28-30.

[13] Lindsey, William C. and Marvin K Simon, Telecommunication System Engineering, Prentice-Hall, Englewood Cliffs, NJ, 1973.

[14] Mengali, Umberto and Aldo N. D'Andres, Synchronization Techniques for Digital Receivers, Plenum Press, New York, 1997.

[15] Miguelez, Phil, and Dr. Robert Howald, *Digital Due Diligence for the Upstream Toolbox*, 2011 Cable Show, Chicago, IL, June 14-16.

[16] Piazzo, L and P. Mandarini, *Analysis of Phase Noise effects in OFDM Modems*, Technical Reprt No. 002-04-98, INFOCOM Dept. University of Rome "La Sapienza", May 1998.

[17] Proakis, Dr. John G, Digital Communications, McGraw-Hill, New York, 2001.

[18] Robuck, Mike, *Cox, Motorola lay claim to new return path speed record*, CedMagazine.com, March 01, 2011.

[19] Stoneback, Dean, Robert Howald, Tim Brophy, and Oleh Sniezko, *Distortion Beat Characterization and the Impact on QAM BER Performance*, 1999 NCTA Convention, Chicago, IL, June 13-16.

[20] Stott, J., *The Effects of Phase Noise on COFDM*, EBU Technical Review, Summer 1998.

[21] Thompson, Robert, *256-QAM for Upstream HFC Part Two,* 2011 SCTE Cable-Tec Expo Atlanta, GA, November 15-17, 2011.

[22] CableLabs, Inc., *Return Laser Characterization Techniques - Results and Recommendations*, Engineering Report, September 21, 1999.

[23] Data-Over-Cable Service Interface Specifications Physical Layer Specification (DOCSIS-PHY), CM-SP-PHYv3.0-I08-090121, January 21, 2009, Cable Television Laboratories, Inc.

[24] DOCSIS Downstream RF Interface Specification (DRFI), CM-SP-DRFI-I10-100611, June 11, 2010, Cable Television Laboratories, Inc.

ACKNOWLEDGEMENTS

The authors would like to thank Stuart Eastman for his invaluable contributions to this paper.