QAM Bulks Up Once Again: Modulation to the Power of Ten

Robert L. Howald, Ph.D.
Motorola Broadband Communications
Director, Transmission Networks Systems Engineering
101 Tournament Drive
Horsham, Pa. 19044
rhowald@gi.com

The use of 256-QAM has become more prevalent in cable systems, as both a video and data modulation. With its eight bits per digital symbol (8 bits/symbol), it offers the highest bandwidth efficiency available today among digital cable signals. Expectations are that 256-QAM will evolve to become a dominant modulation format of the digital multiplex. With the value of bandwidth at a premium, particularly based on bandwidth consumption trends of the past, it is of continued interest to find techniques that increase throughput capability. In HFC systems, there are various ways to create more bandwidth, such as increasing fiber counts, implementing equipment segmentation in the plant, adding wavelengths, or improving compression techniques in digital television signals. The migration of 64-QAM, which represents 6 bits/symbol, to 256-QAM and its 8 bits/symbol, provides 33% more efficient bandwidth usage. Of course, this is at the expense of a higher SNR requirement, as well as increased sensitivity to other impairments, such as phase noise and interference. Fortunately, outstanding features of the cable channel include its very high SNR and its excellent linearity characteristics, each because of the need to support analog video. However, the theoretical capacity of a 6 MHz channel with a 45 dB SNR is actually about 90 Mbps, and, yet, for 256-QAM, the transmission rate works out to only about 40 Mbps, and a fraction of this is consumed by error correction overhead.

A natural question arises as to how much more efficiently bandwidth can be used in the modulation domain. Prior analysis of 256-QAM presented at previous industry-wide technical forums shown it to have higher sensitivity than 64-QAM to nonlinear impairments - higher than the 6 dB theoretical difference in performance in the typical additive thermal noise channel. The next higher order, square-constellation modulation is 1024-QAM. This technique achieves an efficiency of 10 bits/symbol, or a theoretical efficiency of 10 bits/sec/Hz. At this modulation efficiency, it provides another 25% throughput capability above 256-QAM, and an impressive 67% improvement relative to 64-QAM. The means to implement this modulation format is now available in the latest generation of DOCSIS-based cable modem chipsets. This completes several years of modem research and development for this very high order of modulation in an RF channel application, such as HFC. It is now time to understand its practical performance. This paper will examine the required specification of the impairments for successful transmission of 1024-QAM. A discussion of the theoretical needs to support this approach will be presented, as well as a practical look at what an HFC channel will mean for 1024-QAM signals. In particular, the discussion will summarize the effect of HFC-specific impairments on such a scheme and compare the performance to 256-QAM and 64-QAM. Finally, we will draw some conclusions about its suitability of implementation in HFC systems.

Introduction

The use of 256-QAM has become increasingly popular in systems deploying digital signals in the forward band. The expansion of the symbol set provides a 33% increase in bandwidth efficiency over 64-QAM. One of the key reasons that the migration is able to take place is that the CNR in the forward path of a CATV network is capable of supporting this sophisticated modulation format. Delivering CNR numbers in the 40 dB range with very high linearity provides great flexibility in choice of modulation type, and it allows system designers to focus on bandwidth efficient modulations, rather than power-efficient modulations. Perhaps most importantly, modem technology has become robust enough to handle the complexity of tracking, equalizing, and detecting 256-QAM symbols. As the constellation size expands, the sensitivity of the signal to forward path impairments also increases, and generally over and above the 6 dB difference in thermal noise sensitivity that exists between families of M-QAM related by a factor of four. That is, while we can use as a good rule of thumb that 16-QAM, 64-QAM and

256-QAM differ by 6 dB relative to their AWGN performance, this is not so, for example, when we consider other impairment types, such as narrowband interference, group delay variation, phase noise, etc.

In this paper, we will take what we have learned over the years, and presented in prior forums [1][2], and apply it to the next level of modulation sophistication being discussed and developed: 1024-QAM. As might be expected from prior efforts in this area of study [1][2][3], any single impairment can easily escalate into a full discussion topic all by itself. Thus, the focus of the paper will be on the impairment mechanisms of most significance – thermal noise, interference, and phase noise. A brief discussion of other impairment issues will also be discussed.

Description of the 1024-QAM Signal

Waveform

As with other forms of QAM on the cable plant, 1024-QAM is a multiple-amplitude, multiple-phase waveform. In other words, information is carried in both the amplitude and in the phase of the waveform. The RF carrier will change its level, its phase, or both, each symbol period, unless the same symbol is being transmitted. The likelihood of this happening in a well-designed system is small. The difference between 1024-QAM and 64/256-QAM is simply in the number of possible amplitude/phase state combinations which exist. For 1024-QAM, this number is simply M = 1024. A look at the constellation diagram in Figure 1 immediately points out the number of amplitude levels by glancing along either the I or the Q axis. There are 16 possible envelope amplitudes of both I and Q waveforms, and error performance is governed simply by the distance between them. Thus, this constellation density means that either the transmit power must increase to operate at the same BER as a lower order modulation, or for a given transmit power, this modulation will be more likely to have detection errors at the receiver. We will quantify this relationship in a later section.

Bandwidth Efficiency

The 64-QAM signal (M = 64) represents $2^6 = 64$ states, and therefore each symbol represents 6 bits. The QAM is sent over the system with symbol rates of roughly 5 Msps (plus change). The result for 64-QAM is simply the product 5 x 6 = 30 Mbps. Note that this is simply the raw bit rate in the channel, not all of which represents payload information, since forward error correction requires redundancy. For M=256, each symbol represents 8 bits, and thus a roughly 40 Mbps rate, or a (40-30)/30 = 33% increase in bandwidth efficiency. For a 1024-QAM carrier, the bandwidth efficiency is derived by noting that $2^{10} = 1024$ states, resulting in about 50 Mbps per carrier. This represents a 25% increase in bandwidth efficiency above 256-QAM, and a 67% increase above 64-QAM. Each increase in M, therefore, provides more efficient use of available bandwidth. This comes, of course, at the expense of channel fidelity requirements.

Theoretical AWGN Performance

The bit error rate (BER) curves for M-QAM are straightforward to develop in an AWGN-only channel, since symbol-by-symbol, hard-decision decoding is optimal and symbols are uniformly affected. Recognizing that most symbols are bounded on four sides with decision boundaries for large M easily generates upper bounds. However, more accurate solutions are not difficult to develop. Assuming one bit error for every symbol error, a situation that can be assured under what is referred to as Gray encoded mapping of bits to symbols (adjacent symbol differ by one bit only), we can show the following bit error probabilities as a function of SNR, assuming the SNR associated with optimal detection mechanisms. This is a key point to make for CATV, because SNR for analog video references a pre-defined measurement bandwidth, as it is in most analog systems. For digital systems, these BER expressions imply a design that performs matched filtering and optimal detection criteria. The BER expressions are given below.

$\mathrm{Pe(64\text{-}QAM)} = (7/12)\ Q\ [(\mathrm{SNR}/21)^{1/2}]$
$\mathrm{Pe(256\text{-}QAM)} = (15/32)\ Q\ [(\mathrm{SNR}/85)^{1/2}]$
$\mathrm{Pe(1024\text{-}QAM)} = (31/160)\ Q\ [(\mathrm{SNR}/341)^{1/2}]$

Figure 2 shows these expression for theoretical M-QAM BER, M = 64, 256, 1024. Recognizing that signal power is related to energy per symbol as $Es = Ps \cdot Ts$, and that $Eb = Es / \log_2 M$, these expression can be written in the form common to digital communcation theory, which uses Eb/No, as

$Pe(64\text{-QAM}) = (7/12)\, Q\, [(2E_b/7N_o)^{1/2}]$
$Pe(256\text{-QAM}) = (15/32)\, Q\, [(8E_b/85N_o)^{1/2}]$
$Pe(1024\text{-QAM}) = (31/160)\, Q\, [(10E_b/341N_o)^{1/2}]$

In this case, N_o is the noise power *density*, and the Q( ) function is a well-known, oft-tabulated function associated with the solution to the integral under the upper tail of the Gaussian probability density function (PDF), which is thus accounting for the statistical nature of the AWGN.

As mentioned, these expressions represent the hard decision decoding performance of QAM in an AWGN only environment. Better performance cannot be achieved without error correction techniques, although more sophisticated detection mechanisms are used in the demodulator. However, those detection mechanisms are doing their best to estimate the performance when the channel has memory or distortions. The best they can do is achieve this matched filter bound, which is achieved in a memoryless, distortion-free channel. If error correction techniques are considered, however, performance enhancements can be obtained. Forward error correction (FEC) takes advantage of asserting a known, linear algebraic structure on the data, spreading the symbol further apart in linear algebraic signal space, so they are less likely to be “near” one another. A simple way to envision error correction is to consider sending every symbol three times, consecutively, and taking the best two out of three. The likelihood is that there will be much fewer errors this way, but of course there is substantial extra transmission speed required (bandwidth) for the same throughput, or a significant slowdown in throughput for the same bandwidth. This indicates that, while this is a valid code, it is an inefficient one. Other techniques implement a signal-space code whereby, to preserve bandwidth, redundancy is implemented by expanding the signal set from, for example, QPSK (4-PSK) to 8-PSK. Trellis-coded modulation (TCM) is a special case of signal space codes, and recognizes that the functions of modulation and coding do not have to be independently performed. By combining them, significant coding gains and major strides towards capacity are obtained. Advanced physical layer work in support of 1024-QAM involves a TCM format.

Impairment Library

The list of possible channel impairments in a CATV system is larger in breadth, although not necessarily magnitude, than most digital communication systems. All of the following exist on an HFC channel at enough of a level to be considered during link budget development, particularly for M-QAM for large M. The wide variety of possible issues occurs primarily due to the nature of the system infrastructure – it is both an optical link and an RF link – and the nature of the signal carriage – it is a very wideband multiplex of multiple signal types, including both analog and digital. A comprehensive list of impairments is shown below.

Thermal Noise
Reflection paths
Frequency response flatness
Interference - RF nonlinearity
Group delay variation
Phase Noise
Laser Clipping
Spurious AM

For this discussion, focus will be on three major trouble spots, as well as some comments on a few others that are insightful into general issues associated with M-QAM as M increases.

Additive White Gaussian Noise (AWGN)

The CATV forward channel has, as its most attractive characteristics, two key qualities. First, the channel has an exceptionally high signal-to-noise ratio (SNR), at least by digital communication system standards. Of course, this is due to the care and feeding required to deliver high quality analog video signals. Second, the channel is of extremely high linearity – again by digital communication standards. And, again, this is driven by the need to support the needs of analog video, in this case relative to its sensitivity to interfering narrowband spectral components caused by distortion. These two qualities present a digital communication designer's dream, and thus the introduction of some of the most sophisticated, commercially available modulation schemes in use today. Recognizing that the digital communication systems are more aptly described in terms of their error performance, the bit error rate metric is of particular significance, and particularly for continuous transmission channels (in contrast to burst systems, such as the reverse path). For QAM modulation without error correction applied, the SNR and BER are related uniquely in an AWGN-only channel, and these performance curves are the classic waterfall curves present in every digital communications textbook, where they are identified as theoretical performance. These curves are presented in Figure 2. In more advanced studies, these curves represent baseline performance against which it can be compared to the effect caused by other impairments, such as the distortion mentioned above. This paper will introduce some of these issues, and identify how they might be expected to effect 1024-QAM signals, in much the same way as prior studies have discussed the sensitivity of 256-QAM to these issues [1,3]. This section, however, will focus on the basic AWGN impairment.

Additive Noise and Digital Detection

Additive white Gaussian noise, by its definition, is describing all of its important characteristics needed to predict QAM performance. The noise is the aggregation of many contributors, but most notably the optical link, the cascade of amplifiers and their associated noise figures, and the relatively poor noise figures of settop boxes. Recent studies are recognizing the potential for Headend noise associated with the addition of broadband noise floors of many QAM transmitters [3]. This source of the noise is additive, in that it sums with the desired signal. White – a reference to white light's containing all of the wavelengths in the optical spectrum – refers to its spectral content. In the RF world, it indicates a broad, flat noise spectrum, such that noise power in a particular bandwidth is simply the bandwidth, B times the noise density – with density in units of dBm/Hz, for example. This noise density term, No, is often seen in digital communication literature in BER plots, along with the bit or symbol energy term, Eb or Es, rather than SNR. The terms Es/No and SNR are interchangeable when referring to a Nyquist system, which defines a noise bandwidth. Finally, Gaussian refers to the statistical characteristics of the noise, essentially describing how the noise amplitudes vary. This provides a way to mathematically predict the average behavior of the link by quantifying the likelihood of a noise amplitude being large enough to cause a transmission error.

The results of evaluating QAM when imposed upon by AWGN are straightforward to determine. Figure 2 shows the theoretical bit error performance for 64-QAM, 256-QAM, and 1024-QAM. The bit rates associated with a 5 Mbaud transmission rate are, roughly, 30 Mbps, 40 Mbps, and 50 Mbps, consistent with the bits-per-symbol efficiency of 6,8, and 10, respectively. The simple rule of thumb is that for "good" BER, these modulations are about 6 dB apart in AWGN performance. Thus, at a BER = 1E-8, uncoded 64-QAM requires about 28 dB SNR. Thus, without even considering the coding gain (which reduces payload throughput due to overhead required), there is significant SNR margin on an HFC system, even recognizing that digital signals are run at lower per-channel power than the analog signals in the multiplex. Of course, actual transmission includes forward error correction (FEC). With roughly 5 dB of coding gain applied, the margin increase even more.

Considering 256-QAM, the required SNR at the same BER requires about 34 dB. Accounting for coding gain, which is slightly less for this more complex scheme, results in about a 30 dB requirement. Again, even with digital levels backed off 6 dB relative to analog, resulting in SNR values in the high 30 dB range, clearly there is still substantial margin relative to AWGN. Of course, the margin is 6 dB + 1 dB (coding) = 7 dB less, which is not an insignificant loss of robustness. However, with 64-QAM run at levels 10 dB below analog, compared to 6 dB backed off for 256-QAM, some of the AWGN margin is regained. The net result is margin of about 3 dB less than 64-QAM under these assumptions. However, as a result, performance issues with 256-QAM associated with less AWGN margin combined with increased sensitivity to other possible degradations can materialize.

Noise and the Power of Ten

Finally, consider the case of 1024-QAM. Certainly, it is easy to show through classic Shannon capacity theory that diminishing small error rate can be obtained so long as the data rate is kept below 80 Mbps for a 40 dB SNR channel, and slightly less as SNR drops below 40 dB due to digital levels being carried below analog levels. Shannon capacity in this calculation is precisely for the AWGN-only channel. As mentioned, 1024-QAM encodes ten bits per symbol, resulting a single 5 Msps transmission offering 50 Mbps. Per Shannon, then, 1024-QAM is absolutely possible in theory. In fact only 60-70% of the available capacity is being used. It should be noted that, since the development of Shannon theory (1948), it has been quite commonplace through history to not be very close to channel capacity simply because Shannon did not point out how to accomplish it. Only relatively recently – the 1980's and 90's - have quantum strides been made.

Based on the simple rule developed above, a 1E-8 uncoded BER for 1024-QAM requires a 40 dB SNR. This condition is shown in Figure 3. With some coding gain, assume this requirement drops to about 37 dB. There is perhaps some room to further adjust digital levels, provided that, for a large bandwidth digital multiplex, the composite digital power does not increase significantly. In other words, there is not much power loading room to add more signal energy without resulting in digital distortion - composite intermodulation distortion, or CIN - problems elsewhere, particularly in the analog band. The amount of digital bandwidth and power that can be allocated is a simple power loading calculation made with basic numerical tools that can account for the number of channels, powers of each, and slope. Clearly, the additional layer of modulation sophistication is beginning to come up against even the digital-friendly limits of the CATV channel. Even were AWGN the only issue, it is clear that concern over QAM performance margin versus thermal noise immediately exists, a case that for the introduction of 64-QAM was minimal risk.

Noise Figure

Some simple calculations will show what all this means to the signal delivered to the home and to the settop. Assume the end of line analog signal gets to a one-splitter home at 0 dBmV and a 45 dB CNR. Of course, the numbers can be tweaked however desired, but this basic example will set up the issue. Assume that the 1024-QAM signal is 3 dB lower than the analog, and thus arrives at a 42 dB SNR (5 MHz BW reference compared to 4 MHz actually costs a dB, but the loose change is being ignored for the greater cause at this point). For the 1024-QAM signal to achieve a 40 dB SNR – an uncoded 1E-8 – when it starts with a 42 dB, the home side has to deliver a 45 dB SNR. The signal is at –3 dBmv, so the home side noise power must be at –48 dBmV to get to a 40 dB total SNR. Starting with a floor of –58 dBmV/5 MHz, this results in a NF requirement of 10 dB.

Now, several dB is tied up in the splitter and filtering before hitting the frequency conversion in the settop. Tuner noise figures (NF) are in the 7 dB minimum to 13 dB kind of ranges, not due to any technology limit – mainly because better wasn't required. Thus, from the input to the home to the receiver demodulation, the total NF may range from a minimum of 8 dB (no splitting) up to perhaps 19 dB with a couple of splitters. Clearly, there may be issues since only 10 dB can be allocated, and 13-15 dB perhaps if we account for coding gain. No matter how the numbers are toyed with to create scenarios, the point is that there are perfectly reasonable scenarios that exist that the 1024-QAM system will be dealing with that are not very favorable to it without lower noise floor receivers. Amplification prior to frequency conversion, of course, can set a lower noise figure, but the fallout is having to deal with the subsequent distortion issues of this topology.

Interference

Many studies have evaluated the impact of narrowband interference on BPSK, QPSK and, in some case, QAM systems, among others. The impact off an additive interference of a CW tone is to create a rotating phasor (relative to a coherently detected QAM symbol) that is added vectorially to the signal. Figure 4 show this effect for an S/I = 30 dB for 256-QAM. The "doughnut" is the telltale sign of this rotating phasor, and can be viewed with relative clarity on a constellation display in the lab for 64-QAM [3], where the interferer must get pretty large – and thus viewable to the naked eye – to see it on this linear display. The effect is to grow the ideal symbol point out into a ring. As pretty as this might look, the quantitative effect is to move the desired points closer to decision boundaries, rendering them more likely to be kicked over the boundary by a noise sample. This is a relatively straightforward geometry problem to solve and generate error rate estimates.

Considering 1024-QAM, the situation of course gets more difficult to manage. As expected when the density of the constellation points increases, the sensitivity to interference does also. Compare the S/I = 30 dB case for 256-QAM to the S/I = 35 dB case in Figure 5. The increased fidelity is apparent, and this condition closely approximates the case where errors just begin to get counted in a laboratory test on today's 256-QAM modems. Referring now to

Figure 6, it is apparent that the S/I = 35 dB situation is much more serious for 1024-QAM. The BER estimates in Figure 7 that show degradation with increasing S/I support what would be expected when observing the "doughnuts" of Figure 6, and in particular how close their radius comes to the decision boundaries.

CTB as Narrowband Interference

The effect of CTB and CSO on analog video carriers has been studied and is well understood. The effect of CTB on digital carriers is still work in development [2,3]. However, what is known through study and measurement is that QAM carriers are impaired by narrowband interference in the band, without receiver processing that can provide some mitigation. This problem is relatively easily analyzed for CW interference combined with AWGN under certain assumptions. This can be used as a starting point to development of CTB. A way to envision CTB is to consider it as a narrowband interference - tens of kilohertz compared to 6 MHz of bandwidth - of varying amplitude. In order to make use of this information for error rate analysis, how the envelope varies must be statistically characterized, and this can be obtained via measurement and statistical analysis [3]. The simplification in such an analysis is that it ignores the frequency spread of the CTB beats associated with unsynchronized RF carriers. While this simplification does not significantly affect the performance analysis for directly detecting QAM symbols with additive noise and interference, it does have a practical impact for the implementation at hand that uses sophisticated equalization techniques to mitigate channel degradation. The impact is that this processing is more effective for the CW interferer than the interferer with spectral energy spread such as the CTB case. Thus, a demodulator with equalization may measure better than predicted (no equalizer assumed), but when the interference becomes CTB, this may not be the case, as the mitigation ability is lessened.

Let's first examine the interference case to gauge the sensitivity of 1024-QAM to interference, in general. It was found [3] that 64-QAM and 256-QAM, while differing by 6 dB in AWGN, difference in 12-14 dB in sensitivity to narrowband interference. That is, while an S/I of 24 dB performed adequately in tests of 64-QAM, a number such as S/I = 37 dB was required for the same error rate performance for 256-QAM. The significance to this is that common CTB requirements of 53 dB for video fall to 47 dB relative to video. By recognizing that this is an average value, and that the CTB signal is noise-like in amplitude distribution, with a peak-to-average in the 15 dB range, the result is an S/I = 32 dB or worse. Of course, live video lowers the average analog load in real plants, and will drop the actual CTB values. This adds necessary margin into this analysis, dropping peak interference levels accordingly 4-8 dB. However, as mentioned, a CW interference does not tax the receiver equalizer, because the amplitude distribution is deterministic. In other words, the receiver equalizer mitigates the effects of CW interference very well. By contrast, receiver equalization is considerably less effective on real CTB beats, which have frequency content of tens of kilohertz and varying envelope amplitudes. Thus, in practice, nominally adequate CTB performance end-of-line can still yield a condition where symbol errors are counted on 256-QAM. This has been particularly seen in cases with substandard (by a few dB) CTB performance [3]. The effect on 1024-QAM will be magnified further still. The relationship between 64-QAM and 256-QAM for interference was on the order of 12 dB, and it is logical that this same delta would be expected for 1024-QAM compared to 256-QAM. More powerful receiver processing can be expected to be employed. However, the sum of the boost in capability is necessary simply to overcome the additional complexity of the higher order modulation.

Another aspect of distortion is its frequency domain characteristics across the full forward band. Taking advantage of this in a way that may enhance QAM performance will be discussed later in the paper.

Phase Noise

Many studies have generated predictions for digital communication systems with phase referencing error. The systems under consideration here are coherent systems, meaning that a replica of the RF carrier phase must exist at the demodulator in the settop in order to demodulate the signal. This is accomplished through carrier tracking loops, which are based on two primary mechanisms – nonlinear carrier regenerators, and data-aided feedback structures. The high SNR and likelihood of good error rate performance make the latter particularly attractive, since they offer significantly better noise performance in these conditions, because of the noise-free aspect of correct data decisions in making a phase estimate.

Carrier tracking loops are designed to provide phase referencing operation for demodulation. Of course, nature gets in the way, and noise effects impose on the ability to provide an ideal carrier. The nature of the tracking mechanism is that it will follow phase variation imposed on the carrier up to its tracking bandwidth, but any phase noise imposed outside of that tracking bandwidth contributes to untracked phase error. That is, the mechanism high-pass filters the carrier-related phase noise. However, the choice of the bandwidth parameters is limited on the lowpass side by noise and modulation related effects that disturb the tracking. For suppressed carrier tracking systems, such as QAM, methods must be used that create a carrier to coherently detect. For the case at hand, this function is, as mentioned, provided by decision-directed tracking loops during steady-state operation, which significantly enhances noise performance relative to alternatives such as xN recovery or data-aided structures that do not rely on actual decisions. When implemented properly and in low BER environments, tracking loops that implement accurate decision information can approach classic, unmodulated PLL noise tracking performance.

Advantage: HFC

CATV QAM carrier tracking needs and requirements generate some unique requirements compared to what is commonly seen in the literature. The reasons for this are many of the same reasons discussed above. The CATV channel is particularly high SNR, and thus has more flexibility in the design of the carrier tracking loop parameters. Systems with lower SNR must consider more heavily the effects of the additive noise within the loop bandwidth creating jitter. Additionally, the most referenced literature on the topic deal with less sophisticated modulation formats, such as QPSK. The requirements of QPSK are considerably less stringent than those of higher order QAM, because the sensitivity of a QAM scheme is dominated by the effect of the phase error on the outermost symbols in the constellation. As the outermost symbols become "further" away, this sensitivity increases. The exact sensitivity to untracked phase error is quantifiable [1], and, through the recommended values, the various contributors to untracked phase error can be allocated. The key for M-QAM is to recognize that the constellation points are not uniformly effected.

For a given tracking receiver structure, the controllable contributors to untracked phase noise are RF-related impairments that occur during the frequency conversion process. In particular the Headend upconverter and the tuner downconverter (settop) generate random carrier instability that must be quantified. It is not particularly difficult to create RF carriers clean enough not to degrade 64-QAM signals, although consideration should be given to the key components in this regard. However, there is substantial existing converter hardware in the field that was not designed to carry QAM signals. By accident, this type of existing hardware turns out to be compatible with 64-QAM transmission. However, this is not necessarily the case for 256-QAM without taking a closer look at some details of the implementation and RF spectral purity, and it is certainly not the case for 1024-QAM.

The results for 64-QAM show that an untracked error of about 1° rms or less is required to assure small (1-2 dB) degradation of the BER curve for a BER of 1E-8. This condition is shown in Figure 8. Recognizing a convenient rule of thumb that a 35 dBc signal-to-untracked phase noise ratio is equivalent to 1° rms, the 64-QAM requirement can be viewed as an allowable total untracked phase noise power of –35 dBc. When the constellation expands to 256-QAM, for less than 2 dB degradation, this value drops to less than .5° rms, or –41 dBc. It should be about half of that to be considered negligible (< .5 dB). Figure 9 shows this "negligible" condition.

Figure 10 shows how this same amount of phase noise - the threshold, say, for 256-QAM - appears for 1024-QAM. For 1024-QAM, the sensitivity is related via the same relationship to 256-QAM that 256-QAM is to 64-QAM. This translates to the degradation condition of 1-2 dB reducing tolerable phase noise further yet to about .25° rms. This corresponds to a total untracked phase noise power of -47 dBc, or, equivalently, a signal-to-phase noise ratio of 47 dB.

The task of determining the ability of the error correction to handle phase noise-related errors is not particularly straightforward. And, certainly, it is not desirable to use up the FEC "budget" cleaning up channel phase noise problems. However, the key information needed to coarsely evaluate the situation is as follows. First it must be recognized that phase noise-related errors represent a burst-error type of noise mechanism. That is, the phase noise process is slow relative to the symbol rate, such that a phase referencing error will exist during the demodulation process for many symbols consecutively. This taxes the FEC more extensively than random errors would. However, the FEC employed does include both interleaving and Reed-Solomon (RS) coding to provide burst protection, although it is not in place particularly to mitigate phase noise-related issues. The parameters of interest for phase noise degradation are the rate of the phase noise process, which is related to the selected loop bandwidth, and how it compares to the symbol rate, interleaver depth, and burst error capability of the RS code. The ratio of symbol rate to loop bandwidth gives a feel for the number of consecutive symbols that can be in error. The number of consecutive symbol errors can be compared to the depth of the interleaver, which randomizes the symbol sequence, and the particular implementation of the RS code.

For tracking bandwidths of 50 kHz, and symbol rates of 5 Msps, the symbol rate/phase noise rate ratio is 100, which implies, as is often the case with phase noise – many symbols! Interleaver depth has no technology limitations. Instead, it effects latency and required storage (memory). Reed-Solomon codes come in many flavors, and can be designed per the needs of the application, but also within the same practical constraints that interleavers have.

U/C & tuner

It is significant to note that, as in the case of AWGN, there are very reasonable scenarios that exist today that yield levels of phase noise that severely challenge the required untracked error for 1024-QAM. As in the case of settop NF, the issues at hand are not technology limited – very low frequency synthesis is a classic design problem in the RF world, with multiple possible solutions that can be offered. The point is that, for the most part, existing hardware will have to be handled carefully in order to accommodate the requirements. Figure 11 is a case in point on the tuner side. This phase noise curve shows performance at the high end of the tuner range, and this would exceed the rms requirment described above by about a factor of three.

Synthesizer design is typically such that the phase noise signature of the upconverter or tuner will vary as the frequency is changed across the band, and is not always worse at the highest frequencies. As a result, as with distortion, it may be beneficial to consider choice of frequency for the most sensitive QAM signals that take into account the regions of converter performance that minimize the untracked phase noise. Murphy being ever-present, the likelihood is that the best location for phase noise and the best location for distortion effects are negatively correlated.

Other Issues

Forward Path Multiplex

The CTB issue for 256-QAM and 1024-QAM imply some obvious questions. If these schemes are particularly sensitive to narrowband interference, what can be done - other than through added receiver complexity – to mitigate it? Can it be avoided altogether? For the latter, the answer is no, it cannot be avoided altogether. There is little guardband between digital channels, and there are beats spread across the 6 MHz of an NTSC system such that it would be impossible not to be dealing with some of them. However, it is certainly the case that little effort has been made to center digital carriers such that they are affected by distortion beats the least. With CTB being the most aggravating, it is reasonable to suggest that carriers be located such that the worst components of distortion occupy the band edges as much as possible, where the effect is felt the least.

Assuming that it is an NTSC spectrum and digital channels are centered per this plan, the question then becomes where is it best to locate the center frequencies for the signals most sensitive to distortion. The answer to this question is a problem that has been solved many times over using analog beat mapping programs. It can now also be extended to analysis with digital channels. An example is shown in Figure 12 for a particular tilt and channel line-up. Any line-up, including both analog and digital, can be analyzed in this fashion, and the anticipated worst-case distortion frequencies identified. More importantly, the anticipated best-case distortion frequencies can also be identified. In this case, it is clear that the frequency band just above 800 MHz – where the CTB composite takes a dip - is an area where the most sensitive QAM carriers want to be placed relative to interference performance. Note that this plot is normalized, and does not account for the frequency dependence of the distortion, which can also be characterized in a tool like this, and is technology dependent. Moving forward, this type of thought process to carrier allocation may become more useful, rather than just loading the channels from the low end up as the services become available.

Group Delay

The job of the receiver/demodulator is to equalize amplitude response problems and non-flat group delay. When the variation is fixed or slowly varying – such as for these types of channel frequency response issues - many equalizers are well-suited to these tasks. The implementation of decision feedback equalizers strengthens the capability of the receiver against more troublesome impairments, such as the narrowband interference previously discussed.

Nonetheless, it is valuable from the standpoint of understanding high-level QAM to glance at the issue of channel distortions and intersymbol interference (ISI). There have been many studies on the sensitivity of simple modulations, such as QPSK, in an ISI environment. There is less work on QAM, although there are results available for reference. However, it is not difficult to offer some generalizations about how much more sensitive to ISI QAM will become as the number of symbol states increases. If desired, these relationship, in concert with some measured results on simpler modulation schemes, can be used to estimate what would happen to unequalized QAM, and thereby also being indicative of the additional "power" needed from the equalizer.

As in the case of phase noise, the performance degradation is going to be dominated by the impact of the impairment on the symbol states that suffer the most from it. In the case of phase noise, this was the outer symbol points. For the case of ISI, it is the case of small magnitude symbols that are adjacent to large magnitude symbols by the amount of delay associated with the ISI decay characteristics. More generally, it is the ratio of peak symbol energy to minimum symbol energy that coarsely estimates the relative difference in ISI sensitivity for different modulations. This is not difficult to determine, and the relationships are as follows:

Modulation	Peak-Min (dB)
16-QAM	9.5 dB
64-QAM	17 dB
256-QAM	23.5 dB
1024-QAM	30 dB

Thus, it can be estimated that ISI in a 256-QAM system can be about 6 dB lower in amplitude under equivalent dispersion delay (amplitude and delay cannot be separated, but useful thought exercise as a reference point) and yield the same effect on a 1024-QAM system. While it is less likely in the 1024-QAM case that a peak symbol will be positioned to effect the minimum (i.e. there are many more symbol, such that this particular sequence case is less probable), 1024-QAM also has more scenarios that are close to this worst-case possibility, because there are many symbol states close in magnitude to the maximum, as well as many close in magnitude to the minimum.

Again, however, it is reasonable to expect a well-designed equalizer to mitigate this distortion. In principle, the equalizer doesn't care about the amplitude values the symbol states may take on, but it can drive implementation issues, in the sense that numerical range and resolution becomes more significant for high order QAM in digitized receivers. That is, the equalizer must be more "powerful," both due to the additional modulator complexity, and the additional ISI sensitivity.

Conclusion

It isn't necessarily here or even right around the corner, but modem vendors have moved onto the next QAM challenge as MSO's have begun deploying the last one. While 64-QAM presented a relatively smooth roll-out, once the learning curve of having digital in the first place was overcome, the road to 256-QAM is dotted with some potholes that need some attention in order to get the desired bang for the buck. In the hardware deployment phase now, 256-QAM represents both a testament to excellence in modem design, as well as a wake-up call that this HFC forward path surely may be a powerful pipe, but it isn't without a few "gotchas" if it is to be squeezed for every last bit. Marrying HFC and 1024-QAM requires a very careful look at the experiences with 256-QAM. A careful evaluation will be taking place to see whether the benefits of the complexity, and its potential short term and long term plant impacts, make higher order modulation part of the necessary toolkit for capacity expansion of the forward path.

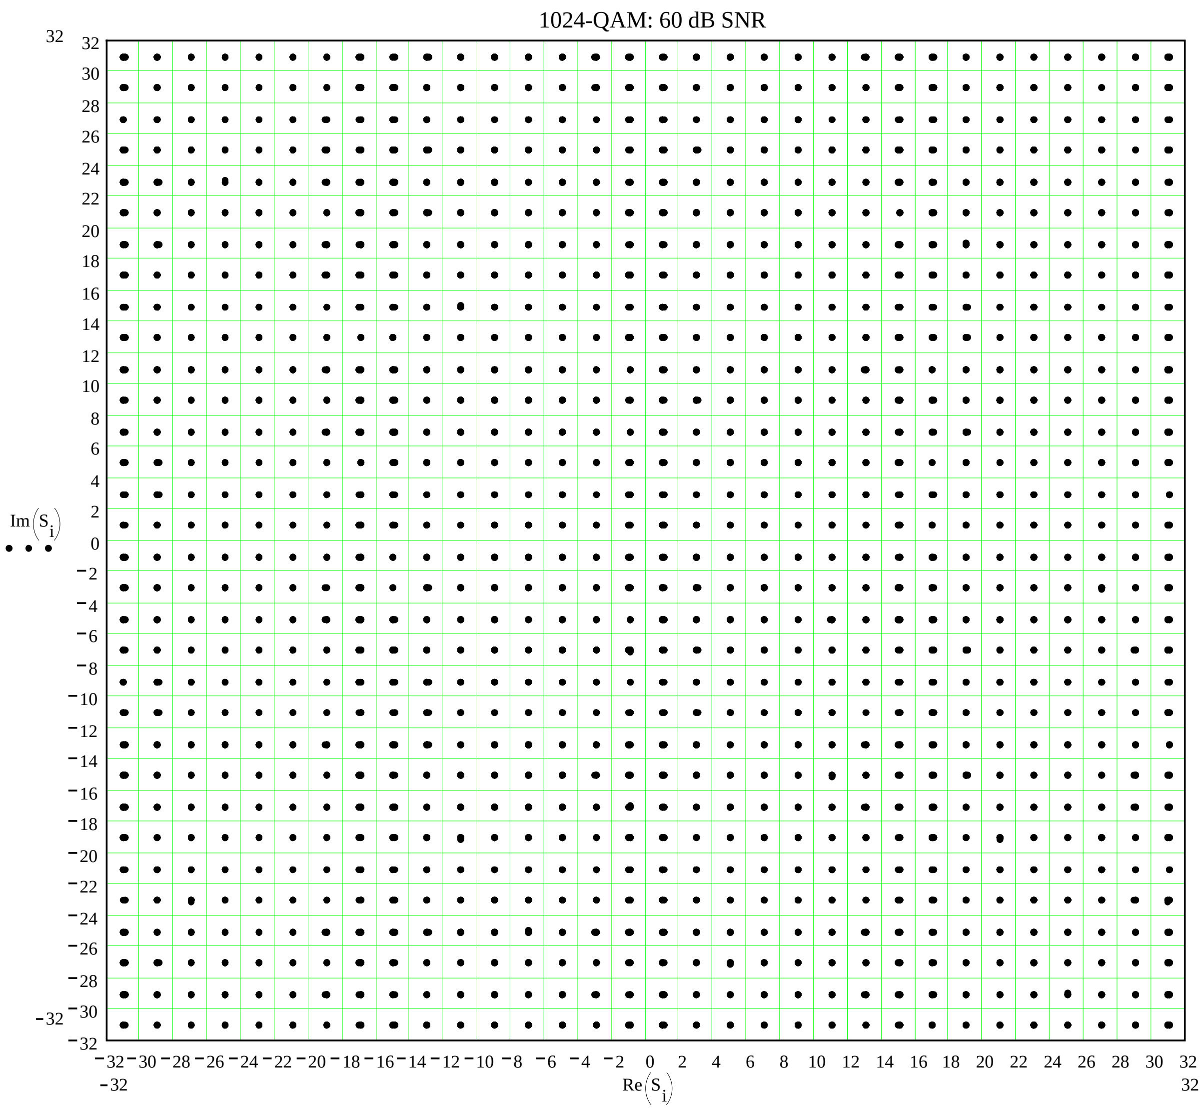

Figure 1 – Ideal 1024-QAM Constellation Diagram

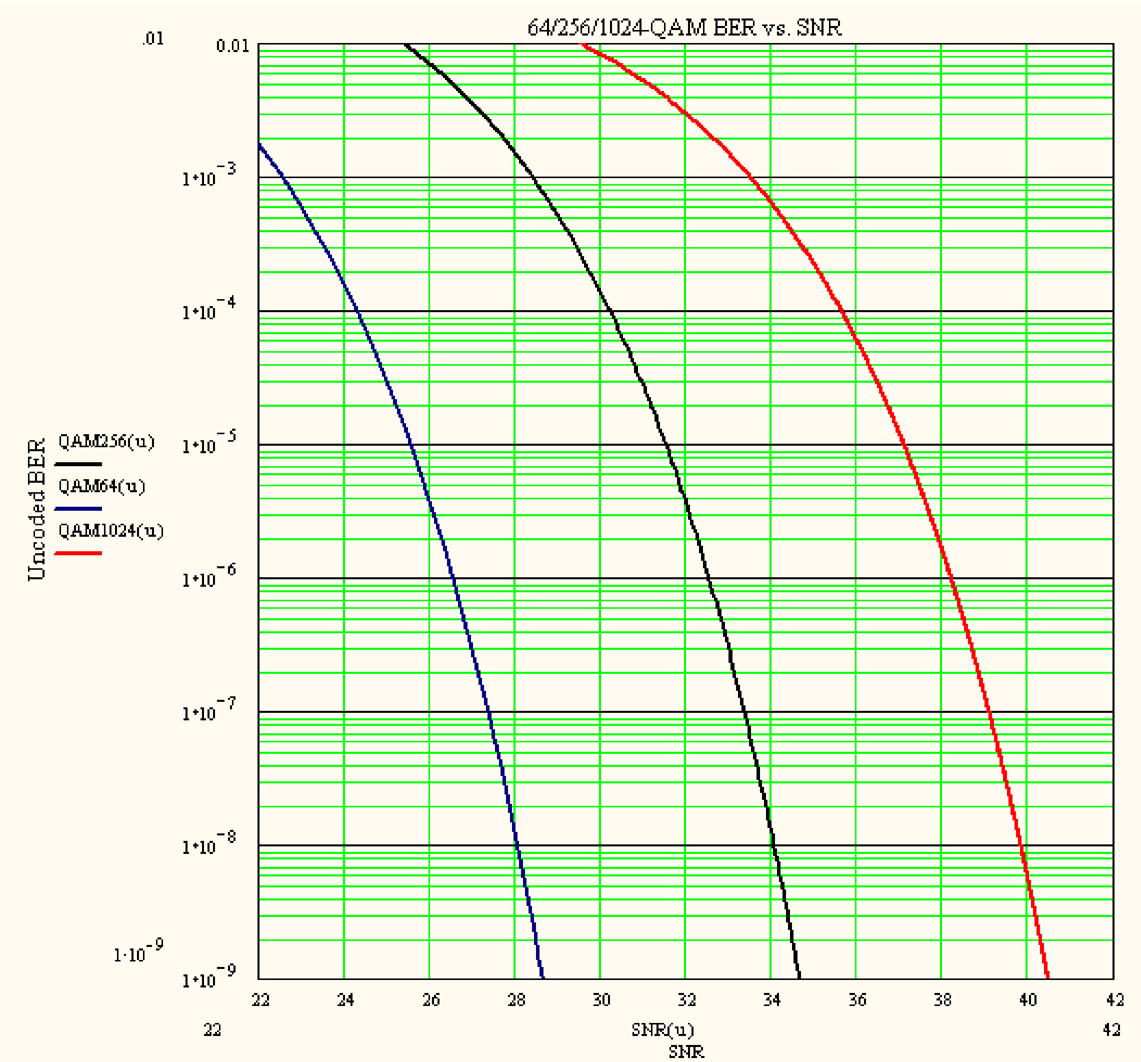

Figure 2 – Theoretical M-QAM BER for M= 64, 256, 1024

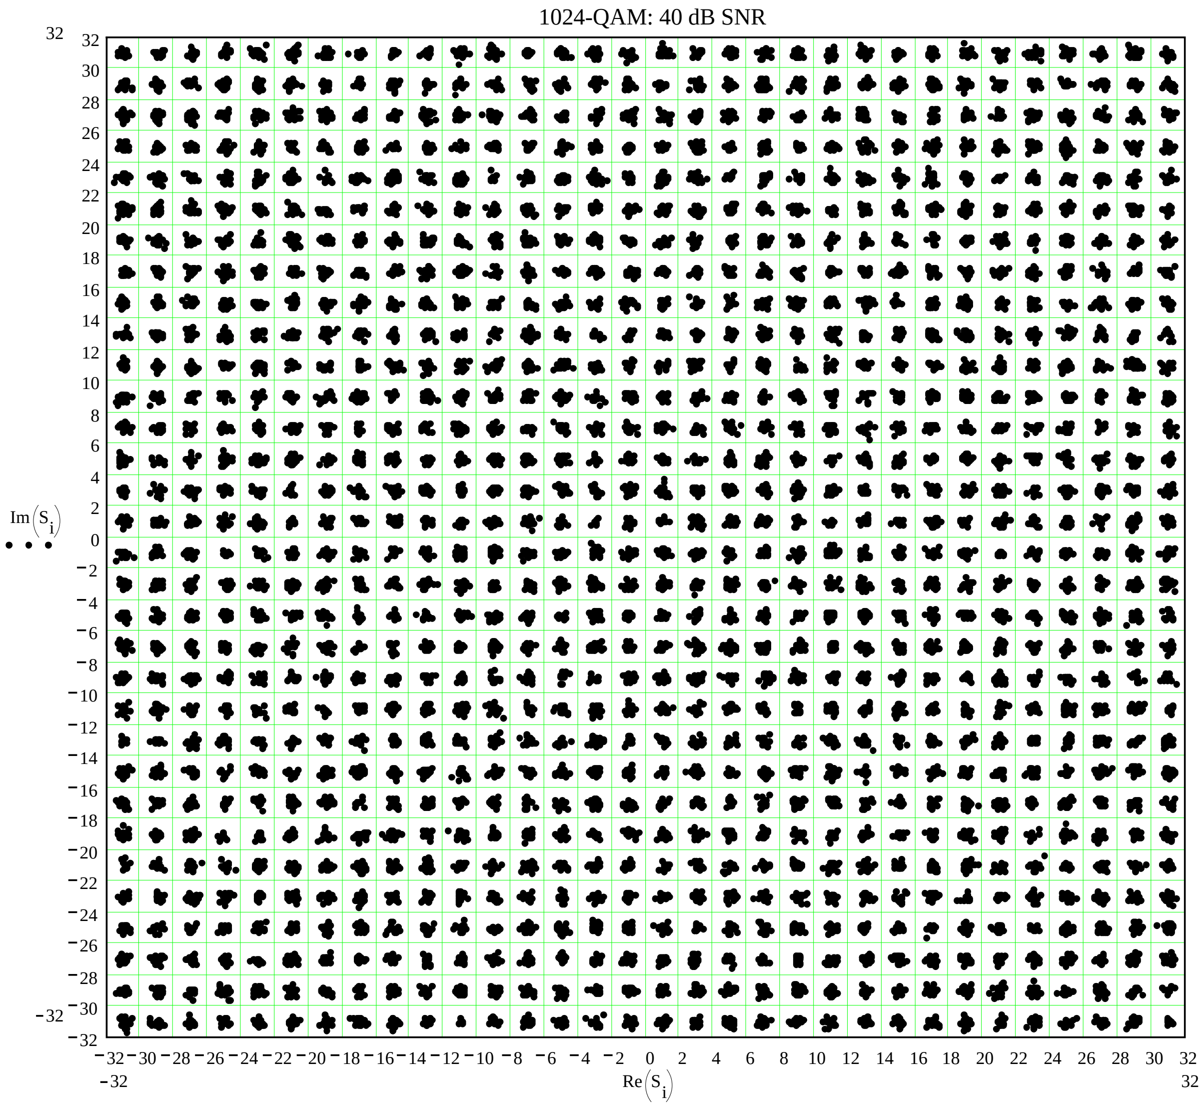

Figure 3 – 1024-QAM at SNR = 40 dB

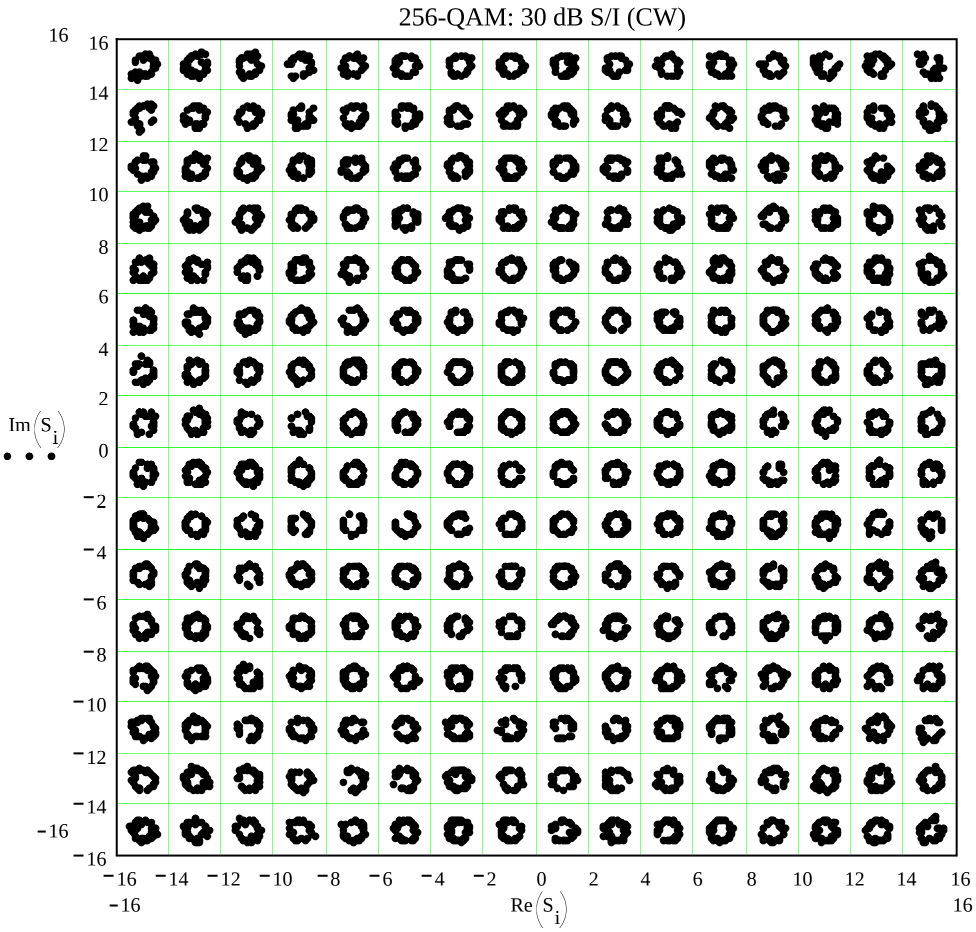

Figure 4 – High SNR 256-QAM with 30 dB S/I

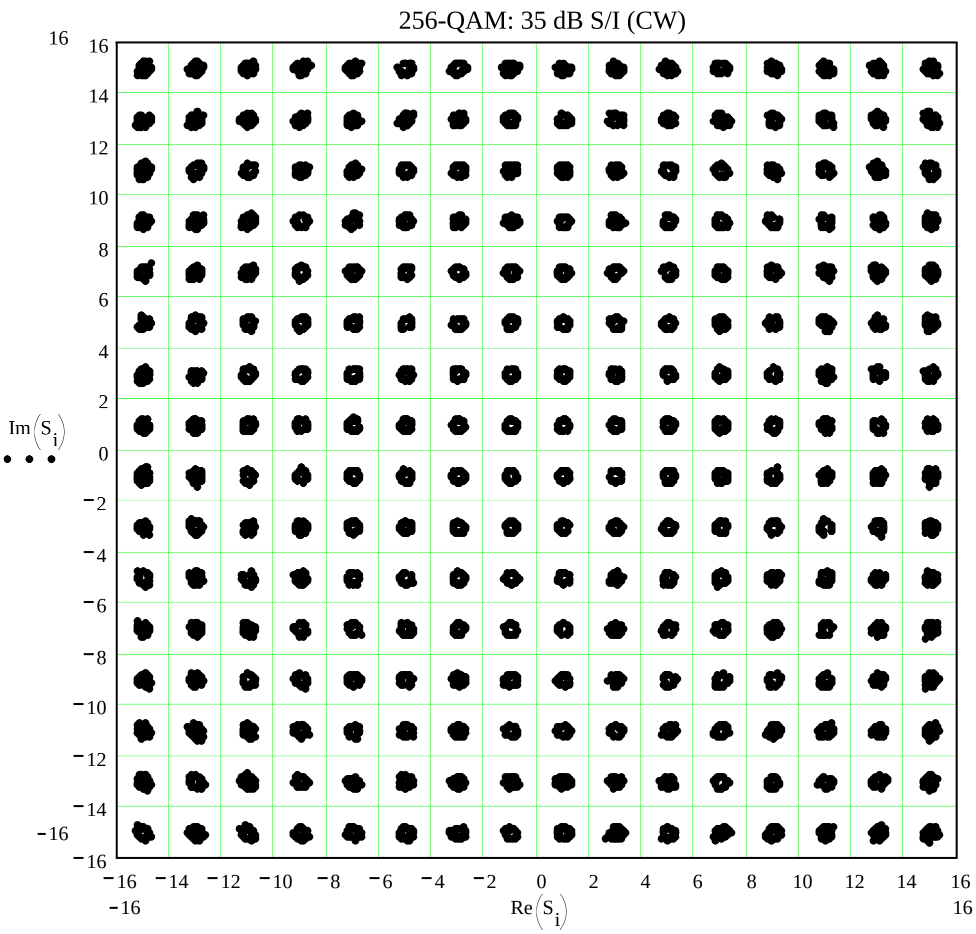

Figure 5 – High SNR 256-QAM with 35 dB S/I

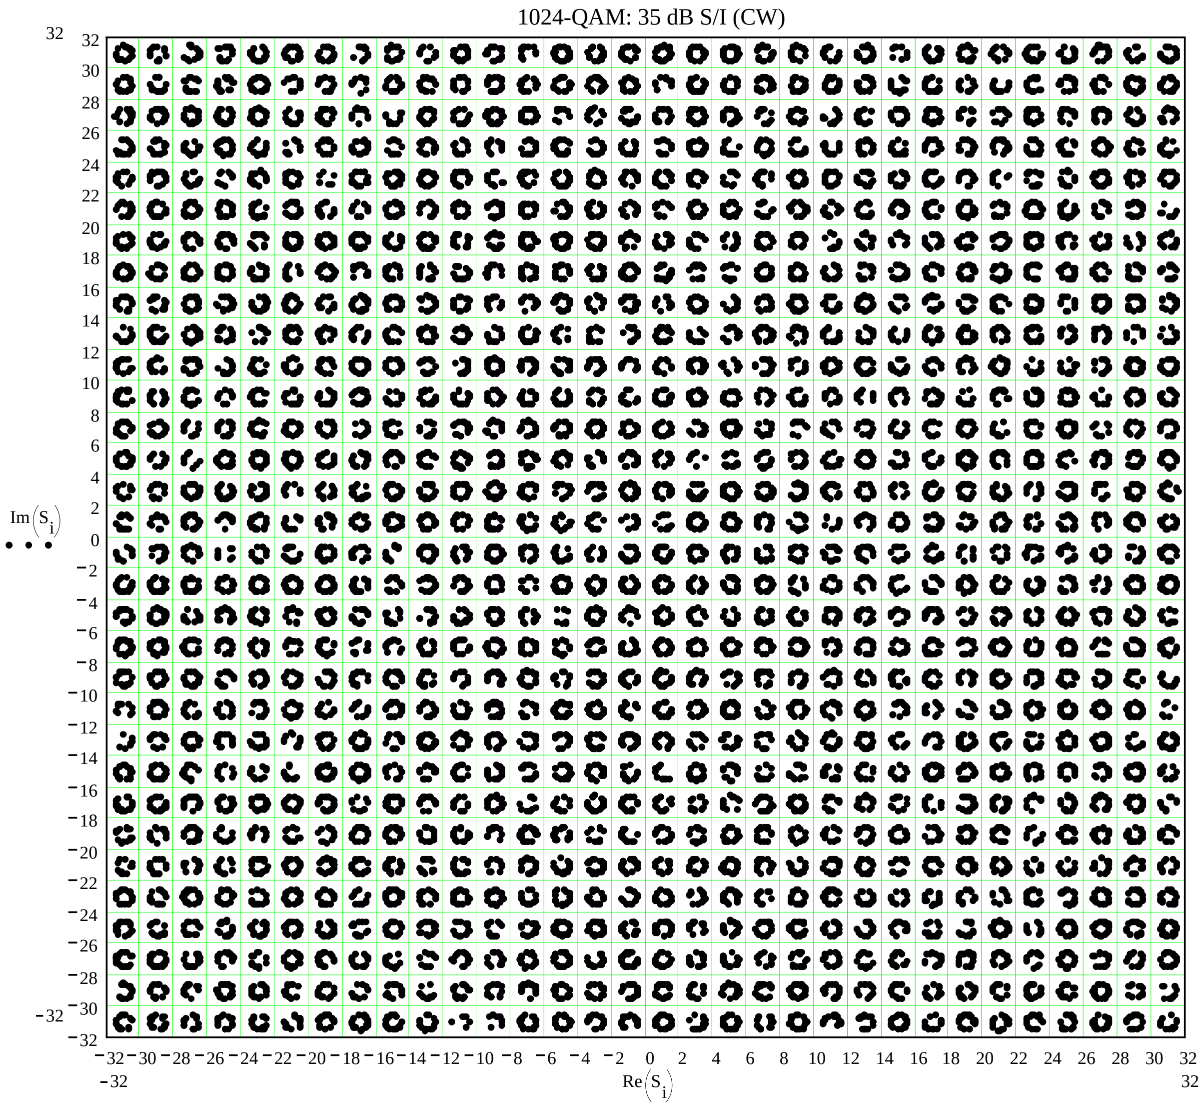

Figure 6 – High SNR 1024-QAM with 35 dB S/I

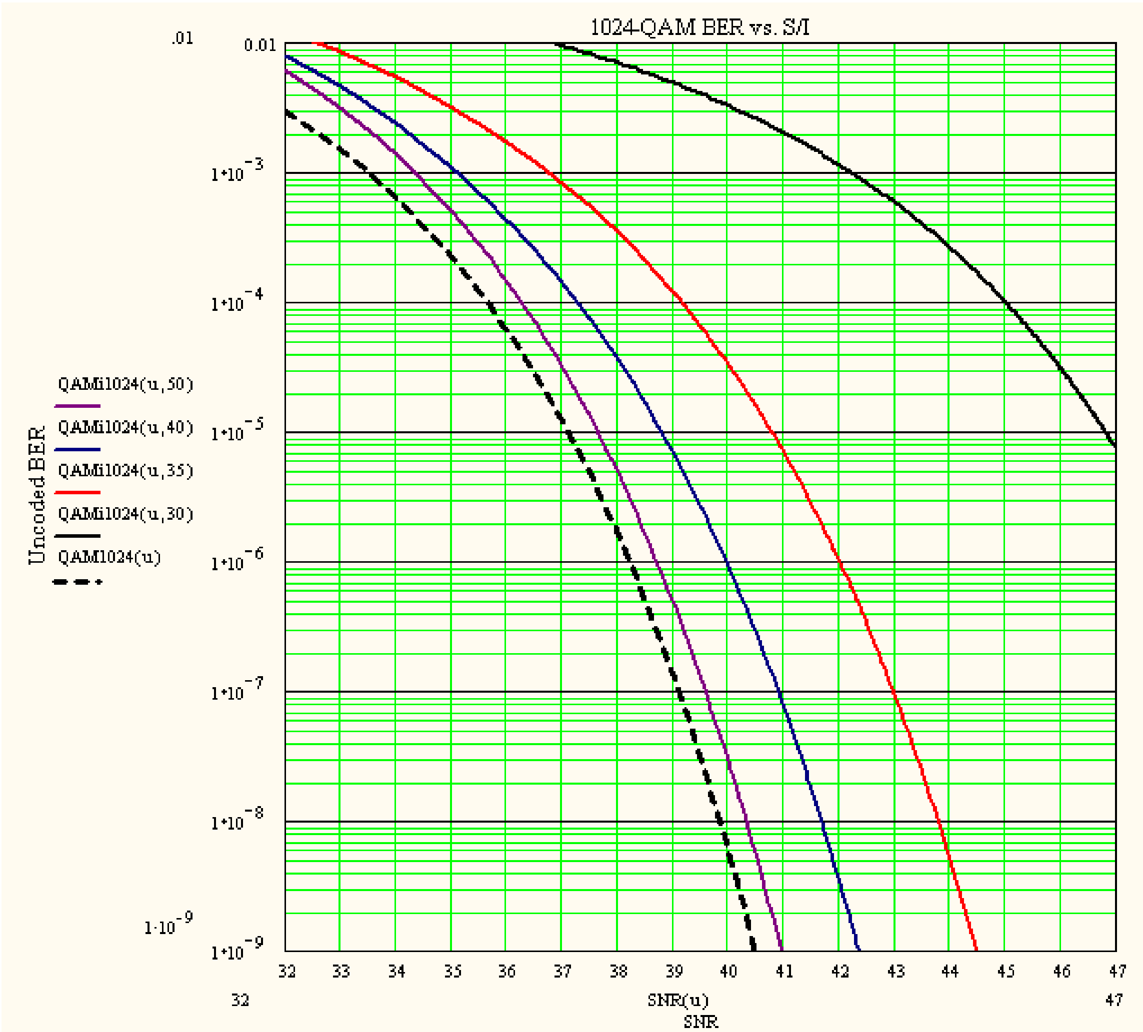

Figure 7 –1024-QAM BER versus S/I (dB)

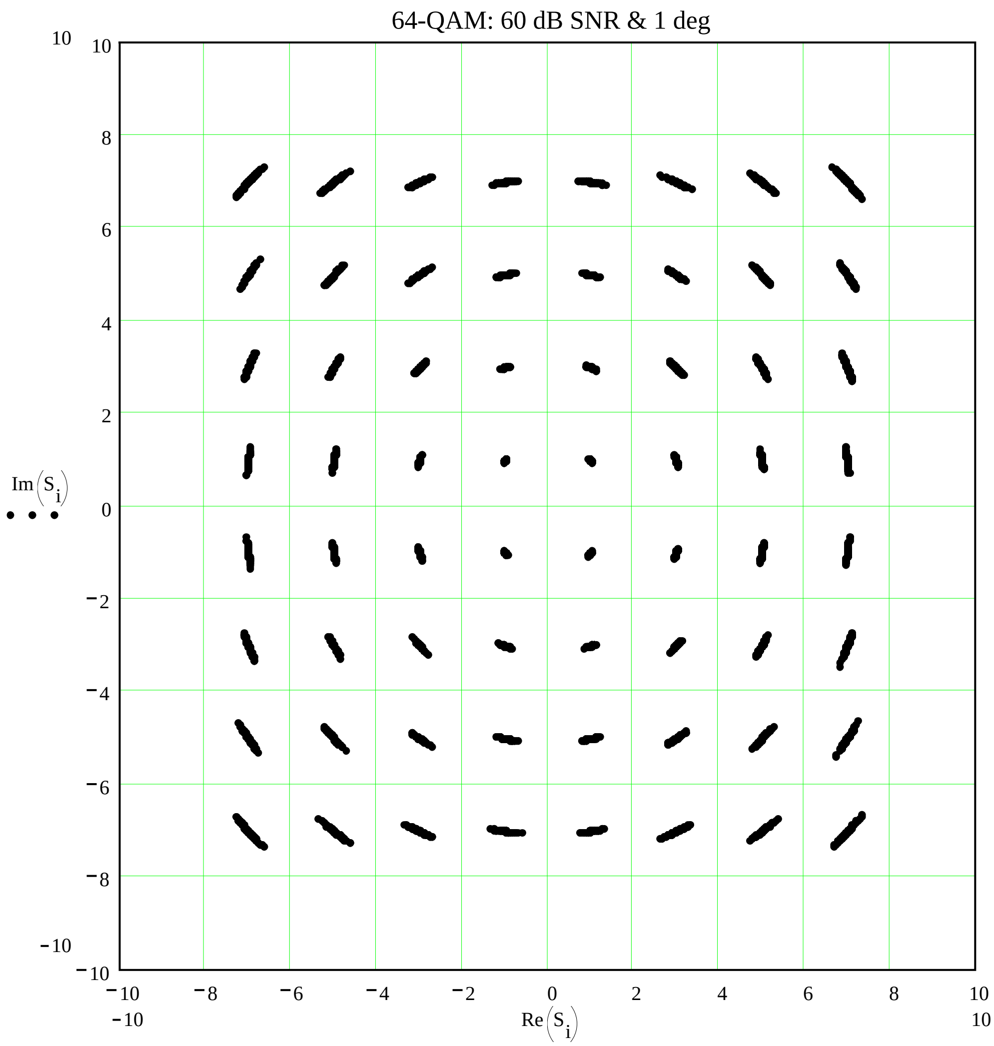

Figure 8 – 64-QAM with 1 deg rms Gaussian Phase Noise

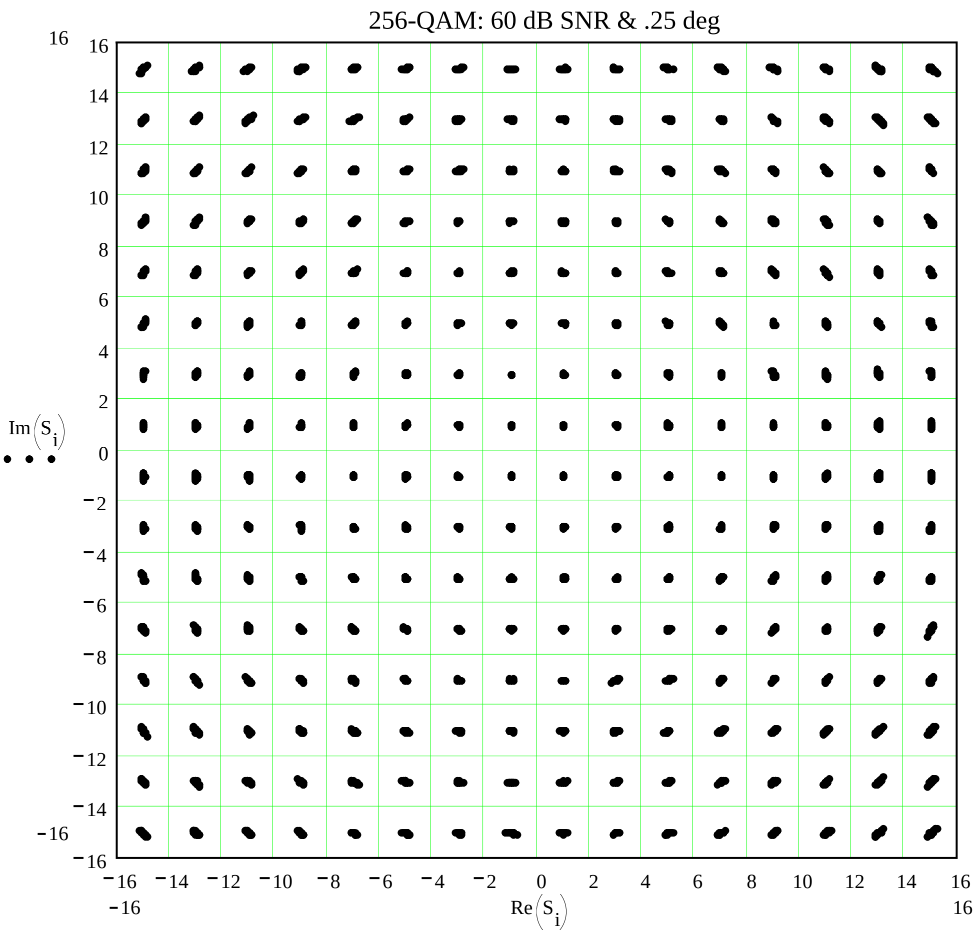

Figure 9 – 256-QAM with .25 deg rms Gaussian Phase Noise

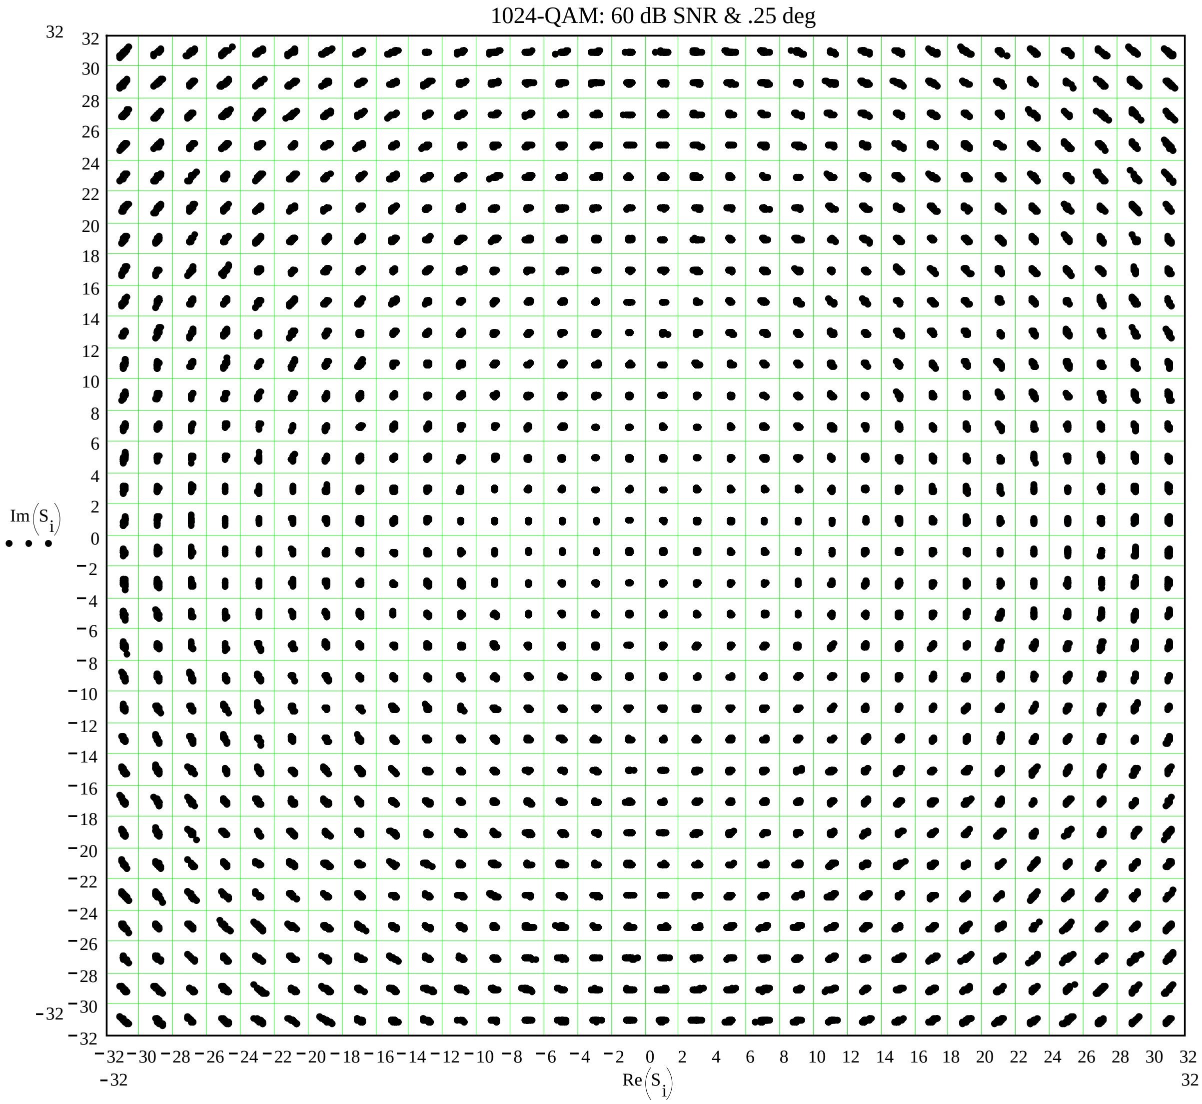

Figure 10 – 1024-QAM with .25 deg rms Gaussian Phase Noise

INTEGRATED NOISE LIST

START FREQ	STOP FREQ	INTEG NOISE	INTEG NOISE-SPURS
10.0 Hz	100.0 Hz	-10.0	-10.2
100.0 Hz	1000.0 Hz	-5.5	-8.5
1000.0 Hz	10000.0 Hz	-31.2	-31.4
10000.0 Hz	100000.0 Hz	-43.1	-43.3
100000.0 Hz	1000000.0 Hz	-48.0	-48.3

Figure 11 – RF Carrier Phase Noise Imposed by Tuner at 855 MHz

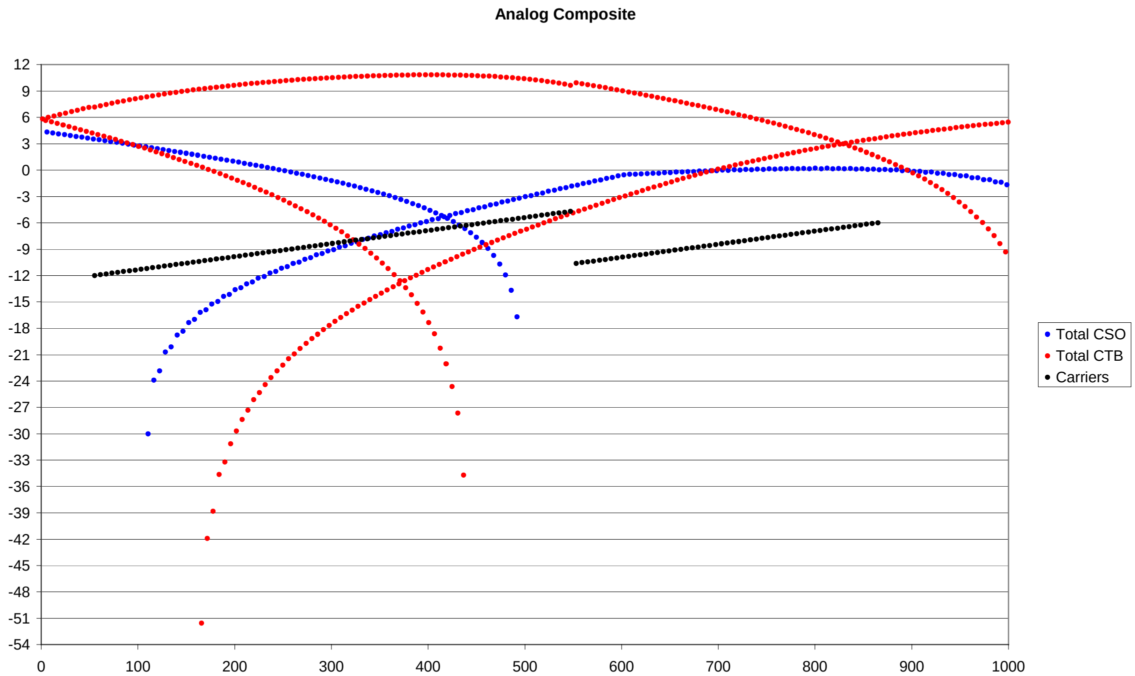

Distortion Map (analog beats) 550 MHz Analog + 320 MHz Digital System @ 12 dB Tilt

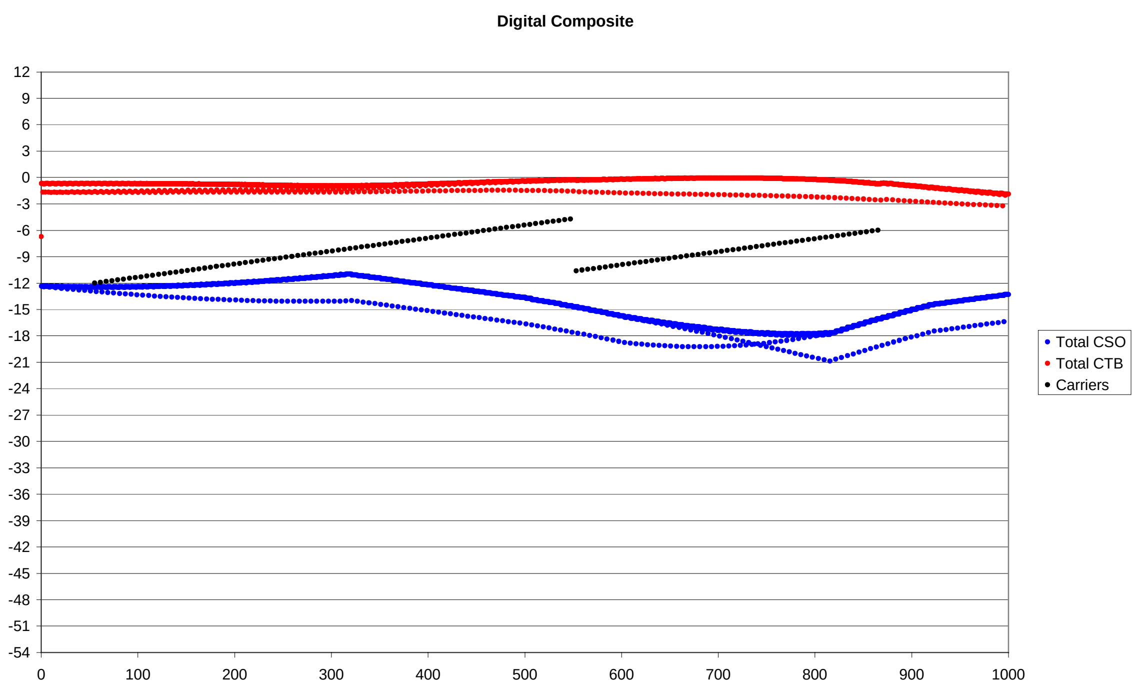

Distortion Map (digital CIN2 & CIN3); 550 MHz Analog + 320 MHz Digital System @ 12 dB Tilt

Figure 12 – Distortion Mapping for Optimal Forward Loading

References

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

[2] Katznelson, R., *Delivering on the 256-QAM Promise: Adopting Aggregate Degradation Measures for QAM Transmitters*, SCTE Cable-Tec EXPO, San Antonio, Texas, June 5-7, 2002.

[3] Stoneback, D., Howald, R., Brophy, T. and Sniezko, O., *Distortion Beat Characterization and the Impact on QAM BER Performance*, 48th Annual NCTA Convention, Chicago, Illinois, June 13-16, 1999.