Welcome to E-XFL.COM

Embedded - System On Chip (SoC): The Heart of Modern Embedded Systems

Embedded - System On Chip (SoC) refers to an integrated circuit that consolidates all the essential components of a computer system into a single chip. This includes a microprocessor, memory, and other peripherals, all packed into one compact and efficient package. SoCs are designed to provide a complete computing solution, optimizing both space and power consumption, making them ideal for a wide range of embedded applications.

What are Embedded - System On Chip (SoC)?

System On Chip (SoC) integrates multiple functions of a computer or electronic system onto a single chip. Unlike traditional multi-chip solutions, SoCs combine a central

Details	
Product Status	Obsolete
Architecture	-
Core Processor	-
Flash Size	-
RAM Size	-
Peripherals	-
Connectivity	-
Speed	-
Primary Attributes	-
Operating Temperature	-
Package / Case	-
Supplier Device Package	-
Purchase URL	https://www.e-xfl.com/product-detail/broadcom/bcm33843mvkfsbg

Introduction

Cable modem termination systems (CMTSs) can report a variety of operating parameters to the end user. For example, CMTSs that use Broadcom® BCM3137, BCM3138, or BCM3140 or Texas Instruments® TNETC4522 series upstream burst receivers can provide an “upstream SNR” estimate. This function is a very useful tool, but it has resulted in much confusion. It is not unusual for a cable company’s network operations center (NOC) staff to report an alarm condition when the reported upstream signal-to-noise ratio (SNR) of the CMTS drops below a defined threshold. A headend technician follows up by checking the upstream RF performance with a spectrum analyzer or similar test equipment, only to find that everything appears normal.

Data personnel in the NOC insist there must be a problem, while outside plant technicians see nothing amiss on their test equipment. What is going on here?

The discrepancy occurs from a lack of understanding about just what the CMTS upstream SNR estimate is—and what it is not.

Further confusion comes from the fact that cable modems and digital set-top boxes (STBs) can provide digitally modulated signal operating parameters such as RF signal level and SNR. These are downstream parameters at the customer premises, not upstream parameters as is sometimes incorrectly assumed. In addition, test equipment used by cable operators to characterize digitally modulated signals can measure downstream—and in some cases upstream—modulation error ratio (MER). Some of these instruments call this parameter SNR.

Also, because of the time-shared nature of the upstream, most of today’s CMTSs can measure parameters on a per-*channel* basis or a per-*cable-modem* basis. Per-channel measurements provide an average of all cable modems or simply a snapshot of the most recently active cable modem(s). It is important to distinguish which type of measurement is presented.

This paper provides a background on several signal quality metrics applicable to CMTS and cable network operation and how they relate to overall performance. The CMTS upstream SNR and cable modem or STB downstream SNR estimates are explained. Noise, as discussed in this paper and unless defined otherwise, refers to *additive white Gaussian noise* (AWGN), or simply *white noise*, also known as thermal noise. Interference such as narrowband ingress and burst or impulse noise is usually treated separately.

First, what the CMTS upstream SNR estimate is not: the SNR estimate from a cable modem or CMTS is *not* the same thing as the carrier-to-noise ratio (CNR) that one measures with a spectrum analyzer.

Here is what the upstream SNR estimate is: an operating parameter provided by the upstream burst receiver used in DOCSIS® CMTSs. Similar information for downstream signals is provided by the quadrature amplitude modulation (QAM) receiver in a cable modem or STB. The SNR estimate, which is derived after the data is demodulated, is more accurately called receive modulation error ratio (RxMER), a term recently defined in the DOCSIS MIB. RxMER includes the effects of the cable network downstream or upstream noise floor, in-channel frequency response (including amplitude tilt and ripple, group delay variation, and micro-reflections), oscillator phase noise, receiver imperfections, and all other impairments that affect the receive symbol constellation. Because it measures the end-to-end performance of the communications link, RxMER is useful for tracking long-term system performance trends.

Interestingly, it is not unusual to have a reported low downstream or upstream RxMER number, yet find the measured CNR and signal levels to be just fine. Why? Because one or more impairments that cannot be seen on a spectrum analyzer—poor in-channel frequency response, including group delay variation and micro-reflections, and even upstream data collisions—may be the cause of the low reported RxMER.

This paper discusses the terms listed in Table 1.

Table 1: Terminology for Various SNR Ratio Concepts

Abbreviation	Term	Definition	Units
CNR (or C/N)	Carrier-to-noise ratio	The ratio of carrier or signal power to the white-noise power in a specified bandwidth, as measured on an RF spectrum analyzer or similar equipment.	dB
C/N_0	Carrier-to-noise - density ratio	The ratio of carrier or signal power to white-noise spectral density.	dB-Hz
CNIR (or C/(N+I))	Carrier-to-noise-plus-interference ratio	The ratio of carrier or signal power to the total noise power (including white noise and interference) in a specified bandwidth, as measured on an RF spectrum analyzer or similar equipment.	dB
E_S/N_0	Energy-per-symbol to noise-density ratio	In digital modulation, the ratio of the average energy of a QAM symbol to the white-noise spectral density.	dB
EVM	Error vector magnitude	The ratio of RMS constellation error magnitude to peak constellation symbol magnitude.	Percent
MER	Modulation error ratio	The ratio of average signal constellation power to average constellation error power.	dB
RxMER	Receive modulation error ratio	The MER as measured in a digital receiver after demodulation, with or without adaptive equalization.	dB
SNR (or S/N)	Signal-to-noise ratio	(a) A general measurement of the ratio of signal power to noise power. (b) In a specific context, a measurement of the ratio of signal power to noise power made at baseband before modulation or after detection or demodulation.	dB
TxMER	Transmit modulation error ratio	The MER produced by a transmitter under test, as measured by an ideal test receiver.	dB

CNR and SNR from a Telecommunications Industry Perspective

Some of the confusion mentioned in the introduction arises from the fact that in the world of telecommunications outside the cable industry, the terms SNR and CNR are often used interchangeably. According to Roger L. Freeman's *Telecommunications Transmission Handbook*, "The signal-to-noise ratio expresses in decibels the amount by which a signal level exceeds its corresponding noise." Another reference, Tektronix's *Measuring Noise in Video Systems*, says "In the most general case, SNR is expressed as the ratio of RMS (root mean square) signal level, S_{RMS}, to the RMS noise, N_{RMS}, ($SNR = S_{RMS}/N_{RMS}$)." [Eq. 1]

Both of the previous SNR definitions can easily be applied to RF CNR measurements (after all, a carrier is a "signal") as well as baseband SNR measurements (baseband video and audio are "signals," too). If the specific measurement is not clearly defined, it is difficult to know whether SNR refers to a baseband or RF parameter. This paper distinguishes between SNR and CNR. In the subsequent sections, each term is defined and explained, and the distinction is illustrated following usage in the cable industry.

CNR and SNR from a Cable Industry Perspective

Modern Cable Television Technology, 2nd Ed., states, "Carrier-to-noise ratio (*C/N*) is defined as follows:

$$C/N(\text{dB}) \equiv 10\log(c/n) \quad \text{[Eq. 2]}$$

where *c* and *n* are the scalar power levels of the carrier and noise, respectively."[1]

When measuring CNR on a spectrum analyzer with thermal noise underlying the carrier, one actually is measuring not C/N but, more precisely, (C + N)/N = 1 + C/N. This distinction is not normally a concern unless the CNR is very low—say, single-digit decibel (dB) values, as we will see later.[2]

The cable industry has long used CNR and SNR to represent quite different measurement parameters, one in the RF domain (Figure 1) and the other in the baseband domain (Figure 2 on page 5). CNR is applied to the transmitted over-the-cable RF waveform, whereas SNR refers to the video and audio signal prior to modulation for broadcast, or after demodulation of the RF waveform at the receiver.

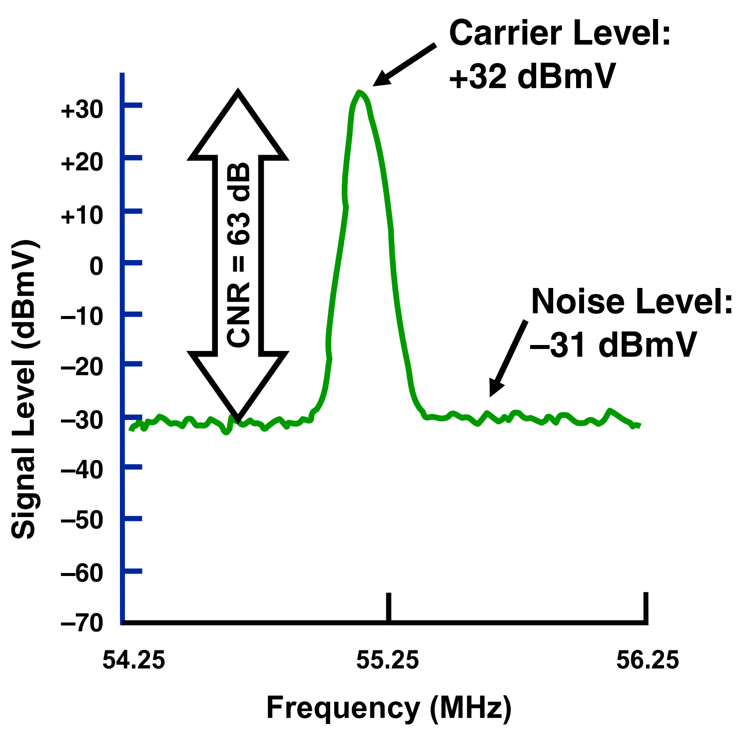

Figure 1: RF CNR Measurement (As Detected in the Test Equipment Resolution Bandwidth)

1. In this paper, all logarithms are base 10.
2. The expression 1 + C/N uses power quantities, not dB; that is, we are not adding 1 dB to C/N.

Analog Video SNR in Cable Networks

What about SNR? SNR is, in cable industry vernacular, a premodulation (at the transmitter or modulator) or postdetection (at the receiver) measurement—one made on a baseband signal such as video or audio. The previously mentioned Tektronix application note says: "In video applications, however, it is the effective power of the noise relative to the nominal luminance level that is the greater concern." It goes on to define video SNR in decibels as:

$$SNR = 20\log(L_{NOMINAL}/N_{RMS}) \qquad \text{[Eq. 3]}$$

where $L_{NOMINAL}$ has a value of 714 millivolts peak-to-peak (100 IRE units) for NTSC or 700 mV peak to peak for PAL. These luminance values do not include sync.

Equation 3 simply states that baseband video SNR is the ratio of the peak-to-peak video signal, excluding sync, to the noise within that video signal. The noise is measured in a bandwidth defined by a combination of low-pass, high-pass, and weighting filters. These filters limit the measured noise to a bandwidth that is roughly the same as the video signal and may be used to remove certain low-frequency noise from the measurement. Weighting filters are used to simulate the eye's response to noise in the TV picture. Various standards such as RS-170A, RS-250B, and NTC-7 specify the characteristics of filters that are used in baseband video SNR measurements.

To recap: CNR is a predetection measurement performed on RF signals. It is the difference, in decibels, between carrier power and noise power in the RF transport path only—for instance, a coaxial cable distribution network or a stand-alone device such as an upconverter or headend signal processor. As such, CNR is ideal for characterizing network or individual device impairments. SNR, when applied to analog video or audio signals, is a premodulation or postdetection measurement performed at baseband. It is equal to the ratio of the peak-to-peak baseband signal to the noise within that signal (refer to Figure 1 on page 4). SNR includes noise in the original signal—say, noise in the video from a TV studio camera—as well as noise contributions from the transmitter or modulator, transport path, receiver, and demodulator. It is ideal for characterizing end-to-end performance—the overall picture quality seen by the end user, in the case of baseband video SNR.

Discrete Versus Modulated Signals and Carrier-to-Noise Density Ratio

A measurement that is closely related to CNR is carrier-to-noise-density ratio (C/N_0), defined as the ratio of carrier or signal power (in watts) to the underlying white-noise power spectral density (in watts/Hz). Noise power spectral density N_0 is the noise power in a 1 Hz bandwidth—that is, watts per Hz. Because of the impracticality of making a 1 Hz bandwidth noise power measurement, noise power spectral density is usually measured in a larger, more convenient bandwidth—the test equipment resolution bandwidth (RBW)—or, to be more precise, the equivalent noise bandwidth of the RBW filter. The measured value in watts is then divided by the test equipment resolution bandwidth in Hz, which yields the power (in watts) in a 1 Hz bandwidth. If the noise power measurement is in dBmV, subtract 10log(RBW in Hz) from the measured value to get the 1 Hz bandwidth equivalent, also in dBmV.

Taking the ratio of units shows that C/N_0 has units of Hz:

$$\frac{C\ (\text{watts})}{N_0\ (\text{watts/Hz})} = C/N_0(Hz) \qquad \text{[Eq. 4]}$$

Is a Digitally Modulated Signal a "Carrier"?

Quadrature amplitude modulation results in a double-sideband, suppressed-carrier RF signal. Before modulation is applied, the unmodulated signal is certainly a carrier—a CW carrier. But when modulation is applied, the carrier is suppressed. From one perspective, it could be correctly argued that the "haystack" of a QAM signal is not a carrier—it is just an RF signal—because the carrier is suppressed. From another perspective, it is legitimate to argue that a QAM signal "carries" information because its modulation was impressed on a carrier frequency. The underlying carrier can be recovered and tracked by a digital receiver. In fact, a QAM signal is *cyclostationary*—its carrier can be regenerated and observed on a spectrum analyzer by first passing the QAM signal through a nonlinearity such as 4th power. To minimize confusion with baseband signals and to follow common industry parlance, this paper refers to digitally modulated RF signals as "modulated carriers," or simply "carriers."

To measure the white noise—the "N" in CNR—underlying a digitally modulated signal, a noise power bandwidth (also called equivalent noise bandwidth or effective bandwidth) equal to the symbol rate should be used. The noise measurement should be performed when the signal is not present, or in an empty band near the signal. Table 2 and Table 3 summarize noise power bandwidth values for DOCSIS downstream and upstream channels.

Table 2: Noise Power Bandwidth for DOCSIS Downstream Channels

Channel RF (Occupied) Bandwidth	Symbol Rate R_s	Noise Power Bandwidth (approximate)
6 MHz	5.056941 Msym/sec	5.06 MHz
6 MHz	5.360537 Msym/sec	5.36 MHz
8 MHz (Euro-DOCSIS)	6.952 Msym/sec	6.95 MHz

Table 3: Noise Power Bandwidth for DOCSIS Upstream Channels

Channel RF (Occupied) Bandwidth	Symbol Rate R_s [note a]	Noise Power Bandwidth
200 kHz	160 ksym/sec	160 kHz
400 kHz	320 ksym/sec	320 kHz
800 kHz	640 ksym/sec	640 kHz
1.6 MHz	1.280 Msym/sec	1.28 MHz
3.2 MHz	2.560 Msym/sec	2.56 MHz
6.4 MHz	5.120 Msym/sec	5.12 MHz

a. In the DOCSIS 2.0 Physical Media Dependent Sublayer Specification, the term "modulation rate in kHz" is used instead of symbol rate in kilosymbols per second. This definition encompasses both TDMA and S-CDMA transmission.

QAM Spectrum Basics

Consider what a QAM signal spectrum looks like, and how we can read the signal power, noise power, and CNR from the spectrum display.

First, some background on transmit and receive filters. Figure 3 shows a communications system, which represents an upstream or downstream cable data network. Both the transmitter and receiver contain "matched filters." The purpose of the transmit matched filter is to band-limit the transmitted spectrum so that it will not interfere with adjacent channels. The purpose of the receive matched filter is to select the desired channel and to reject noise. The "matched" property of the filters implies that they have identical frequency magnitude response $|H_1(f)|$. [Eq. 10]

A time-invariant system such as a filter can be characterized by its impulse response $h(t)$ or by its frequency response $H(f)$, which comprise a Fourier transform pair. The asterisk on the receive filter in Figure 3 indicates that the receive matched filter exhibits complex conjugation of its frequency response, or equivalently, time reversal of its impulse response, relative to the transmit filter. However, this property is not of practical importance in cable systems because the filters are time-symmetric.

Matched filtering is known to maximize receive SNR in the presence of white noise. The cascade of the two matched filters gives the "full" magnitude response $|H(f)| = |H_1(f)||H_1(f)| = |H_1(f)|^2$. A signal having this full magnitude response is seen only at the output of the receive matched filter, inside of the receiver, and is not normally visible to an outside observer. However, the full-filter response is important in that it is designed to have the Nyquist property, described in the next section.

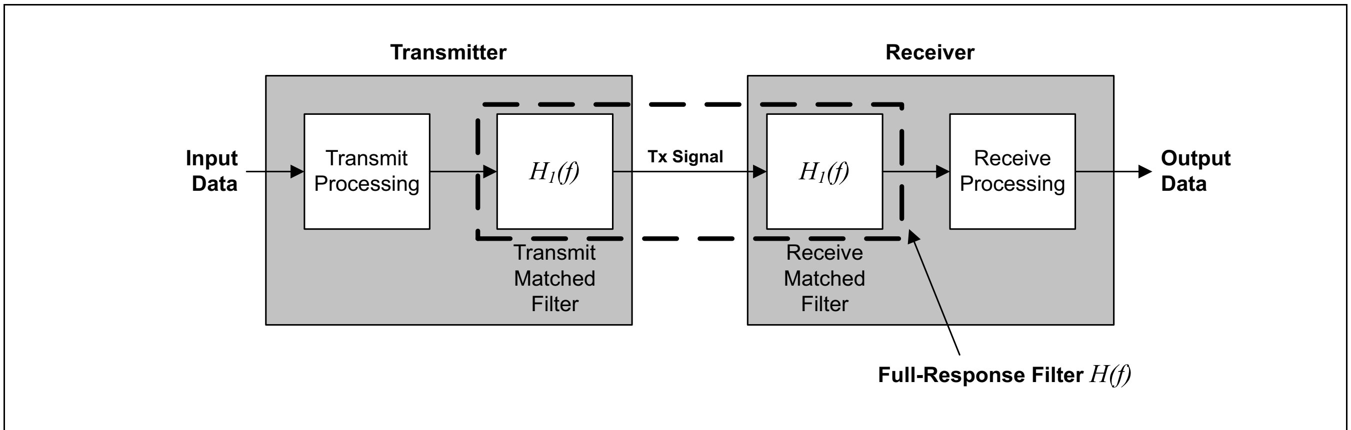

Figure 3: General Digital Communications System Showing Matched Filters in Transmitter and Receiver

The signal marked "Tx Signal" in Figure 3 represents the actual transmitted signal that is observable on the cable plant. (We are neglecting upconversion and downconversion to/from RF and are dealing only with "complex envelopes" in this discussion.) Because the scrambler guarantees that the QAM symbols affecting the transmit filter are white,[4] the transmitted spectrum is given its shape by the transmit filter, and its power spectrum, or power spectral density (PSD), is $|H_1(f)|^2$. Note an interesting "cascading property": The full magnitude response of the cascaded transmit and receive filters, and the PSD of the transmitted signal are both $|H_1(f)|^2$. In the former case, the squaring of the magnitude response comes

4. A "white" spectrum is one that is flat, which means it has a constant magnitude across all frequencies. This constant magnitude corresponds to statistically uncorrelated samples in the time domain.

Practical QAM CNR Measurements

Figure 6 shows an example of measuring downstream digitally modulated signal CNR and E_S/N_0 using a spectrum analyzer marker noise function, which measures noise power in a 1 Hz bandwidth. The carrier is centered on the screen, and the marker noise function is activated. Make a note of the indicated marker noise amplitude of the digitally modulated signal (–80.70 dBmV in the left display), and then tune the analyzer to a frequency that allows measurement of the noise floor of the cable network. The right display shows a noise-floor amplitude of –95.82 dBmV. The difference between the two measurements is the approximate CNR and E_S/N_0, in this case about 15 dB.

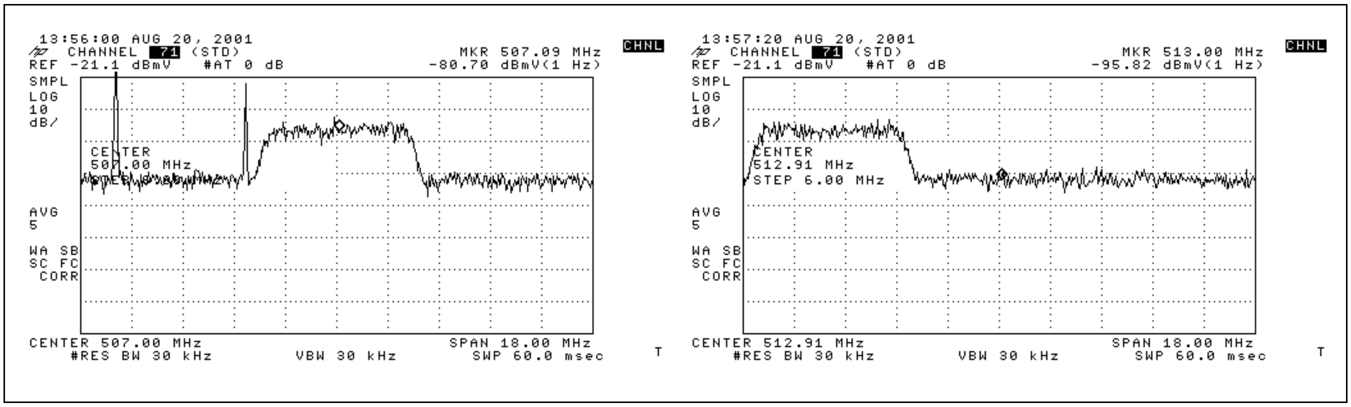

Figure 6: Digitally Modulated Signal CNR (This Example Shows an Approximate CNR of 15 dB.)

For best accuracy, noise should be measured using sample detection, which is typically employed by a spectrum analyzer marker noise function. Marker noise measurements usually encompass a range of frequency samples, starting from the marker and going plus or minus a small number of points out of the total available in the analyzer trace, so proper placement of the marker away from the carrier is important for accuracy. Refer to the analyzer documentation.

Signal and Noise Measurements on a Spectrum Analyzer

A few practical tips are worth noting when measuring modulated signals and noise on a spectrum analyzer.[6]

- Ensure that the spectrum analyzer uses sample detection when measuring noise or noise-like signals. If log/peak detection is used, a detector correction factor (typically 2.5 dB—consult the spectrum analyzer documentation) needs to be applied to the measurement result. Of course, if the signal and noise have the same statistics (Gaussian or nearly so), the same correction factor will be added to both the signal and noise measurements, and the correction will cancel out when measuring the CNR.

6. For a thorough treatment of this subject, refer to Agilent Application Note 1303, "Spectrum Analyzer Measurements and Noise: Measuring Noise and Noise-like Digital Communications Signals with a Spectrum Analyzer."

level for all modems at the upstream port, typically with less than 1 dB signal level difference among modems. Cable network noise funnels back to one point (the CMTS upstream port), so the noise amplitude at the CMTS input will be the same for all modems sharing the return spectrum of that port.

In some situations, upstream CNR may be different for certain modems, but this difference usually occurs because the RF signals of the affected modems reach the CMTS upstream port at levels lower than the set value—often an indication that some modems are transmitting at or near their maximum output level, yet still cannot reach the CMTS at the commanded ranging level because of excessive upstream attenuation. In the vast majority of cases, subscriber drop problems are the cause (splitter or coupler installed backward, corroded connectors, damaged cable, etc.), although occasionally, upstream amplifier misalignment or feeder problems may be at fault. Check your CMTS documentation for the command to determine whether any modems are transmitting at their maximum output level. DOCSIS requires the CMTS receiver to operate properly with up to ±6 dB receive level differences between bursts to account for ranging inaccuracies.

Upstream power levels for cable modem signals can be measured using a spectrum analyzer in zero span mode with the resolution bandwidth set approximately equal to the symbol rate. This method is useful for confirming received levels at the CMTS upstream input, as well as performing CNR measurements. Care should be exercised when using this method with adjacent channels present, because the RBW filter of the spectrum analyzer may not completely reject the adjacent channel energy.

Upstream Digitally Modulated Signal CNR—An Example

Assume that the digital channel power of a 3.2 MHz bandwidth 16-QAM signal at the CMTS upstream port is 0 dBmV. The noise floor is measured with the spectrum analyzer set to 100 kHz RBW, using the analyzer marker (assume sample detection for the noise measurement), and found to be –40 dBmV. At first glance, one might be inclined to think the CNR is 40 dB [0 dBmV – (–40 dBmV)], but the noise measurement must be corrected to the equivalent of one made in a bandwidth equal to the symbol rate of the digitally modulated signal. In this example, the noise power bandwidth of the measurement is only 100 kHz, equal to the spectrum analyzer RBW setting. (For greater accuracy, the RBW or noise bandwidth offset could be included, as discussed previously.) From Table 2, a 3.2 MHz bandwidth digitally modulated signal symbol rate is 2.56 Msym/sec, and its equivalent noise power bandwidth is 2.56 MHz. The correction factor that must be added to the –40 dBmV measurement result is found with the formula:

$$dB_{correction} = 10\log(NBW/RBW) \qquad \text{[Eq. 19]}$$

where NBW is the noise power bandwidth (symbol rate equivalent) for the digitally modulated signal, and RBW is the resolution bandwidth of the test equipment. Both values must be in the same units—for instance, Hz.

$$dB_{correction} = 10\log(2{,}560{,}000/100{,}000)$$
$$dB_{correction} = 14.08$$

To obtain the bandwidth-corrected noise amplitude, add the calculated correction factor to the original measured noise amplitude.

$$-40\text{ dBmV} + 14.08\text{ dB} = -25.92\text{ dBmV}$$

The CNR is 0 dBmV – (–25.92 dBmV) = 25.92 dB.

What Is Inside the Blocks in a Digital QAM Receiver?

Analog and digital front end: Analog and digital front-end components perform tuning, automatic gain control, channel selection, analog-to-digital conversion, and related functions. Their purpose is to preprocess the signal so that the individual QAM RF channels are available for further digital processing.

Matched filter: The square-root Nyquist filter has a response matched to the symbol or S-CDMA chip. An identical filter is located in the transmitter; this "matched-filter" arrangement gives optimal receive SNR in white noise. The cascade of the transmit and receive square-root filters gives a response with the "Nyquist property." This property, expressed in the time domain, ideally results in zero ISI, even when symbols are transmitted so close together in time that their responses significantly overlap.

Adaptive equalizer: This element compensates for channel effects, including group delay variation, amplitude slope or tilt, and microreflections. It adapts its filter coefficients to dynamically varying channel responses so as to maximize the receive MER. In effect, an adaptive equalizer creates a digital filter with the opposite response of the impaired channel.

Ingress canceller: An ingress canceller is normally included in a CMTS burst receiver to remove narrowband interference (including CB, ham and shortwave radios, etc.). It operates by dynamically detecting and measuring the interference, and adapting its coefficients to cancel it.

Acquisition and tracking loops: Tracking loops provide estimates of frequency, phase, and symbol timing, allowing the receiver to lock to the incoming signal. Acquisition refers to the initialization and pull-in process that occurs when the receiver is first powered on or changes channels.

Despreader: (S-CDMA upstream only) Despreading consists of multiplying the composite received signal by a given code sequence, and summing over all 128 chips in the code. There are 128 despreaders, one for each code. The output of the despreader is a soft symbol decision.

Slicer: The slicer selects the nearest ideal symbol, or hard decision, from the QAM constellation.

Trellis decoder: (Downstream and some S-CDMA upstream modes) The trellis decoder uses the Viterbi algorithm to choose the most likely sequence of symbols and thereby reject noise.

Descrambler: The descrambler adds a pseudorandom bit sequence to the received data bits, reversing the scrambling operation performed at the transmitter. The purpose of scrambling is to randomize the transmitted data in order to provide an even distribution of QAM symbols across the constellation.

Deinterleaver: The deinterleaver pseudorandomly reorders groups of received bits, reversing the interleaving operation performed at the transmitter. The purpose of deinterleaving is to break up long bursts of noise so that the errored bits can be corrected by the Reed-Solomon decoder.

Reed-Solomon (RS) FEC decoder: This device processes groups of bits (7- or 8-bit symbols) arranged in codeword blocks, in terms of an algebraic code using Galois field arithmetic. By processing the received code words, which include redundant parity symbols, receive symbol errors can be found and corrected, up to one corrected RS symbol for each two redundant RS parity symbols.

MPEG deframer: The downstream DOCSIS signal is grouped into 188-byte MPEG transport packets, permitting the multiplexing of video and data over the common physical layer. The MPEG deframer removes the MPEG transport overhead to recover the bytes that are delivered to the MAC layer.

MAC: The MAC layer controls the physical (PHY) layer and is the source and sink of PHY data. The MAC layer processes data frames delineated by DOCSIS headers. In the upstream, the MAC layer governs how cable modems share the channel through a request or grant mechanism.

The input and output of the slicer are complex numbers or vectors, each represented by two components: magnitude and phase, or equivalently, real (in-phase or "I") and imaginary (quadrature or "Q") parts, as shown in Figure 12. In an ideal zero-noise, zero-ISI condition, the soft decision would lie exactly on one of the constellation points, and the magnitude of the error between them would be zero. In a real-world receiver, subtracting the hard-decision vector from the soft-decision vector gives the error or noise vector at each symbol time. The implicit assumption is that a low symbol error rate exists – that is, very few decisions are incorrect, ensuring that the "decision-directed" error vector from the nearest symbol nearly always equals the true error vector from the correct reference symbol.

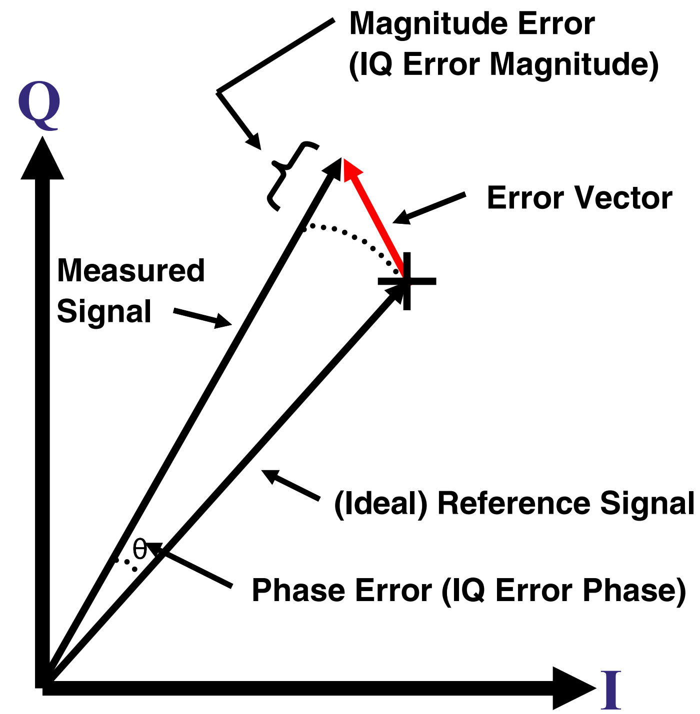

Figure 12: The Error Vector Is the Difference Between the Measured Signal (Soft Decision) and the Reference or Target Signal (Hard Decision). (Source: Hewlett-Packard)

For RxMER, we are concerned with the average power of the error vector, which is computed, as shown previously in Figure 11, by taking the squared magnitude of the complex error vector and accumulating or averaging it over a given number of symbols *N*. This process gives the error vector power (or noise power) at the slicer. Because we want the ratio of signal to noise, we divide the average signal power (a known constant for each constellation, such as 64-QAM or 256-QAM) by the average error vector power. We then take the logarithm to convert to decibels, giving RxMER in dB. To summarize: RxMER is simply the ratio of average symbol power to average slicer error power, expressed in dB.

More About Modulation Error Ratio

Modulation error ratio is digital complex baseband SNR—in fact, in the data world, the terms "SNR" and "MER" are often used interchangeably, adding to the confusion about SNR, especially considering that, as mentioned previously, in the telecommunications world, the terms "CNR" and "SNR" are often used interchangeably.

Why use MER to characterize a data signal? It is a direct measure of modulation quality and has linkage to bit error rate. Modulation error ratio is normally expressed in decibels, so it is a measurement that is familiar to cable engineers and technicians. It is a useful metric with which to gauge the end-to-end health of a network, although by itself, MER provides little insight about the type of impairments that exist.[9]

Figure 13 illustrates a 16-QAM constellation. A perfect, unimpaired 16-QAM digitally modulated signal would have all of its symbols land at exactly the same 16 points on the constellation over time. Real-world impairments cause most of the symbol landing points to be spread out somewhat from the ideal symbol landing points. Figure 13 shows the vector for a *target symbol* – the ideal symbol we want to transmit. Because of one or more impairments, the *transmitted symbol* vector (or received symbol vector) is a little different than ideal. *Modulation error* is the vector difference between the ideal target symbol vector and the transmitted symbol vector. That is:

$$\overrightarrow{\textbf{Modulation Error}} = \overrightarrow{\textbf{Transmitted Symbol}} - \overrightarrow{\textbf{Target Symbol}} \qquad \text{[Eq. 20]}$$

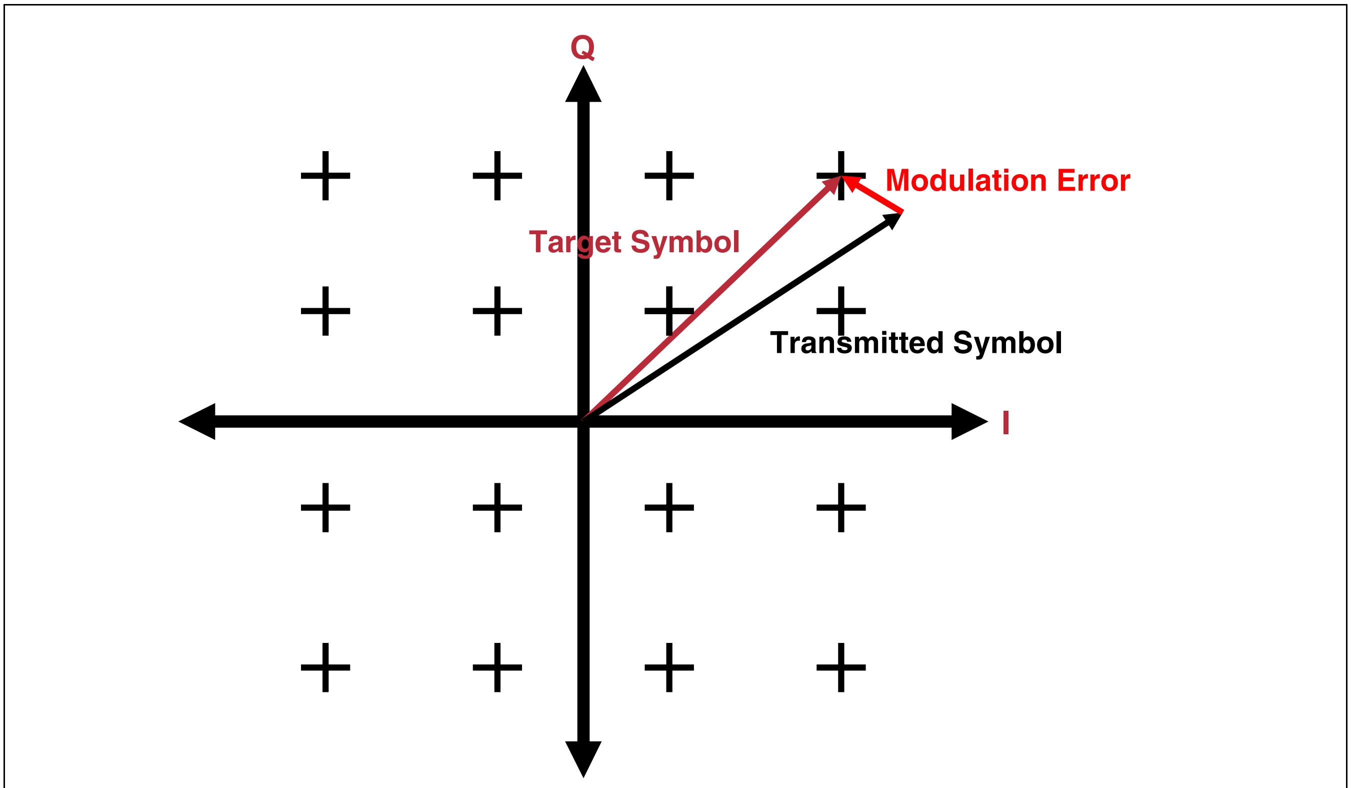

Figure 13: Modulation Error Is a Measure of Modulation Quality. (Source: Hewlett-Packard)

If a constellation diagram is used to plot the landing points of a given symbol over time, the resulting display forms a small “cloud” of symbol landing points rather than a single point. Modulation error ratio is the ratio of average symbol power to average error power (refer to Figure 14 on page 28):

$$\text{MER(dB)} = 10\log(\text{Average symbol power} \div \text{Average error power}) \qquad \text{[Eq. 21]}$$

In the case of MER, the higher the number, the better.

9. The reader is referred to the literature for discussions of visual constellation impairment evaluation.

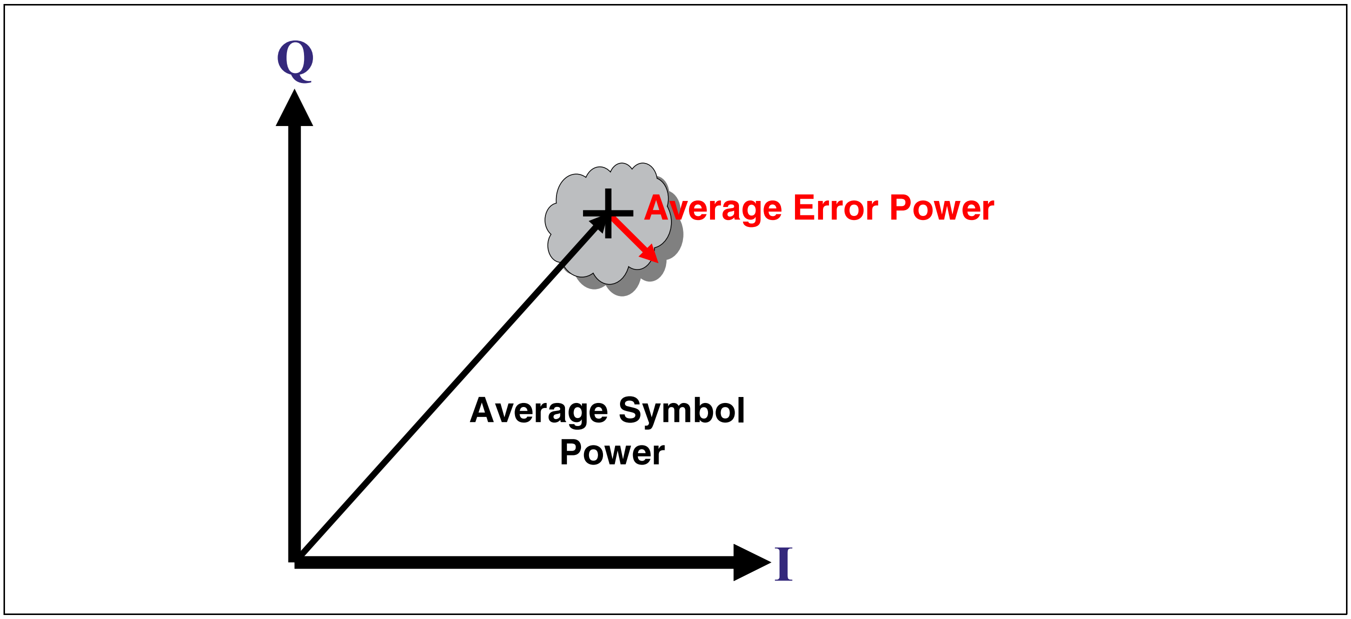

Figure 14: Modulation Error Ratio Is the Ratio of Average Symbol Power to Average Error Power. (Source: Hewlett-Packard)

Mathematically, a more precise definition of MER (in decibels) follows:

$$MER = 10\log_{10}\left[\frac{\sum_{j=1}^{N}\left(I_j^2 + Q_j^2\right)}{\sum_{j=1}^{N}\left(\delta I_j^2 + \delta Q_j^2\right)}\right] \qquad \text{[Eq. 22]}$$

where I and Q are the real (in-phase) and imaginary (quadrature) parts of each sampled ideal *target symbol* vector, δI and δQ are the real (in-phase) and imaginary (quadrature) parts of each *modulation error* vector. This definition assumes that a long enough sample is taken so that all the constellation symbols are equally likely to occur.

In effect, MER is a measure of how "fuzzy" the symbol points of a constellation are. Table 4 summarizes the approximate E_S/N_0 range that will support valid MER measurements for various DOCSIS modulation constellations. The two values in the table for the lower threshold correspond to ideal uncoded symbol error rate (SER) = 10^{-2} and 10^{-3}, respectively. The upper threshold is a practical limit based on receiver implementation loss. Outside the range between the lower and upper thresholds, the MER measurement is likely to be unreliable. The threshold values depend on receiver implementation. Some commercial QAM analyzers may have values of the lower E_S/N_0 threshold 2 to 3 dB higher than those shown in the table.

Table 4: Valid MER Measurement Ranges

Modulation Format	Lower E_S/N_0 Threshold	Upper E_S/N_0 Threshold
QPSK	7–10 dB	40–45 dB
16 QAM	15–18 dB	40–45 dB
64 QAM	22–24 dB	40–45 dB
256 QAM	28–30 dB	40–45 dB

Good engineering practice suggests keeping RxMER in an operational system at least 3 to 6 dB or more above the lower E_S/N_0 threshold.[10] This guideline will accommodate numerous factors that can affect operating headroom, including temperature-related signal-level variations in the coaxial plant, amplifier and optoelectronics misalignment, and imperfect calibration of test equipment. The lower E_S/N_0 threshold can be thought of as an "MER failure threshold" of sorts, that is, when unequalized RxMER approaches the lower E_S/N_0 threshold, the channel may become unusable with the current modulation. Possible workarounds include switching to a lower order of modulation, using adaptive equalization, or identifying and repairing what is causing the low RxMER in the first place.

Transmit MER

Although this paper is mainly concerned with RxMER, transmit MER (TxMER) is also of interest. TxMER is defined as the MER produced by a transmitter under test, as measured by an ideal test receiver. In a real test, the receiver is, of course, not ideal and will introduce its own degradations to MER measurement. The receiver contribution, if small (a few dB) and accurately measurable, can be used to correct the TxMER measurement. In addition, DOCSIS provides the following method for removing the effects of the frequency response of a test receiver on a TxMER measurement. First a near-ideal test transmitter is connected to the test receiver, and the receiver equalizer coefficients are allowed to converge in order to compensate for the frequency response of the test bed. The receive equalizer is then frozen, and the transmitter under test is connected in place of the test transmitter. The MER reading taken from the test receiver is the unequalized TxMER measurement. The test receiver is again allowed to adapt its equalizer coefficients, and the resulting MER reading is the equalized TxMER measurement.

Factors Affecting MER Measurement

Because MER is a digital computation performed on digital quantities in the receiver, it is by nature extremely accurate in itself. However, the measured value can be affected by many things. As a result, MER may not accurately reflect the CNR or E_S/N_0 at the input to the receiver.

- *Statistical variation*: The number of samples N over which the MER (or RxMER) is averaged affects the reliability of the measurement. For independent samples, the standard deviation of a measurement is in general proportional to $1/\sqrt{N}$, so, for example, averaging over 10,000 samples will result in 10 times smaller standard deviation of the MER measurement than using 100 samples. A smaller standard deviation means the MER measurement will appear more stable. Conversely, taking fewer samples can also offer advantages. In a sense, the number of samples N provides a

10. Many cable operators use the following unequalized MER (RxMER) values as minimum acceptable operational values: QPSK ~18 dB; 16-QAM ~24 dB; 64-QAM ~27 dB; and 256-QAM ~31 dB.

control analogous to the video averaging function on a spectrum analyzer. A smaller number of symbols allows the observation of transients in the MER measurement, which can highlight the effects of burst noise, distortion, clipping, and pulsed ingress.

- *Unequal occurrence of symbols*: The average constellation power is a known constant for each constellation, such as 64-QAM or 256-QAM, and does not need to be computed. Some MER implementations nevertheless compute the average constellation power by taking the complex magnitude-squared of the received ideal symbols and averaging it over a given number of captured symbols *N*. For large *N*, this works fine and approaches the average constellation power. However, if the MER measurement is performed over just a few symbols (for example, $N < 100$), the result may be unreliable because, in some cases, many large QAM symbols (near the outer edges of the constellation) will happen to be transmitted, and in other cases, many small QAM symbols (in close to the origin, or center of the constellation) may happen to be transmitted.
- *Nonlinear effects*: Nonlinearities in the signal path, including laser clipping and amplifier compression, can affect outer constellation points more severely than inner points. As an example, in one return-path system with nonlinearities present,[11] the equalized RxMER (24-tap equalizer) of a QPSK signal was measured at 38.0 dB, a level that seemed to promise good margin for higher orders of modulation. However, with 16-QAM, the equalized RxMER was 31.9 dB, and with 256-QAM, the equalized RxMER was 30.2 dB. Hence, when measuring MER, it is important to measure the same constellation that will be used for transmitting data. It is also important to capture a large enough data sample to ensure that all symbols occur with equal likelihood.
- *Linkage of carrier loop bandwidth to capture length*: Some MER measurement equipment does not have an explicit carrier tracking loop. Instead, a block of *N* received symbols is captured and averaged. The averaging produces an effective carrier tracking loop with equivalent one-sided noise bandwidth $BL = R_S/2N$, where R_S is the symbol rate. To achieve the DOCSIS specification value of BL = 50 kHz, for example, would require $N = 54$ symbols at $R_S = 5.36$ MHz. As mentioned previously, measuring MER over such a small number of symbols can give unreliable results.
- *Implementation-loss MER ceiling*: Even if the input E_S/N_0 is very high, the MER reading will saturate at a value reflecting the implementation loss of the receiver. The receiver contributes noise to the MER measurement because of front-end noise figure; imperfect time, frequency, or phase tracking; round-off effects; imperfect equalization; etc. For example, it is unusual in a 256-QAM receiver for the MER measurement to go much above 40 to 45 dB, even when there is no noise at the receiver input, as mentioned previously in the description of Table 4.
- *Symbol-error MER floor*: The slicer produces the hard decision by taking the soft decision and finding the nearest ideal constellation point. If the wrong constellation point is chosen, a symbol error occurs. The error vector magnitude then indicates the distance to the nearest symbol point, which may be closer than the correct symbol, meaning that the error will seem smaller than it really is, and the MER will seem better than its true value. As a general rule, the MER measurement is not valid when the input E_S/N_0 is below the point that produces roughly a 1-percent symbol error rate (before trellis or FEC decoding), as mentioned previously in the description of Table 4.
- *Analog front-end noise*: The analog front end of the receiver contributes thermal noise and possibly spurious products, effectively raising the noise floor of the system and lowering the RxMER relative to the CNR measurement. This effect is most pronounced at low RF input levels.
- *Phase noise*: Phase noise is a slowly varying random phase in the received signal. The analog tuner is a primary contributor of phase noise in the receiver. The carrier phase tracking loop in the receiver

11. Reported by Cooper, et al. in a paper presented at SCTE Cable-Tec Expo 2006.

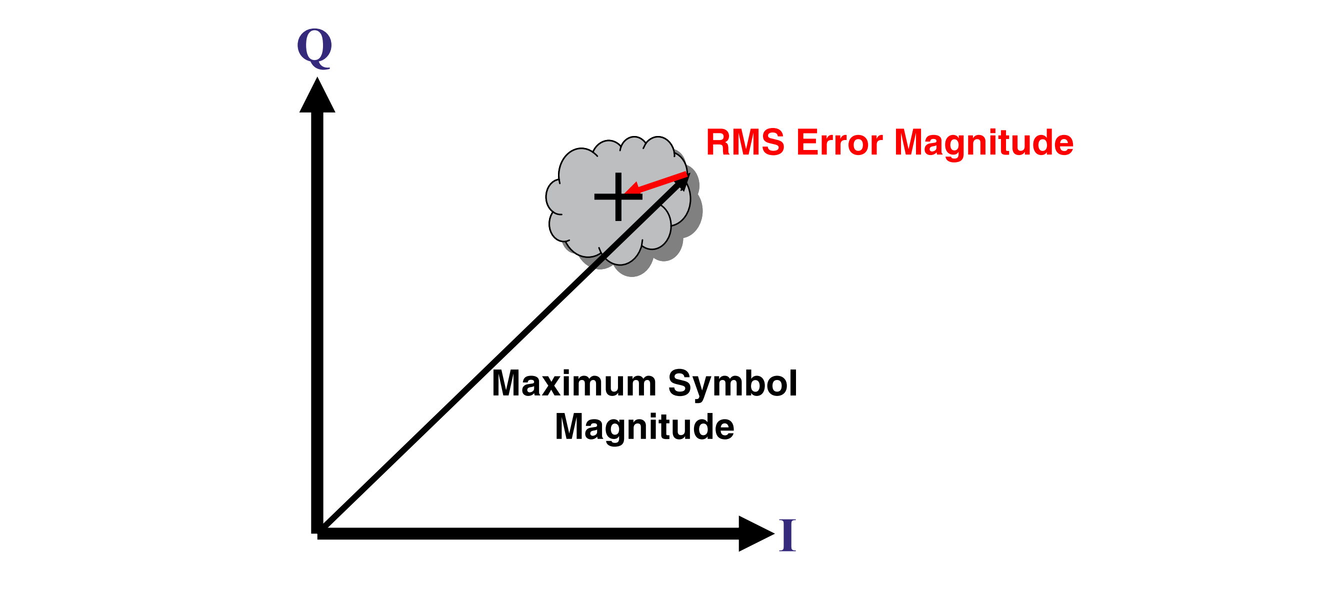

Figure 15: Error Vector Magnitude Is the Ratio (in Percent) of RMS Error Magnitude to Maximum Symbol Magnitude. (Source: Hewlett-Packard)

DOCSIS MIB Definition of Upstream RxMER

In DOCSIS, the upstream RxMER MIB measurement is defined as an estimate, provided by the CMTS demodulator, of the ratio:

$$\frac{\text{average constellation energy with equally likely symbols}}{\text{average squared magnitude of error vector}} \qquad \text{[Eq. 23]}$$

The CMTS RxMER is averaged over a given number of bursts at the burst receiver, which may correspond to transmissions from multiple users. The MIB does not specify whether receive equalization is enabled; this is implementation-dependent.

EVM Versus MER

Another measurement metric that is closely related to MER is error vector magnitude (EVM). As shown previously in Figure 13, EVM is the magnitude of the vector drawn between the ideal (reference or target) symbol position of the constellation, or hard decision, and the measured symbol position, or soft decision. By convention, EVM is reported as a percentage of peak signal level, usually defined by the constellation corner states. The mathematical definition of EVM follows:

$$EVM = (E_{RMS}/S_{max}) \times 100\% \qquad \text{[Eq. 24]}$$

where E_{RMS} is the RMS error magnitude and S_{max} is the maximum symbol magnitude. EVM is illustrated in Figure 15. From this, it is clear that the lower the EVM, the better. Contrast EVM with MER, where the higher the MER, the better.

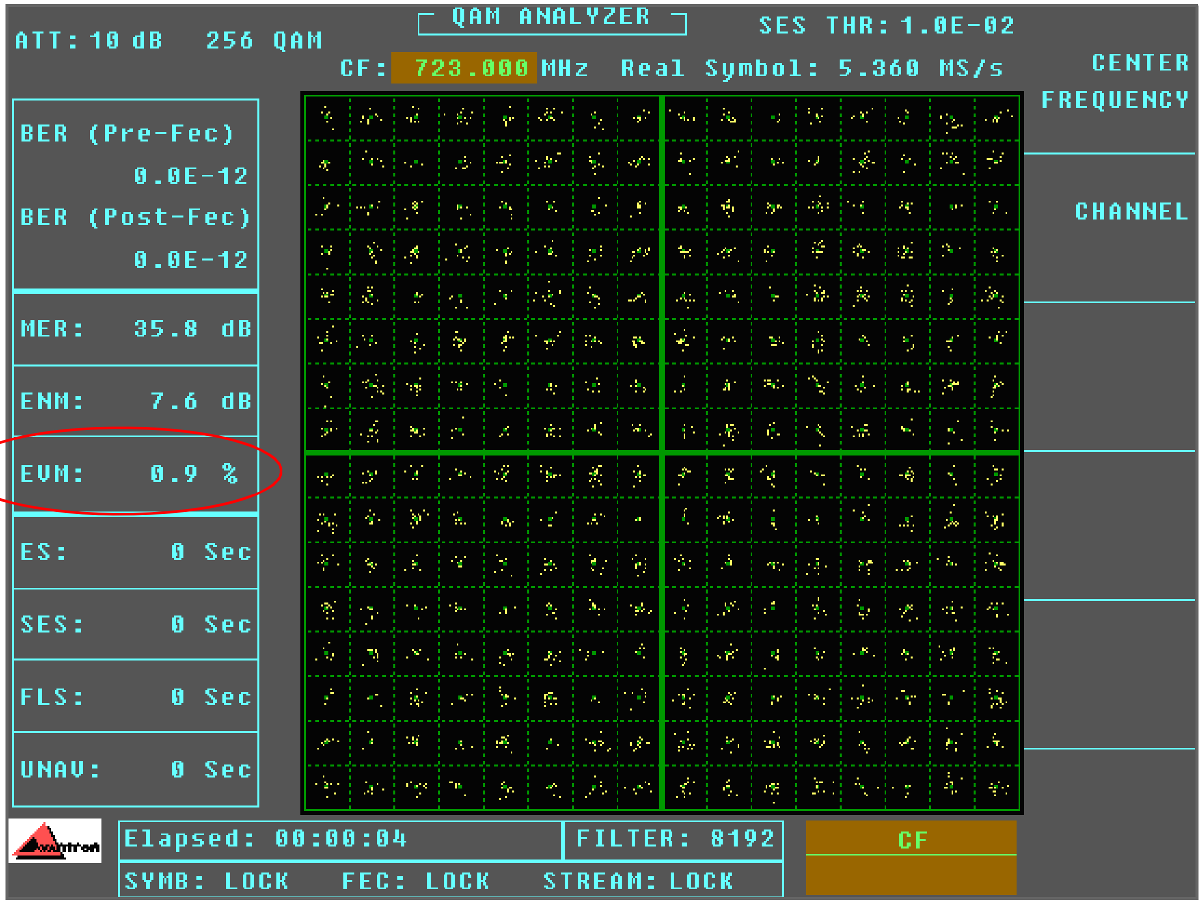

Figure 16: The EVM of This Downstream 256-QAM Digitally Modulated Signal Is 0.9 Percent. (Courtesy of Sunrise Telecom)

Figure 17 and Figure 18 on page 37 show two examples of a 16-QAM upstream digitally modulated signal. The constellation in Figure 17 illustrates a relatively unimpaired signal[14] whose unequalized MER is 27.5 dB. Figure 18 shows an impaired signal, where the unequalized MER is only 19 dB. This level is close to the failure threshold for unequalized, uncoded 16-QAM for some demodulators. In fact, in Figure 18, we can see that some soft decisions are close to the decision boundaries. Note that there is really no way to tell for sure *what* is causing the low MER in Figure 18 by simply observing the display—it could be because of low CNR or perhaps because of one or more linear or nonlinear distortions.

14. The signal has a slight amount of phase noise, observable in the angular spread of the corner constellation points, but is otherwise what could be considered a "clean" signal.

The spectrum analyzer screen shot in Figure 19 shows the upstream spectrum of a cable plant in the band from 0 to 100 MHz. The region from 5 MHz to 22 MHz contains relatively low-level ingress. Clean upstream spectrum is seen in the 22 MHz to 42 MHz range. The upstream DOCSIS carrier is located at 32 MHz under the rightmost red cursor. The diplexer roll-off is seen in the 42 MHz to 54 MHz range. This upstream is relatively clean: The CNR of the carrier is excellent at about 36 dB; there is no visible common path distortion, ingress, or strong impulse noise anywhere near the signal; and we see only a modest amount of ingress well below the signal frequency. Yet although QPSK worked fine in this upstream, unequalized 16-QAM was found to be unusable. The reason? Linear distortions. A severe impedance mismatch about 1100 feet from the node caused a micro-reflection, resulting in amplitude and group delay ripple. These distortions were not visible on the spectrum analyzer display, yet were significant enough to degrade the upstream RxMER and impair 16-QAM transmission. The CMTS reported low SNR —in reality, unequalized RxMER—of 21 dB, which agreed with third-party test equipment measurement of unequalized MER (refer to Figure 20 on page 39). To support 16-QAM transmission on this plant, the operator could enable the preequalization function that is available in DOCSIS 1.1 and higher modems, or increase the level of FEC coding.

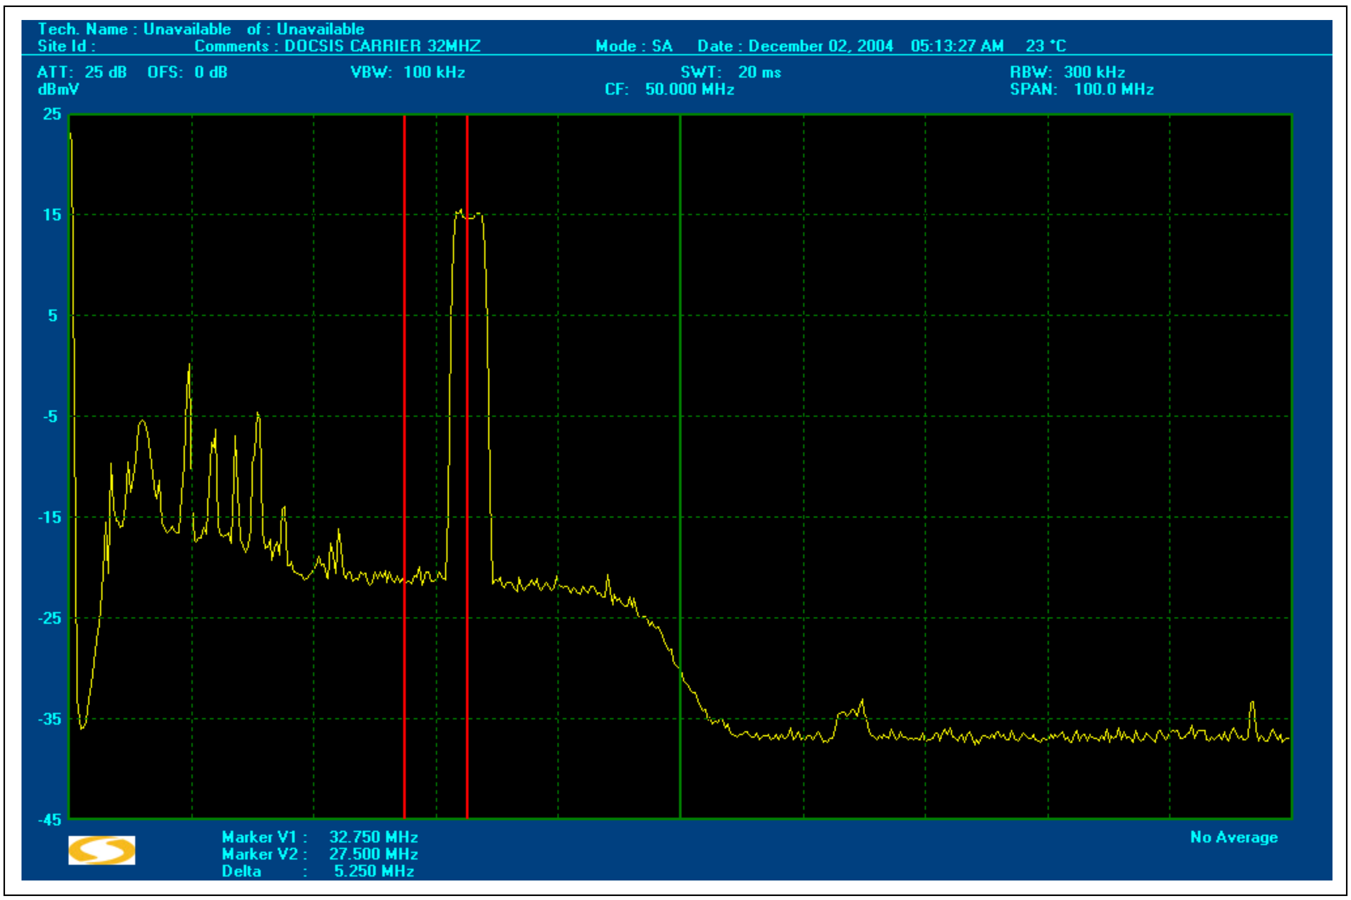

Figure 19: A Spectrum Analyzer Display Shows What Appears to Be a Relatively Clean Upstream, But Unequalized 16-QAM Would Not Work. The CNR Is About 36 dB. (Courtesy Sunrise Telecom)

0.46 dB must be added to the CNR measurement. Another approach is to measure both signal and noise with the spectrum analyzer in maximum-hold mode, resulting in the analyzer trace filling in the areas where the signal was off. However, this method is susceptible to errors in that the highest TDMA user's power is being measured, not the average signal power; and any high-noise excursions are measured, not the average noise.

CMTS burst receivers can measure RxMER on each upstream burst that is received, thereby avoiding the problem of the quiet times when no signal is present. Because the receiver synchronizes to each upstream burst, it ignores the dead times between bursts.

As mentioned previously, the DOCSIS MIB defines the CNIR measurement as the ratio of expected burst signal power to average noise plus interference during the times when no signal is present. This measurement is tailored to the TDMA nature of the upstream.

S-CDMA Transmission and E_S/N_0 or RxMER per Code

S-CDMA introduces a new SNR concept: E_S/N_0 or RxMER per code. First, let's review the basics of S-CDMA. Figure 21 shows the fundamental concept of S-CDMA data transmission: Each symbol of data is multiplied by a spreading code at the cable modem transmitter and summed with other spread symbols originating both from the same cable modem and from other cable modems on the plant. The composite signal travels across the upstream RF channel, where noise is unavoidably added. At the CMTS burst receiver, the signal is applied to 128 despreaders, which reverse the process by multiplying by the respective spreading codes and summing over all 128 chips in the code. This process reproduces the original data symbols at the receiver output, perturbed, of course, by the noise.

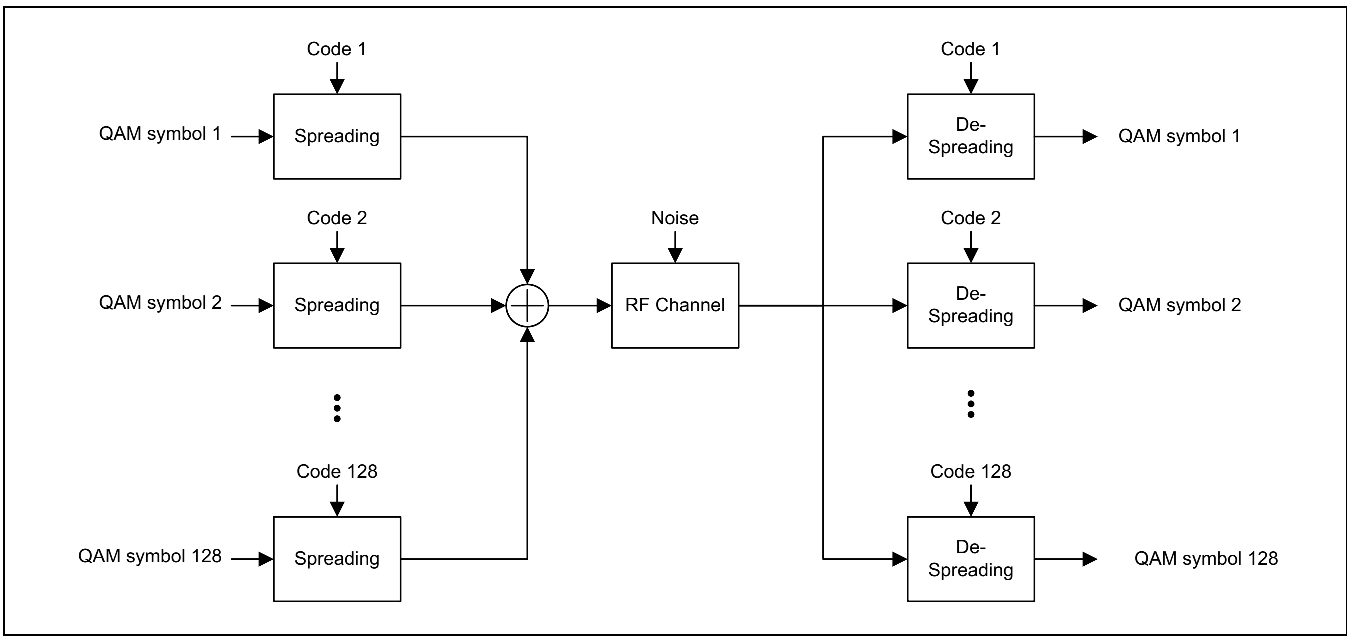

Figure 21: Upstream S-CDMA Data Transmission Consists of Spreading at the Transmitter and Despreading Each Code at the Receiver.

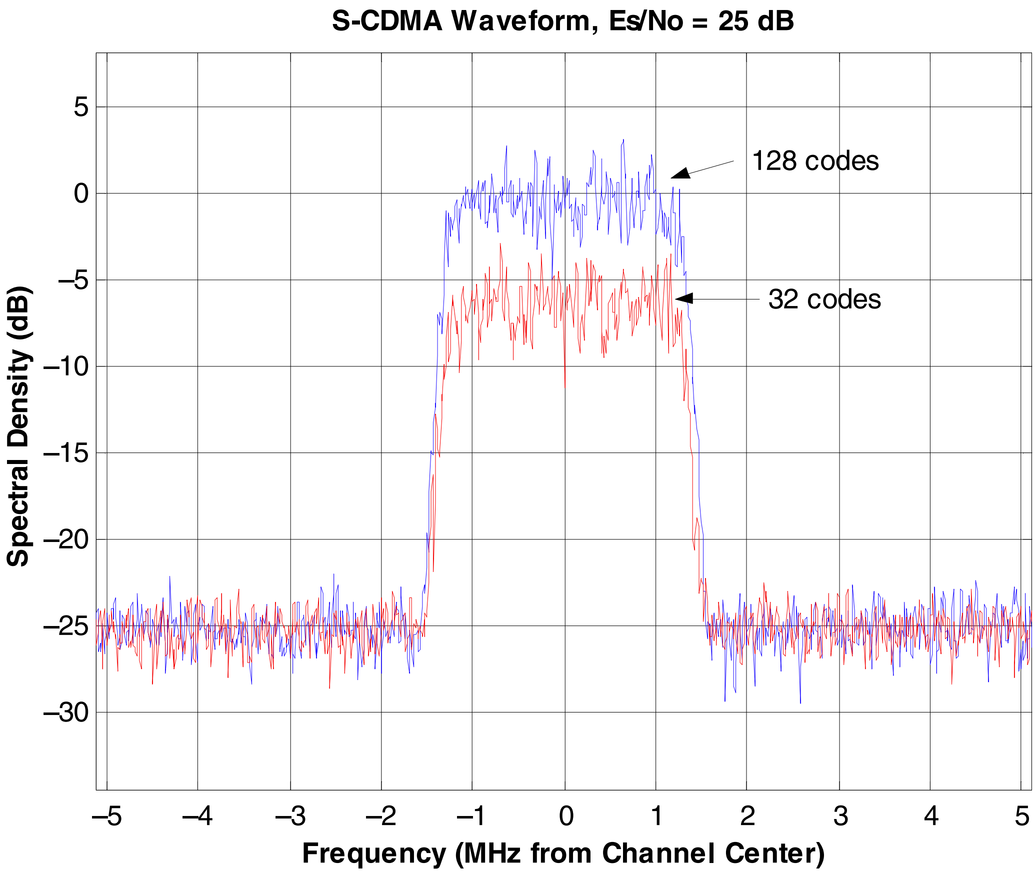

Upper trace: all 128 active codes are transmitted; CNR = 25 dB.
Lower trace: only 32 of the 128 active codes are transmitted; CNR reduces to 19 dB, but E_S/N_0 per code is still 25 dB because of 6 dB spreading gain.

Figure 24: S-CDMA Upstream Spectrum with E_S/N_0 per Code = 25 dB

Spreading Gain in S-CDMA

A "spreading gain" or "processing gain" results when fewer than all the active codes are transmitted; in the lower trace of Figure 24, the spreading gain is 10 log(128/32) = 6 dB. Of course, the throughput is also reduced by a factor of 4 for this frame. This situation might occur temporarily in a lightly loaded S-CDMA upstream channel, in which the CNR might be seen varying up and down dynamically by a few dB, analogous to the carrier pulsing on and off in a lightly loaded TDMA upstream. The RxMER, however, remains steady at 25 dB in both the S-CDMA and TDMA cases. In fact, the RxMER may even show a slight improvement when fewer codes are transmitted, because spreading gain tends to reduce residual ISI effects.

Acknowledgments

Authors: Ron Hranac (Cisco Systems, Inc.) and Bruce Currivan (Broadcom Corporation)

The authors would like to acknowledge the contributions of the following reviewers:

Arris International, Inc.: Scott Brown

Broadcom Corporation: Victor Hou, Tom Kolze (special mention), Jonathan Min, and Rich Prodan

BigBand Networks: Hal Roberts

Cisco Systems, Inc.: John Downey, Rick Meller, and Greg Taylor

HEYS Professional Services: Francis Edgington

Motorola, Inc.: Jack Moran

Sunrise Telecom, Inc.: Gerard Terreault

Xytrans Corporation: Rob Howald