


White Paper

Digital Transmission: Carrier-to-Noise Ratio, Signal-to-Noise Ratio, and Modulation Error Ratio

January 2012



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 .




Figure 2: Baseband SNR Measurement

Analog Video CNR in Cable Networks

Consider CNR, which is generally accepted to be a predetection measurement—that is, one made at RF. When only analog TV channels were carried on cable networks, CNR was understood to be the difference, in decibels, between the amplitude of a TV channel visual carrier³ and the root mean square (RMS) amplitude of system noise in a specified noise power bandwidth. In this application, noise power bandwidth is normally specified as the modulation bandwidth, which is approximately equal to the bandwidth of the baseband modulating signal. It is common practice to express power in terms of RMS voltage across a nominal resistance. For example, *NCTA Recommended Practices for Measurements on Cable Television Systems* defines 0 dBmV (decibels referenced to one millivolt) as the power of a signal of 1 millivolt RMS in 75 ohms, or 13.33 nanowatts = -48.75 dBm.

According to the Federal Communications Commission's (FCC's) cable regulations in §76.609 (e), system noise is the "total noise power present over a 4 MHz band centered within the cable television channel." This latter definition is applicable only to analog National Television System Committee (NTSC) TV channel CNR measurements, and defines the approximate bandwidth of the baseband video that modulates the channel visual carrier.

The FCC does not actually use the term CNR in the rules. §76.605 (a)(7) states "The ratio of RF visual signal level to system noise shall...not be less than 43 decibels." That definition is more or less in line with the general definition of SNR, although it is understood in this specific instance to mean CNR. Even though the FCC's cable rules mandate a minimum CNR of 43 dB, good engineering practice targets end-of-line analog TV channel CNR in the 46 to 49 dB range. More on this topic appears later in this paper.

3. Visual carrier amplitude is the RMS amplitude of the synchronizing peak or "sync." This amplitude corresponds to the visual carrier's peak envelope power (PEP).

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:

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

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 $10\log(\text{RBW in Hz})$ 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:

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

In decibels, C/N_0 is expressed in dB-Hz, which means “decibels referenced to one Hz.” Because C/N_0 is not unitless like other SNR and CNR metrics, care must be taken to reference the noise measurement to a 1 Hz noise bandwidth.

C/N_0 is especially useful for measuring the CNR of a narrowband signal such as an unmodulated or continuous wave (CW) carrier. Consider a spectrum analyzer capture of a CW signal in a white-noise background, with the analyzer RBW set to 100 kHz. Assume that placing the analyzer marker on the CW signal indicates an amplitude of -10 dBmV, and moving the marker to the displayed noise floor shows -40 dBmV. Be careful—this noise reading is *not* in dBmV/Hz, but represents the noise power in the analyzer RBW as mentioned previously, giving a spectral density of -40 dBmV/(100 kHz). To convert to a 1 Hz bandwidth, subtract $10\log(100,000 \text{ Hz}) = 50$ dB-Hz. So the noise density is actually -40 dBmV $- 50$ dB-Hz = -90 dBmV/Hz. Some spectrum analyzers have a marker noise function that provides automatic readings in a 1 Hz bandwidth, eliminating the need for this conversion. The true C/N_0 is then (using dB quantities):

$$\begin{aligned} C/N_0 &= \text{Signal} - \text{Noise} + \text{RBW} \\ &= -10 \text{ dBmV} - (-40 \text{ dBmV}) + 50 \text{ dB-Hz} \\ &= 80 \text{ dB-Hz} \end{aligned} \quad [\text{Eq. 5}]$$

To convert C/N_0 to CNR in a given bandwidth B, we use

$$\text{CNR} = C/N = C/(N_0 B) \quad [\text{Eq. 6}]$$

So in decibels, to convert C/N_0 to CNR, subtract $10\log(B)$. The CW signal in this example would have a CNR in a 6 MHz bandwidth (using decibel quantities) of:

$$\begin{aligned} \text{CNR} &= C/N_0 - 10\log(B) \\ &= 80 \text{ dB-Hz} - 10\log(6 \text{ MHz}) \\ &= 80 \text{ dB-Hz} - 67.8 \text{ dB-Hz} \\ &= 12.2 \text{ dB} \end{aligned} \quad [\text{Eq. 7}]$$

This example illustrates an important principle when measuring a mix of discrete signals (CW or any signals that are much narrower than the RBW of the analyzer) and spread signals (such as noise or modulated signals that are much wider than the analyzer RBW). The spectrum analyzer marker simply measures power in the RBW. For a narrowband signal, this measurement equals the carrier power. For noise, it gives the density referenced to the RBW. Scale to a 1 Hz bandwidth to get C/N_0 , and scale to a desired bandwidth B to get CNR.

Digitally Modulated Signal CNR

What about CNR measurement of digitally modulated signals on a cable plant? The DOCSIS *Radio Frequency Interface Specification* states an assumed minimum 35 dB CNR for downstream digitally modulated signals. If the network analog TV channel CNR is maintained in the 46 dB or higher range, in most cases, there will be little or no problem complying with the DOCSIS assumed minimum for downstream digitally modulated signals. The DOCSIS assumed minimum upstream CNR for digitally modulated signals is 25 dB. Carrier power—the “C” in CNR—is the average power level of the digitally modulated signal, often called digital channel power. It is measured in the full occupied bandwidth of the signal; for example, 6 MHz for a North American DOCSIS downstream signal.

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.

Downstream Digitally Modulated Signal CNR—An Example

Assume that a cable network has been designed to provide a downstream end-of-line CNR of 46 dB with +15 dBmV subscriber tap levels for analog TV channels, and a downstream 64-QAM DOCSIS digitally modulated signal is carried at –10 dBc. Thus the digital channel power of the 64-QAM signal at the tap spigot will be +5 dBmV, or 10 dB lower than the +15 dBmV analog TV channel levels. What is the 64-QAM signal CNR? It is *not* 36 dB, as one might first assume.

Because analog TV channel CNR is 46 dB at the tap spigot, the noise-floor amplitude N_{NTSC} for the analog channels is +15 dBmV – 46 dB = –31 dBmV (4 MHz noise power bandwidth for analog NTSC TV channels). To determine the 64-QAM signal CNR, we have to first calculate what the noise-floor amplitude is for that QAM signal, based on a noise power bandwidth equal to its symbol rate. For DOCSIS 64-QAM signals, the symbol rate is 5.056941 Msym/sec, so the noise power bandwidth is 5.06 MHz (refer to [Table 2](#)). From this, we can calculate the noise-floor amplitude $N_{\text{64-QAM}}$ for the QAM signal with the equation:

$$N_{\text{64-QAM}} = N_{\text{NTSC}} + [10\log(5.06/4)] = -29.98 \text{ dBmV} \quad [\text{Eq. 8}]$$

The 64-QAM signal CNR is +5 dBmV – (–29.98 dBmV) = 34.98 dB.

The CNR of the digitally modulated signal is degraded by more than the 10-dB reduction in signal level because the wider noise bandwidth in the digital signal case allows more AWGN through to the demodulator.

E_s/N_0 and CNR of a Digitally Modulated Signal

E_s/N_0 is the most prevalent parameter used in digital communications to represent the SNR of a signal. It is defined as the ratio of the average energy E_s per QAM symbol to the noise power spectral density N_0 with the noise assumed white. It is a unitless ratio and is normally expressed in dB. Energy per symbol is equal to the average power—that is, the energy in 1 second—of the signal divided by the number of symbols in 1 second.

How can we reconcile E_s/N_0 with CNR? If we multiply the numerator and denominator by the symbol rate R_s , we get:

$$\frac{E_s}{N_0} = \frac{E_s R_s}{N_0 R_s} = \frac{\text{Signal Power}}{\text{Noise Power in Symbol Rate Bandwidth}} = \text{CNR} \quad [\text{Eq. 9}]$$

(assuming no synchronous code-division multiple access [S-CDMA] spreading, which we discuss later). This equation tells us why the measurement of CNR for digitally modulated signals typically uses the noise power bandwidth equal to the symbol rate: It results in CNR being equal to E_s/N_0 .

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 $|H_T(f)|$. [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 $|H(f)| = |H_T(f)||H_R(f)| = |H_T(f)|^2$. 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 $|H_T(f)|^2$. 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 $|H_T(f)|^2$. 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.

from cascading the same filter twice; in the latter case, the squaring results from converting the magnitude spectrum to a power spectrum. In a real system, using a spectrum analyzer, we are used to observing the power spectrum of the Tx signal in decibels, which is $20\log|H_I(f)|$ or $10\log|H_I(f)|^2$.

Ideal QAM Spectra

Now we return to the discussion of QAM spectra. Figure 4 shows four ideal QAM spectra, all with the same symbol rate R_S , which is normalized to 1 in these plots. In part (a) of the figure, a perfect brick-wall rectangular spectrum is shown. As the magnitude response $|H_I(f)|$ of a matched filter, this signal is unrealizable in practice because the pulse-shaping filter $h_I(t)$ would have to be infinitely long in time duration. Part (a) of the figure also represents the power spectrum $|H_I(f)|^2$ of the waveform transmitted using such an ideal filter. Despite its impracticability, it is useful as an illustration of an ideal world in which the occupied bandwidth equals the symbol rate and there is no excess bandwidth.




Figure 4: Ideal QAM Spectrum Plots

Part (b) of Figure 4 shows the full-response magnitude spectrum $|H(f)|$ used in the DOCSIS upstream, representing the cascade of the transmit and receive filters. In order to make the filters realizable, an excess bandwidth of 25 percent ($\alpha = 0.25$) is used, resulting in the S-shaped “raised cosine” roll-off regions shown in red, while the passband ideally remains flat. This spectrum possesses the Nyquist property in the frequency domain: If the frequency response $H(f)$ is replicated many times shifted by multiples of the symbol rate, and the copies are overlaid and added as illustrated in Figure 5, the result is a flat spectrum, which results in zero intersymbol interference (ISI).⁵ Because of the cascading property mentioned previously, part (b) of Figure 4 also represents the power spectrum $|H_I(f)|^2$ of the transmitted signal and of the square-root Nyquist filter, described next.




Figure 5: The Nyquist Property States That When Copies of the Spectrum Are Shifted by Multiples of the Symbol Rate and Added, the Result Is a Flat Spectrum, Which Results in Zero ISI.

In practice, the full Nyquist spectrum $H(f)$ of part (b) in Figure 4 is divided into two identical cascaded “square-root Nyquist” filters $H_I(f)$, one in the cable modem upstream transmitter and one in the CMTS burst receiver, using the matched filtering concept discussed earlier. The square-root Nyquist magnitude response $|H_I(f)|$ is shown in part (c) of the figure. Again, because of the cascading property, the power spectrum $|H_I(f)|^2$ of the square-root Nyquist filter and of the transmitted signal, is given in part (b) of the figure.

5. Historically, H. Nyquist discovered this property and published it in 1928, in the context of telegraphy.

When viewed on a spectrum analyzer, with vertical scale in dB, the square-root Nyquist spectrum looks like part (d) of [Figure 4](#), which has an added white-noise floor giving $\text{CNR} = 35 \text{ dB}$. The -3 dB point of the spectrum occurs at the symbol rate.

Because several of the spectra in the figure have double meanings, it might be useful to summarize [Figure 4](#):

- Part (a) is the magnitude or power spectrum of an ideal brick-wall filter.
- Part (b) is the magnitude spectrum of the full Nyquist response, and, by the cascading property, the power spectrum of the actual transmitted RF signal.
- Part (c) is the magnitude spectrum of the square-root Nyquist response, that is, the transmit or receive matched filter.
- Part (d) is the power spectrum of the actual transmitted RF signal, as seen on a spectrum analyzer in dB, with a noise floor giving $\text{CNR} = 35 \text{ dB}$.

Reading CNR from a QAM Carrier Spectrum

Note that the vertical height from the top of the spectrum haystack down to the noise floor in part (d) of [Figure 4](#) equals the $\text{CNR} (= E_S/N_0)$ value of 35 dB , making it easy to read the CNR of a QAM signal off a spectrum analyzer display. In fact, in an ideal system in which AWGN is the only impairment, the height of the spectrum above the noise floor also equals the equalized RxMER value (discussed later) in the range over which the RxMER measurement is valid. As we will see, the RxMER in a real system is usually somewhat lower than the CNR because RxMER includes receiver imperfections and hidden distortions of the RF input signal.

Why does this simple CNR measurement method work? First, consider the ideal case: the brick-wall signal. Viewed on a spectrum analyzer, the height of a brick-wall power spectrum gives the signal density, S_0 , in units of dBmV in the spectrum analyzer RBW, and we can scale it to the symbol rate bandwidth, giving the total signal power. Similarly, the height of the noise floor gives the noise density, N_0 , also in units of dBmV in the spectrum analyzer RBW, and we can scale it to the symbol rate bandwidth. We subtract one measurement from the other, giving the ratio in dB of total signal (or carrier) power to noise power in the symbol rate bandwidth. This value is CNR. So, this discussion explains why for a brick-wall signal, the CNR is simply the vertical height from the top of the power spectrum down to the noise floor.

But why does this simple measurement technique also work for the square-root Nyquist power spectrum of part (d) of [Figure 4](#)? The key lies in comparing the brick-wall power spectrum in part (a) of the figure to the square-root Nyquist power spectrum in part (b) of the figure. Note that they have the same height and the same area. (The reasoning that they have the same area follows: Overlay the two curves, notice that the S-shaped roll-off curve is symmetric, and imagine cutting and pasting two halves of the roll-off region to match the brick-wall curve. [Figure 5](#) also helps illustrate this point.) We know that, in general, the area under a power spectrum gives the signal power. So, the brick-wall and square-root Nyquist power spectra have the same height and same signal power. What about the noise power? As mentioned previously, when computing CNR, we measure the noise power in the symbol rate bandwidth; so the noise power is, by definition, the same for the brick-wall and square-root Nyquist power spectra. That means the same simple measurement will work for the square-root Nyquist signal: The vertical height from the top of the power spectrum down to the noise floor gives the CNR and E_S/N_0 value.

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.”

- Be sure that the signal being measured (the information signal itself or the system noise, which is also a “signal”) is at least 10 dB above the noise floor of the spectrum analyzer.⁷ Sometimes, a low-noise preamplifier must be added at the spectrum analyzer input, or a test point with greater signal amplitude must be found. If measurement of a signal with very low CNR cannot be avoided, the offset caused by the analyzer noise floor can be subtracted from the raw measurement in order to correct the power readings. Figure 8 in the next section gives the applicable noise-floor correction. Care should be taken when subtracting nearly equal noise power measurements (for example, System noise + Analyzer noise floor – estimated analyzer noise floor), because the result may become zero or negative because of measurement uncertainties. In that case, more smoothing of the measurements may be needed.
- Correct the measurement to account for the ratio of the resolution bandwidth to the noise bandwidth of the analyzer. The RBW is normally expressed as the –3 dB bandwidth of the RBW filter. The equivalent noise bandwidth of the RBW filter is typically 6 to 13 percent wider than its –3 dB bandwidth, requiring a 0.25 to 0.5 dB correction, respectively, to the measurement. Consult the analyzer documentation for the exact values.

More on the Effect of Noise Floor on CNR and Power Measurements

As mentioned earlier, measuring a signal with very low CNR requires a correction to back out the noise underlying the signal. Let's look more closely at the effect of underlying noise on the measurement of CNR, or, in general, the difference between S/N and $(S+N)/N$. Figure 7, which shows a close-up view of a band-limited digitally modulated signal with a CNR value of only 4 dB, illustrates the effect. The blue trace, $S + N$, is observed on a spectrum analyzer. The underlying signal S without the noise N is shown in red in the figure; it would not be visible on the spectrum analyzer because noise is always present in a real system. The top of the blue haystack is about 1.5 dB above the top of the red haystack, showing the measurement error $(S+N)/S$ caused by the noise-floor contribution.

7. Temporarily disconnect the spectrum analyzer RF input. The displayed noise should drop at least 10 dB. If it does not, a significant portion of the displayed noise is the test equipment noise floor adding to the cable system noise floor.




Figure 7: Close-Up View of Signal and Noise As Measured on a Spectrum Analyzer

When we measure the signal power on the analyzer, we observe the absolute height of the Signal + Noise haystack (blue trace). Expressing this measurement in units of power (non-dB) gives us the quantity (S+N). When we take the uncorrected CNR measurement, as discussed earlier, we measure the distance in dB from the flat top of the blue Signal + Noise haystack down to the noise floor. Converting to a unitless ratio (again non-dB) gives us the quantity (S+N)/N. We can manipulate these two quantities algebraically to get the desired values. First, we compute the error or offset to the signal power measurement caused by the noise floor, that is, the ratio (S+N)/S:

$$\frac{S+N}{S} = 1 + \frac{1}{\left(\frac{S+N}{N}\right) - 1} \quad [\text{Eq. 11}]$$

Next, we derive the true power S of the underlying signal:

$$S = \frac{(S+N)}{1 + \frac{1}{\left(\frac{S+N}{N}\right) - 1}} \quad [\text{Eq. 12}]$$

Finally, we compute the true CNR, that is, S/N:

$$\frac{S}{N} = \left(\frac{S+N}{N}\right) - 1 \quad [\text{Eq. 13}]$$

For practical use, we can rewrite the previous equations using dB quantities:

$$\text{signal_pwr_offset_dB} = 10\log\left(1 + \frac{1}{10^{\text{haystack_height_dB}/10} - 1}\right) \quad [\text{Eq. 14}]$$

$$\text{true_signal_pwr_dBmV} = 10\log\left(\frac{10^{\text{haystack_top_dBmV}/10}}{1 + \frac{1}{10^{\text{haystack_height_dB}/10} - 1}}\right) \quad [\text{Eq. 15}]$$

$$\text{true_CNR_dB} = 10^{\text{haystack_height_dB}/10} - 1 \quad [\text{Eq. 16}]$$

where:

- *haystack_height_dB* is the height of the signal haystack (including the noise-floor contribution) above the displayed noise floor, in dB.
- *haystack_top_dBmV* is the power reading on the spectrum analyzer at the top of the signal haystack (including the noise-floor contribution), in dBmV in the analyzer RBW.
- *signal_pwr_offset_dB* is the offset to the signal power measurement caused by the noise floor, in dB.
- *true_signal_pwr_dBmV* is the true signal power reading with the noise-floor contribution backed out, in dBmV in the analyzer RBW.
- *true_CNR_dB* is the true CNR with the noise-floor contribution backed out, in dB.

The formula for signal power offset or error as a function of height above the noise floor (Eq. 13) is graphed in [Figure 8](#). As a general rule, if the signal is at least 10 dB above the noise, the measurement offset will be less than about 0.5 dB. If the signal is at least 15 or 16 dB above the noise, the measurement offset will be less than about 0.1 dB and can be neglected for all practical purposes.




Figure 8: Signal Power Measurement Correction Relative to Noise Floor

Returning to the example in Figure 7, the height of the haystack above the noise as seen on the spectrum analyzer will be:

$$\begin{aligned}
 \text{haystack_height_dB} &= 10\log(1 + 10^{\text{true_CNR_dB}/10}) \\
 &= 10\log(1 + 10^{4/10}) \\
 &= 5.45 \text{ dB}
 \end{aligned}
 \tag{Eq. 17}$$

and, reversing [Eq. 17], the true CNR follows:

$$\begin{aligned}
 \text{true_CNR_dB} &= 10\log(10^{\text{haystack_height_dB}/10} - 1) \\
 &= 10\log(10^{5.45/10} - 1) \\
 &= 4 \text{ dB}
 \end{aligned}
 \tag{Eq. 18}$$

In summary: In the example in Figure 7, measuring a signal with a very low CNR of 4 dB requires subtracting 1.5 dB from the haystack-height CNR reading in order to back out the noise underlying the signal. To measure the CNR of normal QAM signals with CNR greater than about 15 dB, such a correction is not necessary. For measurement of system noise power that is less than 10 dB above the spectrum analyzer noise floor, a correction from Figure 8 is required to back out the analyzer noise-floor contribution.

CNR at the CMTS Upstream Input Port

In most cases, the CNRs for all modems at a given CMTS upstream port are identical, or nearly so. Cable modem upstream transmit levels are managed by the CMTS ranging process to provide the same receive

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:

$$\text{dB}_{\text{correction}} = 10\log(\text{NBW}/\text{RBW}) \quad [\text{Eq. 19}]$$

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.

$$\begin{aligned} \text{dB}_{\text{correction}} &= 10\log(2,560,000/100,000) \\ \text{dB}_{\text{correction}} &= 14.08 \end{aligned}$$

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.

Carrier-to-Noise-Plus-Interference Ratio

Often we need to make a distinction between the underlying thermal noise floor, which is flat, and narrowband interference, which appears as spectral lines or narrow humps on a spectrum analyzer display. In that case, we use the term “carrier-to-noise-plus-interference ratio” (CNIR). In DOCSIS, the CMTS upstream receive CNIR is defined in the MIB as the ratio of the expected commanded received signal power to the noise-plus-interference in the channel. Both the digitally modulated signal power and noise-plus-interference power are referenced to the same point—the CMTS input. The expected commanded received signal power is the power of an upstream burst that has been correctly adjusted by the CMTS ranging process. The noise-plus-interference power is the total power measured when no desired signal is present, that is, during quiet times on the upstream, when only the unwanted noise and interference are present. Note that one reason for not including narrowband interference with white noise when measuring CNR is that devices such as ingress cancellers can reduce the effect of narrowband ingress, whereas white noise cannot be similarly mitigated.

Digitally Modulated Signals and Baseband SNR

RF CNR is fairly straightforward for digitally modulated signals. But what about baseband SNR? Rather than using our previous definition of baseband video or audio SNR, we must look for a different way to measure baseband data SNR.

Views of a QAM Waveform

For further insight, consider the DOCSIS 256-QAM downstream signal shown in [Figure 9](#). In (a), we see the familiar haystack spectrum of a QAM signal. Its shape is flat on the top, except for implementation imperfections and measurement noise, and, in this example, the channel has inserted some slight amplitude ripple and tilt as well. The height of the spectrum above the noise floor is 45 dB, which is the CNR and E_s/N_0 value. Its width at the -3 dB point is equal to the symbol rate $R_S = 5.36$ MHz. Its sides are shaped by the transmit filter, with excess bandwidth equal to 12 percent; that is, the signal occupies 12 percent more bandwidth than an ideal brick-wall signal of the same symbol rate, because of the roll-off region. (As a quick check, 6 MHz occupied bandwidth $\div$ 5.36 MHz symbol rate = 1.12.)




Figure 9: Views of a QAM Waveform (256 QAM, $E_s/N_0 = 45$ dB, with slight channel amplitude tilt and ripple)

In part (b) of Figure 9, we see a time domain trace of either the in-phase (I) or quadrature (Q) component of the signal after the transmit filter. The I and Q components are independent, approximately Gaussian processes, because the QAM symbols occur randomly and are filtered by the square-root Nyquist pulse-shaping filter. When filtered by the receive filter and properly equalized, the I and Q components will describe a smooth trajectory between the constellation points. The constellation point is only a sample at a specific time of that widely varying signal. The trajectories overshoot the constellation points, with the amount of overshoot and ringing being greater for smaller values of excess bandwidth. The near-Gaussian nature of the transmit-filtered signal is seen graphically in part (c) of the figure, which shows

the histogram or probability density function (PDF) of the I or Q component of the QAM signal. The QAM signal approaches a Gaussian distribution, shown as a dotted red line, although the QAM distribution bulges more and has limited tails compared to the reference Gaussian distribution, which has the same standard deviation (signal level) as the QAM signal.

Baseband SNR could be defined as the ratio of the average power in the complex baseband (I and Q) signal to the average noise power in the symbol rate bandwidth. However, this definition would just replace RF CNR with an equivalent complex baseband CNR, because it does not include demodulation of the signal. For a digital SNR, we wish to measure the QAM signal after demodulation all the way down to its received constellation symbols. That is where RxMER comes in.

The SNR of a Demodulated Digital Signal: RxMER

The solution is to define a new quantity to represent the SNR of a digital baseband signal: receive modulation error ratio. RxMER is defined as the ratio of average constellation symbol power to average constellation error power, expressed in dB. As we will see, RxMER looks at the demodulated complex baseband constellation symbols and measures their quality. The RxMER measurement gives the near “bottom line” status of the communications link, because it is these demodulated symbols that will go on to produce correct bits, or bit errors, at the receiver output after processing by the forward error correction (FEC) decoder.

Equalized and Unequalized RxMER

Returning to [Figure 9](#), we see in part (d) the received QAM constellation before equalization. Although the CNR on the spectrum display is 45 dB, the channel imperfections have caused the constellation to exhibit an RxMER value of only 26 dB. After the adaptive equalizer compensates for the channel response, the RxMER is restored to 45 dB, as shown in part (e) of [Figure 9](#). This discussion shows that when we discuss RxMER, it is important to specify whether we are talking about the equalized or unequalized value.

[Figure 10](#) shows a diagnostics page screen capture from a residential cable modem.⁸ The indicated signal-to-noise ratio is 36 dB (circled). Most cable modems provide an equalized RxMER measurement of the downstream 64- or 256-QAM digitally modulated signal. For the example in [Figure 10](#), a QAM analyzer provided a nearly identical equalized MER measurement of the cable modem received 256-QAM signal (35 dB).

8. Many DOCSIS cable modems allow viewing the diagnostics page on the computer connected to the modem. To access the diagnostics page on modems that support this function, type <http://192.168.100.1> in the browser address or URL window.




Figure 10: Cable Modem Diagnostics Screen Showing Downstream Equalized SNR (RxMER)

RxMER Measurement in a Digital Receiver

Before further discussing RxMER, we consider how a digital receiver is implemented, and how RxMER is measured. Figure 11 is a generalized block diagram of a digital QAM receiver. The receiver may reside in the CMTS, in which case it receives time-division multiple access (TDMA) or S-CDMA upstream bursts; or it may reside in a cable modem or set-top box (STB), in which case it receives a continuous stream of downstream digital data. The RF signal from the cable plant enters at the left of the diagram, and is processed by analog and digital front-end components that perform tuning, automatic gain control, channel selection, analog-to-digital conversion, and related functions. The square-root Nyquist filter has a response “matched” to the symbol or S-CDMA chip (a “chip” is a bit in the pseudorandom spreading code used in S-CDMA, as explained later). An adaptive equalizer compensates for channel response effects, including group delay variation, amplitude slope or ripple, and microreflections. An ingress canceller is normally included in a CMTS burst receiver to remove in-channel narrowband interference. Acquisition and tracking loops provide estimates of frequency, phase, and symbol timing, allowing the receiver to lock to the incoming signal. In the CMTS burst receiver, preamble symbols are used as a reference to aid in the acquisition and tracking of each upstream burst. In the case of S-CDMA, the chips are despread. The received QAM symbol, or soft decision, is passed to the slicer, which selects

the nearest ideal symbol, or hard decision, from the QAM constellation. The decisions are passed to the Trellis decoder, descrambler, deinterleaver, Reed-Solomon (RS) FEC decoder and MPEG deframer, and on to the MAC layer, which assembles and outputs received packets to the user.




Figure 11: Block Diagram of Generalized Digital QAM Receiver, Showing Computation of Receive MER

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:

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




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):

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

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:

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

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.

has the job of tracking out (removing) the low-frequency phase noise. In the DOCSIS Downstream RF Interface (DRFI) specification, phase noise below 50 kHz is separated out from the TxMER requirement because the phase noise is assumed to be largely removed by the receiver carrier loop. In general, a narrow (for example, 5 kHz) carrier loop will allow more phase noise to degrade the RxMER, but will pass less thermal noise through to the RxMER measurement. Conversely, a wide (for example, 50 kHz) carrier loop will allow less phase noise to degrade the RxMER, but will pass more thermal noise through to the RxMER measurement. The net effect is that varying the receiver carrier loop bandwidth will affect the RxMER, so the correct measurement carrier loop bandwidth must be carefully specified. The effects of phase noise become more critical as the modulation order increases.

- *Ingress cancellation effects:* Modern CMTS burst receivers have ingress cancellers, which remove in-channel narrowband interference entering the cable plant from the environment. After ingress cancellation, the upstream RxMER will be much higher than the input CNIR because the interference has been removed. The ingress canceller may add some white noise, depending on its implementation; the net result, however, is a dramatically improved RxMER.
- *Burst noise:* Short, strong bursts of noise may have unpredictable effects on the RxMER measurement. When burst noise hits, the RxMER will register a decrease, depending on the amount of averaging in the RxMER measurement and the burst properties of the noise. Downstream or upstream laser clipping—and its accompanying clipping distortion—tends to affect all frequency channels simultaneously. This is known as cross-compression, and usually degrades BER (the symptoms may be similar to burst noise), and if severe enough, RxMER. In some instances, MER reported by instruments such as QAM analyzers will change little, if at all, in the presence of short, infrequent, or weak burst noise, because the instrument averages the measurement over many symbols.
- *Collisions:* In a TDMA upstream, some time slots are “contention slots” in which multiple modems may randomly transmit. When two modems choose the same slot to transmit in, a collision occurs. At the receive slicer, the resulting signal looks like it has been hit by burst noise. RxMER measurements can be designed to exclude these noise contributions because contention slots are scheduled by the MAC and, therefore, predictable.
- *Multiuser nature of upstream:* In both TDMA and S-CDMA upstreams, the channel is shared by multiple users. In an ideal world with perfect ranging and equalization, all users would have equal RxMER. In reality, each upstream signal takes a unique route from modem to CMTS, so there will be slight differences in received power, transmit fidelity, etc., resulting in potentially different per-modem RxMER values. In DOCSIS, the upstream RxMER MIB measurement is defined as the average over a given number of valid bursts—that is, from many users, excluding contention slots.
- *Suboptimal modulation profiles:* Modulation profiles define how upstream information is transmitted from the cable modem to the CMTS. These profiles set modem transmit parameters—burst guard time, preamble, modulation type, FEC protection, and so forth—for request, initial maintenance, station maintenance, and short and long grant messages. Poorly configured modulation profiles can result in degraded upstream RxMER. For example, a preamble that is too short does not provide enough time for the tracking loops of the burst receiver to converge, resulting in lower RxMER. Interburst guard times that are not adequate result in interference from the end of one burst onto the beginning of the next, also degrading RxMER.

RxMER and DOCSIS Upstream Equalization

Let's briefly review DOCSIS upstream adaptive equalization. In DOCSIS 1.1 and later systems, each cable modem contains a preequalizer whose purpose is to predistort the transmit waveform so as to compensate for the upstream channel frequency response. In effect, the adaptive equalizer—the preequalizer of the modem—creates the equivalent of a digital filter with the opposite complex frequency response of the channel through which the upstream signal is transmitted. As defined in DOCSIS 2.0, the preequalizer is a 24-tap finite-impulse-response filter. When a cable modem is first powered on, it sends a specialized ranging burst (distinct from data traffic bursts) to the CMTS. The CMTS adapts its burst-receiver equalizer based on this “sounding” of the unique channel from the modem to the CMTS. The CMTS sends the equalizer coefficients back to the modem, which loads them into its preequalizer. Ideally, the preequalizer exactly corrects the response of the channel, and the data traffic bursts that the CMTS receives from that modem are free of linear distortion from then on. In a real system, the number and spacing of preequalizer taps limit the extent to which impairments can be compensated. The channel also varies with time. The modem sends periodic ranging bursts so that the CMTS can “tweak” the preequalizer coefficients of the modem in a tracking process. The preequalizer coefficients are updated by convolving them with the “residual” equalizer coefficients computed in the CMTS receiver on each ranging burst.

An important parameter in an equalizer is the span, defined as the (Number of taps – 1) times the spacing of the taps. If the channel response contains a significant echo (microreflection) that is further out in delay than the span of the equalizer, the equalizer cannot compensate for it. Hence, the span is a design parameter that depends on the channel model. The spacing of the preequalizer taps is defined in DOCSIS 2.0 as “T-spaced,” or symbol-spaced. At 5.12 Msym/sec, 24 taps gives a span of $(24 - 1)/5.12 = 4.5$ microseconds, which we may compare to the DRFI specification maximum assumed upstream micro-reflection or single echo parameter of >1.0 microsecond at -30 dBc.

Some CMTSs provide equalized RxMER measurements: The burst receiver adapts its equalizer on each data traffic burst as it is received, whether or not preequalization is in use. Other CMTSs provide unequalized RxMER measurements: The burst receiver does not adapt its equalizer on each data traffic burst as it is received. If preequalization is in use, there will be little difference between the RxMER measurements from these two types of CMTSs because the signal is already compensated when it arrives at the CMTS and there is little equalization left to do in the burst receiver. However, if preequalization is not in use, as, for example, with older DOCSIS 1.0 modems, then the RxMER measurements from CMTSs that perform receive equalization will generally be higher than the RxMER measurements from CMTSs that do not perform receive equalization. One might notice in this instance that longer packets show a higher RxMER measurement than shorter packets because the receive equalizer has more time to converge to its steady-state coefficient values on a long packet. Using a longer preamble may also help the receive equalizer to adapt more completely on each burst.




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:

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

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:

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

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.

Error vector magnitude is normally expressed as a linear measurement in percent, and MER is normally expressed as a logarithmic measurement in decibels. Why use EVM instead of MER to characterize a data signal? Many data engineers are familiar with EVM, and for some, linear measurements are easier to work with than logarithmic measurements. Error vector magnitude links directly with the constellation display, and there is a linear relationship between EVM and a constellation symbol point “cloud size” or “fuzziness.”

Maximum-to-Average Constellation Power Ratio and EVM/MER Conversion

Because EVM and MER are referenced differently, in order to relate EVM to MER, we must first compute the ratio of the peak constellation symbol power to the average constellation power. The peak constellation power is the squared magnitude of the outermost (corner) QAM symbol. Its formula for a square QAM constellation on an integer grid follows:

$$P_{peak} = 2(\sqrt{M} - 1)^2 \quad [\text{Eq. 25}]$$

where M is the number of points in the constellation (M = 4, 16, 64, 256, etc.) and the points are spaced by 2 on each axis. For example, for 16-QAM, the I and Q coordinates take on values from the set {−3, −1, 1, 3} and the peak power is $2(4 - 1)^2 = 3^2 + 3^2 = 18$. (Use of the integer grid is for illustration purposes only and does not imply any particular power normalization.)

The average constellation power (averaged equally over all symbols in the constellation) follows:¹²

$$P_{av} = \frac{2}{3}(M - 1) \quad [\text{Eq. 26}]$$

For example, for 16-QAM, the average constellation power is $(2/3)(16 - 1) = 10$. Note that this result happens to equal the power of one of the constellation points; the point (3,1) also has power $3^2 + 1^2 = 10$.

The maximum-to-average constellation power ratio (MTA) is, therefore, the unitless ratio:

$$MTA = \frac{P_{peak}}{P_{av}} = 3 \frac{\sqrt{M} - 1}{\sqrt{M} + 1} \quad [\text{Eq. 27}]$$

which approaches 3, or in decibels, $10\log(3) = 4.77$ dB, for very-high-order QAM. MTA (converted to dB) is tabulated in [Table 5](#), which contains entries for the standard square constellations as well double-square constellations. A double-square constellation is a subset consisting of half the points of the next-higher square constellation, arranged like the black squares on a checkerboard, and contains the same peak and average values as the next-higher square constellation. DOCSIS uses 64-QAM and 256-QAM square constellations for downstream transmission, and specifies both square and double-square constellations from QPSK to 128-QAM for upstream transmission.

12. Simon, Hinedi, and Lindsey, *Digital Communication Techniques*, equation 10.25, page 628.

Table 5: MTA Ratio for Square and Double-Square QAM Constellations

Constellation (DS = Double-Square)	MTA Ratio for Constellation Symbols (dB)
QPSK and BPSK	0
16-QAM and 8-QAM-DS	2.55
64-QAM and 32-QAM-DS	3.68
256-QAM and 128-QAM-DS	4.23
1024-QAM and 512-QAM-DS	4.50
Limit for infinite QAM	4.77

We can now convert from MER to EVM using the formula:

$$EVM_{\%} = 10 \times 10^{-(MER_{dB} + MTA_{dB})/20} \quad [Eq. 28]$$

where:

- $EVM_{\%}$ is error vector magnitude (percent).
- MER_{dB} is modulation error ratio (dB).
- MTA_{dB} is maximum-to-average constellation ratio (dB).

MTA Versus Peak-to-Average Ratio of an RF Signal

It is important not to confuse MTA with the peak-to-average ratio (PAR) of the actual transmitted signal. MTA accounts only for the distribution of the ideal QAM constellation symbols. Because of the subsequent spreading, filtering, and modulation processes that operate on the symbols, the effective PAR of a single modulated RF carrier will typically lie in the range of 6 to 13 dB or more. (The effect of filtering was illustrated previously by the long tails in the distribution in part (c) of [Figure 9 on page 21](#).) The actual PAR value depends on the modulation, excess bandwidth, and whether preequalization or S-CDMA spreading is in use.¹³ The PAR of a combined signal containing multiple carriers (such as the aggregate upstream or downstream signal on the cable plant) can become very large. Fortunately, the peaks occur very seldom, and the aggregate signal can often be treated like a random Gaussian signal.

MER and EVM Equipment and Example Measurements

Specialized test equipment such as a vector signal analyzer is generally needed to measure downstream or upstream EVM in a cable network, although some QAM analyzers support its measurement because the equipment must incorporate a digital QAM receiver in order to demodulate the signal to its complex baseband symbol constellation. [Figure 16 on page 36](#) shows a downstream 256-QAM digitally modulated signal whose EVM is 0.9 percent (circled). This value is representative of what might be seen at a headend or hub site, or at the downstream output of a node.

13. HEYS Professional Services' Francis Edgington has measured practical PAR values in the 6.3 dB to 7.3 dB range for 64-QAM signals and 6.5 dB to 7.5 dB range for 256-QAM signals.



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 signal¹⁴ 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.




Figure 17: Unimpaired 16-QAM Digitally Modulated Signal; the Unequalized MER Is 27.5 dB. (Courtesy Filtronic-Sigtek)




Figure 18: Impaired 16-QAM Digitally Modulated Signal; the Unequalized MER Is 19 dB (Courtesy Filtronic-Sigtek)

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)




Figure 20: Upstream Unequalized MER for the Digitally Modulated Signal in Figure 19 Is Only 21 dB, and the EVM Is a Relatively High 6.4 Percent. (Courtesy Sunrise Telecom)

Figure 19 on page 38 and Figure 20 emphasize the SNR (RxMER) versus CNR confusion that often occurs in cable systems. The cable operator's NOC might find that the CMTS reported upstream SNR (RxMER) is low (~21 dB in this instance), while a spectrum analyzer shows good CNR (~36 dB) and no apparent problems. The fact that unequalized 16-QAM would not work on this particular upstream—and QPSK was fine—indicates something is amiss.

TDMA and S-CDMA Effects on Upstream RxMER Measurement

Beginning with DOCSIS 2.0, both TDMA and S-CDMA modes are included in the upstream. In S-CDMA mode, multiple cable modems can transmit at the same time, while being separated by orthogonal spreading codes. Because various numbers of codes can be transmitted or quiet in a given S-CDMA frame, care must also be taken to properly normalize the signal measurement to avoid errors in S-CDMA MER measurement. In TDMA mode, upstream transmissions have a burst nature: Multiple cable modems share the same upstream frequency channel through the dynamic assignment of time slots. The effects of TDMA and S-CDMA on MER measurement are discussed in the following sections.

TDMA Burst Effects on SNR Measurement

Because a TDMA signal is bursting on and off, care must be taken to measure the signal power only when the signal is on. If an RF spectrum analyzer is used to measure CNR, the duty factor must be accounted for. The TDMA duty factor can be estimated and normalized out of the CNR measurement. For example, if the channel is active 90 percent of the time, the factor $10\log(0.9) = -0.46$ dB results; a correction of

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.

Each S-CDMA symbol is stretched 128 times longer by the spreading process. For example, consider a 3.2 MHz-wide channel (modulation rate 2.56 MHz). An S-CDMA symbol is made up of 128 chips, with each chip duration $1/2.56 \text{ MHz} = 0.39 \text{ microsecond}$ (the same duration a TDMA symbol would have in the same width channel). The S-CDMA symbol duration is $0.39 \text{ microsecond} \times 128 = 50 \text{ microseconds}$. This duration was designed to be longer than most bursts of noise that occur in the upstream, explaining why S-CDMA is resilient to impulse or burst noise.

Orthogonality of S-CDMA Codes

The S-CDMA codes possess the property of *orthogonality*, meaning that each despreader output ideally depends only on its assigned code, and not on what is happening on the other codes. In effect, each code acts like an independent communications channel, with its own noise component and “per-code” E_S/N_0 . The E_S/N_0 in one code does not depend (to a first approximation, assuming perfect orthogonality) on whether the other codes are even being transmitted. This situation is depicted in Figure 22. In a real system, perfect orthogonality is never achieved because of imperfect equalization; phase, frequency, and timing errors; and other implementation effects, which result in intercode interference (ICI).




Figure 22: Conceptually, S-CDMA Can Be Thought of as Having an Independent Upstream Channel for Each Code, Each with Its Own Noise Component.

Assuming ideal ranging, all codes are received at the CMTS at the same level, with the total received signal power divided equally among the 128 codes. The white-noise floor from the channel is also divided into 128 equal parts by the despanders, which function as a bank of 128 filters. Because both the signal and noise are reduced by the same factor of $1/128$, the E_S/N_0 in each code is the same as the overall E_S/N_0 of the channel. Figure 23 shows a received S-CDMA constellation after the despanders, with 25 dB E_S/N_0 or RxMER.




Figure 23: S-CDMA 16-QAM Constellation with 25 dB RxMER

Divergence of CNR and MER in S-CDMA

Because the codes are effectively independent, as shown previously in [Figure 22](#), turning some codes off reduces the total signal power in the channel but does not affect the E_S/N_0 on the other codes, meaning the CNR seen on a spectrum analyzer will appear to fluctuate as some codes are transmitted and others are not, but the E_S/N_0 per code will remain constant. This effect is seen in [Figure 24](#). In the upper trace, all 128 active codes are transmitted. The CNR, E_S/N_0 per code, and RxMER per code are all approximately equal to 25 dB. In the lower trace, all but 32 of the codes have been turned off, while keeping the received power per code unchanged. The CNR is reduced by 6 dB, but the E_S/N_0 per code and received RxMER remain approximately unchanged at 25 dB. Thus, in S-CDMA, the CNR measured on a spectrum analyzer can vary dynamically, and is a valid indication of the E_S/N_0 or RxMER per code only when all active codes are being transmitted in a given frame.




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.

We can also view this concept by building from the bottom up. Consider an S-CDMA transmission of a single code, consisting of 128 chips.¹⁵ When we despread it, we have a spreading gain of $10\log(128) = 21$ dB. Using the same parameters as in Figure 24, the single-code signal has a CNR of only 4 dB, but an E_s/N_0 of 25 dB because of its processing gain. If we add multiple orthogonal codes, the total power goes up proportionally with the number of codes, as does the CNR, but the E_s/N_0 stays at 25 dB because each code is independent of the others.

Maximum Scheduled Codes and MER in S-CDMA

S-CDMA spreading gain can be applied in more than one way. In the example of Figure 24, when the number of transmitted codes was reduced from 128 to 32, the power per code was held constant, thereby keeping the E_s/N_0 per code and RxMER constant while reducing the total transmitted power, and, hence, the CNR, by 6 dB. Conversely, we could use the 6 dB spreading gain differently: We could reduce the number of transmitted codes from 128 to 32, while keeping the transmitted power and CNR constant. This setup would cause the E_s/N_0 per code and RxMER to increase by 6 dB. In fact, we could apply the 6 dB spreading gain to transmit power in any amount not restricted to the previous two cases; for example, reducing the transmitted power and CNR by 3 dB, while increasing the E_s/N_0 per code and RxMER by 3 dB.

DOCSIS provides a mechanism, called Maximum Scheduled Codes (MSC), to give this spreading-gain boost to a cable modem that is otherwise unable to produce the required RxMER at the CMTS receiver, perhaps because of increased cable and splitter attenuation in a multidwelling complex where the modem is located. A cable modem using MSC will have reduced throughput because it is using fewer codes, but will have increased transmitted power per code. If the spreading gains given to the modems using MSC equal their increased upstream path attenuation and all other parameters are equal, the CMTS burst receiver will see equal received power per code and RxMER for MSC and non-MSC modems. The overall upstream throughput is also unaffected by the application of MSC, because the codes not being used by an MSC cable modem can be scheduled by the MAC for use by other cable modems on the plant.

Chip MER in S-CDMA

Referring back to Figure 21 on page 40, the input to the RF channel at each modulation interval consists of the sum of up to 128 S-CDMA chips. Each chip is the equivalent of a TDMA symbol in duration; for example, $T_{\text{chip}} = 1/5.12$ MHz. Normally, a chip is not specified individually, but only in the context of a full sequence of 128 chips, which represents one QAM symbol before spreading or after despreading. That is, although individual chips may have low MER in the presence of AWGN, the signal is restored to high MER by the despreading process. However, DOCSIS also levies a TxMER requirement on each individual chip, called MER_{chip} . Its purpose is to guarantee that fidelity is maintained at low transmit signal levels, for example, when only two of the 128 spreading codes are being transmitted. This requirement is intended to place a limit on the noise-funneling effect that occurs when many modems transmit simultaneously and their transmitted spurious noises sum at the CMTS receiver. The DOCSIS physical layer specification provides details of this requirement.

15. Illustrative example only; in DOCSIS, the minimum number of spreading codes permitted is two.

Conclusion

This paper investigated the common ways that signal-to-noise ratio is defined and measured in digital transmission over cable systems. It shows that RF and baseband measurements of CNR and SNR have to be treated differently, and that digitally modulated signals require their own precise definition and measurement of SNR.

CMTSs can report what has for many years been called upstream SNR, a parameter that is often confused with CNR. In reality, the upstream SNR of a CMTS is equalized MER or, in some cases, unequalized MER—specifically, as defined in the DOCSIS MIB, RxMER. Cable modems and most digital STBs also can report an SNR value. This value is not CNR, but is equalized downstream RxMER. Likewise, QAM analyzers and similar test equipment used by the cable industry can report MER values for downstream—and, in some cases, upstream—digitally modulated signals. These values, too, are not CNR, but are RxMER, as discussed in this paper. Most QAM analyzers report equalized MER measurements, although some also can provide unequalized MER measurements (or the equivalent of unequalized measurements). RxMER provides a “baseline” indication of signal quality, but must be interpreted carefully to gain the full value of this important measurement.

Summary of CNR, SNR, and MER

SNR is a general signal-to-noise-ratio measurement, and can refer to measurements performed at RF or baseband. The cable industry has long used SNR to refer to a baseband measurement; for example, baseband video or audio. CNR is the RF carrier-to-noise ratio seen on a spectrum analyzer. E_s/N_0 is the most common measurement of the quality of a digitally modulated signal. RxMER (often called receive SNR estimate) is the modulation error ratio of the demodulated digital constellation, a "bottom-line" measurement that includes transmitter imperfections, plant distortions, thermal noise, and receiver imperfections. CINR is the carrier-to-interference-plus-noise ratio, which includes narrowband ingress in addition to white noise.

SNR, CNR, E_s/N_0 , RxMER, and CINR are equal in an ideal system with no impairments other than AWGN, and with full traffic loading. When impairments occur in a real system, the differences among these measurements provide clues about what the problem is.

Downstream RxMER is measured by the cable modem after the adaptive equalizer, which cleans up most linear distortions. Upstream RxMER is reported by the CMTS burst receiver, usually also after equalization, although some implementations may report an unequalized value. Most QAM analyzers provide equalized MER measurements.

If upstream RxMER is low but the spectrum analyzer shows high CNR, there is most likely a plant or equipment problem to be addressed. Likely candidates are micro-reflections from impedance mismatches or high group delay variation from the roll-off of diplex filters. If DOCSIS preequalization does not cure the problem, then nonlinear impairments, such as phase noise, spurious, burst or impulse noise, etc., may be involved.

TDMA CNR measurements are complicated by the fact that a signal is not continuously present, but the noise is. The maximum hold function on the spectrum analyzer may help. However, maximum hold will display the signal power of the highest-power cable modem burst and highest noise, if the noise is fluctuating in time, typically resulting in errors in the average CNR measurement of 1 to 2 dB or more.

S-CDMA presents challenges for accurate E_s/N_0 measurement because the power of an S-CDMA burst depends on the number of codes used in the burst. If a small amount of data is transmitted, only a small number of codes will be required and the signal power will be low, and low CNR, which may erroneously be perceived as implying low E_s/N_0 , will result. An accurate measurement of CNR and E_s/N_0 from a spectrum analyzer can best be obtained when all possible codes are transmitted in a burst. Maximum hold on a spectrum analyzer can provide a reasonably accurate measurement as long as at least one burst is transmitted with all codes used during the measurement period.

CNR measurements accurate to within tenths of a dB are difficult and require knowledge of proper measurement bandwidths, conversion formulas, and compensation for instrument imperfections. If measurements are required only to within 1 to 2 dB or so, straightforward delta marker measurements will suffice as long as proper bandwidth compensations are made.

References

- Agilent Technologies, “Digital Modulation in Communications Systems – An Introduction”; Application Note 1298, ©2001.
- Agilent Technologies, “Spectrum Analyzer Measurements and Noise: Measuring Noise and Noise-like Digital Communications Signals with a Spectrum Analyzer”; Application Note 1303, ©2003.
- Agilent Technologies, “Using Error Vector Magnitude Measurements to Analyze and Troubleshoot Vector-Modulated Signals”; Product Note PN 89400-14, ©2000.
- Agilent Technologies, “Cable Television System Measurements Handbook – NTSC Systems”; Technical Note, 1994: <http://www.home.agilent.com/agilent/facet.jspx?to=80030.k.1&cc=US&lc=eng&sm=g&k=Cable+Television+System+Measurements+Handbook>.
- Ciciora, W., J. Farmer, D. Large, and M. Adams, *Modern Cable Television Technology*, 2nd Ed.; Morgan Kaufmann Publishers, 2004.
- Cooper, M., J. Finkelstein, and J. Moran, “Maximizing Throughput in the Return Path via DOCSIS 2.0.” In *Proceedings Manual and Collected Technical Papers*, SCTE Cable-Tec Expo 2006, Denver, Colo.
- Currivan, B., T. Kolze, J. Min, and G. Ungerboeck, “Physical-Layer Considerations for Advanced TDMA CATV Return Path.” In *Annual Review of Communications*, International Engineering Consortium, vol. 55, Chicago, 2002.
- Currivan, B., “Cable Modem Physical Layer Specification and Design.” In *Cable Modems: Current Technologies and Applications*, International Engineering Consortium, Chicago, 1999.
- DOCSIS 2.0 Radio Frequency Interface Specification: <http://www.cablemodem.com/specifications/specifications20.html>.
- DOCSIS 3.0 Radio Frequency Interface Specification: <http://www.cablemodem.com/specifications/specifications30.html>.
- DOCSIS Downstream RF Interface Specification: <http://www.cablemodem.com/specifications/d-rfi.html>.
- Downey, J. “Upstream FEC Errors and SNR as Ways to Ensure Data Quality and Throughput.” Cisco Systems (Document ID: 49780): http://www.cisco.com/en/US/tech/tk86/tk319/technologies_white_paper09186a0080231a71.shtml.
- Edgington, F. “Digital RF Measurements Show Return System Suitability for Increased Data Rate/Bandwidth.” In *Proceedings Manual and Collected Technical Papers*, SCTE Cable-Tec Expo 2004, Orlando, Fla.
- Freeman, R. L., *Telecommunications Transmission Handbook*, 2nd Ed., John Wiley & Sons, Inc., 1981, ISBN 0-471-08029-2.
- Howald, R. L., “QAM Bulks Up Once Again – Modulation to the Power of Ten.” In *Proceedings Manual and Collected Technical Papers*, SCTE Cable-Tec Expo 2002, San Antonio, Tex.: <http://broadband.motorola.com/ips/pdf/QAM.pdf>.
- Howard, D., B. Currivan, T. Kolze, J. Min, and H. Samueli, “Last Mile HFC Access.” In *Broadband Last Mile: Access Technologies for Multimedia Communications*, CRC Press, Boca Raton, Fla., 2005.

- Hranac, R. and M. Millet. "Upstream Power Measurements: Watts Up Doc?" *Communications Technology*, December 2000.
- Hranac, R., "Linear Distortions, Part 1." *Communications Technology*, July 2005.
- Hranac, R., "Linear Distortions, Part 2." *Communications Technology*, Aug. 2005.
- Hranac, R., "CNR Versus SNR." *Communications Technology*, Mar. 2003.
- Hranac, R., "Spectrum Analyzer CNR Versus CMTS SNR." *Communications Technology*, Sept. 2003.
- Hranac, R., "More on CMTS SNR." *Communications Technology*, Oct. 2003.
- Hranac, R., "CMTS Upstream SNR Revisited." *Communications Technology*, Apr. 2006.
- Hranac, R., "Cable Modem Upstream Signal Levels." *Communications Technology*, May 2006.
- Hranac, R., "More on Cable Modem Upstream Signal Levels." *Communications Technology*, June 2006.
- Kolze, T., "An Approach to Upstream HFC Channel Modeling and Physical-Layer Design." In *Cable Modems: Current Technologies and Applications*, International Engineering Consortium, Chicago, 1999.
- NCTA Recommended Practices for Measurements on Cable Television Systems, 3rd Edition*, National Cable & Telecommunications Association, Washington, D.C., 2002
- Nyquist, H., "Certain Topics in Telegraph Transmission Theory." Transactions of the A. I. E. E., pp. 617–644, Feb. 1928.
- Proakis, J., *Digital Communications*. McGraw-Hill, 3rd Ed., 1999.
- Simon, M., S. Hinedi, and W. Lindsey, *Digital Communication Techniques*. Prentice-Hall, Englewood Cliffs, N.J., 1994.
- Tektronix, "Measuring Noise in Video Systems." Application Note 25W-11148-0, 1997.
- Terreault, G., "Signal Impairments: A Closer Look at MER and BER." *Communications Technology*, Dec. 2002.
- Williams, T., "Upstream Data Impairments, Part 1." *Communications Technology*, Nov. 2003.
- Williams, T., "Upstream Data Impairments, Part 2." *Communications Technology*, Dec. 2004.

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

Broadcom®, the pulse logo, **Connecting everything®**, and the Connecting everything logo are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

Connecting
everything®



BROADCOM CORPORATION

5300 California Avenue

Irvine, CA 92617

© 2012 by BROADCOM CORPORATION. All rights reserved.

Phone: 949-926-5000

Fax: 949-926-5203

E-mail: info@broadcom.com

Web: www.broadcom.com

CMTS-WP101-R January 2012

All contents are Copyright © 2006–2011 Broadcom Corporation and Cisco Systems, Inc. All rights reserved.