Welcome to E-XFL.COM

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

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

What are Embedded - System On Chip (SoC)?

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

Details	
Product Status	Obsolete
Architecture	-
Core Processor	-
Flash Size	-
RAM Size	-
Peripherals	-
Connectivity	-
Speed	-
Primary Attributes	-
Operating Temperature	-
Package / Case	-
Supplier Device Package	-
Purchase URL	https://www.e-xfl.com/product-detail/broadcom/bcm33843mrkfsbg

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.

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 $N_{64\text{-}QAM}$ for the QAM signal with the equation:

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

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:

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

(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.

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.

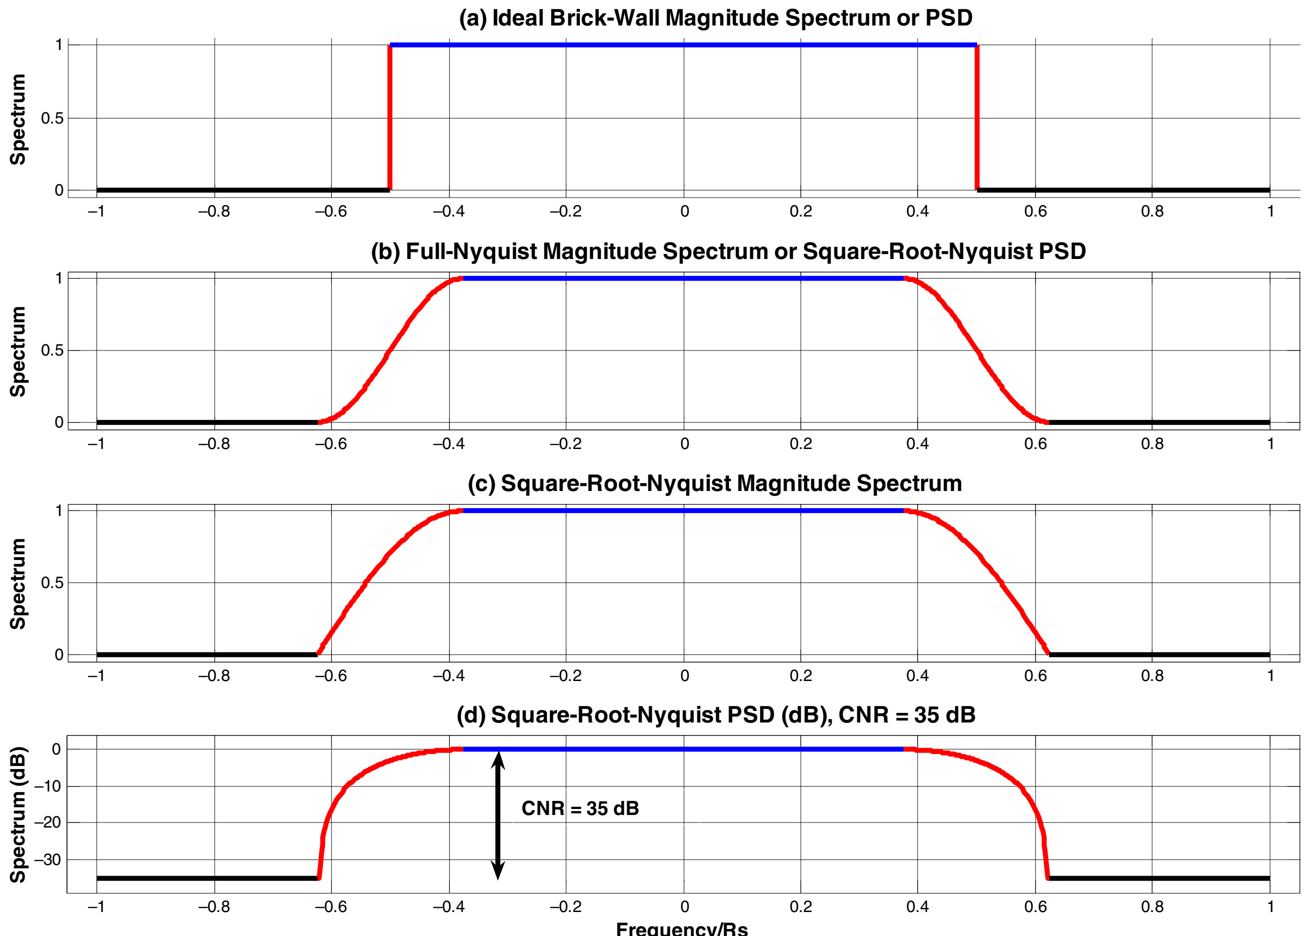

(a) Ideal Brick-Wall Spectrum
(b) Full Nyquist Spectrum
(c) Square-Root Nyquist Spectrum
(d) Square-Root Nyquist Spectrum in dB, with Noise Floor Giving CNR = 35 dB
The passband is shown in blue, transition bands in red, and stop bands in black.

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 $|H_1(f)|^2$ of the transmitted signal and of the square-root Nyquist filter, described next.

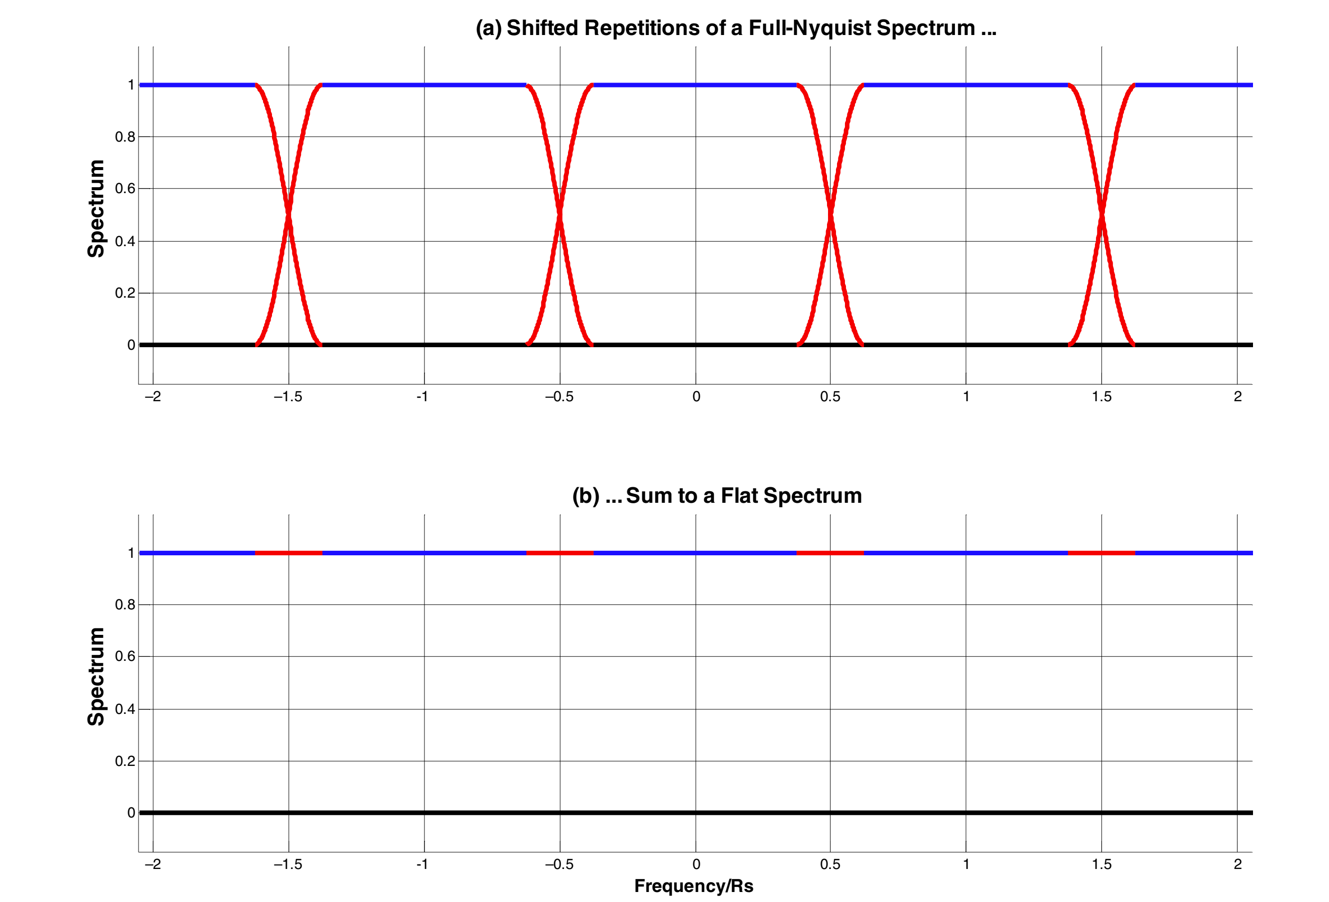

Figure 5: The Nyquist Property States That When Copies of the Spectrum Are Shifted by Multiples of the Symbol Rate and Added, the Result Is a Flat Spectrum, Which Results in Zero ISI.

In practice, the full Nyquist spectrum $H(f)$ of part (b) in Figure 4 is divided into two identical cascaded “square-root Nyquist” filters $H_1(f)$, 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 $|H_1(f)|$ is shown in part (c) of the figure. Again, because of the cascading property, the power spectrum $|H_1(f)|^2$ 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.

- 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.

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.

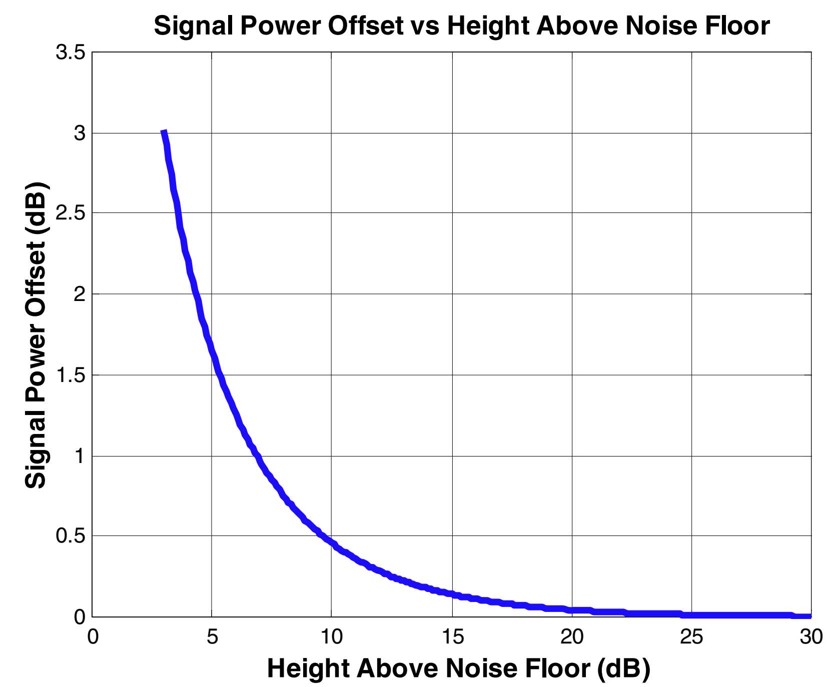

Figure 8: Signal Power Measurement Correction Relative to Noise Floor

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

$$\begin{aligned} &\text{haystack\_height\_dB} = 10\log(1 + 10^{\text{true\_CNR\_dB}/10}) \\ &= 10\log(1 + 10^{4/10}) \\ &= 5.45\text{ dB} \end{aligned} \quad \text{[Eq. 17]}$$

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

$$\begin{aligned} &\text{true\_CNR\_dB} = 10\log(10^{\text{haystack\_height\_dB}/10} - 1) \\ &= 10\log(10^{5.45/10} - 1) \\ &= 4\text{ dB} \end{aligned} \quad \text{[Eq. 18]}$$

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 of the constellation points, and the magnitude of the error between them would be zero. In a real-world receiver, subtracting the hard-decision vector from the soft-decision vector gives the error or noise vector at each symbol time. The implicit assumption is that a low symbol error rate exists – that is, very few decisions are incorrect, ensuring that the "decision-directed" error vector from the nearest symbol nearly always equals the true error vector from the correct reference symbol.

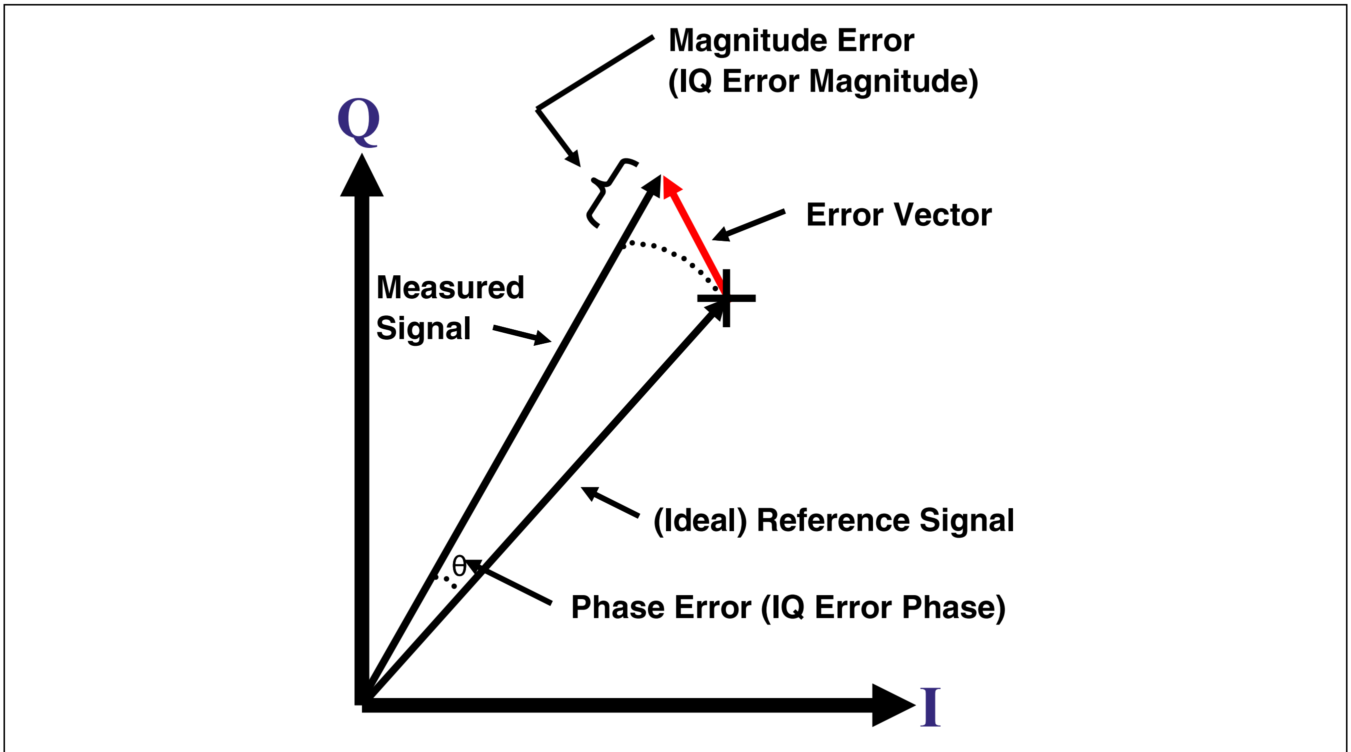

Figure 12: The Error Vector Is the Difference Between the Measured Signal (Soft Decision) and the Reference or Target Signal (Hard Decision). (Source: Hewlett-Packard)

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

More About Modulation Error Ratio

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

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

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

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

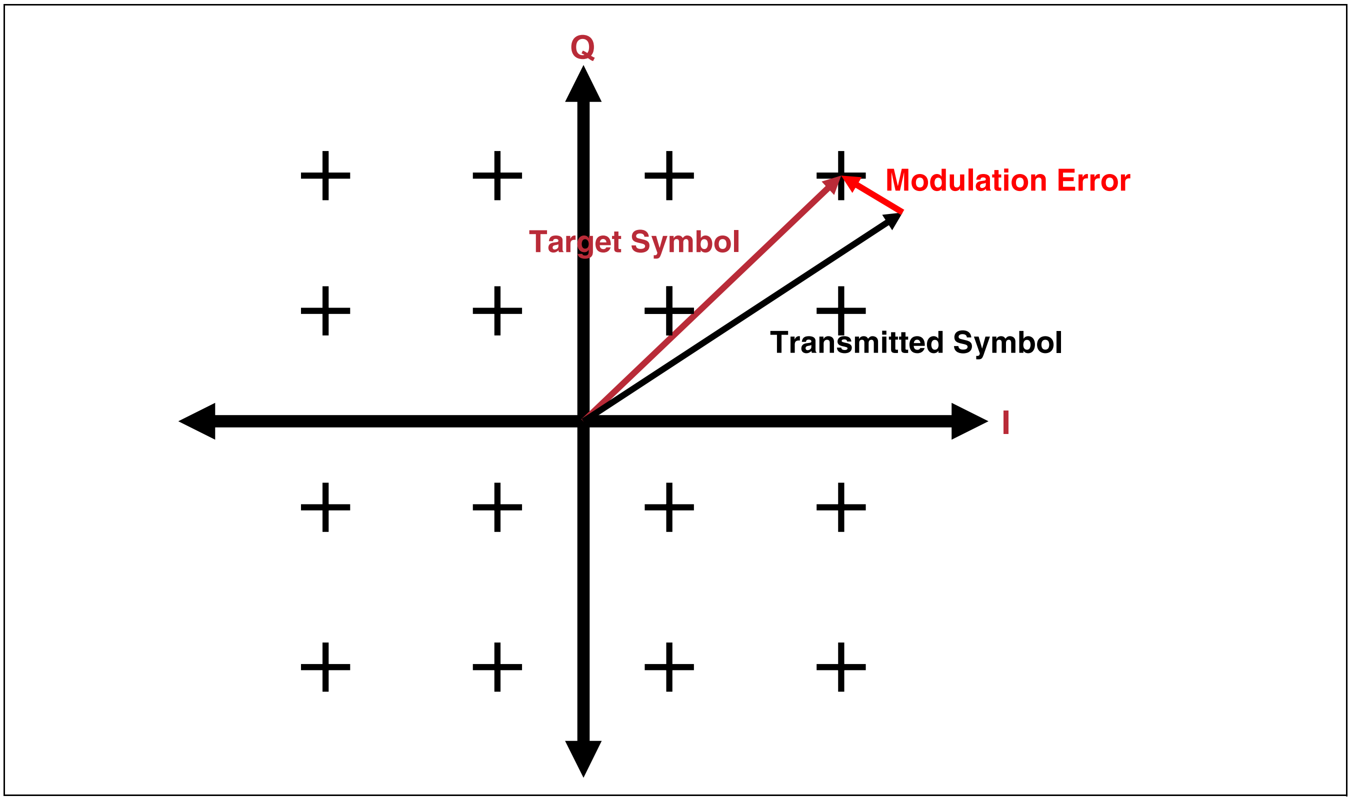

Figure 13: Modulation Error Is a Measure of Modulation Quality. (Source: Hewlett-Packard)

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.

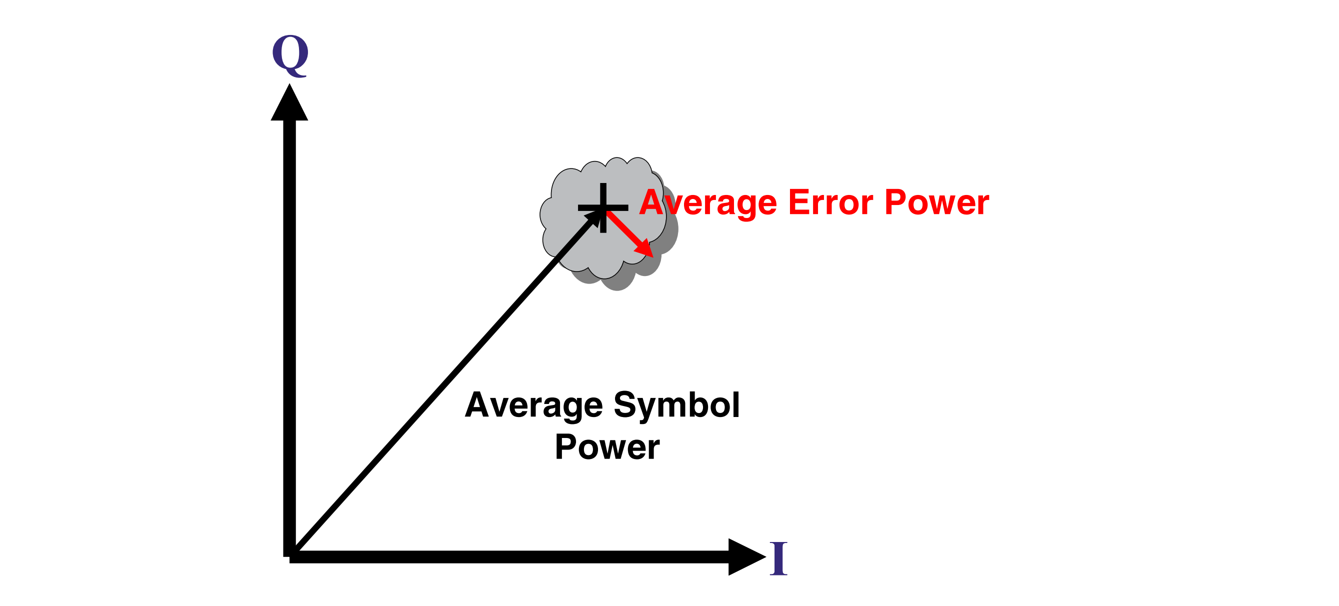

Figure 14: Modulation Error Ratio Is the Ratio of Average Symbol Power to Average Error Power. (Source: Hewlett-Packard)

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

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

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

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

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.

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.

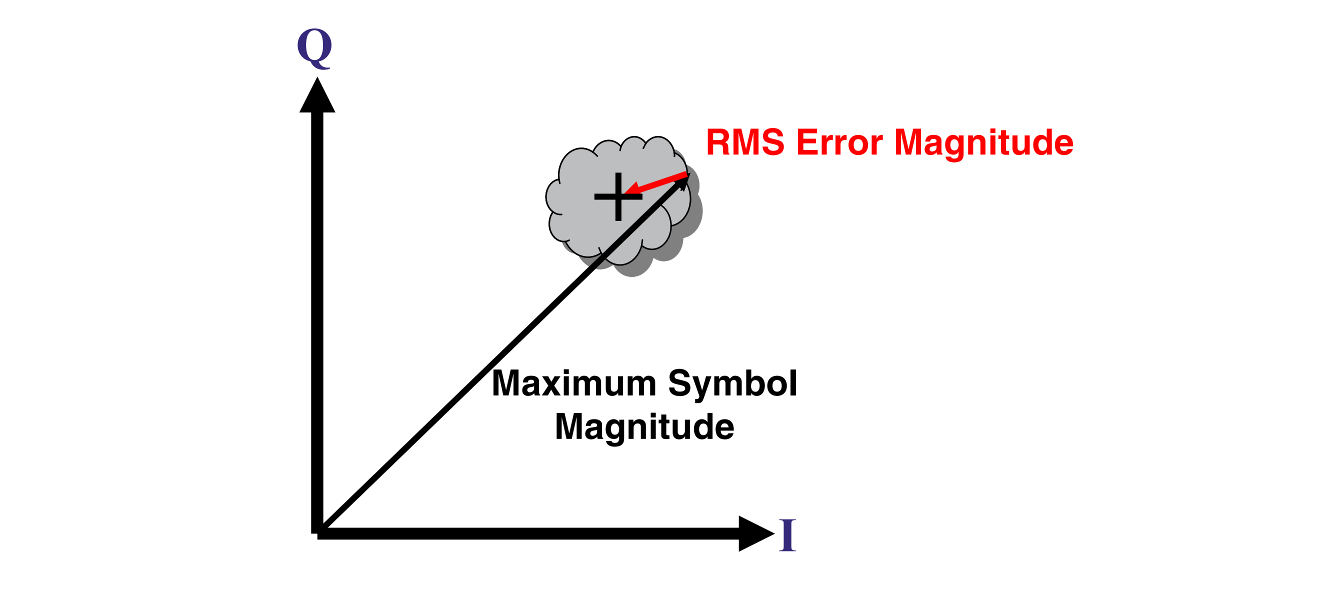

Figure 15: Error Vector Magnitude Is the Ratio (in Percent) of RMS Error Magnitude to Maximum Symbol Magnitude. (Source: Hewlett-Packard)

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:

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

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

EVM Versus MER

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

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

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

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:

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

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]

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

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:

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

which approaches 3, or in decibels, 10log(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.

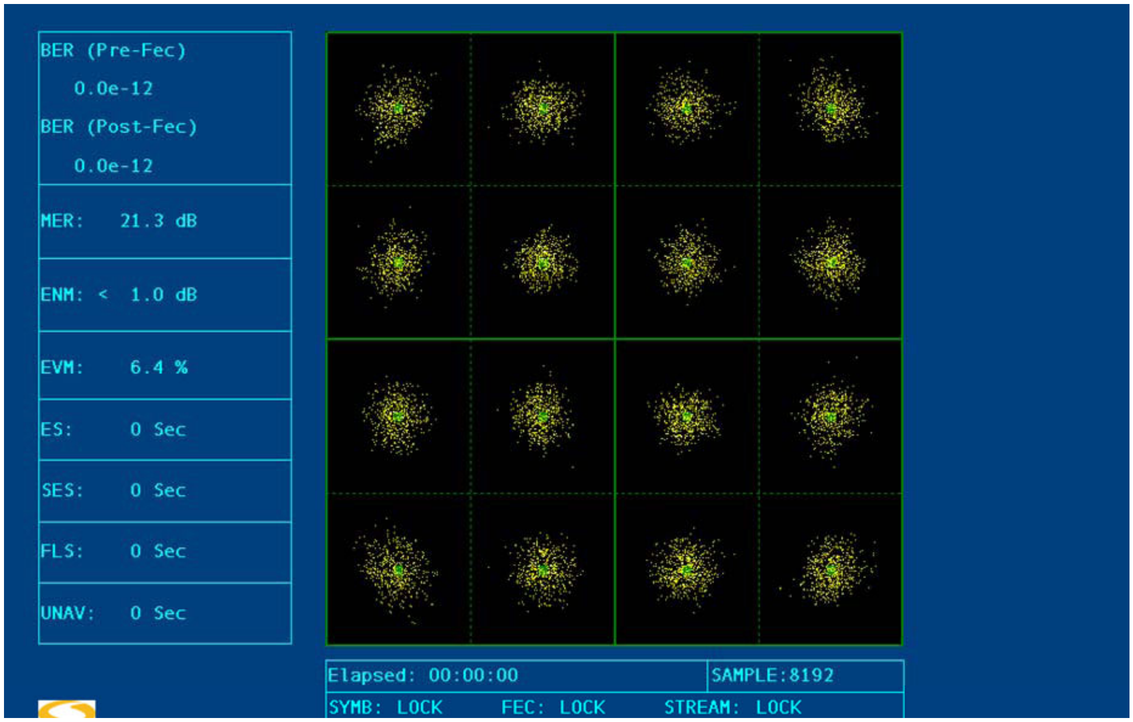

Figure 20: Upstream Unequalized MER for the Digitally Modulated Signal in Figure 19 Is Only 21 dB, and the EVM Is a Relatively High 6.4 Percent. (Courtesy Sunrise Telecom)

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

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 MHz = 0.39 microsecond (the same duration a TDMA symbol would have in the same width channel). The S-CDMA symbol duration is 0.39 microsecond x 128 = 50 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).

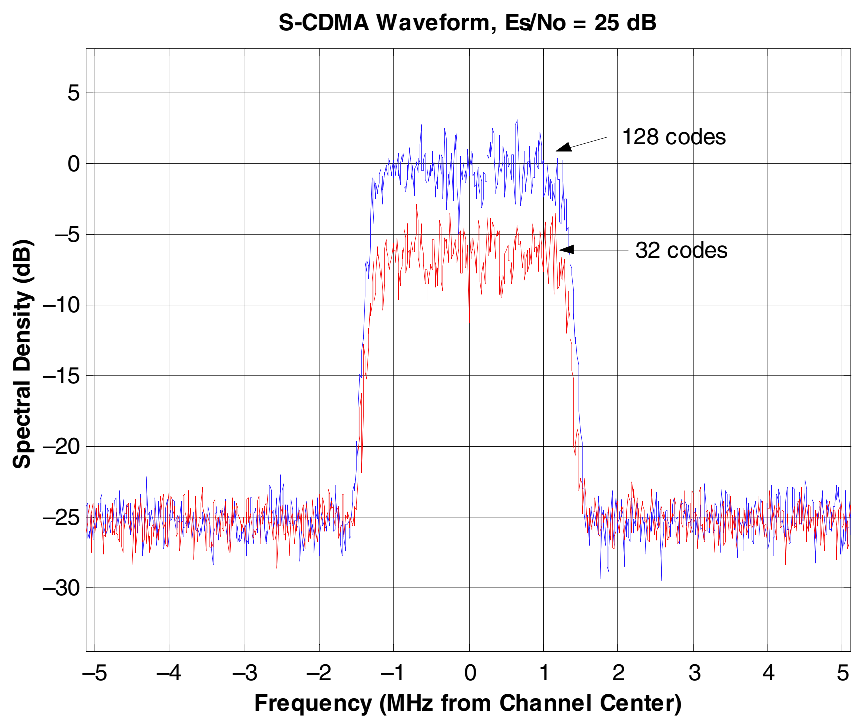

Figure 22: Conceptually, S-CDMA Can Be Thought of as Having an Independent Upstream Channel for Each Code, Each with Its Own Noise Component.

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 despreaders, 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 despreaders, with 25 dB E_S/N_0 or RxMER.

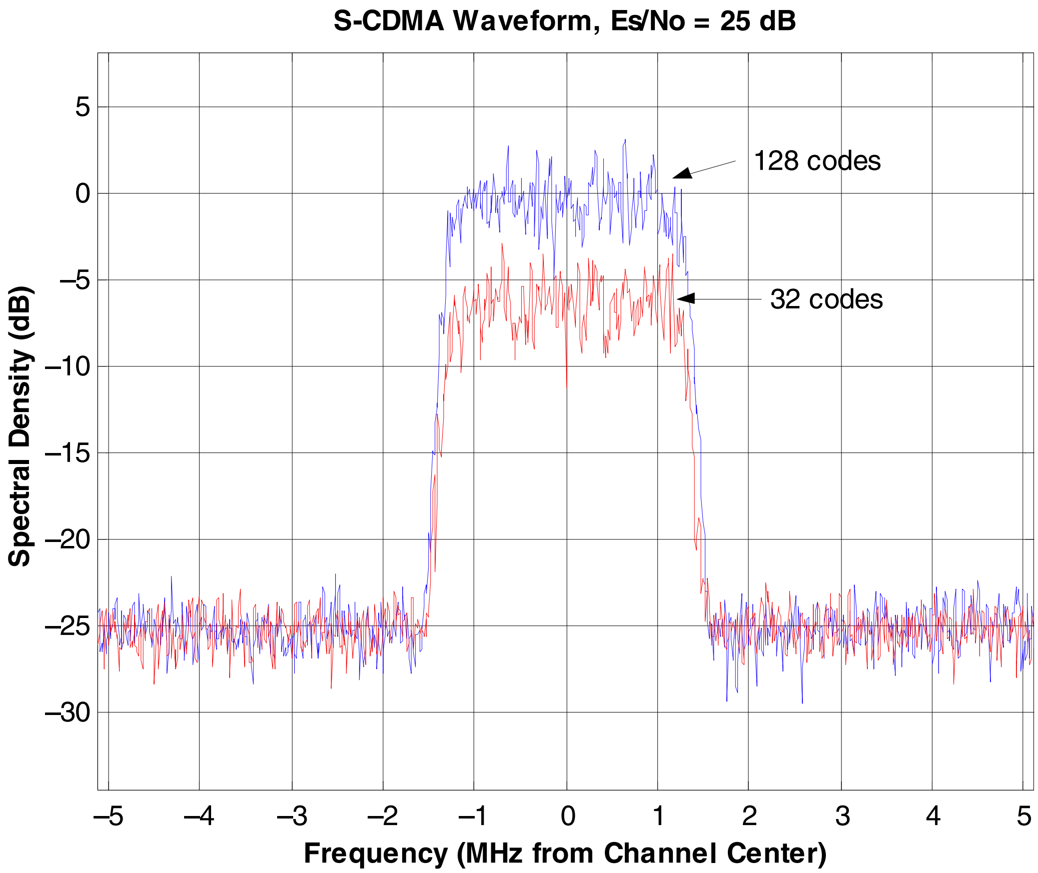

Upper trace: all 128 active codes are transmitted; CNR = 25 dB.
Lower trace: only 32 of the 128 active codes are transmitted; CNR reduces to 19 dB, but E_S/N_0 per code is still 25 dB because of 6 dB spreading gain.

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

Spreading Gain in S-CDMA

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

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 10log(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_{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.

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