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/bcm33843eukfsbg

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.

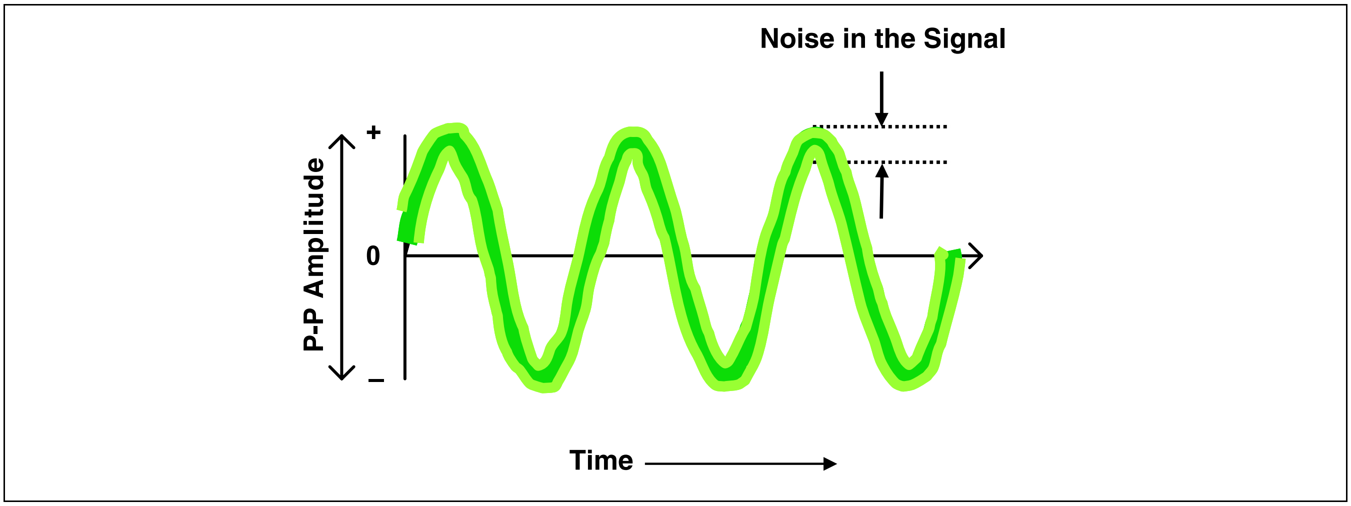

Figure 2: Baseband SNR Measurement

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 carrier[3] 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:

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

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

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

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

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

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

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

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

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.

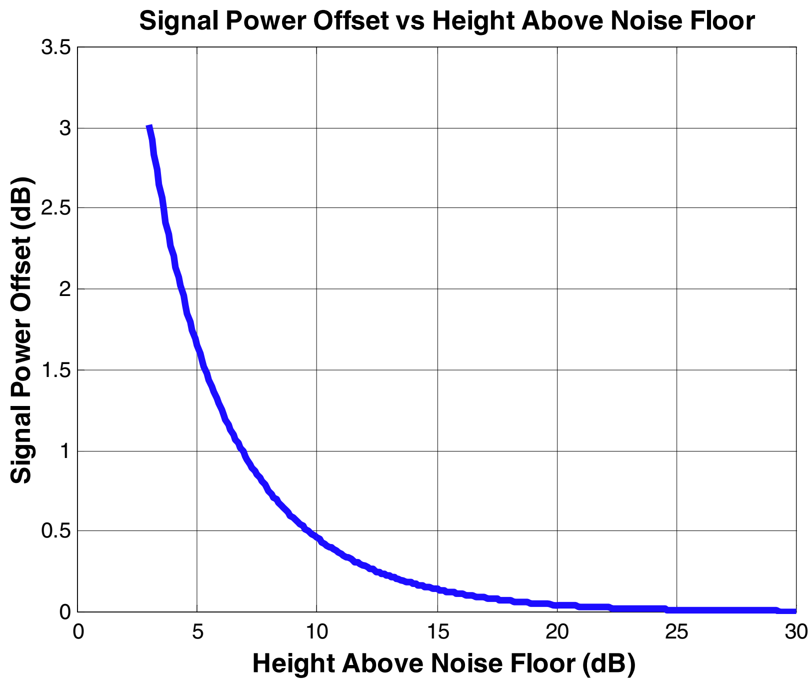

Figure 8: Signal Power Measurement Correction Relative to Noise Floor

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

haystack_height_dB = $10\log(1 + 10^{\text{true\_CNR\_dB}/10})$
= $10\log(1 + 10^{4/10})$ [Eq. 17]
= 5.45 dB

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

true_CNR_dB = $10\log(10^{\text{haystack\_height\_dB}/10} - 1)$
= $10\log(10^{5.45/10} - 1)$ [Eq. 18]
= 4 dB

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

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 ÷ 5.36 MHz symbol rate = 1.12.)

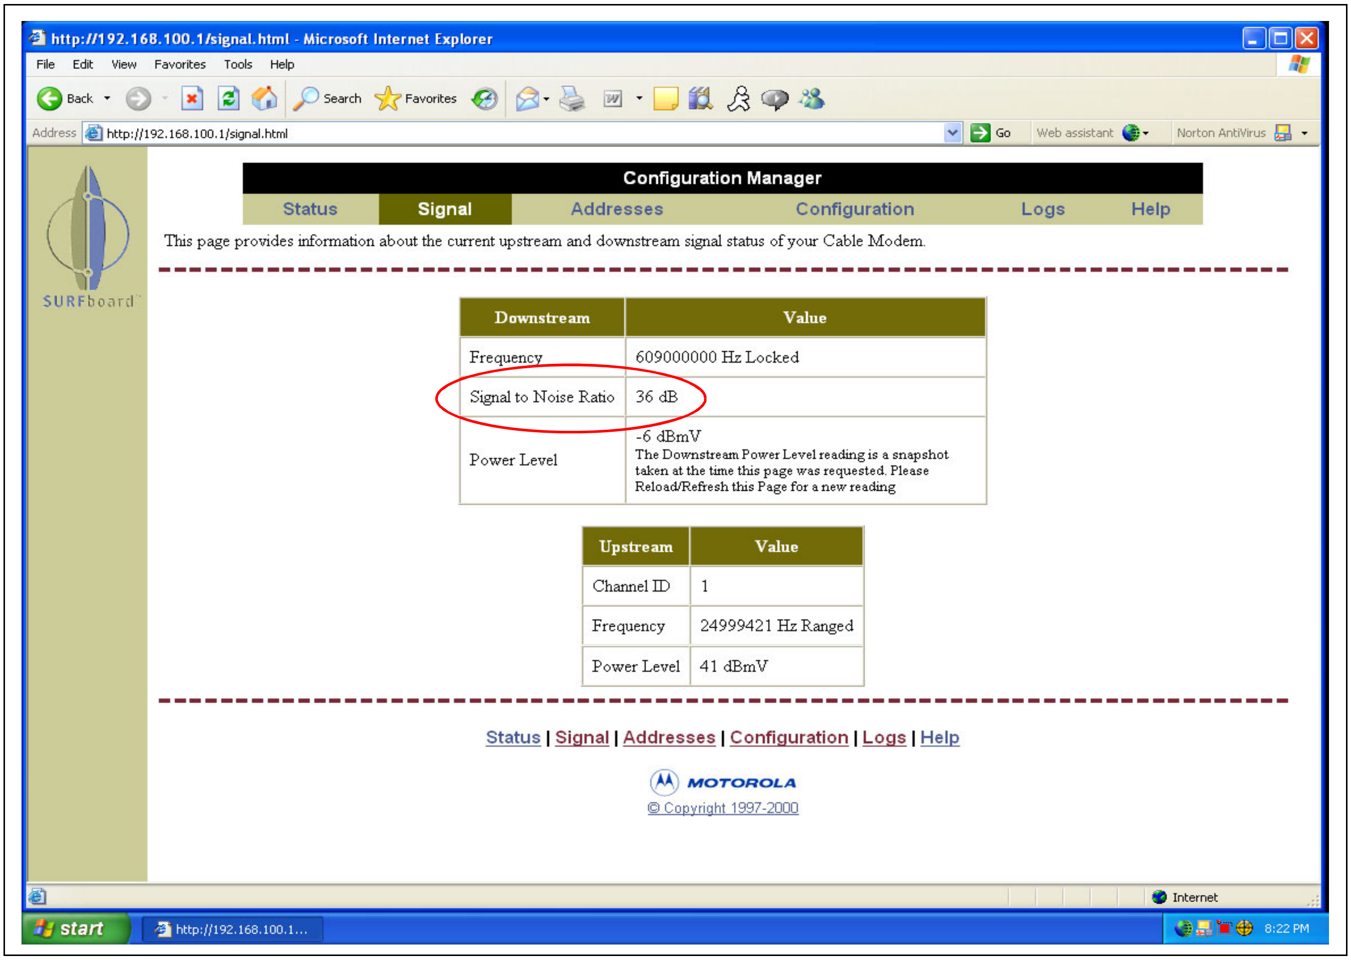

Figure 10: Cable Modem Diagnostics Screen Showing Downstream Equalized SNR (RxMER)

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.

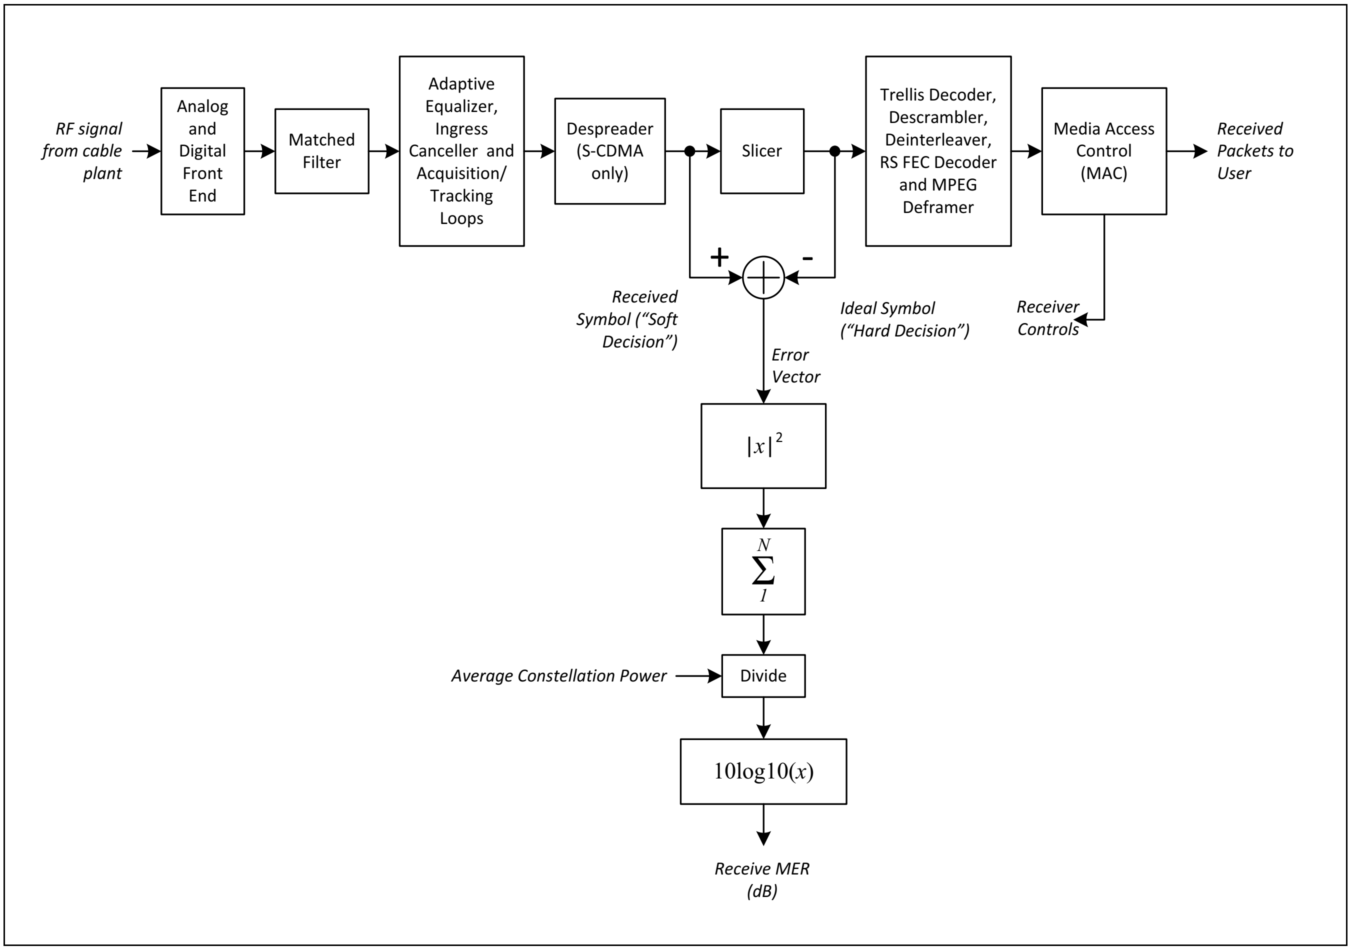

Figure 11: Block Diagram of Generalized Digital QAM Receiver, Showing Computation of Receive MER

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.

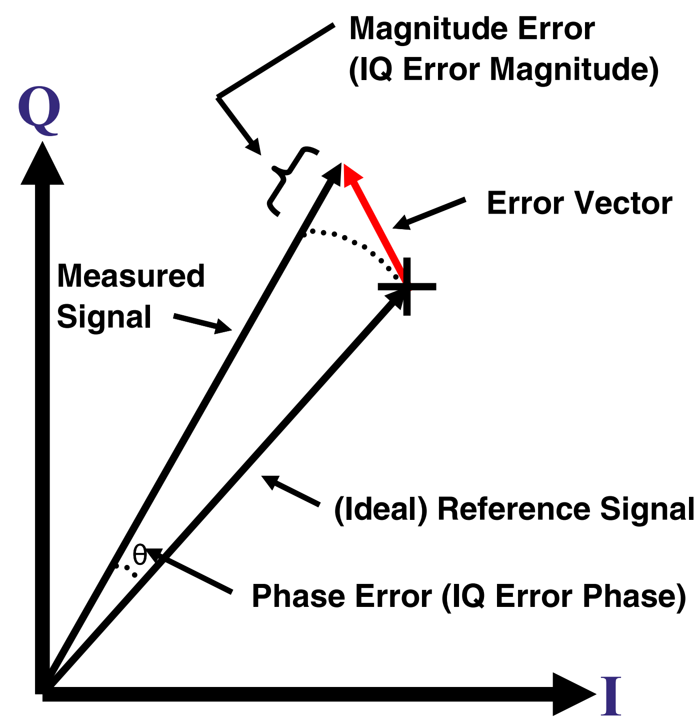

Figure 12: The Error Vector Is the Difference Between the Measured Signal (Soft Decision) and the Reference or Target Signal (Hard Decision). (Source: Hewlett-Packard)

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:

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

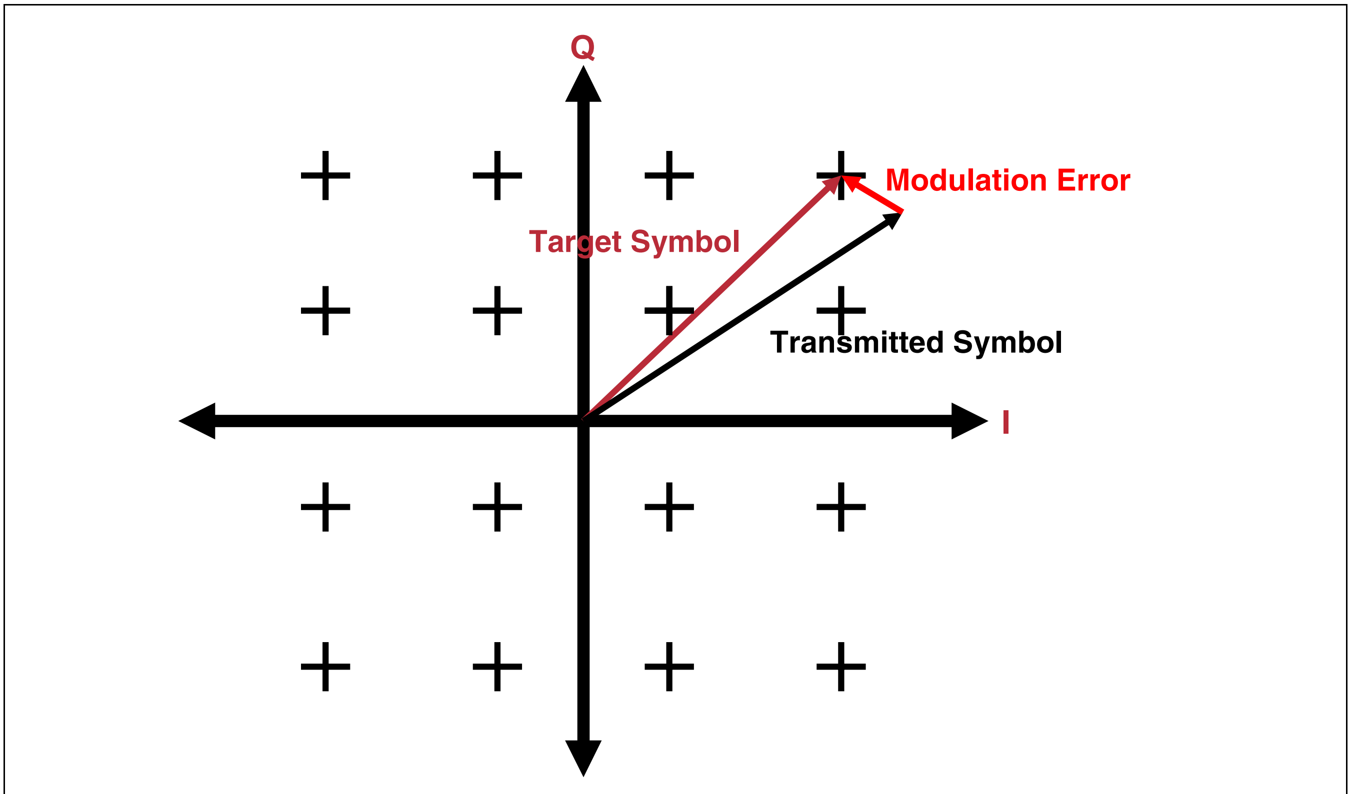

Figure 13: Modulation Error Is a Measure of Modulation Quality. (Source: Hewlett-Packard)

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

MER(dB) = 10log(Average symbol power ÷ Average error power) [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.

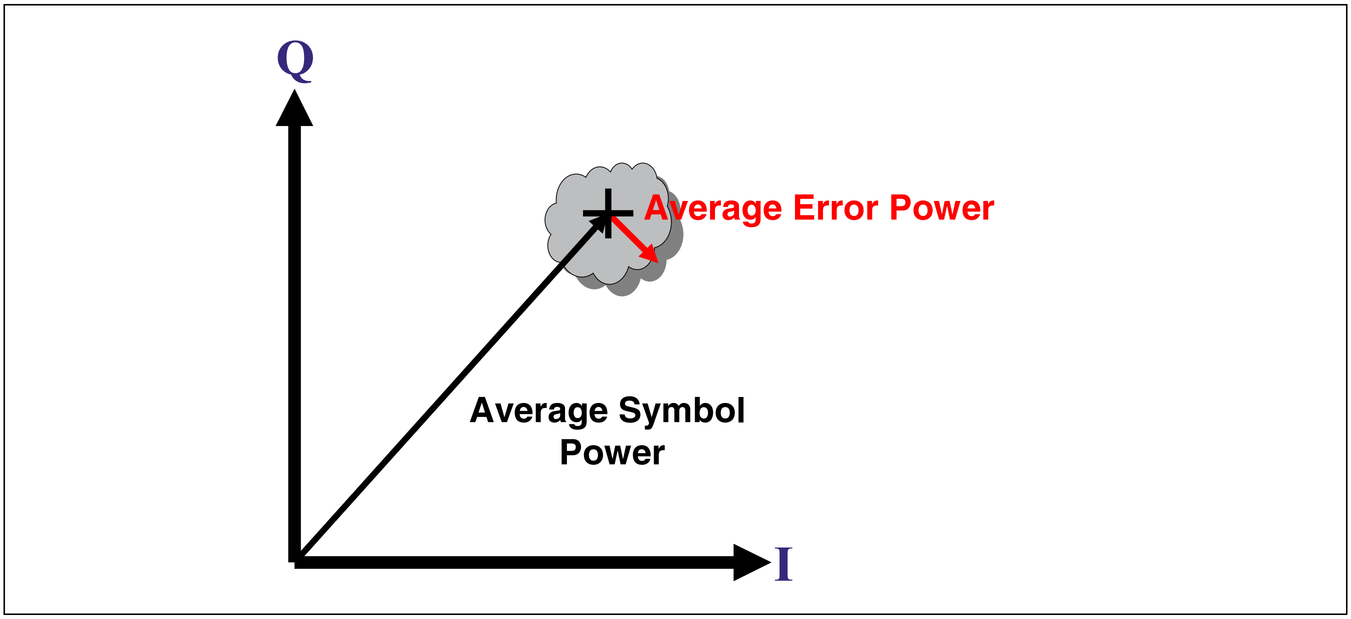

Figure 14: Modulation Error Ratio Is the Ratio of Average Symbol Power to Average Error Power. (Source: Hewlett-Packard)

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

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

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.

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.

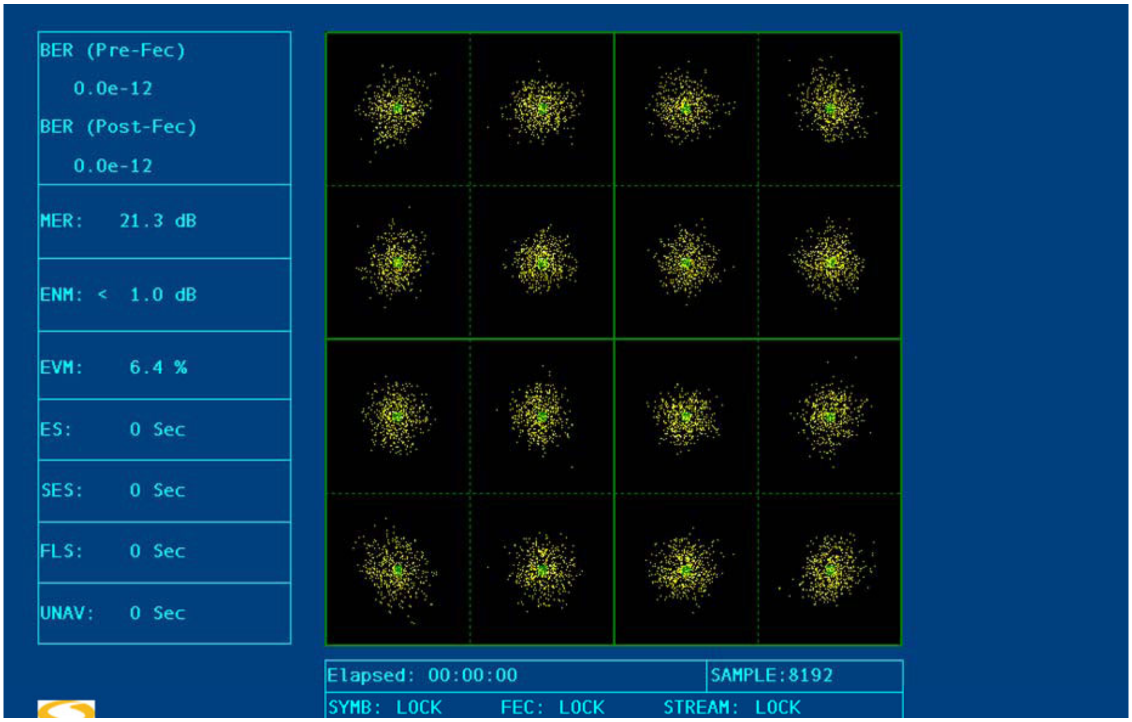

Figure 20: Upstream Unequalized MER for the Digitally Modulated Signal in Figure 19 Is Only 21 dB, and the EVM Is a Relatively High 6.4 Percent. (Courtesy Sunrise Telecom)

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 10log(0.9) = –0.46 dB results; a correction of

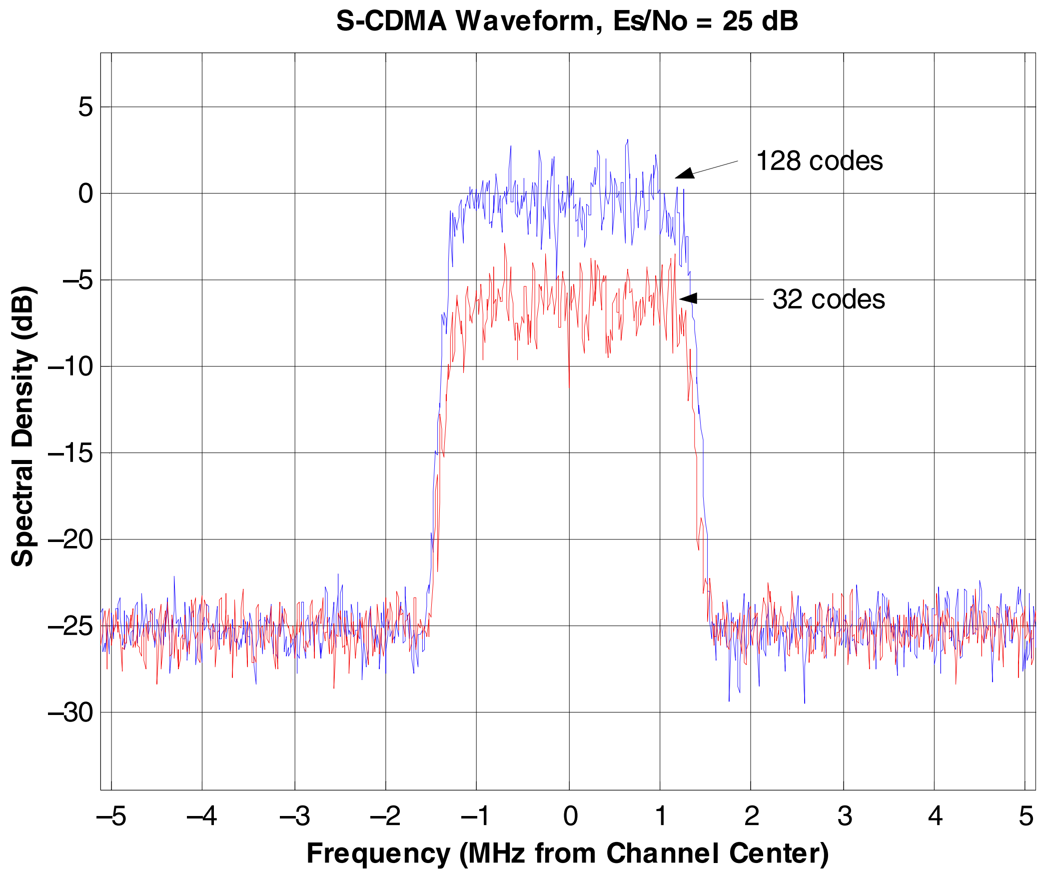

Upper trace: all 128 active codes are transmitted; CNR = 25 dB.
Lower trace: only 32 of the 128 active codes are transmitted; CNR reduces to 19 dB, but E_S/N_0 per code is still 25 dB because of 6 dB spreading gain.

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.

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.

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.

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