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/bcm33843gifsbg

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.

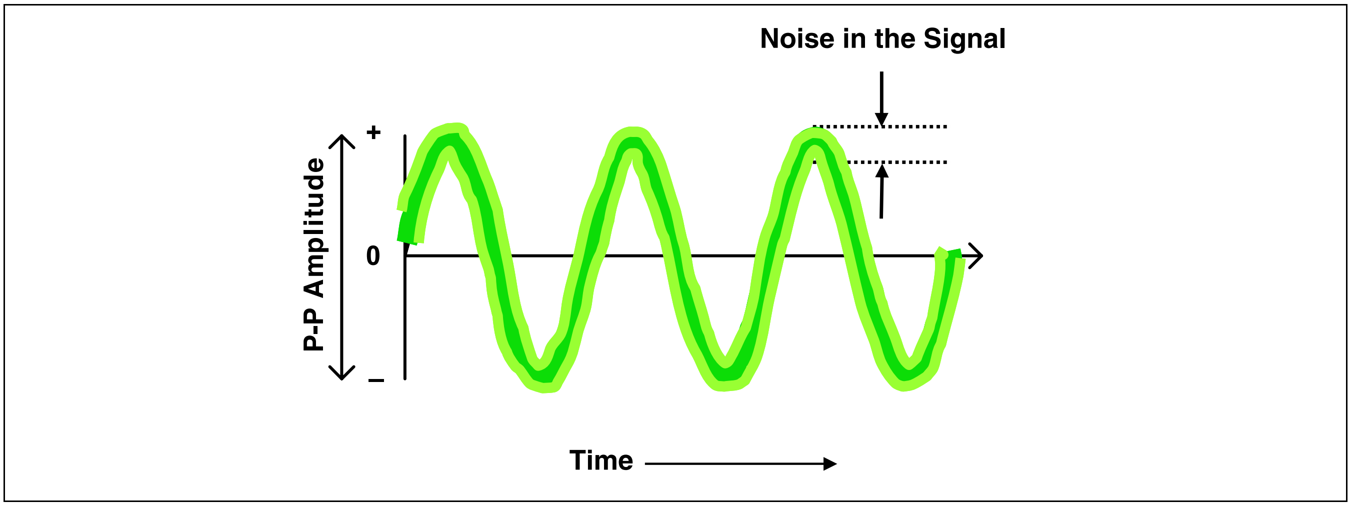

Figure 2: Baseband SNR Measurement

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:

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

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 $N_{64\text{-}QAM}$ 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.

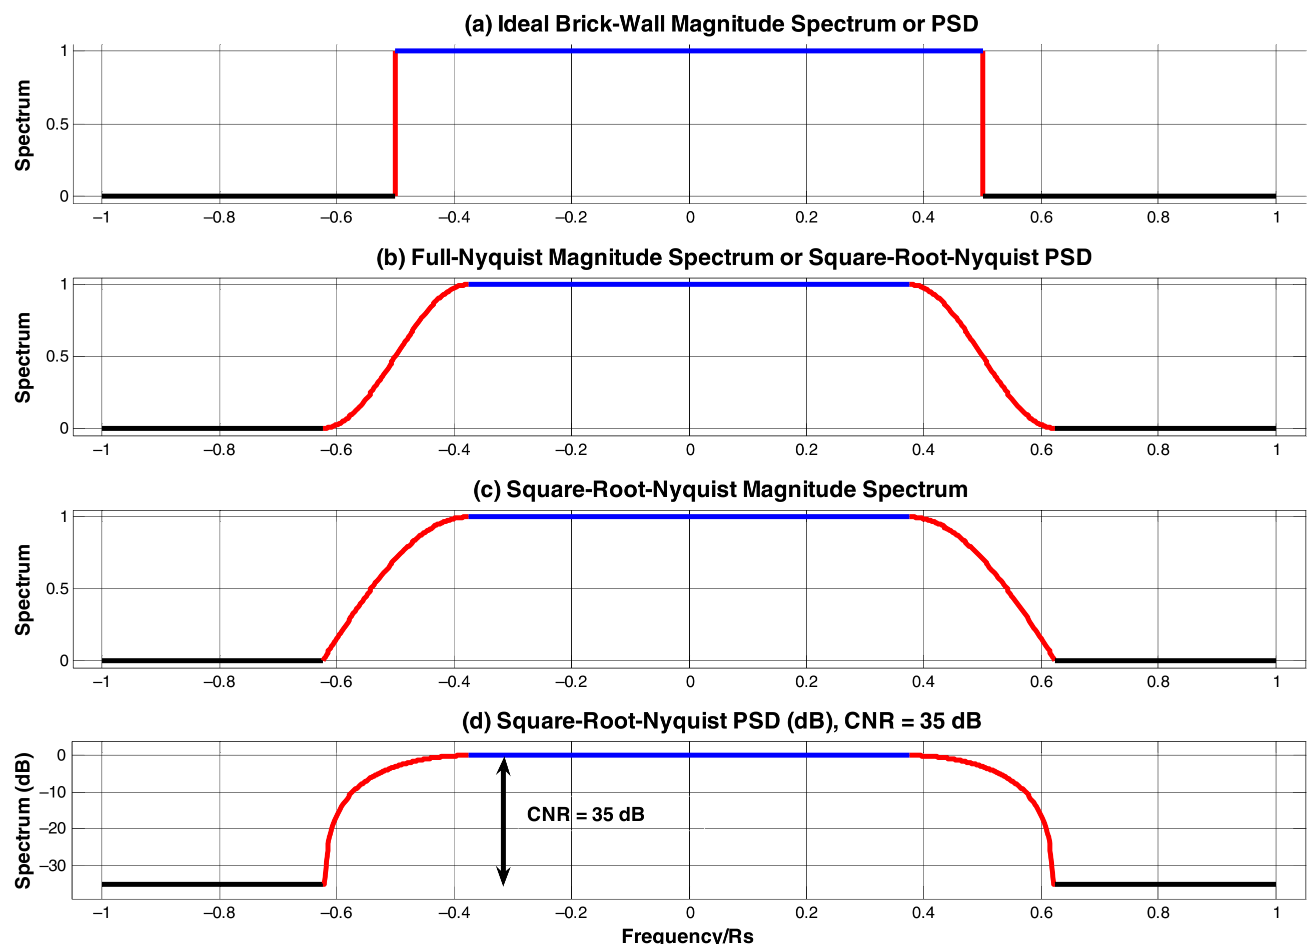

(a) Ideal Brick-Wall Spectrum
(b) Full Nyquist Spectrum
(c) Square-Root Nyquist Spectrum
(d) Square-Root Nyquist Spectrum in dB, with Noise Floor Giving CNR = 35 dB
The passband is shown in blue, transition bands in red, and stop bands in black.

Figure 4: Ideal QAM Spectrum Plots

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 CNR = 35 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 CNR = 35 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 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.

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

$$\text{signal\_pwr\_offset\_dB} = 10\log\left(1 + \frac{1}{10^{\text{haystack\_height\_dB}/10} - 1}\right) \qquad \text{[Eq. 14]}$$

$$\text{true\_signal\_pwr\_dBmV} = 10\log\left(\frac{10^{\text{haystack\_top\_dBmV}/10}}{1 + \frac{1}{10^{\text{haystack\_height\_dB}/10} - 1}}\right) \qquad \text{[Eq. 15]}$$

$$\text{true\_CNR\_dB} = 10^{\text{haystack\_height\_dB}/10} - 1 \qquad \text{[Eq. 16]}$$

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.

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.

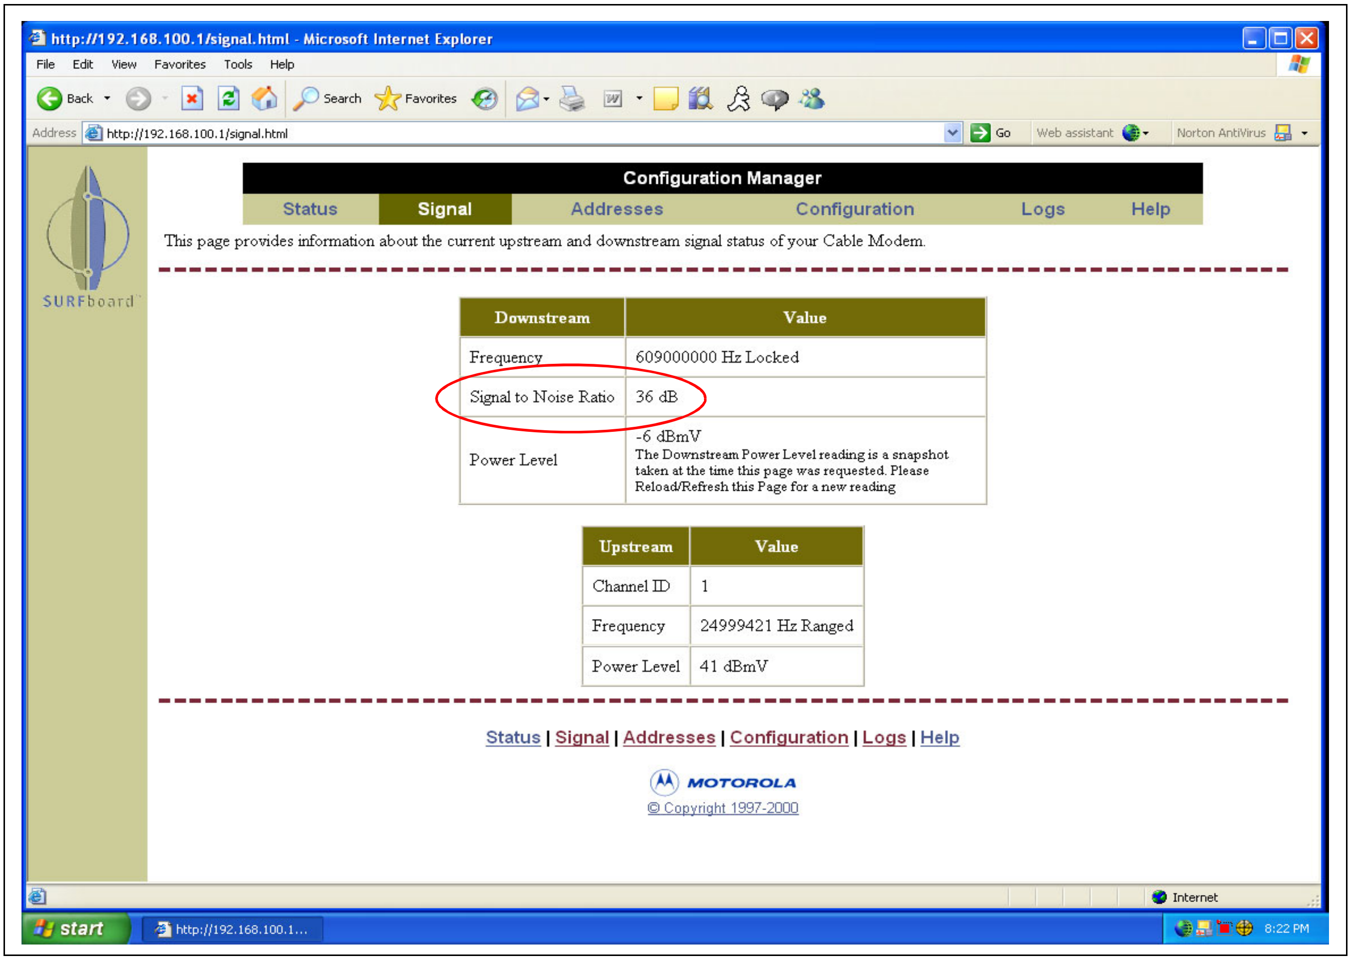

Figure 10: Cable Modem Diagnostics Screen Showing Downstream Equalized SNR (RxMER)

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

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.

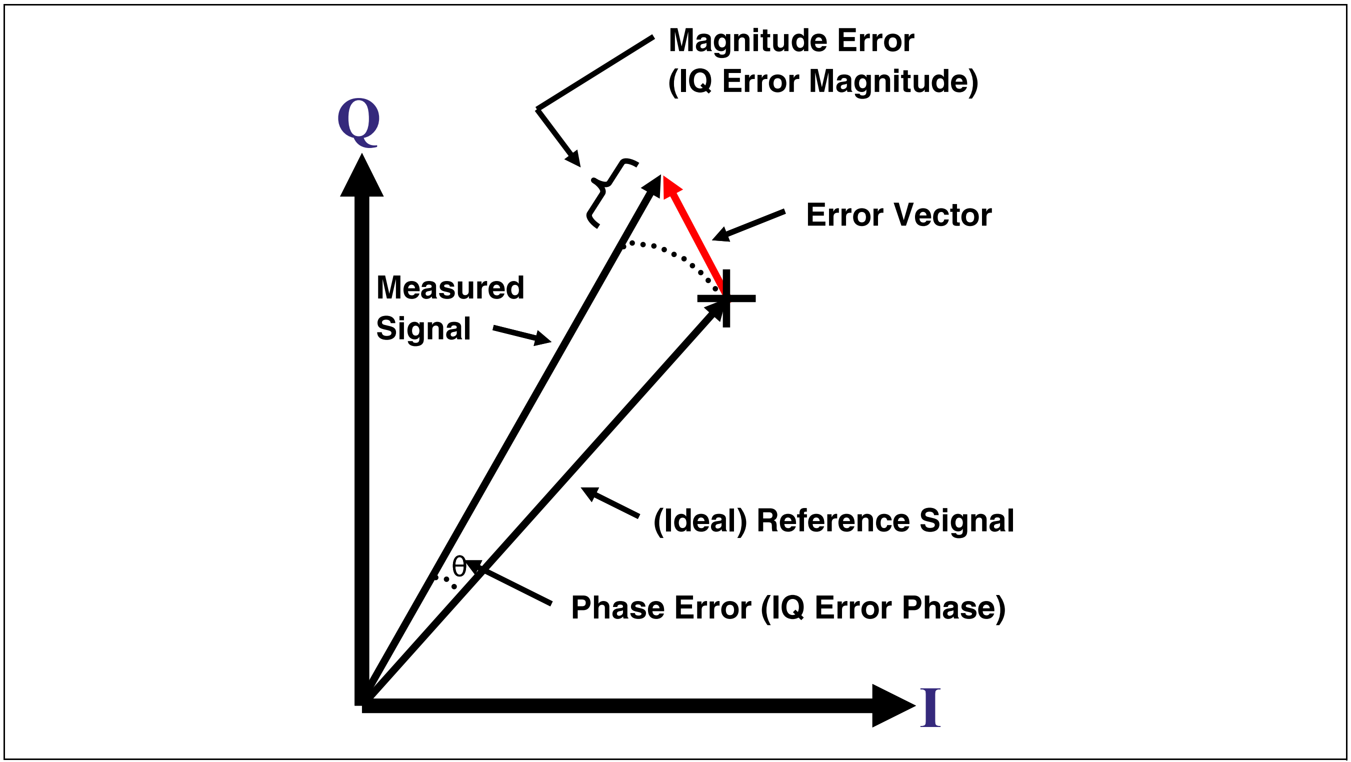

Figure 12: The Error Vector Is the Difference Between the Measured Signal (Soft Decision) and the Reference or Target Signal (Hard Decision). (Source: Hewlett-Packard)

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

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.

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.

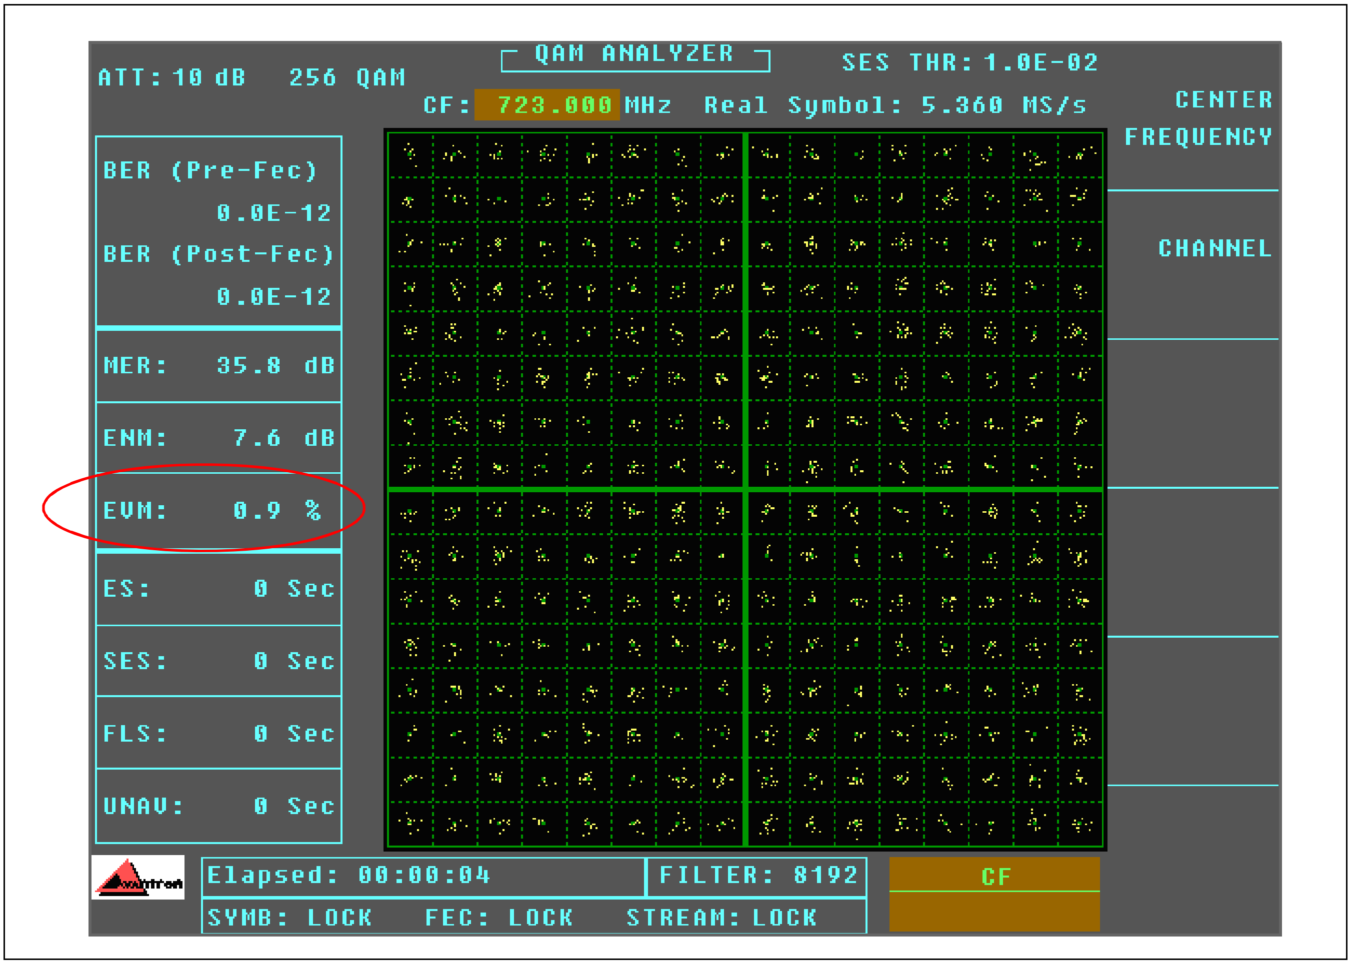

Figure 16: The EVM of This Downstream 256-QAM Digitally Modulated Signal Is 0.9 Percent. (Courtesy of Sunrise Telecom)

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

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

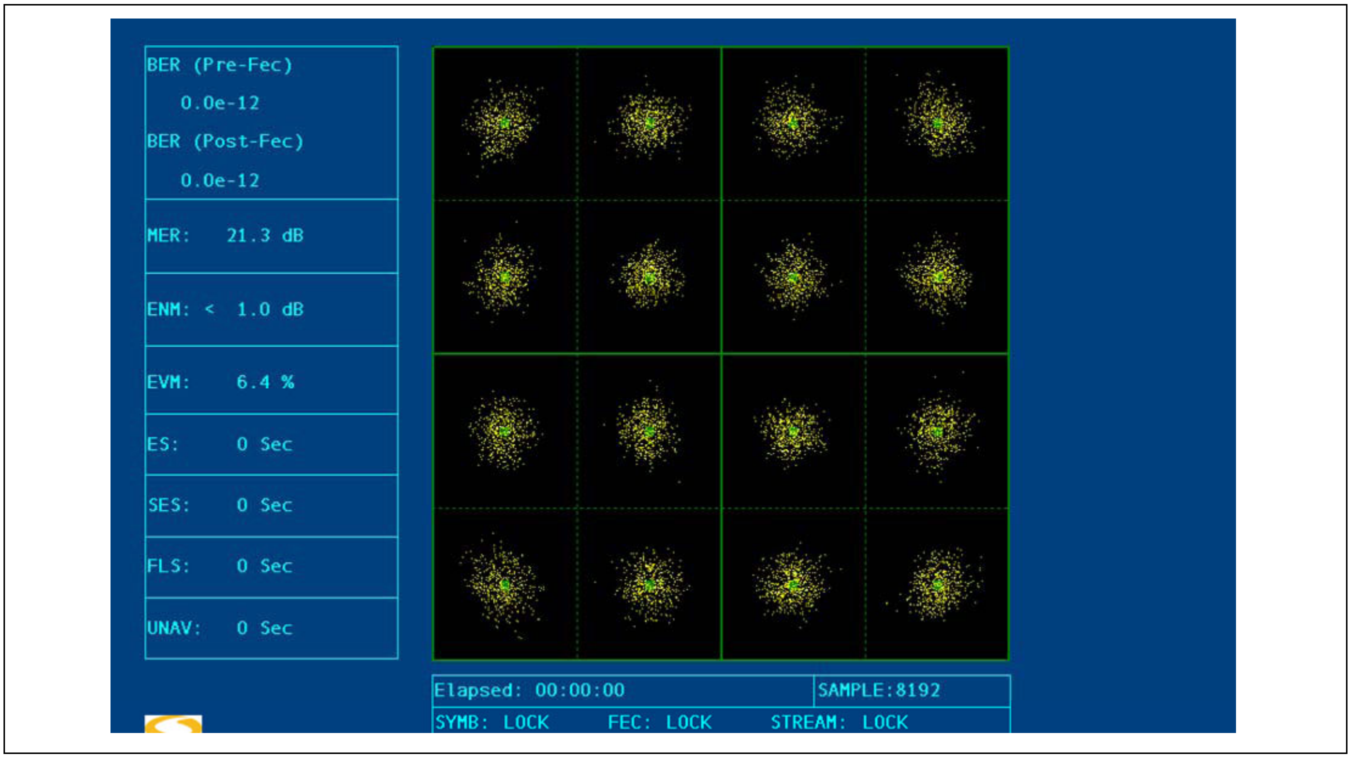

Figure 20: Upstream Unequalized MER for the Digitally Modulated Signal in Figure 19 Is Only 21 dB, and the EVM Is a Relatively High 6.4 Percent. (Courtesy Sunrise Telecom)

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

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).

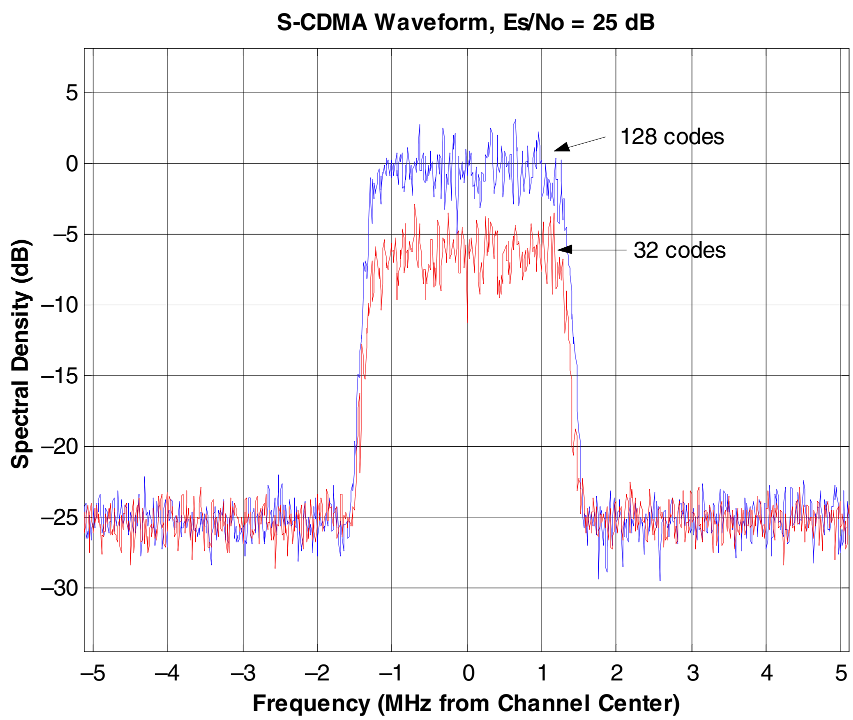

Figure 22: Conceptually, S-CDMA Can Be Thought of as Having an Independent Upstream Channel for Each Code, Each with Its Own Noise Component.

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.

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.

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

Connecting everything®

BROADCOM®

BROADCOM CORPORATION
5300 California Avenue
Irvine, CA 92617

Phone: 949-926-5000
Fax: 949-926-5203
E-mail: info@broadcom.com
Web: www.broadcom.com

CMTS-WP101-R January 2012