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/bcm33843duifsbg

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

CNR and SNR from a Cable Industry Perspective

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

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

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

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

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

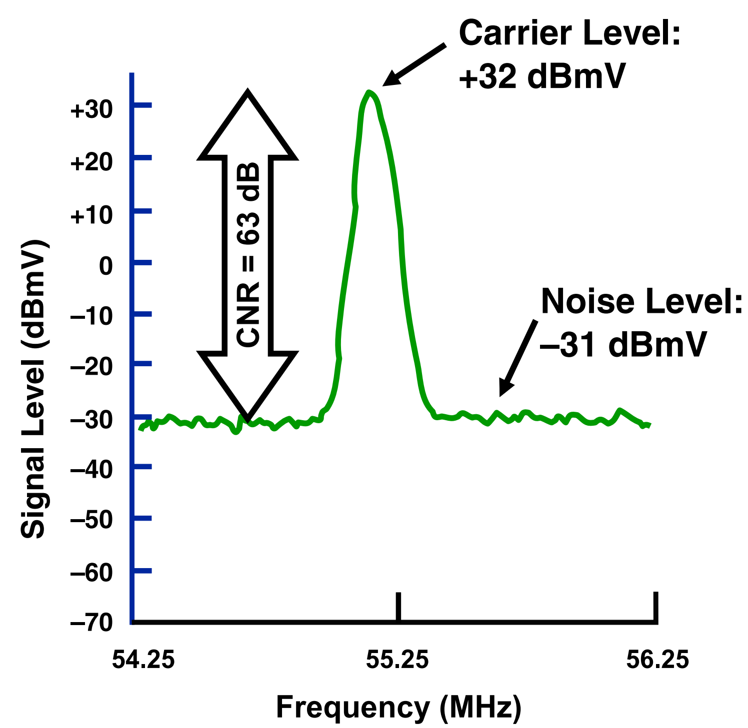

Figure 1: RF CNR Measurement (As Detected in the Test Equipment Resolution Bandwidth)

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

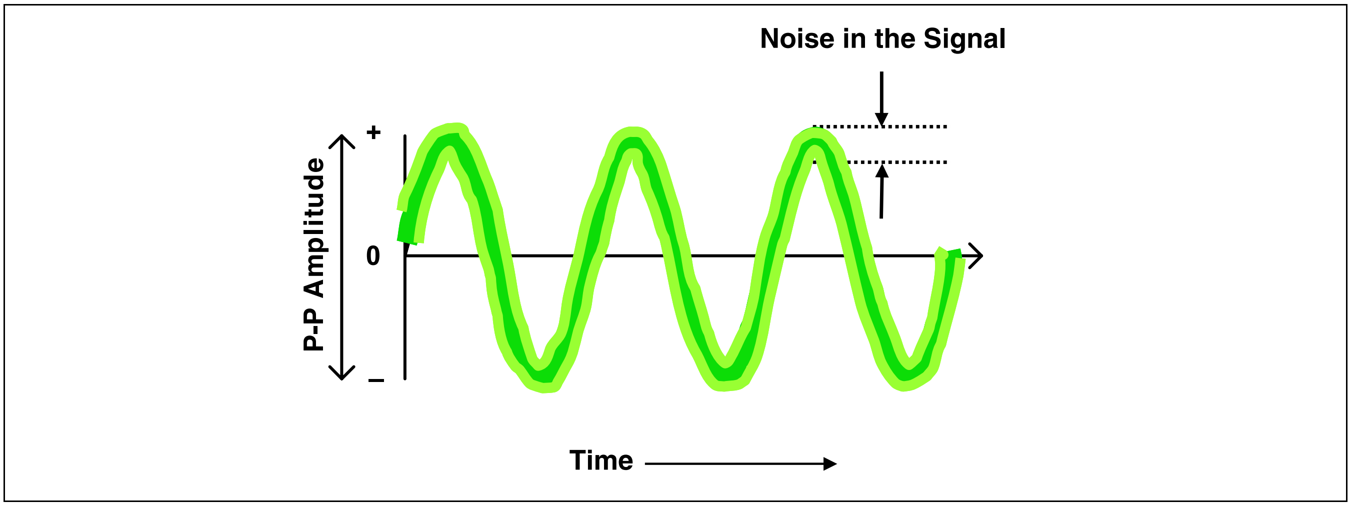

Figure 2: Baseband SNR Measurement

Analog Video CNR in Cable Networks

Consider CNR, which is generally accepted to be a predetection measurement—that is, one made at RF. When only analog TV channels were carried on cable networks, CNR was understood to be the difference, in decibels, between the amplitude of a TV channel visual 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:

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

QAM Spectrum Basics

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

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

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

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

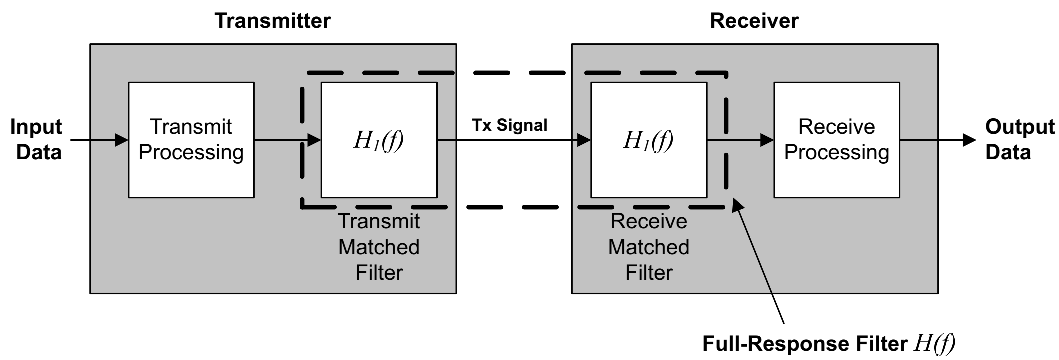

Figure 3: General Digital Communications System Showing Matched Filters in Transmitter and Receiver

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

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

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

Ideal QAM Spectra

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

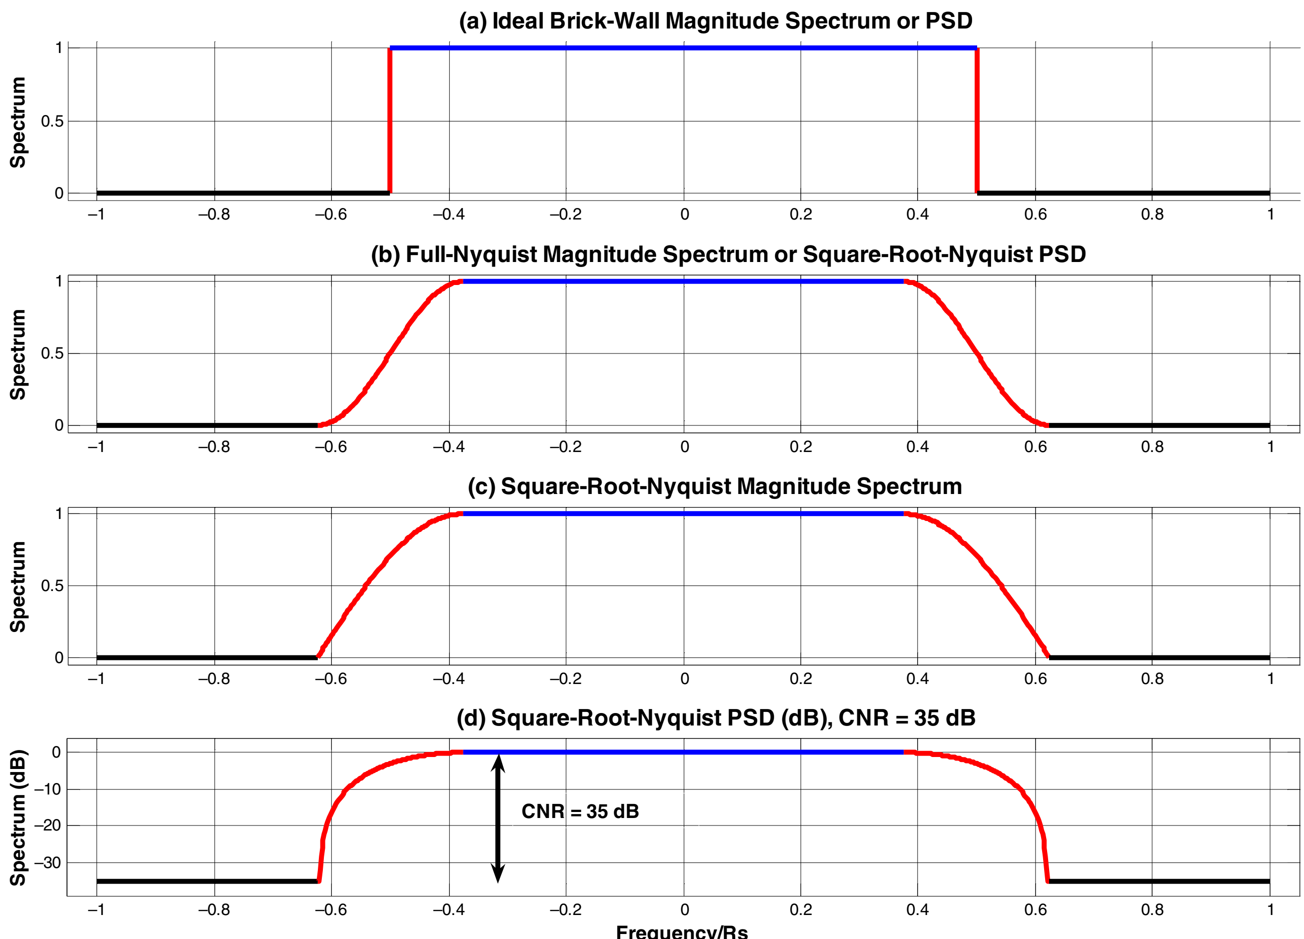

(a) Ideal Brick-Wall Spectrum
(b) Full Nyquist Spectrum
(c) Square-Root Nyquist Spectrum
(d) Square-Root Nyquist Spectrum in dB, with Noise Floor Giving CNR = 35 dB
The passband is shown in blue, transition bands in red, and stop bands in black.

Figure 4: Ideal QAM Spectrum Plots

Part (b) of Figure 4 shows the full-response magnitude spectrum $|H(f)|$ used in the DOCSIS upstream, representing the cascade of the transmit and receive filters. In order to make the filters realizable, an excess bandwidth of 25 percent (alpha = 0.25) is used, resulting in the S-shaped “raised cosine” roll-off regions shown in red, while the passband ideally remains flat. This spectrum possesses the Nyquist property in the frequency domain: If the frequency response $H(f)$ is replicated many times shifted by multiples of the symbol rate, and the copies are overlaid and added as illustrated in Figure 5, the result is a flat spectrum, which results in zero intersymbol interference (ISI).[5] Because of the cascading property mentioned previously, part (b) of Figure 4 also represents the power spectrum $|H_1(f)|^2$ of the transmitted signal and of the square-root Nyquist filter, described next.

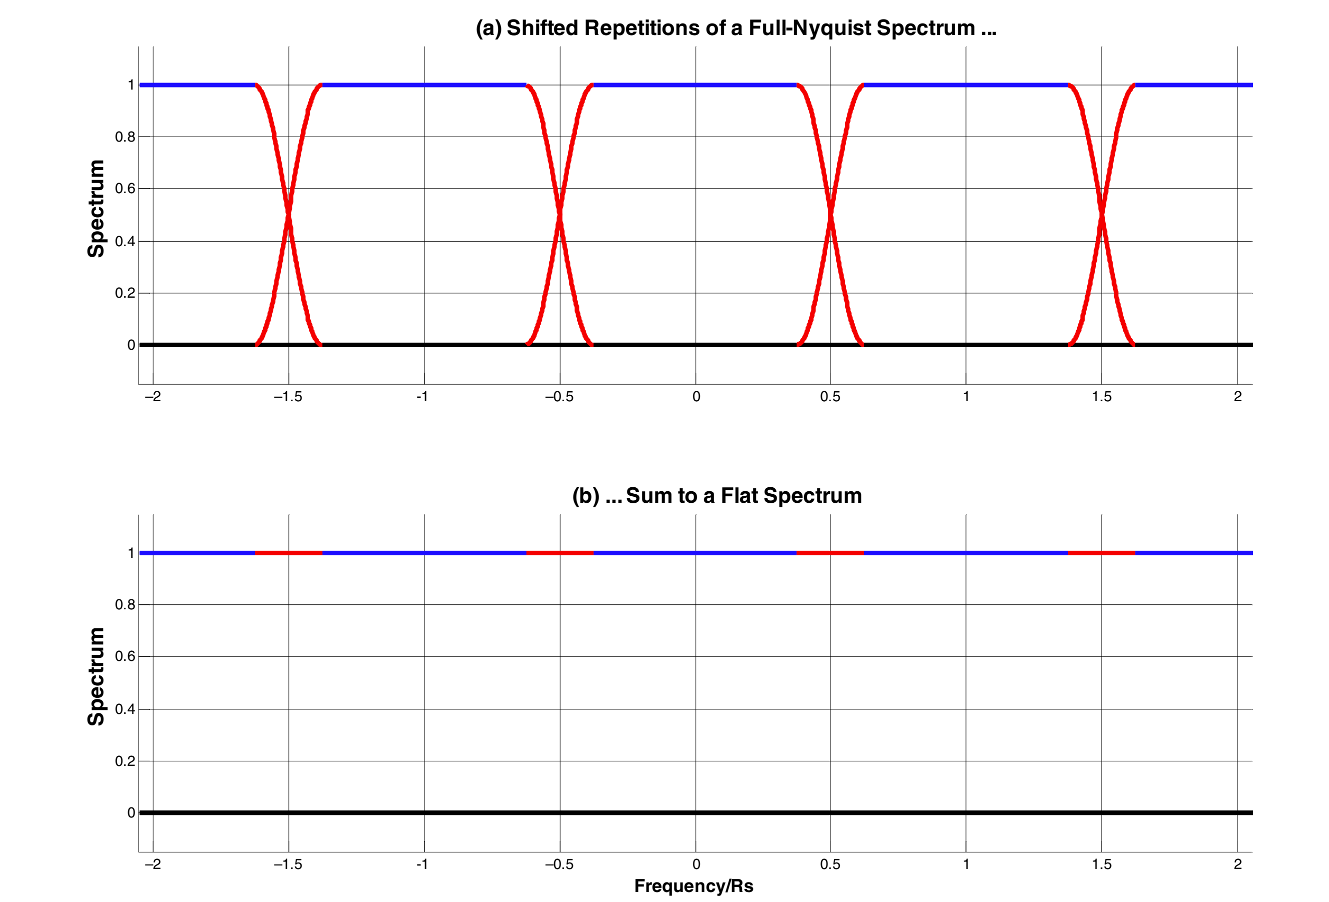

Figure 5: The Nyquist Property States That When Copies of the Spectrum Are Shifted by Multiples of the Symbol Rate and Added, the Result Is a Flat Spectrum, Which Results in Zero ISI.

In practice, the full Nyquist spectrum $H(f)$ of part (b) in Figure 4 is divided into two identical cascaded “square-root Nyquist” filters $H_1(f)$, one in the cable modem upstream transmitter and one in the CMTS burst receiver, using the matched filtering concept discussed earlier. The square-root Nyquist magnitude response $|H_1(f)|$ is shown in part (c) of the figure. Again, because of the cascading property, the power spectrum $|H_1(f)|^2$ of the square-root Nyquist filter and of the transmitted signal, is given in part (b) of the figure.

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

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.

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

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

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

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

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

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

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

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

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.

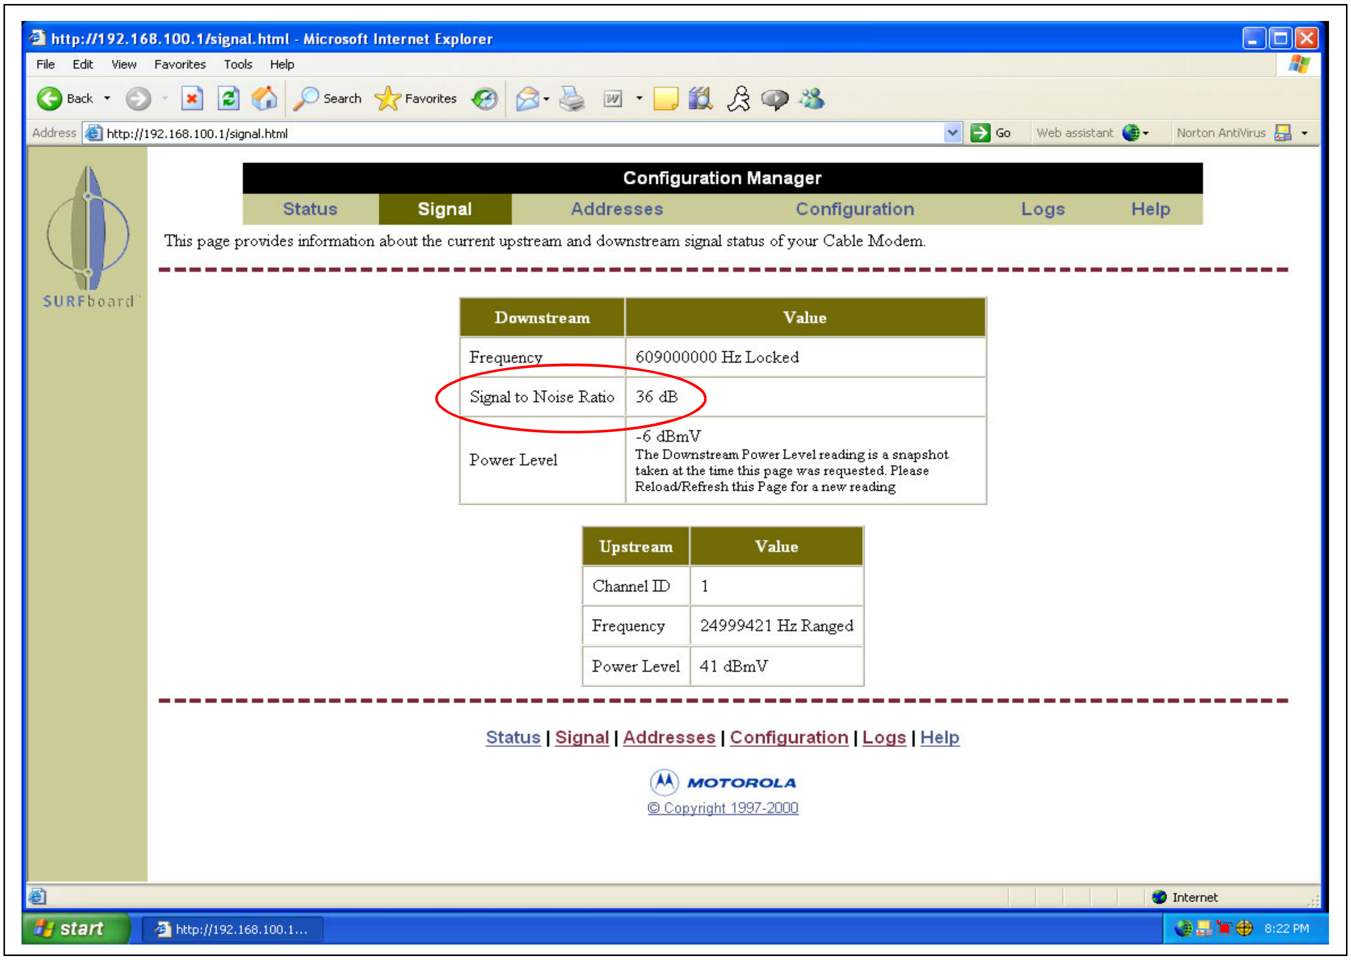

Figure 10: Cable Modem Diagnostics Screen Showing Downstream Equalized SNR (RxMER)

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.

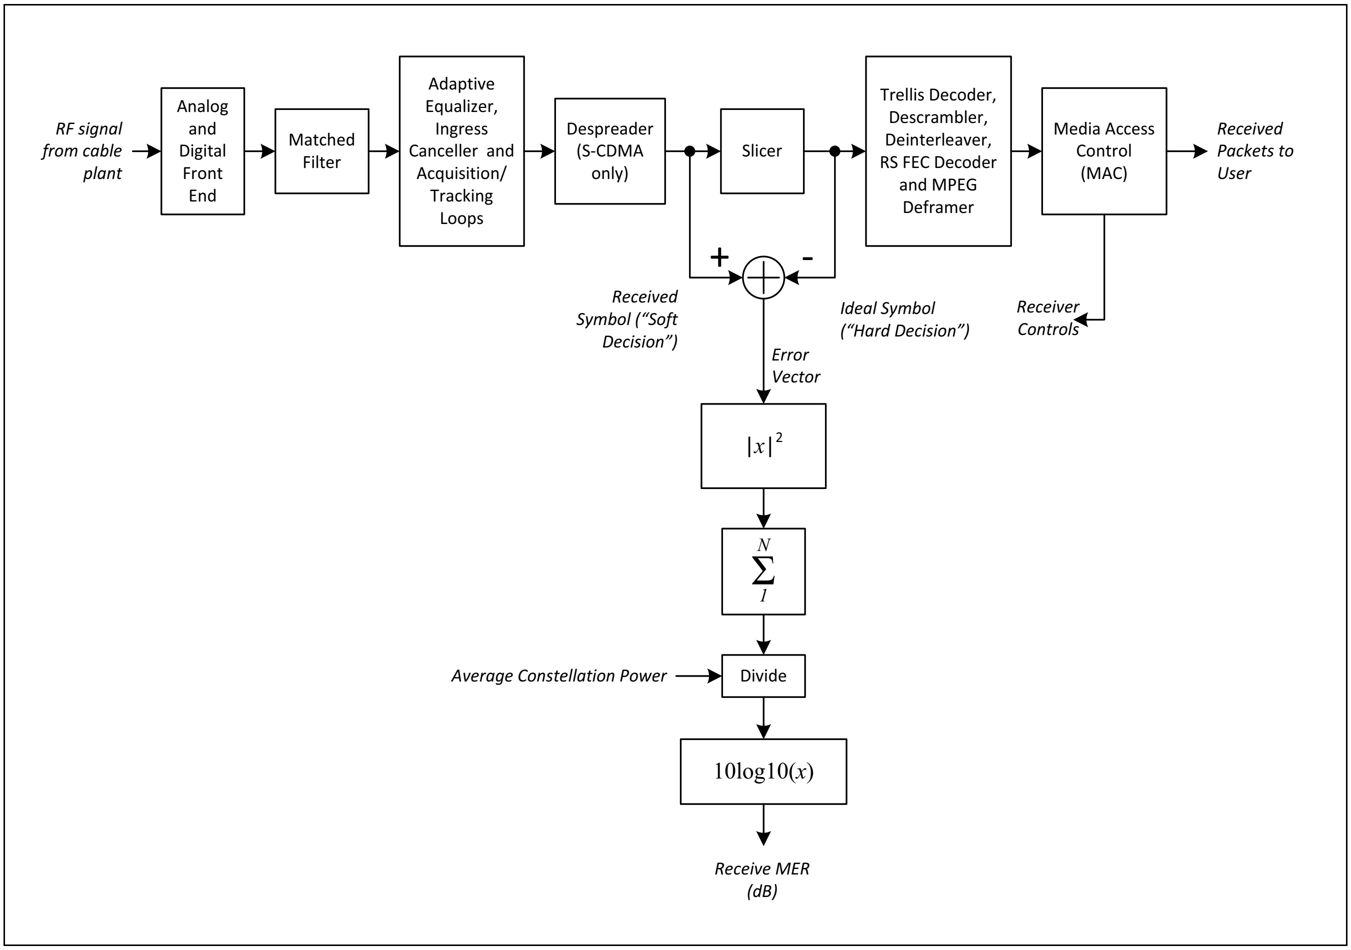

Figure 11: Block Diagram of Generalized Digital QAM Receiver, Showing Computation of Receive MER

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.

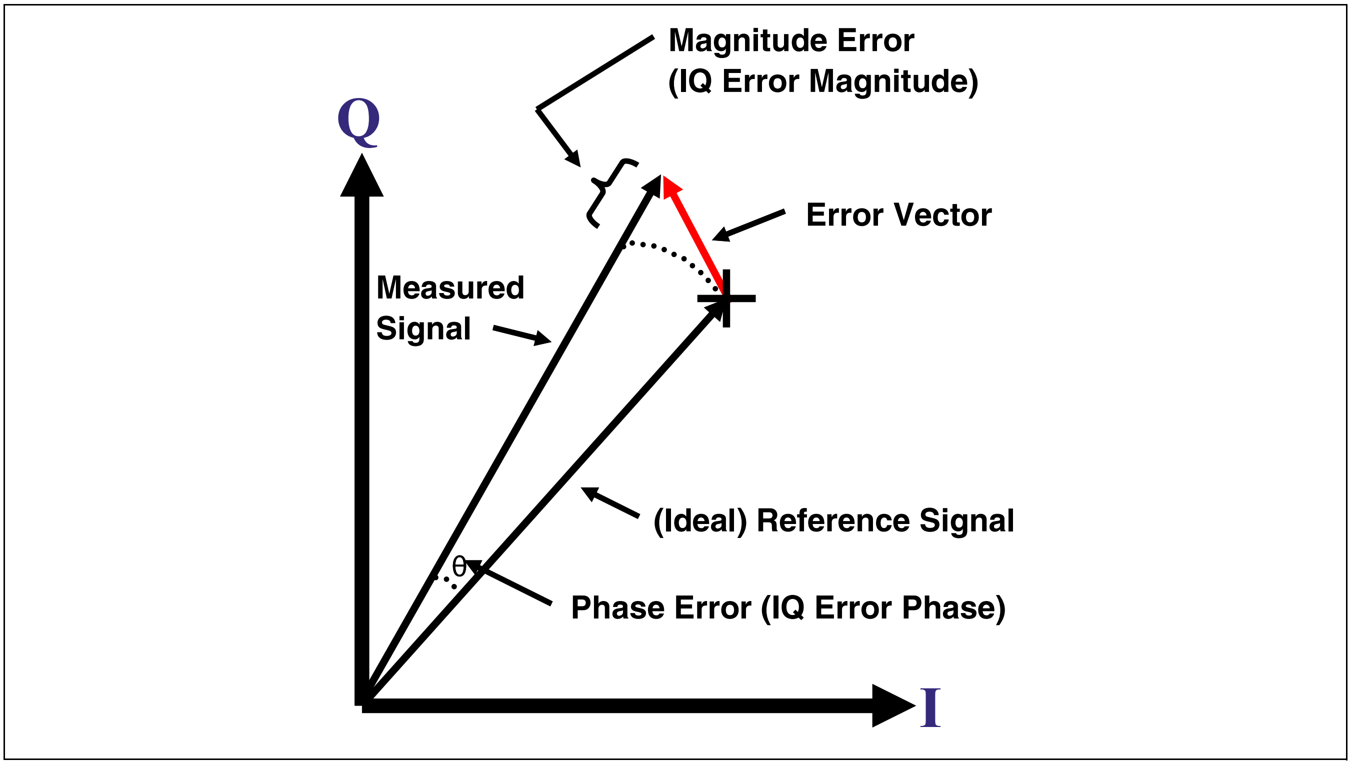

Figure 12: The Error Vector Is the Difference Between the Measured Signal (Soft Decision) and the Reference or Target Signal (Hard Decision). (Source: Hewlett-Packard)

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:

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

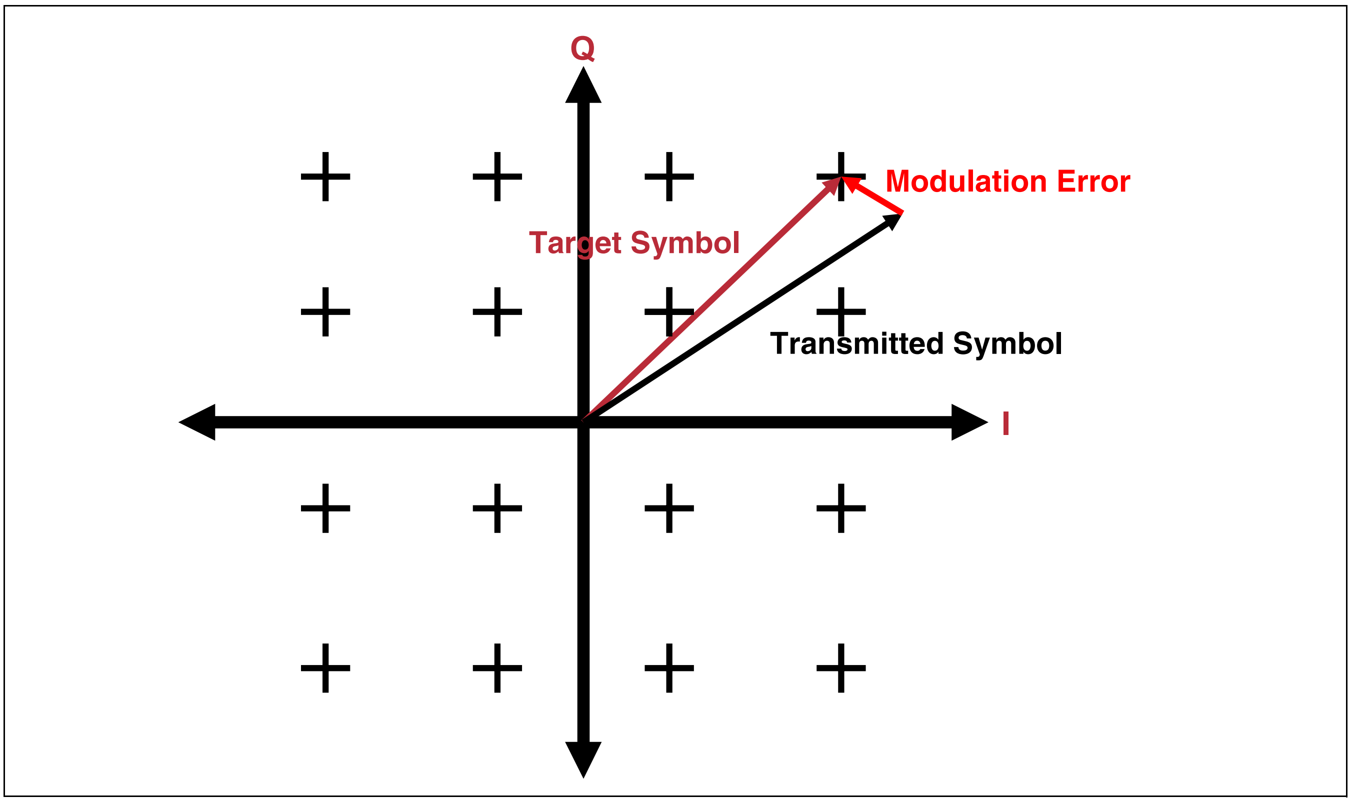

Figure 13: Modulation Error Is a Measure of Modulation Quality. (Source: Hewlett-Packard)

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.

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.

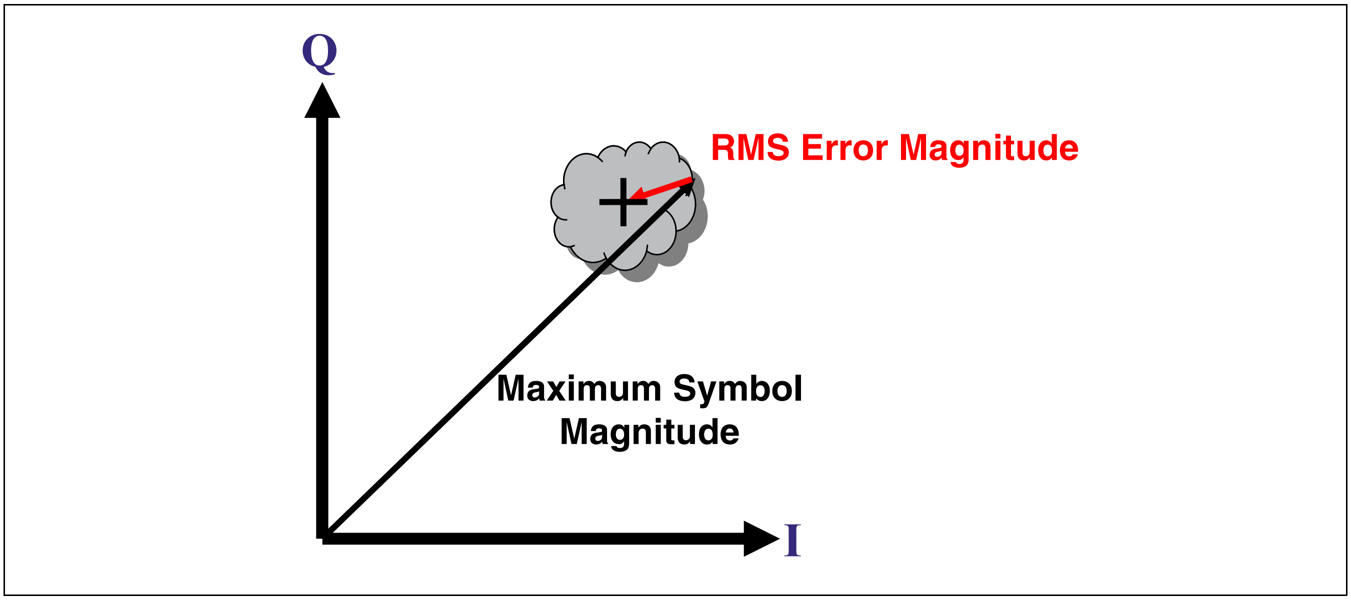

Figure 15: Error Vector Magnitude Is the Ratio (in Percent) of RMS Error Magnitude to Maximum Symbol Magnitude. (Source: Hewlett-Packard)

DOCSIS MIB Definition of Upstream RxMER

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

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

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

EVM Versus MER

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

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

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

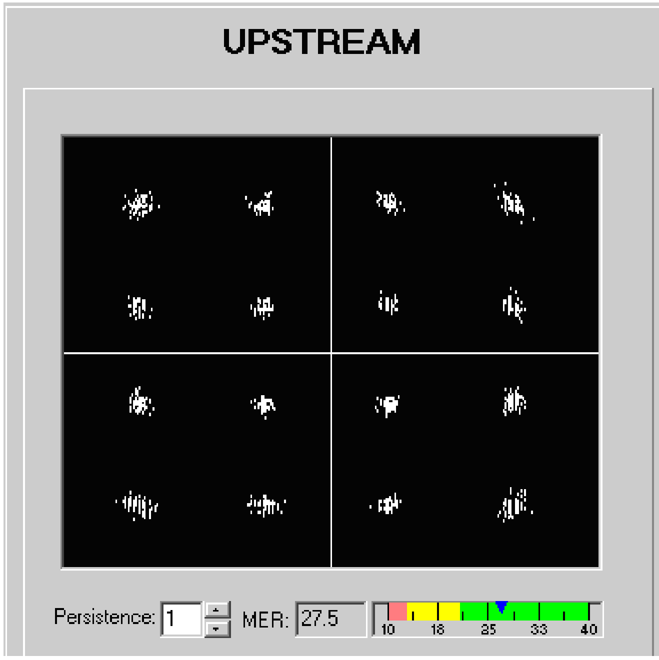

Figure 17: Unimpaired 16-QAM Digitally Modulated Signal; the Unequalized MER Is 27.5 dB. (Courtesy Filtronic-Sigtek)

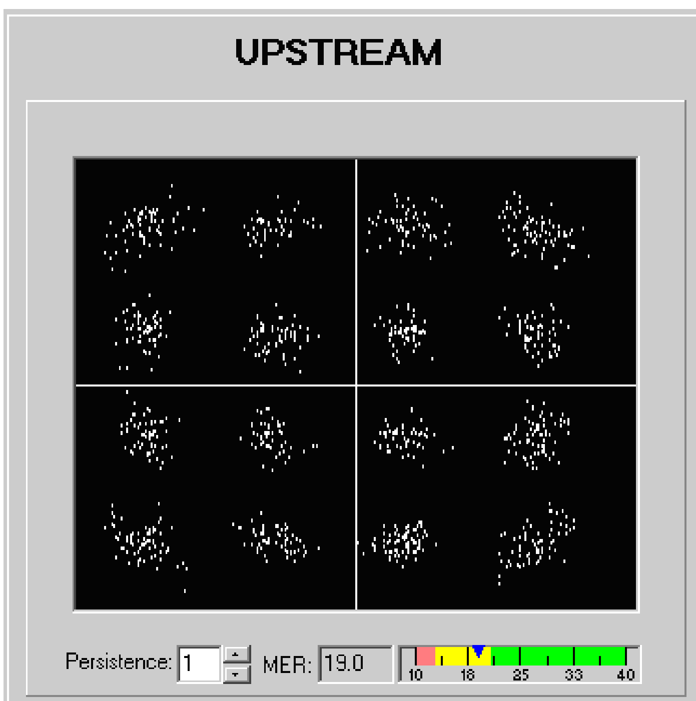

Figure 18: Impaired 16-QAM Digitally Modulated Signal; the Unequalized MER Is 19 dB (Courtesy Filtronic-Sigtek)

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).

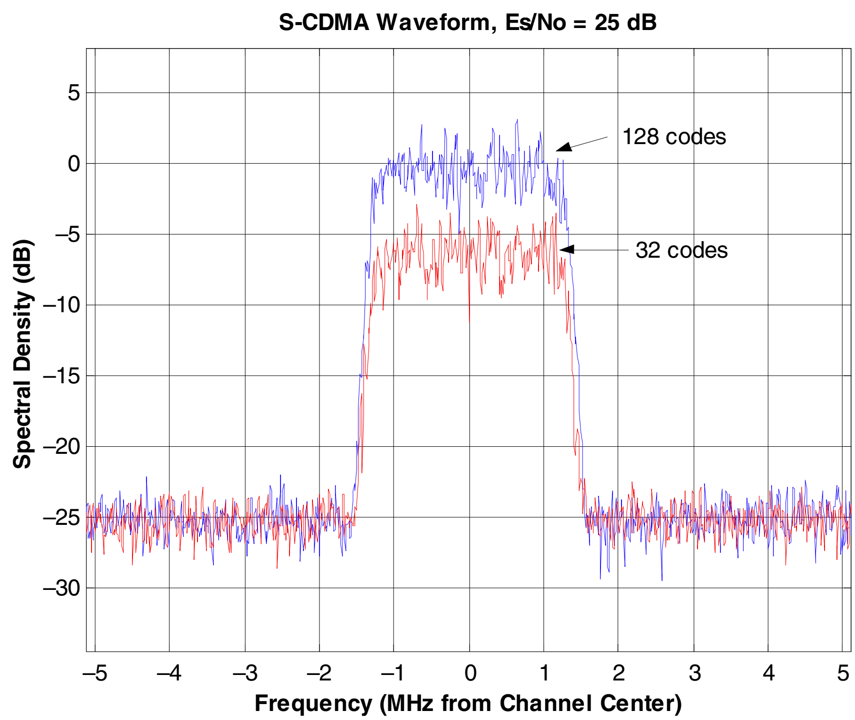

Figure 22: Conceptually, S-CDMA Can Be Thought of as Having an Independent Upstream Channel for Each Code, Each with Its Own Noise Component.

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.

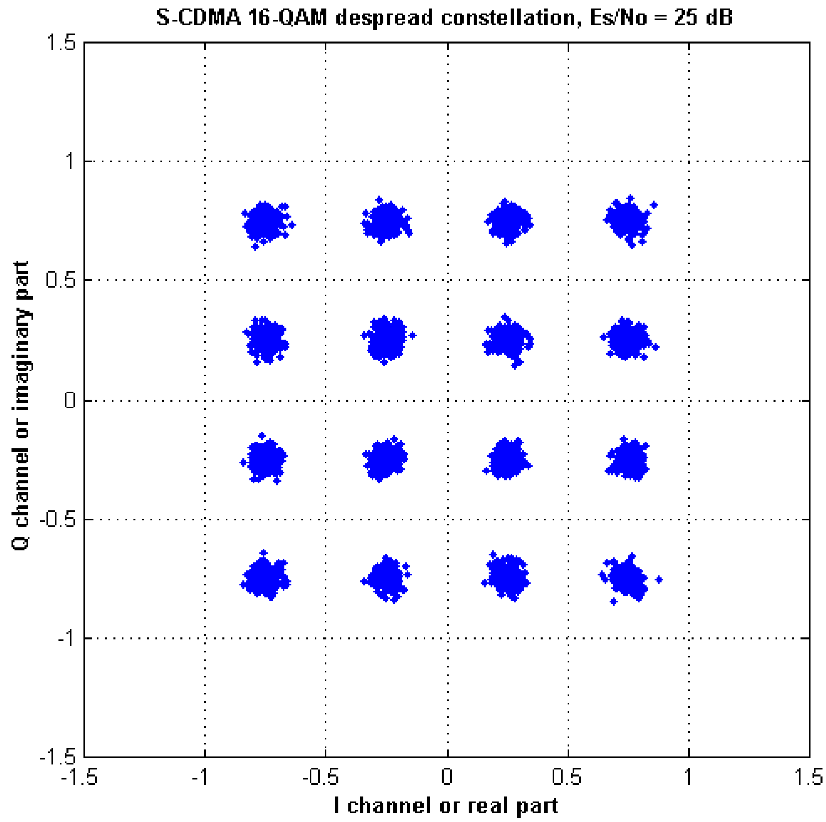

Figure 23: S-CDMA 16-QAM Constellation with 25 dB RxMER

Divergence of CNR and MER in S-CDMA

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

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.

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