Increasing the Performance and Capacity of Digital Reverse Systems

A study of system performance using 4:1 bdr technology

Notices

Trademark Acknowledgments

- Cisco, the Cisco logo, Cisco Systems, the Cisco Systems logo, Scientific Atlanta, bdr, and SciCare are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.
- *All other trademarks mentioned in this document are property of their respective owners*.

Publication Disclaimer

Cisco Systems, Inc., assumes no responsibility for errors or omissions that may appear in this publication. We reserve the right to change this publication at any time without notice. This document is not to be construed as conferring by implication, estoppel, or otherwise any license or right under any copyright or patent, whether or not the use of any information in this document employs an invention claimed in any existing or later issued patent.

Copyright

Improving the Performance and Capacity of Digital Reverse Systems

A Study of System Performance using 4:1 bdr Technology

By David Graves

Abstract

As upstream signals become more numerous and valuable to the cable customer, the need for efficient use of radio frequency (RF) and optical spectrum becomes more acute. Digital reverse systems have increased the capacity of the upstream optical spectrum by populating each optical wavelength of a dense wave division multiplexing (DWDM) system with two reverse data streams using 2:1 time domain multiplexing in the digital space.

The logical extension of this trend is to place more digital information on each optical wavelength. To date, however, this has meant taking unacceptable reductions in performance or unacceptable increases in cost. In this paper, we will examine a way of creating a 4:1 time domain multiplexing system without a substantial reduction in performance, and at the same time decreasing the cost of transporting reverse path signals.

Understanding this new approach to digital multiplexing involves a fresh look at an industry-accepted benchmark, the noise power ratio (NPR) performance curve. Some assumptions and applications that previously have been valid need to be clarified to understand the system application of emerging technologies.

In This White Paper

This white paper contains the following topics.

Introduction

As services become more numerous and take rates increase for these services, the upstream spectrum becomes increasingly crowded. Meanwhile, more downstream content continues to demand additional resources from the plant as well. The result is a strain on transport capacity. It is cost prohibitive to add fiber to achieve more capacity by increasing fiber count. An alternative is to more effectively use the capacity available.

As services become more accepted, penetration rates have increased, creating a higher level of traffic in the upstream path. Additionally, new services such as telephony are rich in upstream data requirements, further exacerbating the situation.

System Impacts

System operators are looking for ways to reduce the size of home pockets served by each reverse path signal. Options open to them are segmenting existing nodes or replacing them with multiple nodes deeper in the network, either of which requires additional fiber capacity. DOCSIS 1.1 provides some relief with the freedom to use additional frequencies in the reverse spectrum, allowing more spectrum to carry payload signals. Still, the need for more reverse path capacity in each fiber remains.

We present in this paper a significant advance in the carrying capacity of reverse path traffic over a fiber link. With advanced digital signal processing algorithms, a single wavelength can transport up to four separate individual reverse path signals, fully recovered and separable at the headend. This represents up to 140 MHz of useable bandwidth from a single node. Previous attempts to provide this level of performance have proven unfeasible. Sacrificing performance by sampling and transporting with fewer bits of information resulted in unacceptably low performance. The option of increasing the bit rate of the transport to 5 Gbps results in an inordinate cost increase. The 4:1 digital transport provides increased capacity with little, if any, performance penalty and a much lower price than other options.

4:1 TDM digital transport doubles the capacity of optical spectrum over 2:1 digital transport systems and quadruples the capacity of standard analog transport.

In this paper, we will first describe the process we use to reduce the amount of bits required to be transported, enabling the low cost transport of 4:1 time division multiplex (TDM) signals. We will then discuss the concept of NPR performance benchmarks and why traditional views of NPR may be misleading when trying to fully understand the performance of digital transport technologies in a realistic system. Finally, we will discuss the implications to a typical system architecture to show a realistic expectation of performance using 4:1 bdr™ technology.

4:1 bdr Bit Rate Reduction

As mentioned previously, transmission rates of 5 Gbps are needed to multiplex and transmit four separate reverse path data streams. The 4:1 bdr products use digital processing algorithms to eliminate unnecessary bits of data and only transport bits that contain useful information. This bit reduction is performed in two processes.

The first process eliminates the unused frequency spectrum that is digitized, but does not contain any useful payload signals. Specifically, the frequencies below 5 MHz and above 42 MHz are not used for return path signals. According to the Nyquist Theory, the rate of sampling needs to be only twice the rate of the bandwidth being delivered. If the reverse spectrum is limited to 37 MHz, then we only need to sample at a 74 MHz rate, meaning that we are oversampling and generating more bits of information than is needed when we sample at 100 MHz. We can digitally “resample” the information collected at the A/D and reduce our effective sampling rate to 75 MHz, eliminating ¼ of the bits normally required to transmit. The information below 5 MHz and above 42 MHz is lost, but no meaningful information is found in that region of the spectrum anyway.

The second process of bit rate reduction is provided by a commonly known process called companding. Appendix A provides a detailed description of the process. In effect, the process takes the 12 bits of information provided by the analog-to-digital converter (ADC) and packs that information into 8 bits for transporting. At the receiver end of the link, the information is converted back to the 12-bit format.

The two processes described here, both accomplished using digital processing algorithms, generate a net bit rate savings of 50%. This allows four data streams to be multiplexed into a bit rate that previously could only occupy two data streams.

NPR Performance Benchmark

NPR was first conceived by the telephony industry in the days of analog telephone transport. With many phone calls simultaneously resident on a single transport system, the composite signal statistics were remarkably similar to Gaussian noise. NPR testing proved to be an accurate measure of the performance of the analog transport system under worst case, very heavy loading. As such, it gained wide acceptance for predicting the capabilities of system designs and components.

Reverse path traffic, being dominated by quadrature phase-shift keying (QPSK) or quadrature amplitude modulation (QAM) type signals, also looks similar to Gaussian noise. There are some distinctions, which we will see shortly, but in many situations these differences are immaterial. Consequently, NPR has gained in popularity in the cable industry as an effective means of comparing performance capabilities between competing technologies and products.

System Application of NPR Performance

Assumptions

To properly understand the application of NPR to system performance, it is useful to keep in mind the assumptions made in the interpretation of the NPR results:

1) The levels of signals in the reverse path are assumed to be of constant power density. That is, every signal present in the reverse path would have the same power level when viewed on a spectrum analyzer. This power level is most generally given in terms of dBmV/Hz, or the amount of power (in dBmV) contained within one Hz of bandwidth.

2) The system noise level remains constant regardless of noise level or spectrum loading. In other words, the noise measured is not affected, whether the system is experiencing heavy traffic loading or light.

3) The distortions as predicted by the NPR results assume worst case, full spectrum loading, representative of 100% capacity utilization. If the system were to experience less than 100% capacity utilization, the distortions, of course, would occur at higher signal levels.

Figure 1 shows a typical NPR curve for an optical transport system. The curve measures the ratio of the level of the signal carrier to the noise plus distortion products (vertical scale), over a range of input signal levels (horizontal scale). The input signal level is expressed as the power spectral density in a 1 Hz bandwidth, assuming the signal level to be constant across the entire system spectrum.

When the signal levels are low (towards the left side of the graph), the system distortions are negligible and the system performance is limited by the system noise. As the signal level is increased, the ratio of the signal to the level of the noise increases at the same rate. This can be seen by the fact that the slope of the curve in this region is linear and rises at a slope of unity.

Continued on next page

At some point, the signal level is increased to the point that the distortions start rising rapidly. This may be due the saturation of an amplifier component, or the clipping effects of an analog laser or ADC.

The useful range of signal levels that can be tolerated while maintaining a minimum system performance can be calculated by the difference between minimum signal level at the point where the carrier-to-noise ratio (CNR) drops below the minimum system requirement and the maximum signal level, at which point the distortions introduced into the system increase above the system carrier to interference requirement. This is usually calculated by determining the width of the NPR curve at the minimum performance level (CNR).

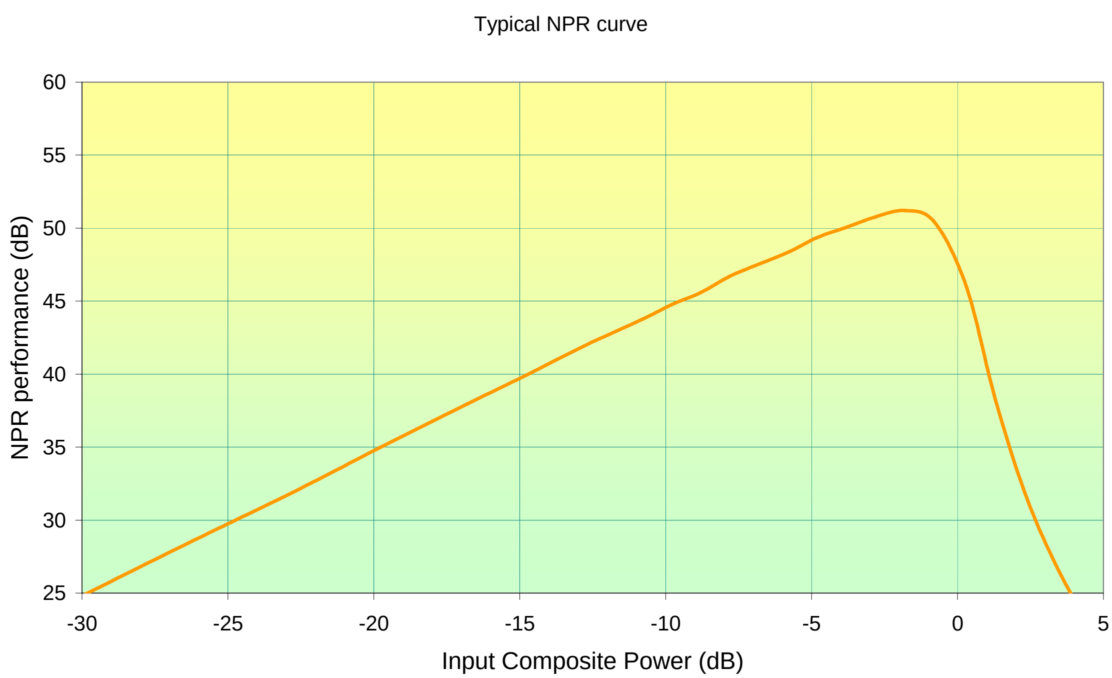

Figure 1. Typical optical transmission system NPR curve

System Performance Requirements

When determining the performance requirements, the highest order modulation scheme expected must be supported. DOCSIS 1.1 supports modulation up to 16 QAM, but the next version of DOCSIS may indeed support as high as 64 QAM. A 16 QAM system requires no more than 25 dB of CNR to achieve good bit error rate (BER) performance. 64 QAM, however, needs to be closer to 30 dB. Some telephony modulation schemes, such as OFDM, also require upwards of 30 dB CNR.

Continued on next page

The CNR requirements are those required by the signal demodulation receiver. This is referred to as the end-of-line performance. Other system degradations must be accounted for in addition to this requirement. RF combining effects in the headend or primary hub, for example, will put more stringent requirements onto individual return path signals to compensate for the noise increase. The CNR requirement is increased by 3 dB when the number of RF paths combined is doubled. Hence, 4:1 RF combining requires 6 dB of additional CNR to account for the noise addition.

Combining Effects

The discussion of the NPR curve so far has been limited to a single reverse path received at the headend. If this were the performance criteria of interest, then the performance of the link above 32 dB would be of little consequence, as even 64 QAM signals require no greater than 30 dB of CNR. However, many systems employ some amount of RF combining at the headend after receiving the optical signal.

When two RF signals are combined, the noise of the two systems are added together creating a combined noise floor 3 dB higher than the original noise floors, assuming the two noise floors were originally equal. The impact to the NPR curve is that the left side of the curve (carrier to noise) is reduced on the vertical scale by 3 dB. The right side of the curve (distortions) does not necessarily change by 3 dB. If the traffic load were always equally split between the two paths, then the clipping level of each leg would be extended to the right by 3 dB. This rarely is the case, however, as the signals could realistically be very heavily loaded to one or the other paths, in which case the distortions created remain as if no RF combining was assumed. In the most general case, we must assume that the distortions as measured by the NPR curve remain the same.

The result of combining two RF return paths is that the NPR curve is reduced by 3 dB in height, and a corresponding 3 dB reduction in dynamic range (width of the curve) occurs.

For a system with 4:1 combining that is designed for telephony and 64 QAM modulation, the performance requirement for an individual transport path will usually be considered to be 38 dB. An NPR width of 10 dB will usually provide ample headroom between a system's nominal operating point and the clipping levels of the system.

Continued on next page

NPR Curve of 4:1 Single Stream

Figure 2 shows the measured NPR curve for the 4:1 bdr product. Notice that the left side of the curve no longer is a straight line and the slope is less than unity. This tells us that an increase in input signal strength does not result in an equivalent increase in CNR. Either the system is no longer a linear system with a constant gain versus input power transfer function, or the system noise does not remain constant with input power. In this case, the latter is true. As mentioned previously, the digital signal processor (DSP) functions that perform bit reduction add increased noise as the composite signal power is increased. Figure 3 shows the noise level as a function of composite input power. Changes in a system's composite input power may come from either increased signal levels or increased traffic and the corresponding increase in used bandwidth.

When a system deviates from the normal expected characteristics of systems we are accustomed to, it forces us to take a closer look at the assumptions being made in the methods used to judge its performance. In this case, we must look more closely at what the NPR curve is telling us about system performance. With the knowledge that the noise floor is a function of input power level, there are a number of factors that now come into play in interpreting the NPR performance curve.

Remember that NPR measurements simulate a system traffic load under extreme conditions, in particular 100% traffic loading. Other system factors influence the end-of-line performance, such as RF combining. These conditions, well understood in previous systems, now have different implications on system performance.

Continued on next page

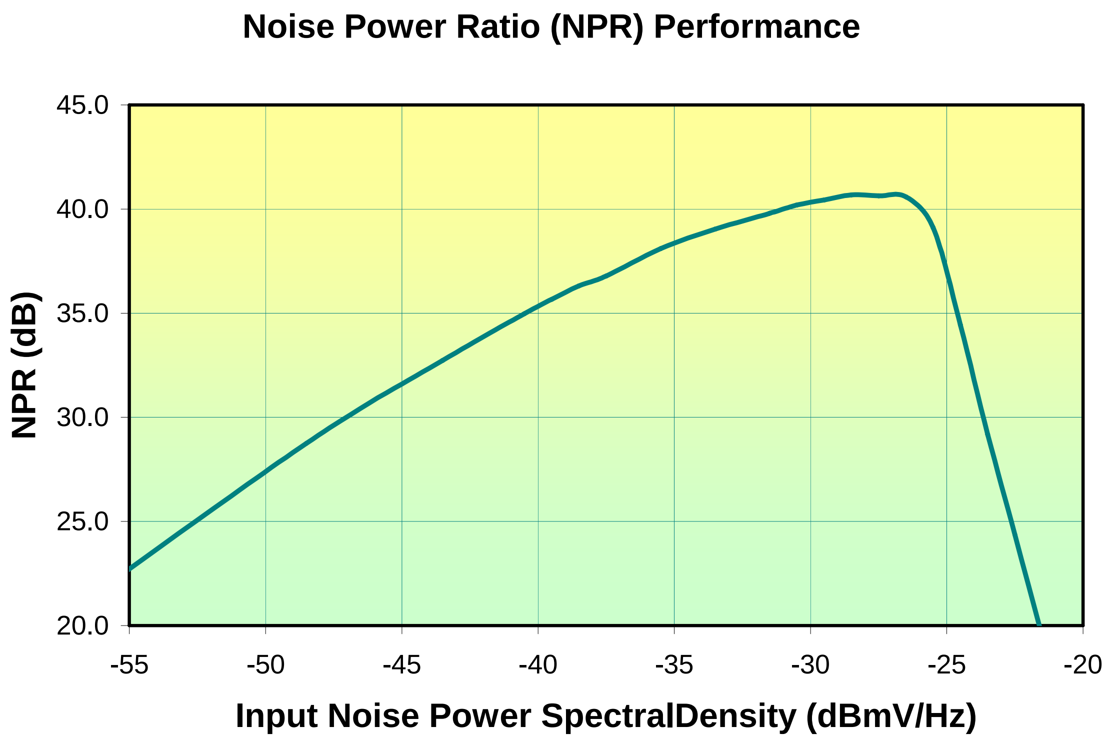

Figure 2. The 4:1 bdr NPR curve differs from a typical NPR curve in significant ways. These differences have a number of significant implications when interpreting system performance.

Continued on next page

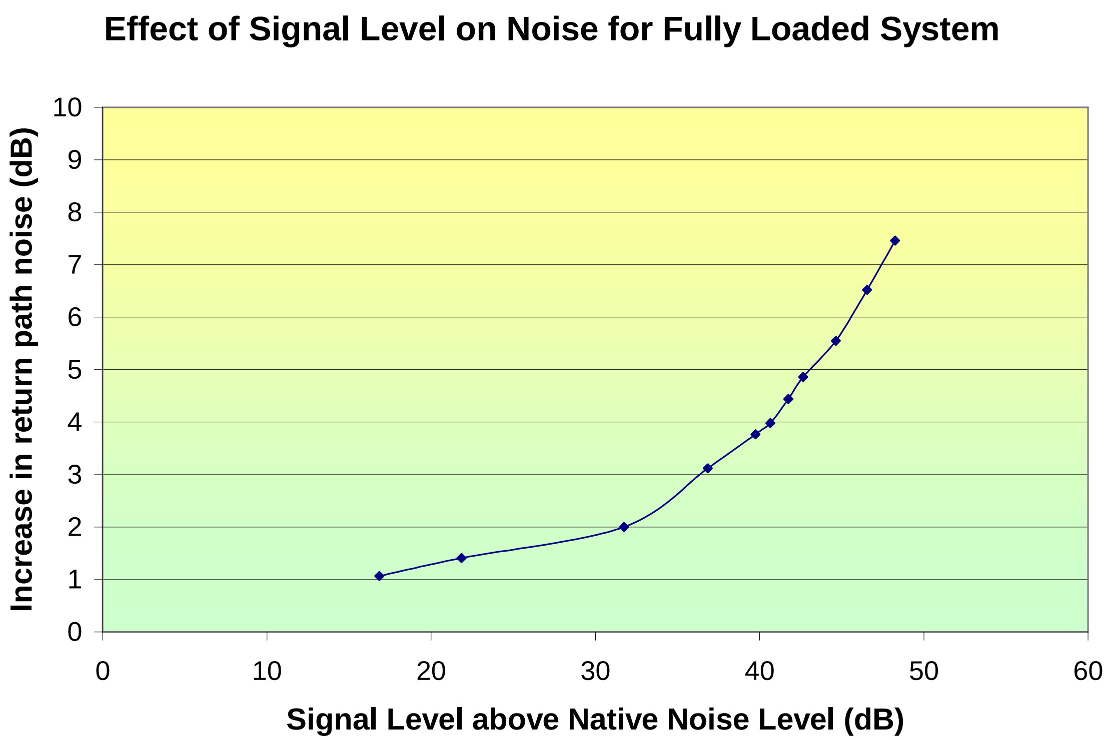

Figure 3. As the input composite power increases, the system noise figure increases due to the DSP functions incorporated in the 4:1 bdr system. This artificially reduces the system performance predicted by NPR measurements.

Factors that Influence Performance in Real System Architectures

1. Less than Full Loading of Reverse Bandwidth will Reduce the Composite Power Contained in the Spectrum

In a conventional transport system, a reduction of composite power only affected system performance if it was related to a change of input power levels. A change in composite power due to an increase or decrease in used bandwidth did not affect the system CNR performance. This is an important distinction between previous digital return systems and the 4:1 bdr system performance. In a realistic reverse path system, we never fully utilize the reverse path bandwidth. Even with a full loading of signals, there would be some unused bandwidth located between active channels. Consider a system that fully loaded the reverse path with 2.5 MHz 16 QAM carriers. These carriers are normally set at 3.2 MHz center-to-center spacing to avoid interfering with neighboring signals. This would allow a maximum of 11 QAM channels in a 35 MHz spectrum. In this very extreme case, the used bandwidth is in reality 27.5 MHz, instead of 35 MHz. A corresponding drop in composite power results, compared to NPR measurements. Without loss of generality, one can see that an improvement in realized performance of a 4:1 bdr system exists compared to that predicted by NPR measurements

Most systems are not going to be able to use the full bandwidth, due to noise pollution and ingress signals. For instance, it is common practice to use the spectrum below 15 MHz very sparingly and for small traffic loads, due to a noisy environment in this spectrum. Although this pollution is significant to affect signal quality in this region of the spectrum, it does not significantly impact the composite power loading of the system. This would further limit a system to about 25 MHz of usable spectrum, or a maximum of about 7 QAM carriers (17.5 MHz of used bandwidth).

Any factor that limits the reverse traffic affects the system CNR in a positive manner. Figure 4 shows a plot of system dynamic range against equivalent signal bandwidth at a CNR of 38 dB. Another view (Figure 5) shows the system CNR for various traffic loadings, compared to a measured NPR curve. The plot shows that system performance improvement can be substantial when this factor is taken into consideration.

Continued on next page

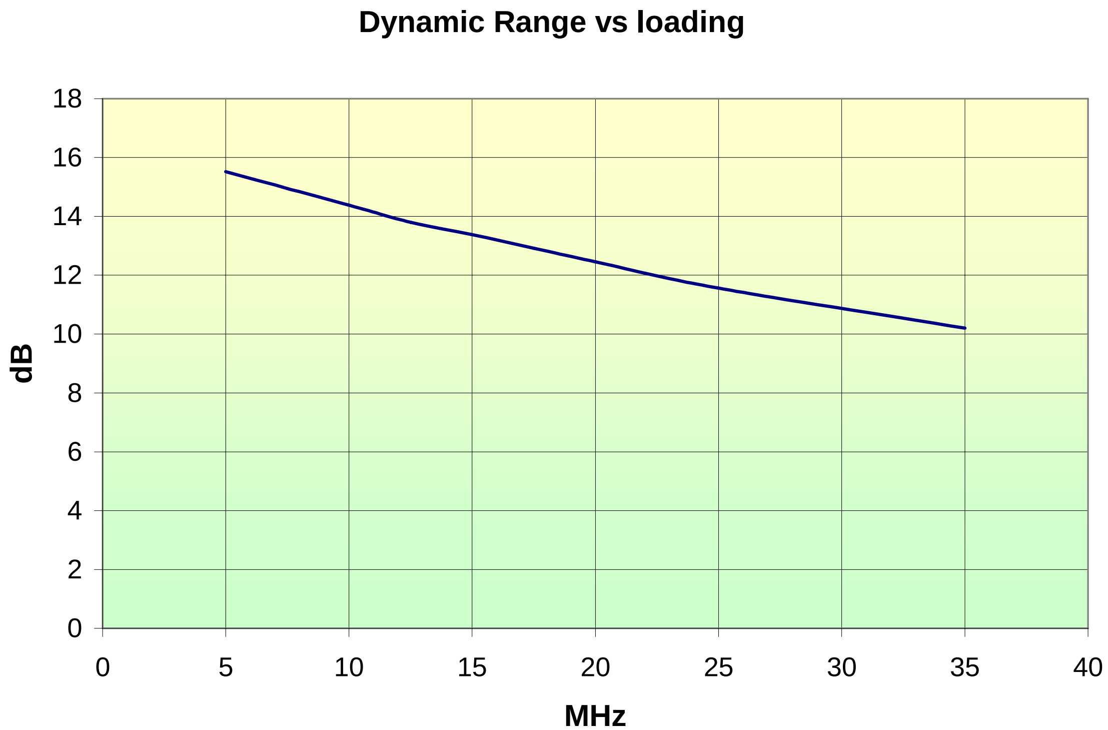

Figure 4. The 4:1 bdr system dynamic range increases as the used bandwidth decreases. This implies that the performance predicted by the NPR measurement is too conservative.

Continued on next page

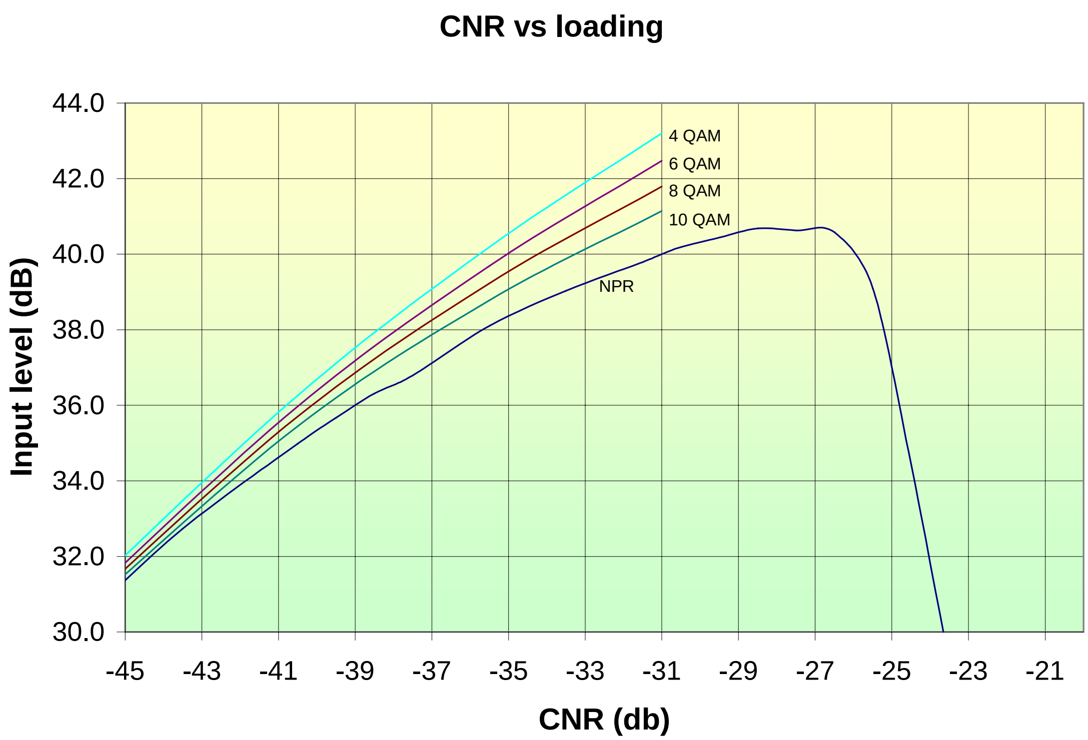

Figure 5. The number of channels active at any given instance in time determines the system performance. The system loading and signal statistics will determine the proper loading that accurately predicts the system performance.

2. Sharing of Service Channels Between Combined Reverse Paths will Necessarily Cause the Instantaneous Loading of Each Path to be Reduced, Further Improving Performance

When a system employs RF combining of multiple signals, it is understood that the average traffic loading on any individual reverse path will be less than 100%. If a service such as modem traffic is shared, then only a single reverse path may have traffic on that service active at any given time. Otherwise, data collision will occur. In this case, the traffic load is shared among all combined paths and the average traffic loading is less than 100%, even if the service is being 100% utilized. This will also contribute to the composite power reduction seen in a reverse path.

In addition, when traffic loading begins to approach 100% utilization for any appreciable fraction of the time, a system operator will in some way relieve the service to provide a more acceptable level of utilization, usually by reducing the number of reverse paths being combined. So, 100% utilization in shared services is more an exception than the rule.

Continued on next page

We can calculate the statistical probability that a given amount of loading exists on a single reverse path. Consider a system that employs 4:1 combining and whose services are utilized 100%. With a fully loaded reverse path, there will be 11 separate QAM channels active in the reverse spectrum. Assuming an average active duty cycle of 25% on each reverse path, the probability of a given number of channels active on that reverse path is:

$$P(C_{active}) = \frac{C_{total}!}{C_{active}!(C_{total} - C_{active})!}(\text{Duty Cycle})^{C_{active}}(1 - \text{Duty Cycle})^{C_{total} - C_{active}}$$

where :

P is the probability,

C_{active} is the number of channels active at any given time,

C_{total} is the total number of channels on the system, and

Duty Cycle is the average amount of time each channel is active

From this equation, we can calculate the probability that a given number of channels are transmitting simultaneously:

# of channels active	Probability of being active
0	0.0422
1	0.1548
2	0.2581
3	0.2581
4	0.1720
5	0.0802
6	0.0267
7	0.0063
8	0.0010
9	0.0001
10	7.86781E-06
11	2.38419E-07

To attain a given level of performance 99% of the time, we can assume that there will be no more than 7 QAM signals transmitting at any given time.

Continued on next page

3. With RF Combining of Shared Services, the Noise Penalty of Combining is Not a Full 3 dB

Conventional approaches dictate that to achieve 32 dB end-of-line performance and 4:1 combining, the individual reverse path CNR must be a full 6 dB higher, or 38 dB to account for the four-path noise combining. The underlying assumption here is that the noise levels of all four paths are the same. This is not necessarily the case with 4:1 bdr, as the noise level is related to the traffic present on the reverse path. If one begins with the performance of one path having relatively heavy traffic, then the paths it is combined with at the headend must necessarily contain lower traffic volume and thus lower noise levels. This will result in a combined noise increase of less than 6 dB. The actual numbers attained will depend on the system traffic modeling. If in the system design the noise level of each return path for a 25% loading is assumed, then the 6 dB penalty would apply from that noise level. It should be appreciated, though, that this level is not the performance predicted by the NPR performance curve.

Final Performance Expectations

Let us consider a system that employs 4:1 bdr as an example. The system uses 8 QAM channels to provide the reverse path traffic. Half of the channels are shared among four reverse paths, employing four -way RF combining on these four channels. Furthermore, we will assume that the traffic loading is equal among the four paths and that we are interested in a 99% performance level (performance is greater than or equal to this level 99% of the time). The traffic is assumed to be less than 70% capacity on all channels. For this exercise, we will ignore the effects of loading on distortion products and focus our attention on the system noise and CNR performance.

Our desire is to determine the effective loading of the paths and determine the CNR performance associated with this loading. The four QAM channels that are dedicated services will all be active 70% of the time on each reverse path. The four shared channels will be active .7*.25 = 0.17% of the time each. The full loading is then four channels active 70% of the time each and four channels active 17% of the time each. For 99% of the time, we can consider worst case a single reverse path to be transmitting on seven or fewer channels. This provides us with a CNR curve as shown in Figure 6.

The final step is to determine the combining effects on the noise level. If we assume worst case all four paths would have their dedicated channels active, and the four shared channels would all be active, as well. However, these channels are shared among the four and distributed among them. The worst-case scenario for a single channel would be to have seven channels transmitting on one path. There would necessarily be a maximum of five channels active on one other path and only four active channels on the remaining two paths. We then have four paths with different noise levels, which need to be summed together to determine the combined noise level. Summing the noise contributions of these four return paths yields a combined CNR given by the dotted curve in Figure 6. This represents the end-of-line performance of the optical path. From the curve it is shown that the width of the curve at 32 dB is 16 dB. Comparing this to conventional method of interpreting NPR curves, the estimated performance would be only 10 dB.

In comparison, the end-of-line performance of most 2:1 digital reverse products also provides a 15 dB dynamic range measured at 32 dB of CNR under the same conditions.

Continued on next page

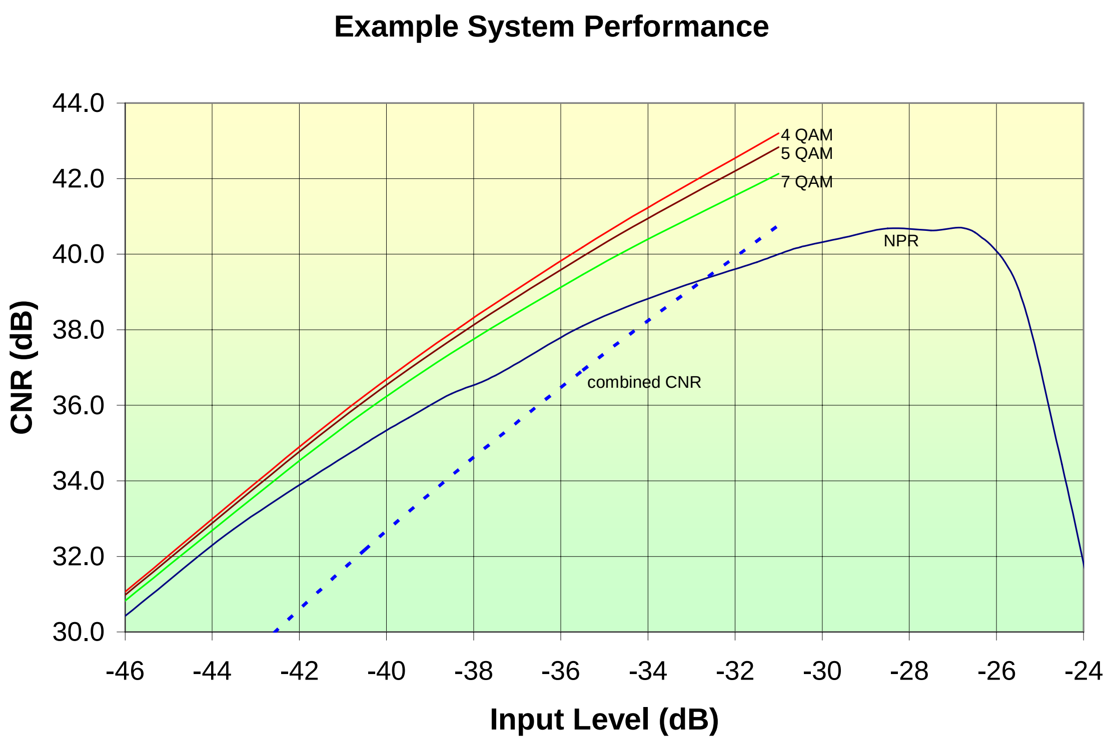

Figure 6. Taking into consideration the system service loading, signal statistics, and noise combining, the end-of-line performance using 4:1 bdr is much better than predicted by NPR measurements.

Conclusion

The signal processing functions in the 4:1 bdr product create a signal dependent noise level that makes the real system performance difficult to predict from the NPR performance curve. NPR measurements provide an overly restricted prediction of system performance. For this reason, the system design parameters and traffic modeling should be considered when predicting the system performance.

Although the NPR curves would predict an end-of-line performance of about 10 dB for a 4:1 optical link, after taking into account more realistic spectrum loading, signal statistics and RF combining effects, the system performance turns out to be closer to 15-16 dB. This shows that for the vast majority of systems, the 4:1 performance will be very close to the performance levels of the standard 2:1 reverse path systems.

There is a minimal, if any, performance penalty when comparing a 2:1 digital reverse system with a 4:1 bdr system. However, the price implications are significant. Consider that the cost of International Telecommunications Union (ITU) optics and passive components are cut in half using the 4:1 bdr system. The quantities of the highest cost components in the optical link are reduced by half. Additionally, there are indirect cost savings with fewer wavelengths required, and fewer fibers required due to the more efficient utilization of each wavelength.

Appendix A: Quantization Noise and Companding

By Lamar West, Ph.D.

Quantization noise results from the inability of an A/D-D/A to reproduce an input voltage with infinite accuracy. For example, a 4-bit A/D-D/A can output only 16 distinct voltages. A transfer characteristic for such an A/D-D/A with an input full-scale range of 1.000 volts centered around 0.000 volts is shown in Figure A1.

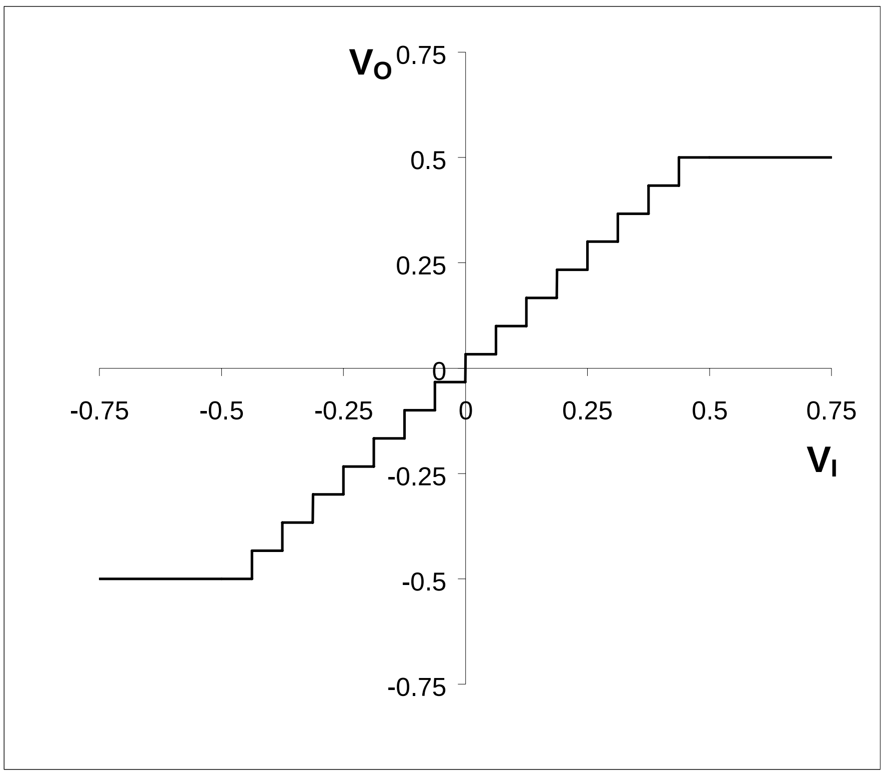

Figure A1. 4-bit A/D-D/A Transfer Characteristic

Consider voltages in the range of 0.0000 to 0.0625 volts. All such voltages will produce an output of 0.03125 volts. If a voltage of 0.01 volts is applied to the A/D-D/A, an output of 0.03125 volts will be produced, resulting in a quantization error of 0.02125 volts. Similarly, if a voltage of 0.0500 volts is applied, an output of 0.03125 volts will also be produced, resulting in a quantization error of –0.01875 volts. Consider now an input less than –0.5 volts. Any such input will result in an output of –0.5 volts. Similarly, an input greater than +0.5 volts will result in an output of +0.5 volts. This, too, is quantization noise and is analogous to clipping distortion in analog devices.

Continued on next page

A plot of quantization error versus input voltage for a 4-bit A/D-D/A is shown in Figure A2. These results may be directly extended to devices with more or fewer bits.

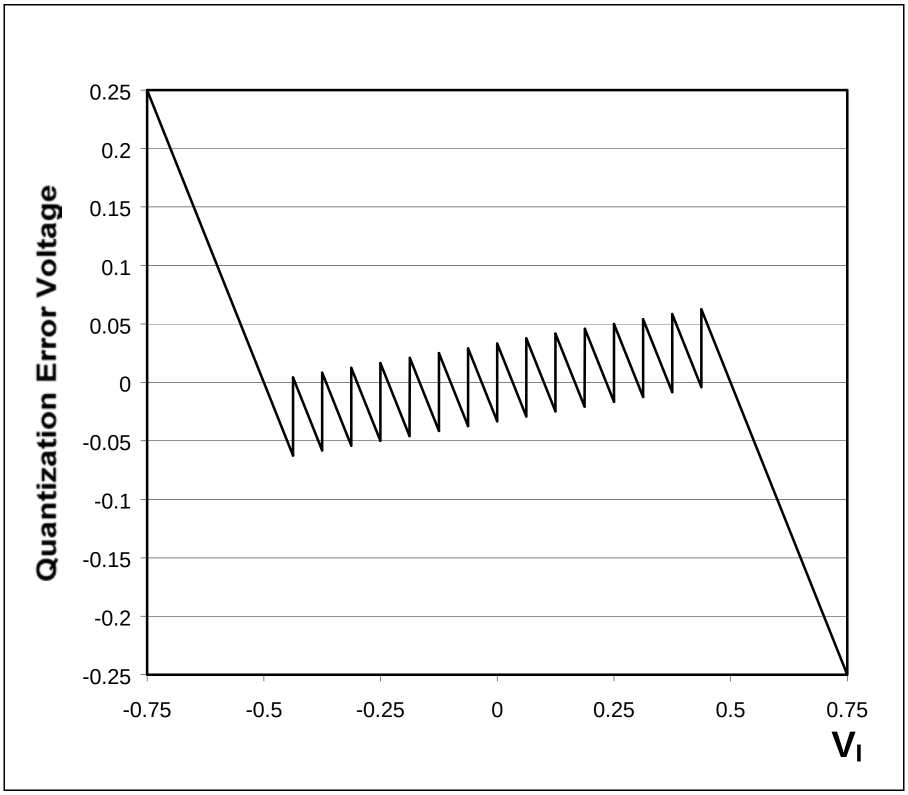

Figure A2. Quantization Error versus Input Voltage

Effects of Companding

Companding (a shortened form of the words “compress” and “expand”) is a technique by which the dynamic range of a link may be increased. Consider a link of the form shown in Figure A3.

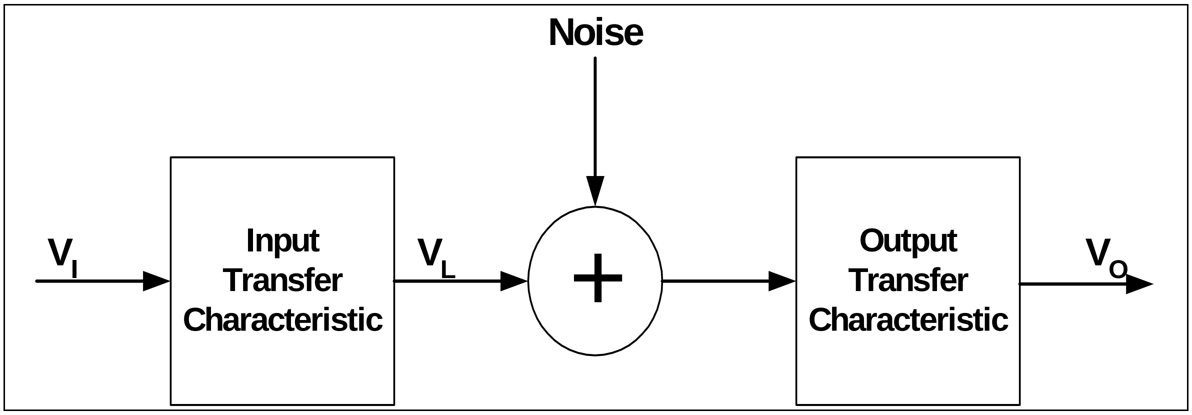

Figure A3. Link Block Diagram

V_I represents the signal at the input to the link. V_L represents the desired signal inside the link. It is assumed that noise is added to the signal as it passes through the link. V_O represents the signal plus noise at the output of the link.

Continued on next page

At the input side of the link, a transfer characteristic similar to that shown in Figure A4 is employed.

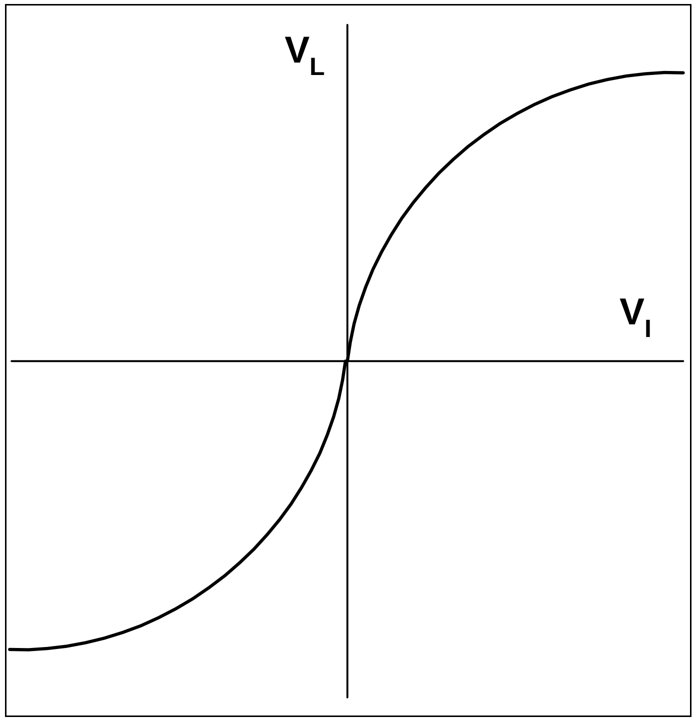

Figure A4. Input Transfer Characteristic

For small input signals, the slope of the transfer characteristic and, hence, the gain is high. As the signal amplitude increases, the transfer flattens out, indicating reduced gain.

Continued on next page

At the output side of the link, a transfer characteristic that is the inverse of that shown in Figure A2 is employed. The output transfer characteristic is shown in Figure A5. Since the output transfer characteristic is the inverse of the input transfer characteristic, signals that pass through the link experience unity gain. However, consider the noise that is added inside of the link. When V_L + noise is small, the output gain is low and, hence, the noise at the link output is attenuated. However, as V_L + noise increases in amplitude, the gain is increased and, hence, the noise at the link output increases.

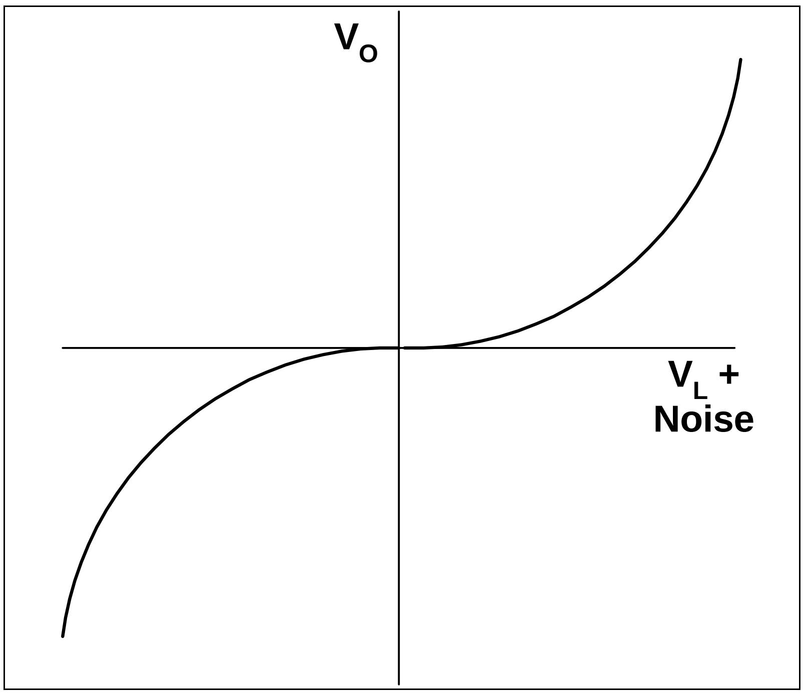

Figure A5. Output Transfer Characteristic

The result is that for small input signals, the link appears to contribute only a small amount of noise. However, for large input signals, the noise contribution from the link increases. In the case of large input signals, increased noise at the output is generally tolerable (that is the larger signals can tolerate more noise being added by the link).

Continued on next page

This technique may be applied to systems that digitize analog signals. A digital companding system digitizes with high resolution (small step size) for small input signals. It uses progressively larger step sizes as the signal swing amplitude increases. A sample transfer characteristic is shown in Figure A6. This characteristic represents the entire link (input, V_I, to output, V_O). Small input signals are digitized with small step sizes and consequently small amounts of quantization noise. However, as the signal swing increases, so does the step size and consequently so does the quantization noise.

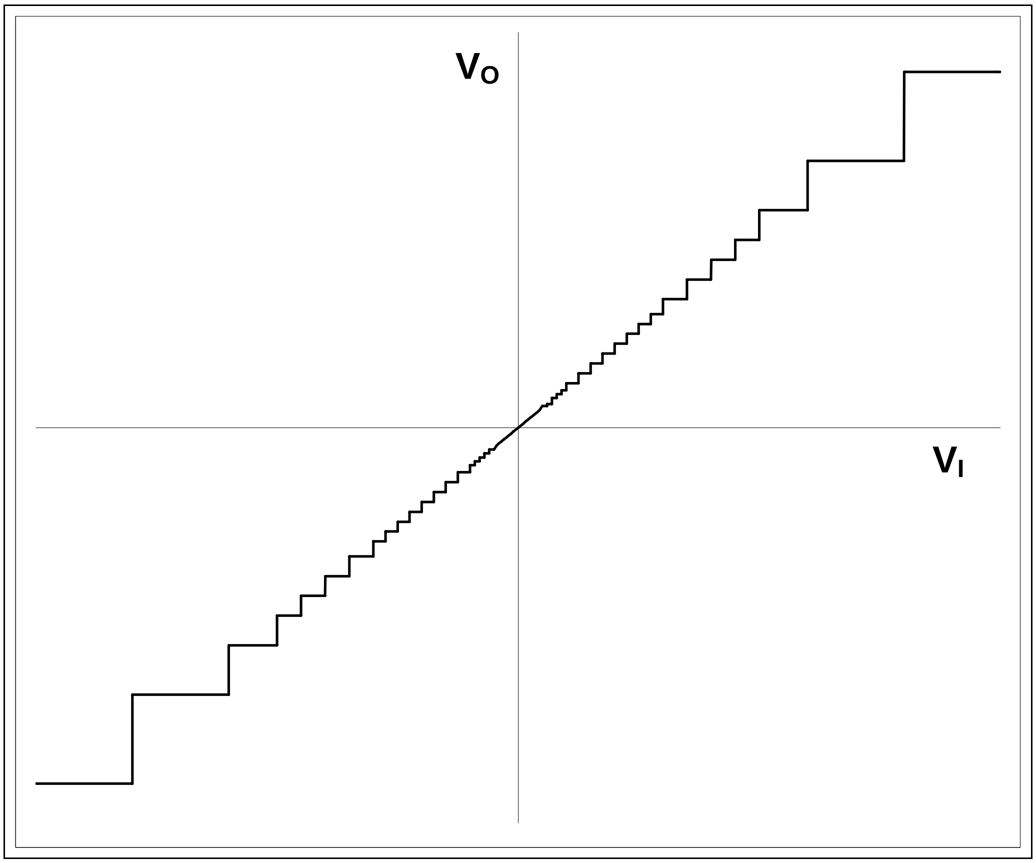

Figure A6. Transfer Characteristic for Digital Companding

If the step size used for small signals were maintained over the entire input range, the number of output states would be increased from that shown in Figure A4. Consequently, the number of bits required to represent these output states is smaller than that required if the small step size were maintained over the entire input full-scale range.

Continued on next page

Consider, for example, a companded system where the step size for small signals was 1/4096th the input full-scale range. If this step size were maintained over the entire dynamic range, 12 bits would be required to represent all of the possible states. However, if the step size were increased for large signal swings such that only 256 possible output states existed, only eight bits would be required to represent the output. For small signals, the output signal to quantization noise would be equivalent to that in a 12-bit system. For large signals that could tolerate more noise, the quantization noise would increase. Hence, performance approaching that of a 12-bit system would be obtained using only 8 bits.

Resulting NPR

The quantization noise in a companded system changes as a function of signal level. This change has an impact on the NPR curve. A typical NPR curve for the 4:1 bdr system is shown in Figure A7.

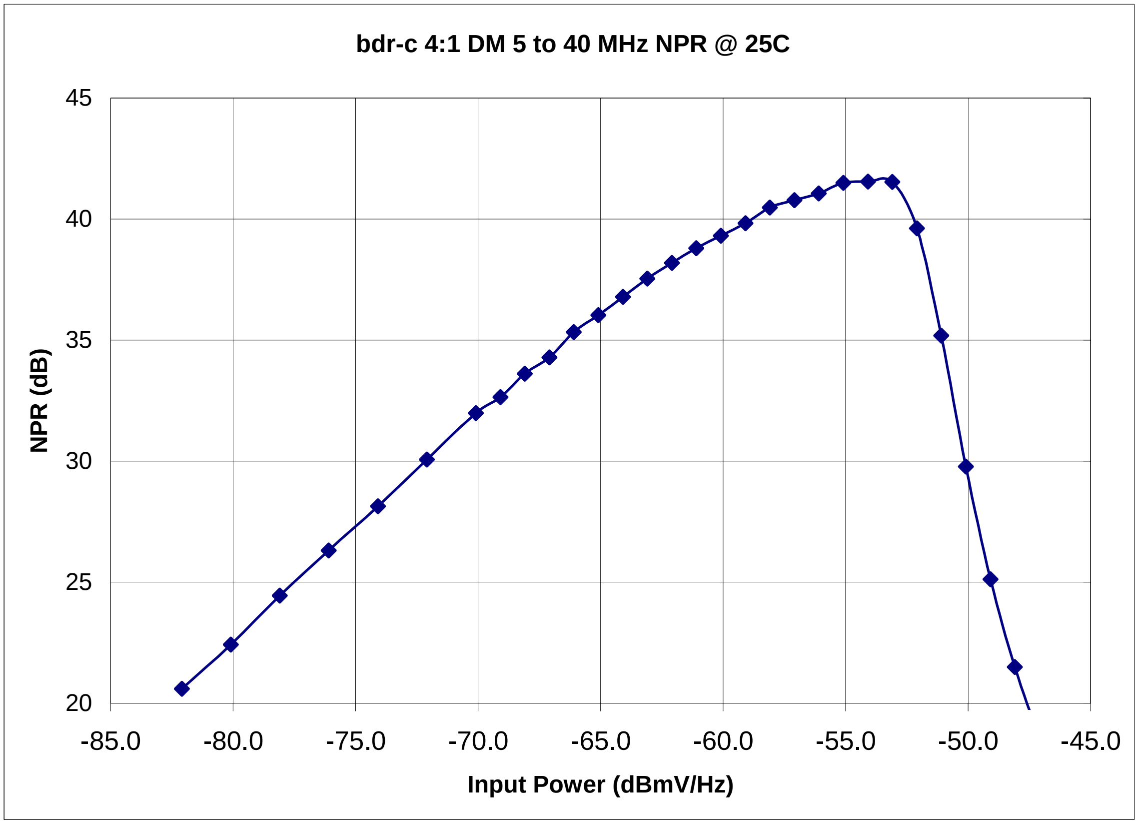

Figure A7. NPR for 4:1 bdr

Continued on next page

An alternative way to present the information in the left half of the NPR curve is shown in Figure A8.

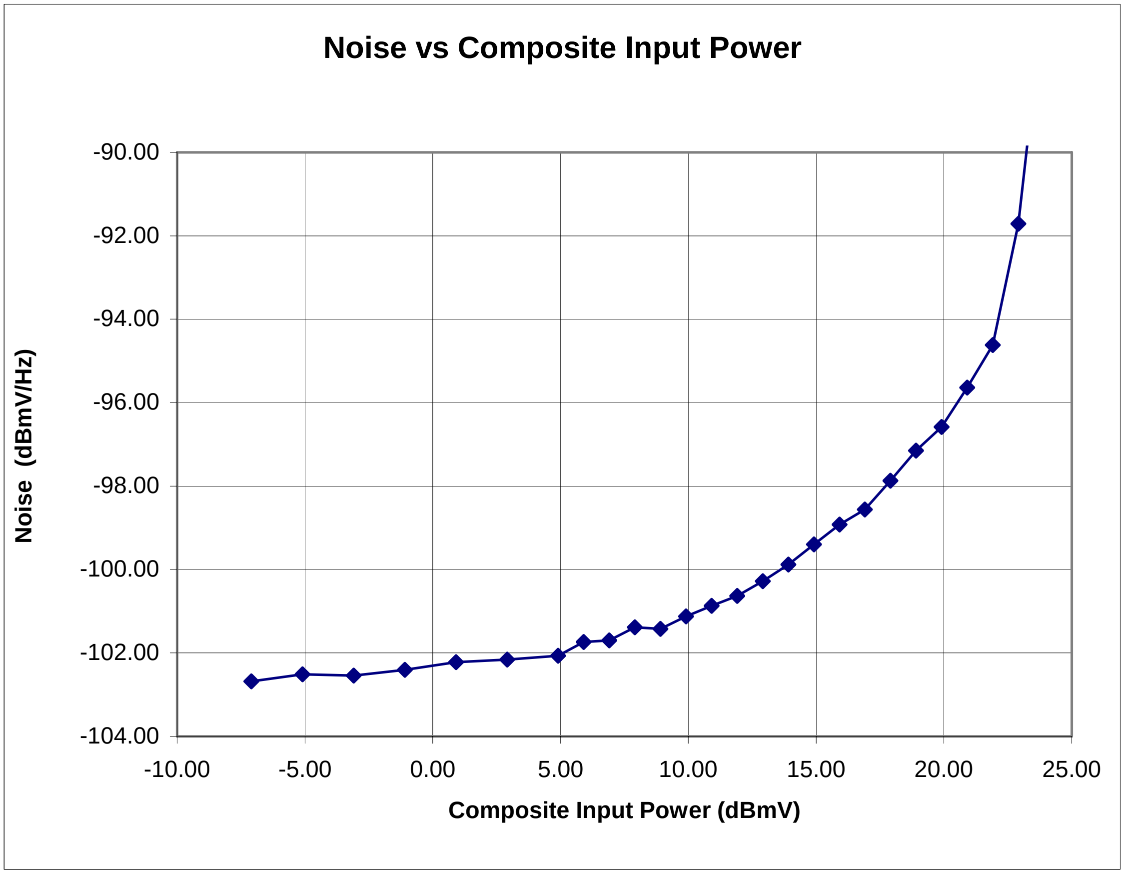

Figure A8. Quantization noise changes as a function of input level

At a composite input level above approximately 22 dBmV, the noise is dominated by clipping distortion. However, below 22 dBmV, the noise is results from quantization error, ***NOT DISTORTION OR CLIPPING***. As Figure A6 indicates, the quantization noise floor changes around 7 dB as the composite input level is increased from –5 dBmV to +22 dBmV. This quantization noise results from the individual steps in the A/D – D/A transfer characteristic.

During the generation of the NPR curve, changing the signal amplitude while maintaining a constant bandwidth varies the composite power. However, the results are the same if the signal amplitude is held constant and the bandwidth is varied. Thus, as loading on an individual link changes, the quantization noise floor changes. This characteristic is critical for understanding overall system performance as multiple links are combined.

Scientific Atlanta, A Cisco Company
5030 Sugarloaf Parkway, Box 465447
Lawrenceville, GA 30042
678.277.1000
www.scientificatlanta.com

August 2008 Printed in United States of America Part Number 4000819 Rev B