ADVANCING RETURN TECHNOLOGY......BIT BY BIT

Dr. Robert L. Howald
Manager, TNS Systems Engineering
General Instrument Corporation

Abstract

Traditional HFC systems have implemented the return path, a 5-40 MHz bandpass in North America, using bi-directional RF amplifiers in the coaxial plant, and a return path laser in the fiber optic node, driving a return fiber for the optical trunking. At the hub or HE, the optical power is converted back to RF. The technology is the same analog AM-based optics approach used to transport the broadcast forward path video signals. There are numerous design and implementation issues that make this approach difficult and costly. These include analog laser specifications, laser second order response characteristics, optical link length constraints, optical receiver specification, and testing of the components. All of these contribute to the overall cost issue of developing high performance analog optics. Node laser issues are exacerbated by the fact that this component must operate in an outdoor environment, specified over a very wide temperature range.

An improved approach implements a digital transport method at the point that the RF plant terminates, and the optical trunk begins inside the node. To do this, the analog laser technology is replaced by a high-speed baseband digital technology of the type that has been used in the telecommunications industry. As a digital signal, immunity from the troublesome analog laser impairments is obtained and longer distances can be covered, potentially avoiding the need for hub repeater hardware required in analog systems, among other benefits.

This paper describes analytical and design issues associated with digitizing the return path at the node. It can be shown that the analog-to-digital (A/D) converter is a mathematical analogue to the AM modulated laser technology traditionally used. We can treat the A/D converter quantization noise as the effective "analog" optical link noise. This can be correlated with the known performance capabilities of the lasers currently used. Additionally, the distortion performance, in particular the laser clipping aspect, is replicated identically by A/D input thresholds. Conveniently, a complete technology upgrade can be achieved, while the key concepts of the mathematical analysis remain virtually unchanged. The second order and third order distortions also can be kept very low in A/D's and, furthermore, the second order distortions do not degrade to poor values as analog laser can. This can be an important issue in broadband applications. Measured performance has been taken of NPR and of BER on a loaded return, showing how the dynamic range analysis above can be applied to the digitized return as it is used in analog systems. Ultimately, the strength of the digitizing technique is

furthered by the functions and processing that can be applied once the information is represented completely as bits. Some of these concepts will be introduced.

THE DIGITAL RETURN PATH CONCEPT

The basic elements of a digitized return path are shown in Figure 1. The idea is quite straightforward. The RF piece of the CATV plant in the reverse path terminates at the node. As such, this is a logical demarcation point to terminate RF signals. Until relatively recently, the 35 MHz (North America) return path bandwidth required an RF solution, and thus employed analog laser technology for reverse transport – essentially the same technology used for broadcast analog video in the forward path. The composite return path waveform AM modulates a laser, and the light intensity varies per the RF signal applied as the modulating signal. The concept being discussed here is about transporting this reverse path signal using baseband serial optical transport – representing the signal by encoding it entirely as ones and zeroes. As shown in Figure 1, there are three main elements to this. They are

1) Converting the composite reverse path waveform to a sequence of digital words whose value represent analog signal samples

2) Arranging the word into a serial stream with appropriate synchronization information to identify the boundaries between words and to recover timing of the bits themselves

3) Converting the electrical digital signal to an optical digital signal, and transmitting the optical ones and zeroes across the fiber

4) Inverting the process at the receive side

This paper will describe the unique parts of each function, which, together, enable this concept. Performance analysis will follow, along with measured results from laboratory prototypes. Finally, the discussion will touch briefly on the next-next generation of digital implementation.

ENABLING TECHNOLOGIES

Analog-to-Digital Conversion

Analog-to-Digital converters (A/D's) have been advancing nearly as rapidly as DSP and VLSI themselves. Digitization of an analog signal for the purposes of allowing efficient transport between locations is obviously an approach with strong roots in the telecommunications industry. In that case, however, the A/D conversion can be quite pedestrian indeed. Voice bandwidths are on the order of 4 kHz, and the channel is baseband and switched in nature, so there is no aggregate build-up of analog bandwidth. In an HFC reverse path, however, both of these simplifications are thrown out. The HFC return is a shared channel, so each user sharing a node consumes bandwidth, effectively aggregating the total possible consumed bandwidth upward with more users. Rather than low frequency baseband, the HFC return is a 35 MHz wide bandpass. Additionally, individual signal bandwidths that can be placed on a cable far outstrip the stingy bandwidth capabilities of the twisted pair to the home

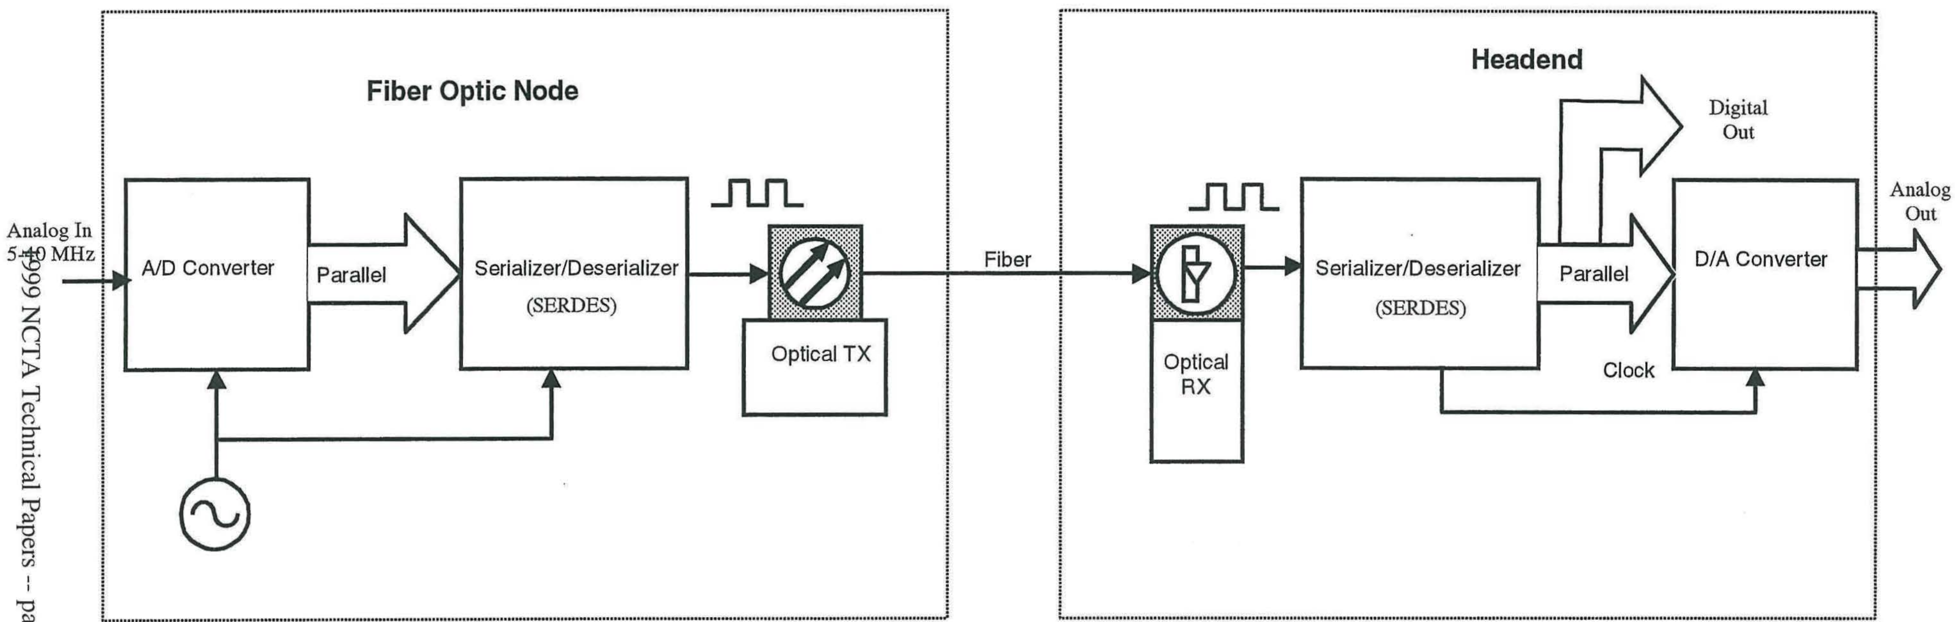

Figure 1 - Digitized Return Path - Basic Operation

originally optimized for voice. This set of differences create the requirements for a significantly different A/D beast than is used in the telco world. Manufacturer's meeting of the requirements for higher performing, robust A/D components necessary for implementation on the HFC return path has come to pass relatively recently. Some key performance items are sampling rate (speed), resolution, and environmental robustness. These first two performance parameters are discussed in more detail below.

Sampling Rate

Nyquist theory tell us that the sampling rate necessary for undistorted signal conversion or reconstruction is twice the signal bandwidth. Most textbooks begin with the most fundamental way to look at this, which is that the sampling rate must exceed twice the highest frequency of the signal. Note, however, that these introductory descriptions assume a baseband (lowpass) signal, and that such a description is really a subset of the more general sampling/ bandwidth relationship. Either view works for this case, and, if we assume that the return path to be converted is 40 MHz worth of analog bandwidth, then an A/D must be sampled at 80 MHz minimum. Now, glancing back into those old signal processing texts, recall that the digital spectrum is periodic, repeating the analog spectrum around every integer related clock component. Filtering that takes place before A/D and after the digital-to-analog converter (DAC) is required to maintain spectral purity through the processing. Practical guidelines for filter design result in the sampling rate to be reasonably higher than the minimum theoretical clock rate, such as 100 MHz clock. For this example, to implement return path digitization, words representing the analog waveform in binary form will come out of the A/D parallel outputs at the rate of 100 million words per second.

Sampling a signal at 100 MHz requires some very intricate design complexities. Solving these design issues has occurred over time. It makes great textbook reading showing nice sampled impulses of some make-believe waveform at discrete time instants. However, there are obvious complications to this picture. Any real sample instant will inherently be *not* infinitely small. Designs must be made to snatch a rapidly varying waveform and hold it steady. A/D terminology associated with these high performance sampling issues are *flash* A/D's and *sample/hold* or *track/hold* technology. Additionally, input conditioning amplifiers requiring very high linearity, but now across RF bandwidths. This is a great challenge outside of an A/D, much less as part of the device. The rapid sample taking puts pressure on the clocking and jitter capabilities, often surfacing in literature as *aperture jitter*. Also, of course, the output of the A/D is a digital word, so advances in fast logic using standard logic families pays off for this device. Lastly, the word formation has moved the digital rates and bandwidths into the realm of RF reality. This means matching, reflections, impedance games – the works. Perhaps most notable in this arena are the circuit design and layout issues. Every RF designer, particularly oscillator designers, know that RF and digital don't mix. Nice, clean RF signals get mucked up quite viciously by voltage-hungry and current flip-flopping logic components. Only now, not only are these sharp-edged waveforms on the same board containing analog inputs, but the clocks and toggling frequencies themselves, as well as the input signal, represent that same class of easily-spreading spectral components. Very careful board design and layout, with specific attention paid to analog and digital return current paths is essential.

Obviously, the evolution of the hundreds of mega-samples-per-second (MSPS) A/D's were critical to the application of the technology for HFC. It has only been within the past few years that such parts have become available at high resolution and with environmentally robust performance. There are a small number of key vendors who play in this arena.

Resolution

Any new technology development that hopes to capture interest must permit a smooth transition into its implementation. From the standpoint of hardware deployed at the node, there are several new pieces. However, because optical modules are often plug-in type packages to allow design flexibility, it is fully anticipated that the hardware aspect of the node will not be very significant. Additionally, as will be discussed below, there is fundamentally no change in alignment or return path maintenance due to the approach.

Beyond the physical implications, however, there are the performance aspects. Current return path technology – AM modulated (analog) lasers – has a certain level of well-characterized performance. It also has selectable grades of performance from low power Fabry-Perot (FP's) lasers to high power Distributed Feedback types (DFB's). There has been substantial return link characterization work going on throughout a host of companies in the industry. Because of this, SNR and distortion numbers achievable by common optical technologies available today are well known. Clearly, all of the effort that has gone into what can be effectively carried on the return path implies that equivalent performance is a logical start. However, it is worthwhile to note that there are also good arguments as to why eight bits and even six bits can effectively handle particular tiers of digital service, which opens the doors to bandwidth savings. Most notably, SNR's in the range of 40 dB have large margin above what is necessary for most anticipated services even ignoring coding gains (consider 16-QAM @1e-8 needs 22 dB uncorrected). Also, most return paths are not constrained in performance by their thermal noise component (AWGN), but instead by transient and burst noise phenomenon. Having increased SNR does not efficiently guard against these phenomenon. It is wise to assure dynamic range headroom in return performance, although overload problems are rarely ingress-related. But the amount of headroom necessary and the cost to provide it, given the link margins available, is a trade-off exercise worth evaluating.

Consider the SNR equation above. Applying a 6-bit conversion, and no oversampling, an SNR less that 40 dB is obtained. Furthermore, this is a maximum (full-scale) number that cannot be achieved in the real application because of the back-off necessary to stay within the A/D dynamic range. On the other hand, an 8-bit conversion lands in the 50 dB range. Again, there is signal back-off from peak required, but the SNR numbers now fall into the range known to be achievable and acceptable with current technology. Going through the noise and distortion numbers as we will below, it can be seen that to achieve customary performance requires the high speed A/D's to have between about 8-12 bits of resolution to be similar to the range of analog capabilities. This will be quantified in a discussion to follow.

Serializer-Deserializer (SERDES) IC's

Major efforts in serial transport technology are currently developing in two arenas – Gigabit Ethernet and Fibre Channel. The results of these standards-based developments can reap rewards, too, for the cable industry. Handled properly, these powerful, new IC's can perform all of the essential functions necessary to implement a digital return from node to HE with complete transparency. The magic in the latest advances for HFC is, again, the Gbps data rates being pushed, and the level of integration being achieved.

The basic functions achieved are fourfold.

Parallel-to-Serial Conversion (and Serial-to-Parallel Conversion)

Parallel digital words at rates around 100 Megawords per second are being delivered to the SERDES chips. The parallel words must be made into a serial stream to transport over the optical link. Thus, the SERDES must latch in the information, multiply up the clock rate, and deliver the increased serialized rate to the optics. For a 10-bit device, the result is serial rates at about 1 Gbps.

On the receive side, the inverse obviously is performed, and the output of the SERDES receiver chain is a digital word handed off to, for example, a D/A converter for analog signal reconstruction. Or, the word is handed off in digital fashion to a direct interface to the application's digital receiver.

Optical Drive

As the standards mentioned develop, compatible optics is also developing. Current technologies can often mix-and-match and be compatible, or very close to it, with just basic modifications, such as signal coupling, pull up/down, impedance transformation, etc. Part of the similarity rests solely on the fact that, at such high data rates, some form of emitter coupled logic (ECL) is unavoidable.

Timing

The parallel-to-serial scenario painted above is a simplistic one at best. The receive IC is delivered only one piece of information – a stream of bits. From this, it must achieve the timing to detect the bits (developing a synchronized clock to sample them at the optimal instant). As might be expected, each side of the link plays a role in this. Clock recovery systems and requirements have become well developed over the years of explosive digital growth. At the transmit side, the role is to encode the data such that there is guaranteed to be enough data transitions for the clock recovery circuit to acquire and hang onto. Thus, the transmitter employs encoding that provides randomization to guarantee transitions in all situations. A simple case that explains the problem would be a very lightly loaded TDMA return, such that the A/D output words are frequently all zeroes. Without transitions, the clock frequency component necessary - embedded in the data sequence through the bit transitions - cannot be tracked effectively by the receive PLL. The role of the PLL is to re-generate this timing to yield effective serial sampling. Subsequent to bit timing, the word timing information - now divided down from the serial rate - is needed at the output for regeneration of the parallel word.

Framing

While it is obviously necessary to achieve bit timing, it is also necessary to know which sets of bits belong together as one analog sample. In other words, it is necessary to make sure that the bits of some 10-bit word representing an analog sample on the transmit side get delivered together to the D/A at the other side. This is another problem with well-developed solutions in the networking world. Solutions center around sending known, predictable patterns signifying word boundaries that can be detected. For HFC, caveats exist. For example, the basic functions of the standards-based chips involves the movement of packets to and from the IC's from other digital sources. This allows, with minimal intrusion, transmission of demarcation characters to be sent between payload packets, particularly with the high sampling rates achievable creating many time slots. For HFC, however, some more intricate design functionality is necessary, as the A/D operation is delivering words constantly in real time to the SERDES function. Additional overhead must be inserted in more creative fashions.

Optics

High-speed baseband digital optics has been around for quite a while. What completes the needed requirement set for HFC is modular, small, wide bandwidth modules, encompassing the data rate needed to capture the RF bandwidths elevated via sampling and serialization. Add field robustness, dropping costs of optics, and increased data rates on the short term horizon, and the units now become both relevant to current needs, and capable of expanding as the need for bandwidth grows.

LINK PERFORMANCE: MATHEMATICAL ANALOGIES

RIN & Quantization Noise

The design complexities discussed above result in the A/D being, of course, non-ideal. There are many ways to express this. There are also ways to express how an ideal A/D *should* behave. The ideal SNR of an N-bit A/D with a full scale sinewave input (peak of sine wave at maximum input threshold limit) is

$$SNR = 6N + 10 \text{ Log } (f_s/2f_{bw}) + 1.76 \text{ dB.}$$

There are a couple of key items to mention here before moving on. Note that SNR moves as 6 dB/bit of A/D resolution, a commonly reference relationship and a good one to remember. The second term points out the benefit of oversampling. The term f_s represents sampling rate, while f_{bw} represents the bandwidth occupied by the input signal of the A/D. The noise associated with A/D quantization is basically flat across the digitized bandwidth. Higher clock rates do not change the noise power, but widen the digitized bandwidth by increasing the clock frequency, thus lowering the noise density. The result is that over the Nyquist bandwidth of interest, the noise power is lower. Finally, the last obscure-looking term actually finds its place by calculating the noise power under the assumption that it is approximated as a Uniformly distributed probability density function PDF. The result is a multiplicative factor associated with the variance of a Uniform PDF that concludes in 1.76 dB. Interestingly, this is a real case of a noise distribution having a Uniform probability density, or approximately so - not an easy thing to find in nature. The usual additive noise processes dealt with are Gaussian in fact or by assumption. Because of the source of this noise, in this discussion we will label the SNR as SQNR - Signal-to-Quantization Noise Ratio - for clarity. It is useful to separate this contribution from other thermal noise contributions in a link, such as noise funneling in the reverse path. In the HFC application, reconstruction of the signal to the analog domain also entails accounting for the noise figure contributions in the analog processing that follows. Despite the nomenclature, the SQNR is not handled any differently in system analysis.

Full Scale Input

Converter SQNR varies as a function of signal level. The reason for this is the absolute nature of quantization noise described above. The possible error of the digital encoding is still half of a quantization step size regardless of whether the signal is large or small. This may sound inconvenient, but in practice turns out not to be so for a couple of reasons. First, the SQNR is based on a CW analog measurement, in keeping with conventional practice across industries. Of course, it is not reference to a 4 MHz bandwidth, an item unique to the television industry. It is instead referenced to the Nyquist bandwidth for the conversion rate. However, this brings us to the second reason. That is, the way SQNR is applied is consistent with how specifications are being written for analog return path products, or at least how they are best written from a purely technical standpoint. That is, the SQNR over the whole bandwidth is specified. For the A/D case, it is given at maximum CW input level. This level is called the A/D *full scale* input level when the input is a sine wave with peak amplitude equal to A/D's maximum input level (in the simplest coding example, the amplitude corresponding to an of 11111111). Sometimes, the manufacturer will specify the SQNR just below full scale.

The Nyquist bandwidth is the well-known maximum signal bandwidth that can be delivered to the A/D without reconstruction problems associated with the frequency domain aspects of discrete processing. Basically, if the clocking frequency of the A/D is not at least twice as high as the A/D input bandwidth (we are treating bandwidth as a lowpass 42 MHz for purposes of this discussion, although lowpass bandwidth is not strictly the requirement for proper processing). The reconstructed output is not uniquely defined by discrete samples. Thus, if a 100 MHz clocking of the A/D is assumed, the Nyquist bandwidth is 50 MHz. In practice, there needs to be some spectral margin below the Nyquist frequency of 50 MHz for reasonable filter design, much like in diplexer design.

The Nyquist variable is not a major issue as far as performance calculations are concerned. The fact that the Nyquist bandwidth varies is accounted for in the SQNR equation given previously. However, in practice, the A/D clock rate never changes in a given system, so there is not a need for repeated manipulation to calculate SQNR. In addition, a manufacturer's specifications of SQNR apply over a range of clock rates, under the assumption that the true user SQNR is determined by the equation above once clock rates for the application are established. For high performance devices, the SQNR may, indeed, be given at several clocking rates. Very high clock rates, and especially the higher resolutions, challenge the A/D performance capabilities for noise and distortion.

Back-Off

As described, the SQNR is typically defined for a full scale sine wave input, or very close to it. Of course, this is not the look of a practical input signal. The stochastic characteristics of the input signal to the A/D are generally related to the loading in the reverse path relative to the maximum recommended design level. However, as in analog optics, under the assumption of constant power-per-Hertz loading, the link design levels and noise and distortion specifications are typically built around the assumption of a fully loaded return. This assures any loading condition will meet the requirements. For the fully loaded case or even the moderately-to-heavily loaded case, the input signal to the A/D looks noise-like for the periods of time that there are simultaneously transmissions. The light loading cases are less threatening from several standpoints in link analysis. First, there is possibility of not using uniform spectral loading if it is desired to dedicate more power to some channels and a straightforward approach to accommodate this is available. This translates to higher SNR. Second, light loading in a constant power-per-Hz approach means clipping is not an issue.

The approach to bounding the anticipated input signal dynamics from a design standpoint is to assume the composite reverse path signal at maximum suggested input looks Gaussian statistically. Single PSK and QAM channels have a noise like look on a spectrum analyzer, but not the same statistical characteristics as noise. Under the Gaussian assumption, the noise and distortion performance of the quantization becomes a tractable problem. In fact, any PDF we can assign would create a definable problem, but the Gaussian quality allows us to draw on some well-known properties to complete the analysis. The result is that it can be shown precisely what optimal input *back-off* is - the number of dB below full scale the input power is run into the A/D. Thus, the noise and distortion performance can be optimized theoretically, and even practically, in the sense that the A/D's *effective number of bits* (ENOB) – the actual resolution of the part in bits degraded by non-idealities. ENOB can be used in place of the performance of an ideal A/D in the analysis. Additionally, the noise and distortion performance can be described in what has become another very useful method of characterizing reverse path performance - noise power ratio (NPR). Elements of the analysis will be described further in a later section.

SQNR, RIN, and Analytical Models

One very convenient side benefit of the digitized return is the ability to focus all impairment analysis and design implementation at the termination of the RF plant at the A/D at the node. This assumes the eventual migration of this family of products into fully digital products in hubs and Headends, a logical and necessary step to fully reap the rewards of having all of the information being moved by transporting bits. Until that time, there still will be noise and distortions added by D/A conversion and processing. The D/A is significantly less expensive that the A/D, but contributes roughly equally to distortion. Following the D/A conversion, sophisticated application receivers now and in near future will immediately re-encode the reconstructed analog output with another A/D conversion in the receiver itself, and obviously wasteful effort. Of course, there will still be traditional boxes and services running on the reverse plant for some time, and some of these may even be analog receivers. Thus, a D/A converter will likely be a necessary option for some return applications. However, it is also the case that these applications are often simple modulation such as FSK, and not significantly effected by additional distortion contributions.

Using the A/D only model, we can conveniently relate quantization noise of the A/D to relative intensity noise (RIN) in an analog laser. In this way, we can relate the performance of an A/D to known analog optical performance. Table 1 shows this comparison using typical return path laser powers for the case of a 100 MHz A/D converter at three different resolutions - eight, ten, and twelve bits. The calculation of densities can be observed from Figure 2. The serial payload (i.e. ignoring overhead associated with any further encoding) for these cases would be 800 Mbps, 1 Gbps, and 1.2 Gbps - all well within the bounds of current optical transceiver technology.

From Table 1, some conclusions can be drawn. An 8-bit A/D provides theoretically about the same performance as the lower quality FP type laser. The 8-bit parts also will tend to provide closer to theoretical performance in terms of ENOB because of their lower resolution. A 10-bit A/D provides an in-between performance category between FP and DFB. The 12-bit device provides nearly the same performance of a return DFB laser.

Table 1 is instructive, but there are two significant items of note with this table that need to be further clarified. First, this table ignores other analog optics impairments in the link (shot noise, receiver noise noise current, iin). Second, the A/D column is based on a full scale input, when in fact the noise like signal characteristics will mean that it will be run backed off (as a laser is). Third, the table masks the distance advantage of the digital approach. Analog link performance degrades about a dB per dB, even worse, at long link lengths. At link lengths that digital optics can achieve, analog link would not be workable without repeater functions.

Table 2 shows how the numbers compare against a sample set of 9 dB optical links that include these impairments using typical receivers. The A/D numbers are adjusted by their performance at the optimum back-off setting minus 3 dB (3 dB of headroom in total power load). Note that this comparison emphasizes the point that the impairments of the digitized link are solely A/D in a complete design.

Table 2 shows that the digitized approach stands up very well to traditional analog performance, without even further considering the distance advantage. By factoring in ENOB rather than theoretical bit resolution, there is still comparable and improved performance for quality A/D converters, where good implies loss of a half a bit to a bit with today's devices.

As another advantage, a similar discussion can take place for the distortion side of the analysis with respect to the clipping impairment. That is, there is a clear relationship to comparing performance to analog optics. In fact, it is even more direct in this case - providing about an exact analytical replication of laser clipping analysis.

Analog Distortion

While the above equation captures the theoretical SNR, the fact that the A/D has an analog path leading up to the digitization means that analog distortions occur in this part. These distortions can be described in very much the same fashion as stand-alone RF or analog parts. However, by the very process of digitization, the distortion energy and its distribution across the spectrum is inherently also a function of the clock frequency. That is, the repetitive nature of the digital version of the spectrum means that traditional harmonic and/or two tone IMD components appear at frequencies related to both the input frequency and clock rate. Second order and third order distortion numbers for this class of A/D converter are, for example, in the mid-60's range for second and third order distortions, all assuming close to full-scale input levels. Like other analog parts, lower levels are generally better for distortion. However, unlike typical analog parts, only considering predictable or low order distortion components can be misleading. The aliasing back in band of intermodulation components can make any distortion component fall in some way into the desired bandwidth, and seemingly meaningless high order components that would otherwise be well out of analog band can show up because of discrete processing. This makes a composite definition for noise and distortion very useful. Examples include noise power ratio NPR, and ENOB.

An interesting interpretation of the quality of an A/D converter is given by its ENOB. Essentially, it takes into account all of the non-idealities of the converter, and expresses this analog performance obtained by the actual part in terms of what number of equivalent bit resolution is obtained. In other words, a 8-bit A/D converter, because of its non-ideal performance, may yield only effectively 7.5 bits of resolution, meaning it can be modeled as a 7.5-bit converter from the standpoint of performance. It is analogous to implementation loss as it is used in modem designs. Modems are often recognized as having some X dB implementation loss, which describes the actual performance by comparing how close it comes to theoretical BER performance. A high-speed 10-bit A/D may have an ENOB of about 9 bits, indicating loss of about one bit due to constraints of the practical implementation. Recall, this would be equivalent to a loss of about 6 dB of SQNR.

There is one other distortion component of an A/D converter, but it is not associated with the non-ideal nature of the part. Instead, it has to do with the nature of the A/D function. It is a concept quite familiar to those in HFC – clipping.

Clipping

The analog laser's clipping impairment effects have been studied quite extensively for both forward and return applications. Products are designed with requirements built around a suggested maximum input range, and this range is

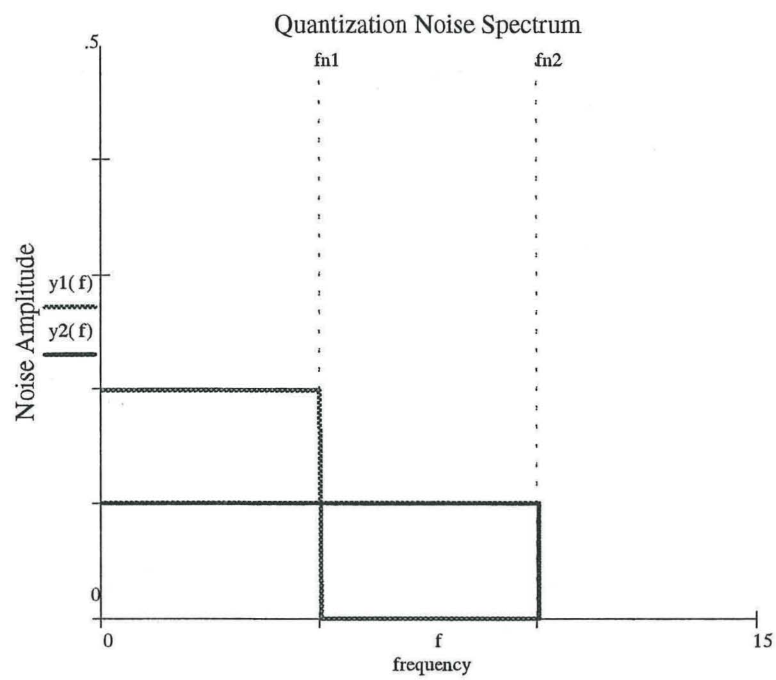

Figure 2 – Effect of Sampling Rate on Noise Spectrum (fn = Nyquist Bandwidth)

Table 1 - Noise Performance: Return Path Laser RIN vs. A/D Quantization Noise

Device	RIN	Quantization Noise (dBc/Hz @ 100 MHz)
8-bit A/D	-------------	-127 dBc/Hz
Fabry-Perot (FP)	-130 dBc/Hz	-------------
10-bit A/D	-------------	-139 dBc/Hz
12-bit A/D	-------------	-151 dBc/Hz
Distributed F'back (DFB)	-155 dBc/Hz	-------------

Table 2 - Noise Performance: Return Path Analog Optical Link (9 dB) vs. A/D

Device	FO Link Noise	Quantization Noise (dBc/Hz @ 100 MHz)
8-bit A/D	-------------	-112 dBc/Hz
Fabry-Perot (FP)	-119 dBc/Hz	-------------
DFB	-122 dBc/Hz	-------------
10-bit A/D	-------------	-123 dBc/Hz
12-bit A/D	-------------	-137 dBc/Hz

the manufacturer's recommendation for providing what they consider the best trade-off between the noise and distortion contributions. Clipping is considered a distortion contribution, although it is the case that on a spectrum analyzer, it appears as a broadband noise floor increase. This is simply due to the time domain characteristics of clipping events - they are typically very short impulsive events, corresponding to broadband events in the RF plant. At the receiver, of course, they are spread in time by the matched filter prior to detection.

Fortunately, none of the work in understanding laser clipping goes to waste when we talk about the return path digitization concept. There is just as much if not more prior analysis related to the clipping effect as it relates to an A/D converter. There is one physical difference in the mechanism of the clipping event. In the analog laser, the event occurs by negative-going amplitude peaks shutting off the laser by dropping the instantaneous current below the bias. By contrast, the A/D converter is limited by maximum and minimum peaks that exceed the all-one and all-zero words. Beyond this, all of the other well-known aspects of the clipping phenomenon apply from an analysis perspective. For a noise like load, the events where the input exceeds the threshold are short and impulsive. Some may be shorter than a sampling period, and the likelihood of sampling during an event are statistical since the signal itself is. We approach the overdrive analysis the same way. From a BER standpoint, clipping events are largely mitigated by FEC if the event is of enough energy to cause an error to begin with. The impulses are spread by a receiver's matched filter, and contribute an additive impairment. Also, the clipping impairment for A/D's creates an NPR curve that slopes steeply on the distortion side because of the hard clipping, higher order distortion components.

Dynamic Range Performance

Dynamic range, as always, describes a unit under test's performance range in a way that captures it thermal noise limitations on one end of the input signal range - its small signal performance - and its large signal performance on the other end. An amplifier, for example, is characterized by noise figure and some measure of distortion performance, such as intercept points or CTB and CSO. Use of these parameters to describe dynamic range is perhaps the most common usage of the term, although it can be generalized to other concepts.

Noise Power Ratio

An A/D converter can be described in terms of dynamic range, and it is necessary to do so to make comparisons among technologies. More importantly, its performance characteristics must be able to be expressed to allow the device to be analyzed for its contribution to link impairments, just like the rest of the pieces of the RF channel. In the return path of HFC networks, a most convenient way of doing this is through use of the noise power ratio (NPR) test.

The NPR idea is very straightforward, and the results very informative at a glance, which is one of the reasons for its widespread use. The basic idea is to use a noise load as an input signal, thus taxing the linearity of the device by delivering a signal of wide bandwidth and high peak-to-average ratio. A bandstop filter is used to create a notch in the white noise spectrum. The depth of this notch in dB is then a measure of the noise of the device only for low total input power. As the total noise power is increased, the notch will get deeper with each dB of increased signal power, until the linearity constraints of the device are reached, at which point distortion components begin to fall in the notch bandwidth, decreasing the notch depth. This curve - notch depth versus input power - is an instantaneous snapshot of the device noise an distortion capabilities.

Figure 3 shows an NPR performance curve of a typical Fabry-Perot (FP) laser through an 8 dB link, recovered through a return path receiver (RPR). The qualities shown are typical. The noise-only impairment region on the left of the peak changes dB-per-dB with input power, and the distortion region on the right side of the peak degrades much more sharply. This effect is dominated by laser clipping, which causes high order intermodulation effects as the maximum NPR is crested, creating the well-known "crash point" phenomenon. Between noise side and distortion side, the curve peaks at the maximum NPR, corresponding to optimal input power point (but with zero headroom). From Figure 3, it is apparent that an 8-bit device closely resembles the performance of the FP link. Figure 4 shows a DFB on a 9 dB link. This figure closely resembles 10-bit performance.

NPR for A/D's

The use of noise power is especially useful for return path, where multiple, random, uncorrelated, digital inputs tend towards a noise-like characteristic, and behave statistically Gaussian. As described, A/D specifications are sinusoidal-based for convenience, as most analog components. As in most of these cases, the practical signal looks much different. And, also like most cases, what happens when things change is important. In amplifiers, where distortions are generally soft limiting effects that create second and third order distortions, the multiple carriers/noise load effect is to degrade common two-tone sinusoidal intermodulation levels further. It is somewhat complex, but an analytically and numerically tractable problem. Forward amplifiers capture the effect with CTB and CSO. Reverse amps do the same for commonality, although NPR is becoming more common. Hard limiting effects like lasers and A/D's create higher order distortions, making them ideally suited to NPR.

Assuming a Gaussian input to the A/D, the nature of the discrete transfer function makes the problem tractable. An optimum back-off can be derived from the analysis, which turns out to be a function of the number of bits, as might be anticipated. Because of the A/D's discrete nature, the point at which the A/D threshold is exceeded by an input is precisely defined. The input is also defined, by assumption, to be Gaussian. The convenient part of this is that the times when we are concerned about clipping effects are primarily times that correspond with heavy loading, which is also when the input looks most noise like, making the assumption the most valid. Without developing the detailed mathematical steps, the analysis follows these steps [1]:

(1) The clipping level of the A/D can be related numerically to the average input power.

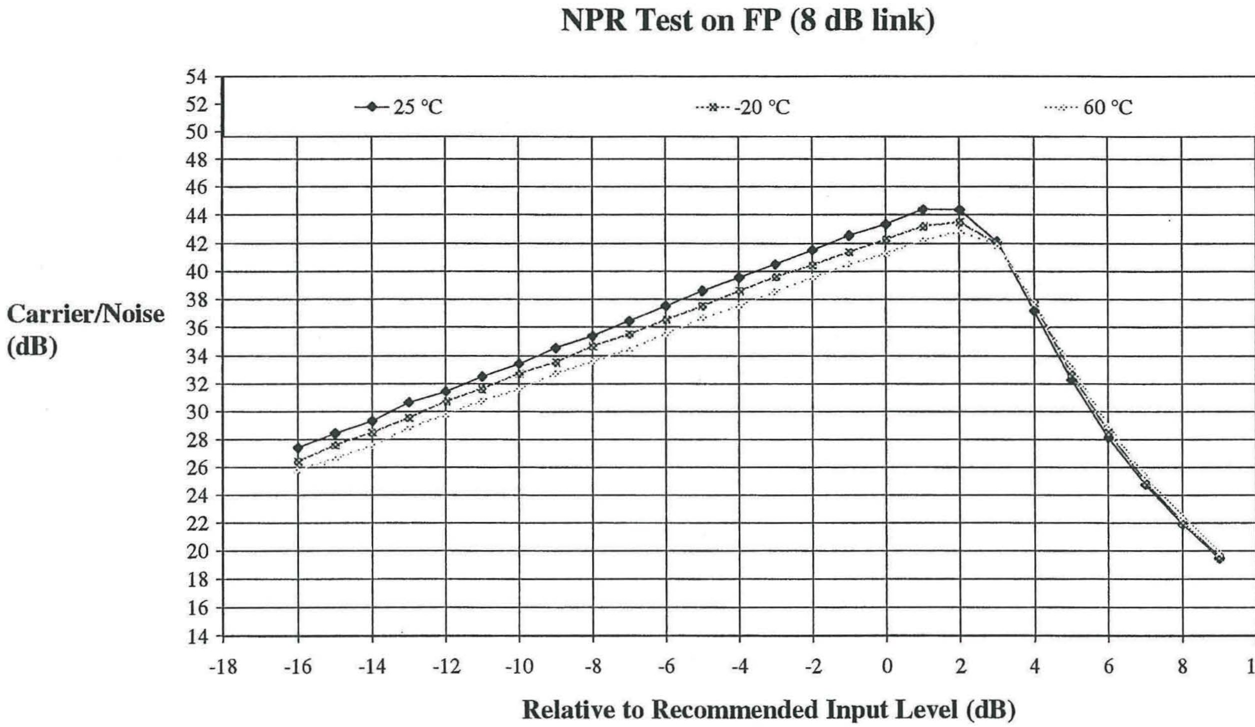

Figure 3 – NPR Performance of an FP Laser on an 8 dB Link

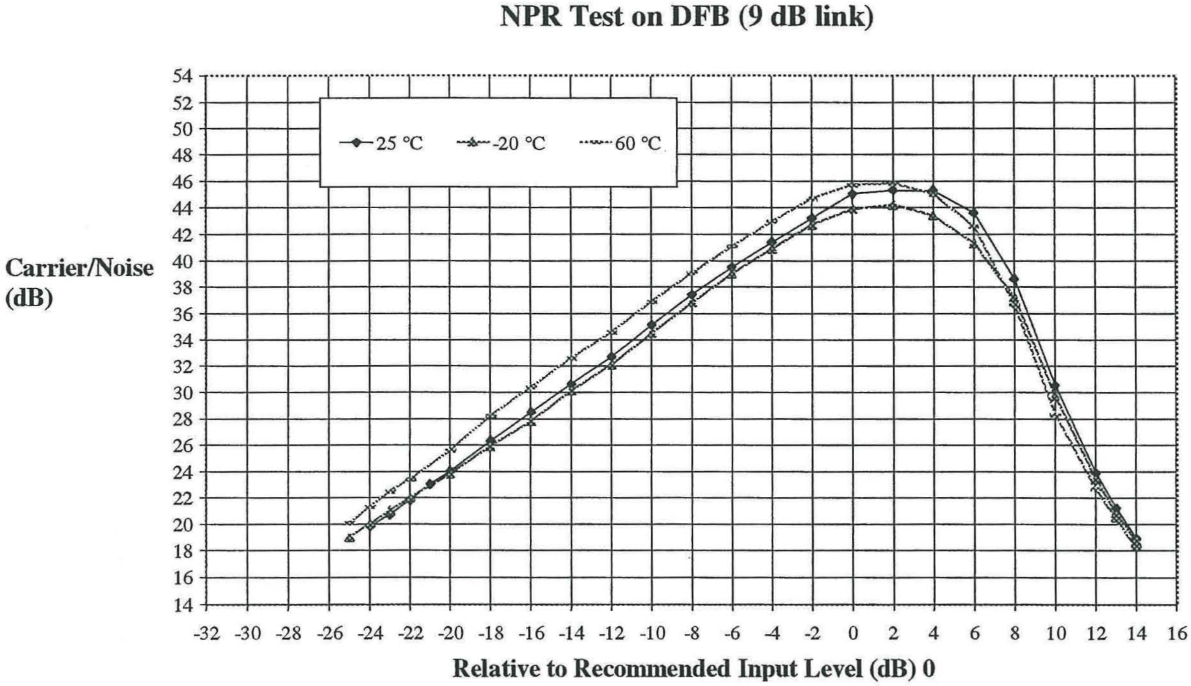

Figure 4 – NPR Performance of a DFB Laser on a 9 dB Link

(2) The input power of a zero mean (AC-coupled) Gaussian signal is its variance.
(3) The rms value of the input voltage is the square root of the variance (the standard deviation).
(4) The A/D's noise power, for a given rms/peak relationship in (1), is calculated from the error voltage squared integrated over the Gaussian likelihood of that voltage being between minimum and peak. The error voltage function of an A/D in the linear region is a sawtooth because of the quantization process.
(5) The distortion power is approached the same way. But, the calculation integrates over the likelihood of the Gaussian input exceeding the peak, and the error voltage linear increases as signal exceeds peak.
(6) The NPR is the input signal power divided by the sum of noise plus distortion power: NPR = $P_{in}/(Q_{noise} + IMD)$.

The results are shown in Figure 5 for 8-bit, 10-bit, and 12-bit A/D converters. Note the 6 dB/bit rule of thumb is easily visualized on the left (noise) side of the curve. Peak NPR's for 8-bit, 10-bit, and 12-bit devices are also about 12 dB apart. For 8-bits, the peak is about 41 dB, for 10 bits about 52 dB, and for 12 bits about 63 dB. The difference is not exactly 12 dB because at the peak of the curve, the distortion contributions are also being felt, not just the 6 dB/bit noise contribution.

Table 3 lists peak NPR, and approximate dynamic range of NPR above 25 dB, 30 dB, 35 dB, and 40 dB. For reference, 16-QAM running at a 1e-8 uncorrected BER requires 22 dB of SNR, where SNR refers to the additive Gaussian thermal noise only contribution. QPSK needs only about 15 dB of SNR. Furthermore, compared to 16-QAM, QPSK behaves more alike when embedded in a distortion floor as in AWGN, because of its robustness and lack of amplitude information in the signal structure.

It is interesting to note as a reference that a back-off of -14 dB corresponds to an OMI (peak) of 20%.

Measured NPR Performance

Figure 6 shows an example NPR measurement, capturing peak NPR performance of an 8-bit A/D (through a high resolution reconstruction D/A). Performance is close to expected, given the ENOB of the tested devices as the "theoretical" consideration. Theoretical performance of an 8-bit A/D degrades from about 41 dB to about 35 dB for a part with a 7-bit ENOB, roughly the specification for the device under test. The actual performance plot, shown in Figure 6, indicates an ENOB of about 7.4 bits. This can be calculated from the measured peak NPR of 37 dB. Thus about a half-bit is lost in the implementation of this A/D. The optimum back-off is 11.9 dB for eight bits. It becomes greater for higher resolution, because the quantization noise contribution is diminished, while the clipping portion stays the same. Thus, it makes sense to back off further into the lower noise for best performance.

THE NEXT DIGITAL WAVE

Processing

By now, most CATV folk are familiar to some extent with DOCSIS, the

Table 3 - Theoretical NPR Performance and A/D Dynamic Range

Device	Max NPR @ Backoff	DR @ 25 dB	DR @ 30 dB	DR @ 35 dB	DR @ 40 dB
8-bit A/D	41 dB @ -11.9 dB	20 dB	15 dB	8 dB	2 dB
10-bit A/D	52 dB @ -13.0 dB	32 dB	27 dB	20 dB	14 dB
12-bit A/D	63 dB @ -14.0 dB	44 dB	39 dB	32 dB	26 dB

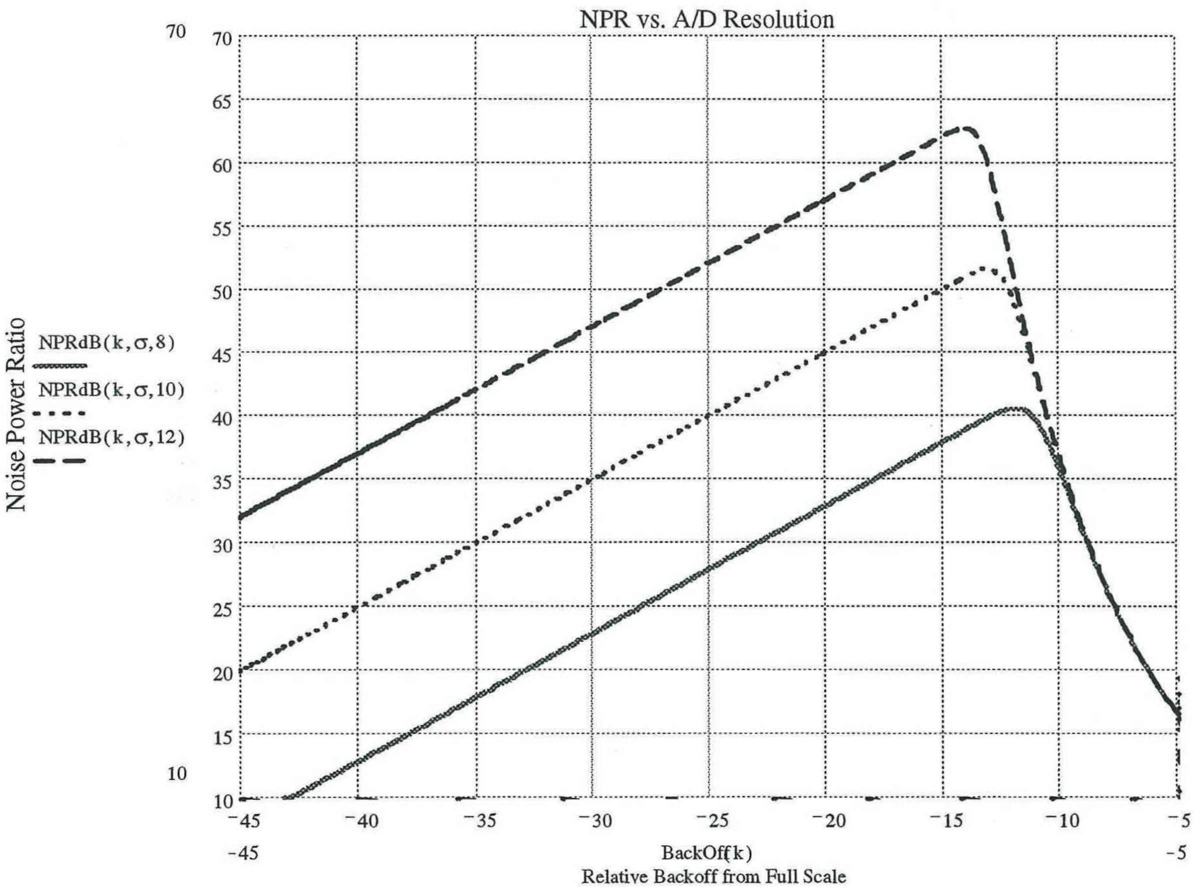

Figure 5 – A/D Dynamic Range Performance vs. Resolution

culmination of cable modem standardization efforts. A detailed read of the specification yields many impressive requirements, and representing truly a state-of-the-art communication system design in all respects. DOCSIS compatible chipsets are a very powerful breed. Similarly powerful chipsets are available in the quickly-losing ground world of ADSL, as the channel being worked with in that case also has a unique set of major impairments to deal with, as well as obvious bandlimiting characteristics. In both cases, functionality available within these chipsets opens up possibilities to network architects to bolster the system flexibility and robustness. And, obviously, they come with the built-in ability to communicate. As the return is turned into all bits, options for processing and taking advantage of processing functionality offers many possibilities for distributing functionality and intelligence throughout the network.

Multiplexers

The current short term transparency of the digital return revolves around reconstruction of the digitized waveform at the HE with a D/A of equivalent resolution (or better). The D/A function, however, introduces the same set of analog impairments as any other analog link – noise and distortion. And, in fact, the high speed products necessary for this application have some specific hardware complexities that make these devices also difficult to implement with high performance (the D/A still, however, is a significantly less costly part than the A/D). Unfortunately, the D/A output will be amplified and split, then delivered to an application receiver. This receiver will then often immediately digitize it and perform the receive, synchronize, and demodulation functions in an all-digital receiver. The wastefulness and illogic of this is obvious.

Looking past the immediate capabilities, it becomes clear that current analog HE equipment will need to take on the look of digital multiplexing products – shipping of the right bit set to the various application receivers, and speaking the right protocols. The fact that digital passage of information is involved opens up many flexibility options, as networking equipment can be controlled and programmed for maximum flexibility.

Transport

The implementation of an all bit stream at the RF-to-optical conversion point immediately brings into the equation the idea of standardized transport, a leveraging these telco-grown technologies for HFC architecture optimization. Use of SONET, SDH, or other standards-based transport simplifies network implementation, accelerates development, and brings with it a comfort level of a proven, robust technology to the HFC infrastructure. The fact that the digitization is pushed out to the node means that hub equipment can become common digital multiplexing equipment. Or, because of the distance advantage, passive equipment of removal completely of hub sites can be a possibility.

CONCLUSION

Here is a piece of non-earth shattering news: the world is going digital. And, what is already digital is becoming advanced digital. Examples abound everywhere – computers, DSP progress, advances in VLSI, even in our own backyard through the evolution of 256-QAM signaling and HDTV. This axiom will apply to the digitized return path also. The list of potential short term benefits are apparent: optical costs, total equipment costs, link distance, distortion performance, component specification, test simplicity. However, perhaps the most benefit is in the years to come, as the conversion of the return path to all bits opens up many networking possibilities associated with processing, transport, multiplexing, and terminating equipment interfaces. This paper discusses the key parameters and definitions for return HFC associated with this approach. It also shows how to relate this technology to performance parameters of interest for proper design of HFC return path systems.

References

Gray, G. and G. Zeoli, "Quantization and Saturation Noise Due to Analog-to-Digital Conversion", IEEE Transactions on Aerospace and Electronic Systems, Jan. 1971.

Acknowledgement

The author is grateful for the contributions of his colleague, Ronald Cambridge. with this approach. It also shows how to relate this technology to performance parameters of interest for proper design of HFC return path systems.

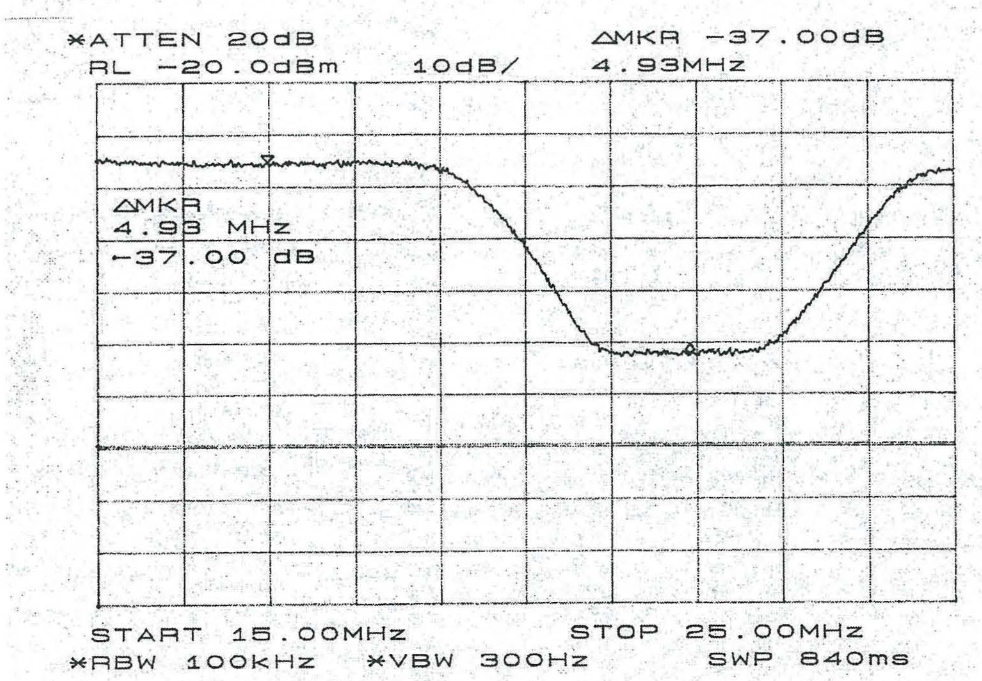

Figure 6 – NPR Performance of an 8-bit A/D