SPECIFYING AND EVALUATING BIT SYNCHRONIZERS

presented at the

1989 International Telemetry Conference

by

John Carlson

currently

Director of Engineering

At

GDP Space Systems

SPECIFYING AND EVALUATING PCM BIT SYNCHRONIZERS

INTRODUCTION

As we enter the 1990's PCM Bit Synchronizers continue to be of major importance to data recovery systems. This paper explains the specification of PCM Bit Synchronizers and provides insight into performance requirements and verification methods. Topics include: bit error rate performance of wideband and prefiltered data, acquisition and sync maintenance, and the effects of jitter, transition density and transition gaps. The merits of multiple and/or adaptive loop bandwidths are discussed, and embedded Viterbi decoders are described. Emphasis is on the new high data rate applications, but the concepts apply to the specification of bit synchronizers in general.

REVIEW OF FUNDAMENTALS

The theoretical Bit Error Rate (BER) performance curves for the standard telemetry PCM codes are shown in figure 1. Theoretical bit error rate performance is based on optimal detection of ideal PCM data in additive white Gaussian noise. Optimal detection implies perfectly synchronous sampling of signal polarity at the output of a data matched filter (optimal symbol-by-symbol detection is assumed for the DM and MDM codes). Ideal PCM data consists of symbols which are square sided (infinite bandwidth) with equal and opposite amplitudes (anti-podal) and equal probability of occurrence.

Bit error rate performance is expressed as a function of signal-to-noise (S/N) ratio. As with all definitions involving a S/N ratio, a noise reference (or measurement) bandwidth must be stated. This does not imply that the noise is restricted to the reference bandwidth; in fact, in most cases the noise bandwidth is considerably greater than the reference bandwidth. For bit synchronizer performance specifications, the noise reference bandwidth is typically chosen to be equal to the bit rate. This gives rise to the expression signal-to-noise in the bit rate bandwidth, or equivalently, energy per bit to noise density ratio, Eb/No. Recalling that energy is the product of power and time, this equivalence is seen as follows.

$$\frac{S}{N} = \frac{S}{No\ BR} = \frac{S\ T}{No} = \frac{Eb}{No}$$

where S is the signal power, N is the noise power in the bit rate bandwidth, No is the noise power density (per hertz), BR is the bit rate, T = 1/BR (the bit period), and Eb is the energy per bit.

Occasionally a reference bandwidth of other than the bit rate is used. This results in a shift of the bit error rate curves along the horizontal axis. However, it is not a change in performance, but simply a change of reference. For example, if a reference bandwidth of one half the bit rate is used, a 3 dB difference in noise power results. That is, a +3 dB S/N ratio with a noise reference bandwidth of one half the bit rate is equal to a 0 dB S/N ratio with a noise reference bandwidth equal to the bit rate. In general the relationship between the signal-to-noise ratio in an arbitrary reference bandwidth, BW, and Eb/No is given by:

$$\frac{S}{N} = \frac{S}{No\ BW}\ \frac{BR}{BR} = \frac{ST}{No}\ \frac{BR}{BW} = \frac{Eb}{No}\ \frac{BR}{BW}\ .$$

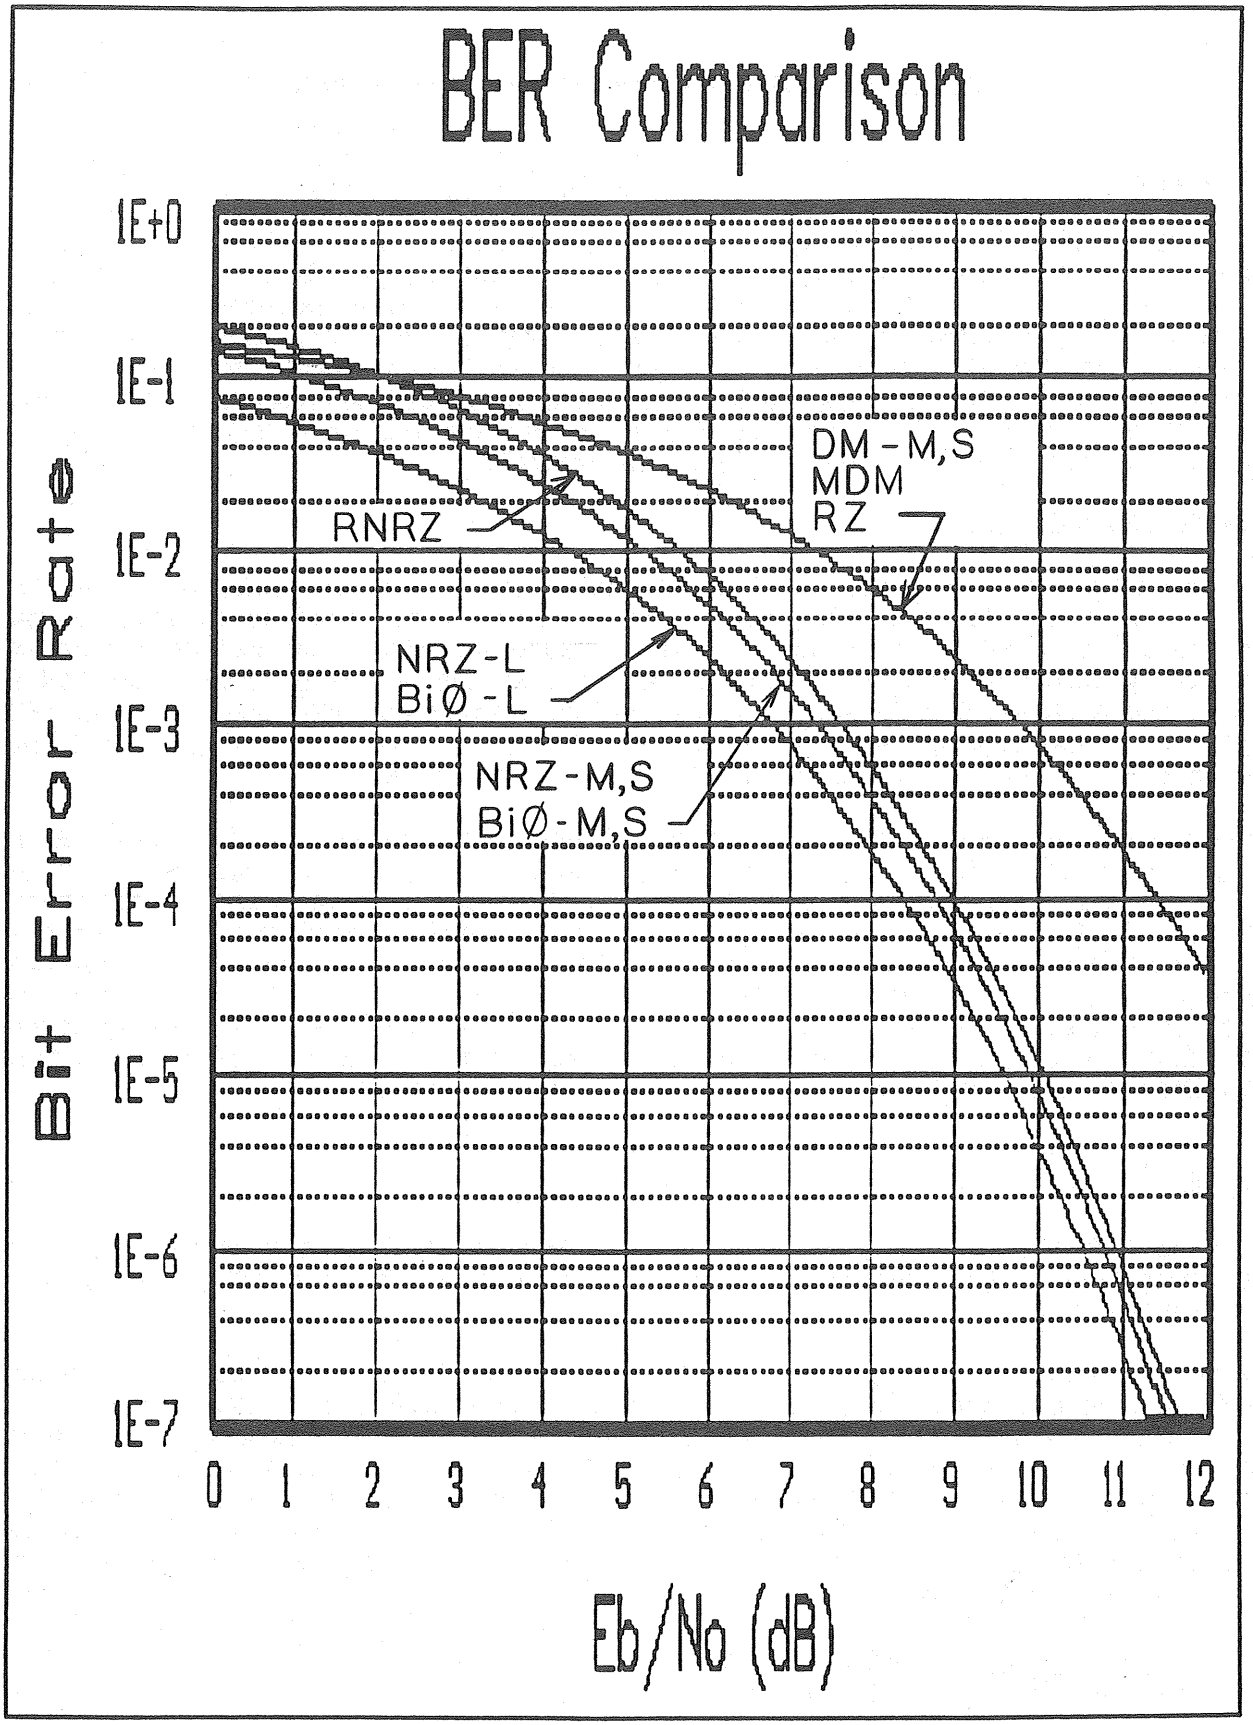

Figure 1. Theoretical Bit Error Rates

As previously stated, theoretical performance applies to the detection of data corrupted by additive white Gaussian noise. In most bit synchronizer applications noise arises from the thermal noise floor of the system front end, as determined by the antenna noise temperature and amplifier noise figures. When adequate bandwidth and amplitude linearity are available, the additive white Gaussian noise assumption is reasonably valid (the main exception to this is the FM discriminator with fluctuation noise, which is non-white (the noise spectrum rises at 6 dB per octave), and pop noise which is non-Gaussian).

Bit synchronizer bit error rate performance is specified as a degradation from theory. A high performance bit sync performs within one dB of theory. That is, its performance is degraded by less than one dB and its bit error rate falls between the theoretical curve and the theoretical curve shifted to the right by one dB.

MEASUREMENT ERRORS

When testing Bit Error Rate performance a measurement tolerance should be added to the specified degradation. Typical measurement tolerances range from ±.2 dB to more than ±.5 dB depending on the accuracy and repeatability of the bit sync test equipment.

A typical bit sync test station is shown in figure 2. The main contributions to measurement uncertainty arise from: noise reference bandwidth inaccuracies, noise power measurement errors, inadequate bit error rate measurement time, and non-white non-Gaussian noise.

Bit error rate measurement uncertainties are directly related to noise reference bandwidth inaccuracies. The bandwidth of a noise reference filter can be calibrated to within ±.1 dB; however, over time and temperature the bandwidth can drift by ±.25 dB or more. It is therefore important to frequently recalibrate the noise reference bandwidth in order to obtain accurate BER measurements.

Errors in noise power measurement are small (±.1 dB) at high data rates when a quality, true RMS power meter is used. However, at low data rates (100 bps) noise power measurements become increasingly inaccurate due to the statistical nature of noise and the typical averaging times of the power meter. Also, as with noise reference bandwidth, noise power should be frequently remeasured to avoid errors due to power variations with time and temperature.

The relationship between bit error rate measurement time and bit error rate is shown in figure 3. The figure shows the measurement time (in terms of bits) required to obtain a given bit error rate with a given accuracy and a given confidence. For example, at a BER of 1×10^{-5}, 50×10^{5} bits must be tested to achieve a ±.2 dB accuracy with a 99% confidence. At high data rates the required measurement times are relatively short, but at lower data rates the required measurements times become prohibitively long. For example, at 5 Mbps a 50×10^{5} bit measurement interval is only 1 second, while at 500 bps almost 3 hours are required.

The most important factor in accurate BER measurement is the noise source. Theoretical performance assumes white Gaussian noise; noise sources are often both non-white and non-Gaussian. The bandwidths of many noise sources are inadequate to support high data rate testing. Bandlimiting the noise reduces the total noise power and therefore reduces the bit error rate. The effect of bandlimited Gaussian noise on ideal bit sync performance is illustrated in figure 4. Note that: 1) the noise bandwidth is the bandwidth of the noise entering the bit sync and is in no way related to the noise reference bandwidth used to measure noise power; and 2) bandlimited noise is a test consideration only; in actual applications the noise and signal receive the same filtering, the noise cannot be bandlimited without also bandlimiting the signal. The IRIG Test Methods for Telemetry

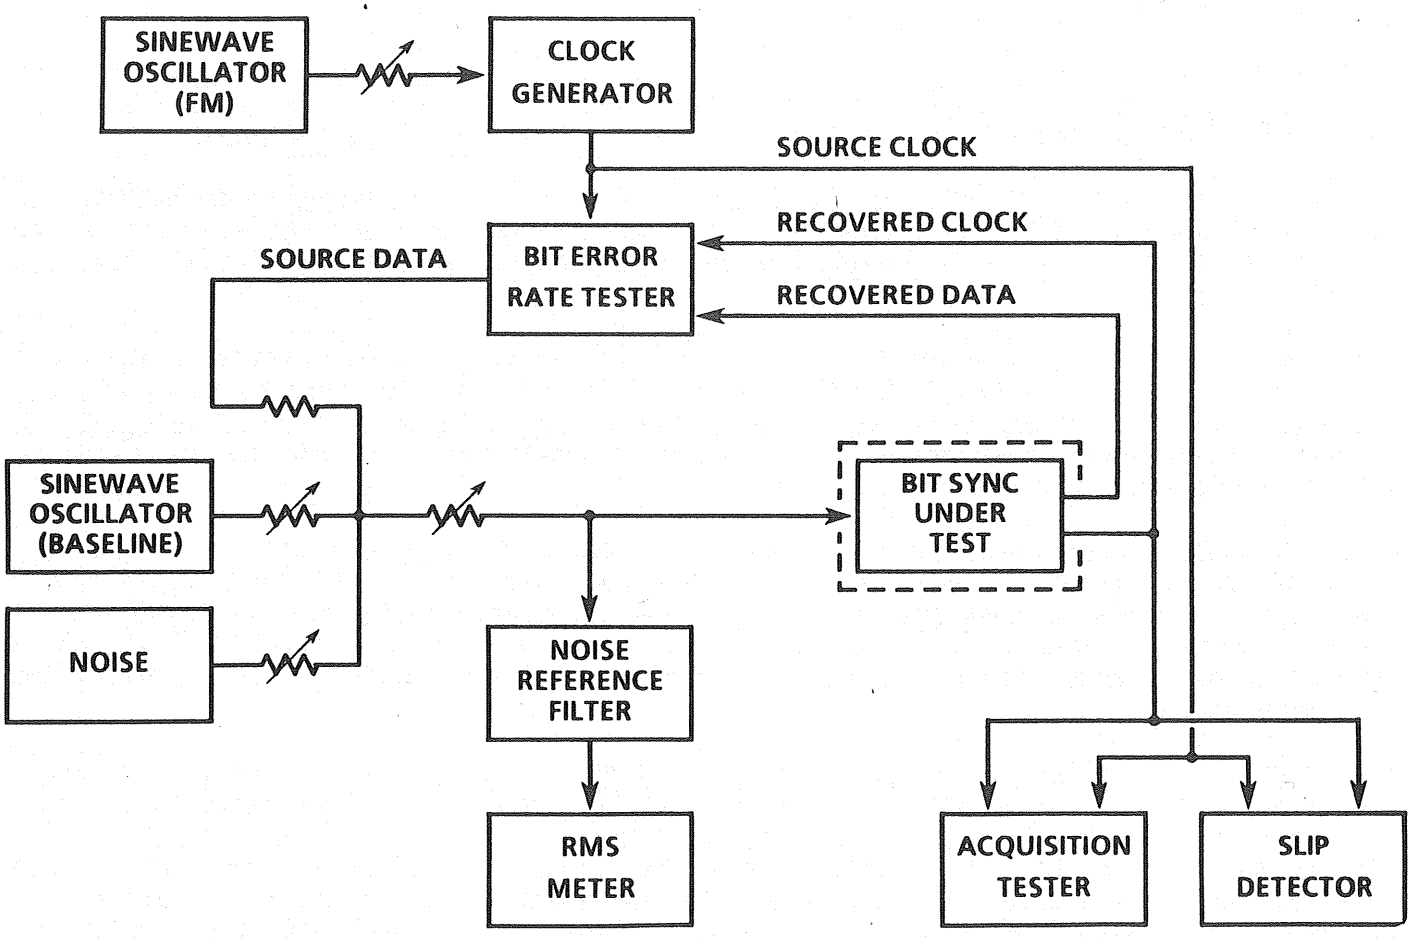

Figure 2. Bit Sync Tester

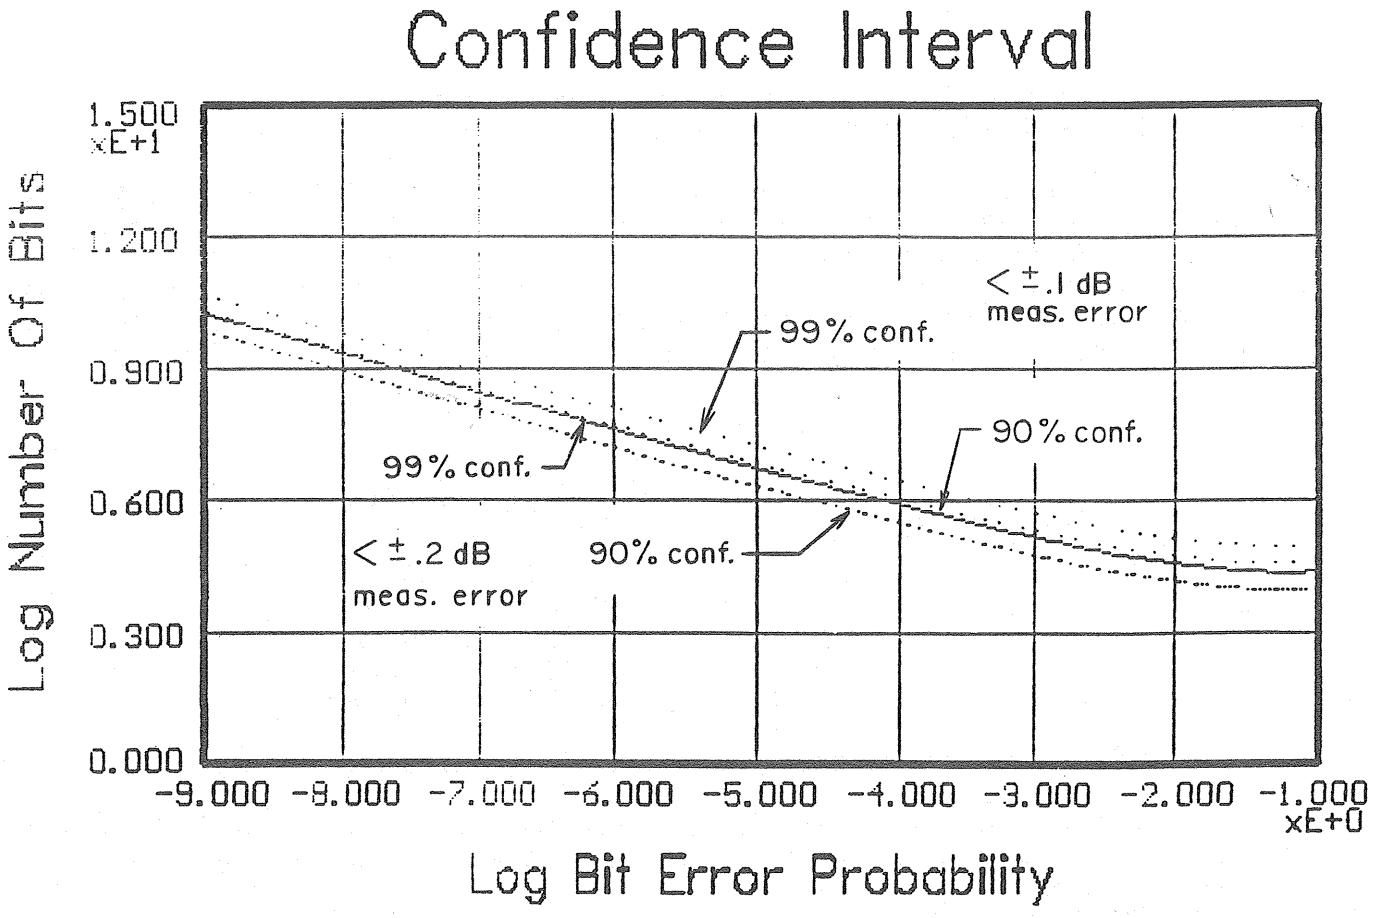

Figure 3. Measurement Time Versus Bit Error Rate

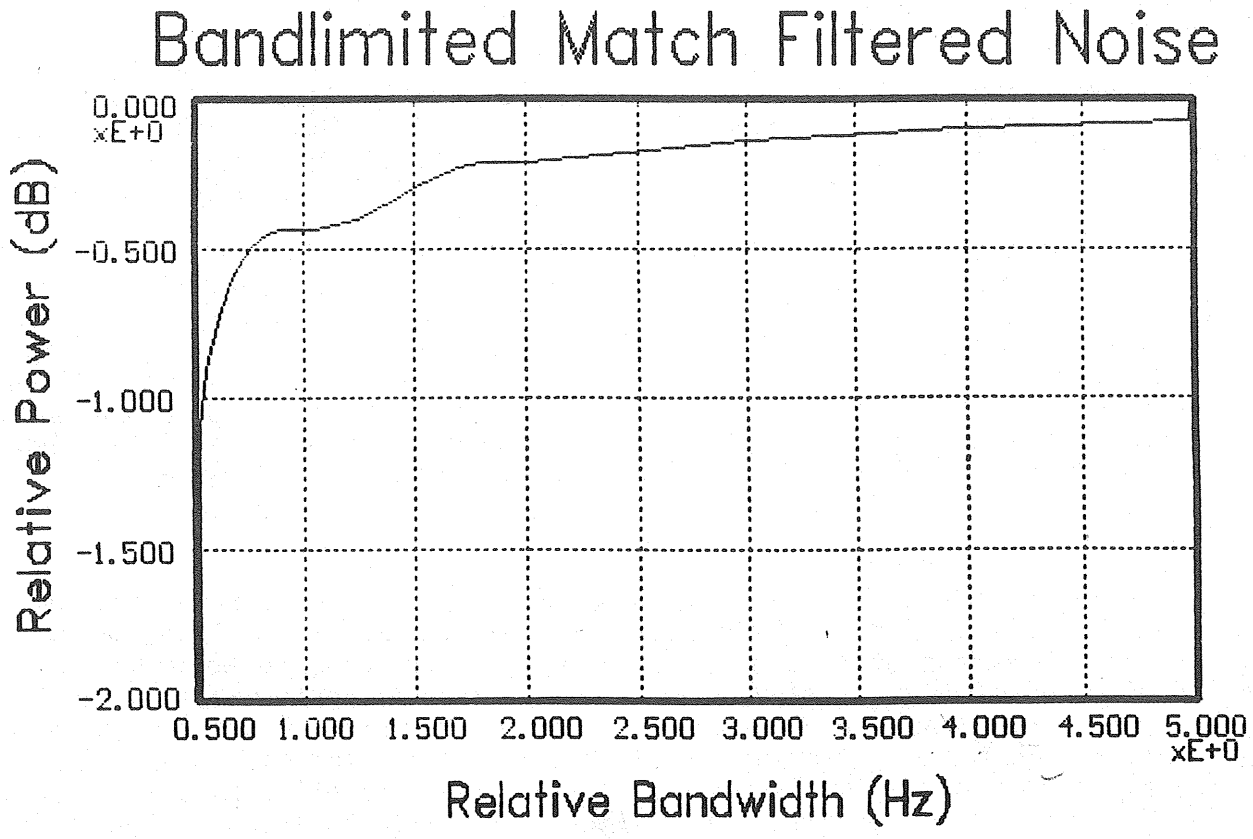

Figure 4. Effect of Bandlimited Noise

Systems and Subsystems specify the use of a noise bandwidth of greater than 5 times the bit rate. From figure 4 this results in a measurement uncertainty of less than .1 dB.

The effect of non-white Gaussian noise on bit error rate performance is summarized in table 1. The cases of rising, falling and sinusoidal power spectrums, as illustrated in figure 5, are treated. Note that these variations are virtually indiscernible when viewed on a spectrum analyzer using the 10 dB per division scale. In all cases Eb/No has been preserved. From table 1 it can be seen that perturbations of ±.5 dB in the noise spectrum result in performance variations of approximately ±.2 dB.

Table 1. Effect of Non-White Noise

	DEGRADATION IN PERFORMANCE		
I or A (dB)	RISING	FALLING	RIPPLE
.1	-0.0414	0.0420	0.0442
.2	-0.0823	0.0846	0.0871
.3	-0.1227	0.1277	0.1285
.4	-0.1624	0.1714	0.1686
.5	-0.2017	0.2157	0.207

NOTE: Ripple has the opposite effect if the first peak is negative.

Not only is it important that the noise spectrum be white, but the noise statistics must also be Gaussian. Noise statistics are most commonly examined in a qualitative manor using an oscilloscope. The noise is displayed on the scope to verify that it looks "like noise" and that there is no evidence of clipping. This is a reasonable check for adequate noise crest factor (peak-to-rms ratio) but is inconclusive as to the actual noise statistics. A more accurate method for determining noise statistics is to sample the noise amplitude (with an aperture compatible with the noise bandwidth) and generate a histogram for comparison with the Gaussian statistics.

The effect of crest factor on bit sync performance is shown in figure 6. Typical crest factors for bit sync testing range from 5 to 10 (14 to 20 dB). From the figure it can be seen that no errors are made after the Eb/No exceeds the crest factor by 3 dB. For example, a crest factor of 15 is marginally adequate for bit error rates down to 10^{-8}.

The effect of non-Gaussian noise statistics is illustrated in figures 7 and 8. Non-Gaussian noise statistics were obtained by passing Gaussian noise through slight non-linearities; a compression for figure 7, and an expansion for figure 8. The resultant non-Gaussian probability distributions are shown in figures 7a and 8a respectively. The noise powers have been normalized to preserve Eb/No. On an oscilloscope, the three noise sources would appear equally "like noise". However, the effect on bit sync performance is considerably different as shown in figures 7b and 8b. Compressed noise statistics improve bit sync performance at large S/N while expanded noise degrades performance.

FILTERED DATA

As previously stated, theoretical performance assumes square-sided (infinite bandwidth) data. In practice, all systems limit the data bandwidth to some extent. Systems with one sided bandwidths in excess of twice the bit rate are considered to be wideband. The performance of wideband data closely approximates theoretical performance; however, not all systems, particularly at high data rates, can support wideband operation. Also, many systems employ premodulation filtering to intentionally limit the data bandwidth. This is particularly common in PCM/FM systems where it is common practice to pre-filter with linear phase filters cut at .7 times the bit rate. To accommodate bandlimited signals, bit synchronizers provide a pre-filtered (also called filter and sample, F/S) data

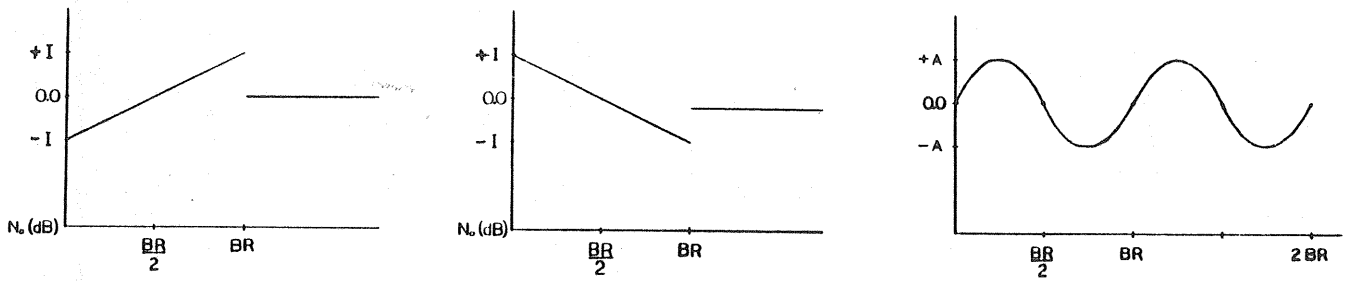

Figure 5. Non-White Noise

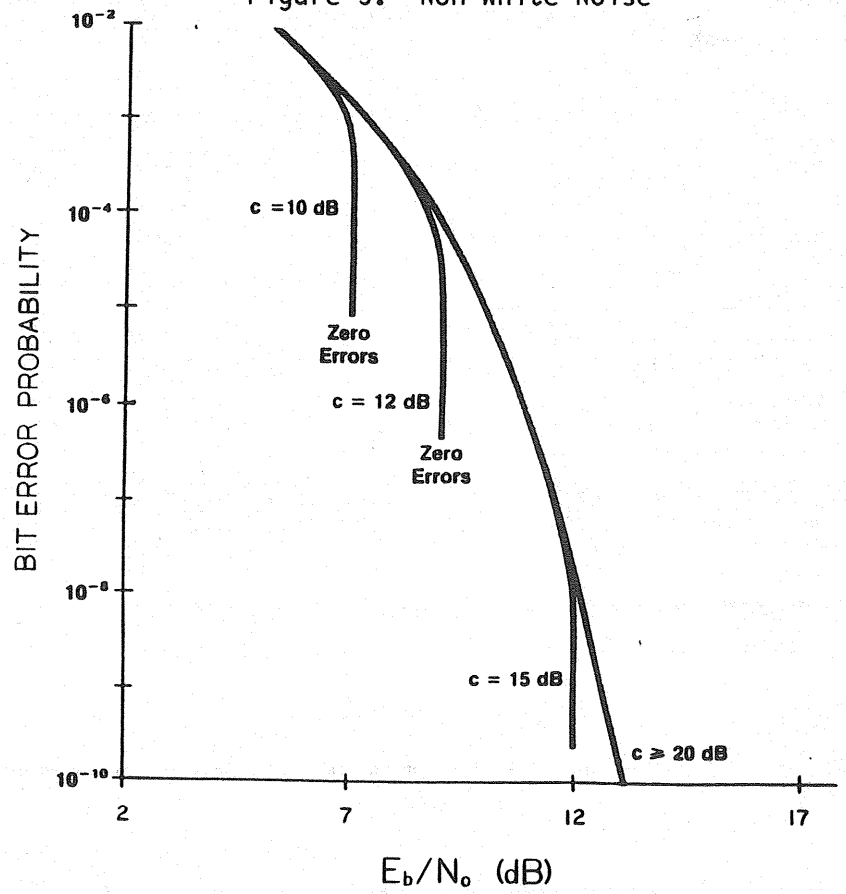

Figure 6. Effect of Crest Factor

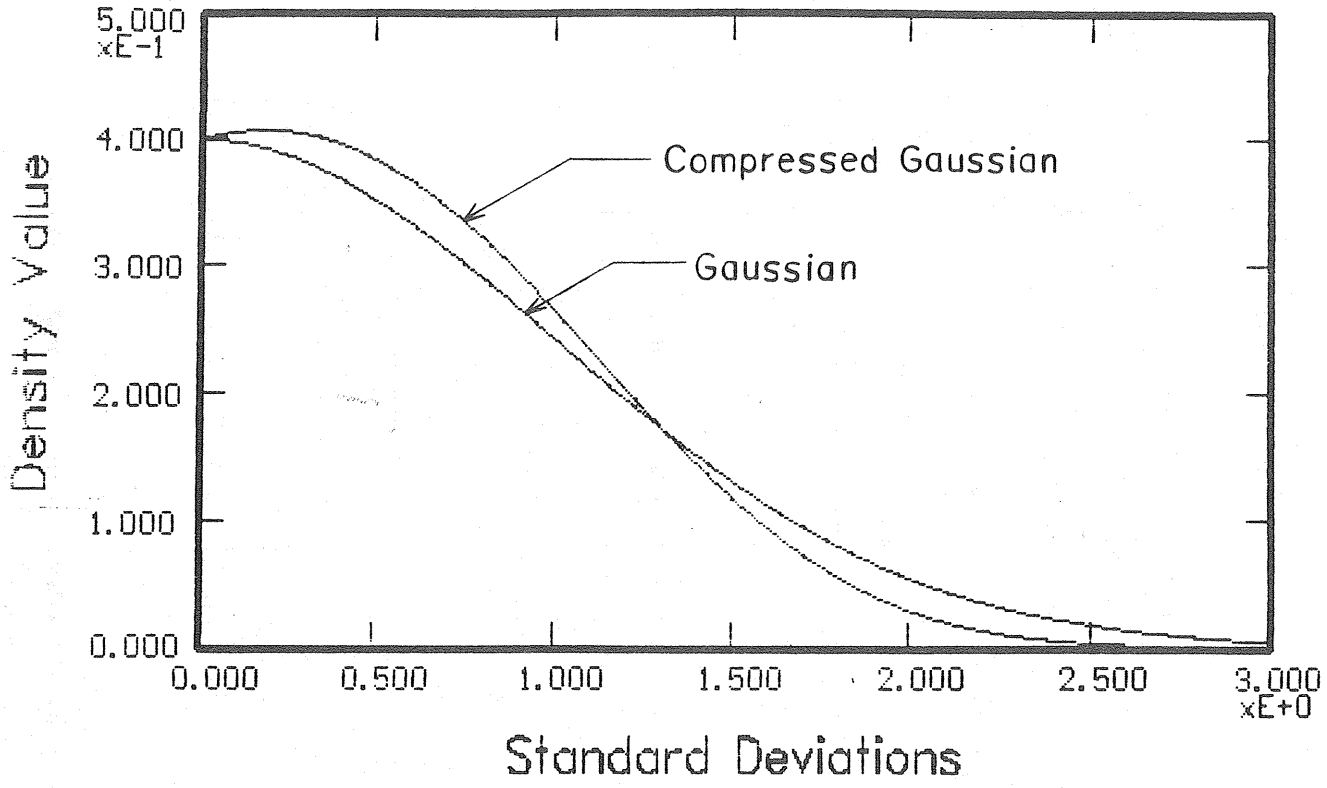

a) Noise Density

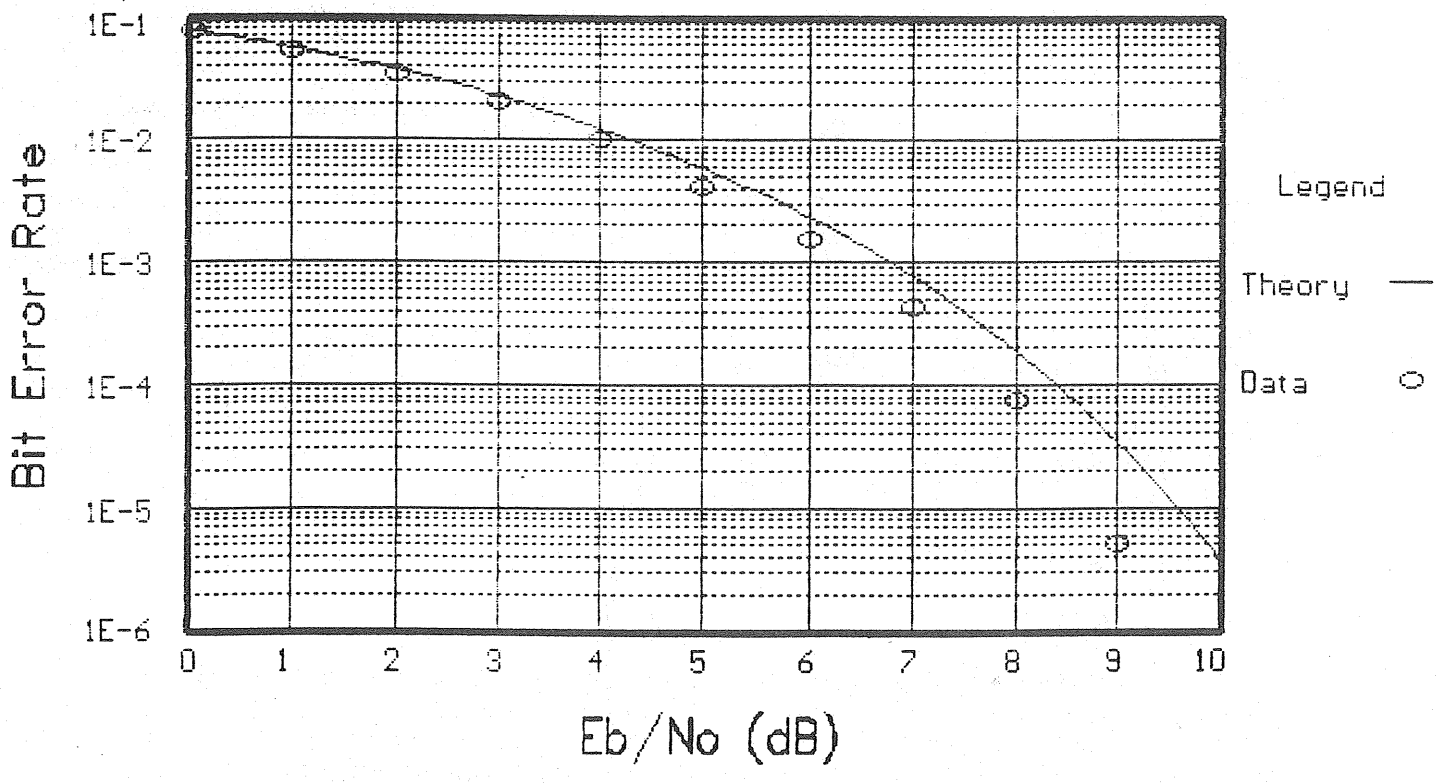

b) BER

Figure 7. Effect of Compressed Gaussian Noise

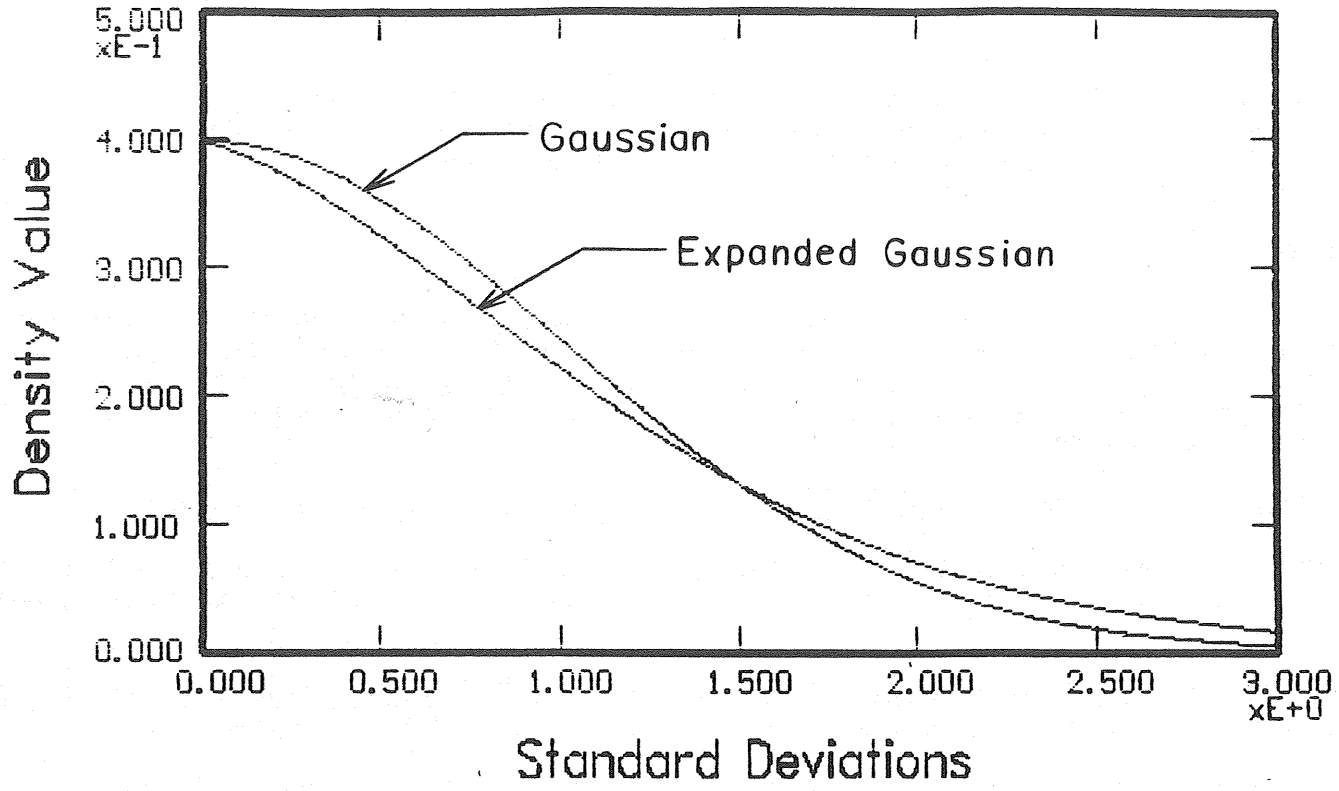

a) Noise Density

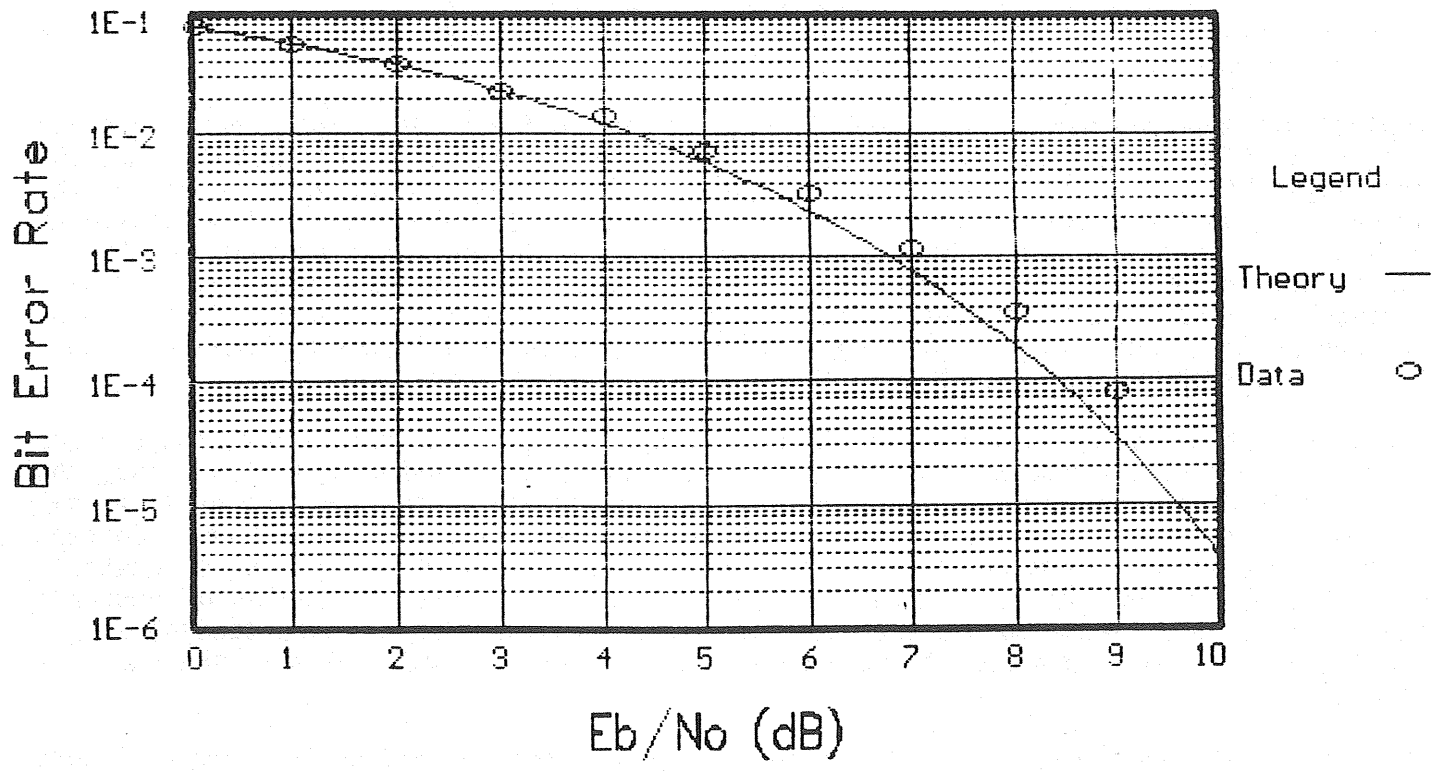

b) BER

Figure 8. Effect of Expanded Gaussian Noise

detection mode. Figure 9 illustrates the benefit of the F/S mode for the recovery of pre-filtered data. It can be seen that although the I/D (integrate and dump, or matched filter) mode provides the best performance for wideband signals, the F/S mode provides the best performance for pre-filtered signals.

ACQUISITION TIME

Acquisition time is a measure of the time required for a bit synchronizer to obtain clock synchronization. Bit sync acquisition time is conventionally specified as an average acquisition time, in terms of the number of bits required to achieve synchronization. A variety of definitions are used to identify when synchronization has been obtained. The two most common are:

1) When the recovered clock no longer cycle-slips relative to the source clock.

2) When the phase error between the recovered clock and the source clock remains less than an acceptable value, or equivalently, when the bit error rate decreases to an acceptable value.

Naturally, bit synchronizers specified using the first definition of synchronization will appear to have shorter acquisition times than those specified using the second definition.

The classic formula for average acquisition time (of a Phase Lock Loop, PLL) is given by [2]:

$$Tacq = \frac{(\Delta\omega)^2}{1.4\omega n} + \frac{1}{\omega n}$$

where $\Delta\omega$ is the initial frequency offset (in radians/sec) between the source clock and the bit sync clock, and ωn is the loop bandwidth (in radians/sec).

Typical bit sync average acquisition times are shown in figure 10. Note that for small frequency offsets, the acquisition time is proportional to the loop bandwidth and that the acquisition time increases rapidly for frequency offsets larger than the loop bandwidth. Average acquisition time increases with decreasing S/N but is relatively insensitive at moderate to high S/N. If bit sync acquisition time is specified without additional qualification, it should be assumed that it is an average acquisition time, measured with a wide loop bandwidth, a frequency offset within the loop bandwidth, and a moderate Eb/No (around 12 dB). Since it is an average acquisition time, acquisition sometimes takes more time and sometimes takes less time.

Acquisition time is actually a statistical quantity; it depends on the initial phase error (between the source clock and the bit sync clock) which is a random variable. Also, the acquisition of a phase-lock loop suffers from a "hangup" phenomenon. Hangup afflicts all phase-lock loops. A PLL has points of both stable and unstable equilibrium. Correct tracking occurs at points of stable equilibrium and hangup occurs at points of unstable equilibrium. An initial phase error can fall arbitrarily close to a point of unstable equilibrium. Near these points the restoring force which drives the loop towards stable equilibrium is low. As a result the acquisition time is extended. This is the basis of hangup.

The probability of acquisition versus time for various loop bandwidths is shown in figure 11. The effect of hangup is evident by the asymptotic behavior of the curve at high probability of acquisition. Note that although the time required to achieve a given probability of acquisition decreases as the loop bandwidth increases, the effect of hangup remains. The hangup phenomenon is clearly illustrated in figure 12, which shows the probability of having not acquired versus time. The effect of frequency offset on the

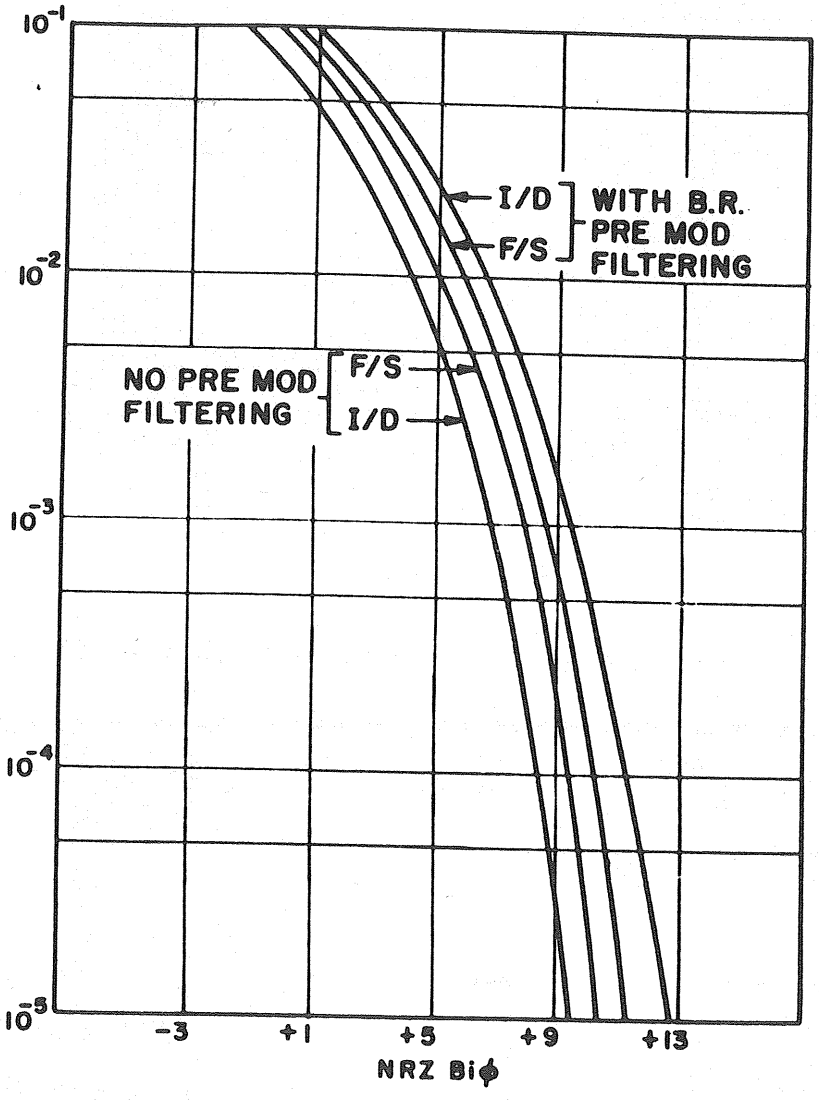

Figure 9. I/D Versus F/S Detection

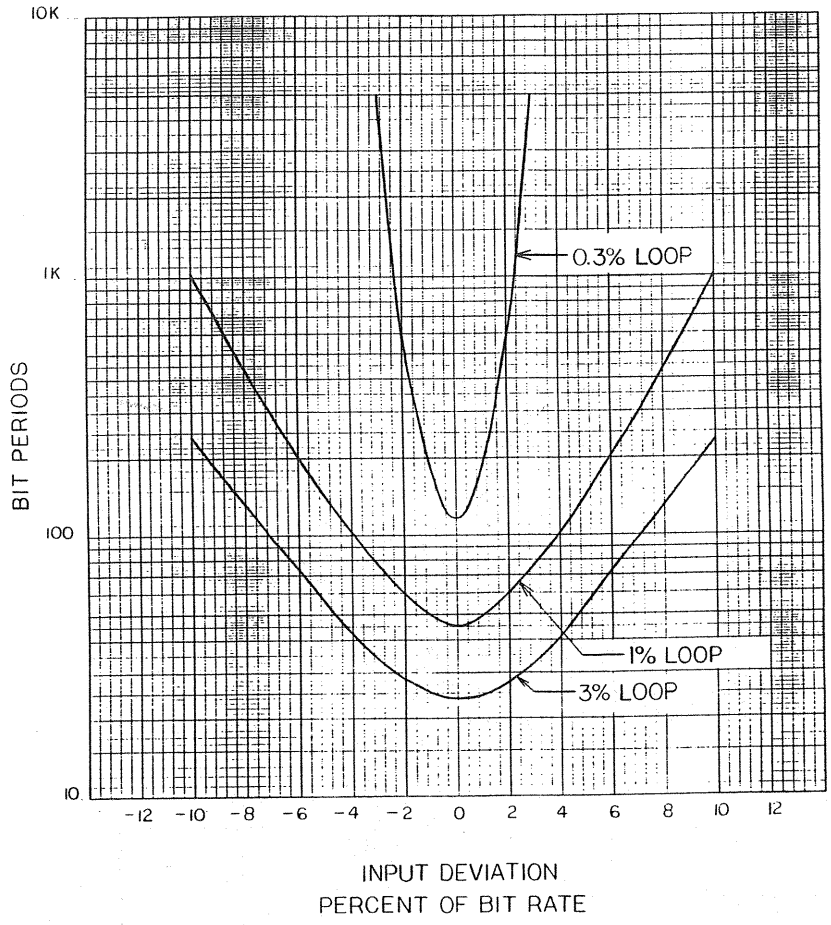

Figure 10. Average Acquisition Time Versus Frequency Offset

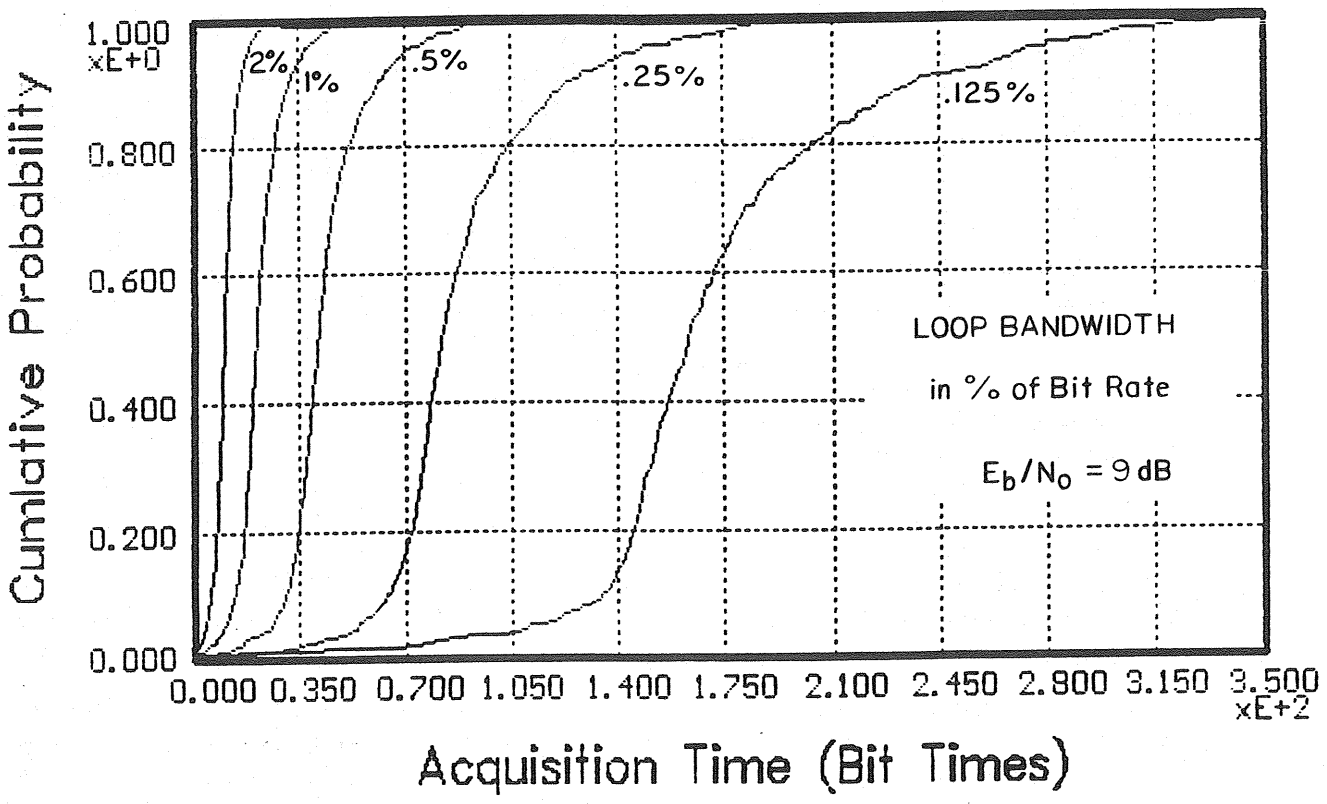

Figure 11. Probability of Acquisition Versus Time

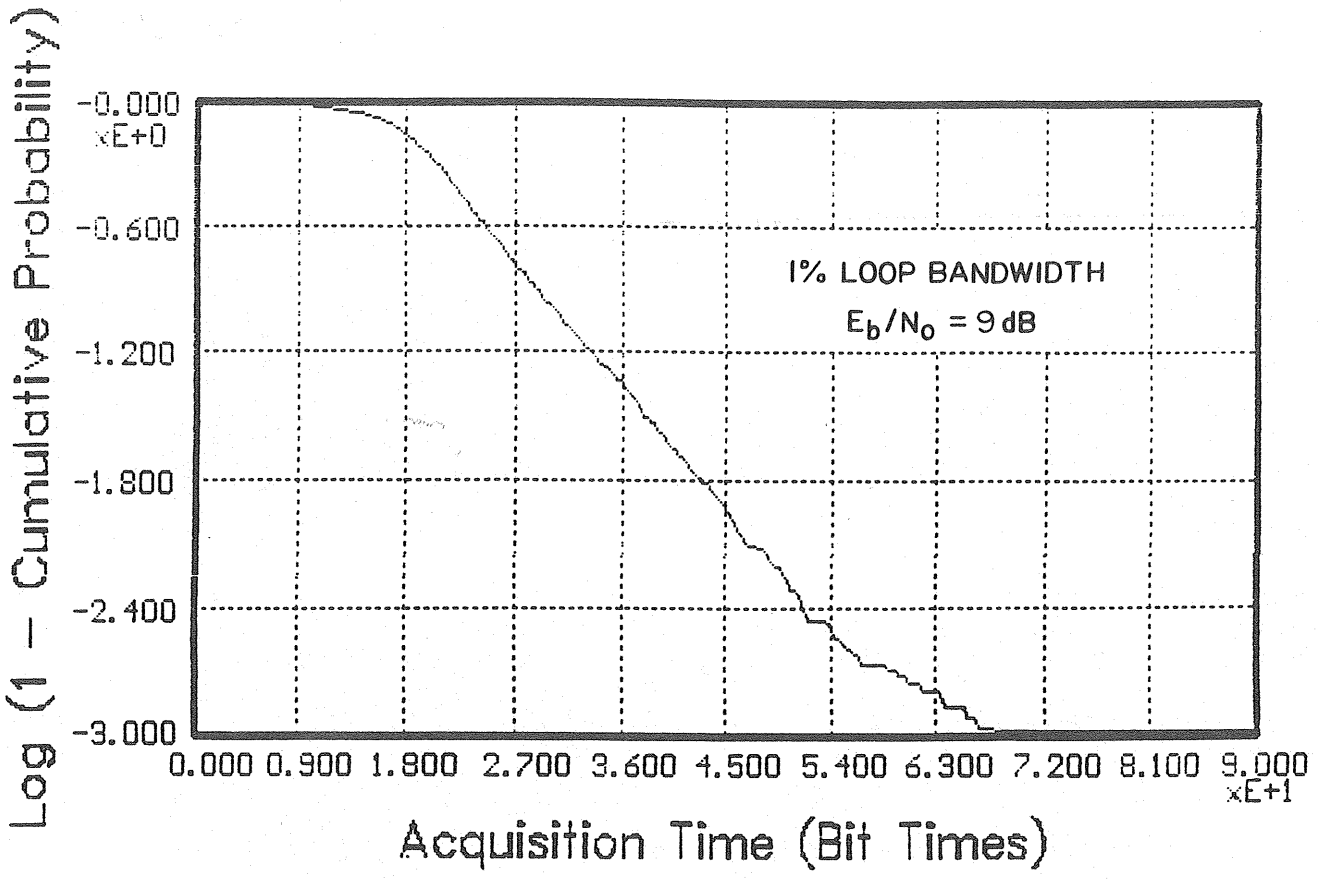

Figure 12. Probability of Not Acquiring Versus Time

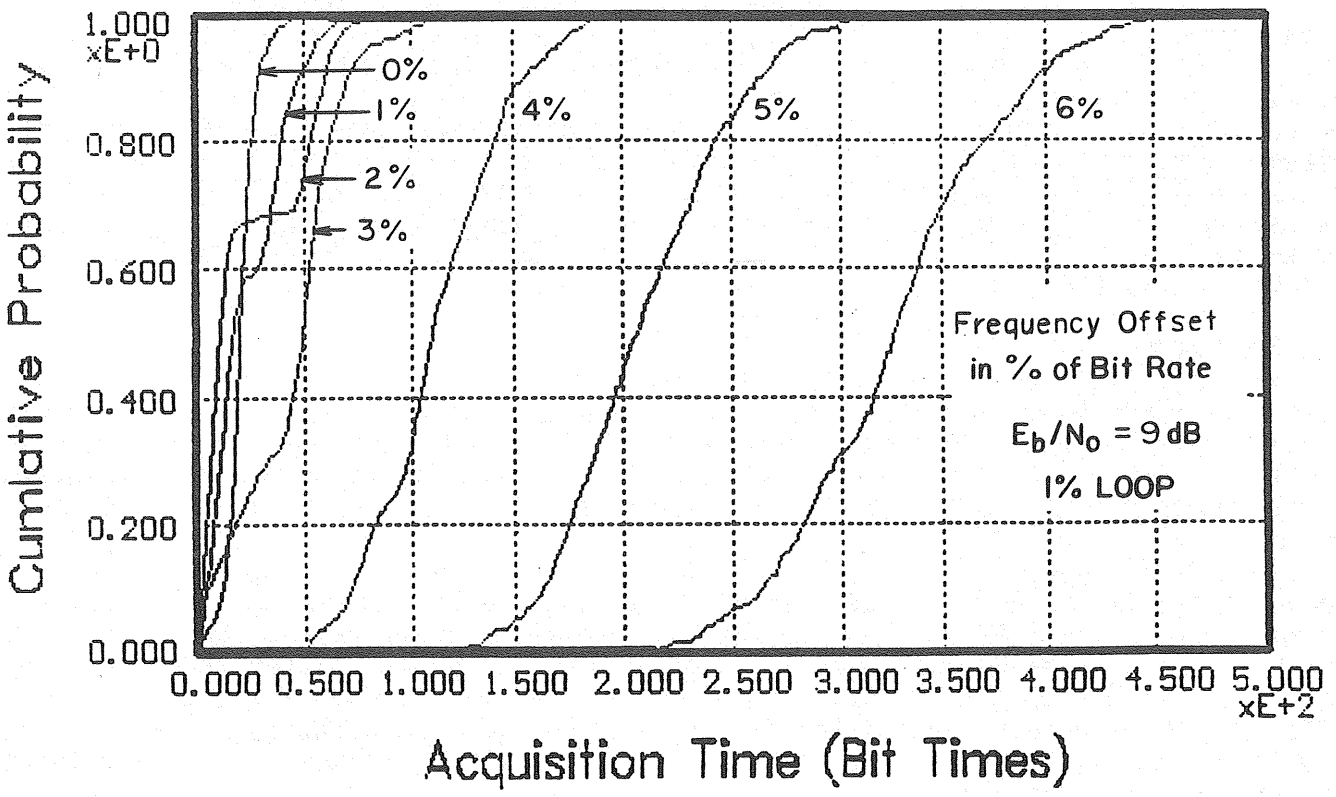

Figure 13. Probability of Acquisition Versus Time

probability of acquisition is shown in figure 13. The occurrence of cycle-slip (areas of nearly constant probability of acquisition versus time) during acquisition are clearly visible. The acquisition time increases as the frequency offset increases. The hangup phenomenon is always present.

All PLLs suffer from hangup and are therefore unsuited for applications where a high probability of rapid acquisition is required. Aydin has developed a "Smart Loop" which overcomes the hangup phenomenon. The smart loop has been employed in the Model 2710 tunable bit synchronizer. The acquisition performance of the smart loop is shown in figure 14. The first figure illustrates the dramatic improvement in acquisition provided by the smart loop by comparing the bit error rate of a smart loop to that of a conventional loop during acquisition. The second figure illustrates smart loop probability of acquisition.

SYNC MAINTENANCE

Sync maintenance addresses the ability of a bit synchronizer to remain in synchronization. Whereas rapid acquisition requires the short time constant of a wide loop bandwidth, good sync maintenance requires the long time constant of a narrow loop bandwidth.

A bit synchronizer maintains sync when the recovered clock does not cycle-slip relative to the source clock. The classic formula for the average time to cycle-slip is given by [2].

$$Tav = \frac{e^{(\pi/SNRL)}}{.5fn}, \quad SNR_L = f(Eb/No, Td, 1/fn)$$

where fn is the loop bandwidth in hertz and SNR_L is the S/N ratio in the loop bandwidth (which is a function of Eb/No, transition density (Td) and 1/fn). Tav rapidly decreases as Eb/No decreases, transition density decreases, or loop bandwidth increases.

Sync maintenance performance is typically specified as a synchronization threshold; the minimum Eb/No for which Tav is greater than 10^6 bits. Typical bit synchronizer sync maintenance performance is shown in figure 15. It is seen that, as predicted by the formula, the sync threshold increases as transition density decreases, or loop bandwidth increases. To maintain a constant sync threshold, as the transition density decreases the loop bandwidth must increase. The effect of transition gaps is similar to that of transition density; however, gaps are somewhat more disruptive since, during a gap, there is no phase-lock loop control signal. This allows the effects of circuit offsets and oscillator instabilities to drive the loop out of sync.

JITTER AND NOISE

The data to be recovered by a bit synchronizer frequently contains jitter. Jitter arises in various ways including source clock instabilities, bit stuffing and deleting in statistical multiplexers, tape recorder wow and flutter, and Doppler effect. To maintain sync and minimize bit error rate the rms phase error of the recovered clock must be minimized.

The phase error due to jitter is reduced by increasing the loop bandwidth. However, the phase error due to noise is increased as the loop bandwidth is increased. As a result the following three statements apply: 1) for a given loop bandwidth, jitter frequency, and Eb/No there is a maximum jitter amplitude which can be tolerated; 2) for a given jitter amplitude, loop bandwidth, and jitter frequency there is a maximum Eb/No which can be tolerated; 3) for any combination of jitter amplitude, jitter frequency, and Eb/No there is an optimum loop bandwidth.

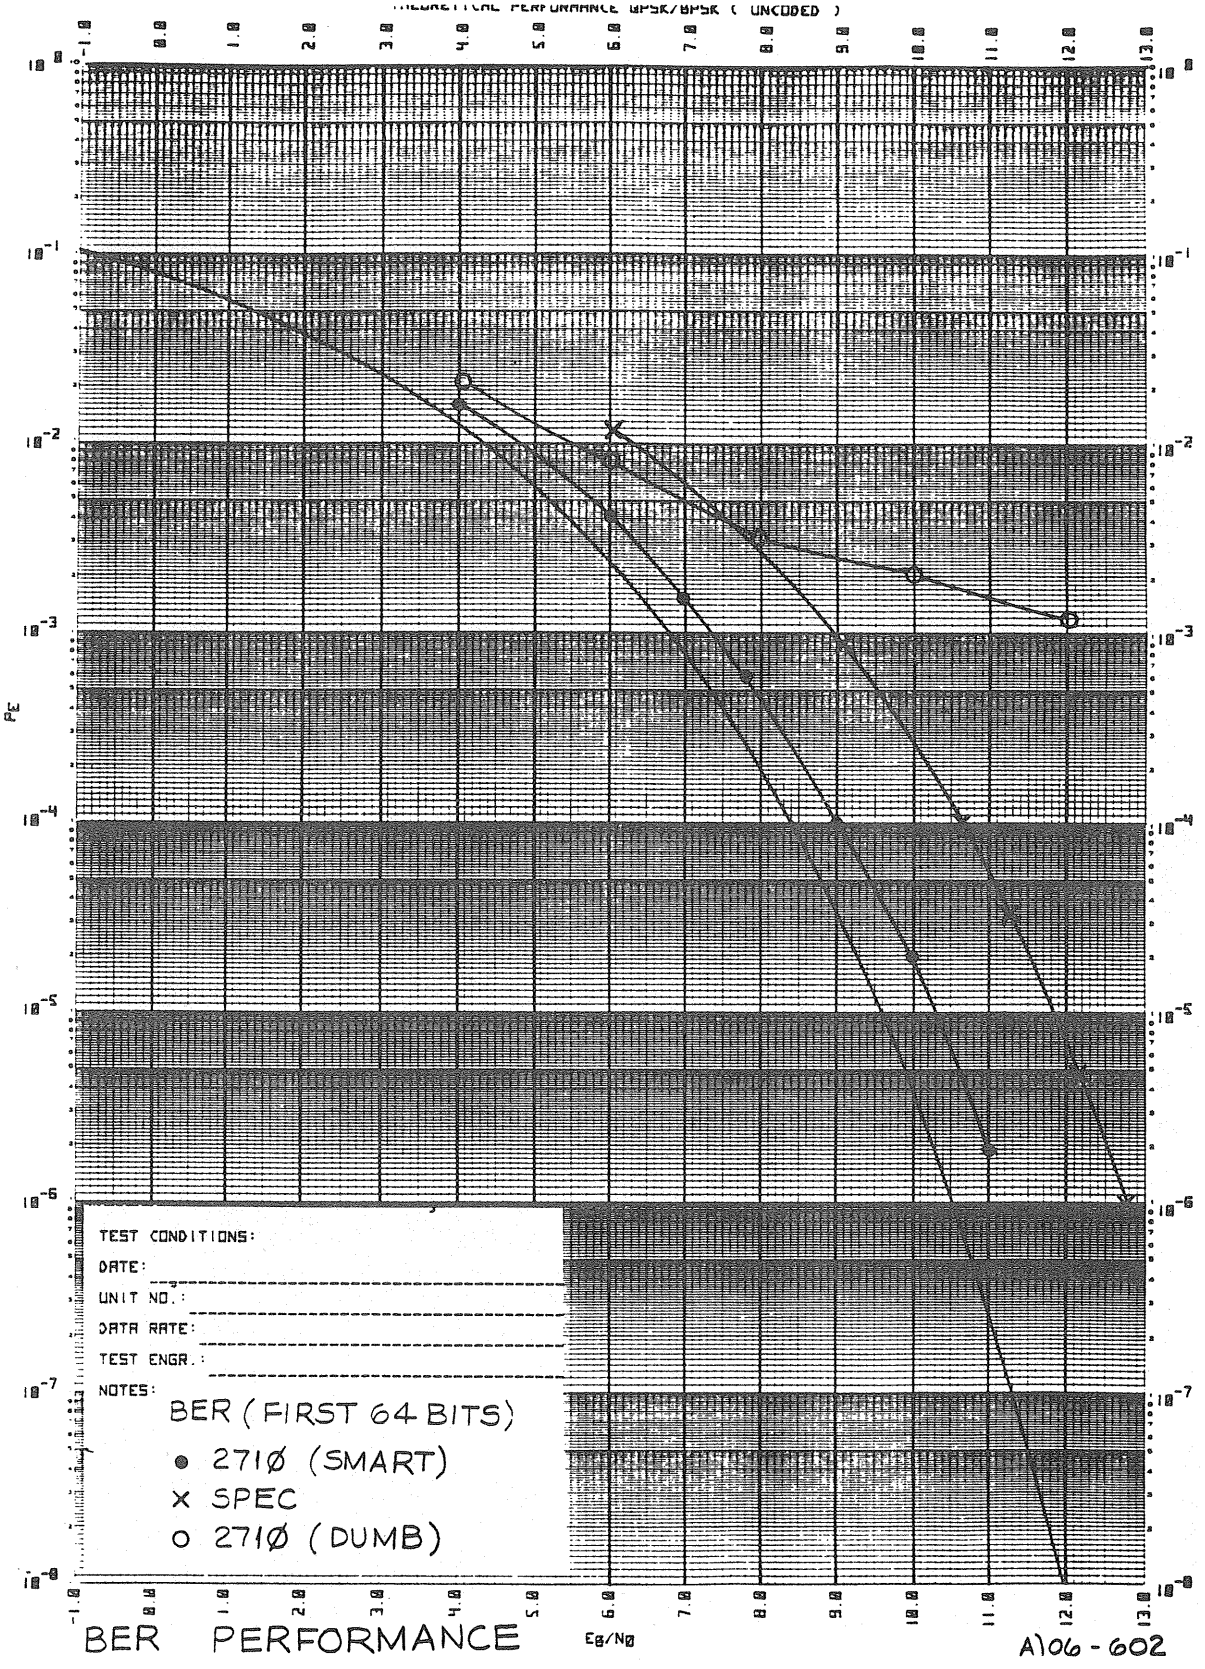

Figure 14a. Acquisition Performance of "Smart Loop" Versus Conventional Loop

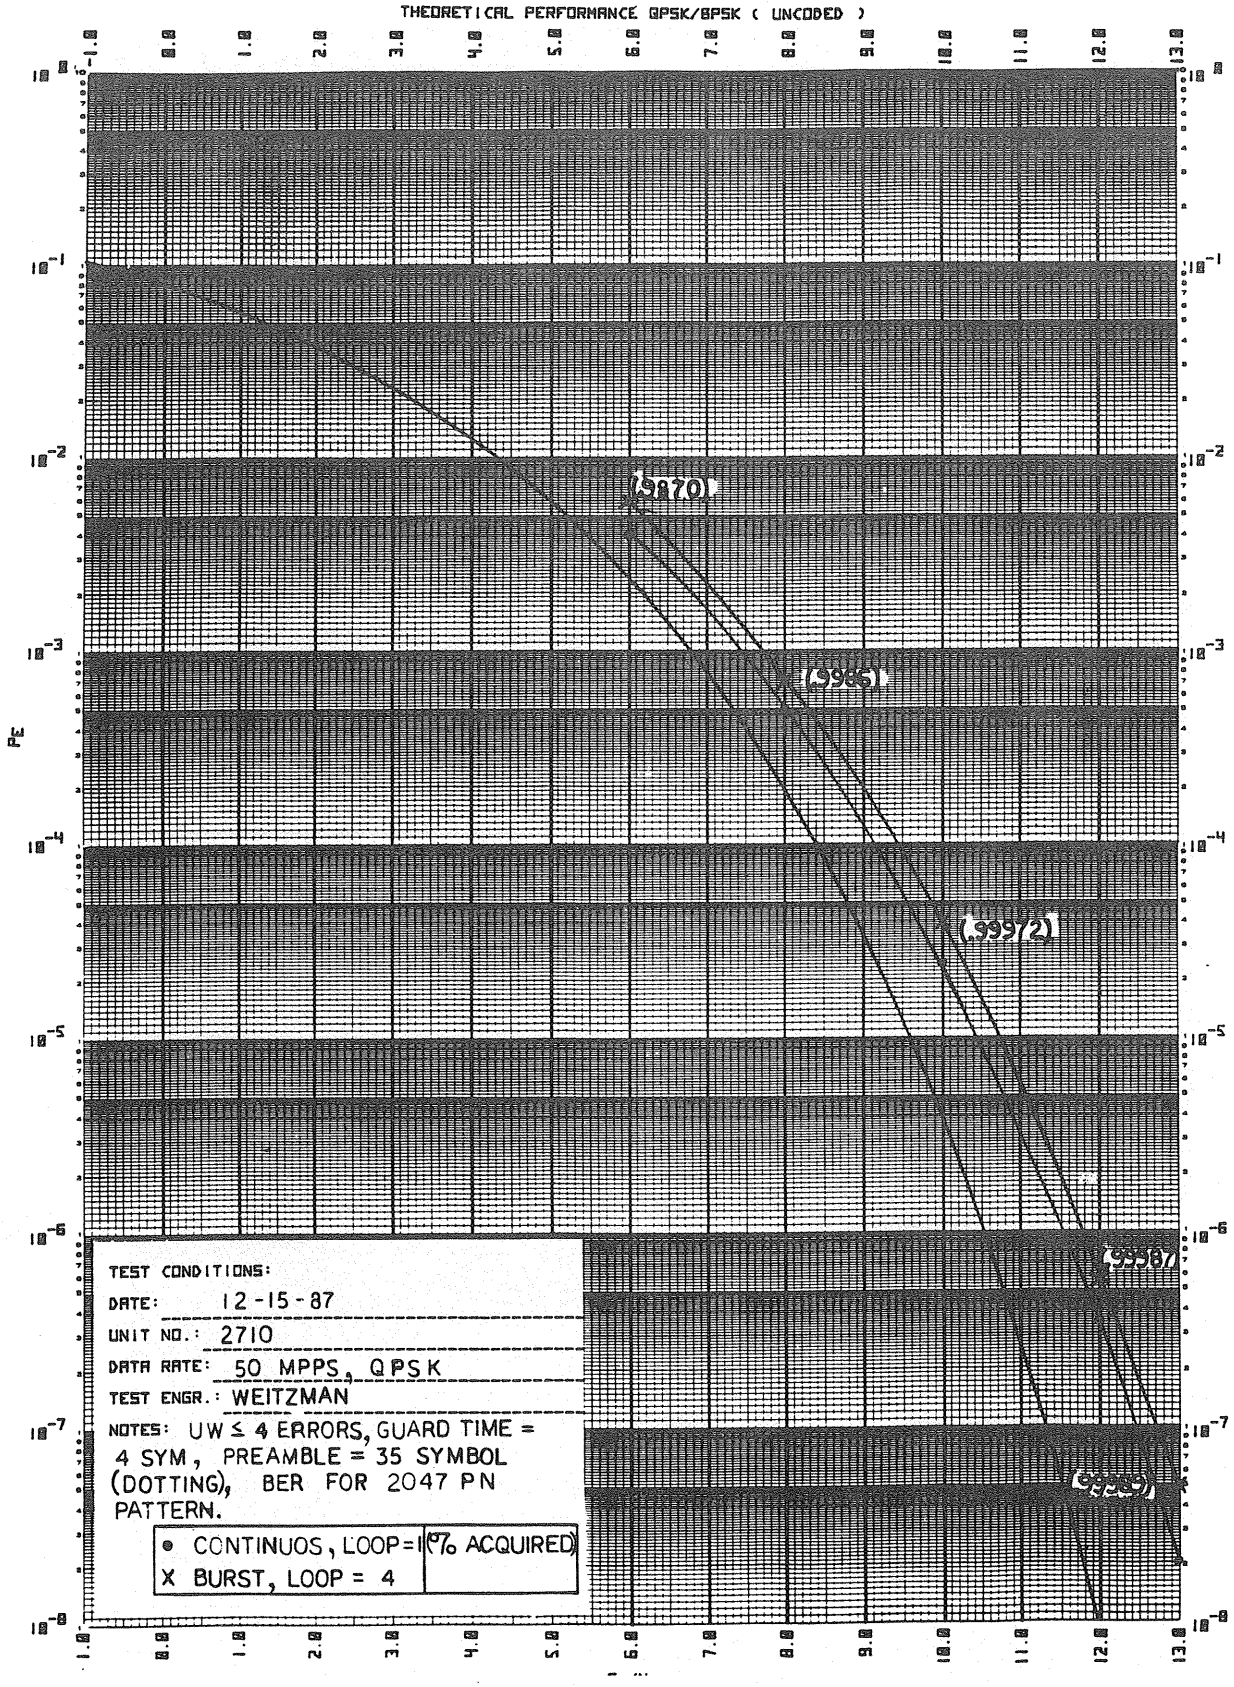

Figure 14b. Acquisition Probability of "Smart Loop"

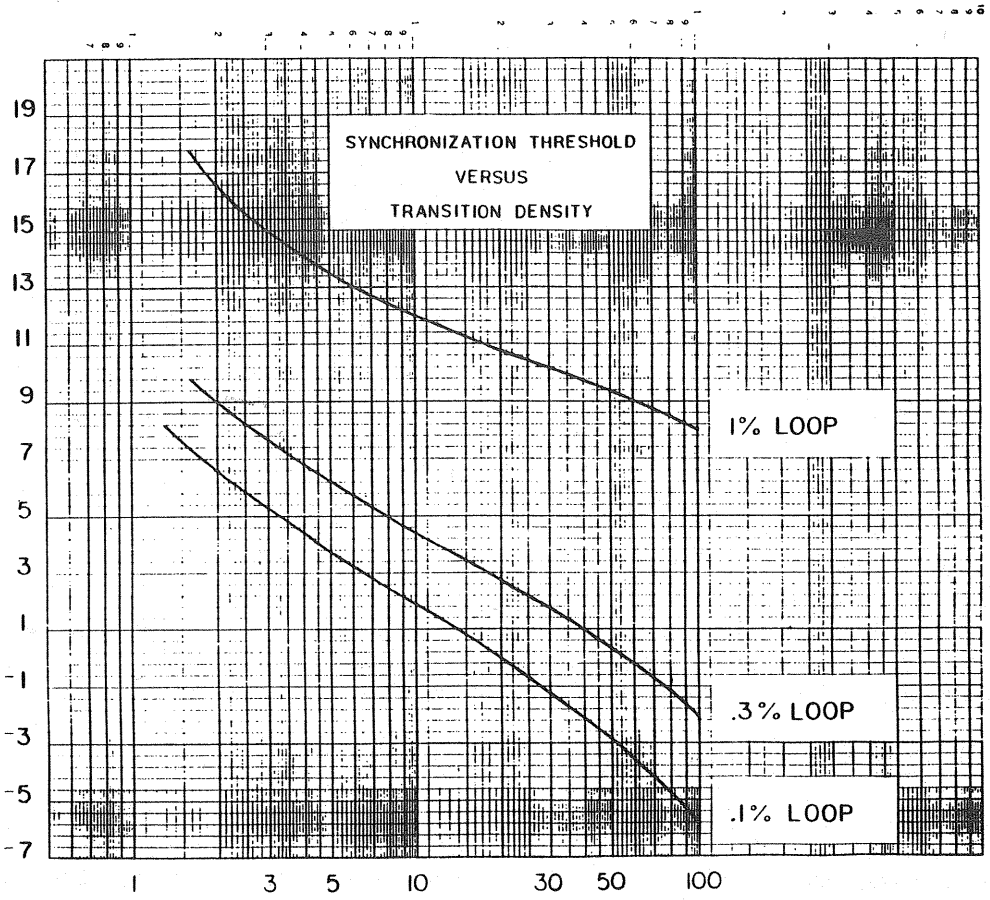

Figure 15. Sync Maintenance Performance

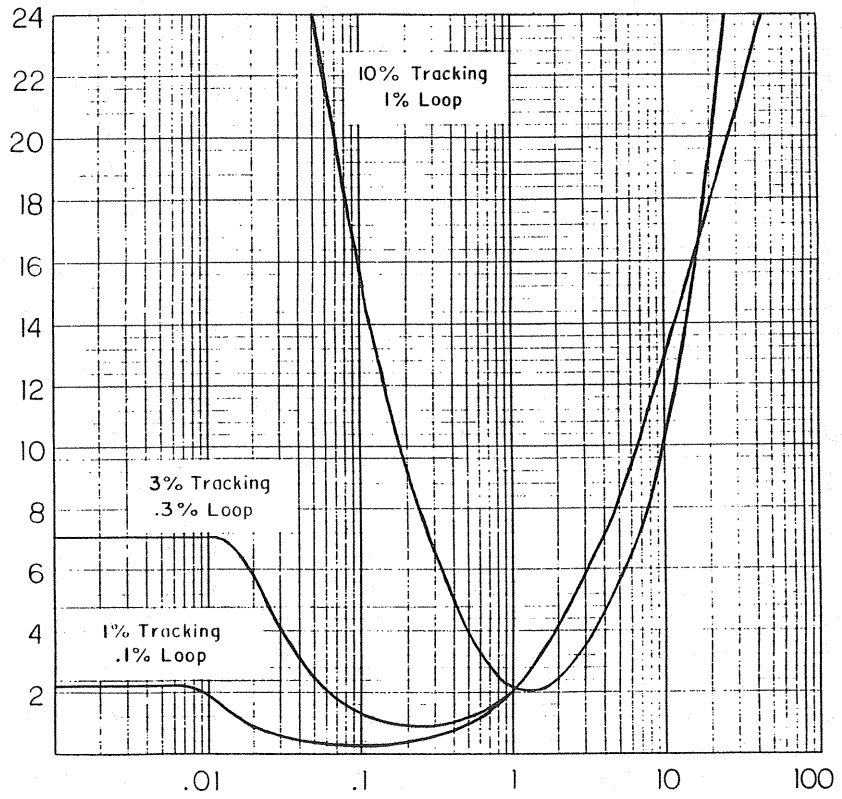

Figure 16. Peak Jitter Amplitude Versus Jitter Rate

Figure 16 presents typical bit sync jitter tracking performance for sinusoidal frequency modulation of the data source clock. The figure shows the peak jitter amplitude versus jitter frequency for which the bit sync can maintain sync at Eb/No = 12 dB. Sync is maintained for the region under the curve. Note that jitter tolerance is lowest when the jitter frequency is approximately equal to the loop bandwidth. As the jitter frequency is reduced, larger jitter amplitudes can be tolerated. As the jitter frequency is increased, larger jitter amplitudes can also be tolerated (this occurs because phase is the integral of frequency; therefore, the phase jitter produced by a sinusoidal frequency jitter is inversely proportional to the jitter frequency). The specified frequency tracking limits are illustrated by the regions of constant jitter amplitude tolerance at low jitter frequencies.

Bit error rate performance versus loop bandwidth for combined (PM) jitter and noise is shown in figure 17 with Eb/No as a parameter. As expected, there is an optimum loop bandwidth. The optimum bandwidth is slightly lower than the jitter frequency and increases as Eb/No increases. Note that the BER degradation is relatively insensitive to loop bandwidths wider than the optimum. Figure 18 shows bit error rate performance versus loop bandwidth for combined (PM) jitter and noise with jitter frequency as a parameter. Note that the optimum loop bandwidth increases as the jitter frequency increases and that, again, the optimum performance is relatively insensitive to a wider loop bandwidth.

LOOP BANDWIDTH SELECTION

The selection of an appropriate loop bandwidth is the key to obtaining the desired bit synchronizer performance. When considering acquisition only (no noise), the widest is best. When considering noise only (no jitter), the narrowest is best. When considering jitter only (no noise), the widest is best. When considering acquisition with noise, an intermediate bandwidth is best. When considering jitter and noise, an intermediate bandwidth is best. The best loop bandwidth also varies as the data transition density varies. With all the possible combinations it would appear that a bit synchronizer would need a large number of loop bandwidths. However, the typical uncertainties and variations in Eb/No, jitter frequency, transition density and component tolerance, obviate this need. Four properly selected loop bandwidths are adequate for most applications, a wide one for rapid acquisition, a narrow one for heavy noise and two in between for combined jitter and noise.

A new feature incorporated in the Aydin Model 3335 bit synchronizer is an adaptive loop bandwidth. The adaptive loop addresses the conflict between acquisition and noise performance. In the adaptive mode, a wide loop bandwidth is used for acquisition and when sync is obtained the unit automatically transitions to a (user-selected) narrower loop bandwidth for tracking.

EMBEDDED VITERBI DECODERS

For many years the bit error rate performance of telemetry and communication systems has been improved through the use of convolutional encoding and Viterbi decoding. In a typical system the data and clock recovered by the bit synchronizers were fed to the Viterbi decoder. Since the Viterbi decoder typically outputs only NRZ-L data and 0 degree clock, the decoder output was routed to a code converter. The code converter was then used to select the desired output codes and clocks.

The Aydin bit synchronizer product line now provides all three functions in a single 5 1/4 inch chassis. This results in a significant reduction in rack space, cabling and system integration time. A test convolutional encoder has also been incorporated to allow testing of the decoder without the modification of existing bit sync test equipment.

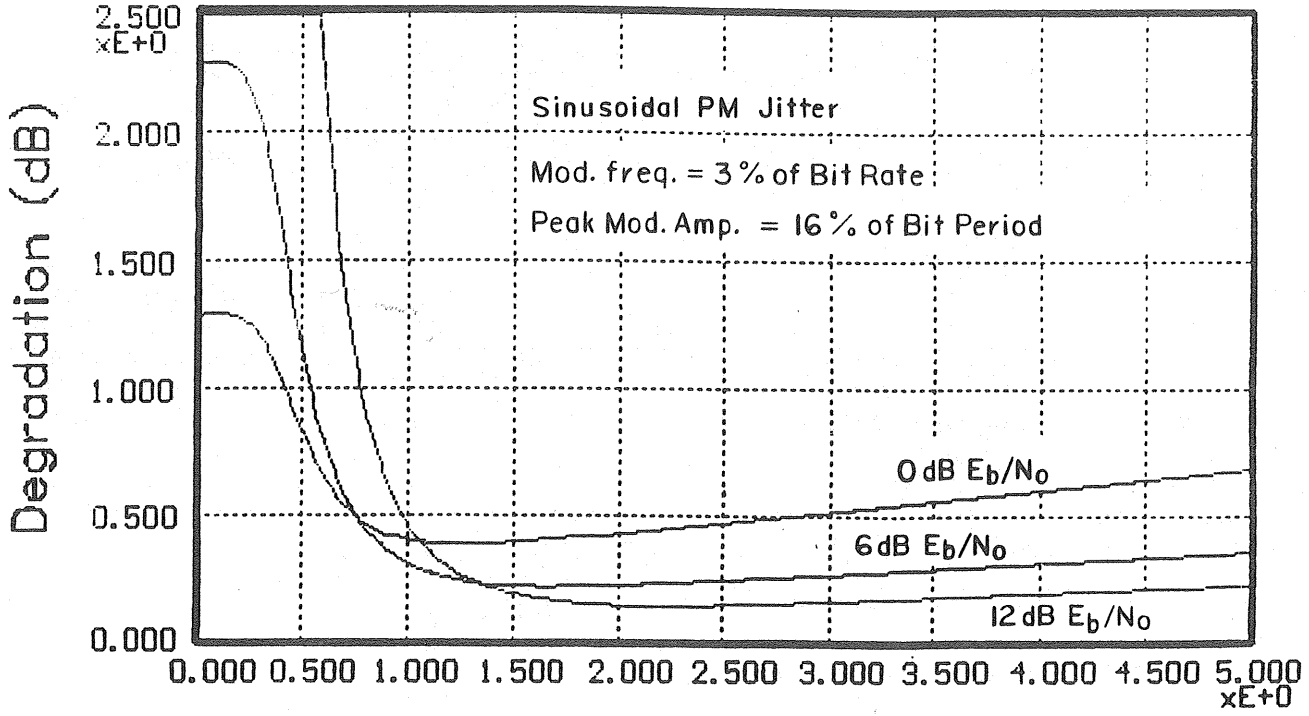

Figure 17. BER Degradation Versus Loop Bandwidth for Combined Jitter and Noise

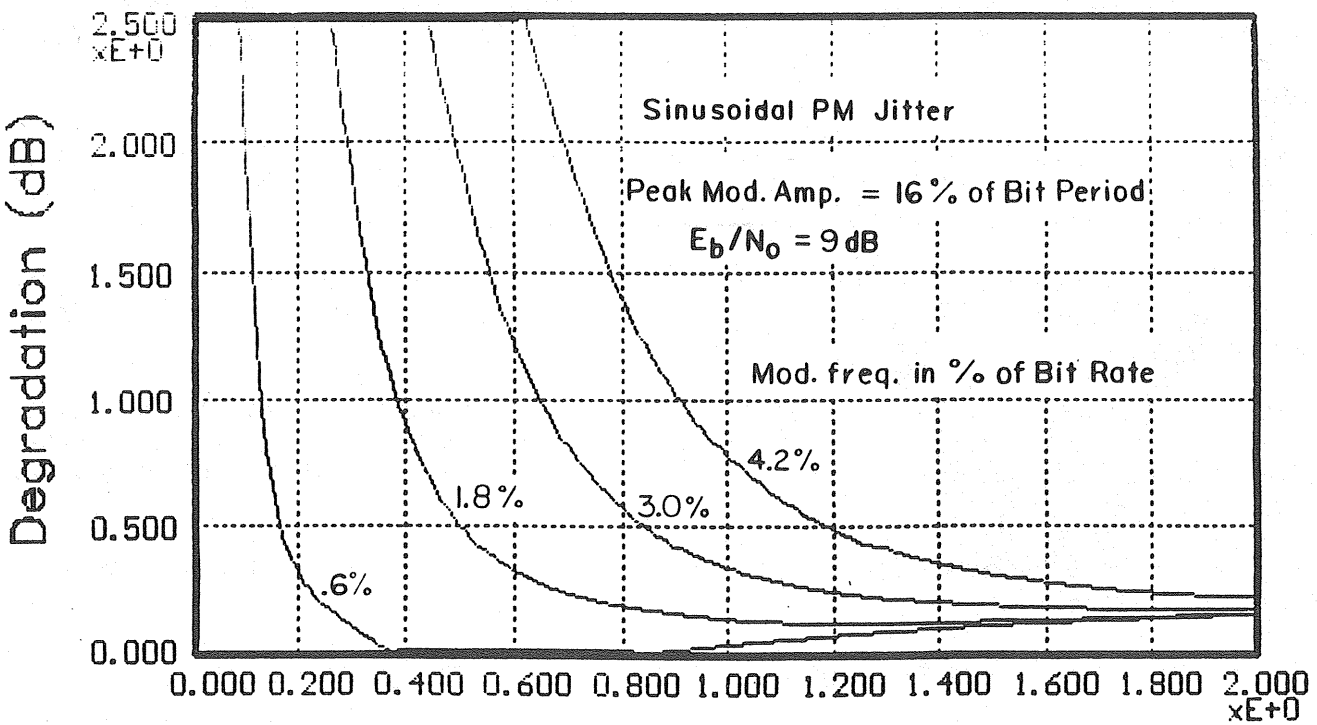

Figure 18. BER Degradation Versus Loop Bandwidth for Combined Jitter and Noise

To maximize performance soft bit decisions are provided to the decoder by the bit synchronizer. Typically, 3-bit soft decision is employed. The soft decisions are obtained by quantizing the output of the data matched filter as shown in figure 19. The automatic gain control of the bit sync maintains the signal at the indicated level. The most significant bit, called the sign bit (and also called the hard decision bit) is the usual data output of the bit sync. The two least significant bits provide additional data quality information to the decoder.

The theoretical performance of constraint length 7, rate 1/2 Viterbi decoding is shown in figure 20 for both hard decision and soft decision. It is seen that a coding gain of over 5.4 dB is available at a bit error rate of 10^{-5} using soft decision. It should be noted that decoder performance is defined in terms of the information rate; the data rate at the input of the encoder and at the output of the decoder. That is, Eb/No as used in the Viterbi decoder performance curse, is the S/N ratio in a bandwidth equal to the data rate at the output of the decoder. The bit synchronizer, however, must recover data and clock at the symbol rate. For a Viterbi decoder of rate R, the symbol rate is higher than the data rate by 1/R. Bit sync performance must therefore be based on the S/N ratio measured in the symbol rate bandwidth, Es/No. This is related to Eb/No by:

$$Es/No = Eb/No + 10 \log R.$$

For a rate 1/2 Viterbi decoder:

$$Es/No = Eb/No - 3 \text{ dB}$$

Viterbi decoders are used to improve BER performance at low Eb/No, often as low as 2 dB. This implies an Es/No of -1 dB at the input of the bit synchronizer. It is therefore essential that the bit sync have a very low sync threshold. Maintenance of sync is extremely important since a cycle slip in the bit sync requires reacquisition of node sync in the Viterbi decoder, a process which requires approximately 150 bits to accomplish. It is also important that the bit sync AGC be designed for operation at very low S/N ratios so as to preserve the relationship of the soft bits and maximize the available coding gain.

SUMMARY

This paper has reviewed the basics of bit synchronizer testing. The noise source has been identified as the main contributor to measurement error. Aydin is currently developing an all-digital noise source for use in a high rate bit sync test set.

The basics of bit synchronizer performance have been reviewed. Conflicting requirements have been identified between rapid acquisition and sync maintenance, and between noise and jitter performance.

New bit sync features have been introduced including: the smart loop for rapid, reliable acquisition, the adaptive loop for accommodating simultaneous acquisition and tracking requirements, and the embedded Viterbi decoder for performance enhancement in low Eb/No applications.

The auther wishes to acknowledge the contribution of Mr. R. Peloso in bit sync modeling, analysis and simulation.

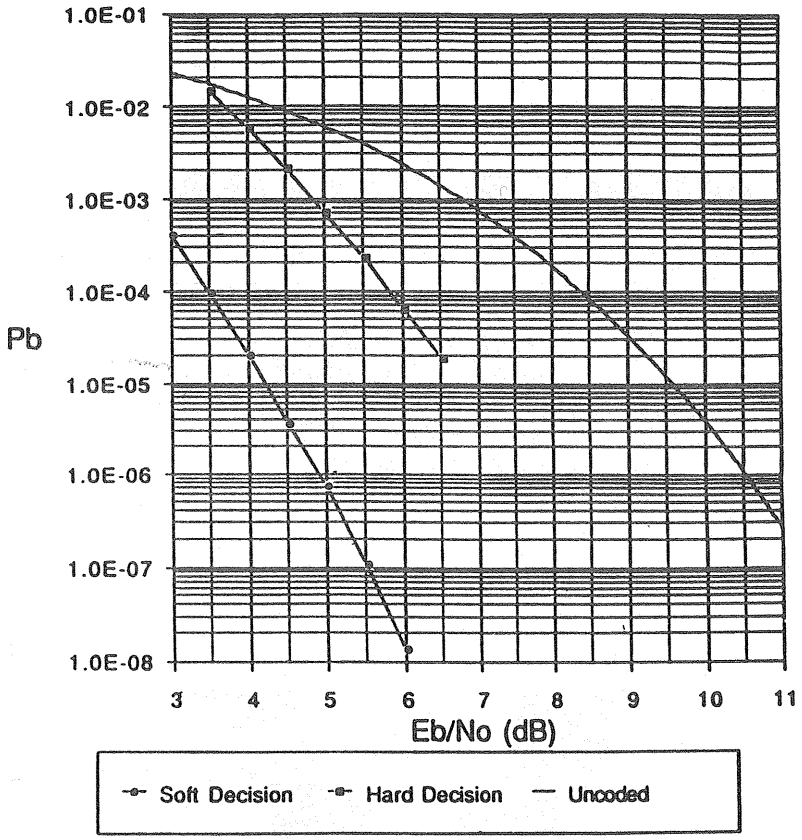

Figure 19. Soft Bit Decision Thresholds

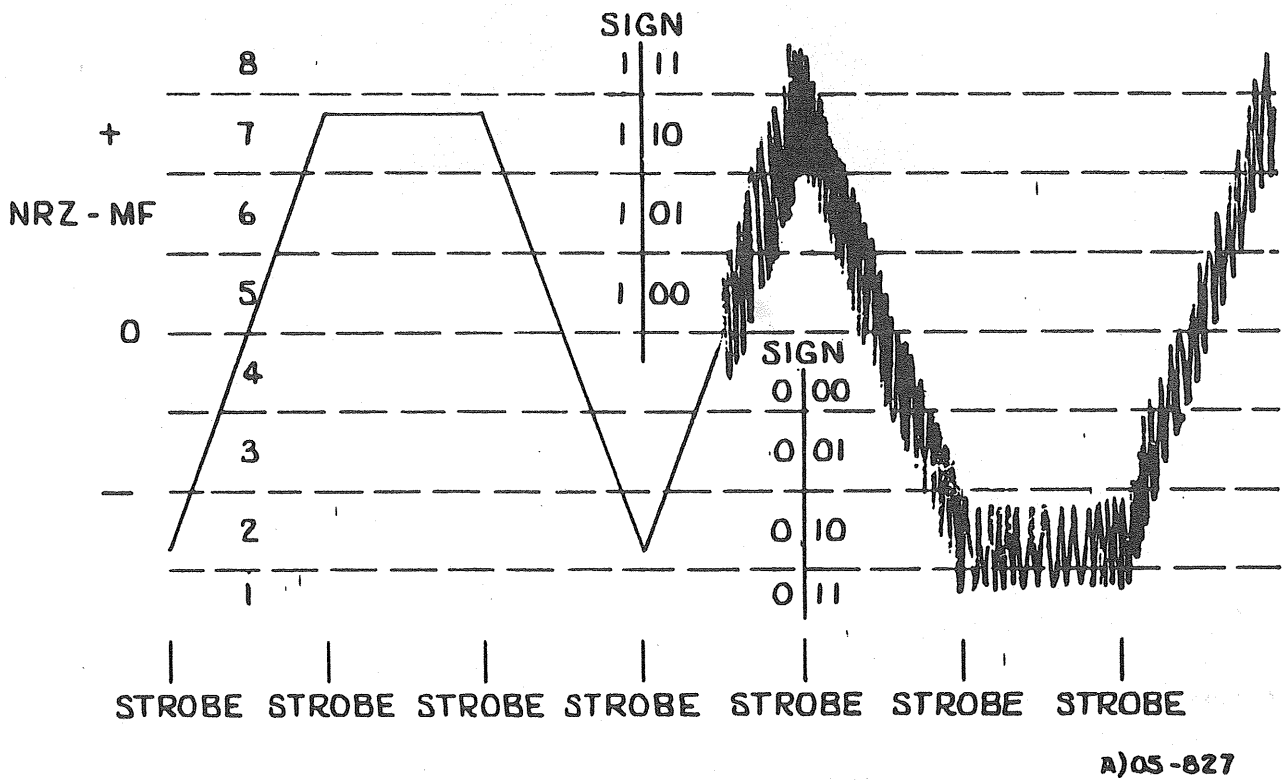

Figure 20. Coding Gain for Rate 1/2 Viterbi Decoder