Digital Demodulation With a Non-Ideal Quantizer

J. W. Layland

Communications Systems Research Section

The use of digital demodulation techniques, or the conversion from analog to digital signal domains before data are demodulated from a radio frequency (RF) carrier, has become increasingly more feasible in recent years as logic circuit speeds have increased. This feasibility has been demonstrated by use in the Dual-Channel (Mu-II) Sequential Ranging System. Further use is contemplated in the telemetry stream, or integrated into the phase-tracking receiver. This article reviews some of the properties which must be considered in the analog-to-digital (A-D) converter to be used in these applications. In particular, the loss to be expected with an A-D unit built from current circuitry is calculated.

I. Introduction

The use of digital demodulation techniques, or the conversion from analog to digital signal domains before data are demodulated from an RF carrier, has become increasingly more feasible in recent years as logic circuit speeds have increased. This feasibility has been demonstrated by use in the Dual-Channel (Mu-II) Sequential Ranging System (Refs. 1 and 2). Further use is contemplated in the telemetry stream, or integrated into the phase-tracking receiver. This article reviews some of the properties which must be considered in the A-D converter to be used in these applications. In particular, the loss to be expected with an A-D unit built from current circuitry is calculated.

II. Quantizer Design Considerations

Analog-to-digital converters for digital demodulation must be fast enough to respond to signals at a reasonable intermediate frequency (IF), say, 10 MHz, and precise enough that a tolerably small degradation is introduced. The speed requirement implies that conversion is performed "in parallel" instead of serially by successive approximation. This in turn imposes a constraint on the precision of the quantizer since a quantizer with b-bits of precision resolves a signal into 2^b separate levels and requires $2^b - 1$ threshold elements. Quantizers with 3 to 4 bits seem relatively uncomplex. Quantizers with 6 or more bits, over half a hundred threshold elements, seem complex enough to be impractical. Fortunately, as

calculations to be displayed later show, a 4-bit quantizer appears adequate for this application.

The signal to be quantized consists of Gaussian noise, white over the band of interest, together with a sinusoidal signal. This sinusoid in turn is phase modulated by a square-wave signal, either containing data, a data subcarrier, or a range code. The phase angle of modulation may be rather large, near $\pi/2$, for a data signal, or exceedingly small, less than 0.1 radian, for a ranging code. The sampling of this signal is performed in phase synchronism with this carrier frequency, and at a rate equal to four times that frequency. The phase of the sampling clock is controlled so that the (unmodulated) carrier component of the signal is sampled at its zero-crossings and peaks. Samples at the carrier zero-crossings correspond directly to the modulation which is in quadrature to that carrier. Alternatively, one could leave the phase of the samples and carrier random, and reconstruct the modulation by appropriate weighting of all samples, as was done in the Mu-II Ranging System. We assume that the bandwidth of the noise is wide enough that each sample of that additive noise is independent. This assumption considerably simplifies quantizer analysis.

III. Quantizer Behavior

An ideal quantizer is a memoryless nonlinear device which is subject to straightforward calculation of input-output relationships. The threshold elements which make up a real quantizer are not memoryless but exhibit hysteresis, i.e., the output will not change state from "above" to "below" until the input has dropped to some specific non-zero level $-\delta$ *below* the threshold, and will not change state from "below" to "above" until the input has climbed to $+\delta$ *above* the threshold. Furthermore, this means that the input could be initially above the threshold, and then drop to slightly below the nominal threshold, but not $-\delta$ below that threshold, and yet continually indicate "above." This is a complicating feature for the calculation of the quantizer input-output characteristic, since if the quantizer input is within $\pm\delta$ of a threshold, the quantizer output depends upon what that output was previously. Assuming that the noise is independent from sample-to-sample, and that the input signal has a short period (4 here) makes the problem tractable for computer solution. For a b-bit quantizer, we calculate $2^{2b} - 1$ probabilities for each of the four samples, which correspond to the probability that a sample falls solidly within one of the 2^b output levels, or in the hysteresis gap within $\pm\delta$ of the $2^b - 1$ thresholds. Whenever a sample is in a hysteresis gap, the output can be associated to the output level either above or below

that threshold, depending upon quantizer output for the previous sample. The probability distribution of the previous sample in turn depends upon the sample previous to it, and so on, until the cycle closes at the period of the input wave, and the calculation can be completed.

Subjectively, a quantizer with hysteresis, given an input of a small amplitude sine wave plus larger-power Gaussian noise, produces an output whose expected value is a phase-shifted sine wave. The amount of phase shift depends upon the width of the hysteresis gap, δ, and upon the relative phase of the sine wave and sampling clock, as is shown in Fig. 1 for total input power equal to 5 (quantizer-steps-squared), and input SNR (ρ) of 1 and 5 dB. Figure 1a for $\rho = 1$ seems representative of low SNR behavior where the noise effectively "dithers" the signal through several quantization steps. Figure 1b only hints at the erratic behavior which appears as the SNR before quantization increases to the point where the additive noise is insufficient to smooth the step between quantization levels.

IV. Calculated Performance

The loss in signal-to-noise ratio for demodulated data extracted from a sine wave carrier in noise has been calculated and is displayed below for a quantizer of the type used in the Mu-II Ranging System (Refs. 1 and 2). This quantizer is a 4-bit, 15-level device with outputs from -7 to $+7$. A zero input corresponds to mid-range on the zero-output level, and numerical output values "make sense" as an approximation to the input voltage. The stepsize between thresholds is 100 mV, and the hysteresis bias δ is specified by the manufacturer as 10 mV maximum for the threshold elements used. The quantizer is preceded by an automatic gain-control amplifier which maintains total signal-plus-noise power constant. Since signal-to-noise ratio is small in the unprocessed samples, this device seems to be a valid approximation to the optimum linear quantizer for decoding (see, e.g., Ref. 3). Figure 2 shows the detected signal SNR loss in dB as a function of sample SNR for various values of total input power for modulation angles of 0.3 radians (Fig. 2a) and 1.5 radians (Fig. 2b) and no hysteresis. Input "power" in this case is actually mean-square input signal level, defined relative to output level number, and represents $\sigma_n^2 + A^2/2$, where A is sine-wave amplitude. A low data rate is assumed so that the quantizer is essentially in steady-state.

Figure 3 shows detected signal SNR loss in dB for modulation angles of 0.3 radians (Fig. 3a) and 1.5 radians (Fig. 3b) and hysteresis δ of 0.1, corresponding directly to

worst-case component values for the Mu-II quantizer. In both Figs. 2 and 3, degradation can be held to less than 0.1 dB over most of the range for power levels in the neighborhood of 5. Hysteresis effects, the motivator for the study described here, seem to contribute at most a few hundredths of a dB degradation in excess of an ideal quantizer at low SNR, and perhaps 0.2 dB at high SNR. The increasing degradation at higher SNRs can be rationalized in that such levels would seldom, if ever, be encountered in the DSN. For example, we could expect a -5-dB predetection SNR for a 10^6-bit/s data stream at 10^{-5}-bit error rate. Higher SNRs could be expected only with higher data rates and correspondingly stronger signals.

Figure 4a-e shows the effect of varying the hysteresis gap width on quantizer loss for several values of modulation angle, and a fixed power level of 5. For low predetection SNR, the effect is relatively insensitive to modulation angle. For hysteresis gap widths up to that encountered in practice, the quantizer degradation remains below 0.1 dB throughout the expected operating range. However, rapid deterioration appears with larger hysteresis gap widths.

V. Conclusions

Digital demodulation of spacecraft radio signals appears viable. No serious distortion problem is evident, and quantizers can be fabricated with existing components which degrade the subsequently detected signal by less than 0.1 dB over the expected operating ranges. Furthermore, it should be noted that a single quantizer used in the receiver replaces several nonlinear analog components which can induce distortion in the current receiver-demodulator implementation.

References

1. Martin, W. L., "System Performance of the Dual-Channel Mu-II Sequential Ranging," in *The Deep Space Network Progress Report 42-26*, pp. 54–68, Jet Propulsion Laboratory, Pasadena, Calif., Apr. 15, 1975.

2. Martin, W. L., and Zygielbaum, A. I., *Mu-II Ranging*, Technical Memorandum 33-768, Jet Propulsion Laboratory, Pasadena, Calif. (to be published).

3. Wax, J., "Quantizing for Minimum Distortion, *IRE Trans. Info. Theory*, Vol. IT-6, pp. 7–12, March 1960.
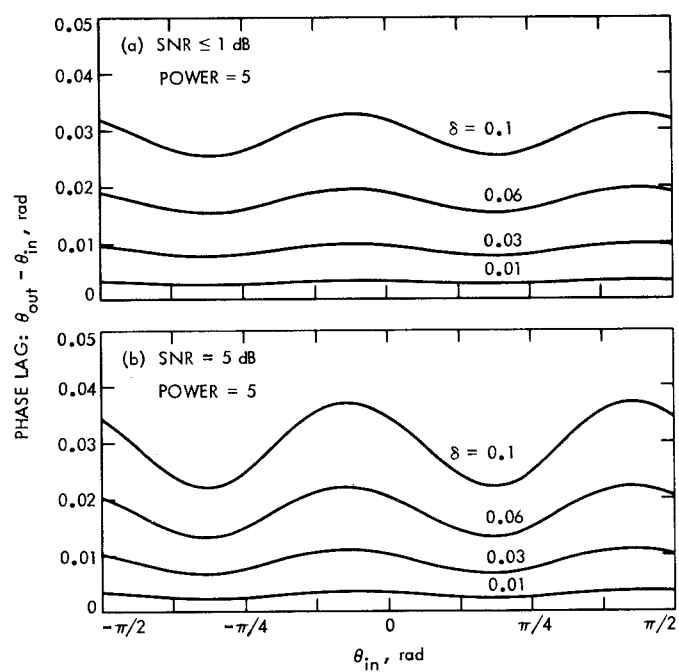
Fig. 1. Phase lag of quantizer with hysteresis as a function of
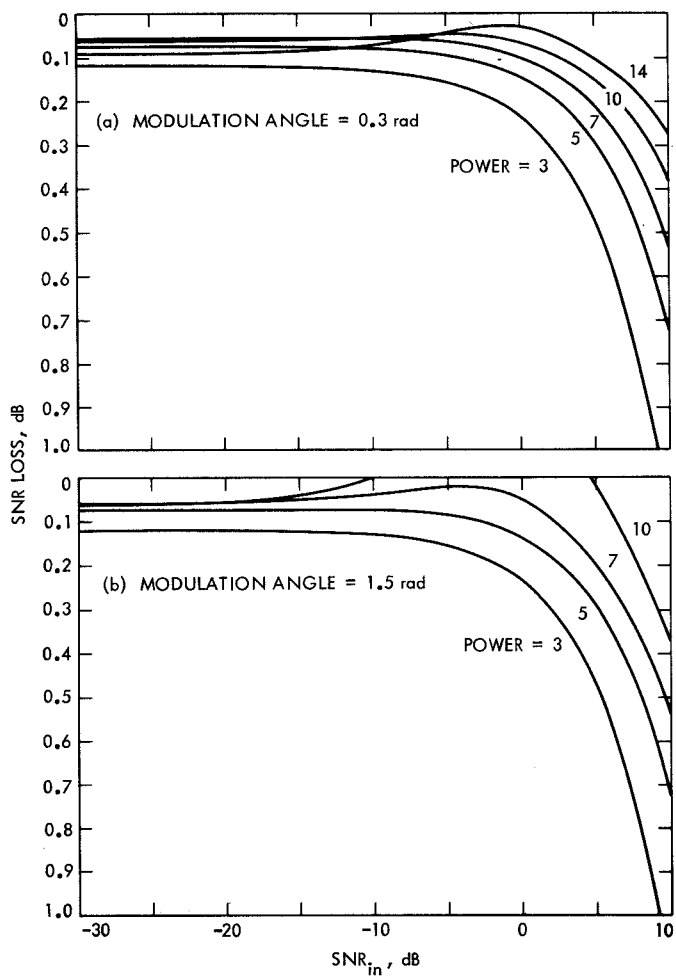sample phase for several hysteresis gap sizes

Fig. 2. Quantizer degradation vs input SNR for several input power levels, no hysteresis
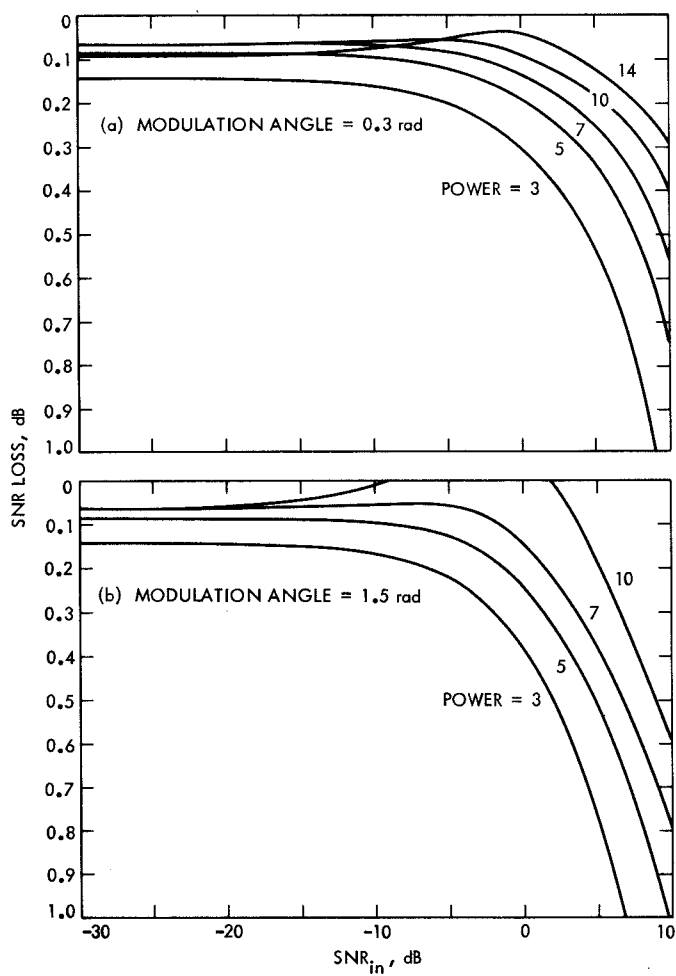

Fig. 3. Quantizer degradation vs input SNR for several input power levels, hysteresis gap = 0.1
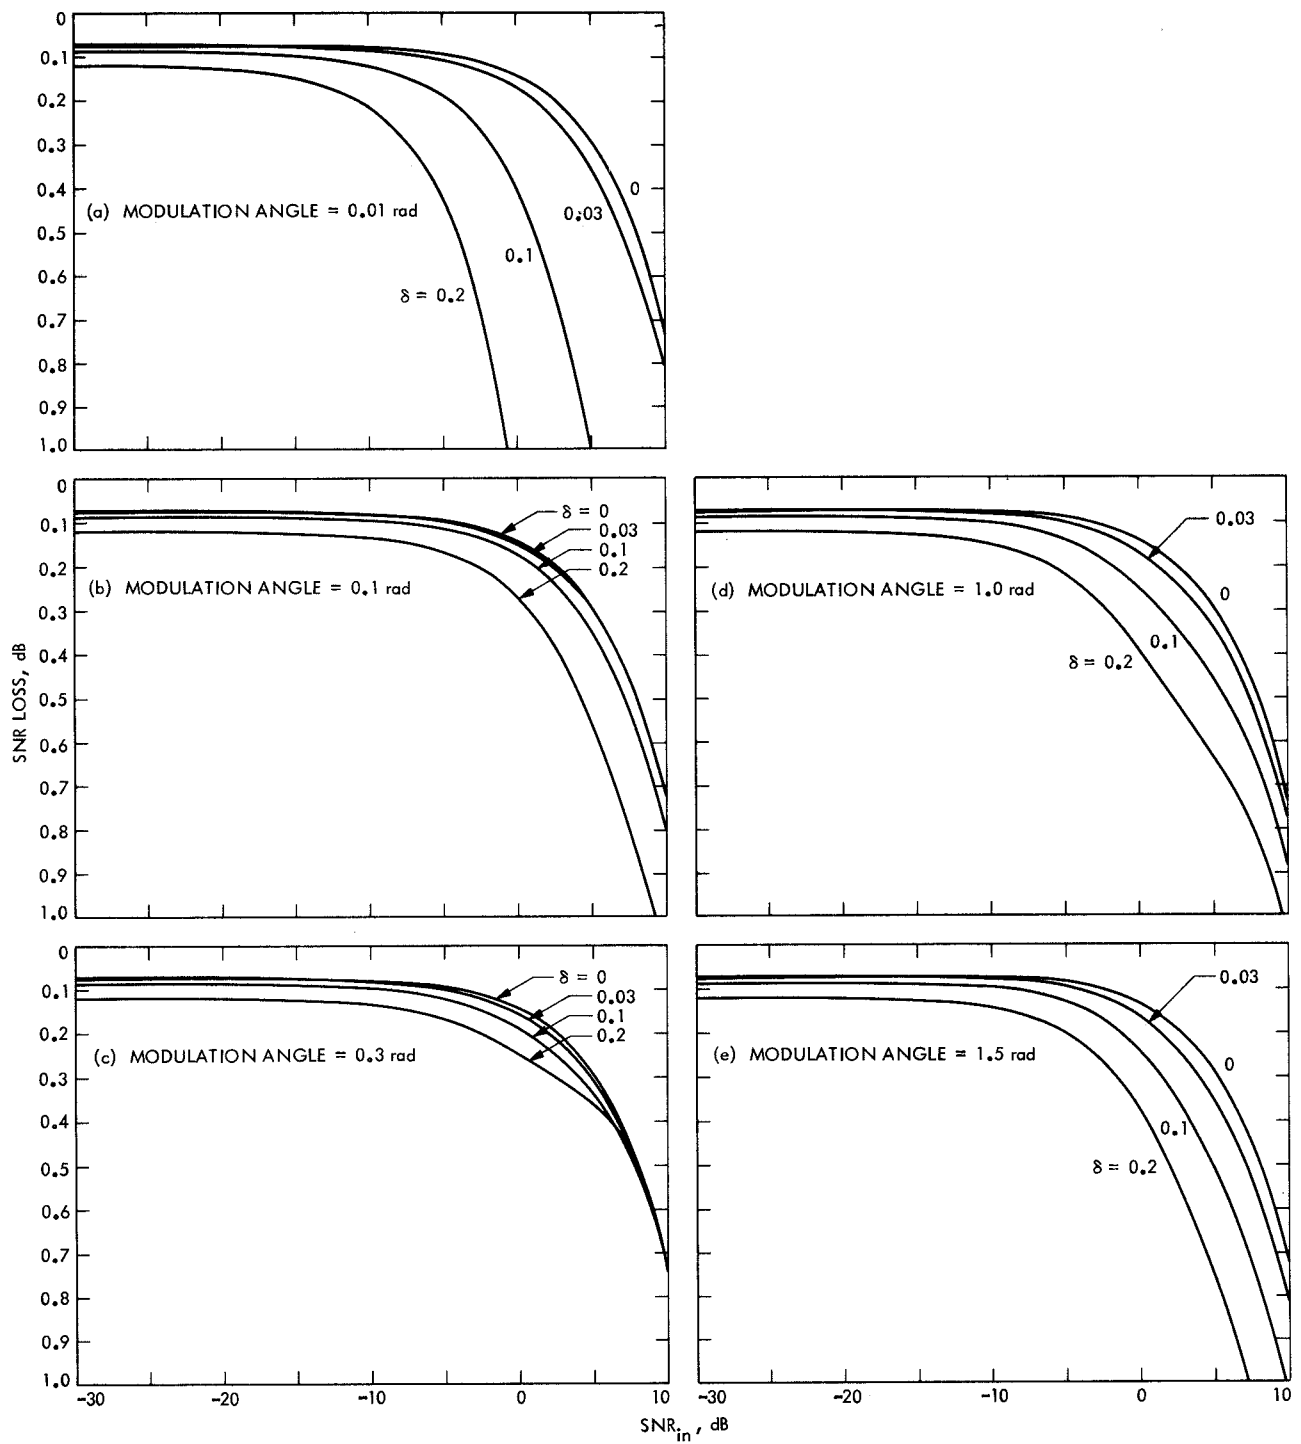
Fig. 4. Quantizer degradation vs input SNR for several values of hysteresis threshold gap, input power = 5