STI TR-4166
April 16, 1976

HIGH RATE BIT SYNCHRONIZER
FINAL REPORT

Prepared by
Richard F. Lyon

Submitted to
NASA Lyndon B. Johnson Space Center

Submitted under
Contract No. NAS 9-14590

STANFORD TELECOMMUNICATIONS INC.
1161 San Antonio Road • Mountain View, California 94043 • (415) 964-9290

STI TR-4166
April 16, 1976

HIGH RATE BIT SYNCHRONIZER
FINAL REPORT

Prepared by
Richard F. Lyon

Submitted to
NASA Lyndon B. Johnson Space Center

Submitted under
Contract No. NAS 9-14590

STANFORD
TELECOMMUNICATIONS INC.

Table of Contents

Section	Title	Page
1.0	INTRODUCTION AND SUMMARY	1-1
1.1	Analysis Summary	1-1
1.2	Experimental Results Summary	1-2
1.3	Recommendation Summary	1-3
2.0	ANALYSIS	2-1
2.1	Bandwidth Limiting	2-1
2.2	Baseline Correction	2-3
2.3	AGC Variations	2-4
2.4	Tracking Loop Delay	2-6
2.5	Lock Detector	2-10
2.6	Drift Due to Offsets	2-12
2.7	Phase Limit-Cycle Oscillations	2-13
2.8	NCO Quantization Effects	2-15
3.0	EXPERIMENTAL RESULTS	3-1
3.1	Bit Error Rate Degradation	3-1
3.2	Clock Jitter	3-1
3.3	Loop Response	3-5
3.4	Acquisition Performance	3-5
3.5	Baseline Variation	3-7
4.0	RECOMMENDATIONS	4-1
4.1	Mid-Phase Lock Detector Scheme	4-1
4.2	Timing Accuracy Improvement	4-1
4.3	Band Limiting Improvement	4-2
4.4	Loop Configuration	4-3
4.5	Fixed-Rate Simplifications	4-3
4.6	Fast Loop Filter	4-3
4.7	Sweep Circuitry	4-5
4.8	Acquisition and Tracking Modes	4-5

List of Figures

Figure No.		Page
2.1-1	Input Filter Step Response	2-2
2.2-1	Baseline Correction Filter Response	2-5
2.3-1	Threshold Spacing Variation	2-8
2.4-1	Loop Delay Timing	2-9
2.4-2	Loop Transfer Functions	2-11
2.7-1	Phase Limit-Cycle Oscillations	2-14
2.8-1	Quantization Envelopes	2-16
3.1-1	Measured Performance, Various Rates	3-2
3.1-2	Measured Coded Performance	3-3
3.2-1	Clock Jitter Without Noise	3-4
4.4-1	Improved Loop Configuration	4-4

List of Tables

Table No.		Page
2.3-T1	AGC Variations	2-7
3.3-T1	Measured Loop Characteristics at R_D = .8 Mbps, E_b/N_o = +3.7 dB	3-6
3.4-T1	Acquisition Conditions	3-6

1.0 INTRODUCTION AND SUMMARY

This report documents the analysis and experimental validation of the performance of various aspects of the STI Model 5000 Bit Synchronizer, which was built for NASA Johnson Space Center under contract NAS 9-14590.

1.1 Analysis Summary

Various non-ideal approaches were used in both the analog and digital subsections. Considerable analysis has been done to roughly optimize the filter parameters and to estimate degradations.

The following main points are derived in Section 2.0:

1) A degradation of less than 0.6 dB is caused by a bandwidth limiting first-order filter with 3 dB point at 2R (R is the symbol rate).

2) A second-order baseline correction filter with Q =.7 and 3 dB point at R/50 provides negligible phase shift and degradation, and the worst case baseline variation (peak amplitude equal to 100% of the signal peak-to-peak amplitude, at sinewave frequency R/100) has no significant degrading effect, even though the filter provides only 12 dB of rejection at that frequency.

3) The detection threshold spacing is not fixed with respect to either the signal level or the noise level, but varies, as a compromise between these two possible schemes. Over the range of SNR's used on coded links, the variation from optimal is negligible. The resulting variation in loop gain is significant, but is not considered a problem.

4) The digital data transition tracking loop components have delays that total about 35T through the proportional path (T is the symbol length) and 55T through the integrator path. This is considerably more than originally estimated, and explains the wide deviation from the nominal loop bandwidth at the higher settings.

5) The lock detector threshold setting can be either fixed or proportional to the level in the lock detector filter, but in either case, increasing the sensitivity, to acquire at lower SNR, increases susceptibility to unbalances in the analog circuitry.

6) Quantizer offsets much smaller than one level can cause a significant frequency drift in the loop with noise-only input. This can degrade the acquisition performance in the case of a frequency offset, and a sweep mode may be required in some applications to avoid this problem.

7) Coarse phase quantization can result in phase limit-cycle oscillations in a loop with delay when no noise is present to provide a random dither. At these high SNR's, no errors occur, so the degradation is negligible.

8) The effect of 7-bit angle quantization, 5-bit sine quantization, and three-pole filtering in the number controlled oscillator is a total clock phase jitter of less than $\pm$ 1%T under most conditions, but slightly more near worst-case conditions (fewer than five samples per cycle).

1.2 Experimental Results Summary

The following main points summarize the results of experimental validation, as discussed further under Section 3.0:

1) The BER degradation is typically 0.8 to 1.0 dB at high SNR, and 1.0 to 2.0 dB at very low SNR.

2) Clock jitter with no noise is about $\pm$ 1%, as predicted and jitter with noise is very low, also, until at low SNR the jitter increases very rapidly due to phase estimation errors.

3) The loop response agrees with the simple linear theory for the low bandwidth settings, but at the higher settings the effect of the delay is to cause higher than anticipated phase errors.

4) The lock detector performs as expected at the low rate settings, but at the higher rates it is too sensitive to imbalances in gains, offsets, capacitor values, and timing waveforms; thus, the adjustment is critical and the performance is compromised.

5) The sensitivity to the worst case baseline variation is negligible.

1.3 Recommendation Summary

Various design changes are recommended for inclusion in future versions of the bit synchronzier. The choice of which changes to implement will depend on the application and the development budget. Future development work recommended in Section 4.0 is summarized in the following items:

1) The two-channel mid-phase integrator implementation is not well suited to the lock detector approach, so one or the other should be changed. If an in-phase lock detector approach is used, further changes will be needed, as discussed in "Lock Detector for NASA Bit Synchronizer," STI TR-8055A.

2) Since a TTL gate delay of 10 ns is 5%T at the highest rate (5 Mbps), it makes sense to use a chain of fast logic (ECL) from the DAC through the I&D circuits. This was not done, due to a tight power goal.

3) The front-end configuration could be improved in one of several ways, to give less degradation from band limiting. One approach is to use a wider input filter and a separate pre-detection filter. The other is to do no filtering until after the ± 1 mixer, and then separately optimize the filters for the MP and IP paths.

4) The basic loop configuration could be changed, with the addition of a second NCO module, to give constant percentage bandwidth and damping with rate, without the gain shifter.

5) For fixed rate applications, simplifications can be made in many parts of the circuit. A single design that would work at any preset rate would be as useful in many applications as one with switch selectable rate.

6) To allow operation with wider bandwidths, a fast loop filter (no prefilter, less delay) could be used. This would be easiest to implement if the need for gain changing were eliminated.

7) Sweep circuitry could be added for more general application.

8) The acquisition and tracking modes could be generalized to allow operation with data formats other than Bi$\emptyset$-L.

2.0 ANALYSIS

Each point presented in the analysis summary, Section 1.1, is expanded in this section.

2.1 Bandwidth Limiting

A lowpass filter preceding the data matched filter decreases the response to the data signal, and also decreases the response to noise. The net result is a degradation if the noise is white and the channel is linear-phase, since the matched filter alone is optimal.

For a first order lowpass with time constant $\tau = 1/\omega_c$, the response to data is easily calculated and the response to noise can be approximated by hand calculation. Figure 2.1-1 shows the response of the filter to a data step, and the areas of missing correlation with respect to the wideband case, assuming optimum timing (integrating between zero crossings).

For simplicity we use signal amplitude and symbol duration equal to unity. Then we can easily see that (assuming $\tau << \frac{1}{2}$):

$$\text{Area 1} = \int_{\tau 1n2}^{\infty} 2 \exp(-t/\tau)\,dt = \tau$$

$$\text{Area 2} = 2\tau 1n2 - \int_{0}^{\tau 1n2} 2 \exp(-t/\tau)\,dt = \tau(2 1n2 - 1)$$

Only adjacent symbols interfere significantly, so we find the response to all four cases of preceding and following symbol interference to be as follows:

1 - 2(Area 1) - 2(Area 2)
1 - 2(Area 1) - Area 2
1 - Area 1 - 2(Area 2)
1 - Area 1 - Area 2

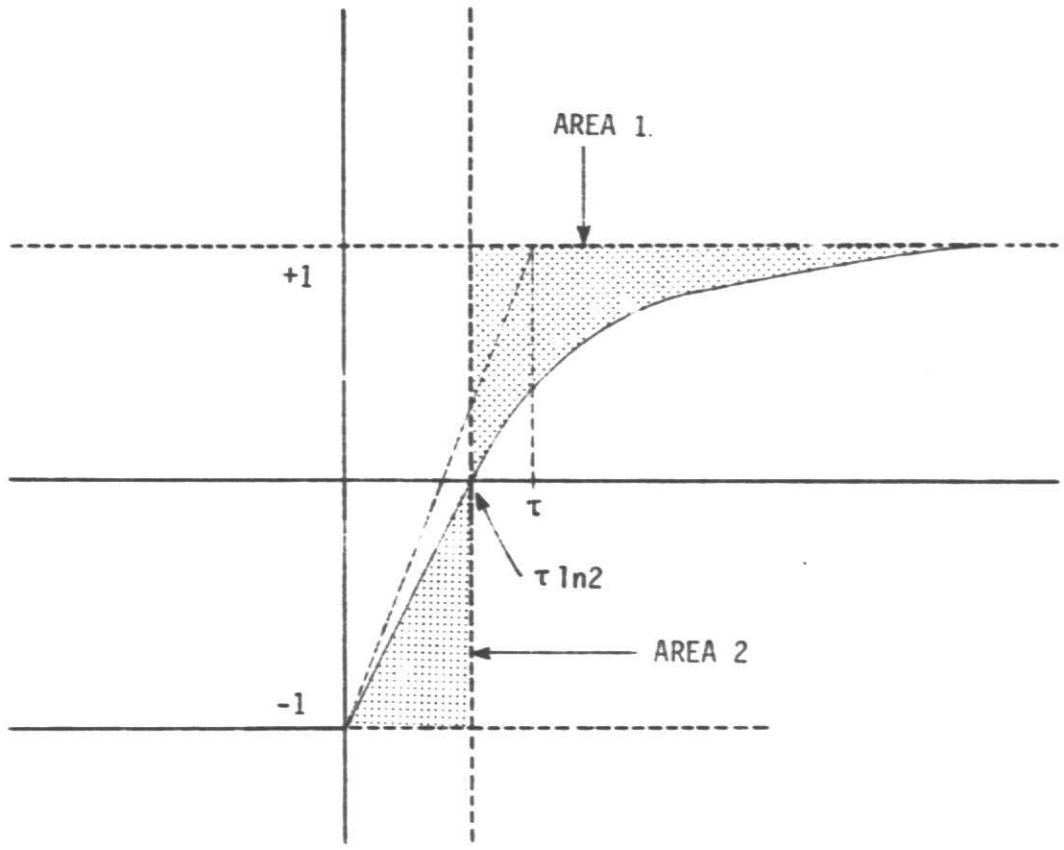

FIGURE 2.1-1 INPUT FILTER STEP RESPONSE

In the case of 3 dB point at twice the data rate, $f_c = 2$, $\omega_c = 4\pi$, and $\tau = 1/(4\pi)$. Then Area 1 = .0796 and Area 2 = .0307, so the four signal degradations are:

.779	(-2.17 dB)
.810	(-1.83 dB)
.859	(-1.32 dB)
.890	(-1.01 dB)

The noise is also reduced at the matched filter output by the amount:

$$\frac{\int |\text{matched filter response}|^2 \cdot |\text{input filter response}|^2 \, df}{\int |\text{matched filter response}|^2 \, df}$$

Using $|\text{input filter response}|^2 = 4/(4 + f^2)$, we can approximate the integral over each lobe of the matched filter response by multiplying each lobe height by the filter power response at the lobe center. Then the noise response is:

$$\frac{\sum_{i=1}^{\infty} \frac{1}{(2i-1)^2} \cdot \frac{4}{4+(2i-1)^2}}{\sum_{i=1}^{\infty} \frac{1}{(2i-1)^2}} \approx \frac{.842}{1.21} \approx .696 \ (-1.57 \text{ dB})$$

The result is that in two cases the effective SNR is improved (by +.56 and +.25 dB) and in two cases the SNR is degraded (by -.26 and -.60 dB). The degraded case dominates, especially at high SNR where the error rate curve is very steep. Hence the degradation approaches -.60 dB.

2.2 Baseline Correction

The maximum specified baseline offset peak value (equal to the signal peak-to-peak value) is enough to completely disrupt

the mid-phase integrals and prevent synchronization. Suppressing the baseline offset by a factor of four (-12 dB) gives random $\pm 45^{o}$ peak phase estimate errors. The resulting rms estimation error of $\pm 32^{o}$ with sinusoidal baseline offset is reduced by the tracking loop to give an actual phase error of about

$$\phi_{err} = 32^{o}\sqrt{\frac{B_L}{2R_D}} \text{ rms}$$

$$= 2.9^{o} \text{ worst case}$$

To provide the 12 dB suppression, a second order highpass filter with Q = .7 is used. The transfer function is

$$H(s) = \frac{s^2}{s^2 + 1.4\, s\omega_c + \omega_c^2}$$

and the step response is

$$g(t) = e^{-.7\,\omega_c t}\left[\cos .7\,\omega_c t - \sin .7\,\omega_c t\right].$$

Using $f_c = R/50$, or $\omega_c = 2\pi\, R/50$, with R = 1:

$$g(t) = e^{-.089t}\left[\cos .089t - \sin .089t\right]$$

which has an initial height of 1 and an initial slope of $-1.4\,\omega_c = .178$. The resulting slant on a square wave at the symbol rate is about $\pm .044$ from the nominal flat level at 1.0. See Figure 2.2-1. The resulting static timing error of .0055T is negligible. For random data patterns the average phase error is somewhat greater than this.

2.3 AGC Variations

As the SNR changes, the signal level out of the AGC changes, since the combined signal plus noise power in the input filter

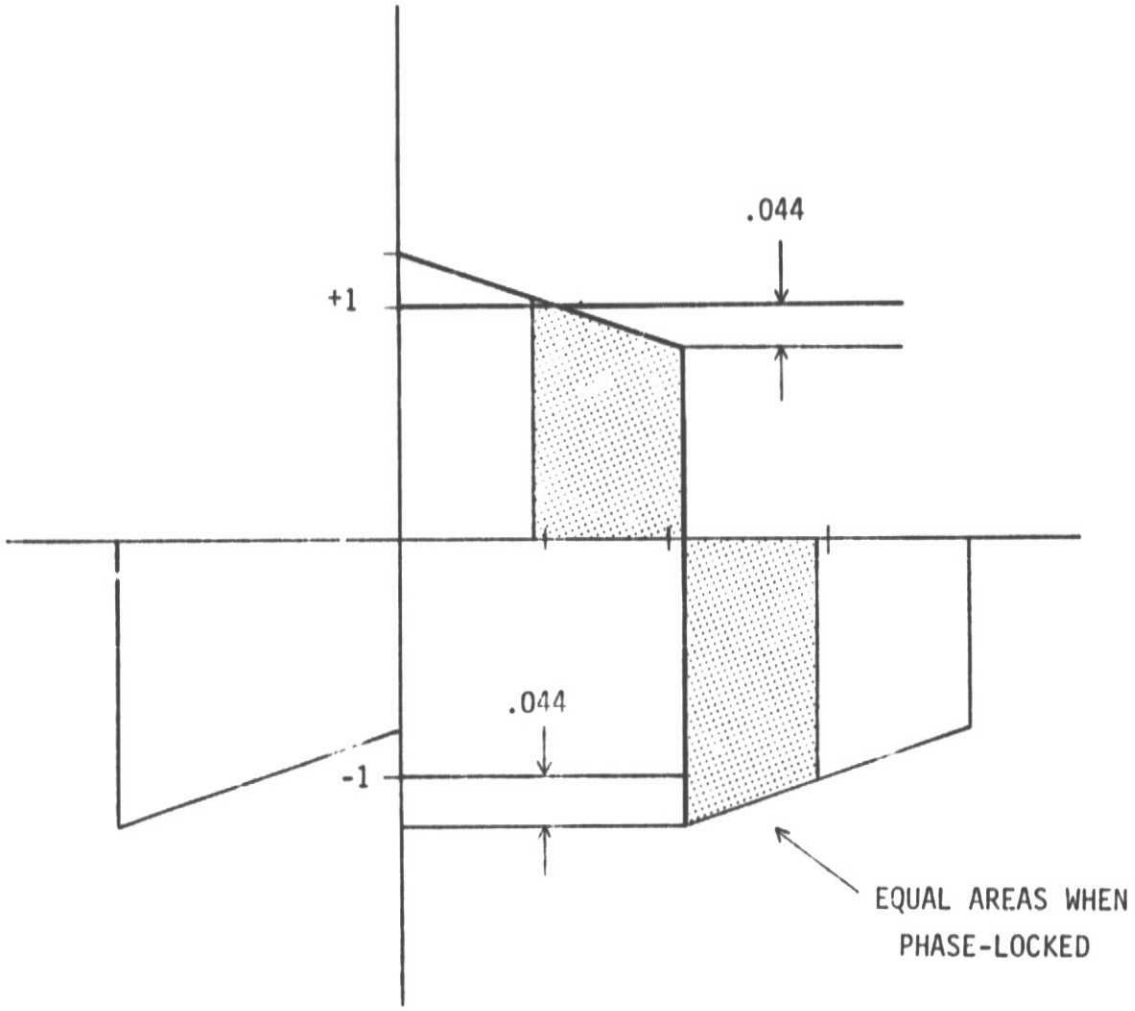

FIGURE 2.2-1 BASELINE CORRECTION FILTER RESPONSE

bandwidth is held constant. The input filter bandwidth for a first order cutoff at 2R is πR. Assuming unity total power we calculate the signal and noise components in Table 2.3-T1. The quantizer threshold spacing, Δ, is often specified in terms of either the signal amplitude A or the postdetection noise standard deviation σ, but is rarely fixed in relation to either. The two normalizations Δ/A and Δ/σ are plotted in Figure 2.3-1 for the above input bandwidth, assuming the nominal spacing $\Delta = .35A$ is set at $E_b/N_o = +3$ dB. That is,

$$\Delta = .35A \ (+ 3 \text{ dB})$$
$$= .35 \times .623 = .218$$

The tracking loop gain is proportional to the signal amplitude, so the loop characteristics also vary with the SNR. Both the loop natural frequency and the damping vary in proportion to the square root of the loop gain, so there is less than a factor of two change over the entire operating range.

2.4 Tracking Loop Delay

A twenty-sample accumulator (prefilter) is used after the phase error logic to reduce the data rate in the loop filter. The result was originally estimated to be simply a delay of 10 bits on the average, which was not considered to be a problem. Closer examination shows that the gain shifter, integrator, synchronizing buffer, and the phase logic itself add delays, for a total of 35 bits in the proportional path and 55 bits in the integrated path; the approximate timing is diagrammed in Figure 2.4-1.

In an ideal second-order loop with damping factor $\zeta = 1$ the loop filter transfer function is a constant times the following:

$$F(\omega) = 2 + \frac{\omega_n}{j\omega}$$

Table 2.3-T1 AGC Variations

E_b/N_o	$\frac{S}{N}$ in πR	S	N	$A = \sqrt{S}$	$\sigma = \sqrt{N \frac{.5R}{\pi R}}$	Δ/A	Δ/σ
-6 dB	.0796	.074	.926	.272	.384	.804	.568
-3 dB	.1591	.137	.863	.371	.370	.589	.590
0 dB	.318	.241	.759	.491	.347	.444	.629
3 dB	.637	.389	.611	.624	.312	.350	.700
6 dB	1.273	.560	.440	.748	.265	.292	.824
9 dB	2.546	.718	.282	.847	.212	.258	1.030
12 dB	5.09	.836	.164	.914	.162	.239	1.348
∞	∞	1.000	.000	1.000	.000	.218	∞

Δ = .35 A (+ 3 dB) = .2183

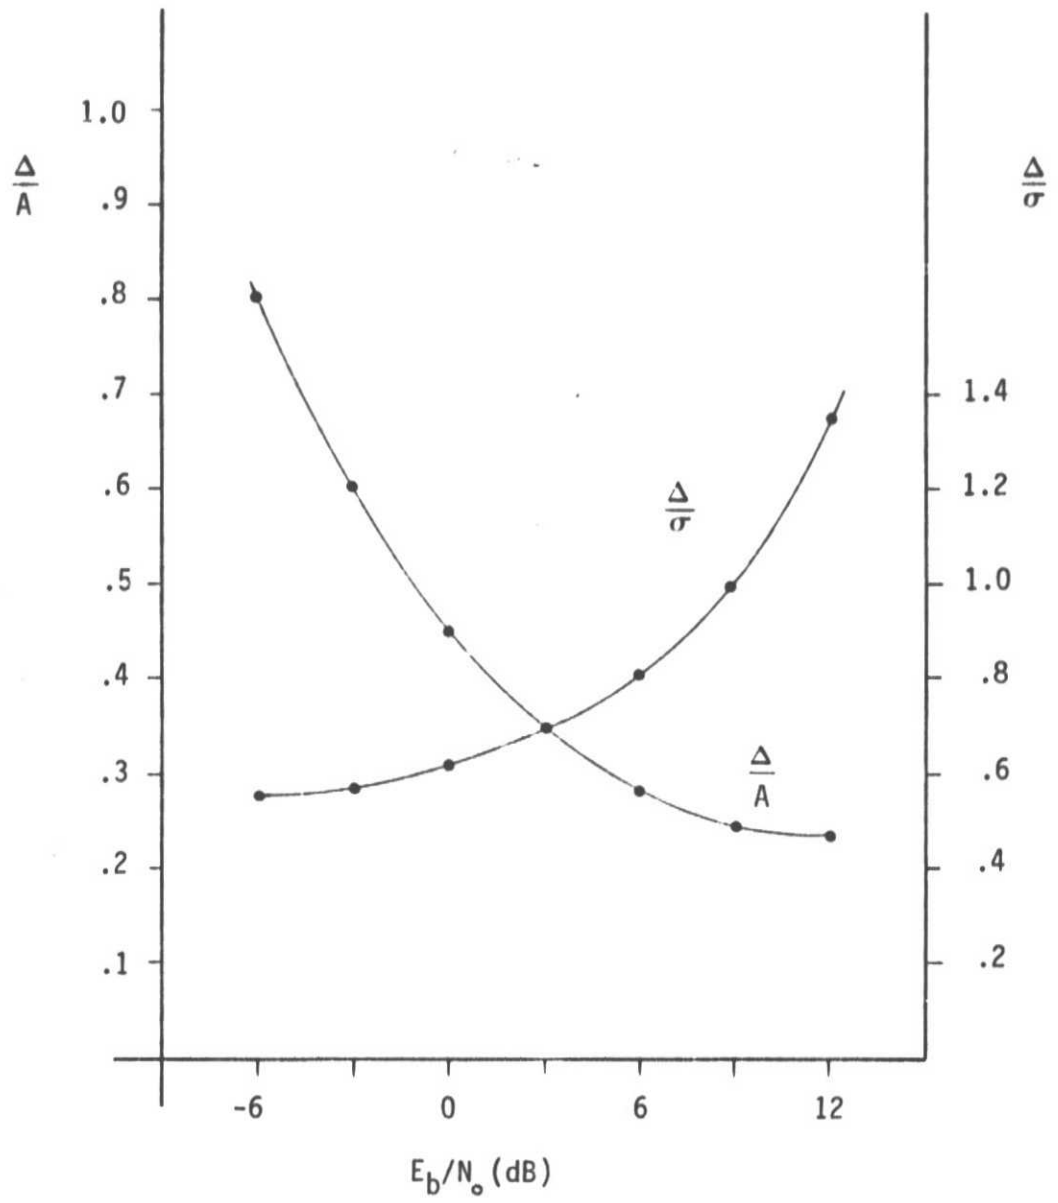

FIGURE 2.3-1 THRESHOLD SPACING VARIATION

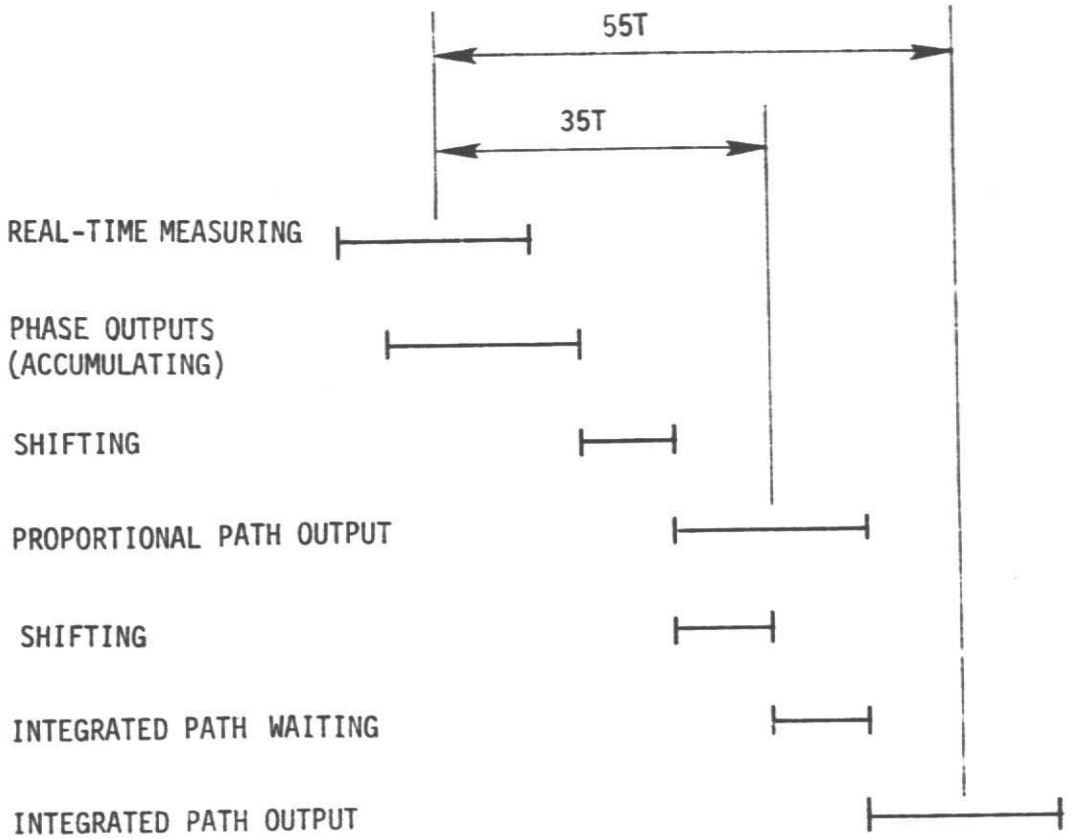

FIGURE 2.4-1 LOOP DELAY TIMING

The two delays add a frequency dependent phase shift in each term as follows:

$$F(\omega) = 2 \exp(-j\omega \tau_1) + j \frac{\omega_n}{\omega} \exp(-j\omega \tau_2)$$

where τ_1 is 35T and τ_2 is 55T.

The loop gain, including the NCO as an integrator, is

$$\frac{\omega_n}{j\omega} F(\omega)$$

which results in a transfer function (from phase input to phase output) of

$$H(\omega) = \frac{\frac{\omega_n}{j\omega} F(\omega)}{1 + \frac{\omega_n}{j\omega} F(\omega)}$$

The transfer function is plotted in Figure 2.4-2 for the five settings actually used, with nominal ω_n = .0005, .001, .002, .004, .008.

2.5 Lock Detector

The calculations of lock detector thresholds done in the report TR-8055A were too pessimistic by 3 dB because of a factor of two error in the half-bit postdetection SNR, which should be equal to E_b/N_o, not $E_b/2N_o$. Thus, changing the lock threshold from 31/32 to 15/16 gave satisfactory results in most cases.

Since the threshold was implemented as a fraction of one data channel, the sensitivity to imbalances was increased at very low input levels (below specified operating range) and a false lock indication could occur. The AGC normally stabilized the

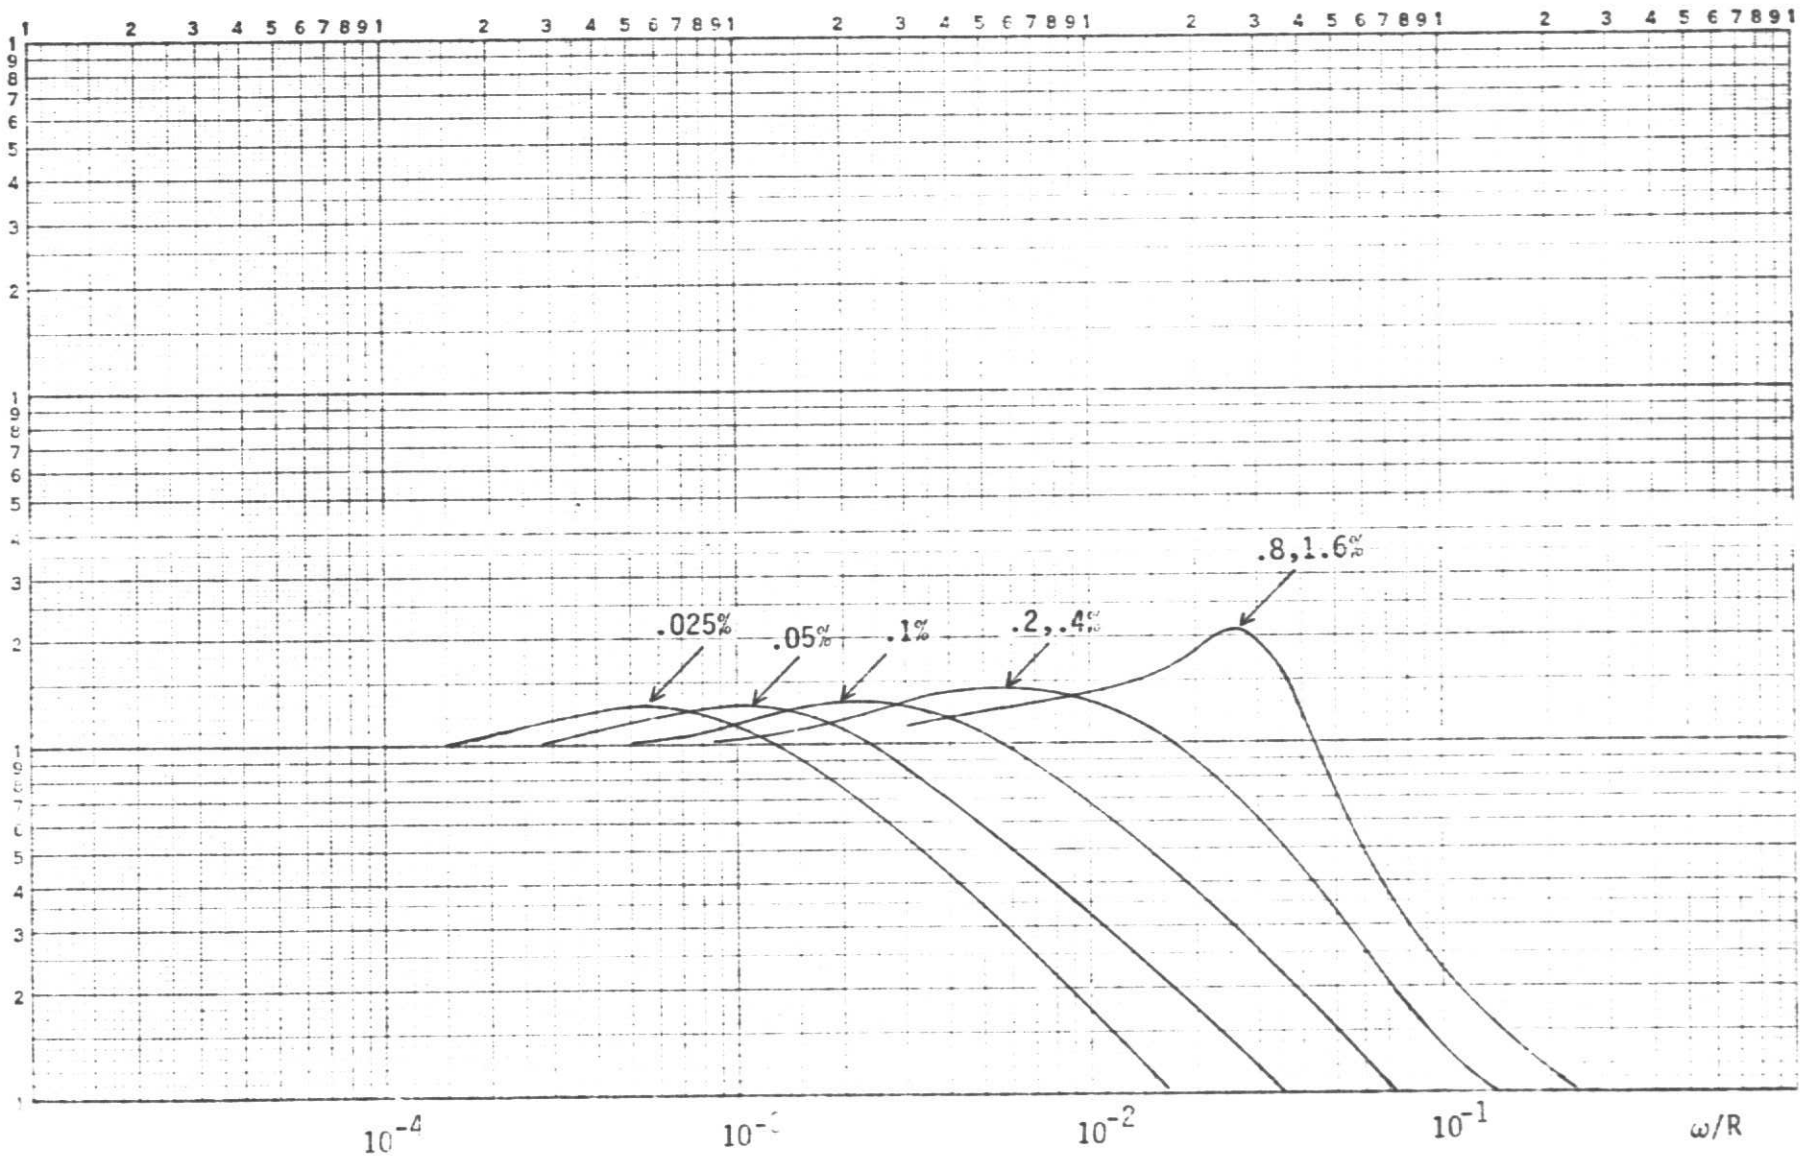

FIGURE 2.4-2 LOOP TRANSFER FUNCTIONS

levels, so threshold did not vary much; a typical threshold on the difference of the two channels was chosen and hard-wired for use at all conditions.

The resulting false-lock and false-unlock probability calculations could be repeated, giving results similar to those in TR-8055A, but the actual performance is limited predominantly by the accuracy of balancing the analog channels over the range of rates.

2.6 Drift Due to Offsets

The acquisition time for a narrow loop and signal with frequency offset is calculated based on ideal balance at the loop filter integrator when no signal is present. An offset in the phase error estimate causes the frequency to drift in one direction, which may be opposite to the direction of the signal.

In the worst case for acquisition, the loop bandwidth is .025%, the signal frequency offset is .05%, and acquisition should occur in 250,000 bit times (one second at the lowest rate). At these settings the loop integrator gain gives a frequency drift of about 1 KHz/sec for one unit (LSB) of offset. An offset as small as 1/16 unit may be enough to swamp out the effect of the signal, which initially pulls very weakly relative to the average rate of 125 Hz/sec that it needs to acquire within specs.

This phase error estimate mean value offset should not be strongly affected by an offset in either the MP or IP quantizer, since their outputs are effectively multiplied together in the phase error logic. However, small offsets in both quantizers will combine nonlinearly to give an offset in the phase estimate. Small timing errors that change the relative correlation between

the MP value and the preceding and following TP values may also cause a net offset.

These effects only become significant when the signal is outside the bandwidth of the tracking loop ($\pm B_L/2$); for operation under such conditions a sweep circuit is recommended.

2.7 Phase Limit-Cycle Oscillations

The output of the phase estimator is quantized to only three bits, but much of the quantization noise is filtered out by the prefilter (20 sample accumulator) if noise is present to provide a random dither. If no noise is present, many successive phase samples may have quantization errors in the same direction, causing a large error in the low-rate sample sent to the loop filter. This will cause an overcorrection of the phase, so that the phase error estimate will oscillate between the two levels bordering on zero, as quickly as the loop dynamics will allow.

The oscillation frequency is high enough that for approximate analysis the integrator path may be neglected. Ignoring the discrete timing of the signals, Figure 2.7-1 shows that the oscillation is stable where the proportional path delay introduces a 90° phase shift. Notice this is independent of the loop gain (or bandwidth) setting, but the amplitude of the actual phase oscillation is proportional to the gain (the proportional path gain is proportional to the loop bandwidth in a first- or second-order loop).

The period of oscillation is then 4 X 35T, which yields 1600 Hz frequency steps at the widest B_L at 1 Mbps. The resulting peak phase errors are $\pm 20^\circ$ by this calculation. Probably due to phase dither from the NCO, smaller oscillations were observed. Fortunately, very small amounts of noise are enough to suppress this effect entirely.

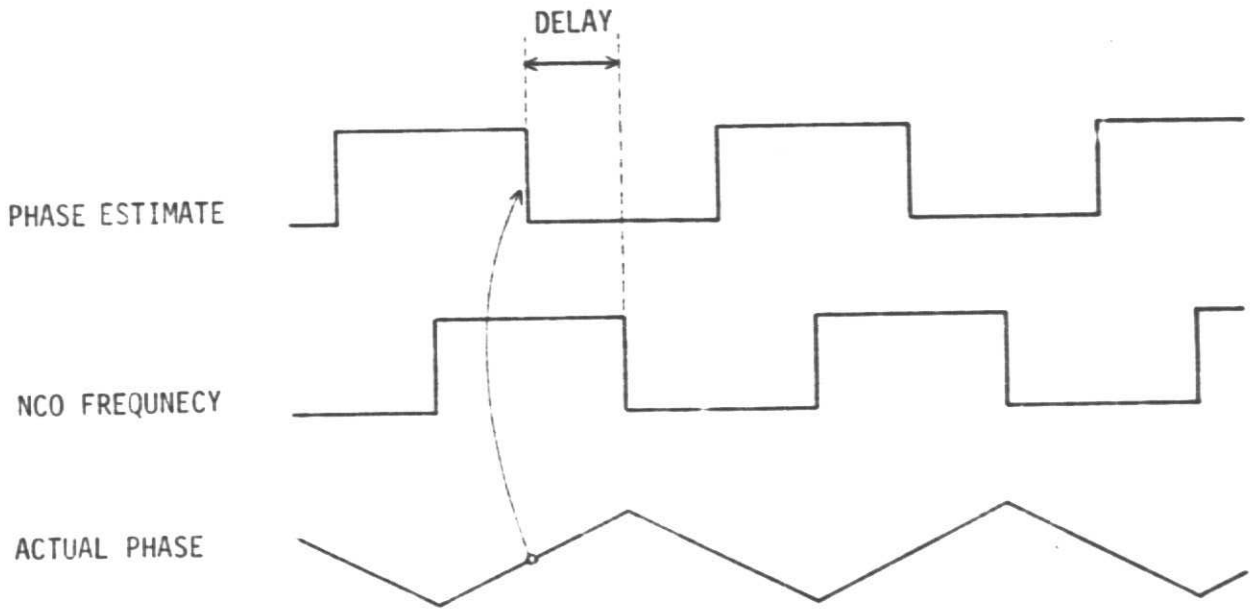

FIGURE 2.7-1 PHASE LIMIT-CYCLE OSCILLATIONS

2.8 NCO Quantization Effects

The signal to noise ratio of a quantized sine wave is usually approximated as 6 dB per bit. For 5-bit quantization, this predicts 30 dB SNR, neglecting angle quantization and imperfect filtering. Figure 2.8-1 shows graphically that the envelope of the quantized waveform has 50% more zero-crossing error with 7-bit angle quantization than with unquantized angle. The total is $\pm 2.8^{\circ}$, or $\pm 0.78\%$ of a cycle. An additional noise is added by the nonlinear response of the digital-to-analog converter during switching.

The image frequencies of the sampled sine wave are suppressed both by the zero-order-hold response of the DAC and by a third order lowpass filter. The first image of a 5 MHz signal is at 15.97 MHz, and is suppressed by sinc (15.97/20.97), or -10.9 dB by the zero-order-hold. The lowpass filter response is not quite 18 dB per octave, yielding close to another 30 dB rejection. This keeps the image power less than the noise power due to quantization.

The various phase errors due to quantization, DAC switching, and imperfect image rejection all depend strongly on the exact frequency setting or data pattern in the NCO. Thus the worst case errors, which occur at various frequencies near the upper limit, occasionally exceed the $\pm$ 1% specification. The rms phase error seems to stay below 0.5%. Errors of this size have negligible effect on the overall performance.

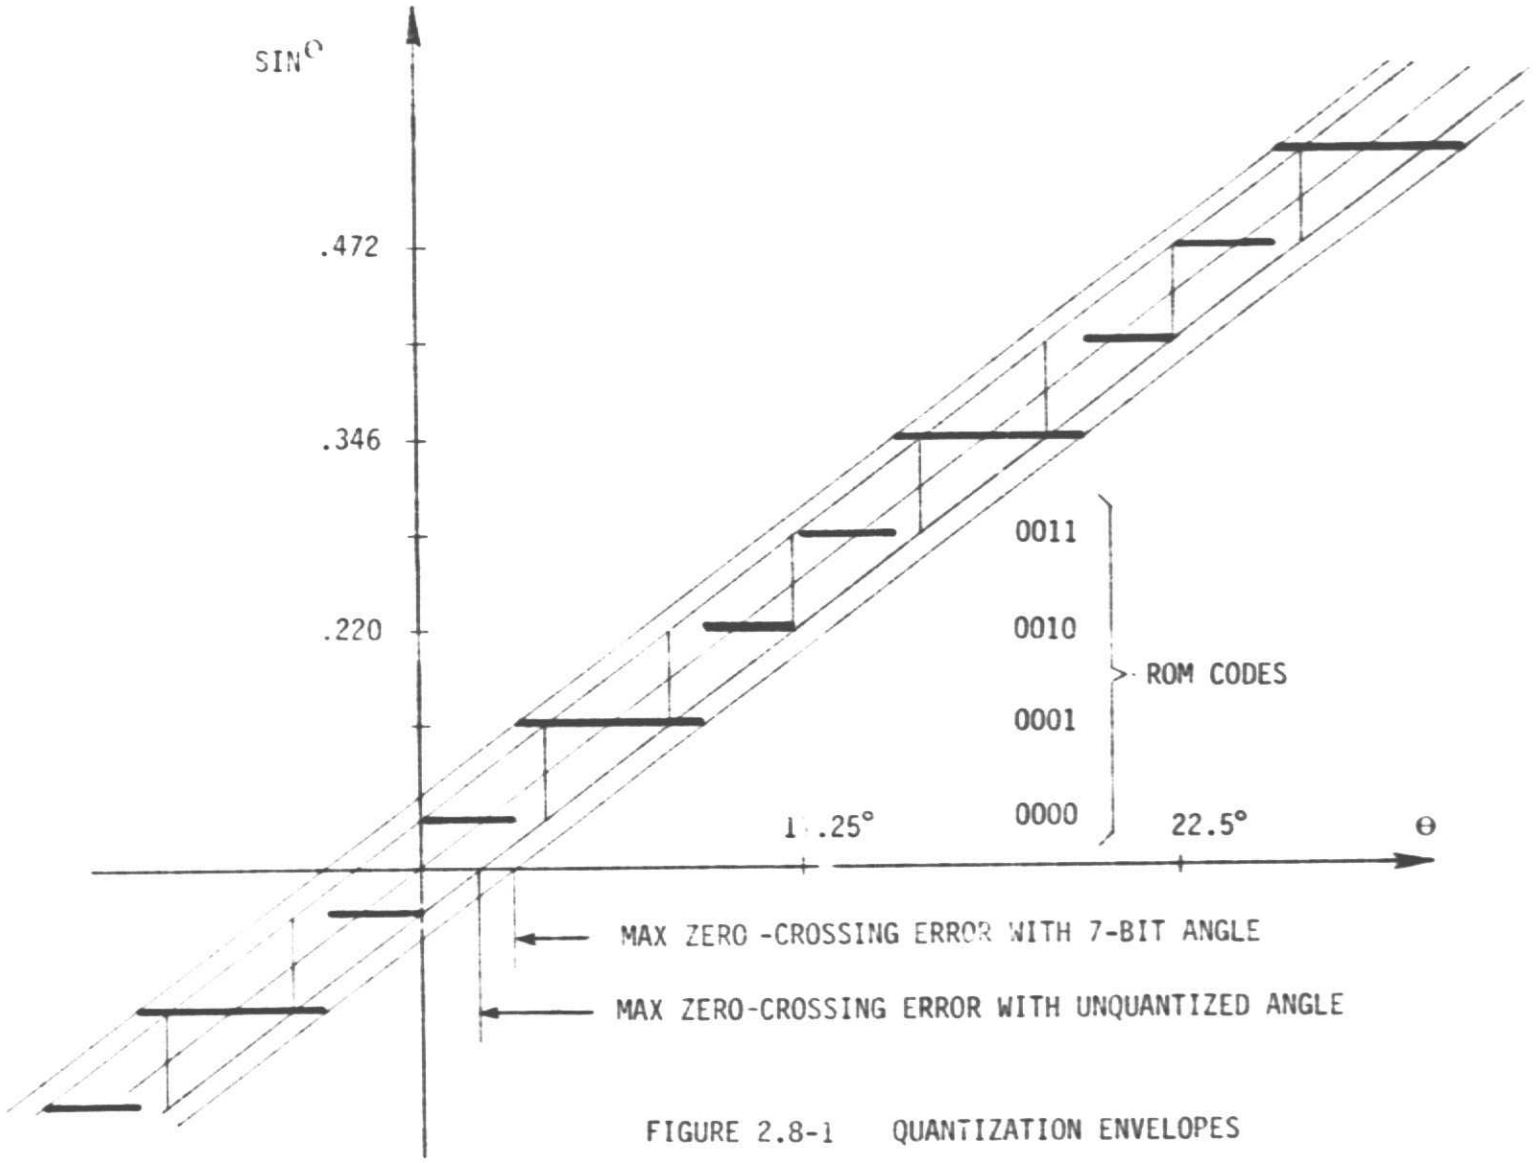

FIGURE 2.8-1 QUANTIZATION ENVELOPES

3.0 EXPERIMENTAL RESULTS

Some experimental data are presented in this section to support the points in Section 1.2.

3.1 Bit Error Rate Degradation

The performance of the bit synchronizer was documented at five rates during acceptance testing. Figure 3.1-1 shows the measured points in relation to the theoretical curve. The exceptionally good performance at the highest rate is an artifact of the measurement process: the power meter response is only calibrated to 10 MHz, so that the signal power is underestimated by about 1 dB.

The soft decision performance was also measured in a coded link with a Linkabit LV7015 LR Viterbi Decoder. The symbol rate was 3 Msps, with E_s/N_o between -3 and 0 dB. Figure 3.1-2 shows the result, compared to the performance with the LV7015 LR built-in digital psuedo-noise source. NASA experience indicates that the degradation with this decoder is always worse than the hard decision error rate degradation of a bit synchronizer.

3.2 Clock Jitter

The absolute and percentage clock jitter was measured on all four clock phase outputs at three rates. These data are plotted in Figure 3.2-1. A low bandwidth setting and a clean signal were used for these measurements.

The performance with noise is not available in quantitative terms, but in general the result is that the phase jitter is negligible for error rates less than 10%, due to the low bandwidths of this synchronizer.

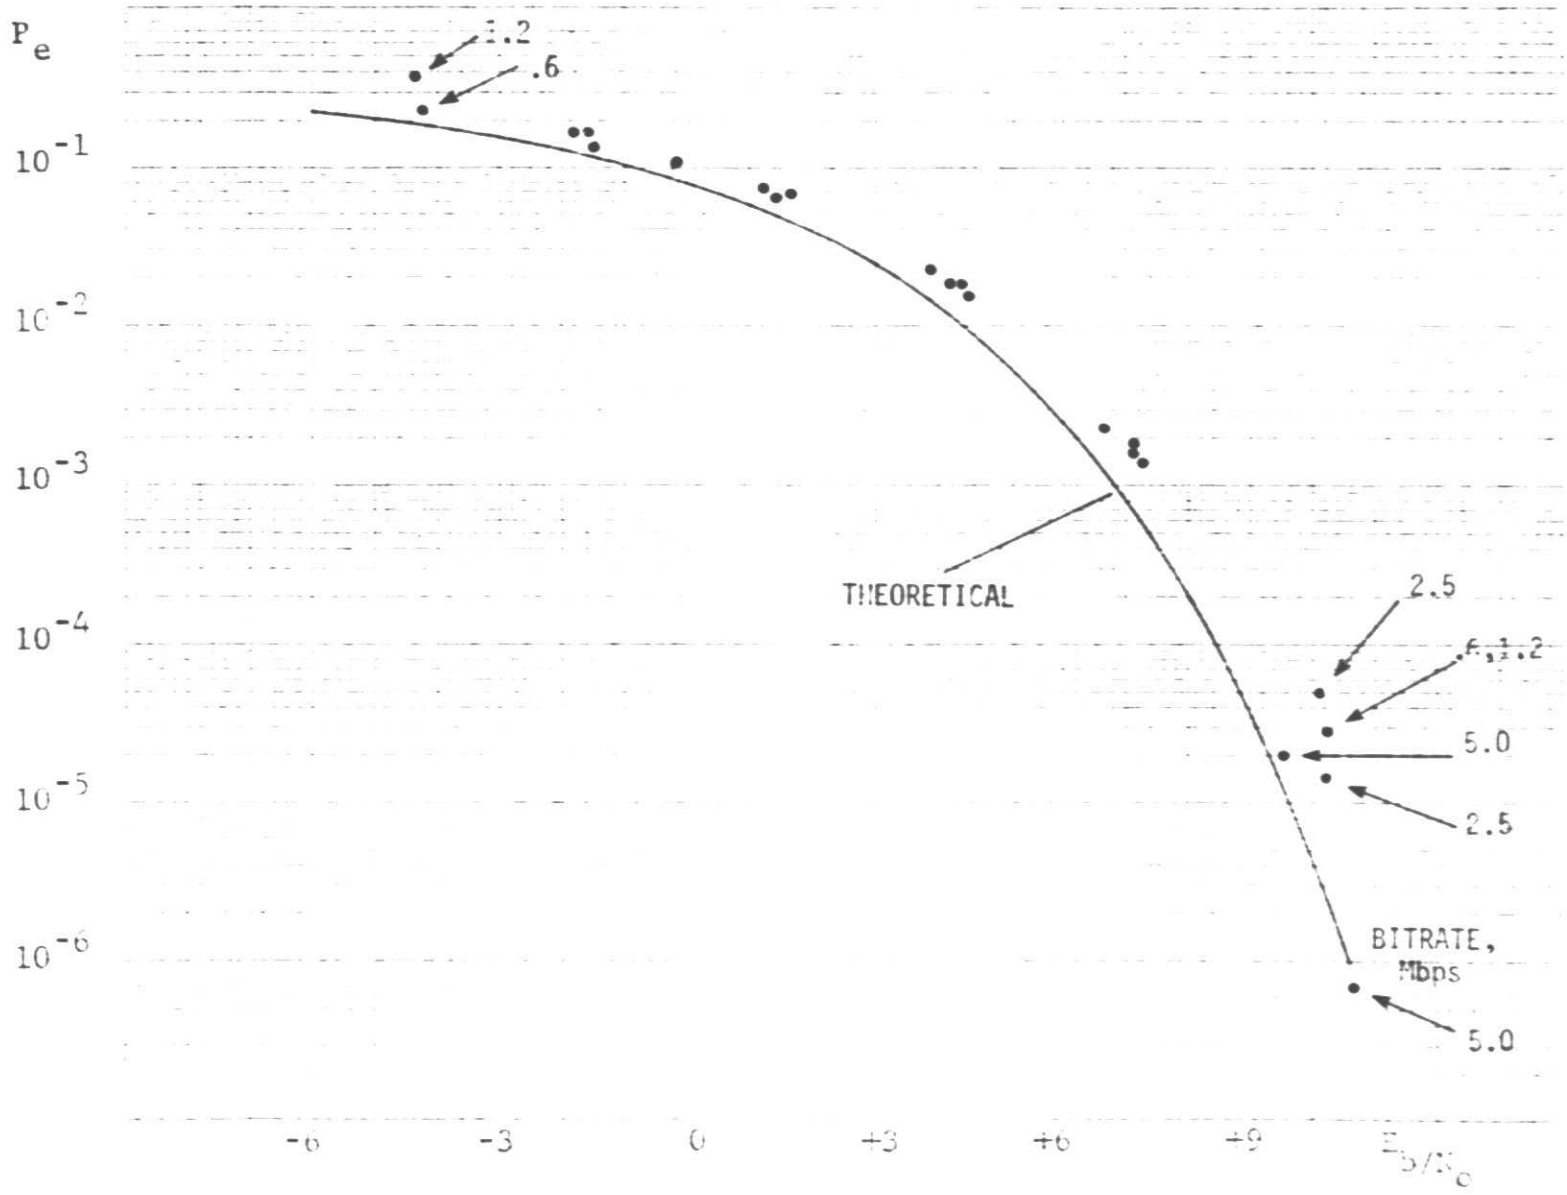

FIGURE 3.1-1 MEASURED PERFORMANCE, VARIOUS RATES

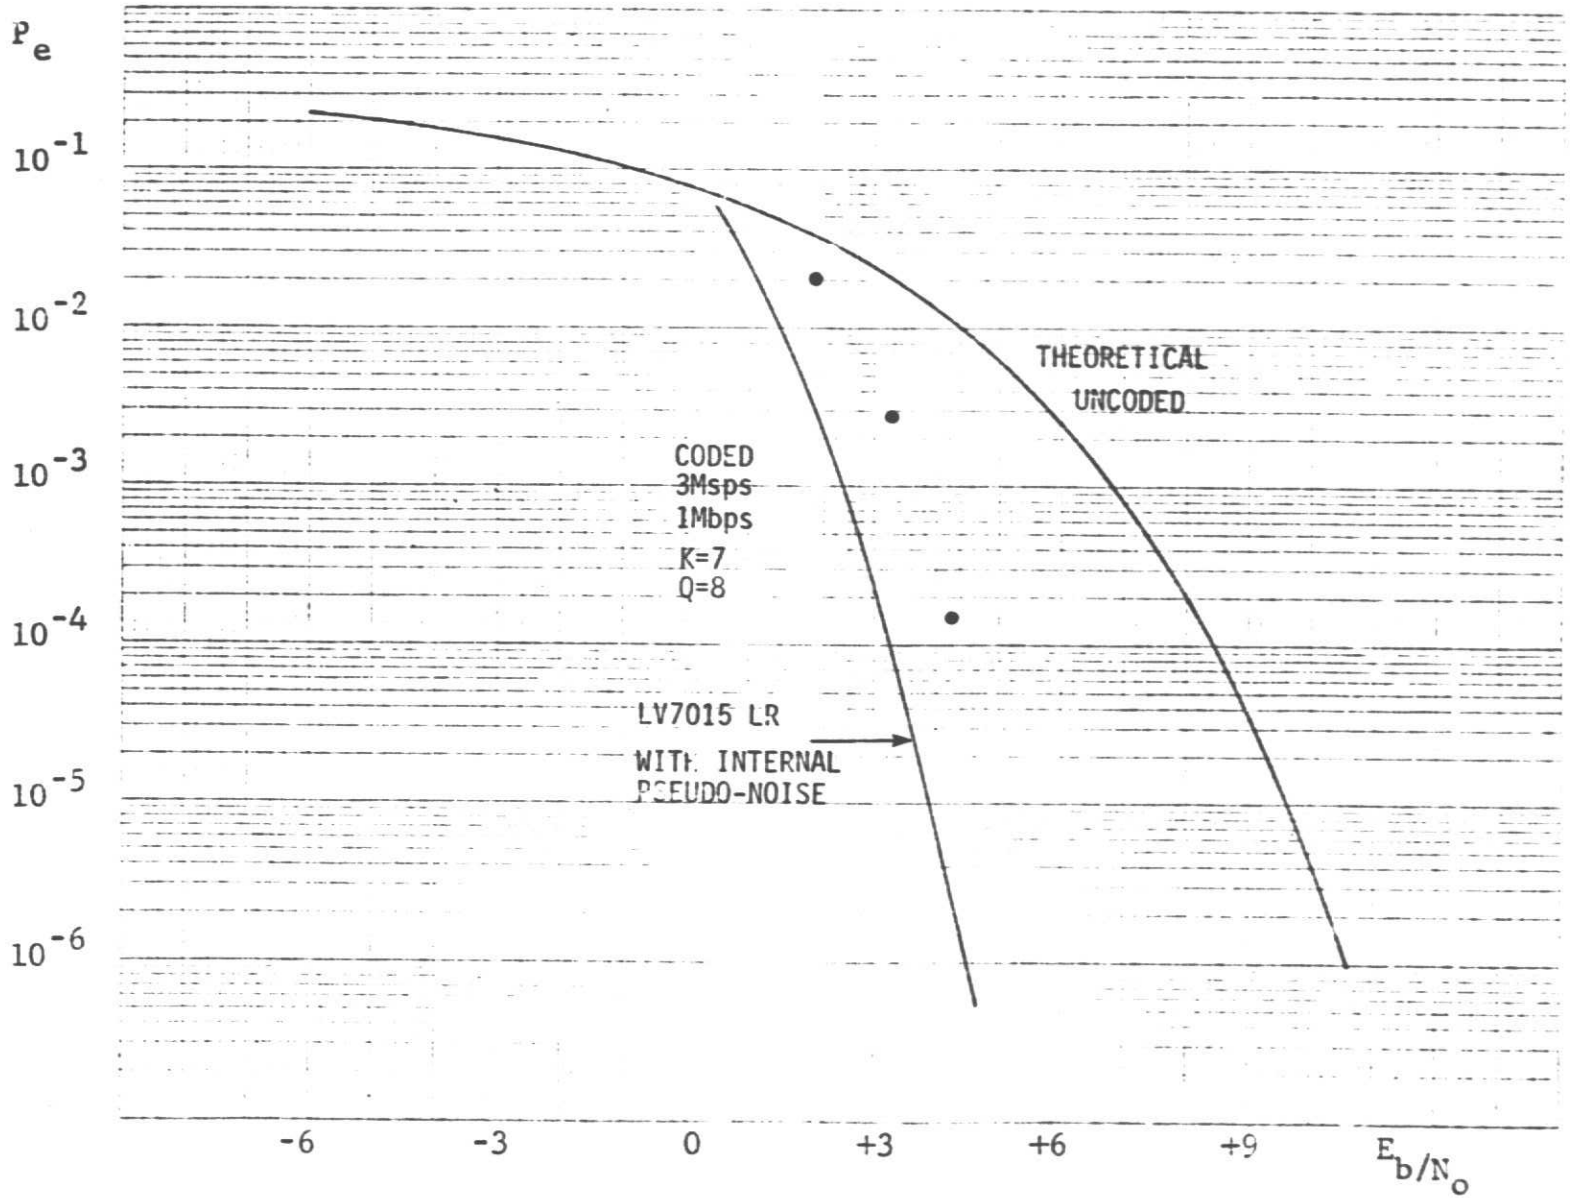

FIGURE 3.1-2 MEASURED CODED PERFORMANCE

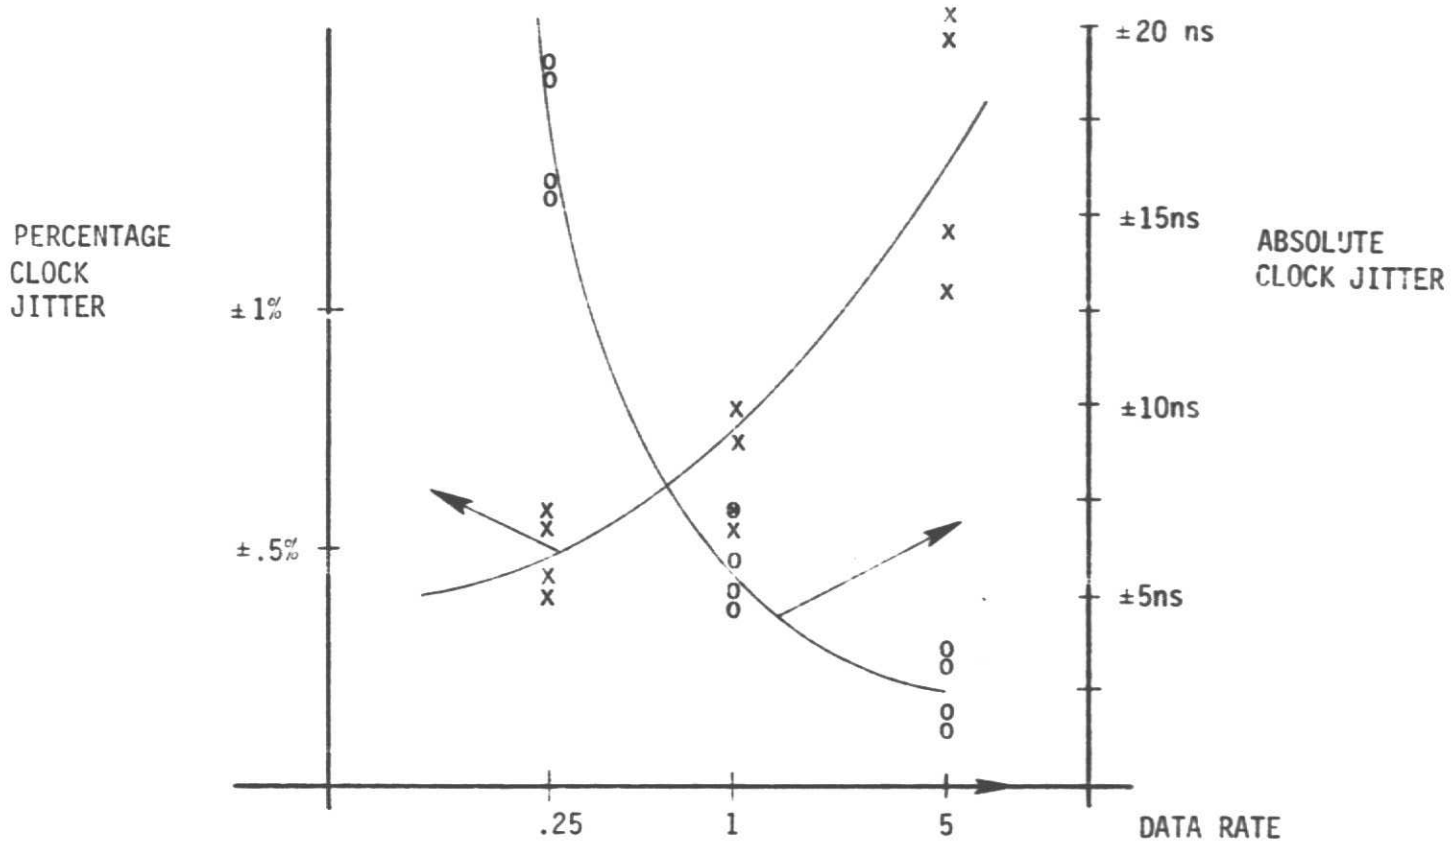

FIGURE 3.2-1 CLOCK JITTER WITHOUT NOISE

3.3 Loop Response

The response of the loop to a sinusoidally FM modulated data signal has been documented at one rate and two bandwidth settings. By analogy with an ideal second-order loop, we measure the frequency for which phase error peaks, and label that the loop natural frequency ($f_m = f_n$). The modulation index (Δf) is adjusted to a peak phase error of 1 rad ($\theta_e = 1$). Then Figure 4-1 in Phaselock Techniques, by Floyd Gardner (Wiley, 1966) relates the parameter $\theta_e/(\Delta f/fm)$ to the damping factor ζ. These numbers may not be particularly significant for the loop with delay, but assuming they are, we calculate the equivalent noise bandwidth using

$$B_L = \omega_n \left(\zeta + \frac{1}{4\zeta}\right)/2.$$

Table 3.3-T1 lists the results.

3.4 Acquisition Performance

Due to the loop frequency drift described in Section 2.6, acquisition usually occurs immediately or on loop reset (every 2^{16} bit times), or not at all. When lock occurs, the lock detector will respond within a few time constants ($\tau = 2^{16}$), or not at all. Therefore, we find that acquisition time is not usually a problem, but that measuring the range of conditions for which lock occurs is more useful.

The lock detector threshold is adjusted to detect lock at least down to -3 dB at all rates. Under some conditions the unit acquires and indicates lock at -5 dB. If it is adjusted for greater sensitivity, the analog imbalances cause a false lock condition at some combinations of noise level and rate setting.

Table 3.3-T1 Measured Loop Characteristics
at R_D = .8 Mbps, E_b/N_o = +3.7 dB

Nominal B_L/R_D	Δf	f_m	θ_e	$\frac{\theta_e}{\Delta f/f_m}$	ζ	ω_n/R_D	Calculated B_L/R_D
0.8%/1.6%	1700	1700	1.0	1.0	0.5	1.3%	0.65%
0.1%	550	200	1.0	0.36	1.5	0.15%	0.12%

Table 3.4-T1 Acquisition Conditions

B_L	E_b/N_o	Allowable Offset
.1%	+3 dB	$\pm$ .06%
.1%	-2.6 dB	$\pm$ .014%
.025%	+3 dB	$\pm$ .016%
.025%	-2.6 dB	$\pm$ .010%

When the data rate is offset from the setting, the unit will sometimes not acquire. The maximum offsets for which acquisition is reliable (with either sign of offset) are listed in Table 3.4-T1 for two bandwidths and two SNR's. All measurements are at the lowest specified signal level (.25 VPP) and the lowest rate (250,000 bps).

3.5 Baseline Variation

Tests at two rates confirm that the degradation in performance due to a baseline variation of sine frequency 1% R and peak amplitude equal to the signal peak-to-peak amplitude is an increase in bit error rate of less than 4%. The equivalent SNR degradation is completely negligible. The bit synchronizer continues to function properly with a baseline variation of twice that amplitude and frequency - the resulting degradation was not measured.

4.0 RECOMMENDATIONS

In this section we discuss various design alternatives that will provide either improved performance or simpler, more economical construction in some applications. Each item should be evaluated in terms of a revised specification of the bit synchronizer requirements, and in terms of the cost of the development work required.

4.1 Mid-Phase Lock Detector Scheme

The lock detector scheme employed compares the energy in the two phases of mid-phase integrals. As discussed above, the two phases are processed through physically separate integrate-and-dump and analog-to-digital conversion circuits, which have to be carefully matched in gain and offset at all rates. If a single circuit, integrating over slightly less than a half bit time (dumping quickly) could handle both channels, there would be no mismatch other than timing (integration duration).

The in-phase lock detector schemes proposed in TR-8055A would give higher performance, but the combination of in-phase half-bit integrals by either digital or analog methods would require the development of new circuitry, with careful attention to balancing.

4.2 Timing Accuracy Improvement

The $\pm$ 1% timing jitter due to the NCO is not a problem, but static offsets and duty cycle variations caused by the threshold and propagation characteristics of TTL circuits and the clock comparators cause problems in the data detectors. Duty cycle variation causes a rate dependent mismatch in the mid-phase channels. Static offset causes non-optimal data timing, also rate dependent.

The use of ECL circuits in the NCO DAC and CLOCK logic would reduce these problems and give more accurate output clocks.

To assure accurate 50% duty cycle, a revised approach is needed, which divides down a double- or quadruple-rate oscillator to derive the clocks. For example, a VCO with a tuning range of 12.8 to 25.6 MHz could be divided down by powers of two and phase-locked to the NCO sine wave output across each rate range. This approach retains the crystal frequency stability of the NCO while improving the short term jitter stability to that of the VCO. The NCO is simplified to have a single output, less critical filter, and optionally a lower operation rate for less critical circuitry. The extra circuitry required is the VCO, its analog PLL circuit (probably a single IC), and the divide by 2^n logic.

4.3 Band Limiting Improvement

The BiØ-L signaling format has a considerably wider spectral distribution than the more common NRZ-L format. Band limiting to the main lobe on NRZ-L data (BW = R) causes only about 0.2 dB degradation, while limiting to the main lobe of BiØ-L data (BW = 2R) causes about 0.6 dB degradation. If a wider input filter is used to reduce the degradation, a wider noise dynamic range is needed in the I&D circuits, and a separate AGC pre-detection filter is needed. Increasing the bandwidth to 4R will decrease the degradation to less than 0.4 dB. Not much can be gained at the highest rate of 5 Mbps, since the specified IF bandwidth is only 15 MHz (lowpass equivalent).

Another approach worth considering is to use a wide dynamic range mixer to remove the biphase coding immediately at the input, before band limiting. Then a single lowpass with BW = R will give little degradation and a very good SNR into the AGC detector. A separate path would then be needed for the mid-phase (transition detection) channel, but it could be separately optimized, without affecting the data detection or the AGC.

4.4 Loop Configuration

The loop model originally proposed had a delay z^{-1} (i.e. proportional to the data period) in the NCO, resulting in percentage bandwidth and damping being independent of rate. To restore this desirable property, the NCO should be clocked at a rate proportional to the data rate, though not necessarily synchronized with the data. A second NCO module could be used to generate a clock proportional to the nominal rate, though a variable bandwidth reconstruction filter would then be needed after the DAC.

An even better approach is to digitally add the two NCO phase outputs and convert to a sine wave at the crystal clock rate, while clocking the synchronizing loop NCO at the recovered bit rate. See Figure 4.4-1. If various bandwidths are required, some kind of gain change is still needed in the loop filter.

4.5 Fixed-Rate Simplifications

A bit synchronizer for fixed applications, such as the Space Shuttle, could be designed to operate at any preselected rate in a wide range by choice of plug-in oscillators and filters. The entire rate select related circuitry could then be eliminated.

4.6 Fast Loop Filter

The loop filter was implemented at a low sampling rate to allow sequential gain shifting and to conserve power. Because of the delay, this approach is not extendable to high bandwidths. If the need for rate range shifting is eliminated by one of the methods discussed above, and the bandwidth selection is limited to one or two settings, the need for complicated range shift circuitry is eliminated. In many applications the

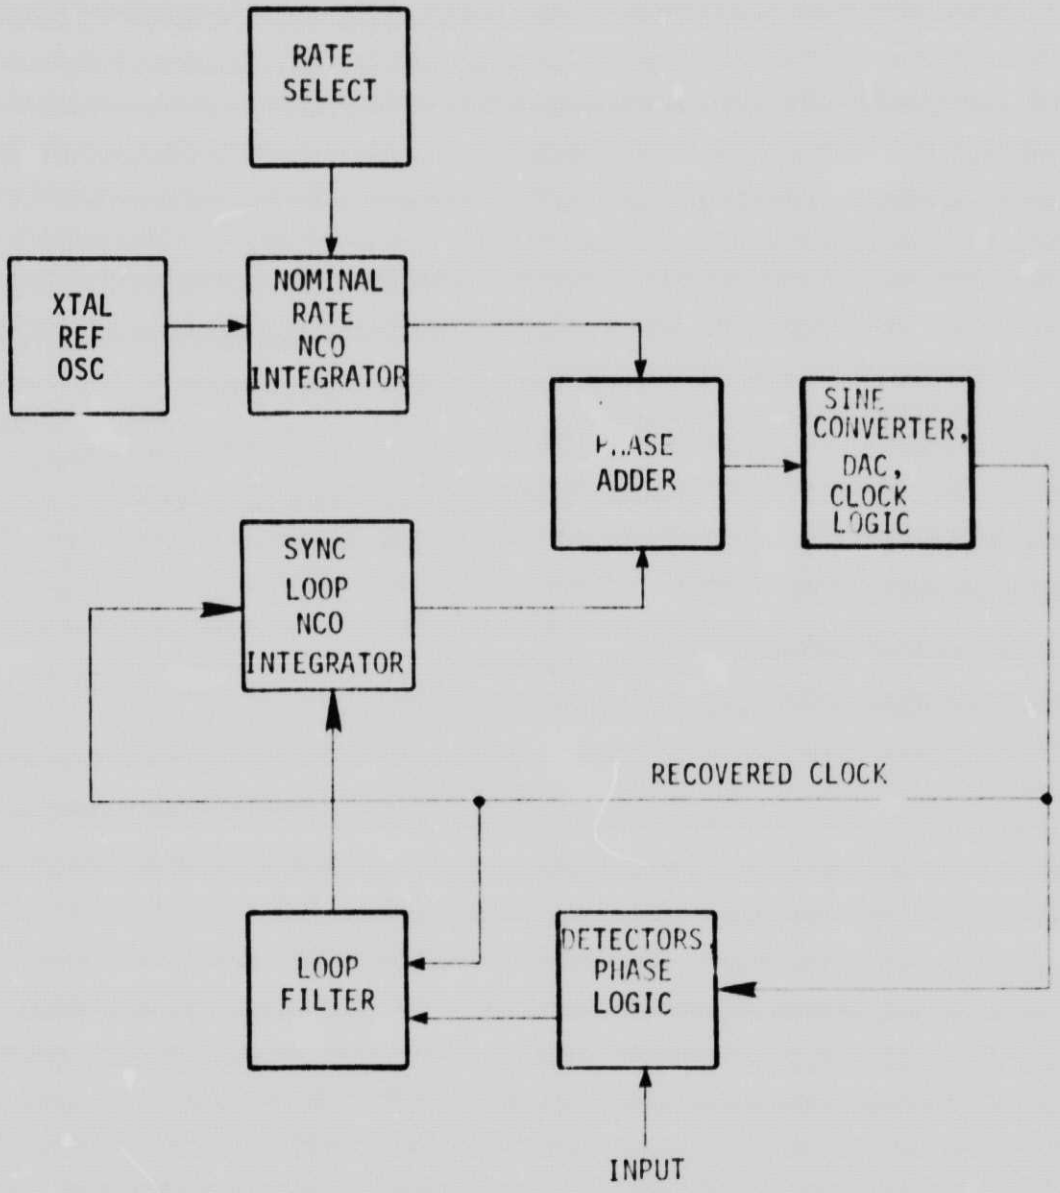

FIG.4.4-1 IMPROVED LOOP CONFIGURATION

extra power used by a TTL loop filter is no problem. Thus a straightforward integrator and adder in TTL could be implemented, giving a total loop delay of about 6T including the phase estimation logic. This would allow a loop bandwidth of about 4% with no problem. In some systems this would be useful for a fast acquisition mode.

4.7 Sweep Circuitry

For systems with typical rate offsets larger than half the desired loop bandwidth, a search mode facilitates acquisition. STI has implemented sweep modes in digital bit synchronizers previously, when the rate uncertainty was many times the bandwidth. If the specifications of .05% offset and .025% bandwidth are realistic, then either a sweep circuit or a wide bandwidth acquisition mode should be used in a later version of the bit synchronizer.

4.8 Acquisition and Tracking Modes

Since an any-rate bit synchronizer is typically used in the laboratory as an experimental tool, it would be useful to include the capability to acquire, resolve phase ambiguity, and track with other types of biphase and NRZ data formats. This would require changing the timing and lock detection and phase algorithms according to the selected data format. This would require a considerable development effort.