NATIONAL RADIO ASTRONOMY OBSERVATORY
Green Bank, West Virginia

SPECTRAL PROCESSOR MEMO NO. 25

MEMORANDUM February 13, 1985

To: Spectral Processor Group

From: R. Fisher

Subj: Some Experiments with an Integer FFT

A number of experiments with the computer simulation of the spectral processor multiplier and accumulator are described in this memo. These experiments were performed to develop a feel for the characteristics of the spectrometer and to try to uncover any peculiarities that might be corrected in the early design stages. If you have any comments on my interpretations of the results or if you have any suggestions for other experiments please let me know.

In all of the following experiments the input levels are referred to the A/D resolution of 1 level = 1/63rd of its total range (6-bit word) from +31 to -32. Unless otherwise noted, the A/D is assumed to produce an output of zero with an input between ± 0.5 levels.

Output power word sizes.

The A/D and FFT multipliers are assumed to use 16-bit, fractional two's complement arithmetic with the ten least significant bits from the A/D set to zero. The fractional arithmetic and one-bit shifts associated with every add operation prevents word overflow anywhere in the FFT from the strongest CW signal that can be accommodated by the A/D converter. A consequence of this is that the average level of white noise in the FFT is half of a bit ($\sqrt{2}$) lower in each successive FFT stage, and because of squaring to obtain the power spectrum each additional FFT stage reduces the average noise level in the accumulators by one bit.

With an input noise rms of 1 level the 1024-channel FFT produced an output power level of about 576 (assuming the binary decimal is at the least significant end) or an average word width of about 9 bits. (The maximum possible output word width is 30 bits.) Dropping 3 FFT stages to produce 128 channels per input will produce an average power word width of 12 bits. The distribution of power products for one FFT output in the 1024 channels is shown in Table 1.

TABLE 1

Single FFT output power statistics, 1024 channels, input noise rms = 1 level.

Power Level Range	Number of Channels	Cumulative Number of Channels
0- 31	60	1024
32- 63	50	964
64- 127	113	914
128- 255	169	801
256- 511	234	632
512-1023	237	398
1024-2047	126	161
2048-4095	34	35
4096-8191	1	1

Effects of rounding power words.

To minimize accumulator memory space the output power word size must be reduced by dropping some of the least significant bits. Even if these words are rounded instead of truncated some bias is introduced by word size reduction because of the asymmetric distribution of values around the mean as shown in Table 1. This distribution should be independent of input amplitude and number of channels. Table 2 shows the results of the average of 256, 64 channel FFT operations on the same gaussian random noise with different numbers of bits rounded in the power words. The average word widths under these conditions was about 13 bits. Two effects of importance to the spectrometer are evident in Table 2.

TABLE 2

Effects of rounding output power word.

Number of Bits Dropped	Average Level	Level Reduction	Measured rms of Spectrum	Rms/Avg Degradation
0	8824.4	----	470.4	----
10 (÷ 1024)	8821.6	0.03%	469.3	1.00
11 (÷ 2048)	8811.5	0.15%	470.4	1.00
12 (÷ 4096)	8762.8	0.70%	486.8	1.04
13 (÷ 8192)	8546.0	3.15%	507.3	1.11
14 (÷16384)	7670.0	13.08%	569.2	1.39
15 (÷32768)	5274.0	40.23%	705.9	2.51

The first is the obvious one of increased noise due to loss of low order bits. The second is the reduction in average level due to rounding an asymmetrically distributed set of numbers. This will cause a distortion of a curved noise spectrum because each channel will have a different average level, and rounding will reduce the level of channels with low output more than those with high output. This distortion is predictable and could be corrected using only the information in the spectrum itself assuming that the noise power is constant during the integration.

Rounding of the power word should be done at least two bits below the average. If we assume that the lowest input noise rms = 1 A/D level under normal operation, a 128 channel spectrometer could drop 10 bits and a 1024 channel configuration could drop 7 bits. This would allow some operation of about half of this input level (-6 dB) if reduced sensitivity and spectrum distortion are less important than dynamic range, and operation at higher input levels would provide a reasonable margin above the rounding limit to produce a stable spectral shape. For example, if the input rms = 5 A/D levels the rounding would be about 6 1/2 bits below the average, and the amplitude distortion would be about 0.001% or about 0.6 mK for a 50 K system temperature which would be reached in 50 hours of integration on a 1024 channel, 40 MHz spectrum.

In memo number 24 I suggested a 24-bit accumulator word width, but the assumption of the lowest order bit just covering the average power level with 1 level input rms has now been changed by 2 bits. The next increment in word width is 28 bits so the extra two bits could be used to cover most of the added range needed by the cases of fewer FFT stages. I now favor a 28-bit accumulator word width.

Sensitivity loss due to quantization of input waveform.

The loss in sensitivity due to quantization of the input waveform has been well documented by others, but out of curiosity I ran a few tests of this with the spectrometer simulation and the results are shown in Table 3. The ratio of output noise power to the standard deviation of this power remains constant with variation of the input quantization, but the relative amplitude of a CW signal decreases with increasing quantization so the signal to noise ratio of the CW signal was measured.

Two cases are shown in Table 3. One uses the normal ± 0.5 level A/D quantization and the other changes the A/D output from zero to one at zero volts input (zero-slicing). The sine wave peak amplitude was approximately equal to the noise rms. At low input levels the zero-slicing A/D has an advantage over the ± 0.5 level case because the state changes still occur at very low signal levels in the former. At normal operating levels the difference between the two cases is very small. The ± 0.5 slicing case has the advantage of a much smaller DC channel amplitude that would be easier to handle in the accumulator. Note that the relative signal to noise ratios in Table 3 are themselves subject to some random error even though the input signal is identical in all cases.

TABLE 3

CW signal-to-noise ratios vs. quantization level.

Input noise rms in A/D level		Relative CW signal to noise: ± 0.5 level slicing	Relative CW signal to noise: Zero slicing
4	.....	1.0 (ref)	1.0 (ref)
1	.....	0.95	0.93
0.5	.....	0.68	0.76
0.4	.....	0.62	0.68
0.3	.....	0.46	0.58
0.2	.....	0.24	0.57

CW harmonics due to input quantization.

A symmetric A/D converter will produce odd harmonics when sampling a pure CW signal due to quantization errors, and these harmonics or their aliases will appear in the spectrum of the A/D output. The addition of noise to the CW signal will tend to destroy the coherence of these harmonics, and to determine the harmonic amplitude reduction that can be expected from noise a number of accumulated spectra were computed with different CW signal and noise levels. The results are shown in Table 4.

TABLE 4

CW harmonics as a function of quantization and added noise.

Harmonic Number	Relative harmonic strength: CW rms in A/D levels (no noise)			Relative harmonic strength: Noise rms in A/D levels (CW rms = 0.7	
	0.7 dB	1.1 dB	3.0 dB	0.2 dB	0.5 dB
Fundamental	0	0	0	0	0
3	-16	-21	-24	-27	<-40
5	-25	-17	-23	<-40	<-40
7	-18	-20	-28	-24	<-40
9	-18			-27	<-40
11	-24			-40	<-40
13	-47			<-40	<-40
15	-26			<-40	<-40
17	-24			<-40	<-40

The measurement of harmonic strengths weaker than -40 dB would have taken too much computer time, but it seems safe to assume that with noise rms values greater than 1 A/D level the harmonics will be much weaker than -50 dB.

Noise amplitude distortion in the presence of a strong CW signal.

With a one bit A/D converter a CW signal with a strength comparable to or stronger than the noise power will severely distort the sample statistics and reduce the measured spectral noise power in proportion to the CW strength. A similar but smaller effect must happen with a multibit A/D so a few relatively extreme tests were performed with noise in the presence of a CW signal with the spectrometer simulation.

The results are shown in Table 5. With noise rms values greater than 1 A/D level there appears to be very little effect on the measured noise power from strong CW signals.

TABLE 5

Noise spectrum distortion from CW signals.

Noise rms in A/D levels	CW rms in A/D levels	Noise power suppression
Zero slicing		
0.3	4.0	15%
1.0	4.0	<0.02%
± 0.5 slicing		
0.3	4.0	37%
1.0	4.0	<0.01%

Overflow in the FFT arithmetic.

I said in the first section that the FFT arithmetic is such that overflows are avoided. This is true for strong coherent signals but not strictly true for strong noise. In the third and following FFT stages there are instances where the sine and cosine components of the rotation coefficients are both near 0.7, and if both real and imaginary components of the B inputs of such a butterfly are near one the sum of the sine and cosine products will be greater than one; hence, an overflow will occur. The probability of overflow depends on the number of large random noise data values.

Table 6 shows the fraction of overflows per butterfly operation in the third FFT stage as a function of input noise level. Because of the decrease in average noise levels in later stages no overflows were seen in any stage but the third. The noise levels shown in Table 6 are much higher than would be used in normal operation, but this simulation could not process enough data to measure very small overflow probabilities so we have to rely on an extrapolation of the high level results.

The third column in Table 6 gives the fraction of data points which are equal to the maximum possible value from the A/D converter. This fraction should be a good predictor of the overflow probability, and, in fact, the number of overflows measured is proportional to the fourth power of the fraction of maximum A/D values as would be expected from the fact that four large input data points in the right places are required to cause an overflow.

TABLE 6

Overflow probability vs. input noise level.

Noise rms in A/D levels	No. of overflows per butterfly operation	Fraction of maximum A/D levels
3000	4.3 (± .7) x 10^{-3}	0.99
600	3.1 (± .6) x 10^{-3}	0.96
300	2.1 (± .5) x 10^{-3}	0.91
200	1.7 (± .5) x 10^{-3}	0.87
100	1.0 (± .3) x 10^{-3}	0.75
50	0.20 (± .05) x 10^{-3}	0.52
25	<0.02 x 10^{-3}	0.20

Using the fourth-power law derived from the data in Table 6 the probability of having an overflow in one second with the spectrometer accepting a 40 MHz bandwidth and a noise rms of 10 A/D levels is 10^{-6}. With a noise rms of 15 levels the probability is 0.13 under the same conditions. Even the 10 level rms is higher than we would normally put into the spectrometer so the overflow probability is acceptably small without reducing the noise level in the third stage.

To see what effect adding a CW signal to the input noise would have on the overflow statistics a 40-level peak-to-peak sine wave was added to the 50-level rms noise case. The result was a 50% decrease in the overflow probability because the noise statistics were biased against having all of the required high level input conditions at the susceptible butterflies to cause an overflow.

Spectral distortion due to A/D quantization.

The effects of severe quantization in one and two bit auto-correlation spectrometers have been studied in considerable detail by others. Since the autocorrelation and Fourier transform spectrometers are equivalent in this respect we can use the work connected to autocorrelators as a guide in predicting the errors that can be expected in an FFT device. My thanks go to John Granlund for some very helpful discussions in this area.

Quantization error can be looked at in two ways. It reduces the relative intensity of low level correlations with respect to perfect correlation at zero delay in the autocorrelation function of an RF signal. This can be seen in Figure 1 which is a plot of correlation values for quantized vs. unquantized signals. Each curve represents a different degree of quantization. Viewed in the spectral domain the quantization nonlinearity generates harmonics and mixing products of all frequency components in the input spectrum. These products or their aliases are laid over the true spectrum. If the unquantized signal is noise with a reasonably flat spectrum within the sampled passband the error spectrum will be very nearly white and will appear mostly as a DC offset in the true spectrum. In the autocorrelation function this means that the slope at low correlations in Figure 1 is not unity and, hence, does not pass through the zero delay correlation point.

The degree to which the error spectrum is not perfectly white depends on the complexity of the true input spectrum and the size of the quantization errors. In autocorrelation terms a complex spectrum will have large correlations at delays beyond zero, and large quantization errors will produce more curvature in the true vs. quantized correlation curve in Figure 1.

Looking for quantization distortions in a simulated noise spectrum takes an enormous amount o computer time to reduce random errors to levels to be expected in the operation of the spectral processor. Autocorrelation functions can be computed much faster than a spectrum can be transformed so the approach taken has been to compute true vs. quantized correlation curves for several degrees of quantization using computer generated noise. These curves give the slope and amount of nonlinearity at low correlation levels. The autocorrelation functions have been calculated for two input spectra, and these functions have then been distorted by the correlation error curves and transformed back into the frequency domain to look for spectral distortions.

Two rather severe passband shapes shown in Figure 2 have been used in these calculations. The top curve has a maximum autocorrelation value of 13% relative to the zero delay value, and the largest value for the bottom curve is 21%. Except for the case of a very strong CW signal in the passband it is unlikely that the autocorrelation values in practice will be greater than the ones used in these tests.

The slopes of the low end of the curves in Figure 1 are shown in Figure 3 where the slope is plotted as a function of noise amplitude relative to the quantization interval. Even when the noise rms is five quantization levels the slope is sufficiently different from unity to require a correction in the zero level and gain of the computed spectrum.

Upper limits on the curvatures of the lines in Figure 1 were estimated by comparing the fits of linear and quadratic curves on the autocorrelation values from 0.0 to 0.5 and higher. For noise rms levels greater than 1.0 no significant difference between the straight line and quadratic fits could be seen up to an autocorrelation level of about 0.7. For these noise levels a conservative upper limit on the quadratic coefficient is 0.001. For noise rms levels of 0.8 and 0.5 the upper limits on the quadratic coefficient are 0.005 and 0.01, respectively. These limits are set by the noise in the calculations of the autocorrelation coefficients, and the linearity of the curves may be considerably better than this.

The autocorrelation functions derived from the two bandpasses in Figure 2 were modified by a transfer function that contained only a quadratic term ($a_3\ x^2$), and with $a_3 = 0.001$ the peak fractional error in the lower bandpass curve was about 10^{-4} except near the left edge of the bandpass where the error rose to 2×10^{-4}. The error in the top bandpass was 5×10^{-5} rising to 10^{-4} near the left edge using the same quadratic coefficient. An error of 10^{-4} is equivalent to the rms noise in a 1024-channel, 40 MHz spectrum after 40 minutes of integration. This size of error would be detectable, but if the input power to the A/D converter does not change very much between signal and reference the error will cancel in the difference. Also, there is a good chance that the autocorrelation transfer function is even more linear than supposed here so the spectrum distortion may be quite a bit less than 10^{-4}.

Zero offset and gain corrections will certainly need to be applied to the FFT spectrum. If a higher order correction for quantization effects is necessary the averaged spectra could be Fourier transformed to an autocorrelation function in the spectral processor computer, corrected, and transformed back into the frequency domain. Both the zero and higher order corrections depend on the input level, and, in the absence of an exact set of equations for the corrections, the autocorrelation transfer function would have to be calibrated with long integrations with the spectral processor hardware under different quantization conditions. At least a few of these calibrations should be done whether the second order corrections are to be performed or not.

JRF/cjd

Attachments
  Figures 1, 2, 3

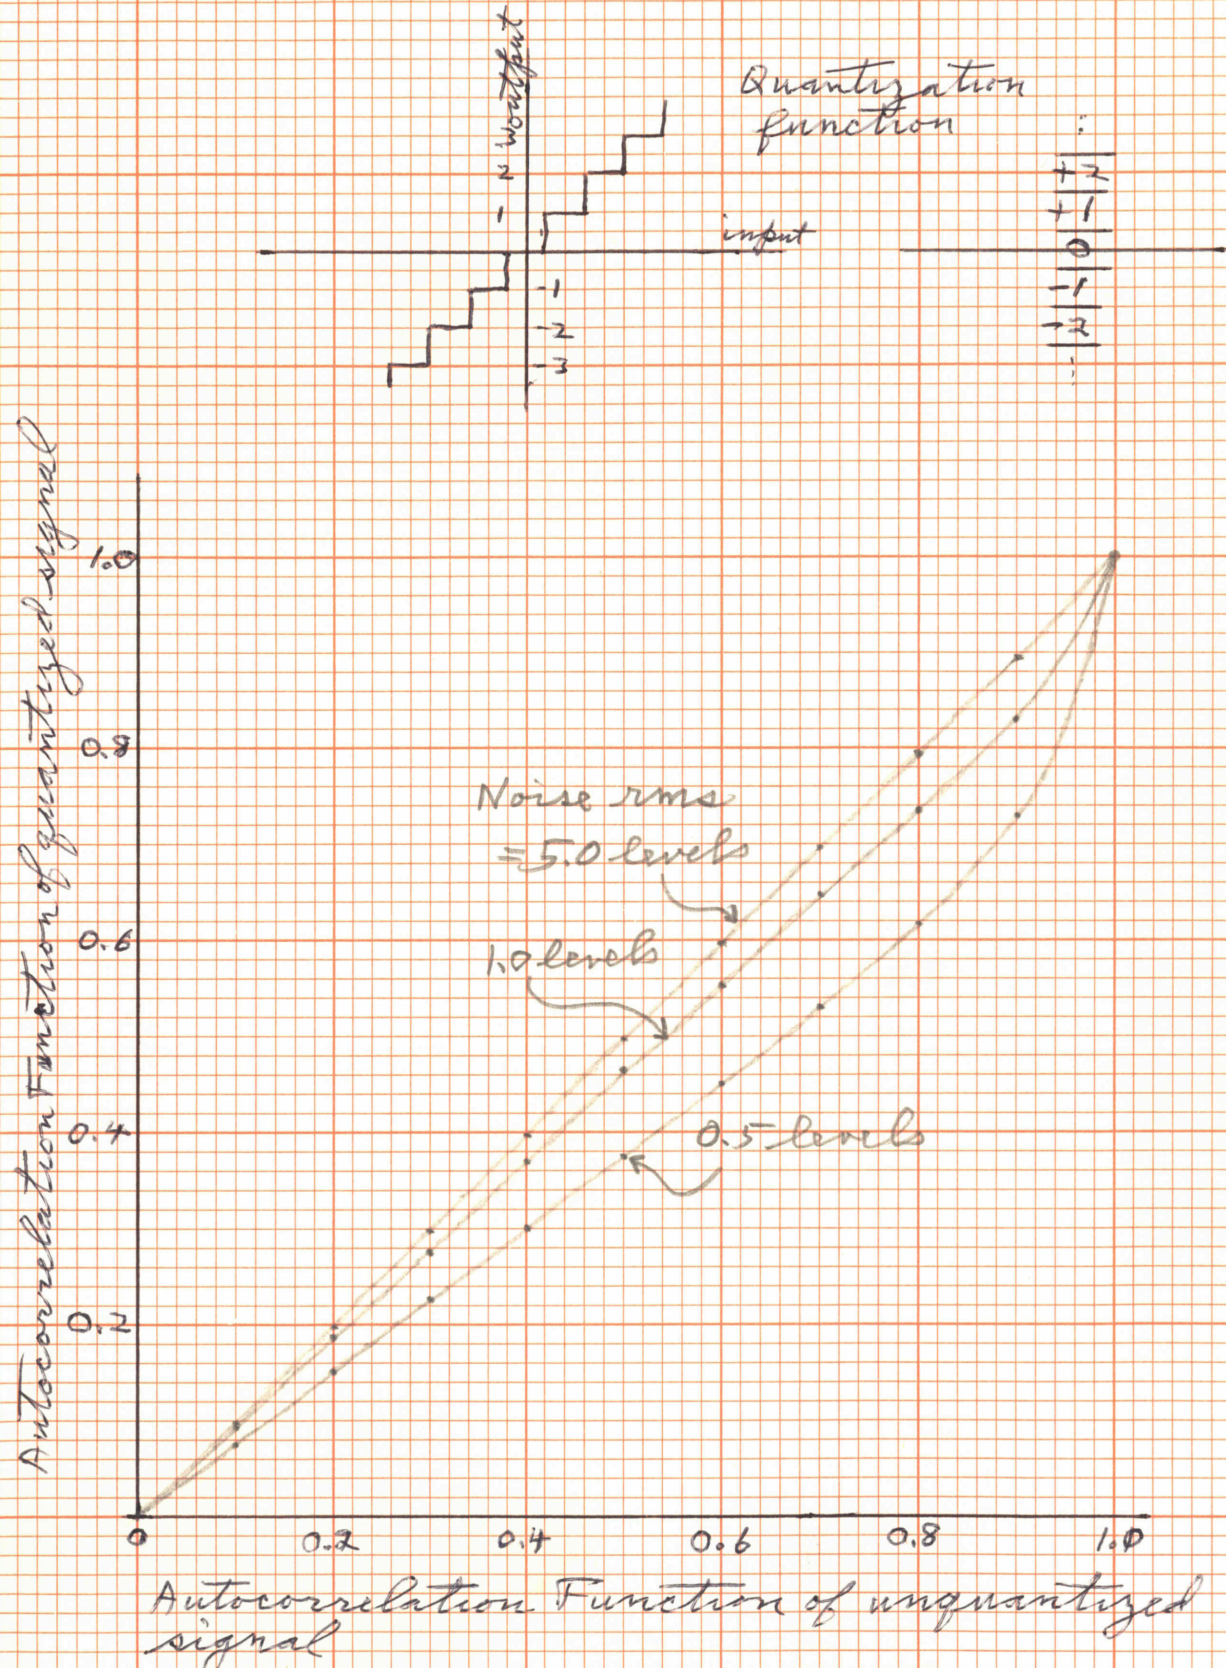

Quantization function
output
input
1
2
-1
-2
-3
+2
+1
0
-1
-2
Autocorrelation Function of quantized signal
1.0
0.8
0.6
0.4
0.2
Noise rms = 5.0 levels
1.0 levels
0.5 levels
0
0.2
0.4
0.6
0.8
1.0
Autocorrelation Function of unquantized signal

Figure 1

Figure 2

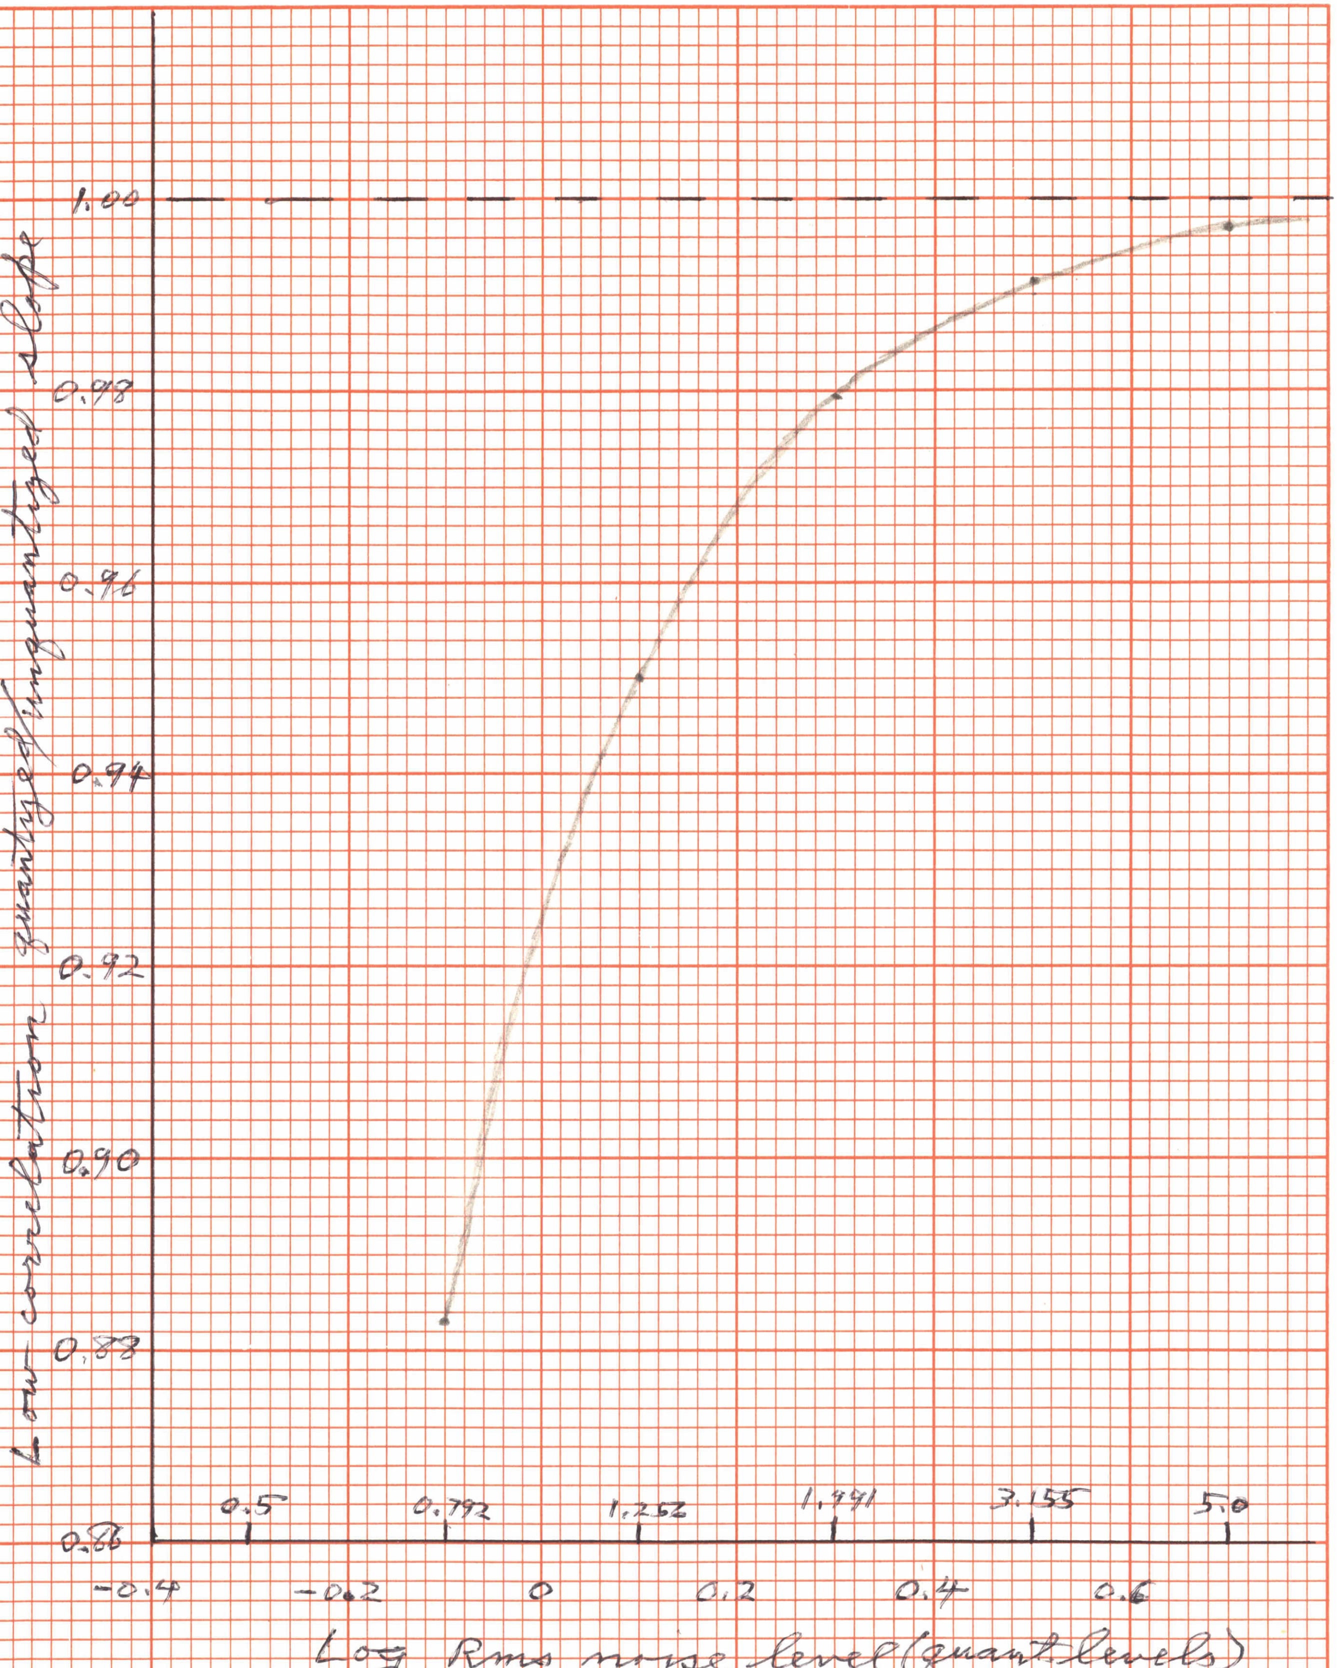

Low correlation quantized/unquantized slope
1.00
0.98
0.96
0.94
0.92
0.90
0.88
0.86
0.5
0.792
1.262
1.991
3.155
5.0
-0.4
-0.2
0
0.2
0.4
0.6
Log Rms noise level (quant. levels)

Figure 3