ILLUSTRATIVE PROBLEM

Illustrative Problem 4.4 [Determining the centroids] Determine the centroids of the quantization regions for a zero-mean, unit variance Gaussian distribution, where the boundaries of the quantization regions are given by $(-5, -4, -2, 0, 1, 3, 5)$.

SOLUTION

The Gaussian distribution is given in the m-file normal.m. This distribution is a function of two parameters, the mean and the variance, denoted by m and s (or σ), respectively. The support of the Gaussian distribution is $(-\infty, \infty)$, but for employing the numerical routines it is enough to use a range that is many times the standard deviation of the distribution. For example, $(m - 10\sqrt{s}, m + 10\sqrt{s})$, can be used. The following m-file determines the centroids (optimal quantization levels).

This results in the following quantization levels: $(-5.1865, -4.2168, -2.3706, 0.7228, -0.4599, 1.5101, 3.2827, 5.1865)$.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.5 In Illustrative problem 4.4 determine the mean-square error.

SOLUTION

Letting $a = (-10, -5, -4, -2, 0, 1, 3, 5, 10)$ we obtain a mean square error of 0.177.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.6 [Uniform quantizer distortion] Determine the mean-square error for a uniform quantizer with 12 quantization levels, each of length 1, designed for a zero-mean Gaussian source with variance of 4. It is assumed that the quantization regions are symmetric with respect to the mean of the distribution.

SOLUTION

By the symmetry assumption the boundaries of the quantization regions are $0, \pm 1, \pm 2, \pm 3, \pm 4, \pm 5$, and the quantization regions are $(-\infty, -5]$, $(-5, -4]$, $(-4, -3]$, $(-3, -2]$, $(-2, -1]$, $(-1, 0]$, $(0, 1]$, $(1, 2]$, $(2, 3]$, $(3, 4]$, $(4, 5]$, and $(5, +\infty)$. This means that in the uq_dist.m function we can substitute $b = -20$, $c = 20$, $\Delta = 1$, $n = 12$, $s = -5$, tol = 0.001, $p_1 = 0$, and $p_2 = 2$. Substituting these values into uq_dist.m we obtain a squared error distortion of 0.0851 and quantization values of ± 0.4897, ± 1.4691, ± 2.4487, ± 3.4286, ± 4.4089, and ± 5.6455.

The m-file uq_mdpnt.m determines the squared error distortion for a symmetric density function when the quantization levels are chosen to be the midpoints of the quantization intervals. In this case the quantization levels corresponding to the first and the last quantization regions are chosen to be at distance $\Delta/2$ from the two outermost quantization boundaries. This means that if the number of quantization levels is even, then the quantization boundaries are $0, \pm\Delta, \pm 2\Delta, \ldots, \pm(N/2 - 1)\Delta$ and the quantization levels are given by $\pm\Delta/2, \pm 3\Delta/2, \ldots, (N-1)\Delta/2$. If the number of quantization levels is odd, then the boundaries are given by $\pm\Delta/2, \pm 3\Delta/2, \ldots, \pm(N/2 - 1)\Delta$ and the quantization levels are given by $0, \pm\Delta, \pm 2\Delta, \ldots, (N-1)\Delta/2$.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.7 [Uniform quantizer with levels set to the midpoints] Find the distortions when a uniform quantizer is used to quantize a zero-mean, unit variance Gaussian random variable. The number of quantization levels is 11, and the length of each quantization region is 1.

Note
Use file: uq_mdpnt.m

SOLUTION

$p_1 = 0$ and density function name, = 'normal' $p_2 = 1$ for the density function parameters, $n = 11$ for the number of quantization levels, and $\Delta = 1$ for the length of the quantization levels. For the parameter b, which chooses the support set of the density function, we use the value $b = 10p_2 = 10$, and we choose the tolerance to be 0.001. The resulting distortion is 0.0833.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.8 [Lloyd-Max quantizer design] Design a 10-level Lloyd-Max quantizer for a zero-mean, unit variance Gaussian source.

SOLUTION

Using $b = 10$, $n = 10$, tol = 0.01, $p_1 = 0$, and $p_2 = 1$ we obtain the quantization boundaries and quantization levels vectors $\boldsymbol{a}$ and y as

Note
Use file: lloydmax.m

$$\boldsymbol{a} = \pm 10, \pm 2.16, \pm 1.51, \pm 0.98, \pm 0.48, 0$$
$$y = \pm 2.52, \pm 1.78, \pm 1.22, \pm 0.72, \pm 0.24$$

and the resulting distortion is 0.02. These values are good approximations to the optimal values given in the table by Max [2].

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.9 [Uniform PCM] Generate a sinusoidal signal with amplitude 1, and $\omega = 1$. Using a uniform PCM scheme, quantize it once to 8 levels and once to 16 levels. Plot the original signal and the quantized signals on the same axis. Compare the resulting SONRs in the two cases.

Note
Use file: u_pcm.m

SOLUTION

We arbitrarily choose the duration of the signal to be 10 s. Then we generate the quantized signals for the two cases of 8 and 16 quantization levels.

The resulting SQNRs are 18.90 dB for the 8-level PCM and 25.13 dB for 16-level uniform PCM. The plots are shown in Figure 4.5.

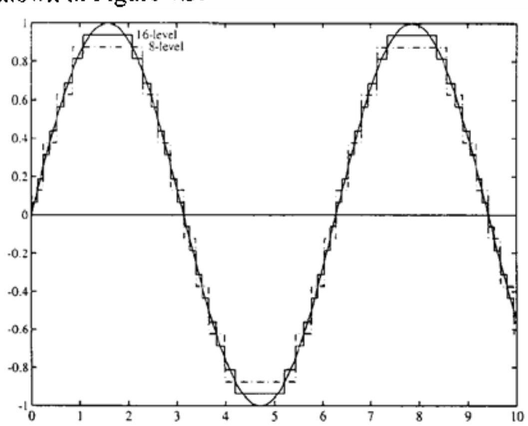

Figure 4.5: Uniform PCM for a sinusoidal signal using 8 and 16 levels.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.10 [Uniform PCM] Generate a sequence of length 500 of zero-mean, unit variance Gaussian random variables. find the resulting SQNR when the number of quantization levels is 64. Find the first five values of the sequence, the corresponding quantized values, and the corresponding codewords.

Note
Use file: u_pcm.m to generate SQNR
Solution in lp_04_10.m

SOLUTION

$$\text{SQNR} = 31.66 \quad \text{dB}$$

$$\text{Input} = [0.1775, -0.4540, 1.0683, -2.2541, 0.5376]$$

$$\text{Quantized values} = [0.1569, -0.4708, 1.0985, -2.2494, 0.5754]$$

$$\text{Codewords} = \left\{\begin{array}{cccccc} 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 \end{array}\right.$$

Note that different runnings of the program result in different values for the input, the quantized values, and the codewords. However the resulting SQNRs are very close.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.11 In Illustrative Problem 4.10, plot the quantization error, defined as the difference between the input value and the quantized value. Also, plot the quantized value as a function of the input value.

SOLUTION

The two desired plots are shown in Figure 4.6.

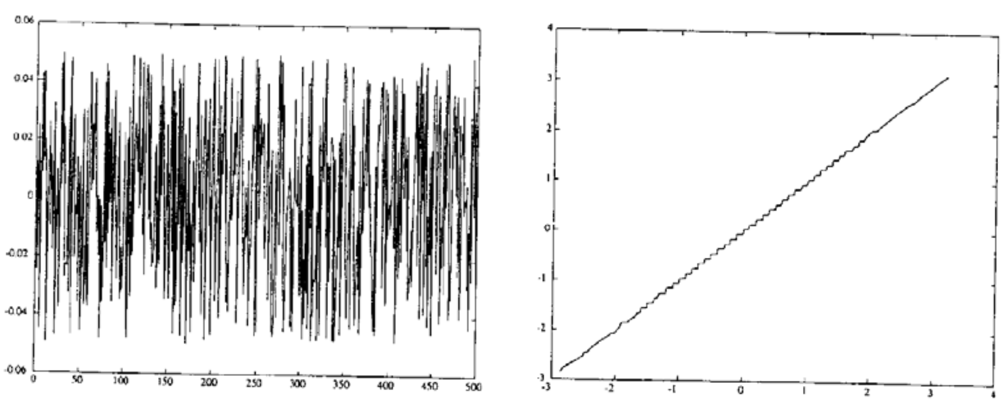

Figure 4.6: Quantization error in uniform PCM for 64 quantization levels.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.12 Repeat Illustrative Problem 4.11 with number of quantization levels set once to 16 and set once to 128. Compare the results.

SOLUTION

The result for 16 quantization levels is shown in Figure 4.7, and the result for 128 quantization levels is shown in Figure 4.8.

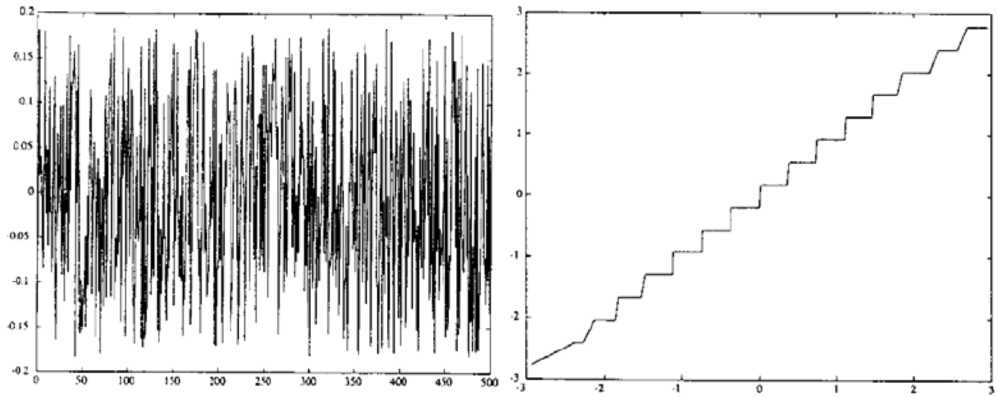

Figure 4.7: Quantization error for 16 quantization levels.

Comparing Figures 4.6, 4.7, and 4.8, it is obvious that the larger the number of quantization levels, the smaller the quantization error, as expected. Also note that for a large number of quantization levels, the relation between the input and the quantized values tends to a line with slope 1 passing through the origin; i.e., the input and the quantized values become almost equal. For a small number of quantization levels (16 for instance), this relation is far from equality, as shown in Figure 4.7.

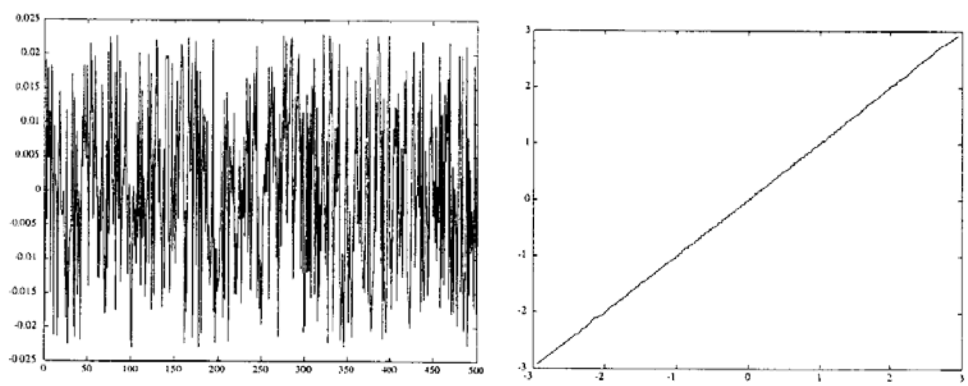

Figure 4.8: Quantization error for 128 quantization levels.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.13 [Nonuniform PCM] Generate a sequence of random variables of length 500 according to an $\mathcal{N}(0, 1)$ distribution. Using 16, 64, and 128 quantization levels and a μ-law nonlinearity with $\mu = 255$, plot the error and the input-output relation for the quantizer in each case. Also determine the SQNR in each case.

SOLUTION

Let the vector a be the vector of length 500 generated according to $\mathcal{N}(0, 1)$, i.e., let

Note: Use a=randn(500)
then the use file mula_pcm.m as under
[dist, a_quan,code]=mula_pcm(a,16,255)

we can obtain the quantized sequence and the SQNR for a 16-level quantization. The SQNR will be 13.76 dB. For the case of 64 levels we obtain SQNR = 25.89 dB, and for 128 levels we have SQNR = 31.76 dB. Comparing these results with the uniform PCM, we observe that in all cases the performance is inferior to the uniform PCM. Plots of the input-output relation for the quantizer and the quantization error are given in Figures 4.10, 4.11, and 4.12.

Comparing the input-output relation for the uniform and the nonuniform PCM shown in Figures 4.7 and 4.10 clearly shows why the former is called uniform PCM and the latter is called nonuniform PCM.

From the above example we see that the performance of the nonuniform PCM, in this case, is not as good as uniform PCM. The reason is that in the above example the dynamic range of the input signal is not very large. The next example examines the case where the performance of the nonuniform PCM is superior to the performance of the uniform PCM.

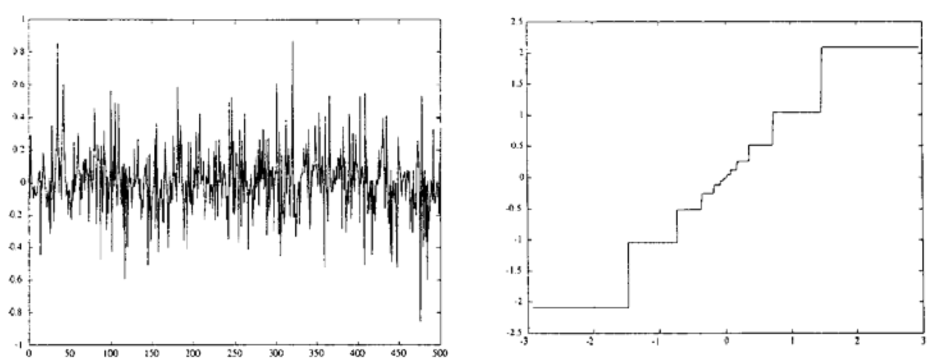

Figure 4.10: Quantization error and input-output relation for a 16-level μ-law PCM.

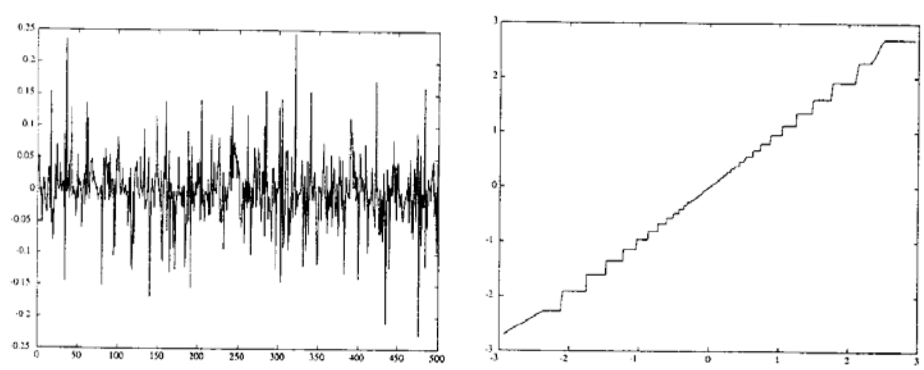

Figure 4.11: Quantization error and quantizer input-output relation for a 64-level μ-law PCM.

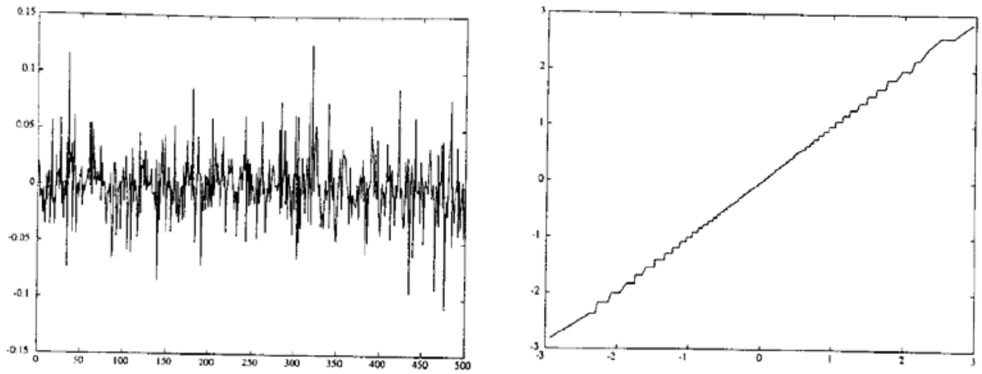

Figure 4.12: Quantization error and quantizer input-output relation for a 128-level μ-law PCM.

ILLUSTRATIVE PROBLEM

Illustrative Problem 4.14 The nonstationary sequence **a** of length 500 consists of two parts. The first 20 samples are generated according to a Gaussian random variable with mean 0 and variance 400 ($\sigma = 20$), and the next 480 samples are drawn according to a Gaussian random variable with mean 0 and variance 1. This sequence is once quantized using a uniform PCM scheme and once using a nonuniform PCM scheme. Compare the resulting SQNR in the two cases.

SOLUTION

The sequence is generated by the MATLAB command

$$a = [20 * \mathrm{randn}(1, 20) \mathrm{randn}(1, 480)]$$

Now we can apply the u_pcm.m and the mula_pcm.m files to determine the resulting SQNR. The resulting SQNRs are 20.49 dB and 24.95 dB, respectively. In this case the performance of the nonuniform PCM is definitely superior to the performance of the uniform PCM.