Chapter 7 Analog-to-Digital Conversion (III)

Nonuniform Quantization (1/10)

- If we relax the condition that the quantization region (except for the first and the last one) to be of equal length, then we are minimizing the distortion with less constraints
- The resulting quantizer will perform better than a uniform quantizer with the same number of levels
- We are interested in designing the optimal mean squared error quantizer with N levels of quantization with no other constraint on the regions

Nonuniform Quantization (2/10)

- The average distortion D is given by

$$D=\int_{-\infty}^{a_1}(x-\hat{x}_1)^2 f_X(x)dx+\sum_{i=1}^{N-2}\int_{a_i}^{a_{i+1}}(x-\hat{x}_{i+1})^2 f_X(x)dx$$
$$+\int_{a_{N-1}}^{\infty}(x-\hat{x}_N)^2 f_X(x)dx. \qquad (7.2.8)$$

- There exists a total of $2N$-1 variables in this expression ($a_1,a_2,\ldots,a_{N\text{-}1}$ and $\hat{x}_1,\hat{x}_2,\ldots,\hat{x}_N$). Minimization of D is to be done with respect to these variables

Nonuniform Quantization (3/10)

- Differentiating with respect to a_i, i=1,2,…,N-1, yields

$$\frac{\partial D}{\partial a_i} = f_X(a_i)[(a_i - \hat{x}_i)^2 - (a_i - \hat{x}_{i+1})^2] = 0, \qquad (7.2.9)$$

$$f_X(a_i) \neq 0 \Rightarrow (a_i - \hat{x}_i)^2 = (a_i - \hat{x}_{i+1})^2$$

$$\Rightarrow (a_i - \hat{x}_i) = -(a_i - \hat{x}_{i+1})$$

which results in

$$a_i = \frac{1}{2}(\hat{x}_i + \hat{x}_{i+1}). \qquad (7.2.10)$$

- This result means that, in an optimal quantizer, *the midpoints of the quantization levels are the boundaries of the quantization regions*
- Because quantization is done on a minimum distance basis, each x value is quantized to the nearest $\hat{x}_i$, i=1,2,…,N-1

Nonuniform Quantization (4/10)

- To determine the quantized values $\hat{x}_i$, $i=1,\ldots,N$, we differentiate D with respect to $\hat{x}_i$ and define $a_0=-\infty$ and $a_N=+\infty$. Thus, we obtain

$$\frac{\partial D}{\partial \hat{x}_i} = \int_{a_{i-1}}^{a_i} 2(x-\hat{x}_i) f_X(x)dx = 0, \qquad (7.2.11)$$

$$\Rightarrow \int_{a_{i-1}}^{a_i} (x-\hat{x}_i) f_X(x) = 0$$

$$\Rightarrow \int_{a_{i-1}}^{a_i} x f_X(x)dx = \hat{x}_i \int_{a_{i-1}}^{a_i} f_X(x)dx$$

- This results in

$$\hat{x}_i = \frac{\int_{a_{i-1}}^{a_i} x f_X(x)dx}{\int_{a_{i-1}}^{a_i} f_X(x)dx} \qquad (7.2.12)$$

Nonuniform Quantization (5/10)

- Eq. (7.2.12) shows that in an optimal quantizer, the quantized value (or representation point) for a region should be chosen to be the centroid of that region
- Eq. (7.2.10) and (7.2.12) give the necessary conditions for a scalar quantizer to be optimal; they are known as the Lloyd-Max conditions
- The Lloyd-Max conditions can be summarized as follows:
  - The boundaries of the quantization regions are the midpoints of the corresponding quantized values
  - The quantized values are the centroids of the quantization regions
- These two conditions are necessary conditions, not sufficient conditions

Nonuniform Quantization (6/10)

- Although the Lloyd-Max conditions are very simple, they do not result in analytical solutions to the optimal quantizer design
- The usual method of designing the optimal quantizer is to start with a set of quantization regions and then, using the second criterion, to find the quantized values. Then, we design new quantization regions for the new quantized values, and alternate between the two steps until the distortion does not change much from one step to the next
- Table 7.2 shows the optimal nonuniform quantizer for a zero-mean, unit-variance, Gaussian source

Nonuniform Quantization (7/10)

TABLE 7.2 OPTIMAL NON-UNIFORM QUANTIZER FOR A GAUSSIAN SOURCE

N	$\pm a_i$	$\pm \hat{x}_i$	D	$H(\hat{X})$
1	—	0	1	0
2	0	0.7980	0.3634	1
3	0.6120	0, 1.224	0.1902	1.536
4	0, 0.9816	0.4528, 1.510	0.1175	1.911
5	0.3823, 1.244	0, 0.7646, 1.724	0.07994	2.203
6	0, 0.6589, 1.447	0.3177, 1.000, 1.894	0.05798	2.443
7	0.2803, 0.8744, 1.611	0, 0.5606, 1.188, 2.033	0.04400	2.647
8	0, 0.5006, 1.050, 1.748	0.2451, 0.7560, 1.344, 2.152	0.03454	2.825
9	0.2218, 0.6812, 1.198, 1.866	0, 0.4436, 0.9188, 1.476, 2.255	0.02785	2.983
10	0, 0.4047, 0.8339, 1.325, 1.968	0.1996, 0.6099, 1.058, 1.591, 2.345	0.02293	3.125
11	0.1837, 0.5599, 0.9656, 1.436, 2.059	0, 0.3675, 0.7524, 1.179, 1.693, 2.426	0.01922	3.253
12	0, 0.3401, 0.6943, 1.081, 1.534, 2.141	0.1684, 0.5119, 0.8768, 1.286, 1.783, 2.499	0.01634	3.372
13	0.1569, 0.4760, 0.8126, 1.184, 1.623, 2.215	0, 0.3138, 0.6383, 0.9870, 1.381, 1.865, 2.565	0.01406	3.481
14	0, 0.2935, 0.5959, 0.9181, 1.277, 1.703, 2.282	0.1457, 0.4413, 0.7505, 1.086, 1.468, 1.939, 2.625	0.01223	3.582

Nonuniform Quantization (8/10)

15	0.1369, 0.4143, 0.7030, 1.013, 1.361, 1.776, 2.344	0, 0.2739, 0.5548, 0.8512, 1.175, 1.546, 2.007, 2.681	0.01073	3.677
16	0, 0.2582, 0.5224, 0.7996, 1.099, 1.437, 1.844, 2.401	0.1284, 0.3881, 0.6568, 0.9424, 1.256, 1.618, 2.069, 2.733	0.009497	3.765
17	0.1215, 0.3670, 0.6201, 0.8875, 1.178, 1.508, 1.906, 2.454	0, 0.2430, 0.4909, 0.7493, 1.026, 1.331, 1.685, 2.127, 2.781	0.008463	3.849
18	0, 0.2306, 0.4653, 0.7091, 0.9680, 1.251, 1.573, 1.964, 2.504	0.1148, 0.3464, 0.5843, 0.8339, 1.102, 1.400, 1.746, 2.181, 2.826	0.007589	3.928
19	0.1092, 0.3294, 0.5551, 0.7908, 1.042, 1.318, 1.634, 2.018, 2.55	0, 0.2184, 0.4404, 0.6698, 0.9117, 1.173, 1.464, 1.803, 2.232, 2.869	0.006844	4.002
20	0, 0.2083, 0.4197, 0.6375, 0.8661, 1.111, 1.381, 1.690, 2.068, 2.594	0.1038, 0.3128, 0.5265, 0.7486, 0.9837, 1.239, 1.524, 1.857, 2.279, 2.908	0.006203	4.074
21	0.09918, 0.2989, 0.5027, 0.7137, 0.9361, 1.175, 1.440, 1.743, 2.116, 2.635	0, 0.1984, 0.3994, 0.6059, 0.8215, 1.051, 1.300, 1.579, 1.908, 2.324, 2.946	0.005648	4.141
22	0, 0.1900, 0.3822, 0.5794, 0.7844, 1.001, 1.235, 1.495, 1.793, 2.160, 2.674	0.09469, 0.2852, 0.4793, 0.6795, 0.8893, 1.113, 1.357, 1.632, 1.955, 2.366, 2.982	0.005165	4.206

(continued overleaf)

Nonuniform Quantization (9/10)

- If a general Gaussian source with mean m and variance σ^2 is used, then the values a_i and $\hat{x}_i$ read from Table 7.2 are replaced with $m+\sigma a_i$ and $m+\sigma \hat{x}_i$, respectively
- **Example 7.2.3.** How would the results of Example 7.2.1 change if, instead of the uniform quantizer shown in Fig. 7.3, we used an optimal nonuniform quantizer with the same number of levels?
- We can find the quantization regions and the quantized values from Table 7.2 with $N=8$, and then use the fact that our source is an $N(0,400)$ source, *i.e.*, $m=0$ and $\sigma=20$.

Nonuniform Quantization (10/10)

- **Example 7.2.3. (Cont'd)** All a_i and $\hat{x}_i$ values read from the table should be multiplied by σ=20 and the distortion has to be multiplied by 400
- The SQNR is

$$SQNR = \frac{400}{13.816} = 28.95 \approx 14.62 \ dB,$$

which is 3.84 dB better than the SQNR of the uniform quantizer

Vector Quantization (1/7)

- In scalar quantization, if we are using a 4-level scalar quantizer and encoding each level into two bits, we are using two bits per each source output. For example, see Fig. 7.4
- The idea of vector quantization is to take blocks of source output, and design the quantizer in the n-dimensional Euclidean space, rather than doing the quantization based on single samples in a one-dimensional space

Vector Quantization (2/7)

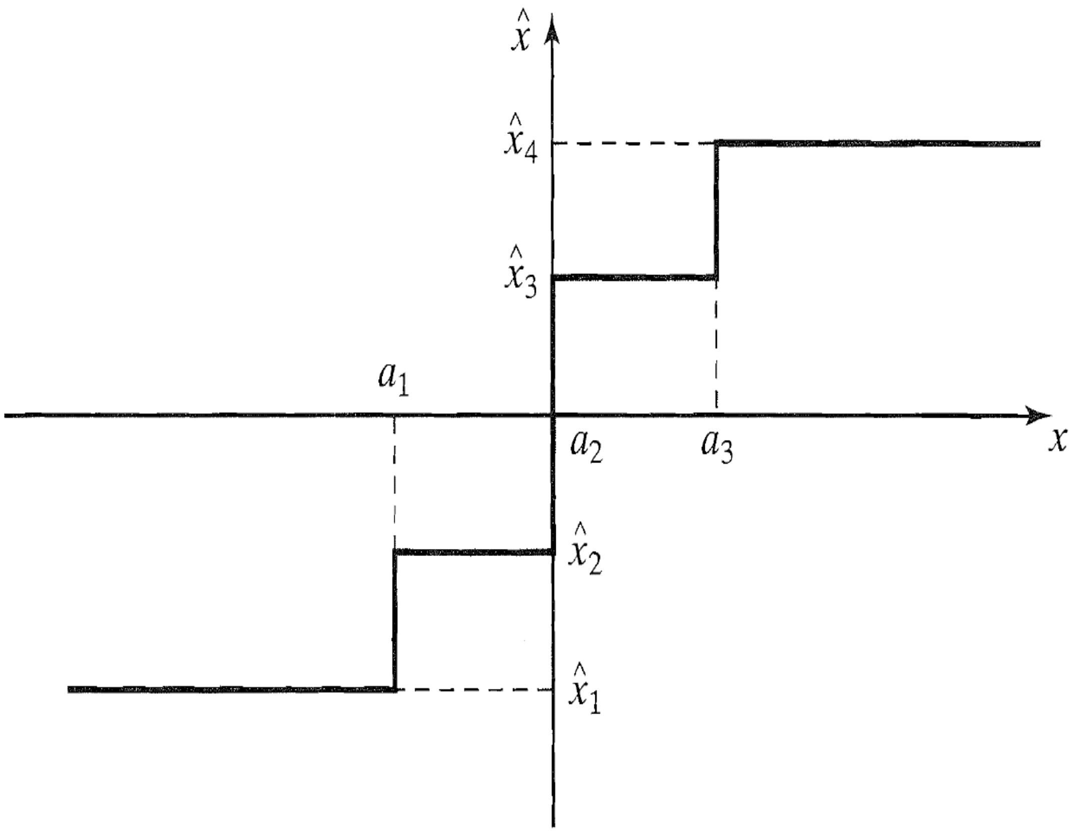

Figure 7.4 A 4-level scalar quantizer.

Vector Quantization (3/7)

- If we consider two samples of the source at each time, and we interpret these two samples as a point in a plane, the quantizer partitions the entire plane into 16 quantization regions, as shown in Fig. 7.5

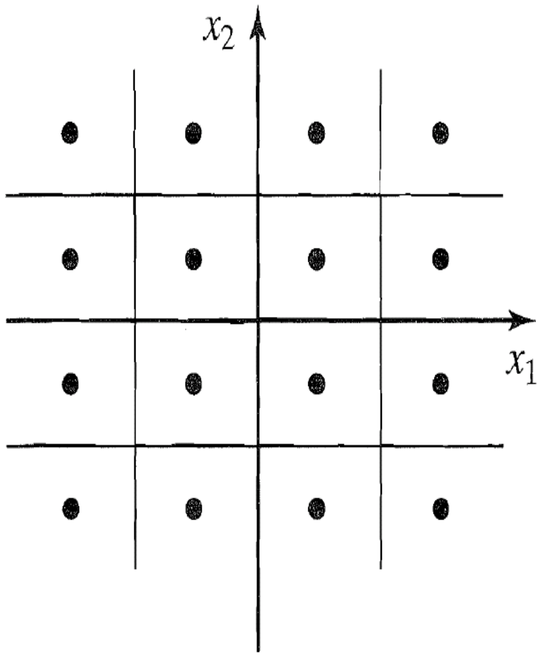

Figure 7.5 A scalar 4-level quantization applied to two samples.

Vector Quantization (4/7)

- If we allow 16 regions of any shape in the two-dimensional space, we are capable of obtaining better results than uniform quantization in the two-dimensional space
- This is equivalent to four bits per two source outputs or two bits per each source output
- If we are relaxing the requirement of having rectangular regions, the performance may be improved

Vector Quantization (5/7)

- Let us assume that the quantization regions in the n-dimensional space are denoted by $\mathcal{R}_i$, $1 \leqq i \leqq K$. These K regions partition the n-dimensional space
- Each block of source output of length n is denoted by the n-dimensional vector $x \in R^n$. If $\boldsymbol{x} \in \mathcal{R}_i$, it is quantized to $Q(\boldsymbol{x}) = \hat{\boldsymbol{x}}_i$
- If there are a total of K quantized values, $\log_2 K$ bits are enough to represent these values. This means that we require $\log_2 K$ bits per n source outputs, or the rate of the source code is

$$R = \frac{\log_2 K}{n} \qquad \text{(bits/source output)}$$

Vector Quantization (6/7)

- Fig. 7.6 shows a quantization scheme for n=2

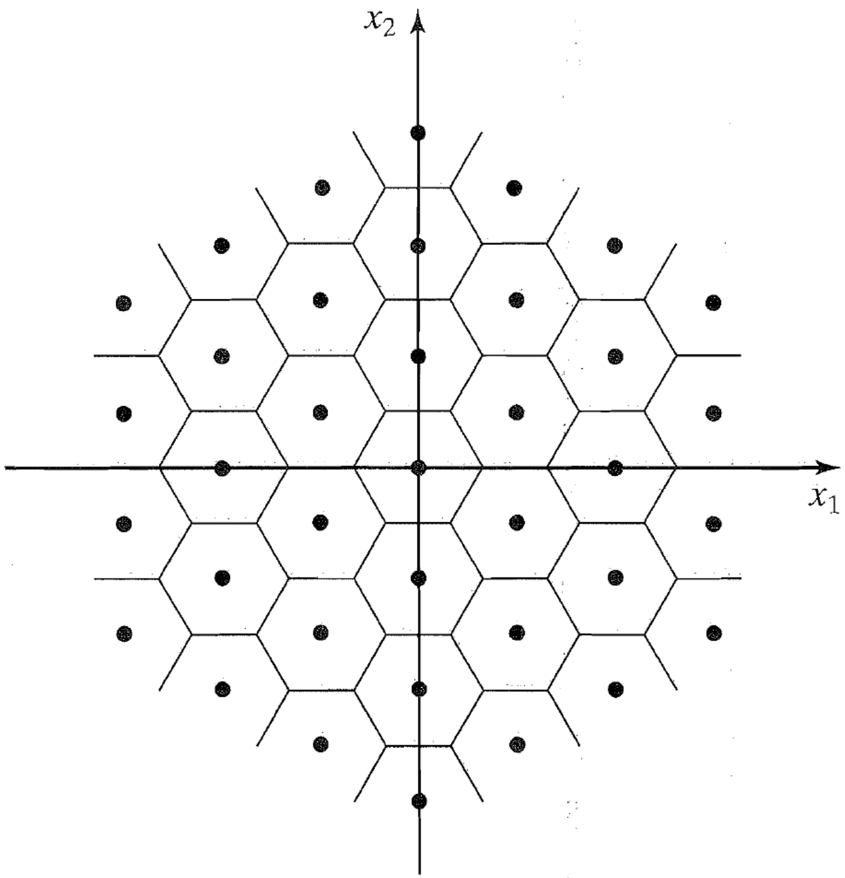

Figure 7.6 Vector quantization in two dimensions.

Vector Quantization (7/7)

- The optimal vector quantizer, of dimension *n* and number of levels *K*, chooses the regions $\mathscr{R}_i$'s and the quantized values $\hat{x}_i$'s such that the resulting distortion is minimized
- Borrowing the idea of optimial scalar quantization, we obtain the following criteria for an optimal vector quantizer design:

1. Region $\mathscr{R}_i$ is the set of all points in the *n*-dimensional space that are closer to $\hat{x}_i$ than any other $\hat{x}_j$, for all $j \neq i$; that is,

$$\mathscr{R}_i = \{x \in R^n : \|x - \hat{x}_i\| < \| x - \hat{x}_j \|, \ \forall j \neq i\};$$

2. $\hat{x}_i$ is the centroid of the region $\mathscr{R}_i$; that is,

$$\hat{\mathbf{x}}_i = \frac{1}{P(\mathbf{X} \in \mathscr{R}_i)} \int \cdots \int_{\mathscr{R}_i} \mathbf{x} f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}.$$

Encoding (1/3)

- In the encoding process, a sequence of bits are assigned to different quantization values
- If there are a total of $N=2^{\nu}$ quantization levels, ν bits are sufficient for the encoding process
- Since the sampling rate is f_s samples/second, we will have a bit rate of $R=\nu f_s$ bits/second
- Natural binary coding (NBC) assigns ν zeros to the lowest quantization level, ν ones to the highest quantization level
- Gray coding is to encode the quantized levels in a way that adjacent levels differ only in one bit

Encoding (2/3)

TABLE 7.3 NBC AND GRAY CODES FOR A 16-LEVEL QUANTIZATION

Quantization level	Level order	NBC code	Gray code
$\hat{x}_1$	0	0000	0000
$\hat{x}_2$	1	0001	0010
$\hat{x}_3$	2	0010	0011
$\hat{x}_4$	3	0011	0001
$\hat{x}_5$	4	0100	0101
$\hat{x}_6$	5	0101	0100
$\hat{x}_7$	6	0110	0110
$\hat{x}_8$	7	0111	0111
$\hat{x}_9$	8	1000	1111
$\hat{x}_{10}$	9	1001	1110
$\hat{x}_{11}$	10	1010	1100
$\hat{x}_{12}$	11	1011	1101
$\hat{x}_{13}$	12	1100	1001
$\hat{x}_{14}$	13	1101	1000
$\hat{x}_{15}$	14	1110	1010
$\hat{x}_{16}$	15	1111	1011

Encoding (3/3)

- Generate Gray codes

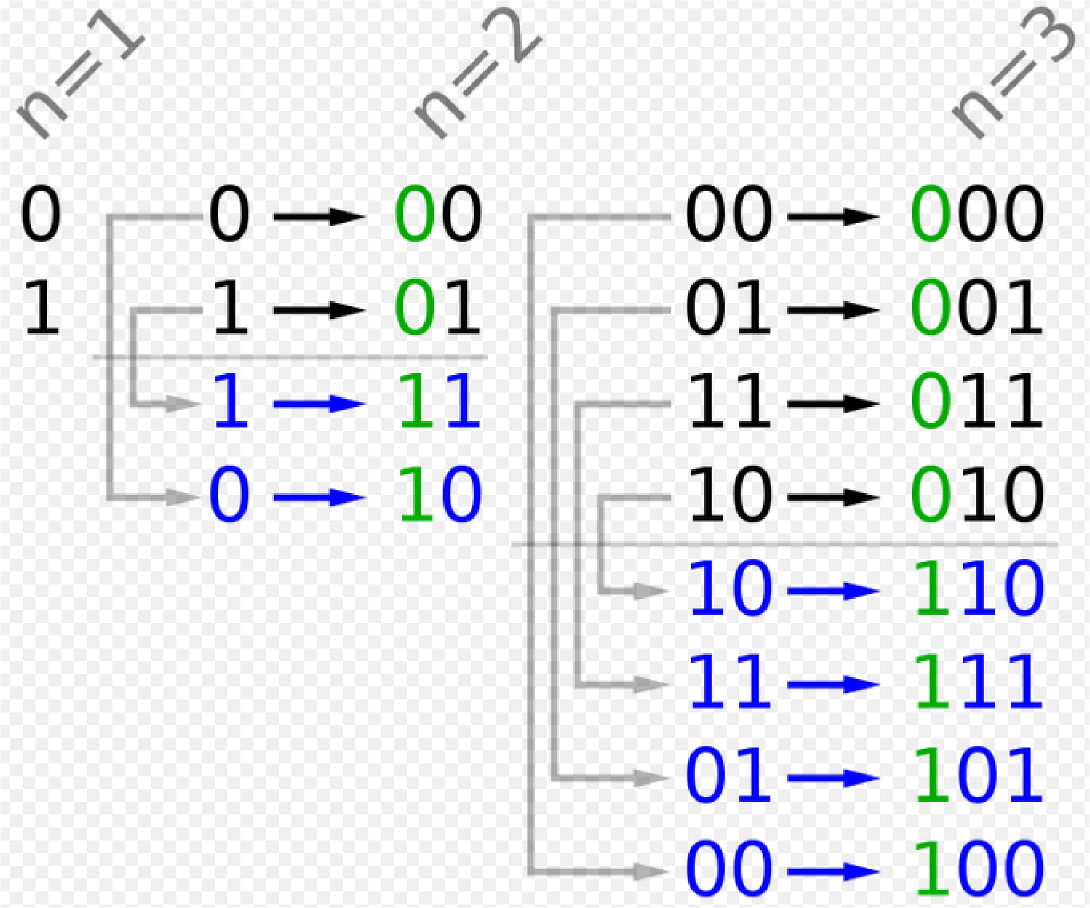

[WiKi]