Proc. IEEE Globecom Conf. (San Diego, CA), pp. 214-218, Nov. 1983

QUANTIZERS FOR SYMMETRIC GAMMA DISTRIBUTIONS

Peter Kabal

Department of Electrical Engineering† *McGill University*

INRS-Télécommunications Université du Québec

Abstract

This paper discusses minimum mean-square error quantization for symmetric distributions. If the distribution satisfies a log-concavity condition, the optimal quantizer is itself symmetric. For the gamma distribution often used to model speech signals, the log-concavity condition is not satisfied. It is shown that for this distribution both the uniformly spaced and the non-uniformly spaced optimal quantizers are not symmetrical. New quantization tables giving the optimal levels for quantizers for the gamma distribution are presented.

1. Introduction

This paper focusses on minimum mean-square error scalar quantizers for symmetric distributions. A number of authors have published tables of quantizers for distributions of interest in the processing of speech or visual signals [1,2,3,4]. These quantizers have been designed for the most part using the iterative methods outlined by Lloyd and Max [1,2]. It is well known that these design techniques applied to general probability distributions may produce quantizers that are only locally optimal. Even so, it seems to be widely assumed that for the symmetrical distributions encountered in practice, the resulting quantizers are also symmetrical. It is shown here that non-symmetric solutions may be optimal for distributions of more than pathological interest.

2. Lloyd-Max Quantizers

Quantization is the process of subdividing the range of a signal into non-overlapping regions. An output level is then assigned to represent each region.

Consider a N level quantizer with output levels $y_1, y_2, \ldots, y_N$. The output level y_k is associated with a decision region specified by its boundaries, the decision levels,

$$y_k \Leftrightarrow \{x_{i-1} < x \leq x_i\}\,, \qquad i = 1, 2, \ldots, N\,. \quad (1)$$

For convenience, the x_i are in increasing order and the two extreme decision levels are chosen to be $x_0 = \infty$ and $x_N = -\infty$. The total mean-square error can be differentiated with respect to x_k and y_k to give the necessary conditions that must be satisfied by a minimum mean-square error quantizer,

$$x_k = \frac{y_k + y_{k+1}}{2}\,, \quad \text{for } k = 1, 2, \ldots, N-1\,, \quad (2)$$

$$y_k = \frac{\int_{x_{k-1}}^{x_k} x p(x)\,dx}{\int_{x_{k-1}}^{x_k} p(x)\,dx}\,, \quad \text{for } k = 1, 2, \ldots, N\,. \quad (3)$$

The conditions (2) and (3) form the basis for an iterative process to determine optimal quantizers.

2.1 Iterative One-dimensional search

A variational technique proposed by both Lloyd and Max involves a one-dimensional search. An initial guess is made as to the value of the first output level y_1. The value of the decision level below this output level, in this case x_0, is known. The next decision level can be determined by finding the value of x_1 which satisfies (3), in this case for $k = 1$. This step will generally have to be carried out using iterative numerical techniques. The next step is to telescope the process to the next interval. This is done by using (2) to determine y_2 from y_1 and x_1. The process continues interval by interval to determine all the output levels. The last output level, y_N, determined in this manner will generally not be the conditional mean of the last interval. The difference between y_N and the conditional mean of the last interval can be used to determine an update for y_1 for the next iteration. The process of determining the output levels continues until sufficient precision has been achieved.

3. Uniqueness

The log-concavity test for a probability density function given by Fleischer [5] is

$$\frac{\partial^2 \log p(x)}{\partial x^2} < 0\,. \quad (4)$$

A probability density function that satisfies this condition has a unique stationary point in the mean-square error in terms of the $2N-1$ variables—the $N-1$

† 3480 University, Montreal, Quebec H3A 2A7.

6.6.1

decision levels and the N output levels. The iterative Lloyd-Max algorithms will find the globally optimal quantizer if the probability density is log-concave. The Gaussian distribution satisfies this condition and hence has a unique stationary point.

Consider the generalized gamma double-sided probability density function,

$$p(x) = \frac{\lambda(\lambda|x|)^{a-1}e^{-\lambda|x|}}{2\Gamma(a)}. \tag{5}$$

The parameter λ sets the variance of the distribution,

$$\sigma^2 = \frac{a(a+1)}{\lambda^2}. \tag{6}$$

The density function (5) becomes Laplacian for $a = 1$. For modelling the measured distribution of speech signals, the general gamma density with $a = \frac{1}{2}$ is often used. The log-concavity test for the general gamma density gives

$$\frac{\partial^2 \log p(x)}{\partial x^2} = -\frac{2(a-1)}{x^2}, \qquad x \neq 0. \tag{7}$$

For $a \geq 1$, a unique stationary point exists.

4. Symmetric Distributions

Consider a distribution which is symmetric about its mean. For every quantizer with a given set of output levels, another with the same mean-square error is generated by simply reflecting the levels about the mean. This argument indicates that if a symmetrical distribution is log-concave, the optimal and unique quantizer will have levels symmetrically placed about the mean.

For symmetrical distributions, a solution which satisfies the necessary conditions for optimality can be obtained by considering the density on one side of the mean. If the total number of levels is even, the problem is solved using $N/2$ levels for the density $2p(x)$, $x \geq \bar{x}$, where $\bar{x}$ is the mean of the distribution. If the total number of levels is odd, the problem can again be solved with half the number of levels but with one level fixed at the mean. In either case, the solution determined for one side of the distribution can be reflected about the mean to produce a symmetrical solution for the distribution. Thus every symmetric distribution has a symmetric quantizer which satisfies the necessary conditions for a minimum mean-square error quantizer. However, this solution may represent a local minimum or a saddle point for distributions which are not log-concave.

5. Gamma Distribution

Consider the general gamma distribution with $a = \frac{1}{2}$, henceforth referred to simply as the gamma distribution. For the gamma distribution, more than one stationary point may exist.

The optimal one level quantizer has an output level at the mean and is symmetric. For the two level quantizer, we can determine the optimal output levels given a decision level x_1 using (3). Because of the symmetry, consider only $x_1 \geq 0$,

$$\begin{aligned} y_1 &= -\frac{Q(\sqrt{u}) + \sqrt{\frac{u}{2\pi}}\, e^{-u/2}}{\sqrt{3}\,[1 - Q(\sqrt{u})]}, \\ y_2 &= \frac{Q(\sqrt{u}) + \sqrt{\frac{u}{2\pi}}\, e^{-u/2}}{\sqrt{3}\, Q(\sqrt{u})}, \end{aligned} \tag{8}$$

where $u = \sqrt{3}\,x_1$ and $Q(x)$ is the integral of the tail of the unit variance Gaussian density function. Combining these terms using (2) gives a single equation,

$$2u - \left[Q(\sqrt{u}) + \sqrt{\frac{u}{2\pi}}\, e^{-u/2}\right] \frac{1 - 2\,Q(\sqrt{u})}{Q(\sqrt{u})[1 - Q(\sqrt{u})]} = 0, \tag{9}$$

Equation (9) has a solution $x_1 = 0$, the symmetric solution, as well as solutions at $x_1 = \pm 0.622$. For the symmetric two level quantizer,

$$y_1 = -1/\sqrt{3}, \qquad y_2 = 1/\sqrt{3}, \tag{10}$$

$$\overline{e^2} = 2/3. \tag{11}$$

The non-symmetric solutions give a mean-square error which is less than for the symmetric solution. For $x_1 = +0.622$,

$$y_1 = -0.266, \qquad y_2 = +1.509, \tag{12}$$

$$\overline{e^2} = 0.599. \tag{13}$$

Fig. 1 is a contour plot of the signal-to-noise ratio (SNR) as a function of y_1 and y_2. In this plot, the decision level is constrained to lie midway between the output levels (see (2)). The contour plot shows two-fold symmetry, since the quantizers (a, b), (b, a), $(-a, -b)$ and $(-b, -a)$ all have the same mean-square error. A symmetric quantizer is restricted to lie on the diagonal line, $y_1 = -y_2$. The optimal non-symmetric and symmetric quantizers are shown as crosses on the contour plot. This view shows that the best symmetric quantizer lies at a saddle point in the y_1–y_2 space. This point is also a saddle point in the 3-space y_1–y_2–x_1 since the decision level x_1 is chosen optimally in the view shown.

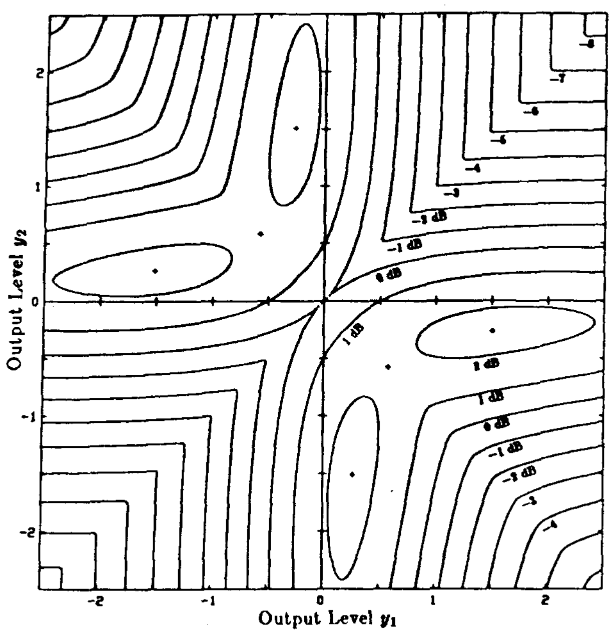

Fig. 1 Contour Plot of the SNR for a Two Level Quantizer

For a three level quantizer, Fig. 2 gives a contour plot of the SNR as a function of the two decision levels, x_1 and x_2. The output levels are constrained to be the conditional means of the decision regions (see (3)) In this case, the optimal solution corresponds to a symmetric quantizer.

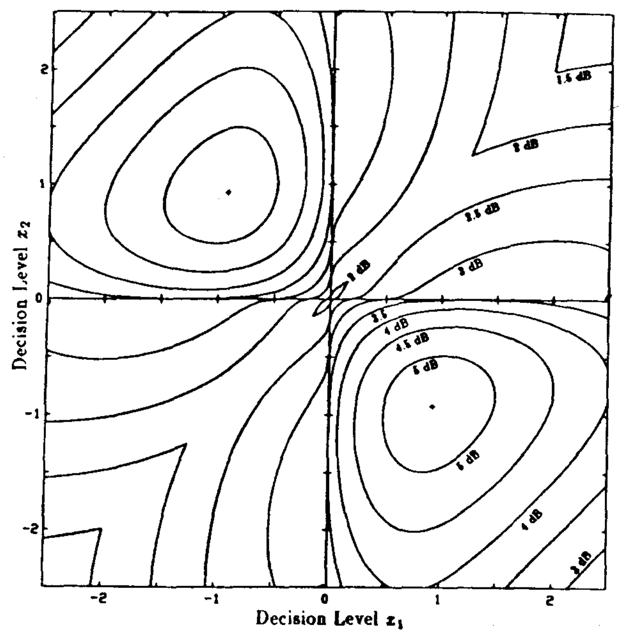

Fig. 2 SNR for a Three Level Quantizer

For higher numbers of levels, a tack suggested by the iterative method given above was adopted. Given an initial output level y_1, subsequent output levels up to y_N are found. The difference between y_N and $\hat{y}_N$, the conditional mean of the last decision region, is plotted. When this difference is zero, the entire quantizer satisfies the necessary conditions for a minimum mean-square error quantizer. The minimum of the mean-square error corresponding to the zero crossings of this difference determines the global minimum. Fig. 3 shows such a plot for a six level quantizer. The plot also shows the SNR as a function of the first output level. Three zero crossings appear. The middle one corresponds to a symmetric solution with $y_1 = -4.773$. The other two correspond to a non-symmetric solution with $y_1 = -3.111$ or $y_1 = -3.818$. The quantizers corresponding to these last two values of y_1 are reflections about zero of each other. A non-symmetric solution gives the best SNR.

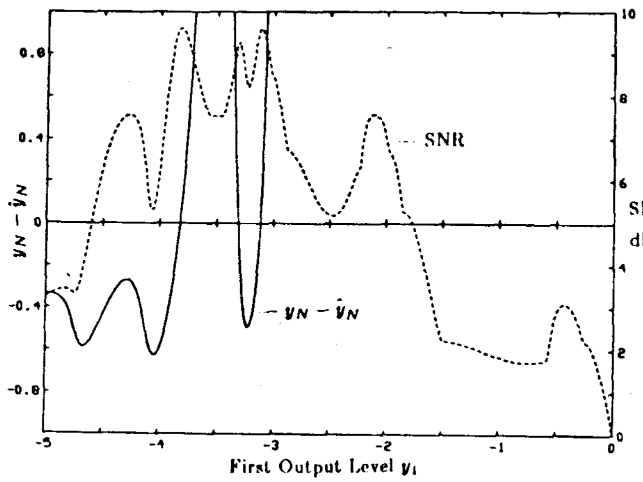

Fig. 3 Last Interval Difference and SNR for a Six Level Quantizer

Previously published tables for the gamma distribution have given only symmetric solutions. Table I compares the best symmetric solution with the optimal solution for selected values of N. The three numbers below the quantizer output levels are the mean-square error (for a unit variance distribution), the SNR (in dB) and the entropy of the quantizer. For odd values of N, the symmetric solution is optimal, although non-optimal non-symmetric solutions satisfying (2) and (3) are possible for $N \geq 5$. For even values of N, both symmetric and non-symmetric solutions are shown. For the larger values of N, several quantizers (apart from those obtained by reflecting the levels about zero) satisfy the necessary conditions of (2) and (3). For example for $N = 14$, three distinct non-symmetric and one symmetric configurations can be found.

For uniformly spaced quantizers, the optimal quantizers are not necessarily symmetrically placed with respect to the mean. This is clear from the two level example above, for in this case the uniform and non-uniform quantizers are the same. Table II compares symmetric and non-symmetric uniformly spaced quantizers. The table entries are the interval between levels, Δ and the offset of the quantizer relative to a symmetrical quantizer, ϵ. Specifically, the output levels are given by

$$y_i = \left(i - \frac{N+1}{2}\right)\Delta + \epsilon, \qquad i = 1, 2, \ldots, N. \quad (14)$$

The three numbers at the bottom of each entry in the table are the mean-square error (for a unit variance distribution), the SNR (in dB) and the entropy of the quantizer. Note that for the symmetric case, adding an additional output level to a quantizer with an odd number of levels actually increases the mean-square error.

Another issue of interest is the convexity of the mean-square error as a function of the number of bits, $\log_2 N$. For integral numbers of bits, the mean-square error for the non-symmetric quantizers (both uniformly and non-uniformly spaced) is convex while for the symmetric quantizers it is not.

Non-convexity of the mean-square error can have interesting consequences. For instance, consider coding a gamma distributed signal with 1 bit per sample. For symmetric quantizers, a lower mean-square error is obtained if samples are coding alternately using a 2 bit and a 0 bit quantizer, than if a 1 bit quantizer is used for every sample.

6. General Gamma Distribution

Plots of the difference between y_N and $\hat{y}_N$, the conditional mean of the last decision region, were also generated for the generalized gamma distribution. Fig. 4 shows such a plot for two level quantizer acting on a density with parameter $a = 0.9$. The plot shows that a non-symmetric solution is optimal for this case. As the parameter a approaches unity, the three zero crossings evident in the plot coalesce to give a single, unique solution for the Laplace density. For values of the parameter a below unity, a non-symmetric solution is optimal. This then indicates that the Laplace distribution occupies a unique place amongst the family of general gamma distributions—on the boundary separating those distributions which have unique minima and those which do not. For the general gamma distribution, log-concavity seems to be both a necessary and sufficient condition for uniqueness.

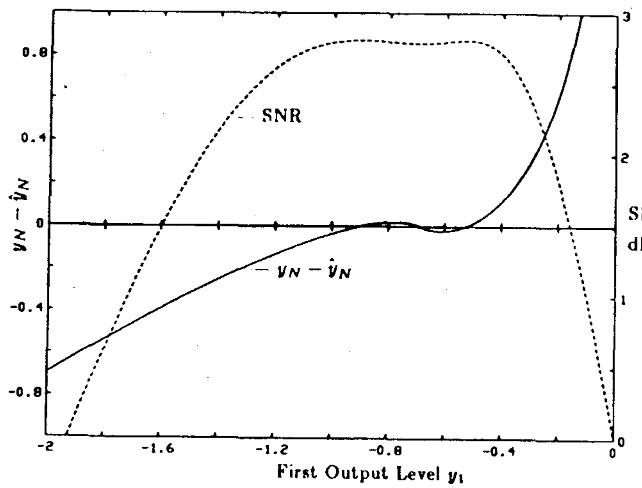

Fig. 4 Last Interval Difference and SNR for a General Gamma Distribution ($a = 0.9$)

References

1. S.P. Lloyd, "Least squares quantization in PCM", Bell Telephone Laboratory Memorandum, 1957, reprinted in *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 129–137, March 1982.
2. J. Max, "Quantization for minimum distortion", *IRE Trans. Inform. Theory*, vol. IT-6, pp. 7–12, March 1960.
3. M.D. Paez and T.H. Glisson, "Minimum mean-squared-error quantization in speech PCM and DPCM systems", *IEEE Trans. Commun. Technol.*, vol. COM-20, pp. 225–230, April 1972.
4. W.C. Adams and C.E. Giesler, "Quantizing characteristics for signals having Laplacian amplitude probability density function", *IEEE Trans. Commun.*, vol. COM-26, pp. 1295–1297, August 1978.
5. P.E. Fleischer, "Sufficient conditions for achieving minimum distortion in a quantizer", *IEEE Int. Conv. Rec.*, 1964, part 1, pp. 104–111.

$N=1$	$N=2$		$N=3$	$N=4$	
0.000	±0.577	−0.266 1.609	0.000	±0.313	−1.981 0.899
1.000	0.6667	0.5990	±1.851	±2.223	−0.108 2.881
0.00 dB	1.76 dB	2.23 dB	0.2961	0.2318	0.2127
0.00 bits	1.00 bits	0.61 bits	5.29 dB	6.35 dB	6.72 dB
			0.94 bits	1.58 bits	1.27 bits

$N=5$	$N=6$		$N=7$	$N=8$	
0.000	±0.210	−3.111 0.635	0.000	±0.155	−3.956 0.487
±1.039	±1.291	−1.110 1.770	±0.710	±0.899	−1.901 1.274
±3.034	±3.307	−0.058 3.818	±1.853	±2.057	−0.755 2.457
0.1395	0.1171	0.1090	±3.905	±4.121	−0.036 4.538
8.56 dB	9.32 dB	9.63 dB	8.075×10^{-2}	7.047×10^{-2}	6.632×10^{-2}
1.50 bits	1.99 bits	1.73 bits	10.93 dB	11.52 dB	11.78 dB
			1.90 bits	2.31 bits	2.08 bits

$N=9$	$N=10$		$N=11$	$N=12$	
0.000	±0.122	−4.631 0.393	0.000	±0.100	−5.193 0.327
±0.534	±0.684	−2.546 0.990	±0.424	±0.548	−3.087 0.806
±1.326	±1.488	−1.359 1.815	±1.025	±1.159	−1.877 1.435
±2.511	±2.682	−0.564 3.023	±1.852	±1.993	−1.049 2.281
±4.595	±4.773	−0.024 5.126	±3.061	±3.207	−0.466 3.505
5.258×10^{-2}	4.704×10^{-2}	4.462×10^{-2}	±5.166	±5.317	−0.017 5.624
12.79 dB	13.28 dB	13.50 dB	3.694×10^{-2}	3.362×10^{-2}	3.208×10^{-2}
2.22 bits	2.57 bits	2.36 bits	14.33 dB	14.73 dB	14.94 dB
			2.48 bits	2.79 bits	2.59 bits

$N=13$	$N=14$		$N=15$	$N=16$	
0.000	±0.084	−5.675 0.280	0.000	±0.073	−6.100 0.243
±0.350	±0.455	−3.554 0.677	±0.297	±0.387	−3.968 0.582
±0.831	±0.945	−2.329 1.183	±0.696	±0.795	−2.730 1.004
±1.462	±1.582	−1.481 2.690	±1.203	±1.307	−1.869 1.525
±2.308	±2.433	−0.850 3.925	±1.851	±1.959	−1.219 2.184
±3.533	±3.662	−0.367 6.056	±2.711	±2.822	−0.711 3.053
±5.653	±5.785	0.013 6.373	±3.947	±4.061	−0.310 4.298
2.736×10^{-2}	2.521×10^{-2}	2.418×10^{-2}	±6.078	±6.195	−0.010 6.437
15.63 dB	15.98 dB	16.16 dB	2.107×10^{-2}	1.961×10^{-2}	1.888×10^{-2}
2.70 bits	2.97 bits	2.79 bits	16.76 dB	17.08 dB	17.24 dB
			2.89 bits	3.14 bits	2.97 bits

Table I Output Levels for Non-Uniform Gamma Quantizers

$N=1$	$N=2$		$N=3$	$N=4$	
-	1.155	1.775	1.851	1.066	1.560
0.000	0.000	±0.622	0.000	0.000	±0.710
1.000	0.6667	0.5990	0.2961	0.3200	0.2330
0.00 dB	1.76 dB	2.23 dB	5.29 dB	4.95 dB	6.33 dB
0.00 bits	1.00 bits	0.61 bits	0.94 bits	1.67 bits	1.14 bits

$N=5$	$N=6$		$N=7$	$N=8$	
1.342	0.912	1.208	1.079	0.796	0.998
0.000	0.000	±0.571	0.000	0.000	±0.480
0.1597	0.1934	0.1346	0.1045	0.1323	9.130×10^{-2}
7.97 dB	7.14 dB	8.71 dB	9.81 dB	8.78 dB	10.40 dB
1.35 bits	1.94 bits	1.48 bits	1.62 bits	2.13 bits	1.72 bits

$N=9$	$N=10$		$N=11$	$N=12$	
0.913	0.708	0.858	0.798	0.640	0.757
0.000	0.000	±0.416	0.000	0.000	±0.369
7.551×10^{-2}	9.756×10^{-2}	6.747×10^{-2}	5.795×10^{-2}	7.560×10^{-2}	5.260×10^{-2}
11.22 dB	10.11 dB	11.71 dB	12.37 dB	11.21 dB	12.79 dB
1.83 bits	2.28 bits	1.91 bits	2.00 bits	2.40 bits	2.07 bits

$N=13$	$N=14$		$N=15$	$N=16$	
0.712	0.585	0.681	0.645	0.540	0.620
0.000	0.000	±0.334	0.000	0.000	±0.305
4.633×10^{-2}	6.071×10^{-2}	4.255×10^{-2}	3.816×10^{-2}	5.008×10^{-2}	3.536×10^{-2}
13.34 dB	12.17 dB	13.71 dB	14.18 dB	13.00 dB	14.51 dB
2.15 bits	2.52 bits	2.21 bits	2.27 bits	2.62 bits	2.32 bits

Table II Step Size and Offset for Uniform Gamma Quantizers