



Quantisation noise

In order to encode a continuous-time signal, we must first quantise it into a finite number of discrete amplitude signals. The quantisation depends on the desired number of different signal levels (the resolution), and the maximum variation of the signal to be represented (the range). Once quantised, the instantaneous values of the signals are lost, and can never be reconstructed exactly.

The quantising operation can be viewed as a function operating on each sample value. In the case of **uniform quantising** there is a uniform spacing between the quantisation levels:




Equivalently, the quantisation operation can be considered in terms of an input-output relation of the form:



The signal is mapped through this function prior to being sampled.

We can define the **quantisation noise**, or quantisation error, as the time function that is the difference $e(t) = f(t) - f_q(t)$, where $f(t)$ and $f_q(t)$ are the signals before and after quantisation:




Since we are discussing discrete-time signals, the error at each sample

point is given by

$$e_n = e(nT_s) = f(nT_s) - f_q(nT_s).$$

Here $f_q(nT_s)$ is the quantised signal value at the time of the n th sample.

Consider the case of uniform quantisation, with a spacing a between quantisation levels. In the absence of any further specific knowledge, a fair assumption is that the quantising error will be uniformly distributed between $-a/2$ and $a/2$:



The average value of this error is zero, owing to the symmetry of the probability density. The mean square error (MSE) is

$$\text{MSE} = \int_{-\infty}^{\infty} p(e_n) e_n^2 de_n = \int_{-a/2}^{a/2} \frac{1}{a} e_n^2 de_n = \frac{a^2}{12}.$$

It is reasonable to model this loss of accuracy as an additive noise random process. The quantised signal can therefore be considered to be generated as the sum of the unquantised sample values, and a zero-mean random noise process with power $a^2/12$:

$$f(nT_s) = f_q(nT_s) + e(nT_s).$$

3 Peak signal SNR

The peak voltage in the input signal is $V_{pp}/2 = na/2$, so the peak signal power is $(na/2)^2$. The ratio of peak signal power to average quantisation

noise power is therefore

$$\left(\frac{S}{N}\right)_{\text{peak quant}} = \frac{(na/2)^2}{a^2/12} = 3n^2$$

or, in decibels,

$$\left(\frac{S}{N}\right)_{\text{peak quant (dB)}} = 4.8 + 20 \log_{10} n.$$

For a binary code with $n = 2^m$ this becomes

$$\left(\frac{S}{N}\right)_{\text{peak quant (dB)}} = 4.8 + 6.02m.$$

The peak SNR therefore increases by 6dB for every additional bit used in a binary system.

4 Mean signal SNR

Essentially similar results are obtained upon defining a mean power SNR. Under the assumption that the signal occupies each of the n quantisation levels with equal probability, the average signal power is equal to

$$S = \frac{2}{n} \left\{ \left(\frac{a}{2}\right)^2 + \left(\frac{3a}{2}\right)^2 + \dots + \left(\frac{(n-1)a}{2}\right)^2 \right\}.$$

This sum can be written as

$$S = \frac{1}{12}(n^2 - 1)a^2,$$

so the mean power output SNR is given by


$$\left(\frac{S}{N}\right) = n^2 - 1 \approx n^2.$$

This only differs by a constant from the peak ratio discussed before.


5 Nonuniform quantisation

In the case of uniform quantisation levels, the quantisation noise power depends only on the spacing between the levels. This is independent of the actual signal level at any instant. In some applications this is undesirable — for example, in a speech system a fixed quantisation noise power will be more objectionable when a quiet speaker is speaking than when a loud one is.

A remedy is to use nonuniform quantisation levels. This can be achieved by using a nonuniform quantiser:




More often, nonuniform quantisation is achieved by first distorting the original signal with a nonlinear **compressor** characteristic, and then using a uniform quantiser on the result:



A given signal change at small magnitudes will then carry the uniform quantiser through more steps than the same change at large magnitudes. At the receiver, an inverse compression characteristic (or **expansion**) is applied, so that the overall transmission is not distorted. The processing pair (compression and expansion) is usually referred to as **companding**.

A uniform quantiser makes sense when the probability distribution of the signal in the range $-V_{\max}$ to V_{\max} is uniform. If we have reason to believe that the distribution is nonuniform, and we know what the actual distribution is, then we can place nonuniform quantisation levels in an optimal manner.

Recall from the discussion on information theory that the entropy is maximised if the probability of occurrence of each level is equal. Thus we should choose the quantisation levels to ensure that this condition is met:



Stremler discusses this in more detail.

Questions

1. A signal voltage in the frequency range 100 to 4000Hz is limited to a peak-to-peak swing of 3V. It is sampled at a uniform rate of 8kHz, and the samples are quantised to 64 evenly spaced levels. Calculate and compare the bandwidths and ratios of peak signal to RMS quantisation noise if the quantised samples are transmitted either as binary digits or as four-level pulses.
2. The sinusoidal voltage $10 \sin(6280t)$ is sampled at $t = 0.33\text{ms}$ and thereafter periodically at a 3kHz rate. The samples are then quantised into eight voltage levels and coded into binary digits. Use MATLAB to plot the original voltage and below it the outputs of the sampler, the quantiser, and the encoder. Calculate the RMS quantisation noise.