ADSP - Exercise 2: When Does Quantization Look Like Noise?

Learning objectives

After this exercise, you should be able to

- explain the additive quantization-noise model and its assumptions,
- interpret the quantization-noise power $P_n = \Delta^2/12$,
- explain the approximate SNR gain of $6\,\mathrm{dB}$ per additional bit,
- assess the effect of signal level and converter utilization,
- distinguish granular quantization error from overload (clipping), and
- use a waveform, error histogram, and error spectrum to judge whether the white-noise model is useful.

Guiding question: Under which conditions may the quantization error be treated as additive white noise?

Notation (lecture notes, Part 2, slides 23-32): $v(n)$ unquantized sequence, $v_\mathrm{q}(n) = Q[v(n)]$ quantized sequence, $e_\mathrm{q}(n)$ quantization error, d_l decision levels, v_l quantization levels, I_l quantization intervals, R range, b word length, L number of levels, Δ step size, P_v signal power, P_n noise power. The signal-to-quantization-noise ratio is written SNR, as in the lecture (often also called SQNR).

1. Reactivation: the quantizer

An AD converter samples and quantizes; this exercise examines the quantizer. A quantizer maps the discrete-time, continuous-valued signal $v(n)$ to

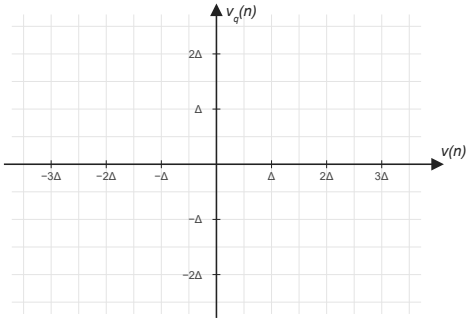
$$v_\mathrm{q}(n) = Q[v(n)].$$

The amplitude range is divided into intervals $I_l = \{d_{l-1} < v(n) \leq d_l\}$, and every sample in I_l is mapped to the quantization level v_l. For a converter range R and word length b, the lecture gives the step size

$$\Delta = \frac{R}{2^b}.$$

In this exercise all 2^b code words are used, so $L = 2^b$ and $\Delta = R/L$.

1. What do b, L, R, and Δ describe?
2. How do a midtread and a midrise quantizer differ at zero?
3. Sketch the characteristic of a uniform midrise quantizer with $b = 2$ ($L = 4$) in the prepared axes. Mark the decision levels d_l, the quantization levels v_l, Δ, and the non-overload range.
4. Why is the 3-bit midtread characteristic of the lecture (slide 26) not symmetric about zero?



Midrise, b = 2 (L = 4): add the decision levels d_l, the quantization levels v_l, Δ, and the non-overload range.

The lecture defines the quantization error as

$$e_\mathrm{q}(n) = v(n) - v_\mathrm{q}(n).$$

State the error interval for samples that do not overload the quantizer.

2. Predict before observing

The demonstration uses the same uniform midrise quantizer for several inputs. Before each case is displayed, predict the error in time, the error histogram, the error spectrum, and whether the white-noise model will be useful:

- well-utilized sinusoid at 997.3 Hz (no short period),
- very small sinusoid,
- periodic sinusoid with exactly 48 samples per period,
- overloaded sinusoid.

Explain why a nearly uniform error histogram does not imply a white error sequence.

Also predict for the well-utilized sinusoid: by how many dB does the SNR change when the word length is increased from $b = 8$ to $b = 9$, and when it is reduced from $b = 8$ to $b = 4$?

3. Interactive demonstration

For every preset in the notebook:

1. make a prediction before the preset is selected,
2. compare $v(n)$, $v_q(n)$, and $e_q(n)$,
3. inspect the distribution and spectrum of $e_q(n)$, and
4. decide which assumption of the noise model is plausible or violated.

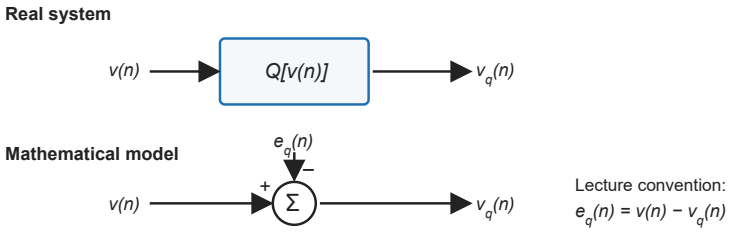
Then check your word-length prediction with the bit comparison in the notebook: the measured SNR is compared with the lecture formula for $b = 4, \ldots, 12$.

Listen to the audio examples with adequate resolution, poor utilization, and clipping. Relate what you hear to the plots.

Change one parameter at a time and find a case in which the histogram looks nearly uniform but the relative PSD is clearly not flat.

4. Interpret the model

The real quantizer and its mathematical model are



Because $e_q(n) = v(n) - v_q(n)$, the model must satisfy

$$v_q(n) = v(n) - e_q(n).$$

Depending on the sign convention of the error, it enters the summing junction with a plus or a minus sign; with the definition of slide 25 it enters with a minus sign, as shown above. Explain why this sign does not change the noise power.

4.1 Quantization-noise power

If e_q is modeled as uniformly distributed over one quantization interval of width Δ, then

$$P_n = \sigma_{e_q}^2 = \int_{-\Delta/2}^{\Delta/2} e^2 \frac{1}{\Delta}\, de = \frac{\Delta^2}{12}.$$

Explain, without evaluating the integral:

1. Where does the density $1/\Delta$ come from?
2. Why does halving Δ reduce P_n by a factor of four?
3. Which part of this argument becomes invalid during clipping?

4.2 SNR, word length, and utilization

In general, the signal power is

$$P_v = E\{v^2(n)\}.$$

If the input has zero mean, $P_v = \sigma_v^2$. For this case, the lecture notes, *Quantization - Parts 9 and 10* (slides 31-32), give

$$SNR/\mathrm{dB} = 10\log_{10}\left(\frac{\sigma_v^2}{\sigma_{e_q}^2}\right) = 6.02\,b + 10.8 - 20\log_{10}\left(\frac{R}{\sigma_v}\right).$$

Use this equation qualitatively to explain:

1. the gain of approximately $6\,\mathrm{dB}$ per additional bit,
2. why a signal that is too small wastes converter range, and
3. why increasing the amplitude is not always beneficial.

4.3 Assumptions and limits

Answer concisely:

> Under which assumptions may the quantization error be modeled as additive white noise, and when does this model fail?

5. Peer oral examination

Work in pairs with examination cards A and B.

- 3 min: Student A examines student B using card A.
- 3 min: Change roles; student B examines student A using card B.
- 2 min: Give feedback using the checklist on the cards.

6. Take-away

Complete and explain these two statements:

1. The additive white-noise model is useful when ...; it can fail when ...
2. More bits reduce granular quantization error, but good SNR also requires ...