ADSP - Exercise 2: Model Solution and Teaching Notes

Teaching focus and timing

The exercise is designed for explanation and diagnosis rather than calculation. Keep the reactivation short. Reactivation: about 90 s individually, then plenary; pairs only in the oral examination. During the demonstration, ask for a prediction before every preset and before the bit comparison. The central distinction is between a useful statistical approximation and the deterministic nonlinear operation that produces the actual error.

Suggested 45-minute teaching sequence (the times are not printed on the worksheet):

Phase	Time	Teaching purpose
1. Reactivation: the quantizer	0-6 min	step size, midrise characteristic, error interval
2. Predict before observing	6-10 min	predictions for the four presets and the bit comparison
3. Interactive demonstration	10-23 min	presets, bit comparison, short audio contrast
4. Interpret the model	23-34 min	noise power, SNR formula, assumptions and limits
5. Peer oral examination	34-42 min	3 + 3 min examination, 2 min feedback
6. Take-away	42-45 min	secure the two central statements

Notation as in the lecture notes, Part 2, slides 23-32: $v_\mathrm{q}(n) = Q[v(n)]$, $e_\mathrm{q}(n) = v(n) - v_\mathrm{q}(n)$, decision levels d_l, quantization levels v_l, $SNR/\mathrm{dB} = 10\log_{10}(P_v/P_n)$.

Planned and extra tasks. The worksheet is ordered by topic and does not mark this distinction, because all tasks are exam-relevant. The 45-minute plan covers the tasks in the middle column; the extra tasks are for self-study, and the answers below cover them.

Section	Planned in class	Extra (self-study, exam-relevant)
1 Reactivation: the quantizer	1.1, 1.2, 1.3, error interval	1.4 (midtread asymmetry)
2 Predict before observing	predictions for the four cases and the word length	–
3 Interactive demonstration	steps 1-4, bit comparison, audio examples	free parameter search (closing paragraph)
4 Interpret the model	4.1, 4.2, 4.3	–

1. Reactivation: the quantizer (0-6 min)

With a word length b, 2^b code words are available, so at most $L = 2^b$ levels can be represented (slide 27: $2^b \geq L$). The step size is

$$\Delta = \frac{R}{2^b} = \frac{R}{L} \qquad \text{for } L = 2^b.$$

A midtread quantizer has zero as a quantization level. A midrise quantizer has zero as a decision level and therefore no zero-valued output level. The notebook uses a symmetric midrise quantizer so that $L = 2^b$ levels fill the range symmetrically.

For $b = 2$ the midrise quantizer has $L = 4$ levels and $R = 4\Delta$:

- decision levels $d_1 = -\Delta,\ d_2 = 0,\ d_3 = \Delta$; the outer decision levels are $d_0 = -\infty$ and $d_4 = +\infty$ as on slide 24, so overloading samples are mapped to the outer levels,

- quantization levels $v_1 = -\frac{3}{2}\Delta,\ v_2 = -\frac{1}{2}\Delta,\ v_3 = \frac{1}{2}\Delta,\ v_4 = \frac{3}{2}\Delta$, each in the middle of its interval,
- non-overload range $-R/2 < v(n) \leq R/2$, i.e. $-2\Delta < v(n) \leq 2\Delta$. The range is half-open in the same way as the intervals I_l: at $v(n) = -R/2$ the error would be exactly $-\Delta/2$, which lies outside $-\Delta/2 < e_\mathrm{q}(n) \leq \Delta/2$.

The dashed diagonal $v_\mathrm{q} = v$ shows that the staircase follows the identity within $\pm\Delta/2$ inside the non-overload range and departs from it in the overload regions.

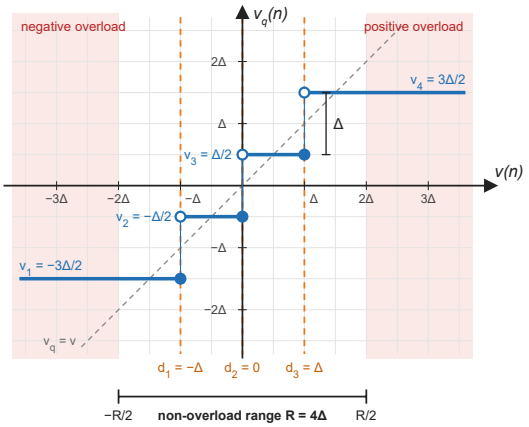

Midrise: zero is a decision level, not an output level.
Filled dot: the decision level belongs to the lower interval, $d_{l-1} < v(n) \leq d_l$.

Optional question: the 2^b code words give an even number of levels. A midtread quantizer uses one of them for zero; the remaining $2^b - 1$ levels cannot be distributed symmetrically. For $b = 3$ the levels are $-4\Delta, \ldots, 3\Delta$, which matches the two's-complement range $-4, \ldots, 3$ (slide 26).

The lecture convention is

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

For a non-overloaded sample, slide 25 gives

$$-\frac{\Delta}{2} < e_\mathrm{q}(n) \leq \frac{\Delta}{2}.$$

The endpoints follow from $I_l = \{d_{l-1} < v(n) \leq d_l\}$: a sample exactly on an inner decision level belongs to the lower interval, so $e_\mathrm{q} = +\Delta/2$. The notebook uses the same rule. During overload the bound no longer applies.

2. Predict before observing (6-10 min)

Expected predictions should distinguish noise-like, periodic, granular, and overloaded error; the word-length prediction is approximately $+6\,\mathrm{dB}$ for $b = 8 \to 9$ and $-24\,\mathrm{dB}$ for $b = 8 \to 4$.

3. Interactive demonstration: expected observations (10-23 min)

Well-utilized sinusoid at 997.3 Hz

Many quantization levels are used (236 of 256 at $b = 8$), so the error is small relative to the signal. Since 997.3 Hz has no short period at 48 kHz, over a sufficiently long record the histogram approaches a uniform distribution and the spectrum looks noise-like. It need not be perfectly white: the quantizer is deterministic, and residual lines or correlations can remain.

Expected judgment: the additive model is a useful approximation, especially for sufficiently large b, but the plots can only support, not prove, the claim. The lecture formula predicts $49.2\,\mathrm{dB}$; the notebook measures $49.3\,\mathrm{dB}$.

Very small sinusoid

The amplitude 0.003 is smaller than $\Delta/2 \approx 0.0039$. Because zero is a decision level of the midrise quantizer, the output only switches between $\pm\Delta/2$ according to the sign of $v(n)$: only two levels are visited. The error $e_\mathrm{q}(n) = v(n) - \frac{\Delta}{2}\operatorname{sgn}(v(n))$ is a deterministic function of the input. Its histogram has a gap around zero and two strongly populated regions, and its spectrum consists of discrete lines at odd harmonics of the input frequency.

Expected judgment: the error is signal-dependent and strongly correlated with $v(n)$ (correlation about -0.64); the white-noise and no-correlation assumptions clearly fail. Low utilization also gives a very poor SNR because σ_v is small while R and Δ remain fixed.

Noteworthy detail: the measured error power is still close to $\Delta^2/12$ (ratio about 0.95), and the measured SNR ($-0.3\,\mathrm{dB}$) is close to the formula value ($-0.5\,\mathrm{dB}$). A correct power estimate therefore does not show that the model assumptions hold. The derivation of $\Delta^2/12$ assumes a uniform error distribution, but here the error amplitude is clearly not uniform (see the histogram); the measured power is close to $\Delta^2/12$ only by coincidence for this amplitude. Moreover, whiteness and independence concern how the error behaves over time and how it relates to the input, and the power says nothing about either.

Periodic sinusoid with 48 samples per period

With $f = 1\,\mathrm{kHz}$ and $f_\mathrm{s} = 48\,\mathrm{kHz}$, one period contains exactly 48 samples. The same finite set of sample phases and quantizer cells is revisited every period, so the quantization error is periodic with period 48. A periodic sequence has a line spectrum (multiples of 1 kHz) rather than a flat continuous noise spectrum, and the histogram consists of at most 48 discrete values.

Expected judgment: good amplitude utilization alone does not guarantee white quantization error. The temporal pattern and the relation between input frequency and sampling frequency matter.

Overloaded sinusoid

With amplitude 1.28 and limits $\pm R/2 = \pm 1$, about 43 % of the samples are clipped to the outer output levels. Around the peaks, $|e_\mathrm{q}(n)|$ exceeds $\Delta/2$ by far. The flattened waveform produces strong deterministic distortion and harmonic components.

Expected judgment: this is overload distortion, not ordinary granular quantization noise. The formula would predict about $52\,\mathrm{dB}$; the notebook measures about $16.5\,\mathrm{dB}$. The derivation of $P_n = \Delta^2/12$ does not describe this case.

Word-length prediction

For the well-utilized sinusoid, each additional bit halves Δ and improves the SNR by about $6\,\mathrm{dB}$: $b = 8 \to 9$ gives about $+6\,\mathrm{dB}$, and $b = 8 \to 4$ gives about $-24\,\mathrm{dB}$. The notebook measures $49.3 \to 55.3\,\mathrm{dB}$ and $49.3 \to 25.1\,\mathrm{dB}$. Small deviations from exactly $6.02\,\mathrm{dB}$ per bit occur because the error of a deterministic signal is not exactly uniform.

Useful supplement for the oral exam: for a full-scale sinusoid with amplitude $R/2$, $\sigma_v = R/(2\sqrt{2})$, and the lecture formula becomes $SNR/\mathrm{dB} \approx 6.02\,b + 1.76$.

Musical signal

The generated musical signal contains several tones, attacks, decays, and changing amplitudes. At adequate resolution and utilization, its quantization error tends to be less obviously periodic than that of a single sinusoid with a short integer period. With few bits, quiet parts become audibly coarse. Poor utilization wastes levels. Clipping creates conspicuous nonlinear distortion during loud sections. In class, use the adequate-resolution case only as a brief reference and contrast poor utilization with clipping; leave the few-bits audio case for later exploration.

The teaching point is not that complex signals automatically guarantee white error. Their richer amplitude evolution often makes the statistical approximation more useful.

4. Interpret the model (23-34 min)

Real system and mathematical model

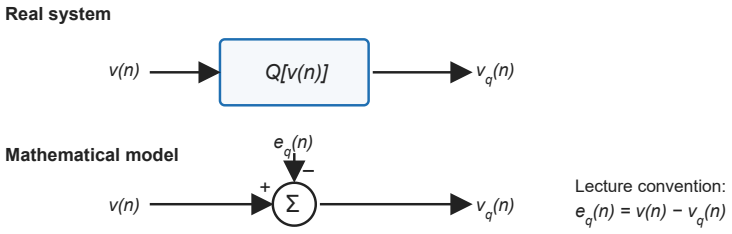

The real system is nonlinear and deterministic:

$$v_q(n) = Q[v(n)].$$

With the lecture definition $e_q(n) = v(n) - v_q(n)$, rearrangement gives

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

Therefore, a summing junction must subtract $e_q(n)$. With the alternative definition $e(n) = v_q(n) - v(n) = -e_q(n)$, the error enters with a plus sign. For the noise model the sign is irrelevant: the modeled error is zero-mean and symmetric, and $E\{e^2\} = E\{e_q^2\}$. In a derivation, however, mixing the two conventions is a sign error.

4.1 Quantization-noise power

The model assumes a uniform probability density over one cell. Using the endpoint convention of the lecture,

$$f_{e_q}(e) = \frac{1}{\Delta}, \qquad -\frac{\Delta}{2} < e \leq \frac{\Delta}{2}.$$

The value $1/\Delta$ normalizes the density: its area over an interval of width Δ is one. The endpoint assignment is irrelevant to the integral. The assumed error distribution has zero mean, so its second moment equals its variance, giving

$$P_n = \sigma_{e_q}^2 = \frac{\Delta^2}{12}.$$

The students need not evaluate the integral. They should explain the scaling. If Δ is halved, Δ^2 and thus the modeled noise power are divided by four. During clipping, the error is not confined to one quantization cell, so both the integration limits and the assumed uniform density are invalid.

4.2 SNR, word length, and utilization

Since $\Delta = R/2^b$, one additional bit halves Δ. The modeled error power then falls by a factor of four. For unchanged signal power,

$$10 \log_{10}(4) \approx 6.02\,\text{dB}.$$

This is the origin of the approximate 6 dB-per-bit rule. It is not an unconditional measured result for every deterministic input.

For an arbitrary input signal, the signal power is

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

The lecture assumes a zero-mean input, so that $P_v = \sigma_v^2$. In *Quantization - Part 10* (slide 32), σ_v is described as the RMS amplitude; for a zero-mean signal the RMS value and standard deviation coincide. With this assumption, slides 31-32 give

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

This formula also shows the role of utilization:

- If R and b remain fixed while σ_v decreases, the SNR decreases. The signal uses only a small part of the available range.
- Increasing the amplitude improves utilization only until the converter range is reached.
- Beyond that point, overload causes signal-dependent distortion, and the granular-noise formula is no longer the relevant model.

4.3 Assumptions and limits

A strong oral answer should state and connect the following ideas. Items 1, 3-5, and 7 correspond to the assumptions listed on slide 30; item 2 is the assumption of slide 31.

1. **No overload:** the samples remain within the quantizer range, so the granular error remains bounded by $\Delta/2$.
2. **Uniform error distribution:** the input visits many positions within the quantization cells sufficiently evenly. This supports $P_n = \Delta^2/12$.
3. **Stationarity:** the statistical properties of the error do not change over the observation interval.
4. **Whiteness:** error samples have negligible temporal correlation, corresponding to an approximately flat power spectrum.
5. **No input-error correlation:** the error should not track the input signal systematically.
6. **Usually sufficient resolution and signal complexity:** large word length and an input that moves through many cells make the assumptions more plausible (slide 30: "fairly well satisfied for large quantizer word lengths").
7. **Zero-mean input for the stated SNR form:** the lecture assumes $\mu_v = 0$, hence $P_v = \sigma_v^2$. Otherwise, use $P_v = E\{v^2(n)\}$.

The model can fail for very small signals, short or exactly periodic sequences, sinusoids with a short integer period, strongly correlated inputs, idle or constant inputs, and overload. In such cases the error can be deterministic, periodic, correlated with the input, nonuniform, and spectrally tonal.

Important distinctions:

- A uniform histogram concerns the marginal amplitude distribution. White noise concerns correlation across time. One property does not prove the other.
- A correct noise power does not confirm the model either: in the very-small-sinusoid case the power is close to $\Delta^2/12$, although the error is tonal and strongly correlated with the input.
- The notebook uses a Welch PSD estimate. A line-free, approximately flat-looking estimate supports the white-noise approximation but does not prove whiteness.
- Stationarity is a model assumption, not something established by the displayed full-record histogram and PSD. This is particularly important for the musical signal, whose envelope changes over time.

Optional extension

Dither can decorrelate the quantization error from the signal and suppress tonal behavior, at the price of additional noise power. It is intentionally not required in this exercise because it would introduce a further topic.

5. Peer oral examination: expected performance (34-42 min)

The student should not merely list assumptions. A strong response links each assumption to observable evidence:

- bounded error and no flattened peaks indicate absence of overload;
- a histogram tests the amplitude distribution;
- a Welch PSD estimate reveals tonal or approximately flat-looking behavior but does not prove whiteness;

- periodic error in time implies spectral lines;
- a correlation measure can support, but not by itself prove, independence.

6. Take-away (42-45 min)

1. Quantization error may be treated as additive white noise only as an approximation under no overload, approximately uniform error, weak temporal correlation, and weak input-error correlation; small, periodic, or strongly correlated signals can violate these assumptions.
2. Each additional bit halves Δ and ideally improves the SNR by about $6\,\mathrm{dB}$, but poor utilization wastes resolution and excessive amplitude causes clipping rather than improved SNR.