

Topic 11

University Electrical Engineering

What this course is, and how to use it

This is the eleventh course in University Electrical Engineering: **communication systems**. EE06 built the signal tools; this course sends signals across real channels – modulated onto carriers, quantised into bits, protected by codes.


- Units 11.1–11.3 are analogue: spectra, AM, FM.
- Units 11.4–11.6 are the digital chain: noise and SNR, PCM, modulation.
- Unit 11.7 sets the limits: entropy, capacity, coding.

The running story is one radio link: a microphone, a transmitter, a noisy channel, a receiver, and a data modem sharing the same wire. Every unit below moves one block of that link forward, and every number is checked the way a communications engineer checks it – every dB claim recomputed as a ratio, every power budget closed (carrier + sidebands = total), and every bandwidth claim tested against the channel it must fit.

Signals, spectra and bandwidth

Why it matters. A radio regulator sells spectrum by the kilohertz; whether a signal fits its licence is decided by its spectrum, and the spectrum is decided by the waveform's shape.

The idea, with a picture. A real signal occupies spectrum. The square wave is odd harmonics at $4/\pi \cdot 1/k$ – the 3rd harmonic at $1/3$ amplitude, the 5th at $1/5$ – so **bandwidth** is a design choice about how much of the tail you keep. For modulated signals the **sideband** structure rules: mixing shifts the spectrum, it never creates new shape, so a baseband signal of width W on a carrier occupies $2W$ of radio spectrum.



Left: the tail you keep is the bandwidth you pay for. Right: modulation doubles the occupancy

The derivation. A Fourier series writes the square wave as $\frac{4}{\pi} \sum_{k \text{ odd}} \frac{1}{k} \sin(k\omega_0 t)$, so a low-pass at B keeps the harmonics with $k f_0 \leq B$, and the power fraction kept follows from Parseval: each harmonic's power is its amplitude squared, so the fundamental alone carries $1.273^2 / (1.273^2 + 0.424^2 + 0.255^2 + \dots) \approx 81\%$.

Worked example 1 – a 1 kHz square through a 5 kHz low-pass. The filter passes the 1st harmonic (full amplitude $4/\pi = 1.273$) and the 3rd ($1.273/3 = 0.424$), blocks the 5th (0.255) and beyond. Check the power kept: $(1.273^2 + 0.424^2) / (1.273^2 + 0.424^2 + 0.255^2 + \dots) \approx 90\%$ – the ear hears the rounding, the meter barely moves, which is why the 300–3400 Hz telephone channel passes speech intelligibly at all.

Worked example 2 – escalate to the radio licence. A 4 kHz music signal is placed on a 100 kHz carrier by mixing: the spectrum becomes the carrier plus mirror sidebands spanning 96 kHz to 104 kHz – $2W = 8$ kHz of occupied bandwidth. If the licence grants 8 kHz, amplitude modulation exactly fills it; grant only 4 kHz and you must transmit one sideband (SSB, unit 11.2's theme). Check: no mix of filters can compress the occupancy below W without deleting signal – the $2W$ law is geometry, not electronics.


Common mistakes. Quoting “bandwidth” without a definition (–3 dB? 98% power? first null? – state which). Believing a filter can narrow a modulated signal below its message width. And forgetting that harmonics of a square wave fall as $1/k$ in AMPLITUDE but $1/k^2$ in POWER – the tail is weaker than it looks.

Summary. Spectra come from shapes; a low-pass tail is a design choice measured in power fraction; modulation shifts and doubles. $2W$ is the occupancy law every licence is written against.

Amplitude modulation

Why it matters. AM is the simplest carrier system that still closes its power budget on paper – and the arithmetic of that budget explains why every broadcast engineer moved on, and why the envelope detector still ships in every radio.

The idea, with a picture. In **amplitude modulation** the carrier's envelope carries the message: $s(t) = A[1 + m \cdot x(t)] \cos(\omega_c t)$ with **modulation index** $m \leq 1$. The spectrum is the carrier plus two sidebands at $f_c \pm f_m$, and the power budget is exact: sideband power = $(m^2/2)P_c$, total $P_c(1 + m^2/2)$.



The envelope is the message; the budget bars show the carrier taking most of the power

The derivation. With a single-tone message $x = \cos(\omega_m t)$: $s(t) = A \cos(\omega_c t) + \frac{mA}{2} \cos((\omega_c + \omega_m)t) + \frac{mA}{2} \cos((\omega_c - \omega_m)t)$ – write the product of cosines as a sum and the two sidebands appear with amplitude $mA/2$, hence power $(mA/2)^2$ each and $\frac{m^2}{2}P_c$ together.

Worked example 1 – the budget at $m = 0.5$. A 100 W carrier: sidebands total $\frac{0.25}{2} \times 100 = 12.5$ W, transmission 112.5 W, efficiency $12.5/112.5 = 11.1\%$ ✓. At the legal maximum $m = 1$: 50 W in sidebands, 150 W total, 33.3% – two thirds of the transmitter power goes into a carrier that carries no information. Both budgets close on $P_c(1 + m^2/2)$.

Worked example 2 – escalate to the detector design. The **envelope detector**

(a diode and an RC) must discharge slowly enough to follow the envelope's fastest dips but fast enough not to bridge carrier peaks: $1/\omega_c \ll \tau \ll 1/\omega_m$. For $f_c = 100$ kHz and $f_m = 5$ kHz, τ must sit between $1.6 \mu\text{s}$ and $31.8 \mu\text{s}$; $\tau = 10 \mu\text{s}$ works with margin at both ends ✓. Push m toward 1 and the upper bound tightens to $\frac{\sqrt{1-m^2}}{m\omega_m}$ – at $m = 0.9$ it is $\frac{\sqrt{0.19}}{0.9} \times \frac{1}{2\pi \cdot 5 \text{ kHz}} = 15.4 \mu\text{s}$: the same RC, less than half the slack, which is why deep modulation distorts before anything else fails.

Common mistakes. Setting $m > 1$ (envelope folds, phase reverses, detector produces garbage). Computing efficiency against total power one way and against carrier power another (state the denominator). And choosing τ by the carrier alone: the carrier sets only the lower bound; the message's highest frequency and its depth m set the upper one.


Summary. AM: envelope = message; sidebands $m^2/2 \times P_c$; efficiency peaks at 33%; the detector's RC lives between $1/\omega_c$ and $\sqrt{1-m^2}/(m\omega_m)$.

Angle modulation (FM/PM)

Why it matters. FM broadcasts through thunderstorms because it trades bandwidth for noise immunity, and the size of that trade – Carson's rule on one side, the SNR gain on the other – is the course's cleanest engineering bargain.

The idea, with a picture. **Frequency modulation** puts the message in the **instantaneous frequency** : $f_i = f_c + k_f x(t)$. The **frequency deviation**

Δf against the highest message frequency f_m gives the **modulation index** $\beta = \Delta f/f_m$, and **Carson's rule** $B \approx 2(\Delta f + f_m)$ bounds the meaningful spectrum.



The message bends the instantaneous frequency; Carson's rule fences the spectral spread

The derivation. Integrating f_i gives the phase $\phi(t) = 2\pi f_c t + \beta \sin(\omega_m t)$; expanding $A \cos(\phi)$ with the Jacobi–Anger identity produces spectral lines at $f_c \pm k f_m$ with Bessel weights $J_k(\beta)$ – significant lines stop near $k \approx \beta + 1$, which is exactly Carson’s $2(\Delta f + f_m)$. Carson’s rule estimates a bandwidth containing roughly 98% of the power in the usual FM model; it is not a strict spectral cutoff.

Worked example 1 – FM broadcast. $\Delta f = 75$ kHz, $f_m = 15$ kHz: $\beta = 5$, $B = 2 \times 90 = 180$ kHz – and the 200 kHz channel spacing leaves a 20 kHz guard band ✓. Narrowband ($\beta \ll 1$): B collapses to about $2f_m$, the walkie-talkie regime. Check the limit: $\Delta f \rightarrow 0$ gives $\beta \rightarrow 0$, $J_0 \rightarrow 1$ – the carrier alone, as it must.

Worked example 2 – escalate to the bargain’s price. Wideband FM buys post-detection SNR: $(S/N)_{out} \approx (3\beta^2/2)(S/N)_{ref}$ in the high-SNR regime, where the reference $(S/N)_{ref}$ is the received power over the noise in the message band f_m (what baseband transmission at the same power would give). For $\beta = 5$: $10 \log_{10}(37.5) = 15.7$ dB of improvement – paid for with 180 kHz of bandwidth (6× the 30 kHz that AM would occupy for the same 15 kHz audio) and lost entirely below the FM threshold near a carrier-to-noise ratio of about 10 dB in the 180 kHz channel, where click noise and capture take over. The ledger: 15.7 dB of gain for 7.8 dB of spectrum.


Common mistakes. Using Carson’s rule as if it were exact (it is a 98%–power bound). Mixing up Δf (a property of the modulator) with B (a property of the spectrum). And forgetting PM: the same machinery with the message in the phase, $\beta_{PM} = k_p A_m$ – frequency-independent, which is why FM won broadcasting while PM lives inside data modems.

Summary. $\beta = \Delta f / f_m$; $B \approx 2(\Delta f + f_m)$; the wideband SNR gain is $3\beta^2/2$, real only above threshold.

Noise and SNR

Why it matters. The link budget of the running story ends at a noise floor that no amplifier can dig below – the thermal floor of the universe – and every receiver stage decides how far above it the signal arrives.

The idea, with a picture. Every receiver adds noise. The **noise figure** F measures how much, and Friis’s cascade law $F = F_1 + (F_2 - 1)/G_1$ says the FIRST stage decides everything: its gain divides every later stage’s contribution.



Same blocks, swapped order, 2.5 dB different – the front-end gain IS the protection

The derivation. Per-stage input-referred noise adds at the input of each stage, and each earlier gain scales the later noise down by its gain, so the cascade’s total referred noise is $F_1 + (F_2 - 1)/G_1 + (F_3 - 1)/(G_1 G_2) + \dots$ – after two stages of healthy gain the tail terms vanish. Thermal floor: $-174 \text{ dBm/Hz} + 10 \log_{10} B$.

Worked example 1 – the LNA wins. LNA with $F_1 = 1.76$ dB ($F = 1.5$) and $G_1 = 10$ dB ahead of a mixer with $F_2 = 4.77$ dB ($F = 3$): the chain computes $F = 1.5 + (3 - 1)/10 = 1.70 = 2.30$ dB ✓ – the mixer’s 4.77 dB collapsed to a 0.5 dB penalty. Swap the blocks: $3 + (1.5 - 1)/10 = 3.05 = 4.84$ dB – 2.5 dB worse for the same parts. In a 180 kHz FM channel the floor sits at $-174 + 10 \log_{10}(180\,000) = -121.4$ dBm.

Worked example 2 – escalate to the link budget. The transmitter delivers -20 dBm at the receiver’s antenna after path loss; the channel is 180 kHz wide. Signal -20 dBm against noise -121.4 dBm is an input $S/N = 101.4$ dB – enormous. The receiver’s own noise raises the floor by its 2.3 dB noise figure to -119.1 dBm, so the carrier-to-noise ratio after the front end is $-20 - (-119.1) = 99.1$ dB, far above the FM threshold. The FM gain of unit 11.3 is quoted against the message-band reference: noise in the 15 kHz audio band is $-174 + 10 \log_{10}(15\,000) + 2.3 = -129.9$ dBm, so $(S/N)_{ref} = 109.9$ dB, and the demodulator’s 15.7 dB lifts the post-detection S/N to about 125.6 dB in the linear regime – far above hi-fi; the practical limit becomes stereo processing, not noise. One budget, three units, closed.


Common mistakes. Adding noise figures in dB (convert to linear F first, add ratios, convert back once). Forgetting the floor scales with bandwidth – 10 dB per decade of B . And quoting a cascade F for a mismatched reference temperature or impedance; F is defined at 290 K.

Summary. $F_{total} = F_1 + (F_2 - 1)/G_1 + \dots$; first stage rules; floor = $-174 \text{ dBm/Hz} + 10 \log_{10} B$; every budget closes in dB arithmetic once the ratios are linear.

PCM and baseband transmission

Why it matters. The running link goes digital in this unit: the microphone's voltage becomes bits, the bits become pulses, and the pulse shape decides how fast the wire can carry them.

The idea, with a picture. Quantisation trades precision for bits: n bits give 2^n levels and $SQNR \approx 6.02n + 1.76$ dB. Sampling at $f_s > 2W$ (Nyquist) then encoding is **PCM** – the whole telephone network's foundation. At baseband, the **raised-cosine** filter with roll-off α trades bandwidth for timing slack: $B = (1 + \alpha)R_s/2$.



Sampling freezes time, quantisation freezes amplitude, shaping fences the spectrum

The derivation. A uniform quantiser with step Δ rounds each sample to the nearest level; the error is uniform on $[-\Delta/2, \Delta/2]$ with variance $\Delta^2/12$. Against a full-scale sine of power $(2^n \Delta)^2/8$, the ratio is $1.5 \cdot 2^{2n} - 10 \log_{10}$ of which is $6.02n + 1.76$ dB. The 6.02 is one bit = one doubling = $10 \log_{10} 4$.

Worked example 1 – the telephone number. 4 kHz voice sampled at 8 kHz with 8 bits per sample: $8000 \times 8 = 64$ kbps; $SQNR = 6.02 \times 8 + 1.76 = 49.9$ dB ✓ (256 levels). Every 1 dB of extra SQNR costs 1/6 of a bit per sample – the 8-bit choice is the toll-quality settlement of that price.

Worked example 2 – escalate to the wire. At $R_s = 1$ Msym/s with $\alpha = 0.35$: $B = 1.35 \times 1/2 = 0.675$ MHz ✓ – the number every modem datasheet quotes. Cut the roll-off to $\alpha = 0$ and the spectrum touches the $R_s/2$ ideal but the pulse's zero-crossings become knife-edge timing. Manchester coding doubles the bandwidth to buy a strong clock (the classic Ethernet choice): its first spectral null sits at $2R$ against NRZ's R , and NRZ has a DC term and no clock. Each line code is a point on the bandwidth-versus-clock plane.

Common mistakes. Sampling exactly at $2W$ (the inequality is strict; at equality the


guard band is zero and any filter ripple aliases). Confusing $SQNR$ with the analog SNR of the whole chain. And forgetting $\alpha = 0$ is unbuildable – it is the bound that real filters approach.

Summary. $SQNR = 6.02n + 1.76$ dB; Nyquist $f_s > 2W$; raised-cosine $B = (1 + \alpha)R_s/2$; line codes trade spectrum for clock strength.

Digital modulation

Why it matters. The link's bits must ride a carrier, and the constellation you draw decides how many of them fit per hertz and how easily noise knocks them off their points.

The idea, with a picture. Bits ride carriers as constellations : ASK moves amplitude, PSK rotates phase, QAM does both. More points per symbol = more bits/Hz and more sensitivity to noise – the two-dimensional trade is the whole story.



Equal minimum distance d : 16-QAM packs $2\times$ the bits per symbol for $5\times$ the average energy

The derivation. For equal minimum distance d , the average symbol energy is the mean of the squared radii, $E_{avg} = \overline{r^2}$, computed directly per constellation: QPSK's four points sit on a circle of radius $r = d/\sqrt{2}$, so $E_{avg} = r^2 = d^2/2$; 16-QAM's grid (4 inner points, 8 edge points, 4 corner points) averages the squared radii to $E_{avg} = (4 \cdot 0.5 + 8 \cdot 2.5 + 4 \cdot 4.5) d^2/16 = 2.5d^2$ – a factor 5, i.e. 7.0 dB, for $2\times$ the bits per symbol.

Worked example 1 – the energy ledger. QPSK: 2 bits/symbol at $E = 0.5d^2$. 16-QAM: 4 bits/symbol at $2.5d^2$ – five times the symbol energy, 7.0 dB. Per bit the ledger softens: $0.25d^2$ against $0.625d^2$ – factor $2.5 = 4.0$ dB, because 16-QAM carries two bits per symbol ✓. Check with the datasheet number: the classic 16-QAM-vs-QPSK gap in E_b/N_0 at 10^{-6} BER is quoted near 4 dB – the per-bit geometry, measured.

Worked example 2 – escalate to the cliff. Under AWGN, BPSK's BER is $Q(\sqrt{2E_b/N_0})$. At $E_b/N_0 = 5$ dB: $\sqrt{2} \times 3.16 = 2.51$, $Q(2.51) \approx 6 \times 10^{-3}$. At 9 dB ($\sqrt{2} \times 7.94 = 3.98$): $Q(3.98) \approx 3 \times 10^{-5}$. At 12 dB: $\approx 10^{-8}$. Three points, one cliff – 4 dB buys two orders of magnitude and the next 3 dB three more, which is why link budgets obsess over the last dB.


Common mistakes. Comparing constellations at equal average energy when the fair comparison is equal minimum distance (or equal peak). Reading E_b/N_0 curves without checking the normalisation. And bits-per-symbol gains scale as $\log_2 M$ while energy grows far faster in M – the cliff gets steeper.

Summary. Equal- d ledger: QPSK $d^2/2$, 16-QAM $2.5d^2$ (7 dB, $2\times$ bits). BPSK BER $Q(\sqrt{2E_b/N_0})$; the curve is a cliff – budget the dB before the link goes out.

Information and coding introduction

Why it matters. The last block draws the map's edges: bits a source deserves, bits a channel allows, and how codes move practice toward theory.

The idea, with a picture. Entropy $H = -\sum p \log_2 p$ measures a source's surprise in bits/symbol; the gap to fixed-length coding is redundancy. Huffman removes it for the source; **channel codes** spend redundancy on protection instead. The ceiling is **channel capacity** $C = B \log_2(1 + S/N)$.



Left: optimal lengths average exactly the entropy for dyadic probabilities. Right: capacity, approached but never beaten

The derivation. For a dyadic source (p 's are negative powers of 2), Huffman's greedy merge produces lengths $\log_2(1/p)$ exactly, so the average equals H – met, not approached; non-dyadic sources sit within one bit.

Worked example 1 – the quartile source. Probabilities $1/2, 1/4, 1/8, 1/8$: the four terms $p \log_2(1/p)$ are 0.5, 0.5, 0.375, 0.375, so $H = 1.75$ bits ✓ against 2 bits fixed-length – 87.5% efficient; Huffman lengths 1, 2, 3, 3 average exactly 1.75 ✓ – dyadic, bound met. **Hamming code** (7,4): 4 data + 3 parity, distance 3, corrects any single error: encode 1011 → 0110011; flip bit 5, the syndrome computes to $101_2 = 5$ ✓. Rate $4/7$ (43% redundancy).

Worked example 2 – escalate to the modem ceiling. A 3 kHz channel at 30 dB SNR: $C = 3000 \times \log_2(1001) = 29.9$ kbps – the dial-up ceiling. V.90/V.92 hit 56 kbps downstream only by going digital to the exchange – inside the law. Check the direction: 10 dB more SNR buys $\log_2(10) \approx 3.3$ bits/Hz – linear in bandwidth, logarithmic in power.

Common mistakes. Computing H with the wrong log base (bits = $\log_2$). Believing Huffman always meets H (only dyadic). And reading C as a guarantee – codes approach it, none reach it.

Summary. $H = \sum p \log_2(1/p)$, met by Huffman on dyadic sources; Hamming (7,4) fixes one error at 43% redundancy; $C = B \log_2(1 + S/N)$ is the ceiling.