

COMS3100/7100

Introduction to Communications

Lecture 28: *Noise, Errors & Synchronisation*

This lecture:

1. Noise.
2. Bit-error rate.
3. Synchronisation.
4. Case Study: 100Base-TX Ethernet.

Ref: CCR pp. 371–372, 382–383, 448–461, 476–485, Kaplan & Nosenworthy, *The Ethernet Evolution from 10 Meg to 10 Gig: How It All Works!*.

Noise

Noise is any unwanted electrical signal encountered at the receiver, especially one that is so hard to predict that it can be deemed random.

- ▶ The signal at the receiver contains both wanted and unwanted parts.
- ▶ The wanted part, bearing information from the transmitter, is often simply called the *signal* (but sometimes confusing, since, technically, the noise is also a signal!).
- ▶ Sometimes the unwanted part is further classified into *noise* and *interference* components.
 - ▶ The boundary is blurred: interference usually refers to man-made noise (especially from other communication systems).

Thermal Noise

One unavoidable cause of electrical noise is the thermal motion of electrons in conducting media—wires, resistors and so on—termed *thermal noise*.

- ▶ When thermal noise is sampled, it can be regarded as a random variable (r.v.).
 - ▶ Typically assumed to be a Gaussian r.v. with zero mean.
- ▶ Noise in communication systems is usually *additive* in the sense that

$\text{received signal} = \text{desired signal} + \text{noise (+interference)}.$


Signal-to-Noise Ratio

Because signal and noise are uncorrelated, their powers add also.

- ▶ An important statistic is the *signal-to-noise ratio (SNR)* of signal and noise power.

Bit-Error Rate

- ▶ We will assume (for now) that the channel is distortionless, gain is equalised at the receiver, and we use a binary line code.
- ▶ A LPF removes excess noise without introducing ISI.




- ▶ In this model, a sample-and-hold unit samples the received noisy process $y(t)$ at times $t_n = nD$.

Bit-Error Rate

- ▶ With $n(t)$ representing additive noise, we have

$$y(t_n) = a[n] + n(t_n).$$

- ▶ We want to set a threshold v to minimise the probability of making a *bit error*.



Bit-Error Rate

Probability of error

- ▶ Let E_0 be the event that a 0 was transmitted in symbol n and let E_1 be the event that a 1 was transmitted.
- ▶ Let D_0 be the event that a 0 is detected: occurs when $y(t_n) \leq v$.
- ▶ Let D_1 be the event that a 1 was detected: $y(t_n) > v$.
- ▶ The *probability of bit error* on symbol n is

$$\begin{aligned}P_e &= \Pr\{(D_1 \cap E_0) \cup (D_0 \cap E_1)\} = \Pr\{D_1 \cap E_0\} + \Pr\{D_0 \cap E_1\} \\ &= \frac{1}{2} \Pr\{D_1|E_0\} + \frac{1}{2} \Pr\{D_0|E_1\} \\ &= \frac{1}{2} - \frac{1}{2}P_{y(t_n)|E_0}(v) + \frac{1}{2}P_{y(t_n)|E_1}(v),\end{aligned}$$

an expression stated in terms of conditional c.d.f.s $P_{y(t_n)|E_i}(\cdot)$.

Probability of Error (BER)

- ▶ Note that P_e is a function of v so to maximise it, the optimal threshold v_{opt} must satisfy

$$\left. \frac{dP_e}{dv} \right|_{v=v_{\text{opt}}} = 0 \quad \Rightarrow \quad p_{Y(t_n)|E_0}(v_{\text{opt}}) = p_{Y(t_n)|E_1}(v_{\text{opt}}).$$

- ▶ This implies that the optimal threshold is a point where the two conditional p.d.f.s intersect.
- ▶ If we assume the events are all independent and identical from one symbol to the next, P_e represents the *bit-error rate (BER)*, the average rate (relative to the bit rate) at which bit errors occur.

Example: Unipolar NRZ with additive Gaussian noise

For a random variable (r.v.) Ξ from a *standard normal* distribution, *i.e.*, a Gaussian r.v. with mean 0 and variance 1, we define

$$Q(\xi) = 1 - P_{\Xi}(\xi) = \frac{1}{\sqrt{2\pi}} \int_{\xi}^{\infty} e^{-x^2/2} dx.$$

- ▶ A common assumption is that the noise $n(t_n)$ is a Gaussian r.v. with zero mean and variance σ^2 .
- ▶ It follows that, under the condition E_0 , $y(t_n)$ is a Gaussian r.v. with zero mean and variance σ^2 .

Example: Unipolar NRZ with AWGN


- ▶ Given $E_{1,y}(t_n)$ is Gaussian with mean A and variance σ^2 .
- ▶ It can be shown that $v_{\text{opt}} = \frac{1}{2}A$ and that the resulting BER is

$$P_e = Q\left(\frac{A}{2\sigma}\right).$$

- ▶ It can be shown that the received signal power at the receiver is $\frac{1}{2}A^2$ and noise power is σ^2 , and so

$$P_e = Q\left(\sqrt{\frac{1}{2}\text{SNR}}\right).$$

For unipolar NRZ with additive Gaussian noise, it is possible to plot BER vs. SNR.




- ▶ An 'acceptable' BER depends on the application.
- ▶ However, typically, acceptable BER is between 10^{-4} and 10^{-7} .
- ▶ This implies the SNR needs to be between 14 dB and 18 dB.
- ▶ This needs to be accounted for when preparing a link budget.

Synchronisation



Observe that, in the discussion on bit-error rates, the *regenerator* or *detector* requires a 'sync' signal.

- ▶ It's usually too expensive to transmit a clock signal separately.
- ⇒ Digital comms signals are designed to be *self-synchronising*.



- ▶ The receiver architecture above incorporates additional block to recover *bit* or *symbol timing*.
- ▶ Bits and symbols are usually arranged hierarchically into *frames*, for which synchronisation must also be maintained.

Bit synchronisation usually involves a non-linear element such as a *square-law device*.




Case Study: Fast Ethernet

Early History of Ethernet

Ethernet is a standard for computer networking in *local area networks (LANs)*.

- ▶ It was invented in 1973 by Robert Metcalfe at Xerox PARC in California.



- ▶ Named to evoke the superseded idea of a *luminiferous ether* through which electromagnetic waves propagate.
- ▶ Originally designed to carry 2.94 Mbps over a thick coaxial cable shared between multiple computers.

Standardisation of Ethernet

IEEE created the 802.3 standard to govern ethernet in 1985.

- ▶ The standard now covers ethernet over a number of media, most notably twisted pair (copper) and optical fibre.
- ▶ Speeds up to 10 Gbps.
- ▶ The standard covers many aspects of how ethernet is set up and deployed, far beyond the signalling to be used on the wire (the *physical layer*).

Structured Cabling

An important aspect of the ethernet standards is its relationship with *structured cabling*.

- ▶ *Structured cabling* is a set of standards for telecommunications cabling in commercial buildings.
- ▶ They define standards for cables such as *unshielded twisted pair (UTP)* and connecting hardware.
- ▶ In particular, they define minimum bandwidths (maximum attenuations) over specified distances.

CAT-5 (and other standards)

Standard	Max Freq	Used in
Cat 3	16 MHz	10Base-T
Cat 5	100 Mhz	100Base-TX [2 pair]
Cat 5e	100 Mhz	1000Base-T [4 pair]
Cat 6	200 MHz	1000Base-T,
Cat 7	1000 MHz	10Ge, 40Ge

- ▶ *Category 5 (CAT-5)* cable has a usable passband of 100 MHz over 100 m.

100Base-TX Example Continued

- ▶ *100Base-TX*, a.k.a. *Fast Ethernet*, is defined with respect to CAT-5 cable.
- ▶ CAT-5 cable contains 4 pairs of UTP.
- ▶ 100Base-TX uses two pairs to achieve full duplex.
 - ▶ Other pairs can be used for other services, such as phone.




Image credit: Wikipedia/Ryan Finnie/Richard Wheeler.

Line code used in 100Base-TX


With 100 MHz bandwidth, we should be able to achieve 200 Mbps with binary signalling (Nyquist consideration).

- ▶ But that limit requires more expensive signal processing: too expensive in 1995 when the standard was developed.

NOTE: with 20dB of SNR (typical on ethernet networks), 100 MHz of bandwidth according to Shannon's limit is ...

MLT-3

- ▶ 100Base-TX uses *MLT-3 encoding* (multi-level transmit, 3 levels).
- ▶ This is a pseudo-ternary line code specially chosen to minimise high-frequency components (without resorting to pulse shaping).
- ▶ To transmit a space, the line level remains unchanged.
- ▶ To transmit a mark, the line level changes according to the pattern $+A \rightarrow 0 \rightarrow -A \rightarrow 0 \rightarrow +A \rightarrow \dots$



4B/5B

In MLT-3, a long sequence of spaces results in a steady output level.

- ▶ This is bad for synchronisation and creates problems if there is any DC block.
- ▶ Therefore, *4B/5B* is used to ensure regular transitions.
- ▶ In *4B/5B*, groups of four bits are translated into groups of five *before* begin MLT-3 encoded.
- ▶ Each group of five bits contains at least two marks (and thus two level transitions).
- ▶ Hence, to transmit 100 Mbps of data, we use MLT-3 at 125 Mbps.

Scrambling

If the bitstream contains a long run of marks or spaces then, even with 4B/5B, the output becomes periodic.

- ▶ The output spectrum then becomes a line spectrum, with power concentrated at specific frequencies.
- ▶ This can cause interference with other communications devices.
- ▶ Therefore, a *scrambler* is used to guard against long runs.
- ▶ A *scrambler* is a simple device, built from a shift register, that pseudo-randomly flips bits.
- ▶ A *descrambler* is similarly simple to build and ‘unflips’ the bits at the receiver.


100Base-T

All modern desktops are sold with 1 Gbit ethernet interfaces.

Improvements over 100Base-TX include:

- ▶ Signaling is PAM-5 (5 level PAM, as described in previous lecture)
- ▶ Uses 4 twisted pairs, and transmits and receivers on each.

Comparison with ADSL



ADSL2 uses 2 MHz of bandwidth.

ADSL1 uses 1 MHz of bandwidth.