Telephones and digitised audio

- Telephone signals are generally "band limited" to $300 < f < 3,400$ Hz.
- In traditional telephony each of these 3,100 Hz channels was moved in frequency by amplitude modulation and filtering to occupy 4 kHz-wide channels, at 4 kHz intervals, with a hierarchy of "groups", "supergroups", "jumbo super groups" etc.
- In modern telephony, the voice channel, still limited to $300 < f < 3,400$ Hz, is sampled at 8,000 samples per second to 8-bit resolution to produce a 64,000 bps digital data stream.
- The analogue signal is pre-distorted to give more bits over low amplitudes – use μ-law in North America and A-law in Europe.

- But why 8,000 samples per second?

Nyquist sampling theorem

- A periodic signal, limited to a bandwidth of W Hz, is totally defined by $2W$ samples per second.

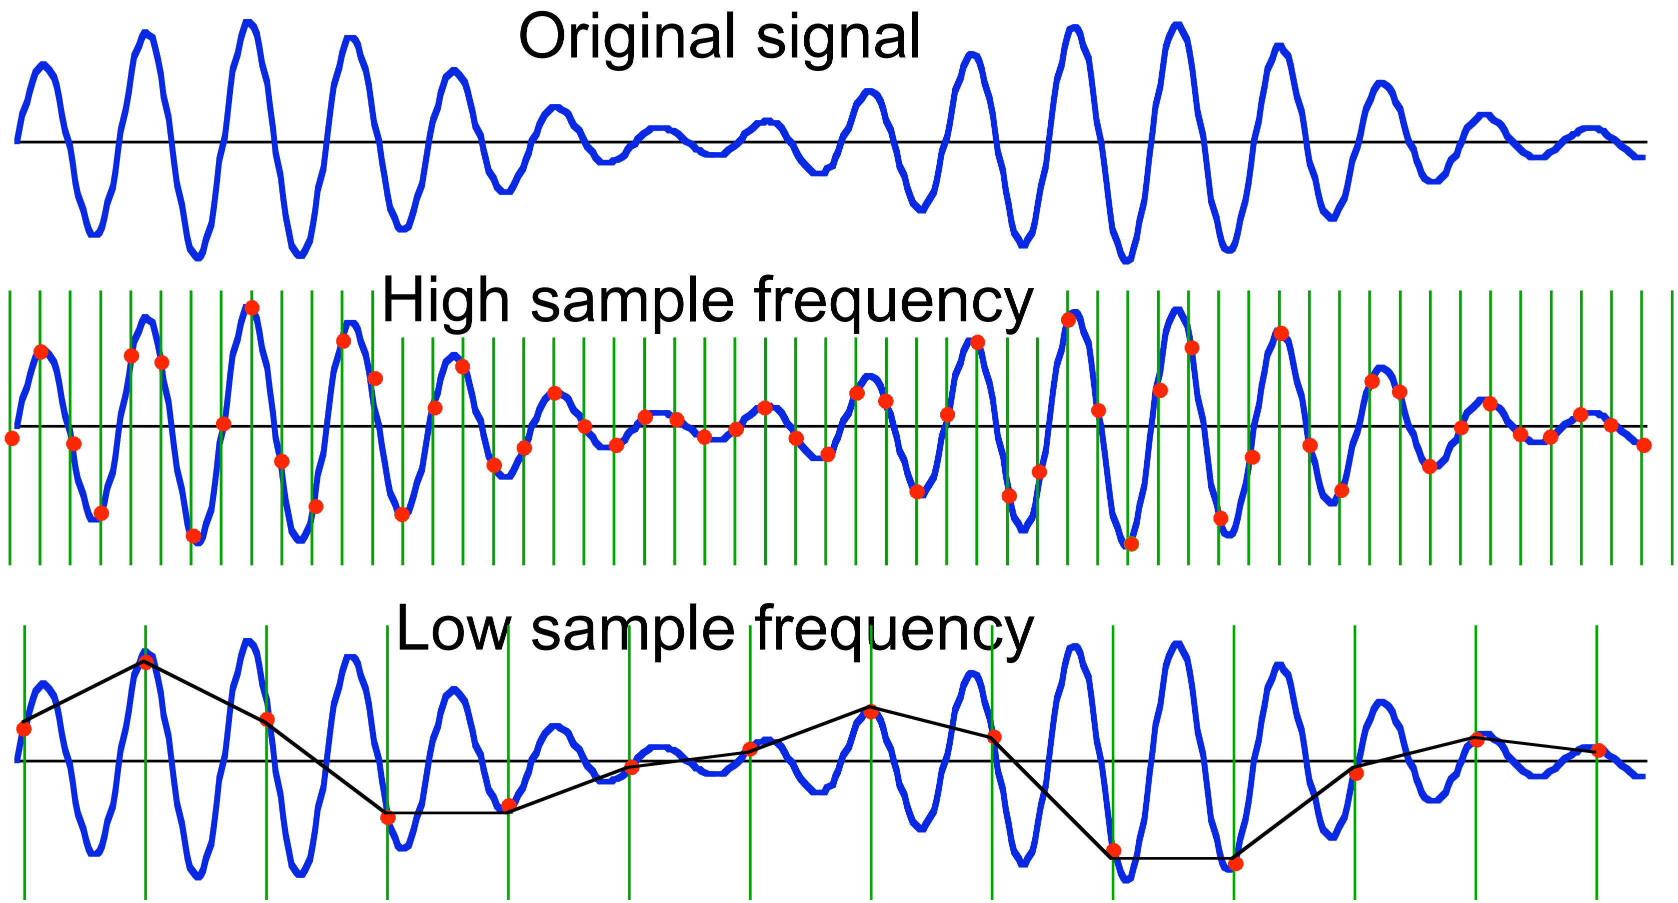

- The low sample rate is not a good representation

- If the sample rate is R, then a frequency $f = R/2 + \mathrm{d}$ is "folded" or "aliased" into $f' = R/2 - \mathrm{d}$ and it is that aliased frequency which appears to be sampled.
- Telephones use a bandwidth W = 3,100 Hz, and sample rate R = 8,000 Hz with 8 bits per sample.
- Audio CDs use a sample rate of 44,200 Hz, corresponding to a maximum bandwidth of 22.1 kHz. With 2 stereo channels, and 16 bits for each sample, the actual data rate is about 1.4 Mbps. (The rate on disk is about twice that.)
- Older CDs sometimes had trouble from aliasing (sounds like cymbals could sound very hollow and empty). Later CD recordings sample much faster than 44 kHz and then use digital filtering to minimise aliasing effects.
- "Oversampling" takes the digital data from the CD and uses a complex interpolation to predict sampling at a faster rate; this predicted sample is used to regenerate the audio.
- A television signal, W = 6 MHz, needs a sample rate R > 12 MHz.

Alternate look at sampling

- A conventional way of representing a modulated signal and its sidebands shows how the baseband signal is displaced by the carrier frequency, with the upper sideband resembling the baseband signal and the lower sideband inverted in frequency

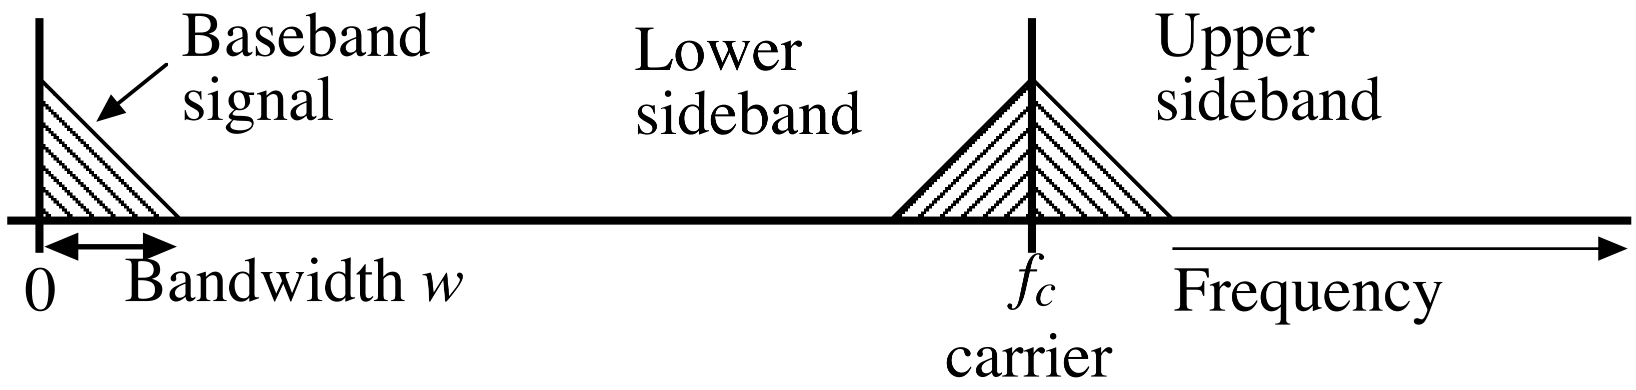

- Sampling can be looked at as a form of modulation where the “carrier” is a regular train of very narrow pulses at the sampling frequency.
- Such a pulse train at frequency *f* has a “comb” spectrum at frequencies *f*, 2*f*, 3*f*, 4*f*, … up to some very high frequency.
- Each “tooth” of the comb is surrounded by a copy of the sidebands.

- With a high sampling rate (short interval *t* between samples) the "teeth" are widely spaced and the sets of sidebands do not interfere. The samples accurately reflect the baseband signal

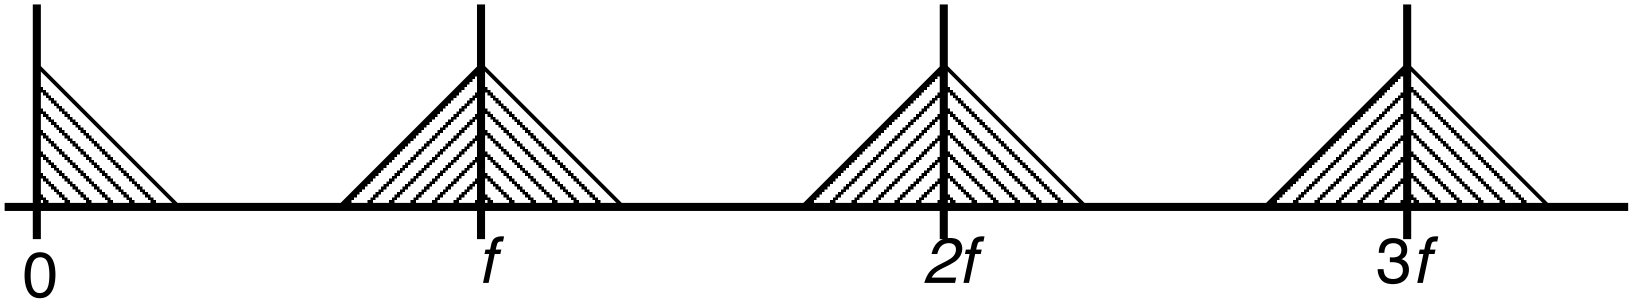

- With a very low sampling frequency the teeth are more closely-spaced and the sidebands interfere. The samples are ambiguous, with a mixture of components from the correct signal and the adjacent sidebands. The data cannot be recovered.

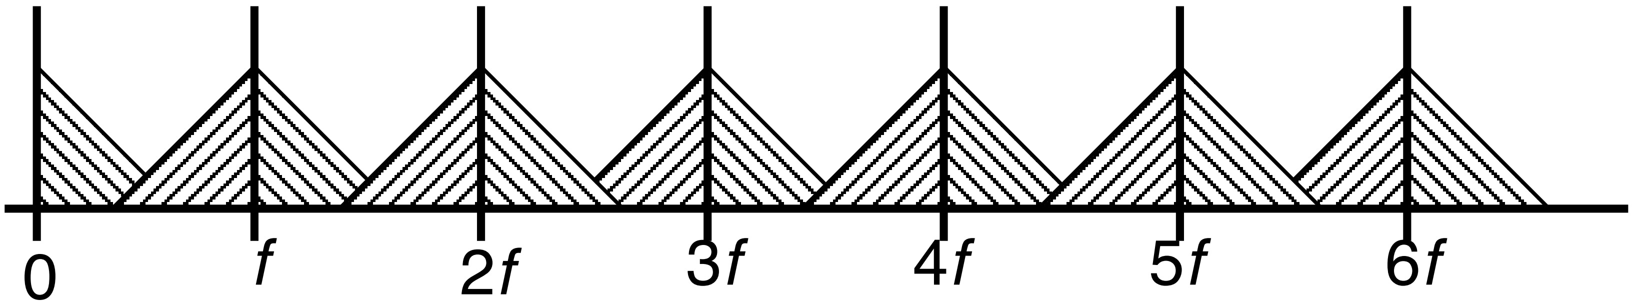

- Clearly, if the sampling frequency $f \geq 2w$ (w is the signal bandwidth) there is no interference; the sampling is unambiguous and the original signal can be reconstructed.

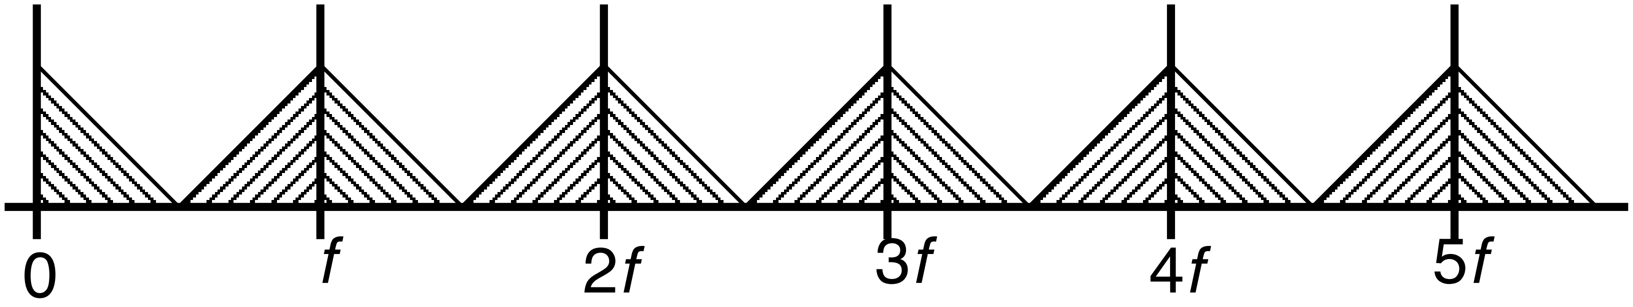

- The minimum sampling rate $f = 2w$, is known as the *Nyquist rate*, after the engineer who first investigated it.

Digital communications

- Digital communication standards are based on telephone standards and a 64 kbps channel (8 bits per sample, at 8,000 samples per second).
- There are two groups of standards, one at around 1–2 Mbps and the other above 50 Mbps.
- The low data rate signals are used for ISDN (Integrated Services Digital Network), and the high rates for B-ISDN (Broadband ISDN).

The North American T1 or DS-1 standard handles 24 telephone channels, each 64 kbps. Each frame has an 8-bit sample for each of the 24 channels, with 1 framing bit ending the frame (193 bits per frame). With 8000 frames per second the bit rate is 1.544 Mbps. The framing bit "flips" on alternate frames; for telephone work (but not data) it corresponds to sampling a 4 kHz signal. (Two channels have their LSB reserved for signalling)

Frame has 193 bits	Channel 1	Channel 2	Channel 3	....	Channel 24	0/1	Each channel has 8 bits
	Channels 6 & 12 have 7 data + 1 signalling bit					↑	Framing bit

The European ITU-T E-1 standard uses a frame with 32 8-bit “slots”, again with 8,000 frames per second, for 32×64 kbps = 2.048 Mbps.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	**16**	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

↑ Slot 0 – Frame alignment ↑ Slot 16 – signalling

Multiplexing, or grouping, of channels

North America			ITU-T		
Circuit	Bit Rate (Mbps)	channels	Circuit	Bit Rate (Mbps)	channels
DS1	1.544	24	E1	2.048	30
DS1C	3.152	48	E2	8.448	120
DS2	6.312	96	E3	34.368	480
DS3	44.736	672	E4	139.264	1,920
DS4E	139.264	1,920	E5	565.148	7,680
DS4	274.176	4,032			

Synchronous Digital Hierarchy (SDH, or SONET)

- The ITU-T (was CCITT) specifies a Synchronous Digital Hierarchy, or SDH, based on a 125µs frame (8,000 frames/second).
- The basic STM-1 frame has 9 rows of 270 columns of bytes (octets), with the bytes transferred in "raster scan" (left to right, top to bottom).
- Higher speed STM-*n* frames are formed by byte-interleaving *n* of the basic STM-1 frames. (For an STM–*n* frame, send the first bytes of the first row, one for each channel in order, then the second bytes, and so on.)

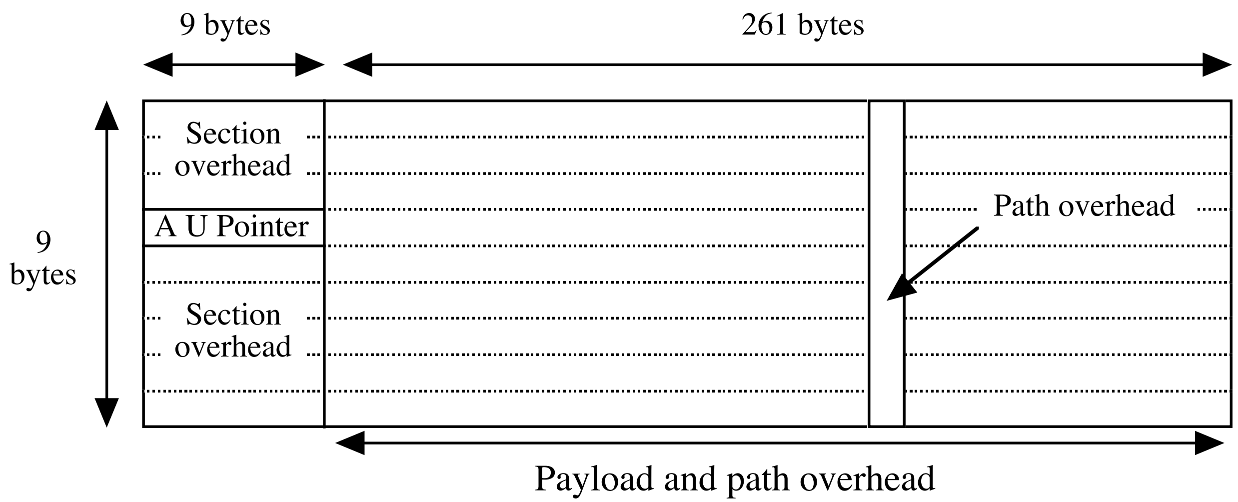

S T M - 1 frame

STS-1 Frame: overhead layouts

	Section Overhead			Path Overhead
Section Overhead	Framing A1	Framing A2	STS-ID C1	Trace J1
	BIP-8 B1	Orderwire E1	User F1	BIP-8 B3
	DataCom D1	Data Com D2	Data Com D3	Signal Label C2
Line Overhead	Pointer H1	Pointer H2	Pointer Action H3	Path Status G1
	BIP-8 B2	APS K1	APS K2	User F2
	DataCom D4	Data Com D5	Data Com D6	Multiframe H4
	DataCom D7	Data Com D8	Data Com D9	Growth Z3
	DataCom D10	Data Com D11	Data Com D12	Growth Z4
	Growth Z Z1	Growth Z2	Orderwire E2	Growth Z5

The actual data frame “floats” within and across transmission frames to accommodate clock variations between the sender and the frame.

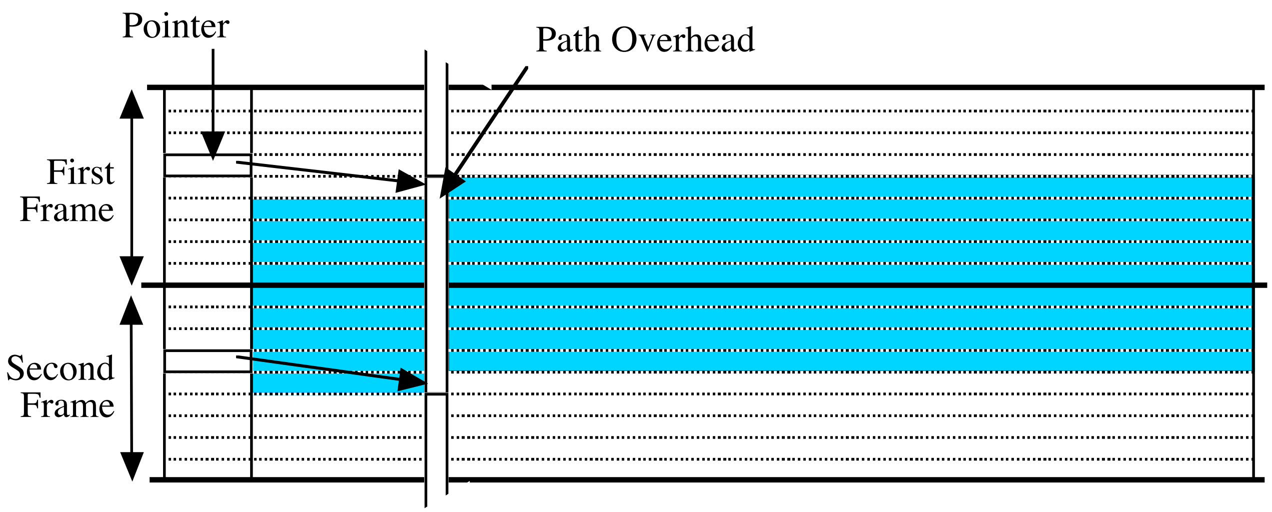

- The AU Pointer of the section header gives the byte offset within the physical frame of the start of the data frame (the first byte of the Path Overhead)

The defined transmission speeds for the North American (SONET) standards and the ITU (CCITT) standards are shown in the table.

SONET designation	**"Optical" designation**	**SONET rate Mbps**	**ITU designation**	**ITU rate Mbps**
STS-1	OC-1	51.84		50.112
STS-3	**OC-3**	**155.52**	**STM-1**	**150.336**
STS-9	OC-9	466.56	STM-3	451.008
STS-12	**OC-12**	**622.08**	**STM-4**	**601.344**
STS-18	OC-18	933.12	STM-6	902.016
STS-24	OC-24	1,244.16	STM-8	1,202.688
STS-36	OC-36	1,866.24	STM-12	1,804.032
STS-48	**OC-48**	**2,488.32**	**STM-16**	**2,405.376**
STS-192	**OC-192**	**9,953.28**	**STM-64**	**9,621.504**

High capacity transmission rate assignments.

Nine bytes of each frame are interpreted by repeaters, 18 by multiplexers and 9 bytes (one column within the frame) by the end-to-end path logic.

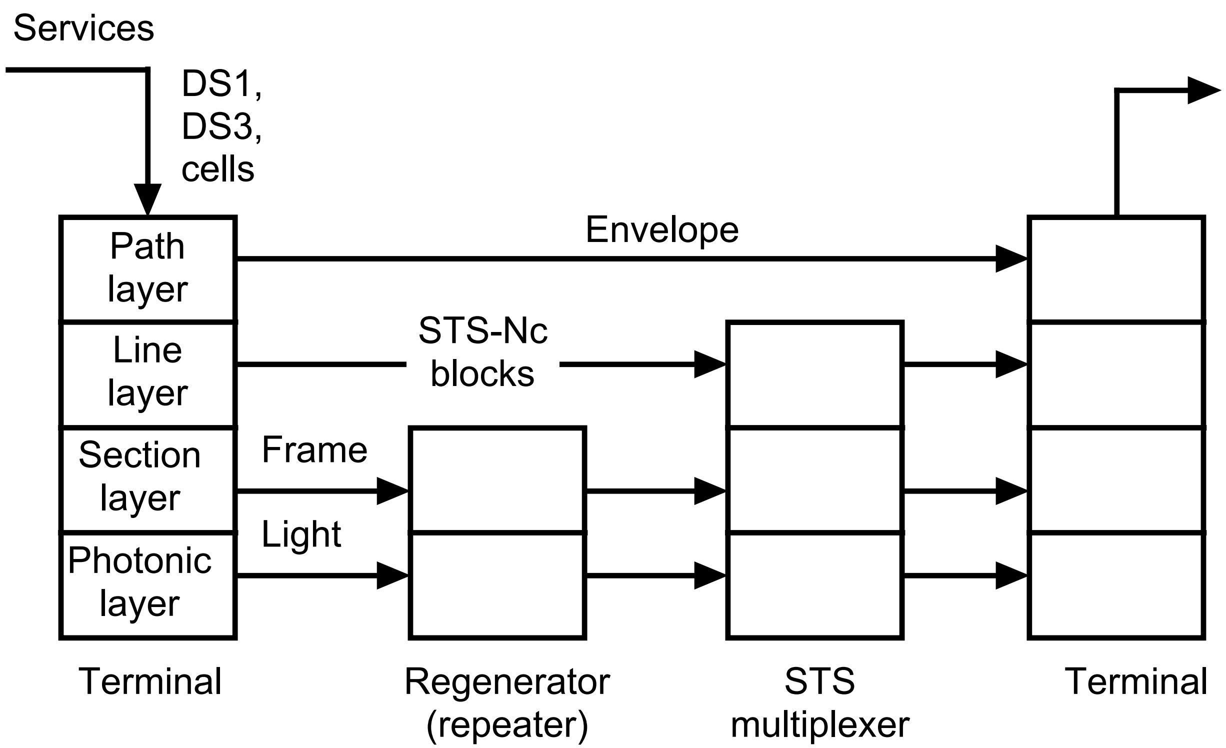

Digital coding

- Here are some of the main digital codes which you may meet.
- The names are not standardised. What is here is as accurate as I can get, but you will certainly find variations and many textbooks differ.
- Most encode 0s and 1s into sequences of 0s and 1s, but some encode into sequences of +1, 0, –1
- All but the simplest (NRZ, NRZI and NRZ1) allow clocks to be recovered from the transmitted (or recorded) bit stream.
- The more complex codes minimise the number of data transitions (baud rate) for given data (bit rate).
  [As the *transmission medium* can handle a limited baud rate, this maximises the *data rate*.]
- You should know NRZ, NRZI, NRZ1 and Manchester codes; the others are for information only.

Return to Zero (RZ), or NRZ-L A pulse for 1, no pulse for a 0. Some versions use +ve and -ve pulses, for 1 and 0 respectively.

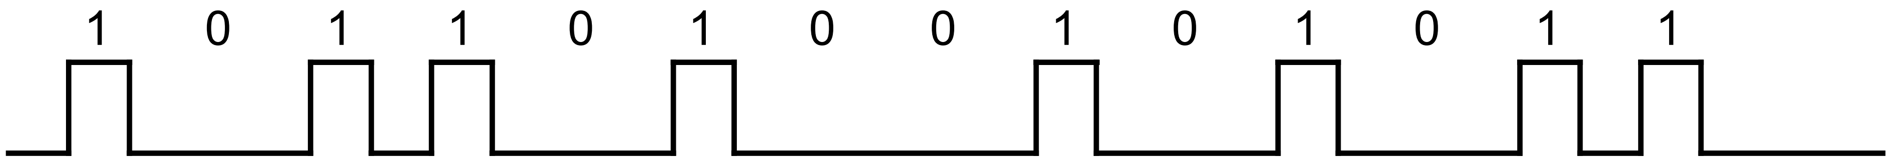

Non-Return to Zero (NRZ) A 0 or 1 level as required

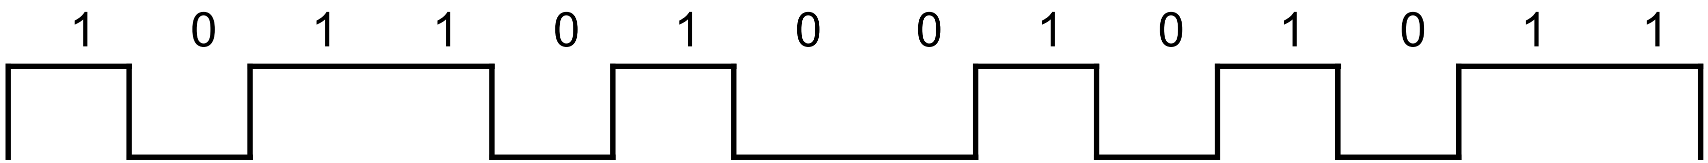

Non-Return to Zero 1 (NRZ1, or NRZ-M) A transition for a 1, no change for a 0

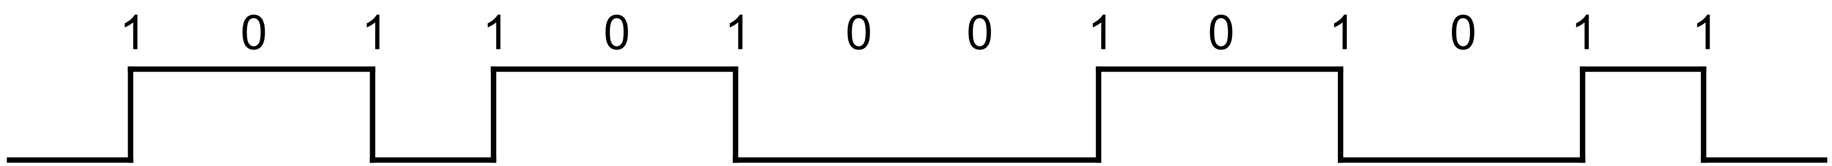

Non-Return to Zero, Inverted (NRZ-I, or NRZ-S) A transition for a 0, no change for a 1

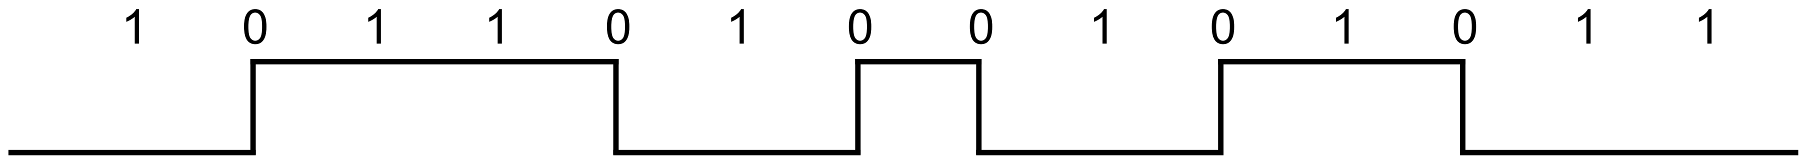

Phase, or Manchester, Encoding (PE) The centre of each bit cell has a transition to the value of that bit; adjacent like bits have a transition on the cell boundary, unlike bits have no transition between cells.

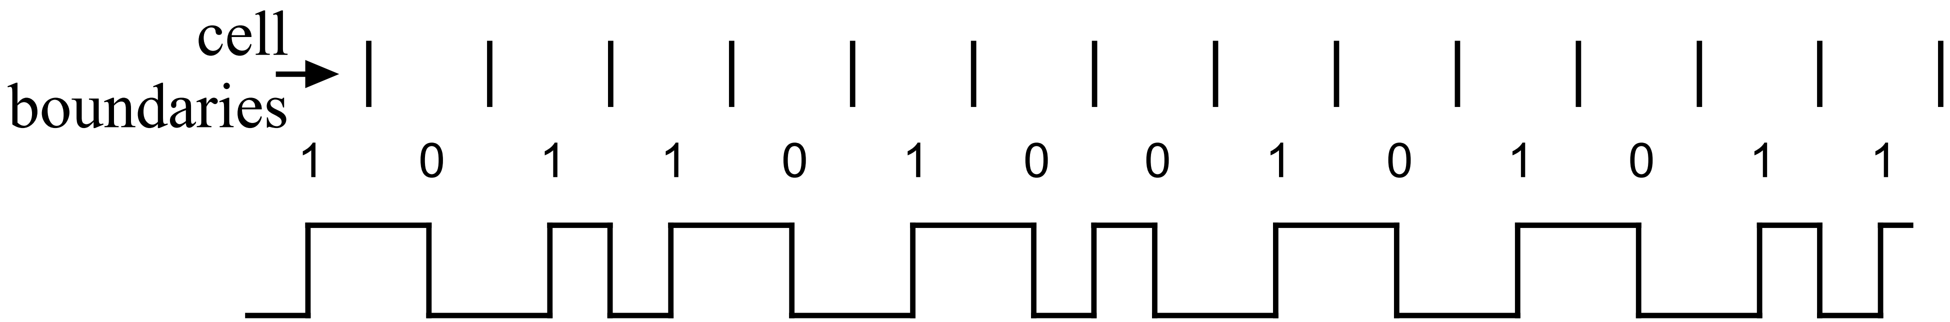

Differential Manchester Encoding The centre of each bit cell has a transition; a 0 bit starts with a transition at the cell boundary, while a 1 bit has no leading transition.

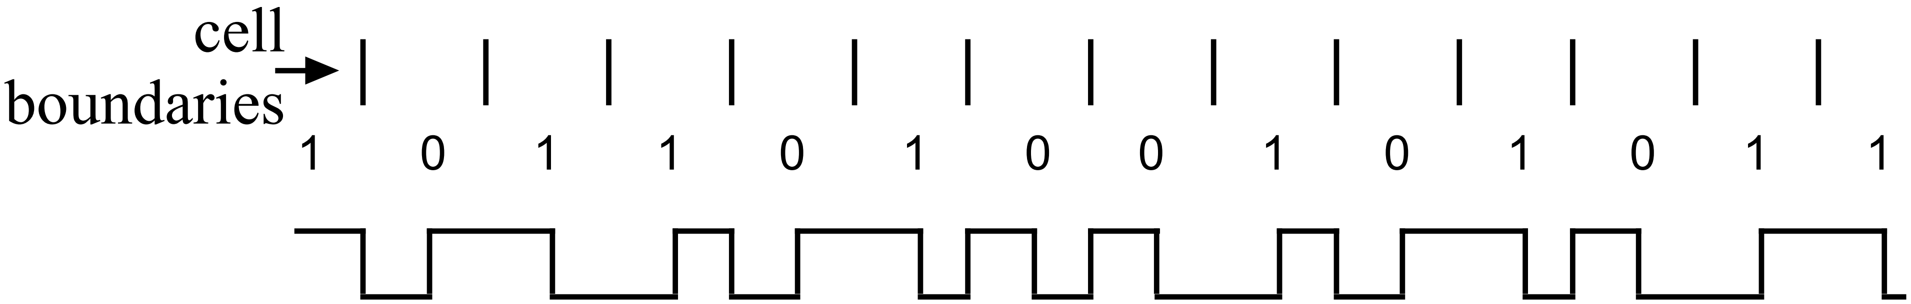

Frequency Modulation (FM) A transition at every cell boundary and at the middle of every “1” cell.
Also known as Binary FM, Phase-Encoding Mark or Double-Frequency

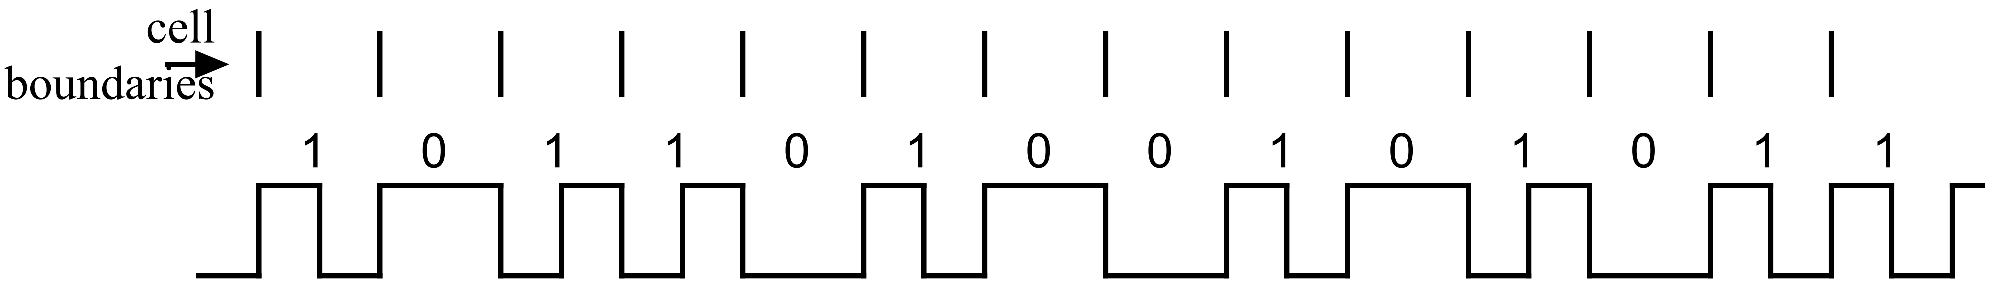

Modified FM (MFM, or Miller code)

A transition at the centre of a 1 cell and between adjacent 0 cells. Derived from FM by deleting “excess” clock transitions.

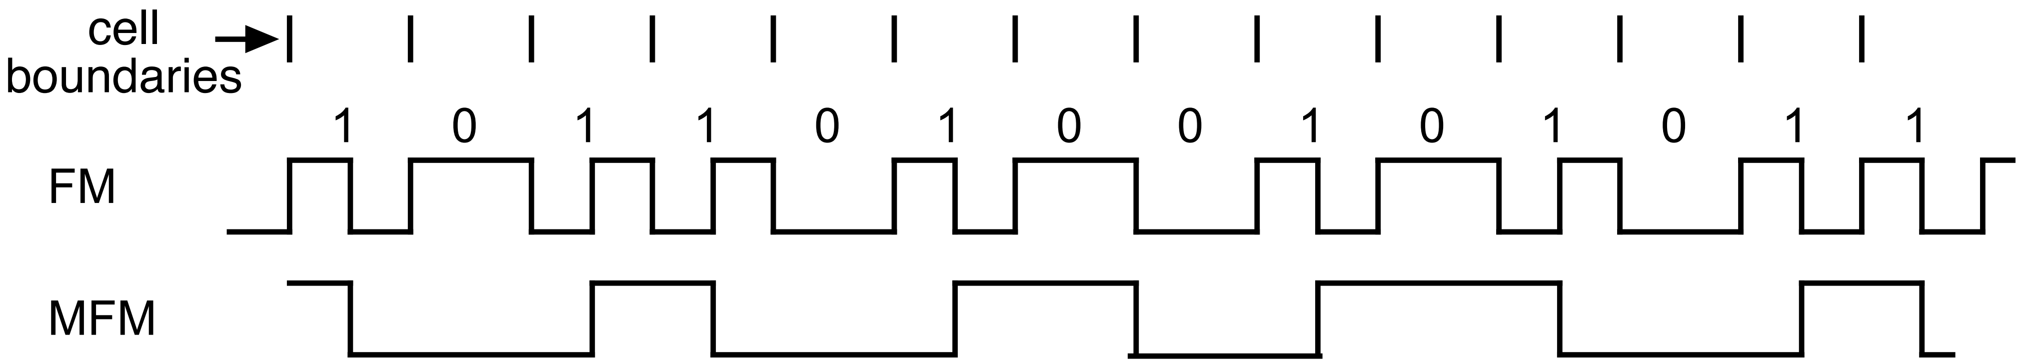

Miller Squared FM (M²FM)

From the MFM code, if an isolated 0 is followed by an even number of 1’s, omit the transition for the final 1. (This controls the average DC level.)

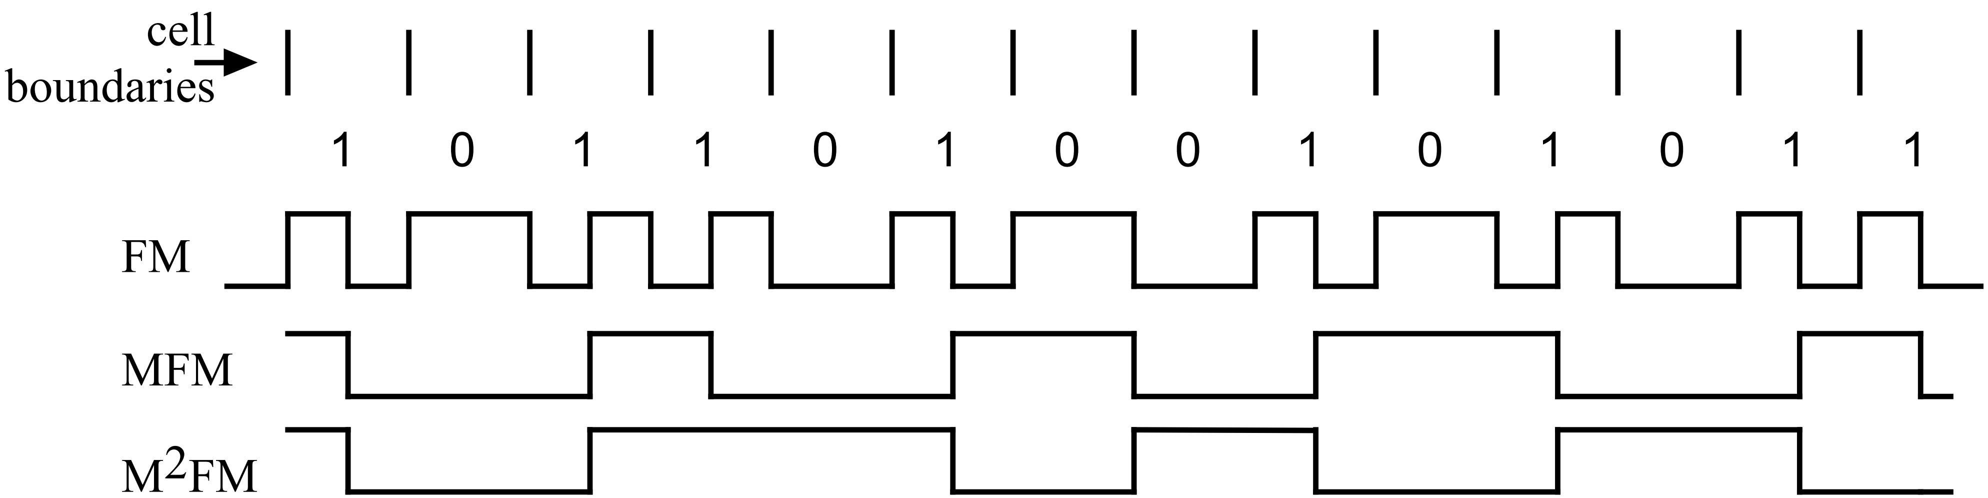

Many codes for very high speed transmission use 3 levels {–1, 0, +1}, with alternating +ve and –ve levels to maintain a DC balance.

These codes are often used in optical communication. Lasers in these cases have varying amplitude, but are *never* turned right off (for a zero etc).

Bipolar AMI (alternate mark inversion) 0 → no line signal (0), 1 → alternate –1 and +1 (This code is used in North America)

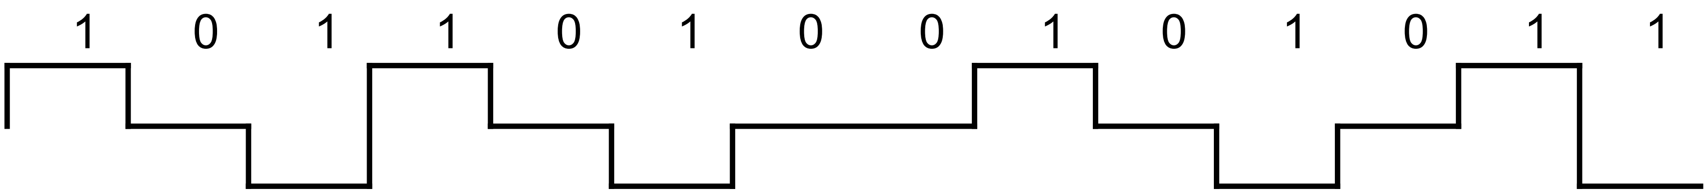

Pseudo Ternary 0 → alternate –1 and +1,
1 → no line signal (0). (Europe and Japan)

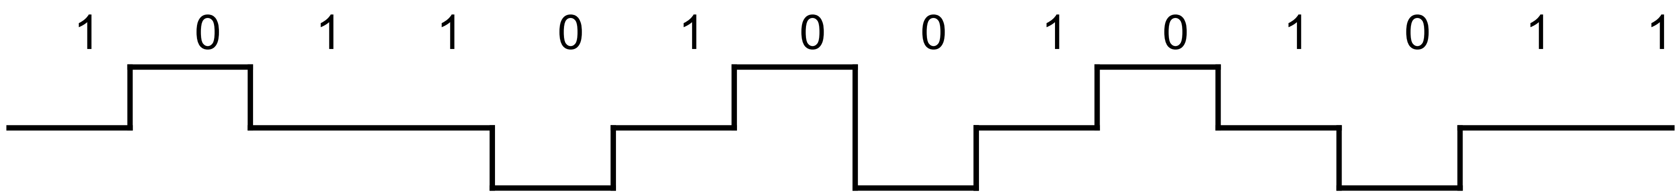

Bipolar, with 8-Zeros Substitution (B8ZS) Same as bipolar AMI, but any string of 8 zeros is replaced by a string with *two* code violations

High Density Bipolar — 3 Zones (HDB3) Same as bipolar AMI, but any string of 4 zeros is replaced by a string with *one* code violation, as 100V, or B00V.

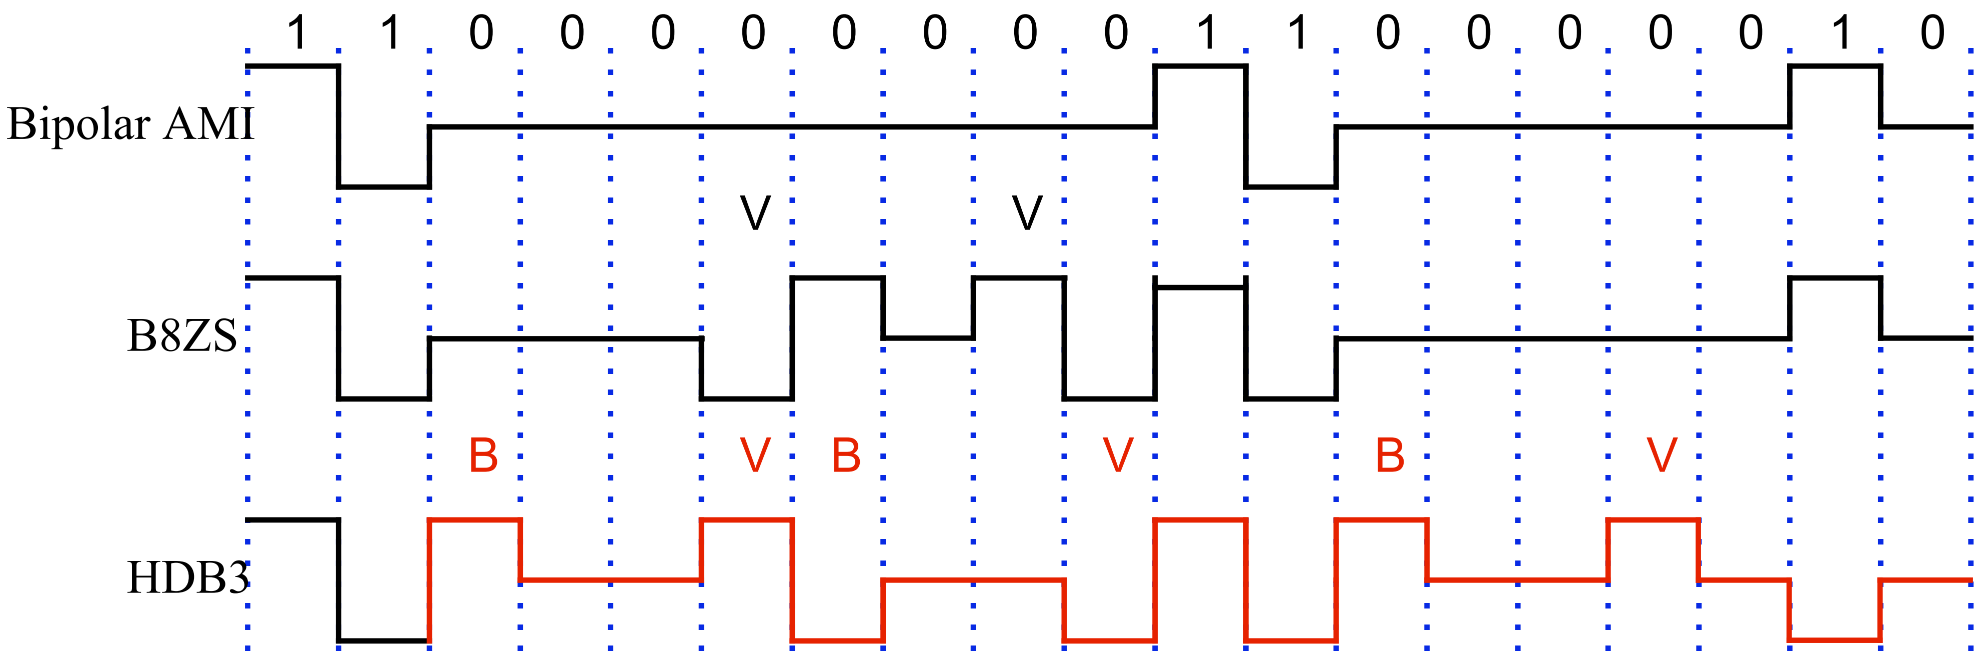

V marks a coding violation (signal not alternating)
B is a "legal" transition (to the opposite polarity)

Optical transmission

Most data transmission of more than a few metres is now over "optical fibre" or, simply, "fibre". Points to note are —

- Most communication uses infrared, so you never see any visible light! (*Never look at the end of a fibre unless you are sure about the other end!! The infrared has a power density like looking directly into the sun.*)
- Light sources may be lasers (for high speed, long distance, expensive) or LEDs (Light Emitting Diodes) slower , shorter distances and lower costs.
- Several types of fibre (details later)
  - Multimode — short distances, low speed, lower cost
  - Graded multimode — medium performance and cost
  - Single mode — long distance, high, expensive
- Core is optically denser (higher refractive index, μ)
- Cladding is optically less dense (lower refractive index, μ)

Characteristics of optical fibres

	Multi-Mode	Graded Index	Single Mode
Light Source	LED or laser	LED or laser	Laser
Bandwidth	20 MHz/km	> 1 GHz/km	>1000 GHz/km
Splicing	Difficult	Difficult	Difficult
Application	Computer data links, LANS	Moderate length telephone	Long distance telecommunications
Cost	Least expensive	More expensive	Most expensive
Core diameter	> 80 µm	50–60µm	1.5–5µm
Cladding diameter	> 160 µm	100–120 µm	15–50 µm
Attenuation	0.5–2.0 dB/km	0.5–2.0 dB/km	<0.15 dB/km

Many fibres now use multiple wavelengths or “colours” (Wavelength Division Multiplexing or WDM)

Optical transmission

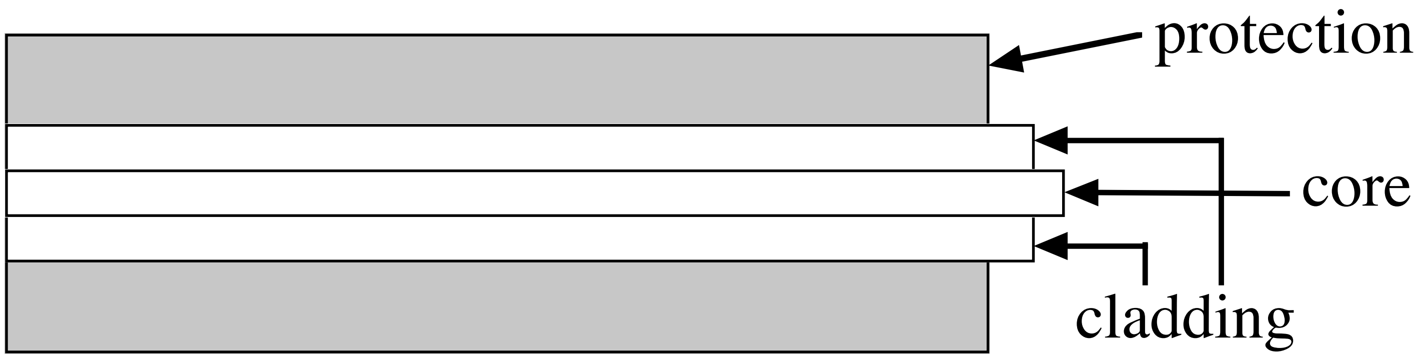

- Light travels slower in the optically denser core ($\mu_{core} > \mu_{cladding}$); the wave front becomes concave.
- As light travels perpendicular to the wave front, it focusses around the core.

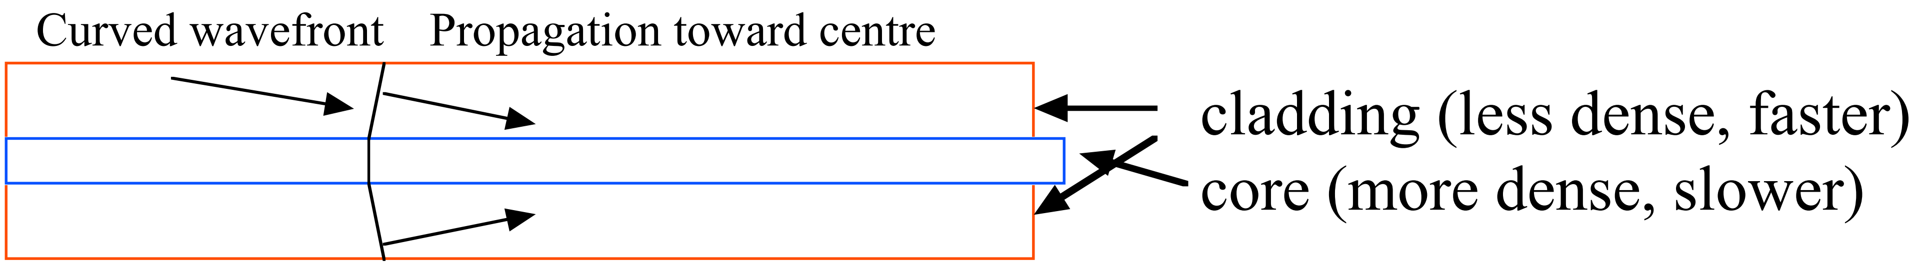

- Light keeps away from the surface of the cladding (and from scratches etc)
- OR light suffers internal reflection from cladding/core boundary; this is the usual explanation and the effect is the same.

Transmission modes

With smaller cores, the light all travels closer to the core. An end-on view of the power density in the cable ("*single mode*") is like

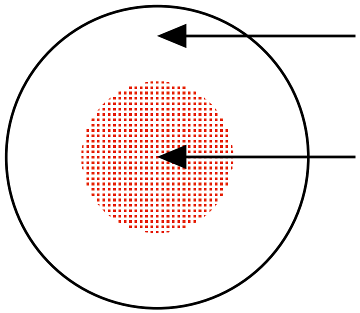

No signal power near edges

Power concentrated near middle

(Depends only on core diameter.)

As the core diameter increases, different power distributions or "modes" become possible, such as

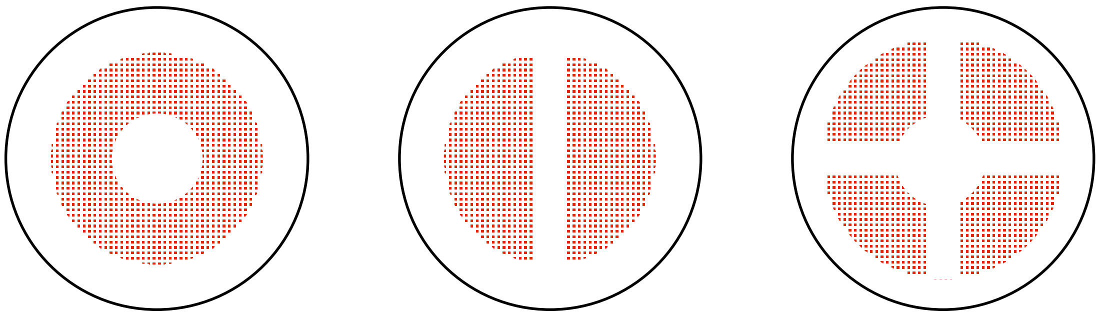

Modes travel at different speeds

Signal swaps between modes, or power "leaks" between modes

Multimode signals "blur" with distance because the modes travel at different speeds and interfere with each other.

Southern Cross Cable

`http://www.southerncrosscables.com/`

As an example of a long distance fibre cable, the Southern Cross cable links Australia, New Zealand, and U S A, in a redundant ring structure.

10	Gbps per colour or wavelength;
16	wavelengths/fibre pair
3	fibre pairs per cable
480	Gbit/s capacity per cable

And some current costings

155 Mbps (STM-1) 2 yr lease	$150,000 US per month
155 Mbps (STM-1) capacity	50,220 GByte/month
Cost per GByte (100% load)	$2.99
Cost per GByte (at 50% load)	$5.97
Cost for one telephone circuit	$0.11 per hour

Two Optical Miscellanea

Optical Amplifiers.

Optical fibres doped with Erbium and “pumped” with a suitable laser power source can amplify optical signals with no need to for conversion electronics. Optical amplifiers are simple and reliable, with very high bandwidth.

Solitons

Most materials are dispersive, changing their refractive index with *wavelength* (prisms, rainbows, etc). A few are also dispersive with signal *amplitude*.

If the refractive index changes suitably with both *wavelength* and *amplitude*, the two types of dispersion interact to allow single pulses to maintain their shape for very long distances, perhaps thousands of kilometres in glass.

Solitons look to become important for very long distance transmission.