which the signals must reside in a given frequency band. Channels are often shared by multiple signals. After all, it is much more convenient to use a single wire to carry several signals than to install a wire for every signal. This kind of sharing is called **multiplexing**. It can be accomplished in several different ways.

2.4.1 Baseband Transmission

The most straightforward form of digital modulation is to use a positive voltage to represent a 1 and a negative voltage to represent a 0. For an optical fibre, the presence of light might represent a 1 and the absence of light might represent a 0. This scheme is called NRZ (Non-Return-to-Zero). An example is shown in Fig. 2.10(b). Once sent, the NRZ signal propagates down the wire. At the other end, the receiver converts it into bits by sampling the signal at regular intervals of time.

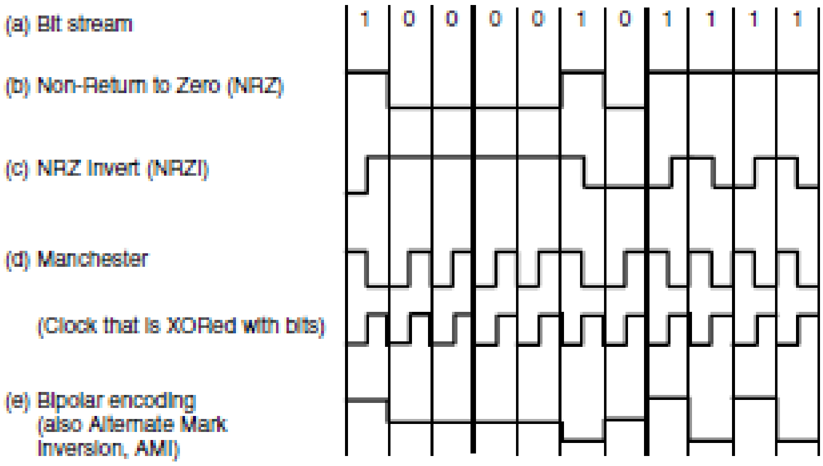

Note that:
The NRZI represents the signal as follows:
One: is represented by a transition of the physical level.
Zero: has no transition.
AMI: In this code, a binary 0 is encoded as zero volts, whereas a binary 1 is encoded alternately as a positive voltage or a negative voltage.

Fig. 2.10: Line codes: (a) Bits, (b) NRZ, (c) NRZI, (d) Manchester, (e) Bipolar or AMI.

This signal will not look exactly like the signal that was sent. It will be attenuated and distorted by the channel and noise at the receiver. To decode the bits, the receiver maps the signal samples to the closest symbols. For NRZ, a positive voltage will be taken to indicate that a 1 was sent and a negative voltage will be taken to indicate that a 0 was sent. NRZ is a good starting point for our studies because it is simple, but it is seldom used by itself in practice. More complex schemes can convert bits to signals that better meet engineering considerations. These schemes are called **line codes**. Below, we describe line codes that help with bandwidth efficiency, clock recovery, and DC balance.

Bandwidth Efficiency:

With NRZ, the signal may cycle between the positive and negative levels up to every 2 bits (in the case of alternating 1s and 0s). This means that we need a bandwidth of at least R/2 Hz when the bit rate is R bits/sec; **How?**

So by increasing the signal levels to four, we can send 2 bits at once as a single **symbol, how?**

We call the rate at which the signal changes the **symbol rate** to distinguish it from the bit rate. The **bit rate** is the symbol rate multiplied by the number of bits per symbol. An older name for the symbol rate, particularly in the context of devices called telephone modems that convey digital data over telephone lines, is the **baud rate**.

Clock Recovery:

To overcome the clock synchronisation limitation, A clever trick here is to mix the clock signal with the data signal by XORing them together so that no extra line is needed. The results are shown in Fig. 2.10(d). The clock makes a clock transition in every bit time, so it runs at twice the bit rate. When it is XORed with the 0 level it makes a low-to-high transition that is simply the clock. This transition is a logical 0. When it is XORed with the 1 level it is inverted and makes a high-to-low transition. This transition is a logical 1. This scheme is called **Manchester** encoding and was used for classic Ethernet. The downside of Manchester encoding is that it requires twice as much bandwidth as NRZ because of the clock.

Further, coding a 1 as a transition and a 0 as no transition, or vice versa, this coding is called **NRZI (Non-Return-to-Zero Inverted),** a twist on **NRZ**. The popular **USB (Universal Serial Bus**) standard for connecting computer peripherals uses NRZI.

Balanced Signals:

Signals that have as much positive voltage as negative voltage even over short periods of time are called **balanced signals**. They average to zero, which means that they have no DC electrical component. The lack of a DC component is an advantage because some channels, such as coaxial cable or lines with transformers, strongly attenuate a DC component due to their physical properties. Also, one method of connecting the receiver to the channel called **capacitive coupling** passes only the AC portion of a signal. In either case, if we send a signal whose average is not zero, we waste energy as the DC component will be filtered out.

Balancing helps to provide transitions for clock recovery since there is a mix of positive and negative voltages. It also provides a simple way to calibrate receivers because the average of the signal can be measured and used as a decision threshold to decode symbols. With unbalanced signals, the average may be drift away from the true decision level due to a density of 1s, for example, which would cause more symbols to be decoded with errors.

A straightforward way to construct a balanced code is to use two voltage levels to represent a logical 1, (say +1 V or −1 V) with 0 V representing a logical zero. To send a 1, the transmitter alternates between the +1 V and −1 V levels so that they always average out. This scheme is called **bipolar encoding**. In telephone networks it is called **AMI (Alternate Mark Inversion)**, building on old terminology in which a 1 is called a ‘‘mark’’ and a 0 is called a ‘‘space.’’ An example is given in Fig. 2.10(e).

2.4.2 Passband Transmission

Often, we want to use a range of frequencies that does not start at zero to send information across a channel. For wireless channels, it is not practical to send very low frequency signals because the size of the antenna needs to be a fraction of the signal wavelength, which becomes large. In any case, regulatory constraints and the need to avoid interference usually determine the choice of frequencies. Even for wires, placing a signal in a given frequency band is useful to let different kinds of signals coexist on the channel. This kind of transmission is called passband transmission because an arbitrary band of frequencies is used to pass the signal.

In the **passband transmission** we can take a baseband signal that occupies **0 to B Hz** and shift it up to occupy a passband of **S to S +B Hz** without changing the amount of information that it can carry, even though the signal will look different. To process a signal at the receiver, we can shift it back down to baseband, where it is more convenient to detect symbols.

Digital modulation is accomplished with passband transmission by regulating or modulating a carrier signal that sits in the passband. We can modulate the **amplitude, frequency, or phase** of the carrier signal. Each of these methods has a corresponding name. In **ASK (Amplitude Shift Keying)**, two different amplitudes are used to represent 0 and 1. An example with a nonzero and a zero level is shown in Fig. 2.11(b). Similarly, with **FSK (Frequency Shift Keying)**, two or more different tones are used. The example in Fig. 2.11(c) uses just two frequencies. In the simplest form of **PSK (Phase Shift Keying),** the carrier wave is systematically shifted 0 or 180 degrees at each symbol period. Because there are two phases, it is called **BPSK (Binary Phase Shift Keying).** ‘‘Binary’’ here refers to the two symbols, not that the symbols represent 2 bits. An example is shown in Fig. 2.11(d).

Further, a combination of these schemes is used to transmit more bits per symbol, such as: QPSK, QAM, where QAM stands for Quadrature Amplitude Modulation, as shown in Fig. 2.12. The phase of a dot is indicated by the angle a line from it to the origin makes with the positive x-axis. The amplitude of a dot is the distance from the origin.

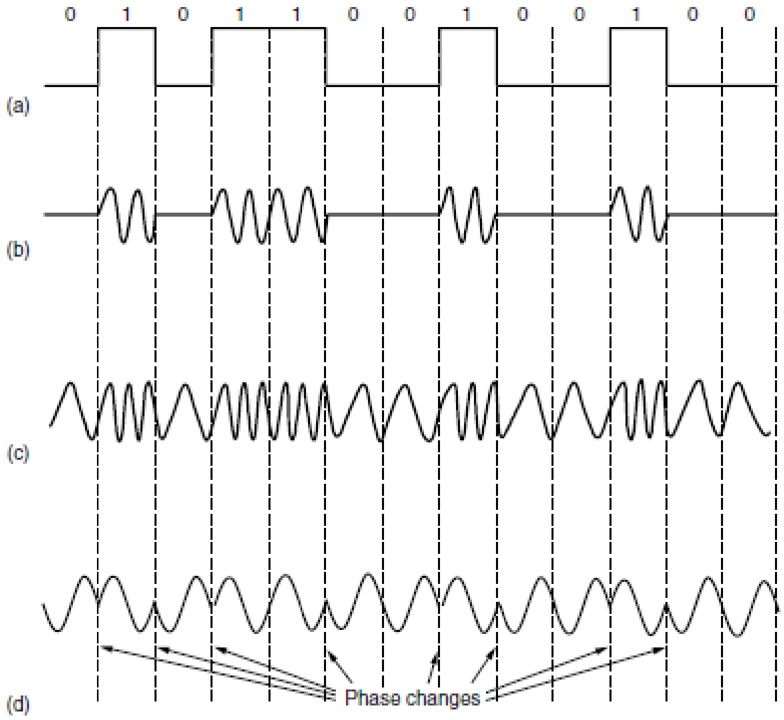

Fig. 2.11: (a) A binary signal. (b) ASK. (c) FSK. (d) PSK.

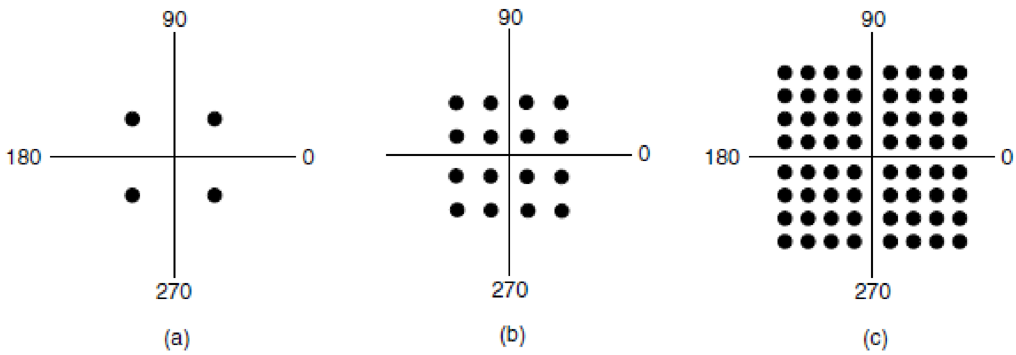

Fig. 2.12: (a) QPSK. (b) QAM-16. (c) QAM-64.

2.4.3 Frequency Division Multiplexing

Multiplexing schemes have been developed to share lines among many signals. **FDM (Frequency Division Multiplexing)** takes advantage of passband transmission to share a channel. It divides the spectrum into frequency bands, with each user having exclusive possession of some band in which to send their signal. Different frequencies are allocated to different logical channels (stations), each operating in a portion of the spectrum, with the interchannel separation great enough to prevent interference.

For a more detailed example, in Fig. 2.13 we show three voice-grade telephone channels multiplexed using FDM. Filters limit the usable bandwidth to about 3100 Hz per voice-grade channel. When many channels are multiplexed together, 4000 Hz is allocated per channel.

The excess is called a **guard band**. It keeps the channels well separated. First the voice channels are raised in frequency, each by a different amount. Then they can be combined because no two channels now occupy the same portion of the spectrum.

In OFDM (Orthogonal Frequency Division Multiplexing), the channel bandwidth is divided into many subcarriers that independently send data (e.g., with QAM). The subcarriers are packed tightly together in the frequency domain.

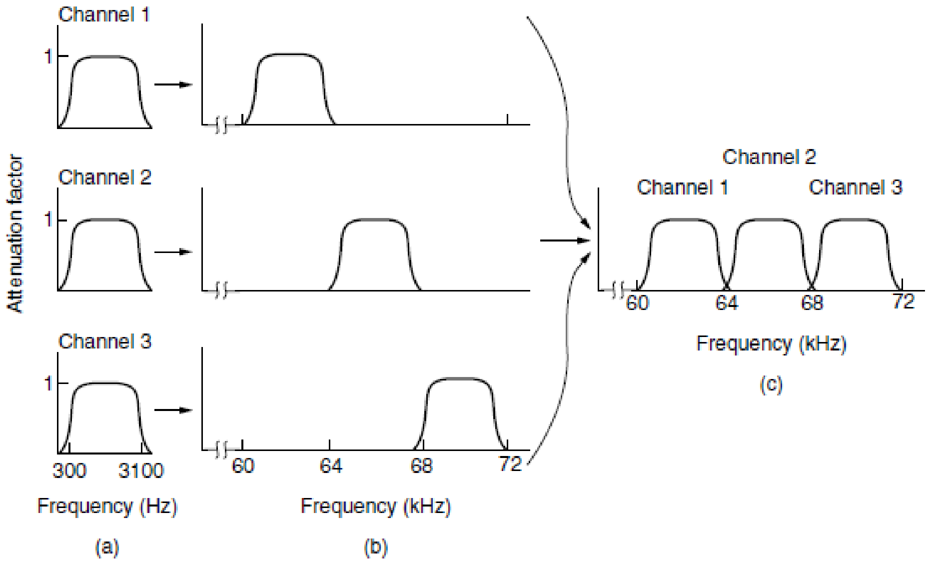

Figure 2.13. Frequency division multiplexing. (a) The original bandwidths. (b) The bandwidths raised in frequency. (c) The multiplexed channel.

2.4.4 Time Division Multiplexing

An alternative to FDM is **TDM (Time Division Multiplexing)**. Here, the users take turns (in a round-robin fashion), each one periodically getting the entire bandwidth for a little burst of time. An example of three streams being multiplexed with TDM is shown in Fig. 2.14. Bits from each input stream are taken in a fixed **time slot** and output to the aggregate stream. This stream runs at the sum rate of the individual streams. For this to work, the streams must be synchronized in time. Small intervals of **guard time** analogous to a frequency guard band may be added to accommodate small timing variations. TDM is used widely as part of the telephone and cellular networks.

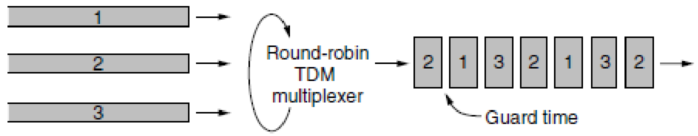

Figure 2.14. Time Division Multiplexing (TDM).

2.4.5 Code Division Multiplexing

CDM (Code Division Multiplexing) is a form of spread spectrum communication in which a narrowband signal is spread out over a wider frequency band. This can make it more tolerant of interference, as well as allowing multiple signals from different users to share the same frequency band. Because code division multiplexing is mostly used for the latter purpose it is commonly called CDMA (Code Division Multiple Access).

2.5 Modems, ADSL, and Fibre

2.5.1 Telephone Modems

To send bits over the local loop, or any other physical channel for that matter, they must be converted to analogue signals that can be transmitted over the channel. This conversion is accomplished using a device called **modem**. Modems come in many varieties: telephone modems, DSL modems, cable modems, wireless modems, etc. The modem may be built into the computer (which is now common for telephone modems) or be a separate box (which is common for DSL and cable modems).

Telephone modems are used to send bits between two computers over a voice-grade telephone line. The main difficulty is that a voice-grade telephone line is limited to 3100 Hz. This bandwidth is more than four orders of magnitude less than the bandwidth that is used for Ethernet or 802.11 (WiFi).

With a perfect 3000-Hz line there is no point in sending symbols at a rate faster than 6000 baud (Why?). In practice, most modems send at a rate of 2400 symbols/sec, or 2400 baud. The humble 2400-bps modem uses 0 volts for a logical 0 and 1 volt for a logical 1, with 1 bit per symbol. One step up, it can use four different symbols, as in the four phases of QPSK, so with 2 bits/symbol it can get a data rate of 4800 bps. The end result is the V.90 and V.92 modem standards. They provide for a 56-kbps downstream channel (ISP to user) and a 33.6-kbps and 48-kbps upstream channel (user to ISP), respectively.

2.5.2 Digital Subscriber Lines

Initially, there were many overlapping high-speed offerings, all under the general name of *x*DSL (Digital Subscriber Line), for various *x*.

The telephone wire runs through a filter that attenuates all frequencies below 300 Hz and above 3400 Hz. The trick that makes *x*DSL work is that when a customer subscribes to it, the incoming line is connected to a different kind of switch, one that does not have this filter, thus making the entire capacity of the local loop available. The limiting factor then becomes the physics of the local loop, which supports roughly 1 MHz, not the artificial 3100 Hz bandwidth created by the filter. A typical ADSL equipment configuration is shown in Fig. 2.15.

The xDSL services have all been designed with certain goals:

- The services must work over the existing Category 3 twisted pair local loops.
- They must not affect customers' existing telephones and fax machines.
- They must be much faster than 56 kbps.
- They should be always on, with just a monthly charge and no per-minute charge.

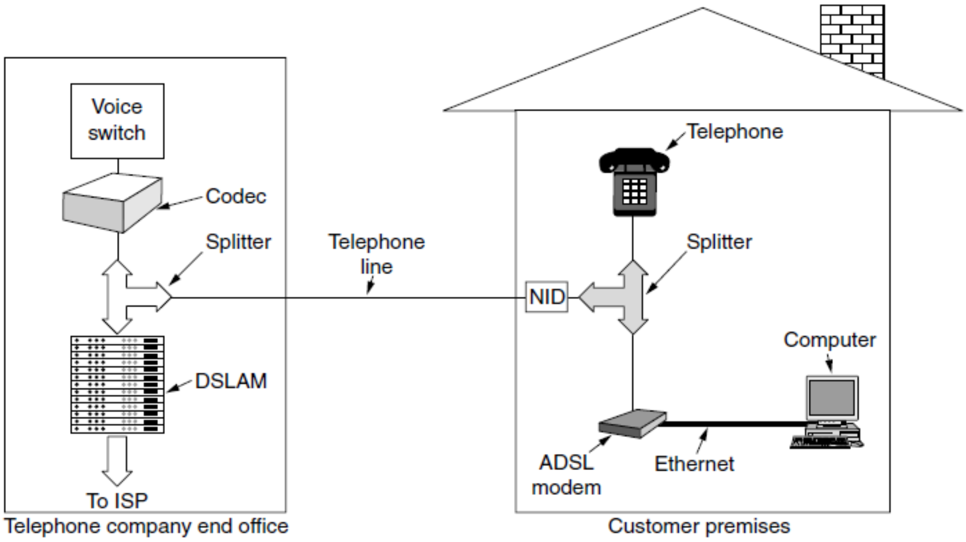

Figure 2.15: A typical ADSL equipment configuration.
DSLAM (Digital Subscriber Line Access Multiplexer)
NID (Network Interface Device)

2.5.3 Fibre To The Home

Deployed copper local loops limit the performance of ADSL and telephone modems. To let them provide faster and better network services, telephone companies are upgrading local loops at every opportunity by installing optical fibre all the way to houses and offices. The result is called **FttH (Fiber To The Home).**

While FttH technology has been available for some time, deployments only began to take off in 2005 with growth in the demand for high-speed Internet from customers used to DSL and cable who wanted to download movies with Internet access speeds of up to 100 Mbps.

Usually, the fibres from the houses are joined together so that only a single fibre reaches the end office per group of up to 100 houses. In the downstream direction, optical splitters divide the signal from the end office so that it reaches all the houses. Encryption is needed for security if only one house should be able to decode the signal. In the upstream direction, optical combiners merge the signals from the houses into a single signal that is received at the end office. This architecture is called a **PON (Passive Optical Network)**, and it is shown in Fig. 2.16. It is common to use one wavelength shared between all the houses for downstream transmission, and another wavelength for upstream transmission. Even with the splitting, the tremendous bandwidth and low attenuation of fibre mean that PONs can provide high rates to users over distances of up to 20 km. The actual data rates and other details depend on the type of PON. Two kinds are common. **GPONs (Gigabit-capable PONs)** come from the world of telecommunications, so they are defined by an ITU standard. **EPONs (Ethernet PONs)** are more in tune with the world of networking, so they are defined by an IEEE standard. Both run at around a gigabit and can carry traffic for different services, including Internet, video, and voice. For example, GPONs provide 2.4 Gbps downstream and 1.2 or 2.4 Gbps upstream.

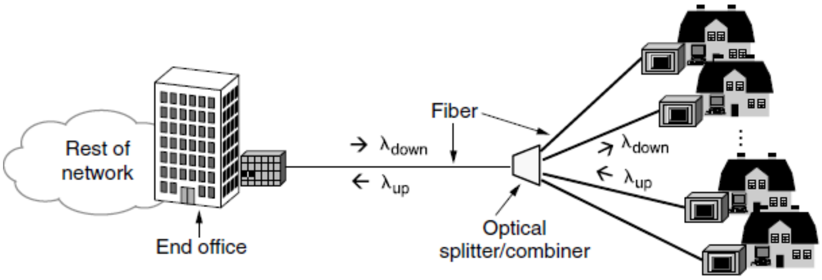

Figure 2.16: Passive optical network for Fibre To The Home.