

Syllabus Content:

1.1.2 Images

-  show understanding of how data for a bitmapped image is encoded
-  terminology associated with bitmaps: pixel, file header, image resolution, screen resolution
-  perform calculations estimating the file size for bitmapped images of different resolutions
-  show understanding of how data for a vector graphic is represented and encoded
-  terminology associated with vector graphics: drawing object, property and drawing list
-  features found in bitmapped and vector graphics software are used in practice
-  justify where bitmapped graphics and/or vector graphics are appropriate for a given task

Graphic images that have been processed by a computer can usually be divided into two distinct categories.

Pixels vs. Vector



Such images are either bitmap files or vector graphics. This section will cover the two main image types: **vector** and **bitmap(pixels)**, along with some compression techniques.

Pixel (Picture element):

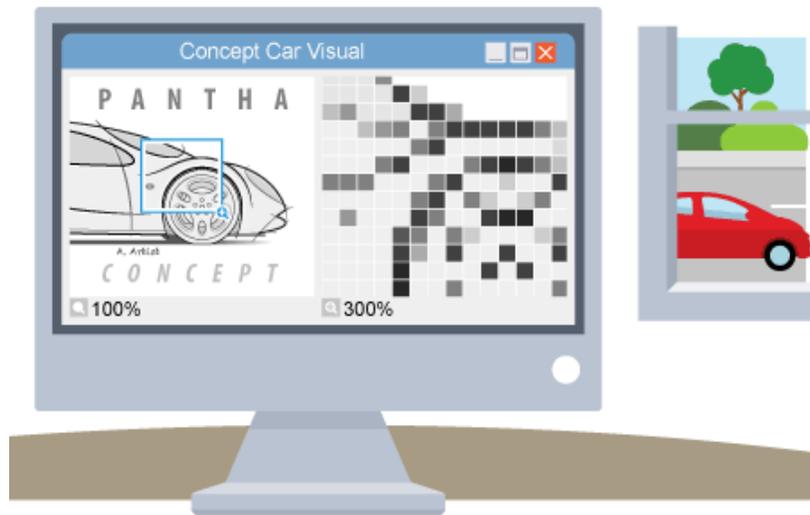
Graphics on a screen are made up of tiny blocks called pixels. The more pixels on the screen, the higher the resolution and the better the quality of the picture will be. The higher the image resolution, the more memory is needed to store the graphic.

Image files can be either **bitmaps** or **vectors**.

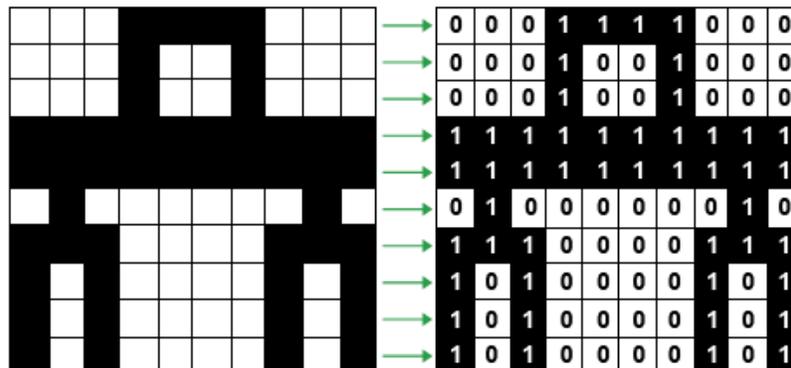
Bitmap Graphics:

Bitmaps images are exactly what their name says they are: a collection of bits that form an image.

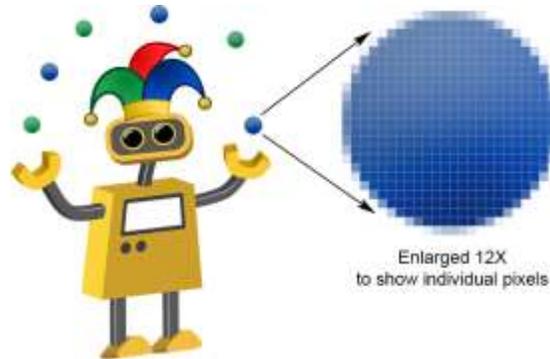
- **Bitmap images** (or raster) images are stored as a series of tiny dots called pixels.
- Each pixel is actually a very small square that is assigned a color, and then arranged in a pattern to form the image.
- The image consists of a matrix of individual dots (or pixels (picture element)) that all have their own color (described using bits, the smallest possible units of information for a computer).
- Bitmap images are organised as a grid of coloured squares called pixels (short for 'picture elements').
- When zooming in or enlarging a bitmap image, the pixels are stretched and made into larger blocks. This is why bitmap images appear as poor quality when enlarged too much.



Each colour of an image is stored as a binary number. In the black-and-white image below, each pixel is either black or white. You need a binary value for each different colour. As each pixel is either black or white, this image can be encoded with a value of 0 for white and 1 for black.



When you zoom in on a bitmap image you can see the individual pixels that make up that image. Bitmap graphics can be edited by erasing or changing the color of individual pixels using a program such as Adobe Photoshop.



Pixel (Picture element):

The smallest possible addressable area defined by a solid color, represented as binary, in an image.

File header:

BMP, PCX, JPEG, FLI/FLC, and AVI files include headers that define the image size, number of colors, and other information needed to display the image. Fastgraph provides functions for reading the image headers and retrieving their more useful items. However, there may be times when you need additional information stored in the file header e.g.

- Confirmation that the file is a BMP
- File size
- Location/offset of image data within the file
- Dimensions of the image (in pixels) // image resolution
- Colour depth (bits per pixel, 1, 4, 8, 16, 24 or 32)
- Type of compression used, if any

Resolution:

Screen Resolution - The screen resolution tells you how many pixels your screen can display horizontally and vertically. It's written in the form 1024 x 768. In this example, the screen can show 1,024 pixels horizontally, and 768 vertically:

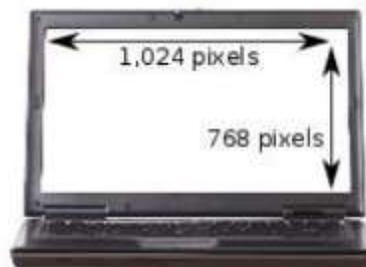


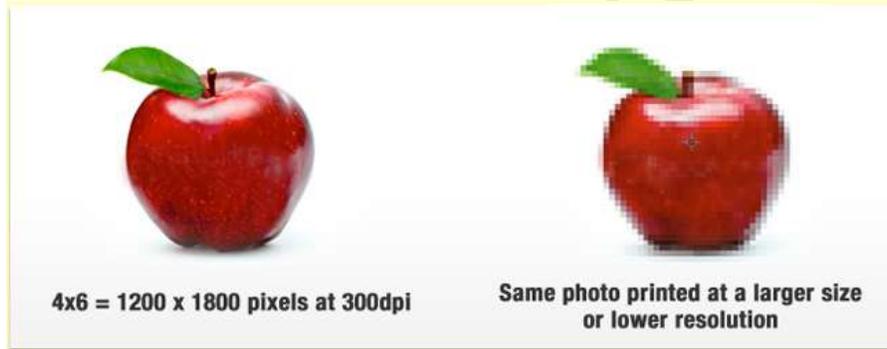
Image Resolution - how many pixels do an image contains per inch/cm?

Measurements of **dots per inch (dpi)** are called image resolution.

-  The number of pixels per unit measurement
-  or the number of pixels in an image
-  or The number of pixels wide by the number of pixels high
-  or number of pixels per row by the number of rows

Resolution refers to the number of pixels in an **image**. Resolution is sometimes identified by the width and height of the **image** as well as the total number of pixels in the **image**. For example, an **image** that is 2048 pixels wide and 1536 pixels high (2048 x 1536) contains (multiply) 3,145,728 pixels (or 3.1 Megapixels).

The higher the resolution, the more pixels are available. Therefore the crisper the picture



Colour depth

The colour depth of an image is measured in bits. The number of bits indicates how many colours are available for each pixel. In the black and white image, only two colours are needed. This means it has a colour depth of 1 bit.

A 2-bit colour depth would allow four different values: 00, 01, 10, 11. This would allow for a range of colours such as:

Binary code	Colour
00	White
01	Light grey
10	Dark grey
11	Black

Greater the colour depth (bits per pixel), the more colours are available.

Colour depth	Available colours
1-bit	$2^1 = 2$
2-bit	$2^2 = 4$
3-bit	$2^3 = 8$
4-bit	$2^4 = 16$
5-bit	$2^5 = 32$
6-bit	$2^6 = 64$
7-bit	$2^7 = 128$
8-bit	$2^8 = 256$

Example: Calculating screen resolutions & Picture file size:

Using the diagram above we are going to work out how many pixels are required to display a single frame on a **VGA** screen.

Checking the resolution:

Height = 480 Width = 640

Area = Width x Height = Total Pixels Area = 640 x 480 = 307200 pixels

Hence, it takes 307,200 pixels to display on a VGA screen!

File size = Total Pixels X Colour Depth

Coloured Images:

- At least **eight bits per pixel are necessary to code a coloured image**. The number of bits per pixel is sometimes referred to as the **colour depth**.

(8 bits or 1 byte for 1 pixel in coloured image) (Computer Science Book)

Hence, it takes 307,200 pixels X 8 bits = 2457600 bits or $2457600/8 = 307200$ bytes or 300KB.

Monochrome or Black & White image:

The simplest option is to use **one bit** to represent the **colour**, so that the pixel is either black or white. (**1 bit for 1 pixel in Black & White or monochrome Picture**)



Greyscale image:

- Storage of the colour in **four bits** would allow simple **greyscale colouring**. (**4 bits for 1 pixel in GreyScale image**)



Greyscale

Coloured Images:

- At least **eight bits per pixel** are necessary to code a **coloured image**. The number of bits per pixel is sometimes referred to as the **colour depth**. (**8 bits or 1 byte for 1 pixel in coloured image**) ([Computer Science Book](#))
- 24 bits** almost always uses **8 bits of each of RGB (Red:8 Green:8 Blue:8)** $8 \times 8 \times 8 = 24$
- As of 2018 **24-bit** color depth is used by virtually every computer and phone display and the vast majority of image storage formats. ([Wikipedia](#))

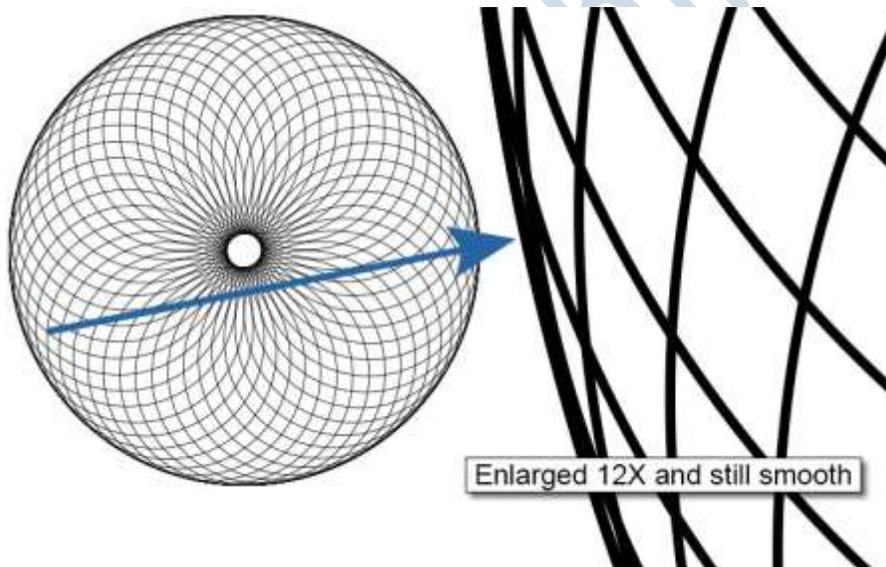


Vector graphics

It is normal for an image that is created by a drawing package or a computer-aided design (CAD) package to consist of a number of geometric objects. The outcome is then usually for the image to be stored as a **vector graphic** file

- Unlike bitmaps, vector images are not based on pixel patterns, but instead use mathematical formulas to draw lines and curves that can be combined to create an image
- Vector graphics are made from geometric objects such as circles and polygons.

- Vector images are edited by manipulating the lines and curves that make up the image using a program such as Adobe Illustrator.
- Vector images tend to be smaller than bitmap images because a bitmap image has to store color information for each individual pixel that forms the image.
- A vector image just has to store the mathematical formulas that make up the image, which take up less space.
- The three most popular image formats used on the Web (PNG, JPEG, and GIF) are bitmap formats.
- The Scalable Vector Graphics (SVG) format comes in a distant fourth due to a legacy of poor support for vector graphics in early browsers. Today however, all major browsers support the SVG (Scalable Vector Graphics) format.
- The dimensions of the vector image objects are not defined explicitly but instead are defined relative to an imaginary drawing canvas. In other words, the image is scalable.



Bitmap formats are best for images that need to have a wide range of color gradations, such as most photographs. Vector formats, on the other hand, are better for images that consist of a few areas of solid color. Examples of images that are well suited for the vector format include logos and type.

Vector graphic: a graphic consisting of components defined by geometric formulae and associated properties, such as line colour and style.

We do not need to consider how an image of this type would be created. We do need to consider how the data is stored after the image has been created.

Drawing list: - a set of commands used to define a vector image. A vector graphic file contains a drawing list. The list contains a command for each object included in the image.

Each command has a list of attributes that define the properties of the object.

Drawing Properties: -The properties include the basic geometric data such as, for a circle, the position of the centre and its radius. In addition properties such as the thickness and style of a line, the colour of a line and the colour that fills the shape, if that is appropriate, are defined. An example of what could be created as a vector graphic file is shown in Figure.

Vector graphics are made up of objects and their properties.

Drawing object: - is a mathematical or geometrically defined construct such as a rectangle, line or circle.

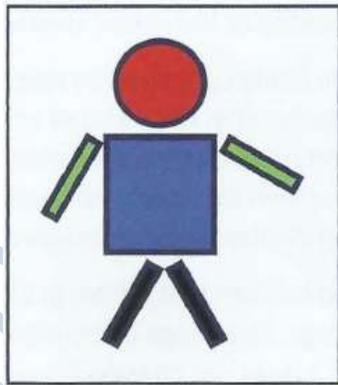
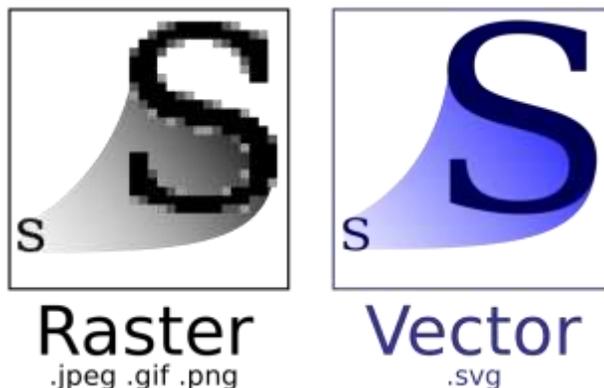


Figure: A simple example of a vector graphic image

Comparison between vector and bitmaps:



This image illustrates the difference between bitmap and vector images. The bitmap image is composed of a fixed set of dots (pixels), while the vector image is composed of a fixed set of shapes. In the picture, scaling the bitmap reveals the pixels and scaling the vector image preserves the shapes.

-  Vector images scale without file size increase / decrease
-  Bitmap images scale resulting in file size increase / decrease
-  Vector images scale without distortion to the image
-  Bitmap images distort (pixelate) when scaling
-  Bitmaps are better for photo editing
-  Bitmaps require less processing power to display

One of the most common vector file formats is **scalable vector graphics (SVG)**. SVG is an open standard for vector graphics.

It is possible to edit SVG images using numbers to change the size and colour variables in HTML. This is often used for graphs and infographics in HTML5.

Vector graphics are used in:

- CAD packages
- AutoShapes in Microsoft Office
- animated movies
- encapsulated postscript (EPS)
- animation programmes such as Blender and Adobe After Effects
- image manipulation programmes such as Adobe Photoshop and GIMP
- Adobe portable document format (PDF)
- Windows meta-file (WMF)

Image display

When a monitor or a printer displays a vector image it is **rasterised**- converted into a grid of pixels. Regardless of the file type, an image will always be outputted onto a screen or printed in pixels.

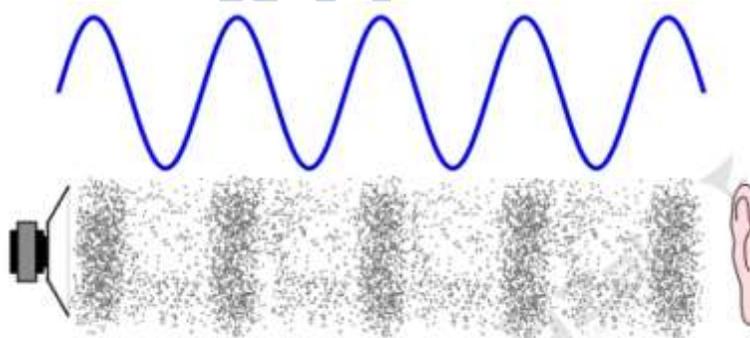
Syllabus Content:

1.1.3 Sound

-  show understanding of how sound is represented and encoded
-  use the associated terminology: sampling, sampling rate, sampling resolution
-  show understanding of how file sizes depend on sampling rate and sampling resolution
-  understanding of how typical features found in sound editing software are used in practice

(1.1.3) Sound:

Sound is an oscillation of pressure transmitted through a solid, liquid, or gas (there is no sound in outer space as space is a vacuum and there is no solid, liquid or gas to transmit sound through!). A speaker works by moving its center cone in and out; this causes the air particles to bunch together forming waves. These waves spread out from the speaker travelling at 340 m/s. If your ear is in the way, then the waves of sound particles will collide with your ear drum, vibrating it and sending a message to your brain. This is how you hear:



Natural sound consists of variations in pressure which is detected by the human ear. A typical sound contains a large number of individual waves each with a defined frequency. The result is a wave form in which the amplitude of the sound varies in a continuous but irregular pattern.

If there is a need to store sound or transmit it electronically the original analogue sound signal has to be converted to a binary code.

A sound encoder has two components.

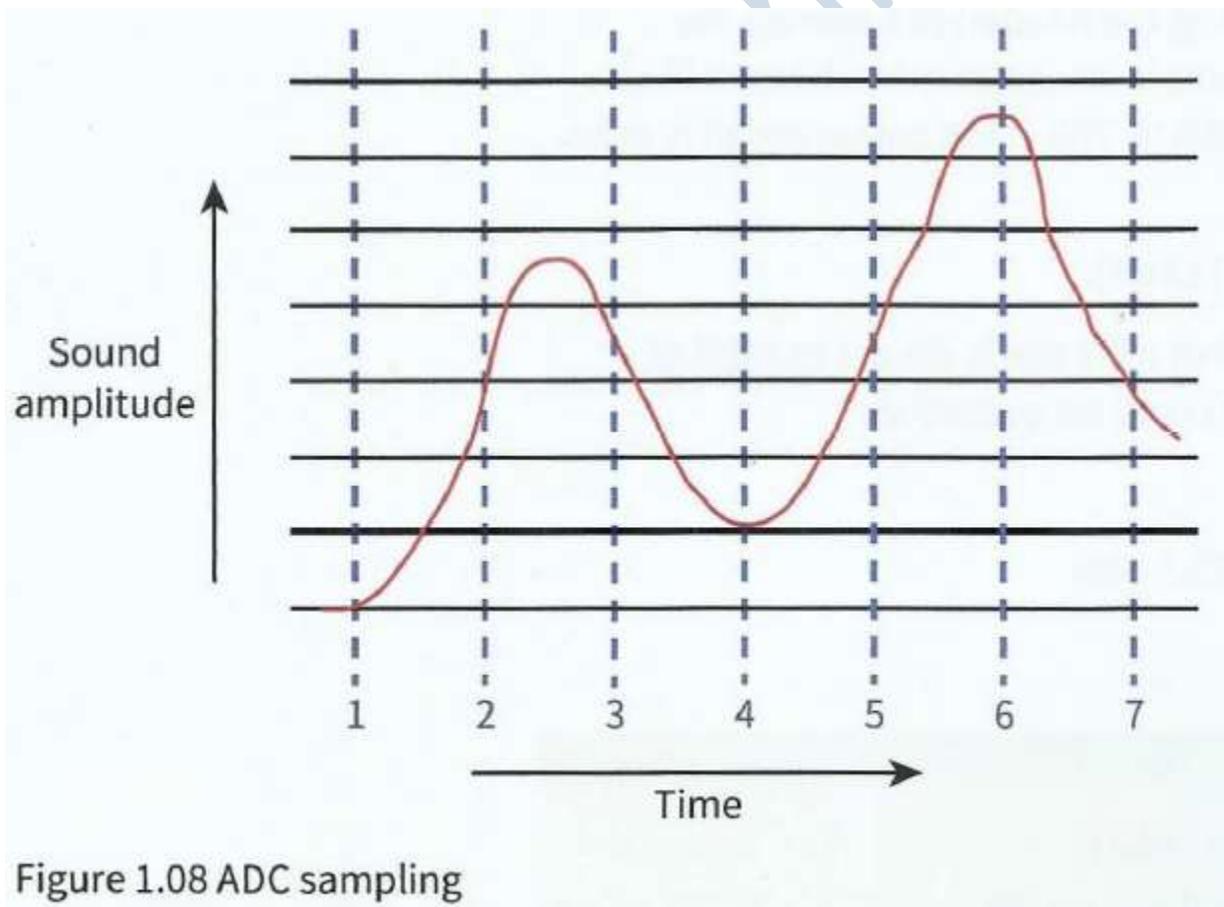
1. Band limiting Filter
2. Analogue to digital converter.

Band-limiting filter:

This is needed to remove high-frequency components. The ear would not be able to detect these and they could cause problems for the coding if not removed.

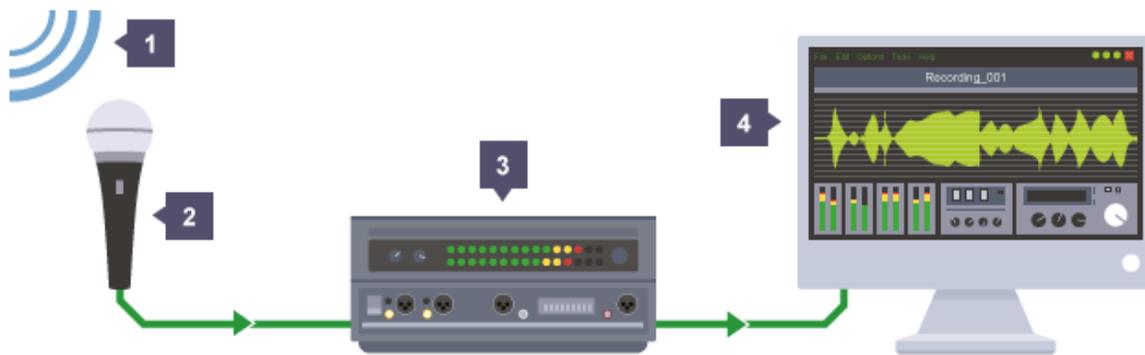
Analogue-to-digital converter (ADC).

The method of operation of the ADC is described with reference to Figure. The amplitude of the wave (the red line) has to be sampled at regular intervals. The blue vertical lines indicate the sampling times. The amplitude cannot be measured exactly; instead the amplitude is approximated by the closest of the defined amplitudes represented by the horizontal lines. In Figure 1.08, sample values 1 and 4 will be an accurate estimate of the actual amplitude because the wave is touching an amplitude line. In contrast, samples 5 and 6 will not be accurate because the actual amplitude is approximately half way between the two closest defined values.



Analogue and digital:

For a computer to store sound files we need to get the continuous analogue sound waves into discrete binary values:



- 1 Microphone measures change in air pressure
- 2 Microphone translates air pressure into electrical voltage
- 3 Analogue to Digital Converter digitises the electrical voltage to bytes of information
- 4 Computer displays the digitised sound for manipulation

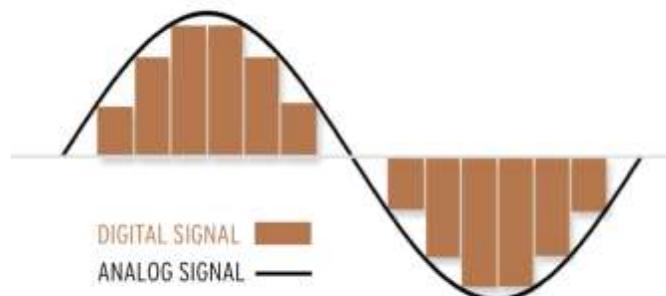
Sound input through a microphone is converted to digital for storage and manipulation

An analogue sound wave is picked up by a microphone and sent to an Analogue to Digital (ADC) converter in the form of analogue electrical signals. The ADC converts the electrical signals into digital values which can be stored on a computer.

Once in a digital format you can edit sounds with programs such as **audacity**. To play digital audio you convert the sound from digital values into analogue electrical signals using the

DAC, these signals are then passed to a speaker that vibrating the speaker cone, moving the air to create sound waves and analogue noise.

Digital audio: Sounds created on a computer exist as digital information encoded as audio files.



Analogue to Digital Converter (ADC)

Converts analogue sound into digital signals that can be stored on a computer.

Digital to Analogue Converter (DAC)

Converts digital signals stored on a computer into analogue sound that can be played through devices such as speakers.

This section will cover how we record, store and transmit sound using computers.

Sound waves in nature are **continuous**, this means they have an almost infinite amount of detail that you could store for even the shortest sound. This makes them very difficult to record perfectly, as computers can only store **discrete** data, data that has a limited number of data points

Audio Bandwidth:

Bandwidth is the ability of a recorded signal to be reproduced at varying degrees of resolution. Think of it like enlarging a low-resolution image versus a high-resolution image. After a certain point, enlarging a lower-resolution image will become pixelated and difficult to see, where the hi-resolution image will resize clearly.

Like images, audio signals can have a limited bandwidth if recorded digitally. Once a digital recording is made, the bandwidth is set in place.

An analog recording is considered unlimited. Therefore, it can move to a higher and higher resolution without losing its original quality.

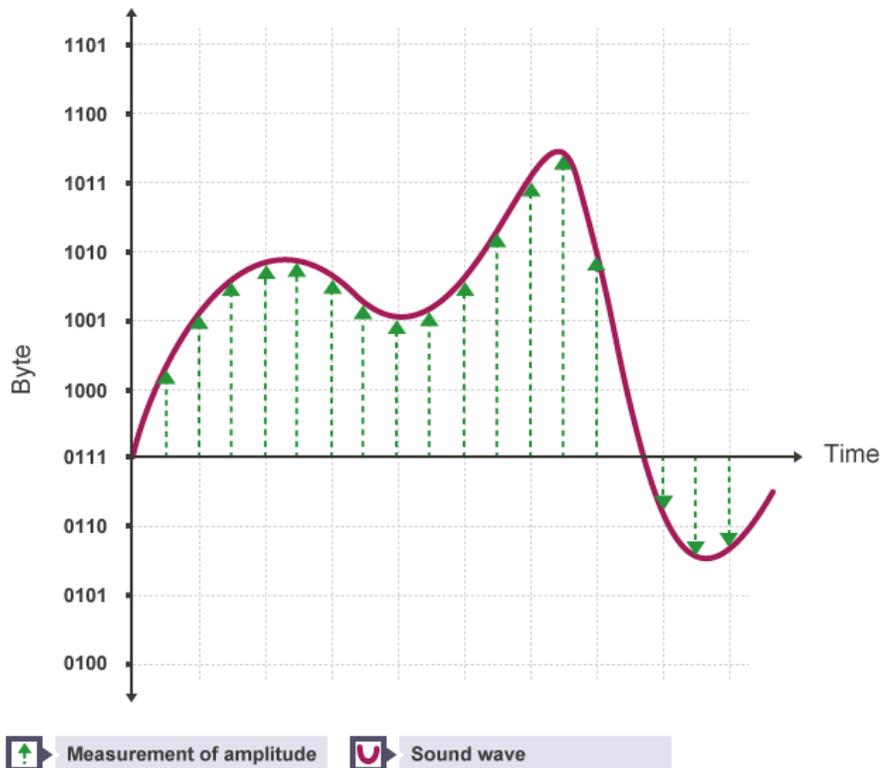
Digital sound is broken down into thousands of samples per second. Each sound sample is stored as binary data.

Digital audio quality Factors that affect the quality of digital audio include:

-  **sample rate** - the number of audio samples captured every second
-  **sample resolution/ bit depth:** - the number of bits assigned to each sample of sound
-  **bit rate** - the number of bits used to store one second of sound

Sound Sampling:

- 📌 amplitude of sound wave taken at different points in time
- 📌 measurement of value of analogue signal at regular time intervals/a point in time

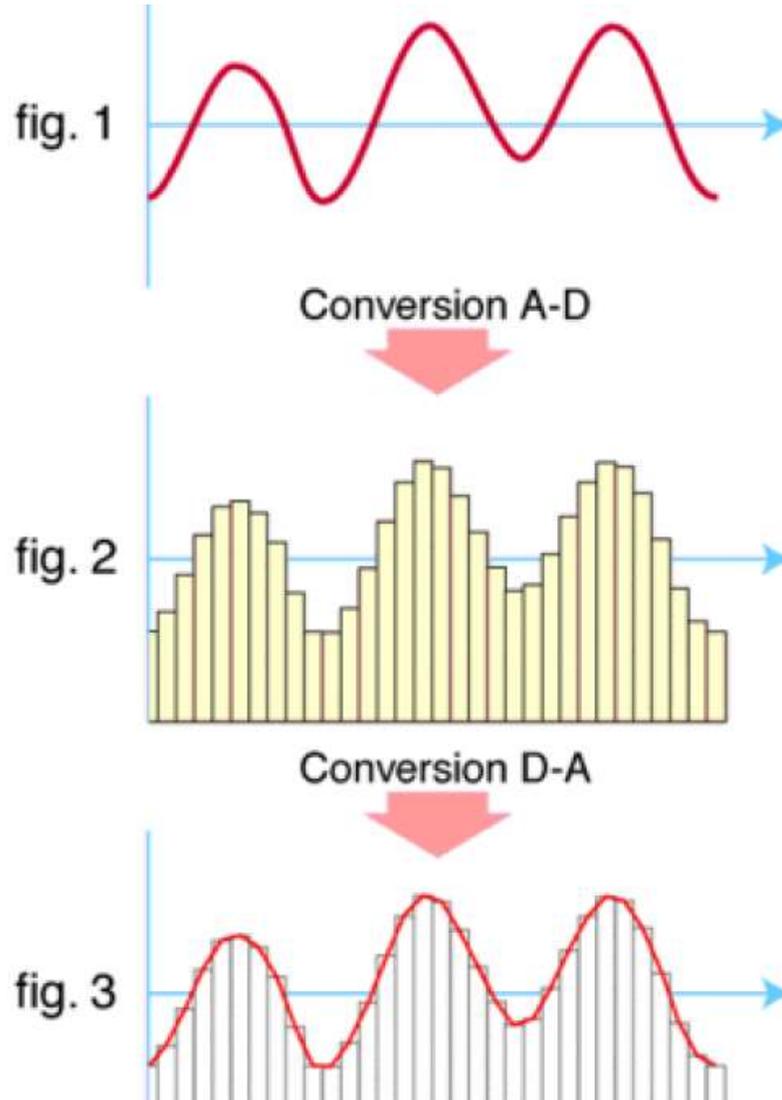


Sample rate

The **sample rate** is how many samples, or measurements, of the sound are taken each second.

The more samples that are taken, the more detail about where the waves rise and fall is recorded and the higher the quality of the audio. Also, the shape of the sound wave is captured more accurately.

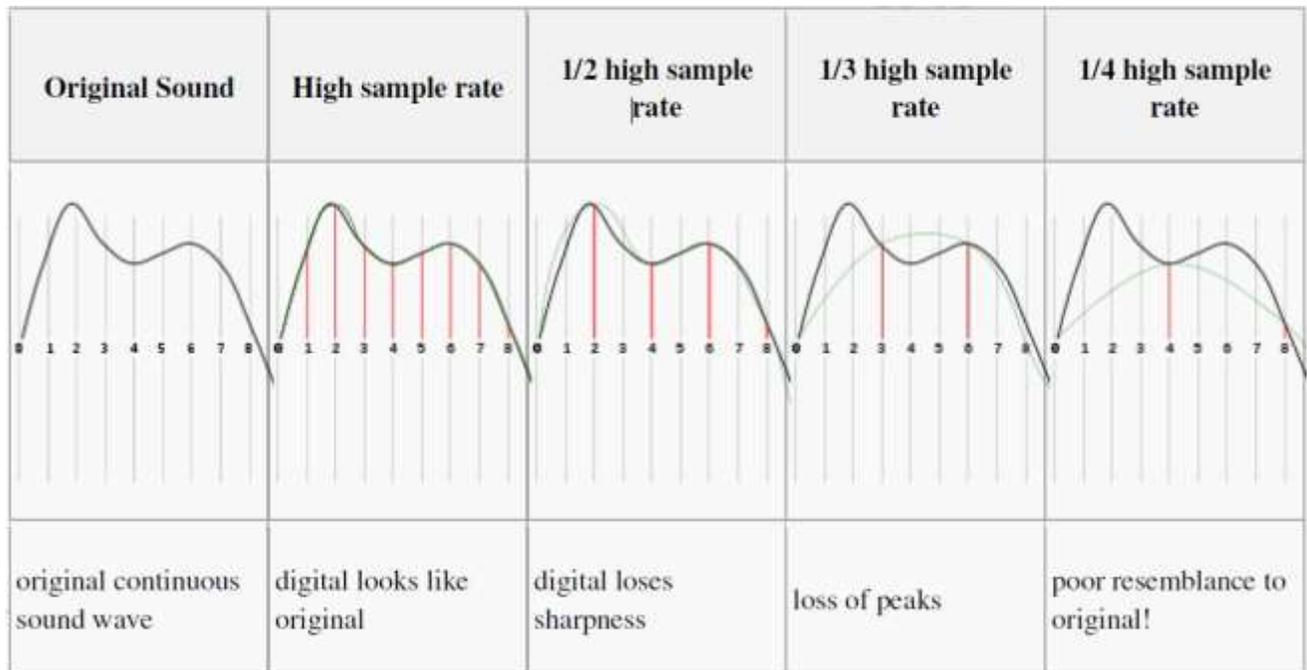
Each sample represents the amplitude of the digital signal at a specific point in time. The amplitude is stored as either an **integer** or a **floating point number** and encoded as a **binary number**.



1. The original analogue sound wave is a continuous set of points
2. ADC converts sound into digital data
3. DAC converts digital data into analogue sound, the analogue wave produced may differ significantly from the original sound wave.

Sound Sampling

So we should know by now that sound waves are continuous and computers can only store discrete data. How exactly does an Analogue to Digital Converter convert a continuous sound wave into discrete digital data? To do this we need to look at how computers sample sound.



Sampling Rate - The number of samples taken per second

A higher sampling rate means the waveform will be converted from analog to digital form more accurately.

Sampling Resolution (bit depth): no. of bits assigned to each sample.

The sampling resolution allows you to set the range of volumes storable for each sample.

The quality of a sound produced using a sound sampler depends on the sampling rate and sampling resolution.

-  The higher the sampling resolution (bit depth), the higher the quality of the audio.
-  Higher sampling rate/resolution means larger file size

Bit depth is the number of bits available for each sample.

Sampling resolution (Bit depth) is usually 16 bits on a CD and 24 bits on a DVD.

Bit Rate: number of bits required to store 1 second of sound.

$$\text{Bit Rate} = \text{Sampling Rate} \times \text{Sampling Resolution}$$

$$\text{File Size} = \text{Bit Rate} \times \text{Length of Sound}$$

File Size = Sample Rate X Sample Resolution X Length of sound

Calculating bit rate

The bit rate is calculated using the formula:

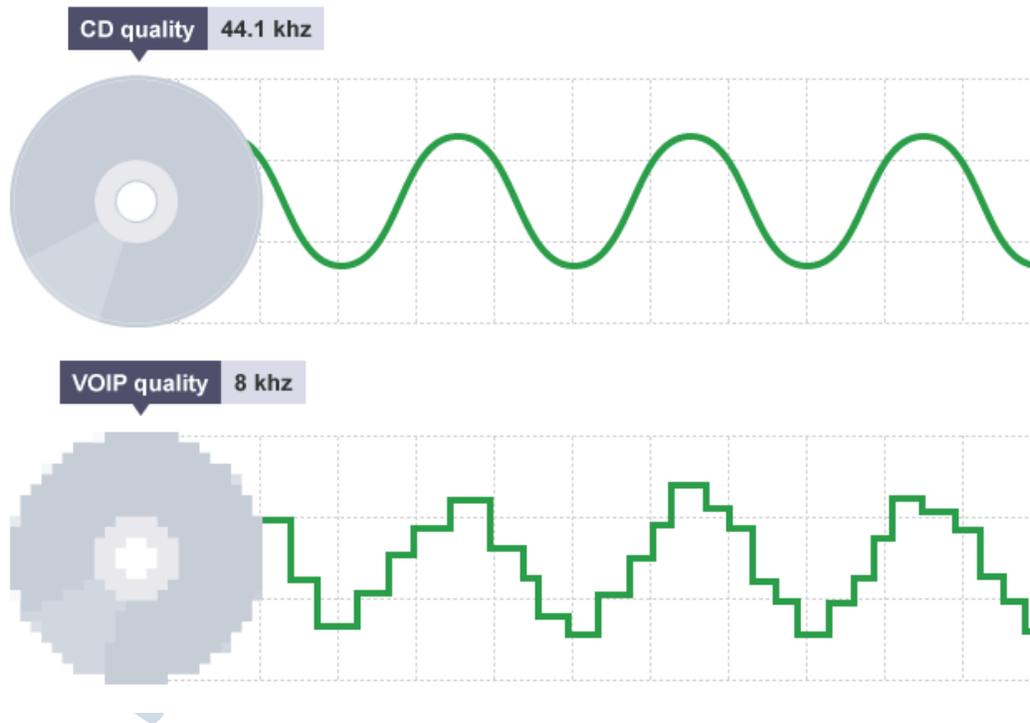
$$\text{Frequency} \times \text{bit depth} \times \text{channels} = \text{bit rate}$$

A typical, uncompressed high-quality audio file has a sample rate of 44,100 samples per second, a sample resolution (bit depth) of 16 bits per sample and 2 channels of stereo audio. The bit rate for this file would be:

$$44,100 \text{ samples per second} \times 16 \text{ bits per sample} \times 2 \text{ channels} = 1,411,200 \text{ bits per second (or 1,411.2 kbps)}$$

A four-minute (240 second) song at this bit rate would create a file size of:

$$1,411,200 \times 240 = 338,688,000 \text{ bits (or 40.37 megabytes)}$$



A common audio **sample rate** for music is **44,100 samples per second**. The unit for the sample rate is hertz (Hz). 44,100 samples per second is 44,100 hertz or 44.1 kilohertz (kHz).

Telephone networks and VOIP services can use a sample rate as low as 8 kHz. This uses less data to represent the audio. At 8 kHz, the human voice can still be heard clearly - but music at this sample rate would sound low quality.

Compression

Compression is a useful tool for reducing file sizes. When images, sounds or videos are compressed, data is removed to reduce the file size. This is very helpful when streaming and downloading files.

Streamed music and downloadable files, such as MP3s, are usually between 128 kbps and 320 kbps - much lower than the 1,411 kbps of an uncompressed file.

Videos are also compressed when they are streamed over a network. Streaming HD video requires a high-speed internet connection. Without it, the user would experience buffering and regular drops in quality. HD video is usually around 3 mbps. SD is around 1,500 kbps.

Compression can be **lossy** or **lossless**.

Lossless compression:

Lossless compression means that as the file size is compressed, the audio **quality remains the same** - it does not get worse. Also, the file can be restored back to its original state. **FLAC** and **ALAC** are open source lossless compression formats. Lossless compression can reduce file sizes by up to 50% without losing quality.

Lossless Compression allows original data to be perfectly reconstructed from compression.

Run-length encoding: compression in which sequences with same data value in many consecutive values are stored as a single data value and count

E.g. **00001234111111** can be written as **(0-4)1234(1-6)**

Lossy compression:

Lossy compression permanently removes data. For example, a WAV file compressed to an **MP3** would be lossy compression. The bit rate could be set at 64 kbps, which would reduce the size and quality of the file. However, it would not be possible to recreate a 1,411 kbps quality file from a 64 kbps MP3.

With lossy compression, the original bit depth is reduced to remove data and reduce the file size. The bit depth becomes **variable**.

MP3 and **AAC** are lossy compressed audio file formats widely supported on different platforms. MP3 and AAC are both patented codecs. **Ogg Vorbis** is an open source alternative for lossy compression.

Lossy Compression is in which file accuracy is low, but file size is smaller than lossless

Perceptual coding: works by reducing certain parts of a sound which are less audible to human hearing.

Sound Editing

Simply recording sound and storing a digital representation is not enough for many applications.

Once a digital representation of the sound has been stored in a file, it can be manipulated using sound-editing software. This will typically have features for:

- combining sound from different sources
- fading in or fading out the sound
- editing the sound to remove noise and other imperfections.

If you are interested in sound editing you can start editing your own music using a program called **Audacity**.

Using Audacity you can create your own sound samples with different sample rates and sample resolutions, listening to the difference between them and noting the different file sizes.

Features

This is a list of features in Audacity, the free, open source, cross-platform audio editor.

Recording

Audacity can record live audio through a microphone or mixer, or digitize recordings from cassette tapes, records or minidisks. With some sound cards, and on any Windows Vista, Windows 7 or Windows 8 machine, Audacity can also capture streaming audio.

Import and Export

Import sound files, edit them, and combine them with other files or new recordings. Export your recordings in many different file formats, including multiple files at once.

Sound Quality

Supports 16-bit, 24-bit and 32-bit (floating point) samples (the latter preserves samples in excess of full scale).

Sample rates and formats are converted using high-quality re-sampling and dithering. Tracks with different sample rates or formats are converted automatically in real time.

Editing:

-  Easy editing with Cut, Copy, Paste and Delete.
-  Unlimited sequential Undo (and Redo) to go back any number of steps.
-  Edit and mix large numbers of tracks.
-  Multiple clips are allowed per track.
-  Label tracks with selectable Sync-Lock Tracks feature for keeping tracks and labels synchronized.
-  Draw Tool to alter individual sample points.
-  Envelope Tool to fade the volume up or down smoothly.
-  Automatic Crash Recovery in the event of abnormal program termination.

Accessibility

Tracks and selections can be fully manipulated using the keyboard.

-  Large range of keyboard shortcuts.
-  Excellent support for JAWS, NVDA and other screen readers on Windows, and for VoiceOver on Mac.

Effects

-  Change the pitch without altering the tempo (or vice-versa).
-  Remove static, hiss, hum or other constant background noises.
-  Alter frequencies with Equalization, Bass and Treble, High/Low Pass and Notch Filter effects.
-  Adjust volume with Compressor, Amplify, Normalize, Fade In/Fade Out and Adjustable Fade effects.
-  Remove Vocals from suitable stereo tracks.
-  Create voice-overs for podcasts or DJ sets using Auto Duck effect. Other built-in effects include:
 -  Echo
 -  Paulstretch (extreme stretch)
 -  Phaser
 -  Reverb
 -  Reverse
 -  Truncate Silence
 -  Wahwah.
-  Run "Chains" of effects on a project or multiple files in Batch Processing mode.

Syllabus Content:

1.1.4 Video

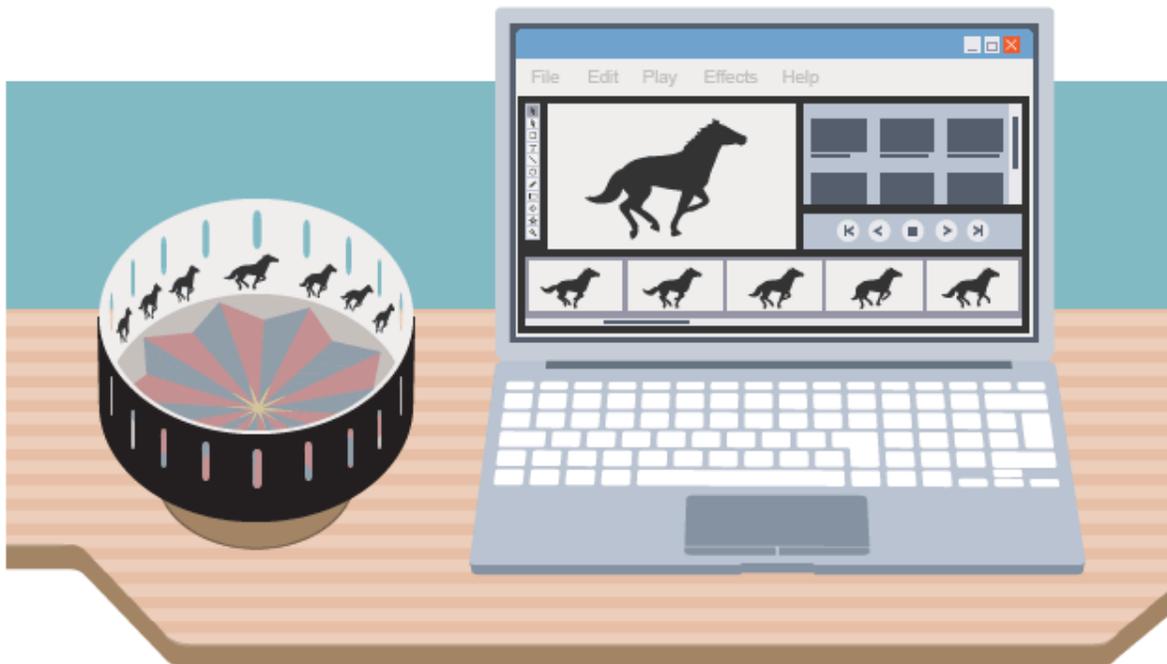
- Show understanding of the characteristics of video streams:
 - the frame rate (frames/second)
 - interlaced and progressive encoding
 - video interframe compression algorithms and spatial and temporal redundancy
 - multimedia container formats

1.1.5 Compression techniques

- show understanding of how digital data can be compressed, using either 'lossless' (including runlength encoding – RLE) or 'lossy' techniques

(1.1.4) Digital video

A digital film is created from a series of static images played at a high speed. Digital films are usually around 24 frames per second but can be anything up to around 100 frames per second or more.



Video is an electronic medium for the recording, copying and broadcasting of moving visual images characteristics of video streams.

The emphasis here is on the visual aspect of a video recording and, in particular, how the image is displayed on a screen.

Video is not stored very simply as a succession of still images or frames and the only concern would be the frame rate defined as the number of frames displayed per second.

In practice the issues are far more complex.

-  The basic principle of operation is that the display of an individual frame is created line by line. One of the issues is the choice of resolution.
-  The resolution can be defined in terms of the number of lines per frame and the number of pixels per line.
-  There needs to be compatibility between the resolution of the stored image and the resolution of the display screen. However, the technology used has to be chosen with regard to the sensitivity of the human eye.
-  One constraint is that unless the screen is refreshed at least 50 times per second the eye will notice the flicker.
-  However, provided that the refresh rate is 25 times per second the eye cannot see that any motion on the screen is not actually continuous.

Interlaced encoding:

The traditional solution to this problem has been to use interlaced encoding. This was used in television broadcasting and then adapted for video recordings.

-  The image for each frame is split into two halves, one containing the odd numbered lines and the other is even.
-  An odd field (upper field) consisting of the odd-numbered lines and an even field (lower field) consisting of the even-numbered lines.
-  The first half is displayed completely then the second half follows.
-  This produces what appears to the eye as being a high refresh rate but is halving the transmission bandwidth requirements.
-  Interlacing was invented as a way to reduce flicker in early mechanical and CRT video displays without increasing the number of complete frames per second, which would have required sacrificing image detail in order to remain within the limitations of a narrow bandwidth.

Progressive encoding:

The alternative approach is to use progressive encoding

-  Full frame is displayed each time.
-  As improved transmission bandwidths become more generally available it is likely that progressive encoding will become the norm.
-  Each refresh period updates all of the scan lines of each frame in sequence.

Interlaced encoding:

-  The data from a single frame are encoded as two separate **fields**.
-  One containing the data for the even numbered rows / lines and the other has the data for the odd numbered rows / lines.
-  The image is rendered by alternating between the even field and the odd field (of each successive frame).
-  The viewer sees data from two frames simultaneously
-  The rate of picture display (the field rate) is twice the rate of image frame display (the frame rate).
-  Originally used in television broadcasting and adapted for video recordings.
-  Produces what appears to the eye to be a high refresh rate.
-  Halves the transmission bandwidth requirements.

Progressive encoding

-  Stores the data for an entire frame and displays all the frame data at the same time.
-  The rate of picture display is the same as the frame rate.
-  Used by traditional film / video digitised from a film camera / computer displays progressive encoding.
-  High bandwidth requirements.

Frame Rate (FPS):

frequency at which frames in a video sequence are displayed

Higher frame rate, better quality video = greater size

Number of frames per second

Frame rate, the number of still pictures per unit of time of video, ranges from six or eight frames per second (frame/s) for old mechanical cameras to 120 or more frames per second for new professional cameras.

-  PAL (Europe, Asia, Australia, etc.) and SECAM (France, Russia, parts of Africa etc.) standards specify 25 frame/s,
-  NTSC (USA, Canada, Japan, etc.) specifies 29.97 frame/s.
-  Film is shot at the slower frame rate of 24 photograms/s, which complicates slightly the process of transferring a cinematic motion picture to video.

The minimum frame rate to achieve a comfortable illusion of a moving images about sixteen frames per second.

Inter-frame compression:

-  Type of video compression and decreases file size
-  It removes neighboring frames which are similar
-  Some change in image data is redundant
-  How redundant the change in image between frames determines the amount of compression possible.

Redundancy: refers to the amount of wasted space consumed by storage media to record picture information in a digital image.

Spatial redundancy: elements that are duplicated within a structure, such as pixels in a still image and bit patterns in a file. Exploiting spatial redundancy is how compression is performed.

Temporal redundancy: pixels in two video frames that have the same values in the same location. Exploiting temporal redundancy is one of the primary techniques in video compression.

Spatial and temporal redundancy in MPEG

Chapter 1 introduced these concepts in a general sense and now they will be treated with specific reference to MPEG. Figure 5.18(a) shows that spatial redundancy is redundancy within a single picture or object, for example repeated pixel values in a large area of blue sky. Temporal redundancy (b) exists between successive pictures or objects.

In MPEG, where temporal compression is used, the current picture/ object is not sent in its entirety; instead the difference between the current picture/object and the previous one is sent. The decoder already has the previous picture/object, and so it can add the difference, or residual image, to make the current picture/object. A residual image is created by subtracting every pixel in one picture/object from the corresponding pixel in another. This is trivially easy when pictures are restricted to progressive scan, as in MPEG-1, but MPEG-2 had to develop greater complexity (continued in MPEG-4) so that this can also be done with interlaced pictures. The handling of interlace in MPEG will be detailed later.

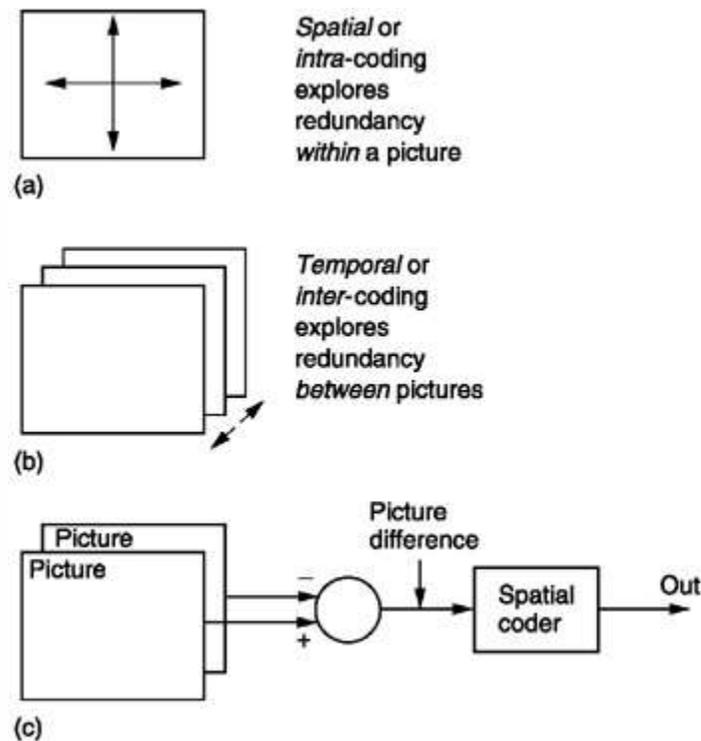


Figure 5.18: (a) Spatial or intra-coding works on individual images. (b) Temporal or inter-coding works on successive images. (c) In MPEG inter-coding is used to create difference images. These are then compressed spatially.

A residual is an image of a kind, although not a viewable one, and so should contain some kind of spatial redundancy. Figure 5.18(c) shows that MPEG takes advantage of both forms of redundancy. Residual images are spatially compressed prior...

Temporal redundancy:

Pixels in two video frames that have the same values in the same location. Exploiting temporal redundancy is one of the primary techniques in video compression.

Spatial redundancy:

Elements that are duplicated within a structure, such as pixels in a still image and bit patterns in a file. Exploiting spatial redundancy is how compression is performed. Contrast with **temporal redundancy**.

Multimedia Container Formats:

- Contains different types of data
- Can be audio or video or both codecs
- This interleaves the different type of data
- The video is compressed into codecs
- E.g. .avi, .mov, .mp4, .ogg, .rm

Lossy Data Compression:

-  These programs eliminate unnecessary bits of data
-  Reduces data file size more than lossless
-  Unable to get back original version

Lossless Data Compression:

-  Breaks data into smaller form of transmission
-  Allows recreation of original

References:

<https://www.bbc.com/bitesize/guides/zgyrq6f/revision/2>
https://en.wikipedia.org/wiki/Vector_graphics
https://en.wikipedia.org/wiki/File:Bitmap_VS_SVG.svg
<https://www.corporate3design.com/blog/112/>
<http://etc.usf.edu/techease/win/images/what-is-the-difference-between-bitmap-and-vector-images/>
<https://www.klipsch.com/blog/digital-vs-analog-audio>
<https://www.bbc.com/education/guides/z7vc7ty/revision/2>
<http://www.globalspec.com/reference/37829/203279/5-7-spatial-and-temporal-redundancy-in-mpeg>
<https://robbfoxx.wordpress.com/2015/07/12/discussion-6-2-1-what-is-redundancy-temporal-redundancy-and-spatial-redundancy/>
<http://www.pcmag.com/encyclopedia/term/56086/spatial-redundancy>
http://christmanwd.weebly.com/uploads/2/5/9/5/25956852/805063_orig.jpg
<https://microscope-microscope.org/microscope-info/image-resolution/>

Cambridge International AS & A level by Silvia Langfield and Dave Duddell.