Topic: 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:

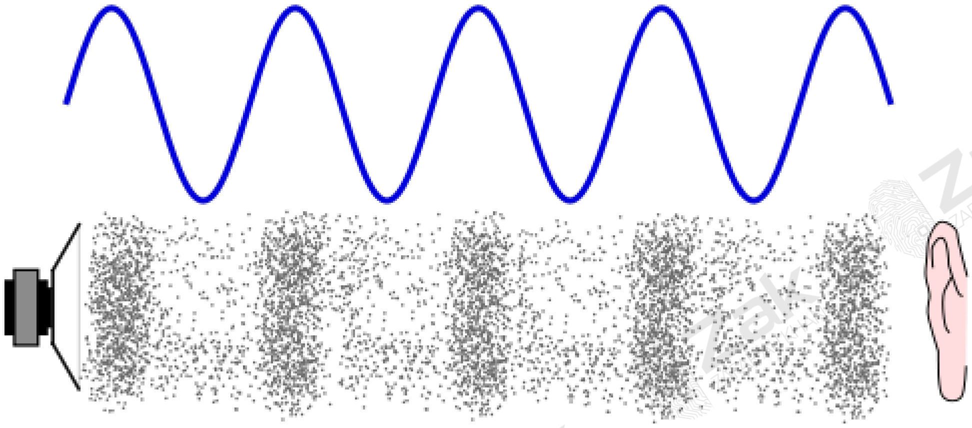

When you hear different volumes and pitches of sound all that is happening is that each sound wave varies in energy for the volume (larger energy waves, the louder the sound), or distance between sound waves which adjusts the pitch, (smaller distances between waves leads to higher pitched sound).

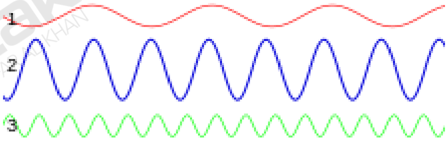

1 - Base volume and frequency
2 - double volume and frequency
3 - same volume treble the frequency

Topic: 1.1.3 Sound

Sound is often recorded for two channels, stereo; feeding a left and right speaker whose outputs may differ massively. Where one channel is used, this is called mono. 5.1 surround sound used in cinemas and home media set ups use 6 channels.

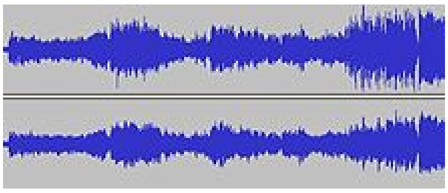

A computer representation of a stereo song, if you look carefully you'll see the volume of the song varying as you go through it

This section of the book 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.

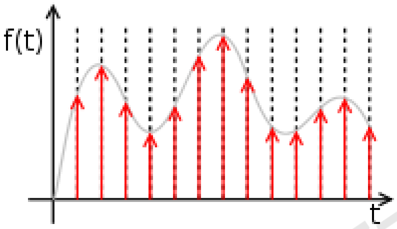 	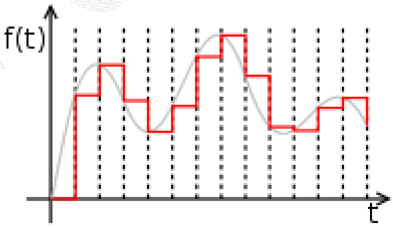
Sound is a continuous set of data points formed by a wave. Computers sample this sound at discrete points to store a digital approximation	The discrete approximations (in red) can be used to recreate the original sound (grey). However, due to limitations in the number of samples we take we are often unable to truly represent a sound wave, though we can get close enough for the human ear not to notice the difference.

Topic: 1.1.3 Sound

Analogue and digital

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

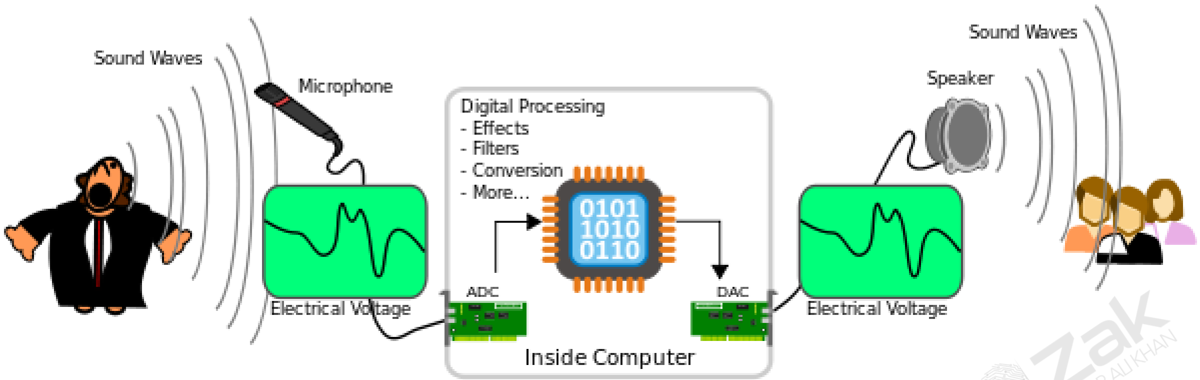

An analogue sound wave is picked up by a microphone and sent to an Analogue to Digital (ADC) converter in the form of analogue electrical signals. The ADC converts the electrical signals into digital values which can be stored on a computer.

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.

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

Topic: 1.1.3 Sound

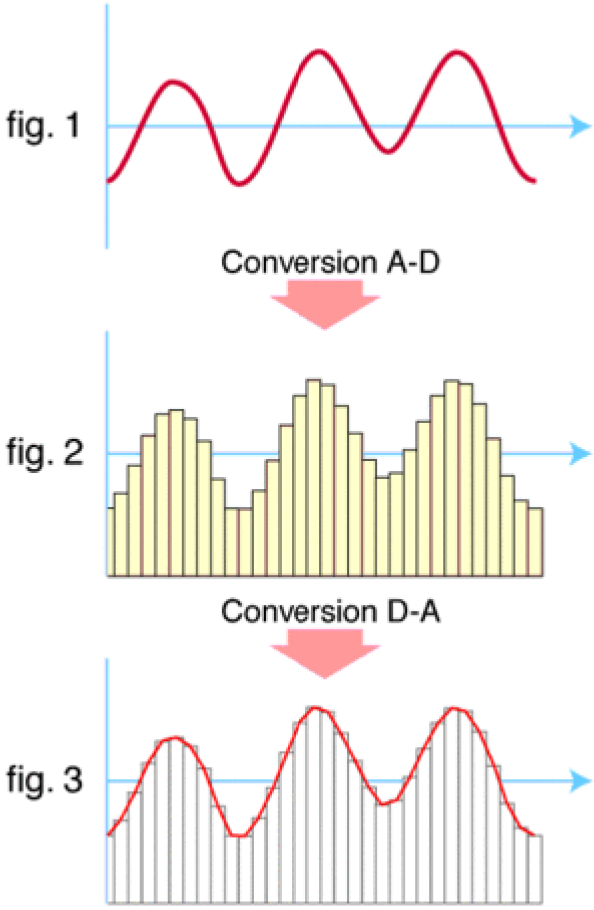

fig 1. The original analogue sound wave is a continuous set of points

fig 2. ADC converts sound into digital data

fig 3. DAC converts digital data into analogue sound, the analogue wave produced may differ significantly from the original sound wave

Topic: 1.1.3 Sound

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

Hertz (Hz) - the SI unit of frequency defined as the number of cycles per second of a periodic phenomenon

To create digital music that sounds close to the real thing you need to look at the analogue sound waves and try to represent them digitally. This requires you to try to replicate the analogue (and continuous) waves as discrete values. The first step in doing this is deciding how often you should sample the sound wave, if you do it too little, the sample stored on a computer will sound very distant from the one being recorded. Sample too often and sound stored will resemble that being recorded but having to store each of the samples means you'll get very large file sizes. To decide how often you are going to sample the analogue signal is called the sampling rate. Take a look at the following example:

Original Sound	High sample rate	1/2 high sample rate	1/3 high sample rate	1/4 high sample rate
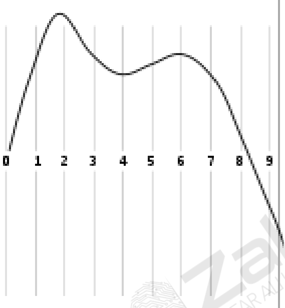	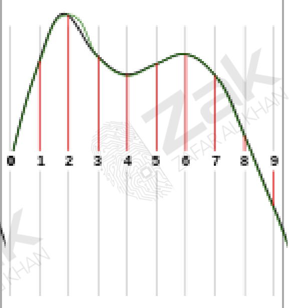	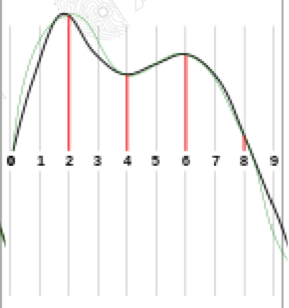	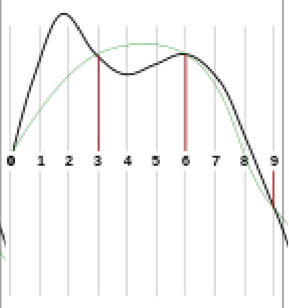	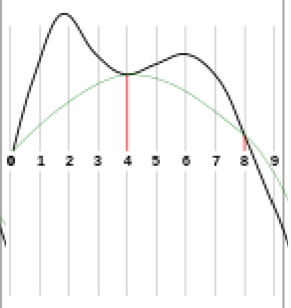
original continuous sound wave	digital looks like original	digital loses sharpness	loss of peaks	poor resemblance to original!

To create digital sound as close to the real thing as possible you need to take as many samples per second as you can. When recording MP3s you'll normally use a sampling rate between 32,000, 44,100 and 48,000Hz (samples per second). That means that for a sampling rate of 44,100, sound waves will have been sampled 44,100 times per second! Recording the human voice requires a lower sampling rate, around 8,000Hz.

Topic: 1.1.3 Sound

If you speak to someone on the phone it may sound perfectly acceptable, but try playing music down a telephone wire and see how bad it sounds.

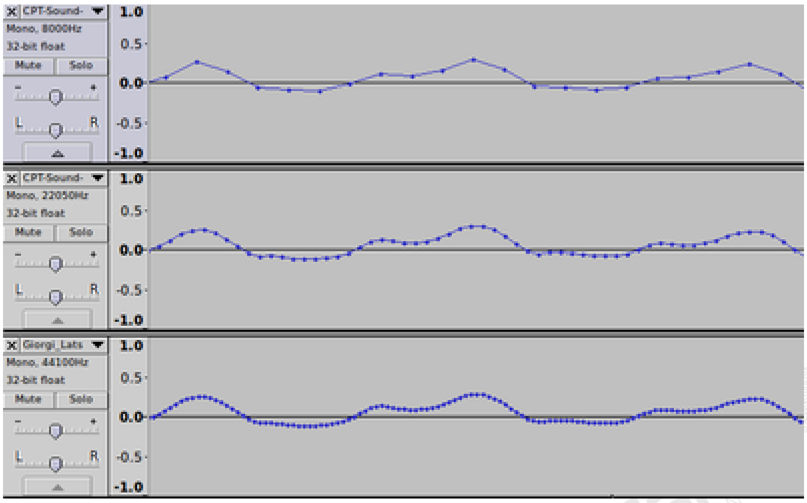

Comparison of the same sound sample recorded at 8kHz, 22kHz and 44kHz sample rate. Note the spacing of the data points for each sample. The higher the sample rate the more data points we'll need to store

Sampling resolution

Sampling resolution - the number of bits assigned to each sample

As you saw earlier, different sounds can have different volumes. The sampling resolution allows you to set the range of volumes storable for each sample. If you have a low sampling resolution then the range of volumes will be very limited, if you have a high sampling resolution then the file size may become unfeasible. The sampling resolution for a CD is 16 bits used per sample.

Topic: 1.1.3 Sound

File sizes

Bit rate - the number of bits required to store 1 second of sound

To work out the size of a sound sample requires the following equation:

```
File Size = Sample Rate * Sample Resolution * Length of sound
```

This is the same as saying:

```
File Size = Bit Rate * Length of sound
```

Let's look at an example:

Example: Sound File Sizes

If you wanted to record a 30 second voice message on your mobile phone you would use the following:

```
Sample Rate = 8,000Hz
Sample Resolution = 16 bit
Length of Sound = 30 seconds
```

Therefore the total file size would be:

```
8,000 * 16 * 30 = 3 840 000 Bits = 480 000 Bytes
```

Sound Editing

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 minidiscs. With some sound cards, and on any Windows Vista, Windows 7 or Windows 8 machine, Audacity can also capture streaming audio.

Topic: 1.1.3 Sound

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 resampling 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.

Topic: 1.1.3 Sound

- 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.

Plug-ins

- Support for LADSPA, Nyquist, VST and Audio Unit effect plug-ins.
- Effects written in the Nyquist programming language can be easily modified in a text editor - or you can even write your own plug-in.

Analysis

- Spectrogram view modes for visualizing frequencies.
- "Plot Spectrum" command for detailed frequency analysis.
- "Sample Data Export" for exporting a file containing amplitude values for each sample in the selection.
- Contrast Analysis for analyzing average rms volume differences between foreground speech and background music.
- Support for adding VAMP analysis plug-ins.

Free and Cross-Platform

- Licensed under the GNU General Public License (GPL).
- Runs on Windows, Mac OS X and GNU/Linux