

[Back to the main page](#)[Back to the Tutorial Page](#)[Digital Audio](#) [Rules of Audacity](#) [Setup, Audio Import and Playback](#) [Recording with Audacity](#)


Tutorial - I.Basics

Part 1 - Digital Audio - Part 1

What is sound?


Sounds are pressure waves of air. If there wasn't any air, we wouldn't be able to hear sounds. There's no sound in space.

We hear sounds because our ears are sensitive to these pressure waves. Perhaps the easiest type of sound wave to understand is a short, sudden event like a clap. When you clap your hands, the air that was between your hands is pushed aside. This increases the air pressure in the space near your hands, because more air molecules are temporarily compressed into less space. The high pressure pushes the air molecules outwards in all directions at the speed of sound, which is about 340 meters per second. When the pressure wave reaches your ear, it pushes on your eardrum slightly, causing you to hear the clap.



A hand clap is a short event that causes a single pressure wave that quickly dies out. The image above shows the waveform for a typical hand clap. In the waveform, the horizontal axis represents time, and the vertical axis is for pressure. The initial high pressure is followed by low pressure, but the oscillation quickly dies out.

The other common type of sound wave is a periodic wave. When you ring a bell, after the initial strike (which is a little like a hand clap), the sound comes from the vibration of the bell. While the bell is still ringing, it vibrates at a particular frequency, depending on the size and shape of the bell, and this causes the nearby air to vibrate with the same frequency. This causes pressure waves of air to travel outwards from the bell, again at the speed of sound. Pressure waves from continuous vibration look more like this:



How is sound recorded?


A microphone consists of a small membrane that is free to vibrate, along with a mechanism that translates movements of the membrane into electrical signals. (The exact electrical mechanism varies depending on the type of microphone.) So acoustical waves are translated into electrical waves by the microphone. Typically, higher pressure corresponds to higher voltage, and vice versa.

A tape recorder translates the waveform yet again - this time from an electrical signal on a wire, to a magnetic signal on a tape. When you play a tape, the process gets performed in reverse, with the magnetic signal transforming into an electrical signal, and the electrical signal causing a speaker to vibrate, usually using an electromagnet.

How is sound recorded digitally ?

Recording onto a tape is an example of analog recording. Audacity deals with digital recordings - recordings that have been sampled so that they can be used by a digital computer, like the one you're using now. Digital recording has a lot of benefits over analog recording. Digital files can be copied as many times as you want, with no loss in quality, and they can be burned to an audio CD or shared via the Internet. Digital audio files can also be edited much more easily than analog tapes.

The main device used in digital recording is a Analog-to-Digital Converter (ADC). The ADC captures a snapshot of the electric voltage on an audio line and represents it as a digital number that can be sent to a computer. By capturing the voltage thousands of times per second, you can get a very good approximation to the original audio signal:



Each dot in the figure above represents one audio *sample*. There are two factors that determine the quality of a digital recording:

- **Sample rate:** The rate at which the samples are captured or played back, measured in Hertz (Hz), or samples per second. An audio CD has a sample rate of 44,100 Hz, often written as 44

KHz for short. This is also the default sample rate that Audacity uses, because audio CDs are so prevalent.

- **Sample format** or **sample size**: Essentially this is the number of digits in the digital representation of each sample. Think of the sample rate as the horizontal precision of the digital waveform, and the sample format as the vertical precision. An audio CD has a precision of 16 bits, which corresponds to about 5 decimal digits.

Higher sampling rates allow a digital recording to accurately record higher frequencies of sound. The sampling rate should be at least twice the highest frequency you want to represent. Humans can't hear frequencies above about 20,000 Hz, so 44,100 Hz was chosen as the rate for audio CDs to just include all human frequencies. Sample rates of 96 and 192 KHz are starting to become more common, particularly in DVD-Audio, but many people honestly can't hear the difference.

Higher sample sizes allow for more dynamic range - louder louds and softer softs. If you are familiar with the decibel (dB) scale, the dynamic range on an audio CD is theoretically about 90 dB, but realistically signals that are -24 dB or more in volume are greatly reduced in quality. Audacity supports two additional sample sizes: 24-bit, which is commonly used in digital recording, and 32-bit *float*, which has almost infinite dynamic range, and only takes up twice as much storage as 16-bit samples.

Playback of digital audio uses a Digital-to-Analog Converter (DAC). This takes the sample and sets a certain voltage on the analog outputs to recreate the signal, that the Analog-to-Digital Converter originally took to create the sample. The DAC does this as faithfully as possible and the first CD players did only that, which didn't sound good at all. Nowadays DACs use Oversampling to smooth out the audio signal. The quality of the filters in the DAC also contribute to the quality of the recreated analog audio signal. The filter is part of a multitude of stages that make up a DAC.

How does audio get digitized on your computer?

Your computer has a soundcard - it could be a separate card, like a SoundBlaster, or it could be built-in to your computer. Either way, your soundcard comes with an Analog-to-Digital Converter (ADC) for recording, and a Digital-to-Analog Converter (DAC) for playing audio. Your operating system (Windows, Mac OS X, Linux, etc.) talks to the sound card to actually handle the recording and playback, and Audacity talks to your operating system so that you can capture sounds to a file, edit them, and mix multiple tracks while playing.

Standard file formats for PCM audio

There are two main types of audio files on a computer:

- PCM stands for Pulse Code Modulation. This is just a fancy name for the technique described above, where each number in the digital audio file represents exactly one sample in the waveform. Common examples of PCM files are **WAV** files, **AIFF** files, and **Sound Designer II** files. Audacity supports WAV, AIFF, and many other PCM files.
- The other type is compressed files. Earlier formats used logarithmic encodings to squeeze more dynamic range out of fewer bits for each sample, like the u-law or a-law encoding in the **Sun AU** format. Modern compressed audio files use sophisticated psychoacoustics algorithms to represent the essential frequencies of the audio signal in far less space. Examples include **MP3**

(MPEG I, layer 3), **Ogg Vorbis**, and **WMA** (Windows Media Audio). Audacity supports MP3 and Ogg Vorbis, but not the proprietary WMA format or the MPEG4 format (AAC) used by Apple's iTunes.

For details on the audio formats Audacity can import from and export to, please check out the [Fileformats page](#) of this documentation. Please remember that MP3 does not store uncompressed PCM audio data. When you create an MP3 file, you are deliberately losing some quality in order to use less disk space.

[Next part - Rules of Audacity](#)

[Back to the main page](#)