



<http://audacity.sourceforge.net/manual-1.2/index.html>

[Introduction](#) [Tutorials](#) [Reference](#) [What's New](#) [FAQ](#) [Credits](#) **v1.2**



Introduction

1. [What's New in Audacity 1.2](#)
2. [How to use this manual](#)



Tutorials

1. [Basics - About Digital Audio](#)
2. [Editing for Beginners](#)
3. [Common Editing Tasks](#)
4. [Effects for Beginners](#)
5. [Selecting and Aligning](#)



Reference

1. [Menus](#)
2. [Toolbars](#)
3. [Tracks](#)
4. [Preferences](#)
5. [Keyboard Shortcuts](#)
6. [Creating and Editing Shortcuts](#)
7. [Effects - How they work](#)

[Frequently Asked Questions](#)

[Credits](#)

[Index of Terms](#)

Obtained from
Omarshauntedtrail.com



Introduction

Experienced users should head down to [this section](#) in this document.

Audacity ... a free cross platform audio editor.

For new users to Audacity, see [the Tutorial in five parts](#) for some basics on digital audio, setting up Audacity, selecting, aligning, and recording.

Audacity is currently in version 1.2.3. This is a stable release of Audacity; it has been tested and works well for most people. However, it comes with no guarantee, and we depend on your feedback in order to continue to improve on it. Please check our website for the latest version. Audacity is a collaborative project, so there are lots of opportunities for you to get involved if you're interested!

<http://audacity.sourceforge.net/>

What can you do with Audacity?

Audacity is a program that manipulates digital audio waveforms. In addition to recording sounds directly from within the program, it imports many sound file formats, including [WAV](#), [AIFF](#), [MP3](#), and [Ogg Vorbis](#). PCM formats of 8,16,24 and 32-bits can be imported and exported.

- [Cut](#), [Copy](#), [Paste](#), [Delete](#), [Silence](#), [Duplicate](#), [Split](#)
- apply [plug-in effects](#) to any part of a sound
- built-in [volume envelope editor](#)
- customizable [spectrogram track view mode](#)
- [frequency analysis window](#) for audio analysis applications
- custom playback [rate](#) for every track, which is live-resampled to the project sample rate

simple to complex [align operations](#) for tracks and groups of tracks

What can't Audacity do?

Although Audacity is a very powerful audio editor that works with an unlimited number tracks of virtually unlimited size, it cannot do everything. It cannot :

- Audacity cannot record more than two channels at once on many systems. Some support for multitrack recording is included in version 1.2.3, but it does not support very many systems yet.

Audacity opens MIDI files, but it is not a MIDI editor, and its MIDI features are extremely limited.

How does it work - a quick intro for experienced users

Audacity works with tracks, that contain one audio file. This file is editable and all actions are undoable. Undo's are instantaneous.

Almost anything is undoable, including importing and deleting tracks!

There are no inserts, auxes or buses yet.

All versions can use LADSPA plugins, and the Windows and Mac versions can use VST plugins using the VST Enabler (a separate download).

There is volume automation via the [envelope tool](#).

Editing functions include [Cut, Copy, Paste, Duplicate](#)(to new track at same timecode), [Split](#)(to new track at same timecode), and [Silence](#).

A grid can be activated for each track, and snapping can be switched on as well.

What makes Audacity unique?

- Audacity is free and the source code is available under the GNU General Public License.
- Audacity is cross-platform - it runs on Windows (98 through XP), Mac OS X, and many Unix platforms, including Linux. Previous versions worked with Mac OS 9.
- No limits on the number of tracks or the length of any track, except the size of your hard disk.
- Import almost anything: WAV, AIFF, Next/AU, IRCAM, MP3, and Ogg Vorbis files are supported natively, but Audacity will also open just about any uncompressed sound file and automatically deduce the format (using the [Import Raw Data...](#) feature).
- Audacity not only includes many high-quality effects built-in, but also lets you use LADSPA and VST plug-in effects. There are dozens of free, shareware, and commercial plug-ins online that do everything from Reverb to Noise Reduction.
- Audacity supports plug-ins written in the Nyquist programming language, a high-level language designed specifically for working with audio.
- Audacity acts like a non-destructive editor, providing multiple levels of undo, but it also writes changes made to the audio to disk, eliminating the need for complicated real-time processing.

- Label tracks allow you to annotate waveforms (for example, transcribing speech) and later export the waveforms to a text file.
- Powerful spectral features allow you to view waveforms as spectrograms or plot the power spectrum of any region of audio, and even export this data to a spreadsheet.
- Pitch-changing and tempo-changing effects


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.

Tutorial - I.Basics

Part 2 - Rules of Audacity - Part 2

If you'd like to get straight playing an imported file or recording something, you can skip this section and come back later.

Whenever you work with Audacity, there are some rules you should remember:

1. One clip per track

A clip is simply a piece of audio material. Imported, recorded, split or duplicated from another track, one track can only carry one piece of audio at a time. You can extend it by pasting material or inserting silence in to it, or cut a piece away, but it will always be one continuous piece of audio.

2. Audacity always records to a new track

This new track is opened at the bottom. You'll have to zoom out and then resize the track view of the bottom most track to see what is recorded. You can actually use the window sliders at the bottom and right to do this after starting to record, but this way no performance will be lost to the windowing system.

I suggest hitting CTRL+F to get an overview of the entire project as well. This only affects the horizontal zoom by the way(left-right zoom). There is no way to zoom out vertically without using the mouse yet.

3. Edit/Duplicate will not create a new audio file

This may not seem a big deal, but it is if you're editing a large live recording.

What Audacity does is reference the original audio material until you actually perform some kind of edit on it, such as cutting a piece away, or using any effect on it. One thing to remember is the **UNDO** function. You can undo/redo stuff as many times as you like, and yes, even after you have saved your project.

You may ask what happens if you do, for example, cut away a piece or mark off a 30 minute piece and split it to a new track. It only writes changed data to disk. Since Audacity works with chunk of audio data of around one megabyte in size, this happens quite fast. Rest assured that the only big waiting period might be the importing of large audio files.


To further explain how Audacity works, flip to the [next page ...](#)

Part 3 - Setup, Audio Import and Playback - Part 3

1. Create a new project

This is very important!

Audacity writes all the changed and recorded audio to a directory called **Projectname_data**, which is located right where you saved the project file itself.



Thus, select and choose a location and filename for your project.

Please note that when you startup Audacity fresh, only the " **Save As...**" menu option is available.

To save your project later on, you can also use the keyboard shortcut : CTRL+S


2. Check the Preferences

Again, this is very important!


Press CTRL+P or go to ...




...then check if the right output is selected :



...set the sample rate of your choice...
(44.1 kHz is the default)



...and here's a crucial screen :



The *File Formats* settings need discussing at this point.

When importing uncompressed audio, there are two ways to do it. "*Make a copy of the original before editing*" means, that Audacity actually copies the entire audio file that you imported in to its project data directory and in the process sets up the little overview graphics, whose descriptions get stored in the project data directory too.

The second way is to use the original imported audio. You may think we're actually editing this file, but no we aren't. In fact, Audacity will now read the imported file once and simply create the graphics overviews for them in the data directory, and

subsequently write to disk all the audio data that you change. The original file is only used for playback. All audio that remains unchanged will be played from the original file.

The advantage of choosing to **make a copy of the original** is that you avoid trouble, should **anything** in the original file change.

For example, should you accidentally delete the original file, you're lost.

You have to make up your mind before you start a project. Choose to make a copy of all imported files, and you'll use more space on your hard disk(s), but it will be easier to back up the project too, because all files that have anything to do with your project will be in the project data directory.

The *Uncompressed Export Format* can be set to WAV or AIFF for now. Please check the [fileformats page](#) for further information on export formats.

We'll ignore the *Spectrogram* settings for now. The *Directories* setting can be ignored as well for now, because all it sets is the directory to use for recordings, undo data and other stuff, if you haven't yet saved your project. Since we already saved our project, this setting is of no importance to us, though you may want to set it properly later on. Initially this is set to a folder called "audacity_temp_1.2" in the system temporary directory.

3. Import an audio file


There are three ways to do this:

1. Simply drag and drop the audio file in to the Audacity window. (If you're using Mac OS 9 or X, drag the audio file to the Audacity icon instead...)
2. Select *Import Audio ...* in the Project menu.
3. Use the keyboard shortcut : CTRL+I

Audacity can import WAV, AIFF, AU, IRCAM, MP3 and OGG files. Please refer to the [fileformats page](#) for further reference on these audio formats.

4. Playback

The imported file should now be displayed in an audio track. The track will look a little like this, depending on what you imported :



Trackpanel and Waveform Overview of the imported Track

If you're not sure where to find audio material, simply rip some off a CD, or in Windows, check the Media folder in the directory of your Windows installation.



Now click on the green Play button at the top and you should hear the file you have just imported.

[Next part - Recording with Audacity](#)

Part 4 - Recording with Audacity - Part 4

1. Create a new project

Save an empty project. Or simply use the one from the [previous part](#). Remember, that if you don't save your project before you start recording or importing, that all recordings, edit and other files will be written to the directory set on the [Directories](#) tab of the [preferences](#).

2. Check the preferences

Make sure your playback and recording device are set. If you're going to record a stereo signal, set the number of channels to record to 2 (Stereo) on the [Audio I/O](#) preferences.

When picking a device to record from, make sure you've set up all the connections properly, such as plugging a microphone in to the **Mic Input**, and any other device in to the **Line In** of your sound card. Then check that the gain level knob (the amount by how much the input should be amplified) of the mixer of your soundcard is set right.

Since most soundcards can mix the inputs back in to the outputs, the easiest way to test your microphone is to speak in to it while playing with your sound card mixer. The sound card mixer is a piece of software either provided by the sound card maker, or by the operating system you're using. The Windows mixer is pretty straight forward, though some soundcards bring their own along. The Mac's mixer is controlled via the Sound

Control Panel, and the Linux users have a variety of mixer applications at their disposal. Just make sure they work before yelling at your screen that nothing works.

3. Hit Record

Click on the red **Record** button  to begin recording.

Click on the blue **Pause** button  to pause the recording. Press it again to continue.

Click on the yellow **Stop** button  to cease recording. The cursor will return to its previous position, before the recording was started.

That's it. You can now play around with your recording and explore the editing capabilities of Audacity. Remember that you can use the Undo function almost without limits whilst the project is open.

Proceed to the [Editing for Beginners](#) section, for a detailed overview of editing functions.

II.Editing for Beginners

Part 1 - Introduction - Part 1

Sound Editing in the "real world"

Sound editors clean up dialogue tracks, cut layers of special effects, place sounds at certain times, create ambiance tracks by cutting out unwanted stuff and mixing in interesting or necessary sounds. Music production engineers may cut pieces of vocals away or shift them to a another spot in a song.

Editing is about cutting, placing, fading, cross-fading, shifting, duplicating and adjusting the volume (also referred to as level) of audio material. Mixing is a form of editing too of course.

Here is an example of what is done in sound editing during the production of a television show or film. In the next part we will run you through a few of those techniques in Audacity.

The Path of Sound in Film and TV Postproduction

Film and TV crews have at least two people present that take care of recording sound during principal photography of a show. Principal photography is usually shooting the scenes with actual live actors or real backgrounds by the way.

Sound in Principal Photography

The first person is the boom pole operator. The boom pole is an extendible stick with a microphone attached to it. This is used to capture dialogue either during filming or not. When not filming, it might be capturing off-scene dialogue or retakes of lines that the actors flunked during actual film takes. The more expensive the show is and the more time there is to do the work, the more people will resort to *looping* those takes, which is recording those lines in a sound studio environment instead of a film studio or location.

The second person is the sound mixer, who usually sits in a place farther off from the shooting and records the sound captured by the boom pole operator, either via cable or wireless devices to tape, optical disks or hard drive.

This is the raw sound material of a show. It is called **production sound** and the only desirable parts it usually contains are dialogue and body sounds. In post production, depending on the complexity, budget size and time, almost everything you hear except for the dialogue and some body sounds, are added later during ...

Post Production

This is where most of the stuff we'll be describing for Audacity will happen. You've got the recording. Now what ?

After the visual part of the show is cut, the first of which usually isn't the final one, it is handed to the sound editor. In TV shows, you'll usually have one or two people for this, for major film productions a whole bunch more, for which tasks will be subdivided on a finer level.

Raw sound - Cleanup time

No shows do without film edits and many have plenty of them.

Scenes may be shot with one or more cameras and mics. Actors might have flunked their lines and picking up shooting prior to the mistake might be chosen or the entire sequence reshot. The film editor may have chosen parts from different takes for the cut of the scene. The action might be moving along at the wrong pace and the film editor shortened or stretched parts of a scene.

The sound editor makes sure transitions between cuts are smooth. He or she removes undesired sounds, such as breathers of the same person that overlap from one film edit to the next.

Material is cut away that contains unwanted sounds, such as creaking chair legs and sharp impacts of objects on tables and floors. Some of these may require looping of dialogue in the studio, because the noise may have been intolerable. Also, material may sometimes be denoised. The most sophisticated methods remove the whirring of the camera motors from takes. It is used as sparingly as possible though.

It's always desirable to get the best possible sound from the start, which is the recording stage.

Adding stuff - the really big deal

After this cleanup is complete, sounds are added.

The first is **ambiance**. Just close your eyes and listen to the sound around you. That's ambience. Sophistication of ambiences rises with budgets. From premixed to over a dozen tracks, you'll find it all in TV shows and feature films. In any indoor scene with a lot of people in the background, nobody except for the actors being filmed will actually talk. That flurry of conversation is added later on.

Next comes **foley**. These are clothes rustling (body sounds), foot steps and objects being handled. People that have the ultimate edition of Terminator 2 will know that all of Arnold Schwarzenegger's footsteps and rustle of his leather clothing were created by a five foot woman. These people are usually called foley walkers or foley artists.

The foley editor then cleans these sounds, chooses the most fitting takes and makes sure they all sync to the picture properly. The foley mixer then does his/her thing.

Next come **effects**. Foley are effects too, but they are a special category and can best be described as live created studio effects. Effects are usually more heavily edited and recorded from all kinds of places. A lot of effects are created by layering sounds on top of each other, changing their pitch and loudness, editing bits out and adding others.

Many effects you'll hear are phone and door bells ringing, mobile phone beeps, doors of houses and cars opening and closing, weapon shots, slaps, car skids, machines of any kind, space ships flying around, explosions, to name a few.

For example, a friend of mine and I created the sound of a small wooden rowing boat hitting a larger wooden sailing ship and scraping along its side by pitching down a knock on a large wooden door for the impact of the rowing boat and ship, and pitching down the sound of a skateboard rolling and scraping along a halfpipe.

The techniques required to properly handle sounds like these are used in all kinds of productions. Audio books, music production, sound effects creation ... you name it.

For more on this subject, read the [Audio Post FAQ](#) at www.filmsound.org.

So let's jump in to the fray and look at how you can handle your sounds in Audacity.

II.Editing for Beginners

Part 2 - Cut, Copy and Paste - Part 2

From here on you may encounter funny letter combinations in boxes like this.

These are keyboard shortcuts to the functions presented to you in the text. These can be either single keys (e.g. SPACE) or combinations that need to be held down at the same time (e.g.CTRL+C). You can usually create your own. Check out the [this page](#) for more details.


The most basic editing step is cut and paste. It's what people did with tape and it's easy with data in computers, so take a look at these basic operations, referred to as **Cut, Copy and Paste**. The [next page](#) will handle **Silence, Duplicate and Split**. You may also want to check out the [reference](#) section, so you'll know where to find all the tools and how to resize tracks for example.

It is assumed that you have a project open and that at least one track of audio material is present.

Let's take a look at this example of an Audacity window:

The View

The Audacity Window



As you can see by the graphics above, the [time shift tool](#)  is selected. It is used to move the entire audio clip around inside its track.

The cursor (little blinking line across a track and the timeline) will remain at its position, so effectively you'll be sliding your audio material underneath the cursor.

Let's say we want to cut out that bit in the middle then. First we've got to **select it**.


Making a selection

To select the part you wish to cut, copy or paste to, use the [selection tool](#) . If it's not activated, do so now by clicking on it in the toolbar.

Now press and hold the left mouse button while you drag the mouse to mark an area.

This area is darker than the surrounding area of the clip. Note, that even though you can mark an area larger than or extending beyond the actual audio clip in the track, the operations will only work on the actual clip. Playback however will work outside the clip.

Press the space bar to listen to the audio in the marked area.




To **extend** or **contract** your selection, hold down the SHIFT button and click on the area you wish your selection to extend or contract to.

If you click at a spot that is on the right hand side from the middle of the current selection, you will set the right hand boundary of your new selection.

Cutting the selection

Cut the selection by selecting "[Cut](#)" from the [Edit menu](#) ... or press CTRL+X.



To undo this operation, select Undo in the Edit menu or press CTRL+Z

Copy will copy the selection to the clipboard.

You can then **paste** that data back in to any track by clicking where you want this audio to be inserted and select Paste in the Edit menu,

or press CTRL+V.

Thus pasting is the opposite of cutting. You can also copy material, make another selection with the mouse and then paste. This will replace the selected material with the contents of the clipboard, no matter how short or long either of them are.

During all operations of this kind, the bottom row of the screen will display two things, namely the start time and the end time of your selection. The display to the left of that called "Project rate:" and its value, defaulting to 44100, can be changed by clicking on that number and selecting another from the drop-down menu. This sets the sample rate of everything you produce in audacity.

All files, no matter which will be played at the project rate, and exported at that rate. Should the sample rate of a track be different from the **Project Rate**, these tracks will be resampled to the **Project Rate** as the project is played back or exported.

Audacity will not change the sample rate of any imported audio. If you want to change the rate of an imported track for any reason this can be done using the Rate option on the track pop-down menu.

II.Editing for Beginners

Part 3 - Silence, Duplicate and Split - Part 3

Silencing unwanted sources

This operation flattens the selection. It essentially is a cut operation without deleting the selection completely. After all, if you cut a second away, nothing remains. Using the Silence operation will still leave you with a flatlined area.

When silencing parts between vocal lines, please keep in mind that a sudden drop in background ambiance can have a bad effect, so at the very least fade the area around the silenced part, to minimize that effect. Rules to start with are, **fade in quickly** and **fade out slowly**.


Alternately, use the [envelope tool](#) to lower the volume in that area. That way, you can comfortably change it later.

Keyboard Shortcut : CTRL+L

Duplicate

The selected area gets **copied**, a new track is created and the copied material is pasted in to that new track at the same point in the timeline.

To illustrate, here's the image from the [menu reference](#):



The benefits of a duplicate are many. One of these is experimentation with effects.

Some of you may say "*I can do that with the original track too*". But you can't change the volume of your effect and original audio separately. If you put some Reverb on to your audio, you can only lower that processed audio in volume later on. If you duplicate the audio first and use the reverb on that (with 100% reverb and 0% original signal), you can freely change the volume for both the original and reverb signal.

Also, you can do weird and wonderful things to your duplicates to create special effects. You'll have two pieces of the same audio to work with. Silence parts, reverb another,

phase a third, filter another and see how that sounds. It is so easy to duplicate a piece of audio and do weird things to it, so try it. Combining sounds produces magic.


A special note on performance :

The new piece of audio isn't actually copied on the hard disk. Audacity will still play from the original audio file(s) until you change a piece of it.

Keyboard Shortcut : CTRL+D

Split

This performs the same as Duplicate, but it also silences the selected material, after copying it to a new track. Again, here's the illustration from the [menu reference](#):



There are plenty of good uses for this function, but I'm not going to tell you about them here. You'll have to go to the [next part](#) for the meat of this tutorial.

Keyboard Shortcut : CTRL+Y

II.Editing for Beginners


Part 4 - Splitting and Submixes - Part 4

Moving bits of an Audio track

In all projects you'll be pushing your audio around at some point. Otherwise, what are you doing here ?


There are techniques, easily achievable with Audacity, to cover almost any kind shifting you'd want to do. In our example, we have a small sentence of speech, where the speaker made a pause after the first word. We'd like to eliminate that pause.

The part after the pause is selected




Select

Then the [split function](#) is used to pop the selected audio to a new track



Split

[The Time Shift Tool](#)  is selected and the audio on the lower track is moved left.



Move Audio and select for fading

Now, it's a good idea to listen to the two tracks individually for breathing sounds for example.

Use the solo button of the tracks for this. Then listen to them both in the mix. Again, you can use the solo buttons for this.

If you have a lot of other tracks playing at the same time, press the solo buttons on both tracks. There should be no over lapping or cut-off breathing sounds.

When you're satisfied, fade out the last two thirds of the overlapping upper part of the track, and fade in the first two thirds of the lower overlapping audio.

Two thirds, and not the whole overlapping audio, are chosen to keep the level of audio constant. If the whole overlapping parts were faded, you would get a level drop of 3dB in the middle of those fades.

You can check this out by taking a piece of music, duplicating it, and then fading the tracks, one fading out, the second fading in. In the middle of those fades, the level of the mix will drop audibly. Do a fade over last two thirds for the fade out and first two thirds for the fade in, and you probably won't notice any change in level.


Two thirds is a guideline, but not the law, so you may have to experiment a little.

Mixing it back together again

!!!Remember!!!

The final mix is done with the [Export as WAV](#) function in the [File Menu](#). Here we'll be looking at creating submixes with the [Quick Mix](#) function.

You've done a lot of edits and now have dozens of little tracks with little bits and pieces here and there. It might look like this:




**Bits and Pieces spread all over the screen
First four tracks selected for quick mixing**

We can use the [Quick Mix](#) function in the [Project menu](#) to bring down the number of tracks. However, you don't need to mix everything in to one new track.

Select the tracks you want to mix together by SHIFT+click 'ing on the track panels. In the graphics above, the first four tracks are selected.

Then select [Quick Mix](#). In this example I have quick mixed everything down to two tracks :



Bits and Pieces, quick mixed down to two tracks

And thus, two **submixes** were created. Remember though, that we did this for convenience of not having to organize a large number of tracks.

If you still want to shift bits around later on, you should make sure that the parts being mixed to a track **do not overlap**, so you can split it away and edit it again later.

Tutorial - III.Common Editing Tasks

Part 1 - Quickies - Part 1

This page will give you an idea of some of the possible things you can do with Audacity and walk you through the process.

- [1. Splitting an MP3 into two separate files](#)
- [2. Mixing background music with a voiceover](#)
- [3. Recording harmonies with yourself](#)
- [4. Recording two sound sources on separate tracks](#)

Splitting an MP3 into two separate files

(Before you try to export MP3 files, read the section on [Exporting MP3 Files](#) for some important information on steps you need to do first.)

- Open the MP3 file.
- Select the part of it that you want to be the first file. Listen to it by clicking the Play button.
- While this part is selected, choose Export Selection as MP3... from the File menu.
- Now select the part you want to be the other song and Export again.

Mixing background music with a voiceover

Audacity makes it very easy to mix two different sounds together.

- Open one sound (for example, the background music).
- Select Import Audio... from the Project menu and open the other sound (for example, the voiceover).
- Listen to your sound using the Play button. **Audacity** automatically mixes them together.
- Choose the [Time Shift](#) tool and adjust the position of one track or the other until they're synchronized the way you want them. You can even move tracks around while they're playing.
- If you hear clipping which wasn't present in either of the original files, it means that the combined volume of the two tracks is too loud. Use the gain controls on the tracks to reduce the volumes until you don't hear clipping anymore.

Export as a WAV or MP3 file.

Recording harmonies with yourself

- Open the Preferences, click on the Audio I/O tab, and check the box marked "Play other tracks while recording new one".
- Click the Record button. Record yourself singing. Click Stop.
- Click the Record button again. The first track you recorded will play, but **Audacity** will also record a new track at the same time, allowing you to sing harmony with yourself.
- When you play the two tracks you recorded together, they probably won't be synchronized. This is normal and is not the fault of **Audacity**. To fix it, you will need to grab the [Time Shift](#) tool and slide one of the tracks around until it sounds right.

Recording two sound sources on separate tracks

- Open the Preferences, click on the Audio I/O tab, and make sure "Record in Stereo" is checked.
- Connect one sound source to the left channel of your soundcard and the other on the right. If you don't have an external mixer with a pan control, use a stereo/mono splitter cable.
- Record your stereo track.
- On the button with the track name is a drop-down menu. Select "Split Stereo Track" from this menu.
- Use the drop down menu again to change each track to mono.
- To position sounds in the stereo image use the pan control on each track.
- To export your recording in stereo you need to include at least one Stereo track. If necessary add a blank one using Project > New Stereo Track.

Repeat as necessary for further tracks (see [Recording harmonies with yourself](#) above). If mastering from tape, a click recorded across all tracks will help you synchronize them in **Audacity**.

IV. Effect for Beginners Part 1 - A Virtual Room with Reverbs - Part 1

An example by the numbers

As an example I'd like to show you how to create a reverb for small club with a band.

Let's see what we have to do with parameters in the reverb fx in sample editors. First we determine some basics about the room we want to simulate.

Our room (the club) is about eight meters on all sides (for simplicity) and three meters high. The band will be playing no more than four meters away from us. The room volume will be about 200 cubic meters.

This means that the DIRECT SOUND will get to you within about 11 ms (milliseconds). This is the amount of time we'll need to take off the other values.

The first EARLY REFLECTIONS (off the left and right walls) travel around nine meters before reaching you, so they'll get to your ears about 15 ms later! (9 meters at 343 meters/sec -> 26 ms minus the 11 ms the DIRECT SOUND takes to travel to your ears).

I'm not taking the floor into account because of the obstacles, such as your own legs and a table, that prevent reflections of sound from the band from reaching you. It is also probable that your sitting close to a wall and EARLY REFLECTION time may drop considerably. Experiment and listen very carefully.

The reason I'm deducting 11 ms from the 26 ms is that in a sample editor we already have the DIRECT SOUND. In Cool Edit an EARLY REFLECTIONS parameter doesn't exist. In Soundforge however there are a few ER times to choose from. We choose whatever is closest to 15 ms. Freeverb unfortunately does not have an Early Reflections feature.

Please note that the timing of the EARLY REFLECTIONS are defined by the rooms characteristics such as shape and size. The sound color of the EARLY REFLECTIONS is defined by the surfaces in the room. These are parameters that are available to professional reverbs only (hard- or software), so don't worry about not finding that stuff in great detail in either Soundforge's, CoolEditPro's or plugins of other editors.

The PRE-DELAY and REVERB TIME depend on more than just room size. Wall and floor material, as well as anything else that reflects, disperses and dampens sound will influence the reverb in quite a few ways.

Our club has lots of wood on it's walls and floor and people are sitting at tables and the air's probably filled with a little smoke (except in California :).

The REVERB TIME, a parameter found in almost all plugins, should be set to no more than one second. In small, crowded rooms where little REVERB occurs it is quite important to set the EARLY REFLECTIONS right. Experiment with the reverb time keep it lower than 1.2 seconds and above 0.7 seconds. Freeverb puts all this information in to the SIZE parameter.

The PRE-DELAY can be set to 30 ms or more. Larger rooms will develop REVERB later since the reflecting sound will travel a while before the reflections become hard to identify but it also depends on the reflecting materials. Freeverb once again doesn't know any Predelays, but since you're reading this to learn how things work, it's still worth spelling out.

Both Cool Edit and Soundforge have parameters referring to the DENSITY of the REVERB part of the sound. Materials such as polished marble will reflect the sound very well and disperse the sound very little. Wood disperses the sound a lot and a carpet practically swallows a large chunk of it. For our Club we need to set the density high since our material disperses the sound a lot (humans, wood and obstacles). No parameter like that in Freeverb and yet it sounds great. It's easy to use as well, so I don't really mind when using it.

Higher frequencies are swallowed a lot faster than low frequencies. Many reverb algorithms refer to this in some form or another. Cool Edit has the 'High Frequency Absorption Time' as a parameter for this. Sound Forge gives you a choice, whether to do this at all for both low and high frequencies and even lets you set them but doesn't provide any time parameters. Professional Reverb hardware such as Lexicon PCM 70, 80 & 90 let you control frequency and time of this, for low and high frequencies. Freeverb has the DAMPEN parameter for this.

This can effectively add to room's liveliness. A room full of people will absorb high frequencies a lot faster than an empty room, in which only the materials would make the difference, so in Cool Edit we'll set the 'High Frequency Absorption Time' very low, since we're in a crowded club. For Soundforge activate the 'Attenuate high frequencies above' box and set the frequency parameter to around 2000-2500 Hz. Experiment a bit. The higher this frequency and the longer the Absorption Time parameter in Cool Edit, the emptier and colder your room will seem. In small rooms this can cause the room to sound very metallic. Sewer pipes and oil pipelines e.g.

The Reverb EQs are the final processing stage and can change the sound dramatically. These Highshelf and Lowshelf EQs are found in almost ANY Reverb plugin and hardware unit. Freeverb has both in form of Hipass and Lowpass filters. These are not EQs. The parameters change the target frequencies of the filters. Again, listen.

The last thing you have to do is set the balance between the WET signal and your original DRY signal. Here you can do as you please but remember that only very smoothed hard surfaces in big cavernous space create very loud reverbs. A church might do that if there aren't a big number of people and wooden benches in there, but a club doesn't :). Remember that in Audacity it is smarter to DUPLICATE the part you wish to add reverb to, and run the effect over that duplicate. The volume of that reverbed material is easier to control, so set the WET parameter in Freeverb to 0dB (- &rt;100%) and the DRY parameter to -endless (->0% left most slider position).

So when you create your next 'soon-to-be-the-hit' song or want to put things in perspective, be sure to check the possibilities of adding dimension and scope to your material through clever use of reverbs. Put your snare to the front, the piano a little to the back, and the sax just a little in front of the drums. Experiment. With these basics you'll be able to get better results in less time.

Btw, Snare always gets reverb. How much and what kind up to you, but it's generally good thing. It's also important to decide whether you're trying to simulate a real live band or doing a multitrack sound, which boasts all kinds of nice ideas, like doubling(another take of the same thing on top of the first take), extreme panning and in general NOT reality :).

Dry versus Wet

Have you heard of those silent rooms, where they test speakers and microphones ? They're great if you don't happen to be inside of one. After a few minutes you hear the blood in your ears roar like an ocean. You hear only your inner voice when you speak. The room swallows almost everything else. Most people can't stand being in a room like this for more than 10-15 m

V.Selecting and aligning Part 1 - Selecting - Part 1

Basic rules

Selecting can be done with the keyboard, the mouse or by combining the keyboard and mouse.

You need to have the [Selection Tool](#)  activated for making a selection.

The **exception** is selecting the entire track by clicking on the track panel. The [align](#) functions all still work the same way no matter which tool is activated. However, making selections and placing the cursor still requires the [Selection Tool](#).

The selection modifier key is SHIFT.

In a word processor or text editor, the SHIFT key is the modifier to turn any character that you enter via the keyboard in to an UPPERCASE character. In Audacity and most audio editing applications, holding down the SHIFT key means that any cursor movement or placement will result in a selection.

What follows are ways of making selections.

How do I ...

Select audio on one track ?

Place the cursor on the waveform of a track, hold down the mouse button, drag the mouse and release the mouse button.

You can select audio with the keyboard as well, but may need to place the cursor with the mouse first. You need to hold down the SHIFT key and use the left/right cursor keys to expand the selection. Combine the former with the CTRL key (STRG on the German keyboard) to shrink the selection.

Expand/shrink my selection ?

Hold the SHIFT key, then click and hold the mouse button. If you clicked to the left of the center of the current selection, you will now be able to move the left selection boundary with your mouse. The same goes for the right hand side of the selection center. Release the mouse button when you are satisfied. You can use the cursor keys as well. It works the same way as described in the previous item.

Select from the cursor to the start of the track ?

Use the [Start to Cursor](#) function in the [Edit](#) menu.

Select from the cursor to the end of the track ?

Use the [Cursor to End](#) function in the [Edit](#) menu.

Select from the cursor to the start of the project ?

Hold the SHIFT key and click on the [Skip to Start](#) button.

Select from the cursor to the end of the project ?

Hold the SHIFT key and click on the [Skip to End](#) button.

Select all the audio of a track ?

Click on the track panel. The resulting selection will be confined to the contents of that track.

Select all the audio of two or more tracks ?

It cannot be done precisely right now, if one track isn't within the range of the other. It actually doesn't have to be, because of the group functions in the Align menu, but more of that later.

If the start and end of one track is within the start and end of the other, simply select the longer track, then SHIFT + click the second track panel.

If the start and end of one track is not within the start and end of the other, select the one track, then expand the selection by SHIFT + clicking in to the second track until the selection encompasses all the desired material.

Select all the audio of the project ?

Use the *Select All* function in the Edit menu, or hit CTRL+A.



Reference

[Menus](#)

[Toolbars](#)

[Tracks](#)

[Keyboard Shortcuts](#)

[Create or Edit Shortcuts](#)

[Effects](#)

[Preferences](#)

[Sound File Formats](#)

[Exporting MP3/OGG Files](#)

[Back to the Reference page](#)

File Menu

New...

Creates a new empty project window.

Open...

Selecting *Open* presents you with a dialog where you can choose a file to open. If the current project window is empty, the new file will take over the current window, otherwise a new project window will be opened.

The file formats recognized by Audacity are

[WAV](#), [AIFF](#), [NeXT/AU](#), [IRCAM](#), [MP3](#), [Ogg Vorbis](#), [MIDI](#), and the Audacity Project file format ([AUP](#)).

In addition, Audacity can also attempt to open an audio file in a different format, as long as it is uncompressed - to do this, use the [Import Raw Data](#) command.

Close

Closes the current project window.

Save Project

Saves the current Audacity project ([AUP](#)) file. Audacity projects are *not* intended to be read by other programs, but they are extremely fast to load and save within Audacity. When you are finished working on a project and you want to be able to use it in another program, select one of the [Export](#) commands instead.

Note that most of the audio data for an Audacity project is not stored in the [AUP](#) file, but in a directory (folder) with the same name as the project. For example, if you save a project as `chanson.aup`, there will be a directory called `chanson_data` created to store the actual audio tracks of the project. For more information on the project file format, see [the file formats page](#).

Save Project As...

Saves the current Audacity project ([AUP](#)) file, allowing you to give it a different name or move it to a new location if you have already saved it in one location.

Audacity projects are *not* intended to be read by other programs, but they are extremely fast to load and save within Audacity. When you are finished working on a project and you want to be able to use it in another program, select one of the [Export](#) commands instead.

For more information on the project file format, see [the file formats page](#).

Export As WAV...

Exports the current Audacity project as a standard audio file format such as [WAV](#) or [AIFF](#).

You can change the format of exported files in the [Preferences](#) dialog.

If there are multiple tracks in your project, they will be automatically mixed in the exported data. For more information about mixing, see [Quick Mix](#).

To export only a single track or part of a track, use [Export Selection](#).

Export Selection As WAV...

This is the same as [Export](#), but it only exports the part of the project that is selected. This is very useful if you want to save a small clip from part of a track as a separate file.

Export As MP3...

Exports the current Audacity project as an MP3 file. Audacity does not encode MP3 files directly, but instead requires that you download a separate MP3 encoder. See [Exporting MP3 Files](#) for details.

To export only a single track or part of track, use [Export Selection as MP3](#).

Export Selection As MP3...

This is the same as [Export MP3](#), but it only exports the part of the project that is selected. This is very useful if you want to save a small clip from part of a track as a separate file. See [Exporting MP3 Files](#).

Export As OGG...

Exports the current Audacity project as an Ogg Vorbis file.

To export only a single track or part of track, use [Export Selection as Ogg](#).

Export Selection As OGG...

This is the same as [Export As OGG](#), but it only exports the part of the project that is selected. This is very useful if you want to save a small clip from part of a track as a separate file.

Export Labels...

If you have any [Label Tracks](#), this command will export them as a text file. This feature is commonly used in Speech Recognition research to annotate a speech utterance and export the annotation to be later processed by another program.

To import these labels into a different project later, use the [Import Labels](#) command.

Export Multiple...

This allows you to do multiple exports from audacity with one command based either on multiple tracks in the project or labels in a single audio track. It's great for splitting up long recordings into CD tracks

Preferences

Opens a dialog window that lets you configure Audacity. For more information on the different preferences, see the section on [Audacity Preferences](#).

Exit/Quit

Closes all project windows and exits Audacity. It will ask you if you want to save changes to your project. It is not necessary for you to save changes if you just exported your file as WAV or MP3 and you are finished working with it. On the other hand, if you are working on a mix and plan to continue where you left off, saving an Audacity Project will let you restore everything as it is later.

Edit Menu

Undo

This will undo the last editing operation you performed to your project. Audacity currently supports full unlimited undo - meaning you can undo every editing operation back to when you opened the window.

Redo

This will redo any editing operations that were just undone. After you perform a new editing operation, you can no longer redo the operations that were undone.

Cut

Removes the selected audio data and places it on the clipboard. Only one "thing" can be on the clipboard at a time, but it may contain multiple tracks.

Copy

Copies the selected audio data to the clipboard without removing it from the project.

Paste

Inserts whatever is on the clipboard at the position of the selection cursor in the project, replacing whatever audio data is currently selected, if any.

Trim

Deletes everything but the selection.

Delete


Removes the audio data that is currently selected without copying it to the clipboard.

Silence

Erases the audio data currently selected, replacing it with silence instead of removing it completely.


Split

Moves the selected region into its own track or tracks, replacing the affected portion of the original track with silence. See the figure below:



Duplicate

Makes a copy of all or part of a track or set of tracks into new tracks. See the figure below:



Select... > All

Selects all of the audio in all of the tracks.

Select... > Start to Cursor

Selects all the audio from the start of the timeline to the current cursor.

Select... > Cursor to End

Selects all the audio from the cursor to the end of project, which is as far as audio exists in the project.

Find Zero Crossings

Moves the cursor or the edges of the selection to the nearest point where the audio waveform passes through zero. Use this to help achieve click-free cuts in audio, as the two ends of the cut will exactly line up.

Selection Save

Saves the current selection and position. This function only saves the selection boundaries, but not the actual audio data.

Selection Restore

Restores the selection to the project. You need to have an existing selection for this to work.

Move Cursor...

These commands provide quick and accurate ways to manoeuvre the cursor around the project to the start and end of tracks and selections.

Snap-To...

Turns snapping of the cursor to a grid of time values on or off. When off, your cursor can be positioned on any sample in the audio file. When on, it is restricted to certain points in time, determined by the time format in use. This is useful if your project must follow an external time code exactly.

View Menu

Zoom In

Zooms in on the horizontal axis of the audio, displaying more detail about less time. You can also use the zoom tool to zoom in on a particular part of the window.

Zoom Normal

Zooms to the default view, which displays about one inch per second.

Zoom Out

Zooms out, displaying less detail about more time.

Fit in Window

Zooms out until the entire project just fits in the window.

Fit Vertically

Adjusts the height of all the tracks until they fit in the project window.

Zoom to Selection

Zooms in until the selected audio fills the width of the screen to show the selection in more detail.

Set Selection Format

Sets the format in which selections are measured in at the bottom of the application window. Combined with setting *Set Snap-To Mode* to ON, this is ideal to cut stuff up in to exact measures.

Set Snap-To Mode

History

Brings up the history window. It shows all the actions you have performed during the current session, including importing. The right-hand column shows the amount of HD space your operations used. You can jump back and forth between editing steps quite easily by simply clicking on the entries in the window.

The history window can be kept open at all times. It doesn't interfere with any other operations.

Plot Spectrum

To use this feature, first select a region of audio from a single track, then select "Plot Spectrum".

It opens up a window which displays the Power Spectrum of the audio over that region, calculated using the Fast Fourier Transform. The graph represents how much energy is in each frequency.

This window can also display other common functions that are calculated using the Fast Fourier Transform, including three versions of the Autocorrelation function.

The Enhanced Autocorrelation function is very good at identifying the pitch of a note.

Float or Dock Control Toolbar

Toggles between displaying the tool docked at the top of each project window, or in a separate floating window. You can also float any toolbar by grabbing the handle on the left side and dragging it to the location you want.

Once they are floating these options change to docking the toolbars back into the main window.

Float or Dock Menu Toolbar

Float or Dock Mixer Toolbar

Float or Dock Meter Toolbar

Do the same thing for the other toolbars

Project Menu

Import Audio...

This command is used to import audio from a standard audio format into your project. Use this command if you already have a couple of tracks, and you want to add another track to the same project, maybe to mix them together.

Use this command to import anything except MIDI files. You cannot yet use this option to import Audacity Projects.

Import Labels...

This command takes a text file which contains time codes and labels, and turns them into a [Label Track](#).

Import MIDI...

This menu command imports MIDI files and puts them into a [Note Track](#). Audacity can display MIDI files, but cannot play, edit, or save them yet.

Import Raw Data...

This menu command allows you to open a file in virtually any format, as long as it is not compressed. When you select the file, Audacity will look through it and try to guess its format. It will guess correctly about 90% of the time, so you can try just pressing "OK" and listening to it. If it is not correct, however, you can use the dialog presented to try all of the different possibilities.

At the beginning of your imported track(s), you may notice a little bit of noise. This is probably the file's header, which Audacity was not able to parse. Just zoom in and select the noise with the [Selection Tool](#), and then choose [Delete](#).

Edit ID3 Tags...

Opens a dialog allowing you to edit the ID3 tags associated with a project, for MP3 exporting.

ID3 Tags are used to store the Title, Artist, Author, and other information in MP3 files. When you import an MP3 file, Audacity remembers this information and stores it with your project. When you export an MP3 file, Audacity will prompt you for the ID3 tags, showing you any tags it collected from an original MP3 file, or that you typed by selecting "Edit ID3 Tags" from the Project menu.

Audacity allows you to store ID3 tags in either the ID3v1 or the ID3v2 format. In general, you should use the ID3v2 format, because ID3v2 tags go at the *beginning* of an MP3 file, whereas ID3v1 tags go at the end. Having the title and author first means that someone downloading an MP3 you created will know that information before they've finished downloading the song.

Quick Mix

This command mixes all of the selected tracks down to one or two tracks. The channel of a track being mixed affects whether it will be mixed into the left channel of the resulting track(s), the right channel, or both (mono).

For example, if you have four tracks:

- Track 1: left channel
- Track 2: left channel
- Track 3: right channel
- Track 4: mono channel

and you select them all and perform a Quick Mix, you will end up with two tracks: the first will contain a mix of tracks 1, 2, and 4 (the new left channel) and the other will contain a mix of tracks 3 and 4 (the new right channel).

Your tracks are implicitly mixed whenever you hit the [Play](#) button and whenever you select [Export](#).

Note that if you try to mix two very loud tracks together, you may get clipping (it will sound like pops, clicks, and noise). To avoid this, you should use the track gain controls to reduce the amplitude of all of your tracks.

New Audio Track

This creates a new empty audio track. This command is rarely needed, since importing, recording, and mixing automatically create new tracks as needed. But you can use this to cut or copy data from an existing track and paste it into a blank track.

New Stereo Track

Creates a stereo version of the new audio track above. You may need this if you start with a load of mono tracks and want to export a stereo mix.

New Label Track

This creates a new Label track, which can be very useful for textual annotation. See the discussion on [Label tracks](#) for more information on how to use label tracks.

New Time Track

Creates a special track that can be used to speed up and slow down playback over the course of the project. This affects all tracks in the project.

Remove Track(s)

This command removes the selected track or tracks from the project. Even if only part of a track is selected, the entire track is removed. You can also delete a track by clicking the X in its upper-left corner. To cut out only part of the audio in a track, use [Delete](#) or [Silence](#)

Align functions

All the align functions work on whole tracks or groups of tracks, NOT on selections, even if they span across multiple tracks. Please read the [tutorial on how to use the align functions](#). You will also find some creative uses and solutions to common problems there.

These functions operate on each track individually. If more than one track is selected, these functions do not preserve their position to one another.

Align with Zero

Aligns the start of any selected tracks to the start of the project.

Align with Cursor

Aligns the start of any selected tracks to the cursor position.

Align with Selection Start

Aligns the start of any selected tracks to the start of the current selection

Align with Selection End

Aligns the start of any selected tracks to the end of the current selection

Align End with Cursor

Aligns the end of any selected tracks to the cursor position.

Align End with Selection Start

Aligns the end of any selected tracks to the start of the current selection

Align End with Selection End

Aligns the end of any selected tracks to the end of the current selection

Align Tracks Together

If you select multiple tracks, this command adjusts their time offset so that they start at the same time. It adjusts their starting time to the average of all of the tracks' original starting times. To adjust the time offset of one track, use the [Time Shift](#) tool.

Align and Move Cursor

These commands work just like the align commands with the [move cursor](#) command immediately following it.

Add Label at Selection

This menu item lets you create a new label at the current selection. You can title the label by typing with the keyboard and then hitting "Enter" when you're done. The label saves your current selection - so when you click on the label later it will return the selection to the state it was in when you created the label. See the discussion on [Label Tracks](#) for more information on how to use label tracks.

Add Label at Selection

Like [Add Label at Selection](#) but the label is added at the current position during playback.

Generate Menu

Silence

Inserts silence. The length is determined by the length of your selection and the

position by the **left boundary of your selection**.

If no selection is made, the default length inserted at the cursor position is **30 seconds**.

For example, to insert a two seconds of silence at a certain point, place the cursor at your chosen spot, hold the SHIFT key down, click and drag your mouse to the right to mark roughly two seconds, release the mouse button and activate the Silence function in the Generate menu. Silence, the length of your selection, will be inserted at the **left boundary of your selection**.

White Noise

Inserts white noise. The length is determined by the length of your selection and the position by the **left boundary of your selection**.

If no selection is made, the default length inserted at the cursor position is **30 seconds**.

Tone

Inserts a wave of chosen type, frequency and amplitude. The length is determined by the length of your selection and the position by the **left boundary of your selection**.

If no selection is made, the default length inserted at the cursor position is **30 seconds**.

Parameters:

Waveform Sine, Square, Sawtooth

Frequency(Hz) 1-20000 (value below 1.0 and above 20000 will be corrected to 1.0 and 20000!)

Amplitude 0-1 - This is the volume of the generated wave.

Pluck

Inserts a synthesized pluck tone. The length is determined by the length of your selection and the position by the **left boundary of your selection**.

If no selection is made, the default length inserted at the cursor position is **one second**.

Pitch 0-127

The number corresponds to the value of a midi note, which is assigned to a pitch. The higher the number, the higher the note.

Effects Menu

Audacity includes many built-in effects and also lets you use a wide range of plug-in effects. You can download many free plugins for Audacity at [our Website](#). Loads of information is also available at <http://www.kvr-vst.com>.

To apply an effect, simply select part or all of the tracks you want to modify, and select the effect from the menu. Titles which end in an ellipsis (...) will bring up a dialog asking you for more parameters.

There are four classes of effects and they are all [applied](#) the same way. These are

[Internal Effects](#), [VST Effect Plugins](#)(Windows/Mac only), [Nyquist Effects](#) and [LADSPA Effects](#)

Internal Effects

Amplify...

This effect increases or decreases the volume of a track or set of tracks. When you open the dialog, **Audacity** automatically calculates the maximum amount you could amplify the selected audio without causing clipping (from being too loud).

Bass Boost...

This is a safe, smooth filter which can amplify the lower frequencies while leaving most of the other frequencies alone. It is most effective if you don't try to boost too much; 12 dB is usually just right.

Echo...

A simple delay line.

This effect repeats the audio you have selected again and again, softer each time. There is a fixed time delay between each repeat.

First select the audio you want to apply the effect to. You may want to first add silence to the end of your track(s) so that the echo has plenty of time to die out. When you select "Echo..." from the Effect menu, Audacity will ask you for two numbers.

The first number is the amount of delay between the echos, in seconds. The second value is the decay factor, which is a number between 0 and 1. A decay factor of 0 means no echo, and a decay factor of 1 means that each echo is just as loud as the

original. A value of 0.5 means that its amplitude is cut in half each time, so it dies out slowly. Smaller values will make it die out even more quickly.

The Echo effect is very simple and is not intended to be used in place of a Reverb effect, which simulates the sound of a room, concert hall, stage, or other natural environment. Audacity for MacOS and Windows comes with Gverb, a free LASPA Reverb plug-in.

Note that if you set the decay value to 1.0, you can use Echo to create loops that repeat as long as you want any never change volume.

Fade In

Applies a linear fade-in to the selected audio. For a logarithmic fade, use the [envelope tool](#).

Fade Out

Applies a linear fade-out to the selected audio. For a logarithmic fade, use the [envelope tool](#).

FFT Filter...

This is the most general type of filter. You define a curve that shows how much louder or quieter each frequency in the signal should be made.

If you're careful, you can use it to highlight exactly the frequencies you want. However, doing an FFT filter is more likely to result in artifacts, especially if the filter you draw is not smooth.

Invert

This effect flips the audio samples upside-down. This normally does not affect the sound of the audio at all. It is occasionally useful, for example when the left and right channels of a song both contain equal amounts of vocals, but unequal amounts of background instruments. By inverting one of the channels and not the other, the vocals will cancel each other out, leaving just the instrumentals. Obviously this only works if the exact same vocal signal is present in both of the channels to begin with.

Noise Removal...

This effect is ideal for removing constant background noise such as fans, tape noise, or hums. It will not work very well for removing talking or music in the background.

Removing noise is a two-step process. In the first step, you select a portion of your sound which contains all noise and no signal, in other words, select the part that's silent except for the noise. Then choose **Noise Removal...** from the Effect menu and click Get

Profile. Audacity learns from this selection what the noise sounds like, so it knows what to filter out later.

Then, select all of the audio where you want the noise removed from and choose **Noise Removal...** again. This time, click the "Remove Noise" button. It may take a few seconds or longer depending on how much you selected.

If too much or not enough noise was removed, you can **Undo** (from the Edit menu) and try **Noise Removal...** again with a different noise removal level. You don't have to get a new noise profile again if you think the first one was fine.

Removing noise usually results in some distortion. This is normal and there's virtually nothing you can do about it. When there's only a little bit of noise, and the signal (i.e. the voice or the music or whatever) is much louder than the noise, this effect works well and there's very little audible distortion. But when the noise is very loud, when the noise is variable, or when the signal is not much louder than the noise, then the result is often too distorted.

Future versions of Audacity may include improved versions of this effect.

Phaser...

The name "Phaser" comes from "Phase Shifter", because it works by combining phase-shifted signals with the original signal. The movement of the phase-shifted signals is controlled using a Low Frequency Oscillator (LFO).

Reverse

This effect reverses the selected audio temporally; after the effect the end of the audio will be heard first and the beginning last. Some people reverse small portions of audio to make inappropriate language unintelligible, while others believe you can hear subliminal messages if you listen to speech backwards. You can also create interesting sound effects by recording natural events and reversing the audio.

Wahwah...

Just like that guitar sound so popular in the 1970's.

This effect uses a moving bandpass filter to create its sound. A low frequency oscillator (LFO) is used to control the movement of the filter throughout the frequency spectrum.

The WahWah effect automatically adjusts the phase of the left and right channels when given a stereo recording, so that the effect seems to travel across the speakers.

VST Plug-ins

To use a VST plug-in effect, put the effect in the directory (folder) called "Plug-Ins", which should be in the same directory where Audacity resides. If there is no directory, create one. You will also need to download and unzip the [VST Enabler](#) to the same directory. The next time you launch Audacity, any plug-ins you added will appear in the Effect menu.

Freeverb2...

Audacity for Mac OS and Windows comes with a VST plug-in called Freeverb, which is in version 2 (hence "Freeverb2"). This effect implements a versatile and high-quality reverb effect.

Getting a good reverb sound depends a lot on the source audio and can take a lot of experimentation. One good strategy is to select a small portion of audio (a few seconds) and try to add reverb. Listen to it, then Undo and try it again with different settings. Keep doing this until you've found the settings you like the sound of best, and then Undo one last time, Select All, then apply the effect to your entire recording.

There are a lot of parameters to Freeverb2: Room size, Damping, Predelay, Lowpass, Highpass, Wet level, and Dry level.

Let's start with the wet and dry levels. Freeverb works by taking your audio signal and modifying it to create the reverberated sound, i.e. the sound you hear echoing off the walls of the room. This is called the "wet" part of the effect.

This is mixed together with the original sound, which is called the "dry" part of the effect, to produce the combination of direct (unreverberated) and indirect (reverberated) sound, which is what you would naturally hear. If you set the dry level to -infinity and the wet level to 0 dB, the result is something like standing outside of a concert hall: you can only hear the reverberated sound.

If you set the dry level to 0 and the wet level to -infinity, it's like standing right in front of the singer in a tiny room - all you can hear is the singer, and no reverb. A good place to start is to set both the dry level and wet level to 0 dB. However, you may want to experiment with lowering the dry level while you are experimenting with the sound of the reverb.

The room size parameter is self-explanatory. The smallest room size setting creates a quick, bright reverb, while the largest setting creates a long, drawn-out, and dark reverb. The damping parameter controls how the sound bounces off the walls - i.e. if it is mostly reflected or absorbed.

The predelay controls the delay between the dry signal (unreverberated) and the wet signal (reverberated) - usually there is some predelay because of the time it takes

sound waves to travel from the sound source to the nearest wall, and to the microphone. Larger predelays are suitable for creating an effect of a larger room.

Finally, the lowpass and highpass filters can be used to make the reverberated sound lower or higher. Increasing the lowpass filter filters out the high frequencies, and similarly increasing the highpass filter filters out the low frequencies.

Nyquist Plugins

To use a Nyquist effect, put the effect in the directory (folder) called "Plug-Ins", which should be in the same directory where Audacity resides. Upon installation, some nyquist scripts are included, so you needn't worry about creating this directory. The next time you launch Audacity, any plug-ins you added will appear in the Effect, Generate or Analyse menus.

LADSPA Plugins

These started out as a plugin format for Linux, but can now be used on Windows and OS X as well. There are lots of free plugins available, and no enabler is needed. More information can be found [on the web](#).

Analyze Menu

Note: plugins that accept input but produce no output will be placed in the Analyze menu. By default, this menu is often empty.

Beat Finder


no documentation available yet

Sample Analyze Nyquist Plugin ...

Will place labels that mark the beginning and end of a selection.

Toolbars

[Main Toolbar](#)



Obtained from
Omarshauntedtrail.com

[I/O Sliders and Input Selector](#)



[Edit Toolbar](#)



Main Toolbar



I Selection Tool

This is the main tool you use to select audio. Click in a track to position the cursor, or click and drag to select a range of audio. If you drag from one track to another, you can select multiple tracks. Extend a selection by shift-clicking a new point in the track.

Playback will always begin at the position of the selection cursor. If a range of audio is selected, only the selected range will play.

Envelope Tool

The envelope tool gives you detailed control over how tracks fade in and out, right in the main track window. When the envelope tool is selected, the amplitude envelope of each track is highlighted in a green line, with control points at the beginning and end of each track. To change a control point, click it and drag it to a new position. To add a new point, click anywhere in the track where there is not already a control point. To remove a point, click on it and drag it outside of the track until it disappears, then release.

Zoom Tool

This tool allows you to zoom in or out of a specific part of the audio. To zoom in, click anywhere in the audio. To zoom out, right-click or shift-click. If you have a middle button, you can middle-click to do the same thing as [Zoom Normal](#).

In addition, you can zoom into a region by clicking and dragging the mouse to highlight the region you want to see, then releasing the mouse button.

Enables the user to draw in to the actual waveforms. This is especially useful to eliminate small pops and clicks from material.

ALT + click smoothes a area of audio

CTRL + click & hold edits only one sample, no matter whether you move the mouse left or right.



Time Shift Tool

This tool allows you to change the relative positioning of tracks relative to one another in time. To use this tool, simply click in a track and drag it to the left or right. To align two tracks together or reset their time shift back to zero, use the [Align Tracks Together](#) or [Align with Zero](#) commands.



Cursor to Start

Places the cursor at the start of the project.

SHIFT + click expands the selection to the **start** of the project.



Play Button

Press the play button to listen to the audio in your project. You can also hit the spacebar to start or stop playback. Playback begins at the current cursor position. If a region of audio is selected, only the selected region will play. To quickly play the entire project, execute [Select All](#) before playing. If there are multiple tracks going to the same channel in your project, they will be mixed automatically for playback.



Record Button

Press the record button to record a new track from your computer's sound input device. Use the [Preferences](#) to configure the recording options. In particular, there is an option to record stereo or mono, and there is an option to play the other tracks while recording.

Recording always happens at the project's sample rate.



Pause Button

Will pause during playback, or during recording. Press again to unpause.



Stop Button

Press the stop button or hit the spacebar to stop playback immediately.



Cursor to End

Places the cursor at the end of the project.

SHIFT + click expands the selection to the **end** of the project.

Mixer Toolbar



Output

Input

Input Selector

These sliders control the mixer settings of the soundcard in your system.

The sliders will jump to zero (left most position) if the drivers do not expose these functions. This is the case with some digital I/O cards from RME and possibly others as well. Generally most cards that have a custom mixer application rather than the standard OS volume control won't work with this function of audacity. On these cards you will not be able to control the anything with these sliders, and you will need to set up the recording levels and the recording source using the card's mixer program.

Output Slider

This is the left hand slider that lets you control the output level of your soundcard. It actually controls the output setting of the soundcard driver.

Should your output sound distorted, this is not the first place to look. You might get lucky though, so give it a quick slide. Usually however, the distortion comes from the mix of your project **clipping**. You actually need to reduce the level of your tracks. It is suggested you reduce every track by the same amount.

Input Slider

This is the right hand slider that lets you control the level of the input selected in the Input Selector. It actually controls the recording level setting of the soundcard driver.

What you record is also determined in the mixer settings on your soundcard. There you can switch certain sources on and off, such as the Line Input, the Microphone Input or the CD Audio.

After selecting the corresponding input source with the Input Selector, you can control the level of the recording with this slider. Use the VU Meters on the Meter Toolbar to get the correct recording level.

Should your input sound distorted, you should lower this slider until no distortion can be heard, and the VU Meters do not hit the end of their travel.

Input Selector

Pick the input source you wish to record from. All these items are exposed by the soundcard driver, so the this of options will vary with different soundcards.



Here is an example of an ESS 1969 soundcard with the Line In input selected.

Edit Toolbar

All these tools perform the exact same function, as those accessible through the [Edit menu](#), [View menu](#) and [keyboard shortcuts](#).



Cut

Obtained from
Omarshauntedtrail.com

Removes the selected audio data and places it on the clipboard. Only one "thing" can be on the clipboard at a time, but it may contain multiple tracks.



Copy

Copies the selected audio data to the clipboard without removing it from the project.



Paste

Inserts whatever is on the clipboard at the position of the selection cursor in the project, replacing whatever audio data is currently selected, if any.



Trim

Deletes everything but the selection.



Silence

Erases the audio data currently selected, replacing it with silence instead of removing it completely.



Undo

This will undo the last editing operation you performed to your project. Audacity currently supports full unlimited undo - meaning you can undo every editing operation back to the last time the document was closed and re-opened.



Redo

This will redo any editing operations that were just undone. After you perform a new editing operation, including simply resizing a track, you can no longer redo the operations that were undone.



Zoom In



Zoom Out



Zoom to Selection




Zoom to entire Project

Tracks

[Audio Tracks](#)

Audio Track Controls

[Label Tracks](#)



[Track Pop-Down Menu](#)

[Time Tracks](#)



[Solo Button](#)


[Note Tracks](#)




[Mute Button](#)



[Track Delete Button](#)




[Gain and Panning Controls](#)



[Vertical Track Resizing](#)

Audio Tracks


Audio tracks contain digitally sampled sounds. In Audacity, a channels of sound is represented by one mono audio track, a two channel sound by one stereo audio track. Each audio track has a *sample rate*: common values include 8000 (telephone), 16000 (speech), 22050 (many games of recent years), 44100 (CD Audio), and 48000 (DAT).



From Audacity 1.2 on, you can specify a different sample rate for each track. You can import audio of any sample rate or bit depth and Audacity will resample and convert it to the project rate and bit depth on the fly, as well as upon Export. You can choose different sample rate conversion qualities for playback and Exporting.

(Note that in Audacity 1.0, all tracks must be at the same sample rate as the project.)

Label Tracks



Label Tracks can be used to annotate an audio file. They can be used for lyrics, markers, or notes, and they can even be used to save selections.

To **create a Label Track**, select New Label Track from the [Project Menu](#). Alternatively, simply click or select where you would like to place a label, and choose Add Label at Selection from the [Project Menu](#), and a Label Track will be created automatically if one doesn't already exist.

To **add a new label**, click or select where you want the new label to appear, then select Add Label at Selection from the [Project Menu](#), then type the name of the label, and finally press Enter or click outside of the label.



To **edit the name of a label**, click anywhere in it. Zoom in first if there are too many labels crowded together and you are unable to click on the one you want. When a label is selected for editing, it looks like the first label below:

When editing the name of a label, you are limited to using just the backspace key for editing. There is no way to move the insertion point to append a letter to the beginning of a label.

Once a label is selected, you can **move to the next label by pressing Tab**, and move to the previous label by pressing Shift-Tab.



To **delete a label** or multiple labels, select the area containing the label flags you wish to delete, and choose Silence from the [Edit Menu](#). Alternatively you can delete an individual label by clicking on it and pressing Backspace until you have deleted all of the characters in the label, then pressing Enter.

To **move labels**, use the normal editing commands like Cut, Copy, Paste, Delete, and Silence.

You can **save a selection in a label**. When you create a new label, the left selection edge determines the position of the label's flag. However, the right selection edge is also stored in the label, and when you click on it, the full original selection will be restored.

If you wish to apply this selection to only a subset of the tracks, shift-click in the label area to the left of each track's waveform to change whether each track is part of the selection or not.

Exporting and Importing label tracks

Label Tracks are saved when you save an Audacity Project, so if you just want to continue using the labels along with the same file, just save a project. But you can also export the labels to a simple text file, and import them.

This allows you to save information about the locations of events in an audio file for use in another program, and also provides a way to mass-edit the labels if necessary.

To export a Label Track, choose Export Labels... from the [File Menu](#). The exported file will contain one line per label, starting with the time offset in seconds, then a tab, and then the name of the label, **for example** :


```
1.217995 Bass intro
3.921073 Guitar enters
7.584454 Drums enter

11.070002 Chorus
```

To import a Label Track, choose Import Labels... from the [File Menu](#).

Time Tracks

This track consists of a single envelope.



It controls the playback speed within a given range, i.e. it facilitates **envelope-based resampling**. The envelope handling is exactly the same as for the volume envelope.



Clicking on the track panel reveals a few options, similar to the ones of the [track panel of an audio track](#).

"Set Range" lets you set the minimum and maximum pitch change. It is expressed in percentage values relative to 100%.

Start by choosing **New Time Track** from the **Project Menu**. Then change the range of time warping by clicking on the Track Pop-Down Menu and choosing Set Range (an option which only appears for Time Tracks)

The converter used will be the one(s) set up on the [Quality](#) tab of the [Preferences](#). You can make separate choices for Realtime (playback) and High Quality (exporting), trading speed against quality.

Older computers may be incapable of playing audio in real-time that makes use of a Time Track, however you can always Export the file as a WAV and then playback the WAV file later.

To edit the time warping, choose the envelope tool from the Control Toolbar - the same tool you use to edit amplitude envelopes in Audio Tracks. Click to create and edit control points. Drag points off the track to delete them.

The Time Track has a ruler. As you edit the warping, the ruler will warp to show you at what time playback will reach each place in the audio.

Note Tracks


Note tracks display data loaded from a MIDI file. They are currently not editable or playable. It's best to ignore the note tracks for now until further notice.

Track Delete Button

Clicking on this button **deletes** the track. This action is undoable.

Vertical Track Resizing


For *horizontal* resizing, which is called zooming, visit the [View Menu](#) page.



Vertically resizing a track is fairly simple.

- **Click and hold the mouse button on the resizing handle** (see figure),
- **move mouse up or down** until the track has the wanted vertical size
- **release mouse button**
- Note:
- **Stereo tracks** have two resize handles. The one beneath the upper track lets you resize both tracks in relation to one another.
- The resize handle beneath the lower track resizes both tracks equally the same way. It changes the overall vertical size of the whole stereo track.

Pop-Down Menu



- **Name** - Edit the name of the track
- **Move Track Up/Down** - Move Track Up or Down in the display
- **Waveform** - Traditional display of audio material. It displays the amplitude of the audio over time. This is the default display mode.
- **Waveform (dB)** - Like Waveform, but logarithmic instead of linear vertical units . It displays the amplitude in dB of the audio over time.
- **Spectrum** - Displays the frequency spectrum of the audio over time.
- **Pitch (EAC)** - Tries to detect the pitch of the current audio and displays that information over time.
- **Mono** Set playback of this single channel track on the left and right channels.
- **Left Channel** - Set playback of this single channel track on the left channel.
- **Right Channel** - Set playback of single channel track on the right channel.
- **Make Stereo Track** - The selected track and that beneath it is turned in to one stereo track.
- **Split Stereo Track** - Turn one stereo track in to two single channel tracks.
- **Set Sample Format** - Pick the sample format for this track. The result is that all newly written data is written in this format, no matter

what format the original material was in. Newly written data can be produced by effects and generator plugins for example.

- **Set Rate** - Set the sample rate of this track. Upon import, the track sample rate is set to that of the imported material. On playback or exporting the audio will be re-sampled to the project rate on the fly.

This option can be useful to change the speed and pitch of audio, or if the sample rate of the source was wrong.

Track Solo Button

 Clicking on this button will switch this track to **solo mode**. This action is **not** reversible with the Undo function, but you can un-solo a track simply by clicking it again.

 In solo mode, only tracks that have the solo button activated (**turning the button red**), will be played.

To deactivate the solo mode for that track, click on the button again. (button turns grey again)

The state of the solo buttons is NOT saved in the project file, and it does not affect Exporting or Mixing.

Track Mute Button

 Clicking on this button will **mute** this track , i.e. this track is switched off without deleting it. This action is **not** reversible with the Undo function, however you can un-mute a track just by clicking the mute button again.

 Muted tracks are easy to spot, because their **Mute buttons are green**.

To unmute the track, click on the button again. (button turns grey again)

The state of the mute buttons is NOT saved in the project file, and it does not affect mixing (i.e. muted tracks are still mixed when you export or Quick Mix).

Gain and Panning Controls



Gain (Volume)

This slider controls the track volume, or rather the overall gain of that particular track.

Its default position is in the middle, as shown in the picture to the left. This setting represents **0 dB**.

Tip: Normally the control moves in increments of 3 dB. Hold down Shift to adjust by 1 dB at a time.



Pan


This slider set the panning position of the track in the stereo field.

Its default position is in the middle, as shown in the picture to the left.

Tip: Normally the control moves in increments of 10%. Hold down Shift for finer control.

Vertical Track Resizing

For *horizontal resizing*, which is called zooming, visit the [View Menu](#) page.



Vertically resizing a track is fairly simple.

- **Click and hold the mouse button on the resizing handle** (*see figure*),
- **move mouse up or down** until the track has the wanted vertical size
- **release mouse button**

Note:

Stereo tracks have two resize handles. The one beneath the upper track lets you resize both tracks in relation to one another.

The resize handle beneath the lower track resizes both tracks equally the same way. It changes the overall vertical size of the whole stereo track.

Preferences

Audio I/O

- The two drop-down list controls in this pane is where you configure which device **Audacity** should use for sound input and output. This basically means which sound card, USB sound device, etc you want to use. If your sound card support multiple sound inputs, e.g. has a microphone input and a line-in, then you select these on the [mixer toolbar](#). No need to adjust this if **Audacity** is already playing to and recording from the correct place.
- **Record in Stereo**: If this is checked, recordings you make in Audacity will be stereo, otherwise they will be monaural. Note that the interpretation of mono recordings depends on your system - on most systems making a mono recording of a stereo input will result in the two channels being mixed, but on other systems it could give you just the left channel, for example.
- **Play other tracks while recording new one**: Check this box if you want to record multiple tracks over top of one another! Record the first track as usual. Then, with this box checked, click record again. The first track you recorded will play, but **Audacity** will also record a new track at the same time, maybe allowing you to sing harmony with yourself or add real-time sound effects to a recording. You may notice that when you play the two tracks you recorded together, they aren't synchronized. This is normal and is not the fault of **Audacity**. To fix it, you will need to grab the [Time Shift](#) tool and slide one of the tracks around until it sounds right.
- **Hardware Playthrough (Mac Only)**: This sets up your computer to send the sound coming in to be recorded straight back out of the speakers, so you can hear what you are recording. Hardware playthrough has no latency and takes no CPU, but is only possible if your input and output are on the same audio device, and that device supports it.
Windows and Linux users can enable this using your operating system's Mixer, by simply turning up the level for the input you are using and unmuting it.
- **Software Playthrough**: This enables you to hear what you are recording by taking the data about to be recorded by audacity, and playing it back out of the sound card. This works with any duplex sound card, or combination of sound devices, so you can record in through an external USB device and play back through the computer's built in sound device or vice versa.
Because of the delays in processing the audio on the computer, there will be some delay or **latency** in the sound that you hear. It also needs some CPU power, so if you have an old machine and are getting breaks in your recordings try turning it off.

Quality

This panel sets the default sample rate and bit depth for new projects. It does not affect any existing project - to change the sample rate of an existing project, just click on the rate at the bottom of the window and select a new rate from the pop-up menu.

The bit depth affects imported material, if it's copied, as well as newly recorded material. The advantage of the 32-bit float resolution is that effects return material at a this resolution, which improves the final mix, because the computer has more precise material to work with. If you aren't doing any processing then you may want to save some space by setting this to 16 bit, if your hardware supports this.

File Formats

- **When importing uncompressed audio files...**
 - Make a copy of the file to edit: selecting this means that **Audacity** will run a little bit slower, but it will always have its own copy of any audio you are using in a project. You can move, change, or throw away your files immediately after you open or import them into Audacity.
 - Edit the original in place: selecting this operation means that **Audacity** depends on your original audio files being there, and only stores changes you make to these files. If you move, change, or throw away one of the files you imported into **Audacity**, your project may become unusable. However, because **Audacity** doesn't need to make copies of everything, it can import files much more quickly this way. If you want to send projects to other people or archive them, then select this option.
- **Uncompressed Export Format:** This lets you select the format that Audacity will use when you export uncompressed files, i.e. anything other than lossy formats, such as OGG and MP3. Audacity can output a huge amount of different formats with the *Other* option, which lets you define any aspect of the multitude of different formats.
For most users, 16 bit WAV (Microsoft) or AIFF (Mac) files will be the most compatible.
- **Ogg Export Setup:** Use this control to set the quality of Ogg Vorbis exporting. Ogg Vorbis is a compressed audio format similar to MP3, but free of patents and licensing fees. A normal quality Ogg Vorbis file is encoded with a quality setting of "5". Note that unlike MP3 encoding, Ogg Vorbis does not let you set a bitrate, because some audio clips are easier to compress than others. Increasing the quality will always increase the file size, however.
- **MP3 Export Setup:** Use these controls to locate your MP3 encoder and set the quality of MP3 encoding. Higher quality files take up more space, so you will need to find the level of quality you feel is the best compromise. For more information, see [Exporting MP3 Files](#).

Spectrograms

You can view any audio track as a Spectrogram instead of a Waveform by selecting one of the Spectral views from the track's pop-up menu. This dialog lets you adjust some of the settings for these spectrograms.

- **FFT Size:** The size of the Fast Fourier Transform (FFT) affects how much vertical (frequency) detail you see. Larger FFT sizes give you more low frequency resolution and less temporal resolution, and they are slower.
- **Grayscale:** Select this for gray spectrograms instead of colored ones.
- **Maximum Frequency:** Set this value anywhere from a couple of hundred hertz to half the sample rate (i.e. 22050 Hz if the sample rate is 44100 Hz). For some applications, such as speech recognition or pitch extraction, very high frequencies are not important (visually), so this allows you to hide these and only focus on the ones you care about.

Directories

Use this panel to set the location of Audacity's temporary directory. Audacity uses this directory whenever you work on a project that you haven't saved as an Audacity Project (AUP file) yet. You have to restart Audacity (close and open it again) for changes to the temporary directory to take effect.

This directory can contain a lot of files and is used a lot, so choose a disk with plenty of free space to put it on. **Never** put it on a network drive.

Interface

- **Autoscroll while playing**
page flips to the right when the cursor travels to the right hand end of the wave display, either during playback or recording. Always keeps the cursor on screen, but may cause breaks in recording or playback if you have a slow machine.
- **Always Allow Pausing**
- **Update spectrogram while playing**
Tries to keep the spectrogram updated while playing. This requires a lot of cpu power, and really isn't a good idea unless you're playing a really small piece or have a powerful CPU and plenty of RAM. In combination with *Autoscroll while playing* this can cause pauses during playback or recording.
- **Enable Edit Toolbar**
Gives you a toolbar of a few useful functions. Check out the page on the [Edit Toolbar](#) for more information.
- **Enable Mixer Toolbar**
Turns on display of the [Mixer Toolbar](#) used to control soundcard mixer settings.
- **Enable Meter Toolbar**
Turns on display of the Meter Toolbar, with VU meters you can use to monitor the recording and playback levels of your project.

- **Quit Audacity upon closing last window**
Sets Audacity to shut down when the last project window you are working is closed. If you turn this off then audacity will keep re-spawning new blank project windows until you do File > Quit.
- **Enable dragging left and right selection edges**
Lets you manipulate selections by clicking on the edges and dragging them along the timeline.
- **Enable dragging left and right selection edges**
Sets the transport control button order as:Pause, Play, Stop, Home, End, Record. Rather than: Home, Play, Record, Pause, Stop, End.
- **Language**
Lets you change the language of the menus and other onscreen text to the language of your choice. Simply choose the correct option from the drop-down list. Note that for languages with unusual character sets you need the correct fonts installed. If your menus come up all ????? that's what's wrong.
- **Minimum of dB mode display range**
This sets up the scale used for logarithmic waveform displays. Because Decibels are logarithmic, it is impossible to have a 0 point on the scale, so some very quiet sounds will display as silence. If you you reduce this effect by increasing the range, then the louder part of the scale gets squashed, so there is a choice of settings for different uses. -48dB is a good compromise for general work.

Keyboard

- **Commands/Key Combination List**
All commands in Audacity that can be triggered by a keyboard shortcut are listed here. Click on the shortcut and the assigned keystroke is highlighted in the command/shortcut list. Click on the "Clear" button to erase the keystroke assigned to the function selected in the Commands column.
- **Adding new shortcuts/changing shortcuts**
Click on the command in the list. Click in to the field next to the "Set" button and press any key or combination of command key(Shift,CTRL,ALT) and alphanumeric key(everything else). The description of that key or key combo will appear in the field or replace the existing keystroke. By clicking on the "Set" button, you can add or change this key or key combo as a shortcut to the function selected in the Commands list.
- **Saving your setup** You can export your keyboard shortcut settings for safekeeping or to import into audacity on another machine. Use **Save...** to save you settings to a file, **Load...** to load a saved settings file, and **Defaults** to restore the default audacity settings.

Head over to [this page](#) for step-by-step description of the dialog and how to edit the command/keyboard combination list with a text editor.

Keyboard Shortcuts

!!Apple MacOS X users!! CTRL = Command

Obtained from
Omarshauntedtrail.com

File Commands

New Project CTRL+N

Open Project CTRL+O

Close Project CTRL+W

Save Project CTRL+S

Preferences CTRL+P

Editing Tools

Selection Tool F1

Envelope Tool F2

Editing Tool F3

Zoom Tool F4

Timeshift Tool F5

Multi Tool F6

cycle tool forward D

cycle tool backward A

Navigation, Playback & Recording

Play/Stop SPACE

Loop SHIFT+SPACE

L

Pause P

Record R

Preview 1 second 1

Play From Cursor To Selection B

|

|

Zoom In CTRL+1

Zoom Normal CTRL+2

Zoom Out CTRL+3

Fit In Window CTRL+F

Fit Vertically CTRL+SHIFT+F

Zoom to Selection CTRL+E

Import Audio CTRL+I

Edit Commands		Create Label	CTRL+B
Undo	CTRL+Z		
Redo (Windows)	CTRL+Y		
Redo (Mac/Unix)	CTRL+SHIFT+Z	Repeat Last Effect	CTRL+R
Cut	CTRL+X		
Copy	CTRL+C		
Paste	CTRL+V		
Delete	CTRL+K		
	DELETE		
Silence	CTRL+L		
Duplicate	CTRL+D		
Select All	CTRL+A		
Find Zero Crossings	Z		

Obtained from
Omarshauntedtrail.com

Edit or Create Keyboard Shortcuts

Where?


In the [preferences dialog](#), go to the [keyboard section](#) or edit the keyboard configuration file directly with a text editor. You will find this file in the directory where you installed Audacity. Its filename is **Audacity-keys.xml** .

Audacity will NOT load this keymap upon startup, but save them in the registry(Windows). Should you make changes to the Audacity.xml file with a text editor, you'll need to load that shortcut file with the "Load" button to make Audacity use them. Thus you can save your keymaps under any filename with an XML extension and keep different keymaps around for different purposes.

Audacity will NOT store the current keymap in the file Audacity.xml. Should you want to make changes to the current keymap with a text editor, you'll have to use the "Save" button to save that keymap to an XML file first. After editing it you'll have to "Load" it again.

Editing and Creating Shortcuts with the Preferences Dialog

Changing a shortcut:



1. Click on the command in the list. The command to be changed is now highlighted.
2. Click in to the box to left of the **Set** button. If a shortcut was assigned to the command previously, it will be displayed in this box.
3. Hit the key or key combination you wish to assign to this command. Then processed to the next step.

Click on the Set button to assign the chosen key or key combination to the chosen command.

Clearing a shortcut:

1. Click on the command in the list. The command to be cleared is now highlighted.
2. Click on the clear button to clear the command.

Editing and Creating Shortcuts with a text editor

The keyboard shortcuts you were editing with the dialog are stored in a simple XML file. Thus, you can edit it with any text editor.

Each keyboard map starts with

```
<audacitykeyboard audacityversion="1.2">
```

and ends with

```
</audacitykeyboard>
```

Each command is defined with the **command** tag. An example :

```
<command name="Undo" label="Undo" key="Ctrl+Z" />.
```

name

The actual command as Audacity understands it.

label

Command description as seen in the Preferences keyboard map editor.

key

This is the actual shortcut. Check out what the format is in the Preferences dialog.



[How do effects work in Audacity ?](#)

1. Select some audio

Use the mouse or [keyboard shortcuts](#) to select some audio. You can listen to that selection as many times as you like and adjust the boundaries of the selection.

2. Select an effect

You can pick any of the internal effects or a LADSPA or Nyquist plugin loaded by audacity. Windows & Mac users can also use VST plugins via the VST Enabler from www.audacityteam.org/vst/.

3. Tweak the parameters to your liking

Play around, go wild, honk your horn. Make those tough decisions without telling anyone why.

4. Let the software work

Hit the OK button. Or hit the Cancel button if you're unsure about something, like what time you were supposed to take the pizza out of the oven.

5. Not happy ? UNDO

That effect didn't cut it ? Did that delay show up too late and you wondered where the summer went ? No problem.

Hit UNDO in the Edit menu (or press CTRL+Z).

Then go back to 3. In fact you can hit undo so many times, YOU might disappear. Or perhaps just all your edits. Who knows.

[Basic Calculus of Acoustics](#)

Basic Calculus of Acoustics

This is so terribly useful for so many things. You can use this stuff to calculate just how bad parallel walls are for sound, because of *Standing Waves*. What's that ? Coming attractions include exciting info on how to improve the acoustics of your room and what standing waves, bass traps and reflecting will do for your listening karma.

You may enjoy finding out that a \$1000 subwoofer won't work in a small 3x3 meter room, 'cause the big energy bass action will happen beyond your walls. After all, the first energy peak happens at 1/4 of wavelength. Check it for yourself and get smart the next time you buy those 'gotta get buried with/under/in them' speakers.

This all happens in plain fresh air:

The speed of sound is 343 meters per second at 20 degrees Celsius

+/- 0.6 m/s for every degree Celsius more/less (-> 340 m/s at 15 degrees Celcius)

$$c = \lambda * f$$

*speed of sound = wave length * frequency*

$$x = (1000 / c) * y$$

*time [milliseconds] = time[ms] for one meter * distance [meters]*

Examples (at 20 degrees Celsius > c = 343 m/s):

distance [meters] time [milliseconds]

0.34	1
3.43	10
6	18
10	29
50	188

[Reverb](#)

What is Reverb ?


A Reverb simulates the component of sound that results from reflections from surrounding walls or objects. It is in effect a room simulator. Some people think it's just a delay effect with some filters, but it's way more complex than that. Reverb effects (software plugins or external hardware units) provide an interface to their changeable parameters that need some explaining. Let's look at a simple room first.

Basic Simulation of a Room

Our model is a simple room with four straight walls, a sound source and a listener. In *Figure 1* the arrows stand for the path of traveling sound.

The listener hears the **DIRECT** signal first. The **DIRECT** signal is also referred to as the **DRY** part of the signal when using any effect. Most digital reverbs produce two parts: The **Early Reflections** and the **Reverb** component.

I. Early Reflections




The first Early Reflection reaches the listener milliseconds after the direct signal does.

The path of the Early Reflections is longer. The difference in time between the arrival of the direct signal and the first Early Reflections is measured in milliseconds.

Figure 1

II. Reverb



The sound reflects off the walls and objects in the room, and in time individual reflections disappear and the **Reverb** develops.

Predelay

The time between the reception of the **DIRECT** signal by the listener and start of the **Reverb** portion of the effect is called **Predelay**. This is a parameter in many digital reverb effects, and it is expressed in milliseconds (ms).

Reverb Time

The time difference between switching off any sound generator and the level of the reverb resulting from that sound dropping by 60dB is called **RT60**.

Figure 1

This is usually referred to as the **Reverb Time**. When anyone refers to the reverb time of a real room or that of a digital reverb, RT60 is what they're talking about.

Most digital reverbs feature this as a parameter.

Common parameters of a digital Reverb

Early Reflections

Usually controlled via **Room type** and/or **Room Size** parameters.

The **EARLY REFLECTIONS** enable the human brain to quickly identify the room size. They are therefore the most critical part of a reverb effect, if a room simulation is the goal.

Predelay

Parameter Range : 0 to x milliseconds

Time between Direct Signal and start of Reverb portion.

The more space between you and walls or other reflective objects, the greater the Predelay.

Reverb Time

Parameter Range : 0.1 to x.x seconds

The greater the enclosed(!) volume of the space you're in, the longer the reverb time. This is of course true only to a certain extent. The larger the room or chamber, the louder your source signal needs to be to produce any significant level of reverb tail. A good way to demonstrate this is to visit some huge cathedral and talk quietly. Then yell a bit and you'll actually hear some significantly loud reverb. The RT60 value is measured with a noise impulse, but the level at which that impulse was sent out should be known too.

Many reverb effects let you set different reverb times for different frequency bands. Most reverb units on PC and Mac architectures provide two bands, so you get to set the reverb time for each frequency band individually.

Reverb Damping

Usually a Highshelf and/or a Lowshelf EQ with variable Frequency.

These EQs are applied continuously to an evolving reverb of a piece of audio. Over time, as the reverb fades, so do certain frequencies, either in the lower or higher frequency range. The more less-reflective surfaces the sound bounces around from, the more the reverb is dampened the higher frequencies.

Freeverb features a Damping parameter that dampens the upper frequencies.


Reverb EQ

Usually a Highshelf and/or Lowshelf EQ with variable Frequency.

These EQs are only applied to the output of the reverb or perhaps even just the reverb portion of the Reverb effect. Some reverbs have Highshelf and Lowshelf EQs for the Early Reflections as well.

Freeverb has both a Hicut and Lowcut parameter.

Here's a screenshot taken from the reverb plugin Renaissance Reverberator by [Waves](#).



Reproduced with permission

Here we have the Highshelf and Lowshelf EQs for Damping(left) and for the final output(right). The upper middle display depicts the Early Reflections placement and level, as well as the Reverb portion start and fading. Above this display is the *Reverb Type* selector. This changes the placement and level of the Early Reflections. The *Size* slider has influence on the Early Reflections as well.

Blank concrete will reflect sound quite well across the frequency spectrum. A large thick curtain or padded walls will reflect sound poorly. A large club filled with people sounds completely different from an empty one. People, comfy chairs, curtains, carpets, beds and many more objects reflect only part of the sound that hits them.


Delay

What is a Delay ?



This effect uses a RAM buffer to store the audio for a certain amount of time, after which it outputs this audio.

The RAM buffer will not always be just large enough to store the amount of audio set by the current delay time, but will actually be considerably larger sometimes, because many delay effects allow for their delay parameter to be automated.



Multi-Tap Delays

This effect simply incorporates several discrete delay lines, each with its own delay time, level(amplitude) and panning. The feedback parameter is usually shared by all delays in a multi-tap delay. Some multi-tap delays feature a modulation section that modulates the delay times, which can be used to create [chorus](#) effects.

Common parameters of a digital Delay

Delay Time

Parameter Range : 0.000 to xxx.xxx seconds

The upper boundary of this parameter(xxx.xxx) depends on the amount of RAM in the hardware unit, or the amount of RAM reserved by the software plugin. For software plugins, the programmers or designers of the plugin set this upper boundary.

The delay software plugin Supertap by Waves for example comes in two flavours on Protools TDM systems. One allows for delays of up to two seconds, but another version accommodates delay times of up to six seconds. This is due to the availability of RAM on the DSP hardware of a TDM system. The native(->CPU) version of this plugin for Windows for example only allows for delays of up to four seconds in length, but there are delay plugins that allow for delays of several minutes.

Feedback

Today's delay units and plugins will feature one of several different types of feedback parameters :

a) Percentage Range : 0 to 99(or 100)%

This is used by most hardware units and some software plugins. It determines the amount of a delayed signal to be sent back to the RAM buffer. The same delay time is applied. This is the most common parameter type for feedback.

b) Decay Time : 0.1 to xx seconds

A slightly different approach. The plugin or hardware unit determines the amount of delayed signal, that it sends back in to its RAM buffer, itself. Some may find this more intuitive. It isn't widely used however.

c) Feedback Delay Time : 0.000 to xxx.xxx seconds

This is like an additional delay line that takes the already delayed signal as input. This kind of parameter also requires a level setting, which is usually a percentage parameter, like a).

EQs & Filters

Many delay plugins feature highshelf and/or lowshelf EQs, or they feature highpass and/or lowpass filters. These EQs or filters usually process each delay. The idea behind this is usually to simulate distance, because with increasing distance, the lower and higher frequencies are dampened.

Hi- & Lowshelf EQs

- Gain : -xx to xx dB

Frequency: 20 Hz - 20 kHz

Low- & Highpass Filters

- Frequency: 20 Hz - 20 kHz

Filter Strength : 6, 12, 18 or 24 dB/Octave

PANNING

These are the most common types of stereo panning parameters you'll find.

LEFT RIGHT LEFT RIGHT LEFT RIGHT


-63 +64 0% 100% -100% 100%

Simply a panning control for the delay. The feedback parameter may also have its own panning control.

Chorus

What is a Chorus ?

A chorus is a [Delay](#) for which the delay time is modulated by an LFO (**L**ow **F**requency **O**scillator). The most used LFO waveforms are sine and triangle.



Most modern chorus effect unit, be they plugins or hardware based, also spread the effect signal across the stereo field. Sometimes a parameter to control the depth and other behaviours may be encountered.

The effect can produce interesting sounds, one of which is to produce the illusion of more than one instance of the instrument/vocalist being present. What you get out of this effect, largely depends on your desire to experiment.

Play with the parameters one at a time and experiment.

Common parameters of a Chorus

Delay Time

Parameter Range : 0.0 to xxx.xx milliseconds

This is your normal delay parameter.

Modulation Range

Parameter Range : 0-100% or sometimes in milliseconds

This determines the maximum amount of modulation applied to the delay time.

Modulation Rate

Parameter Range : 0 - 20 Hz (frequency)

Also sometimes referred to as the LFO frequency. Higher frequencies produce extreme effects.

Frequently Asked Questions

I. General

1. [Is Audacity really free? Why?](#)
2. [Who developed Audacity?](#)
3. [How can I help to improve Audacity?](#)
4. [I found a bug! Now what?](#)
5. [Can I include or sell Audacity on a CDROM ?](#)

II. Installation

1. [Windows: What do I do with the .exe file?](#)
2. [Windows: How do I uninstall Audacity?](#)
3. [Mac: What do I do with the .sit file?](#)
4. [Unix: Why do I get failed dependencies when installing from the RPM?](#)
5. [Unix: Why do I get an error while loading the shared library libwx_gtk.so?](#)

III. Using Audacity

1. [How do I mix two tracks together?](#)
2. [Can I remove the vocals from a recording?](#)
3. [When I try to record harmonies with myself, why are the two tracks out of](#)

[sync?](#)
4. [Why doesn't Audacity ship with an MP3 encoder?](#)

General

1. Is Audacity really free? Why?

Yes, in almost all senses. Audacity is not only distributed free of charge, but you are also free to do almost anything you want with it.

Audacity is distributed under the terms of the [General Public License \(GPL\)](#).

You are free to use this program for personal or commercial purposes. You are also free to give it away or sell it.

The source code to this program is freely available on the web, and you are free to modify it for your own use, however any changes you make must also be distributed under the GPL.

Audacity was build using wxWidgets, a software library which is released under the less-restrictive LGPL. For more information, please visit <http://www.wxwidgets.org>.

The authors of Audacity decided to release it under the GPL for many reasons. Some of us do it out of generosity. Some of us do it for moral reasons, because we feel that all software should be free, while others of us believe that there is a need for both free and commercial software in the world.

One reason Audacity is free is so that it will be more popular. Many of us would rather see a million people happily using Audacity for free than have a thousand people pay us. Most users are more friendly when they get something for free.

Yet another reason is that it encourages collaboration. If Audacity was shareware, it is unlikely that dozens of

Installation

1. Windows: What do I do with the .exe file?

From the browser, select the option to "Run from current location". This launches the Audacity setup program, which will install the Audacity program files and create an entry in the start menu.

If you already downloaded the file, just double-click on it to start the installer.

If you prefer not to use the Installer program, you can also download Audacity as a ZIP file, which you can decompress using WinZip or any other extraction program.

2. Windows: How do I uninstall Audacity?

Audacity can be uninstalled by opening Add/Remove Programs in the Control Panel. Select "Audacity" from the list and click the button "Add/Remove". This launches the Audacity uninstall program.

3. Mac: What do I do with the .sit file?

You need a recent version of StuffIt Expander. StuffIt Expander comes with all Macintoshes and is usually configured by default with all Mac web browsers. If Audacity does not decompress automatically, drag "audacity.sit" to StuffIt Expander to decompress it.

There is no need to install Audacity. Just drag the Audacity folder to your Applications folder, or wherever else you would like to put it. To uninstall, just drag the entire folder to the

4. Unix: Why do I get failed dependencies

people around the world would have contributed code, bug fixes, documentation, and graphics.

2. Who Developed Audacity?

Audacity was started in the fall of 1999 by Dominic Mazzoni while he was a graduate student at Carnegie Mellon University in Pittsburgh, PA, USA. He was working on a research project with his advisor, Professor Roger Dannenberg, and they needed a tool that would let them visualize audio analysis algorithms. Over time, this program developed into a general audio editor, and other people started helping out.

Today, Audacity is developed using Sourceforge, an online site that allows people around the world to collaborate on free software projects. See <http://www.sourceforge.net> for more information. Dozens of people have contributed to Audacity, and progress is continually accelerating.

3. How can I help to improve Audacity?

There are many ways you can help. If you are a programmer and you know C++, we can always use more developers. We especially need more Windows and MacOS programmers, as many of the current developers prefer Linux.

If you are bilingual, you can help us translate Audacity into another language.

If you are good at writing, you can help us write documentation and tutorials.

Finally, anyone can submit bug reports and suggestions. The more detailed,

when installing from the RPM?

First, ensure that the wxGTK is installed; a link can be found from our Linux page - <http://audacity.sourceforge.net/download/linux>. Be sure that you use the appropriate version of wxGTK - Audacity 1.2.x requires wxGTK 2.4.x, and it will not work with any other version.

5. Unix: Why do I get an error while loading the shared library libwx_gtk-2.4.so?

This could be one of two problems:

- If you installed the wxWidgets RPM, be sure to run "ldconfig" (as root) so your system knows about the new library.

- If you installed wxWidgets from source, it may have been installed in /usr/local, but /usr/local/lib may not be in your library path. Add /usr/local/lib to the file "/etc/ld.so.conf" and then run "ldconfig". You may also want to add "/usr/local/bin" to your path, so that programs can find the "wx-config" utility.

Using Audacity

1. How do I mix two tracks together?

Audacity mixes automatically. All you have to do is import two tracks into the same project, which you can do using the "Import Audio" command in the Project menu, or simply by dragging the two audio files to Audacity. When you press the record button, Audacity generates a new track automatically, also.

In order to save your mix, you can either Export the project, which will automatically mix all of your tracks together, or you can select your tracks and use the "Quick Mix" command in the Project menu.

2. Can I remove the vocals from a recording?

the better.

If you would like to help out in any one of these ways, please email the developers at audacity-devel@lists.sourceforge.net. You may also want to go to the Sourceforge web site and join the mailing list yourself.

4. I found a bug! Now what?

The most important thing when reporting a bug is to be as specific as possible. Give us enough information that we can reproduce the bug ourselves, otherwise it's unlikely that we'll be able to fix it.

Be sure to let us know what operating system you run (such as Windows 98, MacOS 9.1, Fedora Linux 3, etc.) and any other information about your computer that you think might be relevant.

Then, can you reproduce the bug? If it happens consistently, tell us the exact sequence of events which causes the bug to occur. If you get an error message, make sure you send us the exact text of the error message.

We want to squash all of the bugs! Thanks for taking the time to help us track them down.

5. Can I include or sell Audacity on a CDROM ?

Basically the same as the [first question](#), so that is a big YES. You are also free to do almost anything you want with it. We do like to hear about it of course, but as long as you comply with the [GNU Public License](#), you'll be ok. This means that you need to make any

With some stereo recordings, it is possible to remove the vocals because of the way in which the recording was mixed at the studio. Often, the vocals are placed in the exact center of the recording, while all other instruments are slightly off-center. If you subtract the right channel from the left channel, the vocals get completely canceled out, leaving only the other instruments.

This only works on some recordings!

To attempt this in Audacity, import a stereo recording, then click on the track pop-up menu (the little down-arrow next to the name of the track) and select "Split Stereo Track". Now select the lower track (the right channel) and use the "Invert" effect (from the Effect menu). Finally, use the track pop-up menus to make both channels Mono channels, and then mix them together using Quick Mix. If you're lucky, the vocals will be gone.

3. When I try to record harmonies with myself, why are the two tracks out of sync?

This is normal and it happens because there is a small delay between when Audacity starts playing sound and when it actually reaches your speaker. Audacity does not automatically try to correct for this delay, and you must do it manually.

To correct this delay, use the [Time Shift](#) tool to slide one of the tracks over until they line up. Note that you can use the tool while you are listening to the recording.

4. Why doesn't Audacity ship with an MP3 encoder?

Unfortunately, the algorithm to encode or create MP3 files is patented, and **Audacity** would not be able to include an MP3 exporting algorithm without either charging a fee or violating the law in many countries.

changes you make to Audacity available, including the source code. The source code doesn't need to be on the CDROM, but it should be available via the internet for example.

As a compromise, Audacity comes with the ability to use other MP3 encoders, which you must download separately. It is up to you to make sure you are in compliance with any licensing restrictions imposed by MP3 encoders.

For more information, see the page on [Exporting MP3 files](#).

Exporting MP3/OGG Files

Audacity has OGG export and import functionality built in. You can set the quality of the exported OGG files in the preferences. OGG is vastly superior to MP3, especially at lower bitrates.

OGG should be preferred to MP3, except in situations in which MP3 is required. This can be the case for external players, which do not support OGG.

Audacity cannot encode MP3 files by itself, because the MP3 encoding algorithm is patented and cannot legally be used in free programs. However, Audacity has been programmed to recognize other existing MP3 encoders that you can download separately. All you have to do is obtain the appropriate MP3 encoder for your computer and then show Audacity where it is located.

If you use... You need to...

Windows Download **LAME** and look for the file called **lame_enc.dll**. You can find binary distribution on the net by Google searching for "lame compile download binary" or visiting [this site](#).

Linux/Unix Download **LAME** from its [home page](#) and compile it as a shared object, then look for the file called **libmp3lame.so** (usually in /usr/local/lib).

Macintosh Download the appropriate **LAMELib** from [here](#).

For links to these MP3 encoders, go to the Audacity web page (<http://audacity.sourceforge.net>) and click on the link for your operating system.

The first time you try to export an MP3 file, Audacity will ask you to locate your MP3 encoder. Locate the file indicated above. From then on, Audacity will not need to ask you again and you will be able to export MP3 files easily!

Support for more MP3 encoders may be coming in the future. Let us know if there is a particular MP3 encoder you want supported.

What is an MP3 file and how does it differ from WAV and AIFF files?

Why doesn't Audacity ship with an MP3 encoder?

MP3 (MPEG II, layer 3) is a popular format for storing music and other audio. A typical MP3 file is one tenth the size of the original WAV or AIFF file, but it sounds very similar. MP3 encoders make use of psychoacoustic models to, in effect, "throw away" the parts of the sound that are very hard to hear, while leaving the loudest and most important parts alone.

Unfortunately, the algorithm to encode or create MP3 files is patented, and Audacity would not be able to include an MP3 exporting algorithm without either charging a fee or violating the law in many countries.

Unfortunately, no MP3 encoder is perfect, and so an MP3 file will never sound quite as good as the original. Still, most people find that the quality of an MP3 file is virtually indistinguishable from a CD when played on headphones or on small computer speakers, which is why the format is so popular.

Credits

The Audacity manual was written by **Anthony Oetzmann**, with major contributions by **Dominic Mazzoni**.

Anthony Oetzmann is the primary documentation writer for Audacity. He wrote most of the manual you are reading now, and has also influenced the program design significantly. He is a professional audio engineer who lives and works in Germany.

Dominic Mazzoni is the Audacity project leader and Audacity's primary programmer. He began the project while a graduate student at Carnegie Mellon University in Pittsburgh, PA. He is currently a research programmer at the NASA Jet Propulsion Laboratory in Pasadena, CA.

The Audacity logo was designed by **Harvey Lubin** of Agrapha.