

Sound Studio 4

Record, Edit, and Add Effects to Audio

Users Manual for Sound Studio 4.1

April 5, 2011

Written by Lucius Kwok

Edited by Agnes Chan



Contents




Introduction	1
Getting Started.....	2
How-To Tutorials	3
Reference for Sound Studio.....	8
File Formats	8
Preferences	9
Windows	11
Document Window	11
Keyboard Shortcuts	12
Input Levels Window	17
Commands	18
File Menu	18
Edit Menu.....	19
Audio Menu.....	21
Insert Menu	23
Filter Menu	24
View Menu.....	31
Window Menu	33

Introduction

About Sound Studio





Sound Studio is an easy-to-use Mac OS X application for recording and editing audio digitally on your computer. Digitize vinyl records and tapes, record live performances, create your own mixes with cross-fades, tweak the levels and EQ, apply digital effects, and save in several file formats with Sound Studio.

System Requirements

-  Operating System: Mac OS X v.10.5.8 or later (10.6 recommended)
-  CPU: Mac computer with an Intel processor
-  Audio Hardware: Built-in audio or other Core-Audio-compatible sound hardware

About This User's Manual

The four main sections of this manual are:

-  **Introduction** introduces you to Sound Studio, states the software system requirements, and provides an overview of this user's manual.
-  **Getting Started** tells you what you need to know to run Sound Studio.
-  **How-To Tutorials** takes you step-by-step through common tasks in Sound Studio.
-  **Reference for Sound Studio** is an item-by-item description of each menu and window item in Sound Studio. It contains more detailed and complete information than the other parts of this document.

Conventions Used In This Manual

In this manual, you will see some words using special typefaces:



Words like this denote menu items (“File ▶ Open”) and file names.

“App” and “application” refer to the Sound Studio application software.

“Operating system” refers to the Mac OS system software.

Getting Started

Installing Sound Studio for the First Time

-  **Installing from the Mac App Store.** If you purchase Sound Studio from the Mac App Store, Sound Studio will automatically be downloaded to your computer.
-  **Installing from the Felt Tip website.** Sound Studio can be downloaded from the Felt Tip website at: <http://www.felttip.com/ss/>. To purchase a Sound Studio license from the Felt Tip store, go to: <https://sites.fastspring.com/felttip/instant/soundstudio4>. When you open the application, you will be able to enter the code to register your copy of Sound Studio.

Audio Hardware

In order to record and play back audio, Sound Studio needs audio hardware. Most computers have built-in audio hardware for both recording and playback, but some only have playback capabilities, and some have no built-in capabilities. Please check the manual that came with your computer.

You may also need other audio equipment such as a tape deck, phonograph, mixer board, or a microphone. You will also need the proper cables to connect the equipment together.

Typically, your computer will have a pair of 1/8-inch mini-jacks, one for audio input (recording), and one for audio output (playback). These are usually located near or on the back of the computer or on the side of the computer. Some computers also have a built-in microphone.

In order to use this equipment, you will have to set up the Sound Studio preferences to tell it what hardware to use. To do this, launch Sound Studio, open the Sound Studio menu, and select Preferences. Then select the appropriate items in the Audio section of the Preferences for Input and Output.

How-To Tutorials

This chapter shows you step-by-step how to do common tasks in Sound Studio.

Sound Studio is designed to let you experiment with the various editing commands and filters. You can undo your actions multiple times by using **Edit ▶ Undo**, and you can revert to your last saved version in case something really goes wrong.

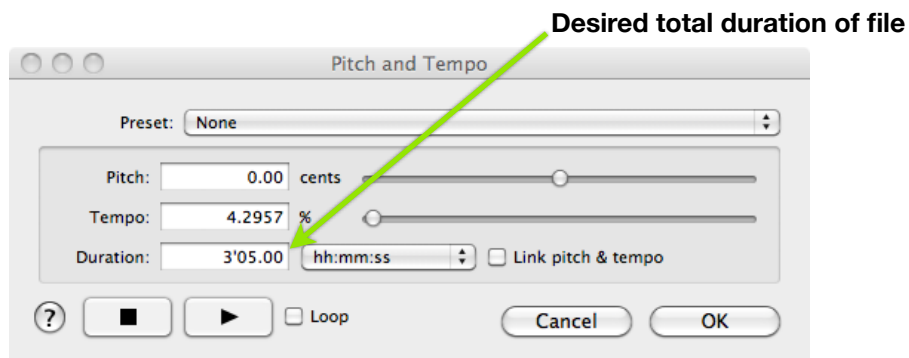
It is a good idea to keep a copy of your original files, in case you want to start over, and to save your work frequently.

How to Open a CD Audio Track

To open a track on an audio CD, select **File ▶ Open**, navigate to your Audio CD, and open your CD track directly. This file can then be edited and saved to the hard disk drive. All CD audio tracks will open up as 44.1kHz, 16-bit, stereo files.

How to Make a File Fit Within a Specific Time

To change the duration of an entire audio file, select **Filter ▶ Pitch and Tempo** and enter the desired total duration of the file. The duration is in the format “hours”minutes’sseconds.fraction” so **3’05.00** is three minutes and five seconds. (The symbol represents the prime mark for minutes and double prime mark for hours, but is entered with the apostrophe and double quote keys.) The pitch of the audio will stay the same.



How to Record an Internet Audio Stream (or From Another App)

There are several applications or situations where your computer is generating sound but there isn't a means to record what is playing, such as Apple Speech synthesis, MIDI data played using QuickTime musical instruments, and streaming audio from RealPlayer or Apple's iTunes.

The simplest way to capture this audio in Sound Studio is to connect the Sound Output port to the Sound Input port of your computer. While this will result in some degradation of the signal due to the digital-to-analog and analog-to-digital conversion, the amount of noise is insignificant for most purposes.

You can also record from the system sound output such as the audio from iTunes or Skype into Sound Studio if you also install a piece of software that lets you re-route the audio output from the other app into something that Sound Studio can see in its Preferences. You can use [Lineln](#) or [Soundflower](#) to do that.

You can also use apps such as [Audio Hijack Pro](#) or [Nicecast](#) (both from Rogue Amoeba) to record directly from the other app into a file.

How to Cross-Fade Two Files

You can cross-fade from one file to another in Sound Studio. First, open two files. You will want to make sure they both have the same sample rate, sample size, and number of channels, but it's not necessary. This example uses a cross-fade of five seconds, but you can use a different duration.



Apply the **Fade Out** filter to the last five seconds of the first file. Apply the **Fade In** filter on the first five seconds of the second file.

Now, select all of the second file and select **Edit ▶ Copy**. This will put it on the clipboard.

In the first file, click once at a point five seconds before its end, and use **Edit ▶ Mix Paste** to mix the clipboard audio with the first file. The audio on the clipboard will be mixed on top of the first file starting from the insertion point, and the file will be extended to hold the new mix.

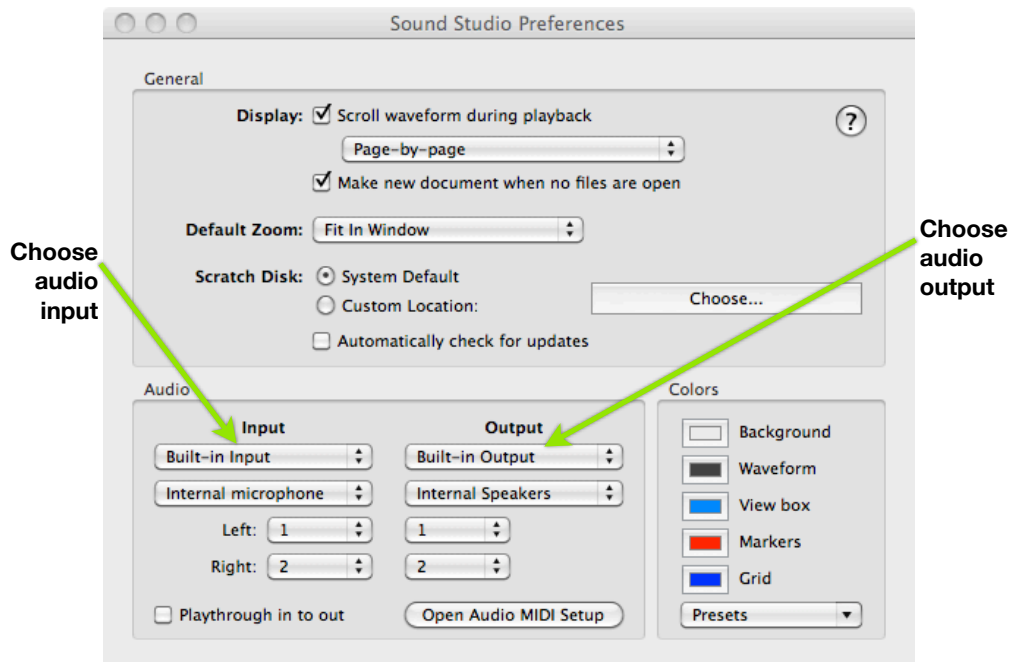
By repeating this cross-fade on several songs, you can create a mix with seamless transitions between songs. For more advanced DJs, you can use the **Pitch and Tempo** filter to match the beats of different songs.

How to Turn a Record or Tape into a CD

To make a digital audio CD out of an analog source, you will need several things:

- 🔦 **Sound input device.** If your computer does not have built-in audio hardware, you will need a USB or FireWire audio capture device.
- 🔦 **CD burner (CD-R or CD-RW).** Most computers come with a CD burner built in, but check that you have one.
- 🔦 **For phonographs:** If you have a turntable, you will also need an amplifier with phonograph inputs to bring the signal up to line levels and to apply the proper RIAA equalization curves.
- 🔦 **The proper audio cables.** This will depend on your set up, but usually you'd use an RCA to 1/8-inch stereo cable to connect your amplifier to your computer.

Make sure you have connected your tape deck or amplifier to the computer with the proper cables. Launch Sound Studio and go to **Sound Studio > Preferences** and make sure the appropriate audio input is selected. This may be Built-In Input or the name of your USB audio device. Also check that the audio output is properly selected.



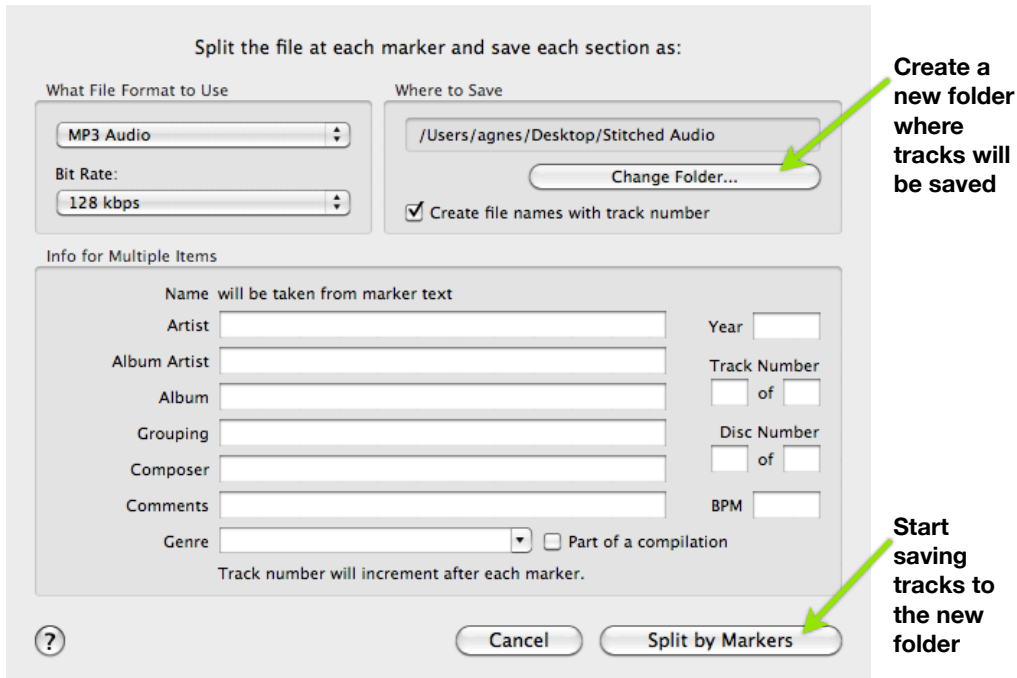
Turn on the playthrough option, and close the preferences window. Start playing the tape or record and see if the level meters start working and sound starts coming through the computer. You want to adjust the levels so that it goes into the yellow occasionally, but not into the red. (See the Input Levels section for the advanced discussion about levels.) If you hear no sound, check your connections.

Once you've set up the levels, you can start recording. **Select File > New** to create a 44.1kHz, 16-bit, stereo file. Press the record button in Sound Studio, and press play on your tape or start your record playing.


Wait until the tape or record is finished, then press the stop button in Sound Studio. You should see the window fill up with the waveform of the audio you just recorded. You can save the file to your hard disk.

You will want to turn each song or section of the recording into a separate file so that your CD burning software can make them into separate tracks. The fastest way to do this is to insert markers at the beginning of each track. First, zoom all the way out by selecting **View ▶ Fit All in Window**. To insert a marker, click once where the first song starts, and select **Insert ▶ Marker**. Then find the beginning of the next track and insert a marker there.

A shortcut for dropping markers is to press the “M” key after clicking where you want the marker to be. Repeat inserting markers for the start of each track. Optionally, you can name each marker after the track that comes after it, by double-clicking in the name of the marker (usually “Marker 1”) and typing a new name.



Then you can split the one long file into several shorter files by using **Edit ▶ Split by Markers**. In the dialog, click on the **Change Folder** button and create a new folder where you want to save the tracks. Click on the **Open** button to choose the newly created folder. Then click the **Split by Markers** button to start saving the tracks to new files in that folder.



Once all the tracks are saved, you can drag the folder full of tracks into iTunes or any other CD burning software, and use create a CD in that software. Please see the instructions from that software if you need help with creating CDs.

Instead of creating CDs, you can save in AAC or MP3 format by going to **File ▶ Save**, however the MP3 file format is available only if the LAME encoder is installed.

Sound Studio does not include an MP3 encoder due to licensing issues. You will need to download and install the [kJams LAME.framework](#), which is 32-bit only as of December 15, 2010. Unless you have a 64-bit version of LAME.framework, you'll need to run Sound Studio in 32-bit mode.

To run Sound Studio in 32-bit mode, right-click on the Sound Studio icon in the the Finder window, select “Open in 32-bit mode,” and relaunch the app. The MP3 option should appear when you save.



Run Sound Studio in
32-bit mode

Reference for Sound Studio

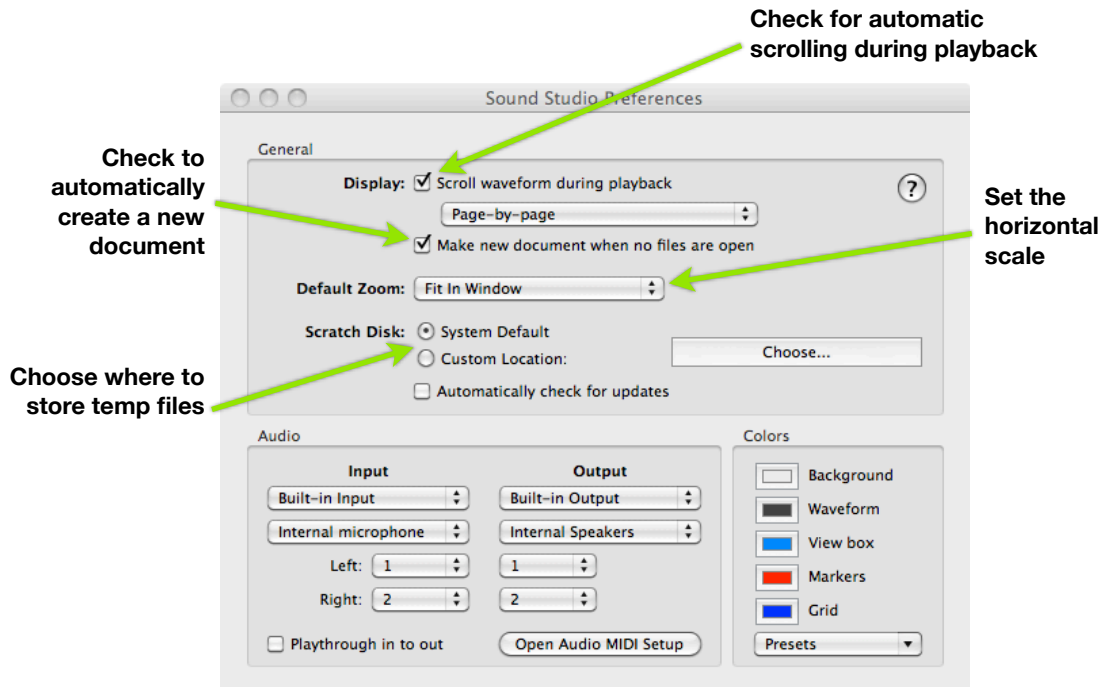
This chapter provides details on specific features and commands, and is organized according to the order of the main menu bar in Sound Studio.

File Formats





Sound Studio supports the following file formats:

<i>File Format</i>	<i>File Extensions</i>	<i>Notes</i>
AAC (m4a) Audio	.m4a, .m4b, .mp4	MPEG-4 audio file with AAC encoding. Compatible with iTunes. The preferred lossy encoding format for iTunes.
ADTS AAC Audio	.aac	Audio data transport stream with AAC data.
AIFF Audio	.aif, .aiff	Audio Interchange File Format. Commonly used for uncompressed audio on Macs. Compatible with iTunes.
AIFF-C Audio	.aif, .aifc	Audio Interchange File Format with Compression. Supports some data compression formats including IMA 4:1.
Apple Lossless Audio	.m4a	MPEG-4 audio file with Apple Lossless audio encoding. Compatible with iTunes.
CAF Audio	.caf	Core Audio Format, preferred format for saving audio data in 32-bit floating-point format. Also allows for saving in IMA 4:1 format inside of a CAF file.
FLAC Audio	.flac	Free Lossless Audio Codec, format suitable for archiving audio collections, ensuring that an exact copy of the original data can be recovered at any time.
NeXT/Sun Audio	.au	An older standard for audio files.
MP3 Audio	.mp3, .mpg3, .mpg, .mpeg, .mp2	MPEG-1 Layer 3 audio file. Compatible with iTunes. The ubiquitous file format, with lossy encoding, and preferred for sending audio over the Internet. Requires the “LAME” framework for saving.
Ogg Vorbis Audio	.ogg	Ogg Vorbis is a free standard for lossy encoding of audio data.
Sound Designer II Audio	.sd2, .L, .R	An older standard for audio files. Uses resource forks, therefore should not be used on non-Mac systems.
Tab-delimited Text	.txt	Sound Studio can only export to text. This is an advanced option which generates very large files for use in scientific data analysis applications.
WAVE Audio	.wav	A file format often seen on Windows PCs. Also supports Broadcast Wave File (BWF) chunk metadata.

Preferences



General

-  **Scroll waveform during playback:** When this is checked, the window will automatically scroll to keep the green playback marker visible when it reaches the right edge. You can still scroll manually to other parts of the waveform. Turn this option off to prevent the window from scrolling automatically.
-  **Make new document when no files are open:** This feature automatically creates a new document when there are no other files open. Click on the application icon in the dock or double-click the application icon in the Finder.
-  **Default Zoom:** This sets the horizontal scale of newly opened and created documents.
-  **Scratch Disk:** This sets the location in your file system where Sound Studio stores its temp files, which contain audio data that is being recorded and processed. Ideally, this would be a large and fast hard disk. Since audio is stored as 32-bit float at the file's sample rate, this data can be very big. One hour of stereo, 44.1 kHz audio takes up about 1.2 GB of hard disk space. Since every change is stored and kept on disk to support unlimited undo, the scratch disk will fill up over time as changes are made and filters are applied, until the file is closed or Sound Studio is quit. By default, the scratch disk is your system disk, and the temp files are in a directory randomly chosen by the OS.

Audio

By default, Sound Studio uses the audio input and output hardware selections from the System Preferences. You can specify that Sound Studio use a specific hardware device for input and another for output. You can also specify which channels to use for each device.

Most computers have built-in audio hardware, and this would show up as “Built-In Audio.” If you have installed or plugged in some other audio hardware, such as a USB or FireWire audio

device, you can select it in the preferences. The selection here only affects Sound Studio, so you can have Sound Studio play audio to one device while other applications play audio on another device.

The input column controls the audio recording or capture hardware. The output column controls the audio playback hardware. The first two lines select the device and source. The source usually refers to the port to use on that device. On some devices, there is only one port, and the source pop-up button may be disabled to show that you can't change it.

The left and right channel selections determine which channels in the device map onto Sound Studio's left and right channels. Sound Studio can record up to two channels at one time from the input device. On playback, it mixes down each track in a file based on its pan settings to create a stereo mix that is sent to the selected channels on the output device.



Playthrough in to out: Turning this on will preview the audio on the inputs by immediately playing input audio on the audio outputs. Be careful about feedback: if the speakers and microphones are too close or pointed at each other, you may get feedback. Playthrough is useful for monitoring your recording through the computer, to make sure you have the levels and connections set up properly.



Open Audio MIDI Setup: This opens the Apple-supplied application "Audio MIDI Setup" which is used to configure audio hardware connected to your computer. It can be used to set the native sample rate and format of the hardware, which is important because you may want to record at a higher than normal sample rate. Sound Studio will continue to work when there is a mismatch between the hardware sample rate and the file sample rate, by resampling the audio.

Colors

The colors used to draw the main waveform views, markers, and grid can be changed. They can also be reset to their factory defaults.

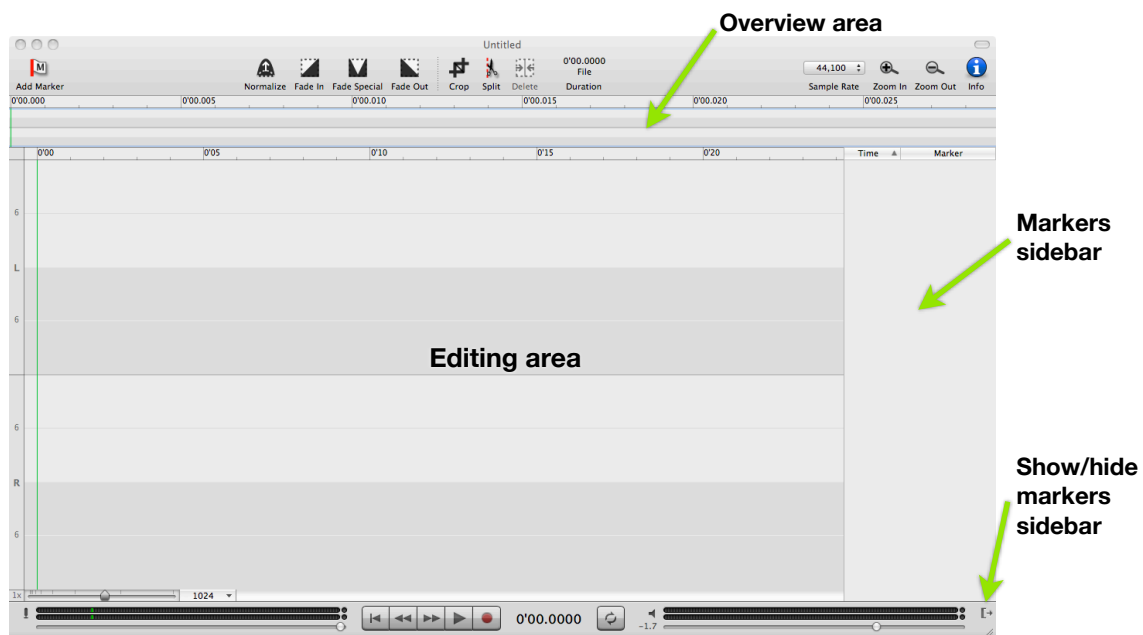
Windows

This section describes the way the document window and input level palette window work. It also contains a table of how modifier keys affect selection in the document window's waveform views.

Information on dialog boxes is found with their corresponding commands. Information on saving presets is found in the “Filter Menu” section.

Document Window

The document window is the window you see when you create a new file or open an existing one. It is the main working area where you will record, edit, and play back your audio.



At the bottom of this window is the input level meter and input gain slider control; the transport controls (return to zero, rewind, fast forward, play, and record); the time readout; the loop playback control; and the output level meter and output gain slider control.

The output level meter shows the actual levels of the stereo mix of the audio data, and is not affected by the output gain slider. It shows the levels of the audio before the output gain amplifier.

To turn on and off the input level meter, click the microphone icon to the left of the input level meter.

The box just below the buttons is the “Overview” waveform view. It always shows the waveform of the entire file and is *not* the editing area. The editing area is the large waveform view below this one. The overview only has a box that is drawn around the part of the file that is shown in the editing area. This box is in a different color from the waveform. Clicking and dragging in the overview will move the box. Command-click-dragging here will let you draw a new box, and change the area that is shown in the editing area.

The editing area is below the overview and stretches almost to the bottom of the window. This is where you make your selection, edit the audio, put down markers, and do other editing tasks. There is an amplitude ruler to the left, marked in decibels, and clicking in this area will pop up a menu for selecting the vertical scale or zoom, which allows low level audio to be more easily visualized. Next to that is a set of controls for the horizontal zoom level, with 1:1 being zoomed all the way in, and higher numbers meaning you have zoomed further out.

Markers Sidebar. The icon above the window resize widget shows and hides the Markers sidebar, which contains a list of all the markers in the current file. You can add, delete, and edit markers in this sidebar.

Keyboard Shortcuts

Here is a list of some keyboard shortcuts that are not shown in the menu system:

General		View		Filters	
Space-bar	Play/Stop	Q	Center on selection start	A	Amplify/Volume
0 (zero)		W	Center on selection end	D	Fade In
Escape	Stop	E	Previous zoom	F	Fade Special
Return	Return to Zero	R	Zoom out	G	Fade Out
· (period)		T	Zoom in	X	Dynamics Expander
, (comma)	Pause	home	Scroll to start	Alt + X	Noise Gate Expander
1	Rewind	end	Scroll to end	C	Dynamics Compressor
2	Fast Forward	page up	Scroll left one page	Edit	
3	Record	page down	Scroll right one page	enter	Insert Marker
4	Toggle loop-playback			M	
				Alt + arrow	Jump by marker
				9	Paste mixing

Here is a list of modifier keys and their effect when clicking the mouse in the overview and editing areas:

Modifier Key	Effect on Editing Area	Effect on Overview
None	Normal selection	Moves view box
Shift	Modifies selection edge closest to pointer	Same as above
Option	Hand Tool: Click-drag to scroll	Same as above
Command	Normal selection	Click-drag to draw a new view box
Control	Start playback at pointer	Start playback at pointer

The audio controls are the space bar or zero key to start and stop audio playback, the escape key for stop, and the return or period key to jump to the beginning of the file. During playback, press the **1** key to rewind and the **2** key to fast forward. To start recording, press the **3** key. To toggle looping playback, press the **4** key.

To move the cursor, use the left and right arrow keys.

To jump to the beginning of the file, press the up arrow key, or the return key.

To jump to the end of the file, press the down arrow key.

To move to the previous marker, press option-left-arrow, and **to move to the next marker**, press option-right-arrow.

To create or extend a selection, hold down the shift key in combination with the above arrow key shortcuts.

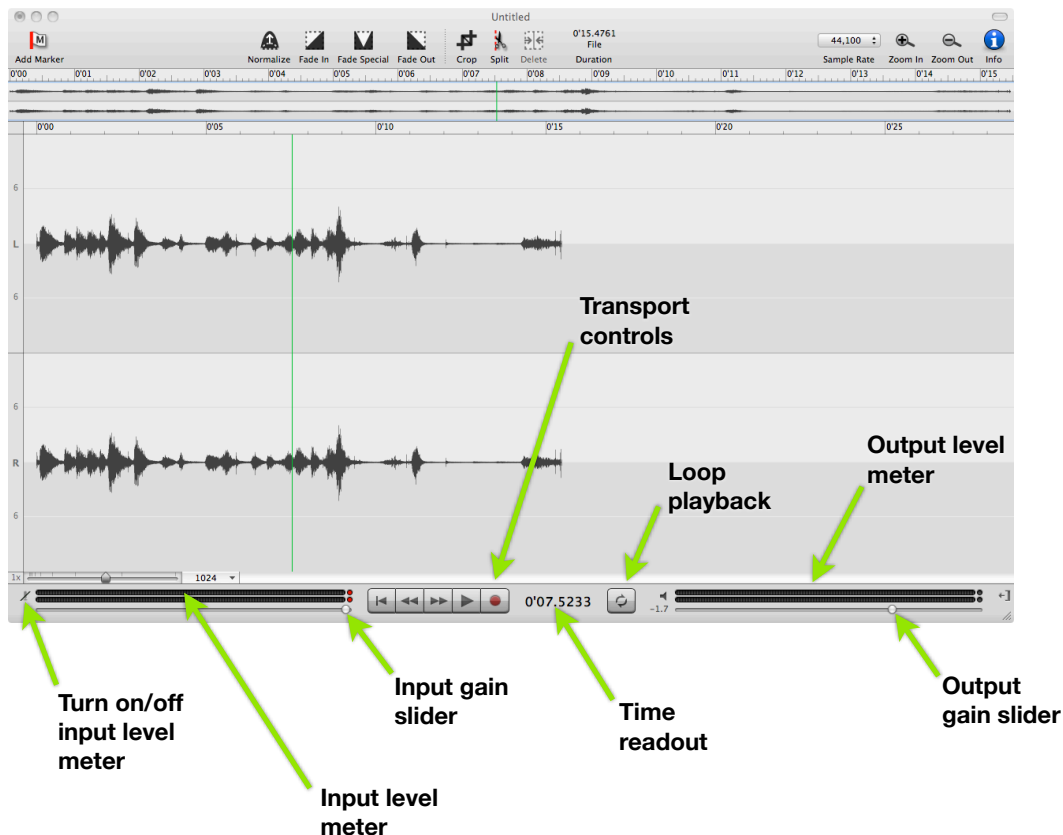
To quickly change the waveform view, press the **Q** key to show the beginning or left edge of the current selection, and press the **W** key to show the end or right edge. Press the **R** key to zoom out, the **T** key to zoom in, and the **E** key to jump back to the previous zoom level.

For editing, use the delete or backspace key to remove the selected audio. Press the **Enter** key or the **M** key to drop a marker at the cursor or playback head, and use the **9** key to bring up the Paste Mixing dialog.

To apply an effects filter, press the **A** key for Amplify, **D** for Fade In, **F** for Fade Special, **G** for Fade Out, **X** for Dynamics Expander, **alt-X** for Noise Gate, and **C** for Compressor.


Scroll wheel, trackpads, and the Magic Mouse. Some mice come with a scroll wheel. Turning this wheel in the Y-axis will cause the window to zoom in or out, with the zoom focused at the mouse pointer. Turning the scroll wheel in the X-axis will scroll the window normally. Some mice do not support turning the wheel in the X-axis. Trackpads and other devices like the Magic Mouse use finger gestures in place of a scroll wheel and can also be used to zoom in and out or to scroll normally.

Single-channel selection. For stereo files, clicking in the top half of the left channel selects the left channel, and clicking in the bottom half of the right channel selects the right channel. Clicking in the middle of the view selects both channels. When you have made a single channel selection, subsequent selections will generally be kept in the same channel.







Transport (Return to Zero, Rewind, Fast Forward, Play, Record)

The playback and recording transport controls are in the bottom middle of the window. They allow you to manipulate playback and recording. Playback and recording are mutually exclusive on a single file, which means you can play or record the file, not both at the same time. However, you can play one file while recording to another file.

 **Playback** starts when you press the play button. Playback will start from the insertion point, or the beginning of the current selection. Pressing the play button again will stop playback and return to the insertion point. If you have **Audio > Loop** turned on, the selection or the entire file will be repeated, which means when playback reaches the end of the audio to be played, it will automatically go back to the beginning of the selection or file and replay the selection.

You can use the right mouse button, or right-click gesture, to start playback from any point in the file. Click the right mouse button, or use the right-click gesture, anywhere in the main view or in the overview, and playback will start from that point. If you have looping enabled, the looped area will be the current selection if you clicked inside the selection, or the entire file if you clicked outside the selection or if there is no selection.

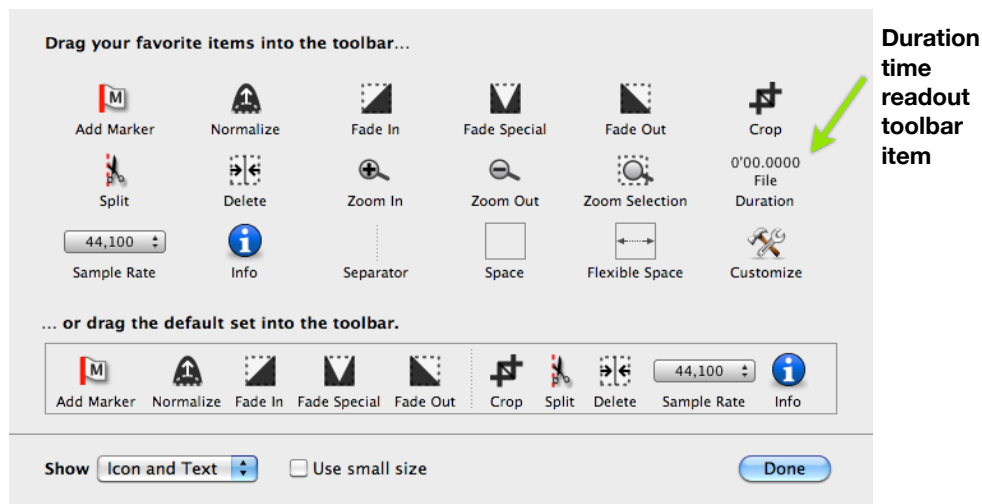
 **Recording** starts and stops when you press the record button. Any recording options you have set up in **Audio > Auto Start/Stop Recording** will apply to this recording. If you have it set to start recording at a certain peak level, the recording will enter the paused mode and wait for the peak level to rise above the threshold. Otherwise, it will immediately start recording. When you press the recording button to stop recording, the recorded audio will be inserted in the file at the insertion point and will replace any selection, the same as pasting audio into a file.

-  **Return to Zero** starts playback from the beginning of the file. The Return to Zero button is used only during playback.
-  **Rewind** moves the cue one second back during playback. Hold down the Rewind button to continue rewinding through the audio file.
-  **Fast Forward** moves the cue one second forward during playback. Hold down the Fast Forward button to continue fast-forwarding through the audio file.

Time Readouts (cue and duration)

In the bottom middle of the window is a time readout for the current cue position. The cue represents the point where the playback and recording head would be, and updates while playing or recording.

To show the duration of the entire file or a selection of the file, you can add the Duration time readout tool to the toolbar at the top of the window by going to **View ▶ Customize Toolbar**. If you select a portion of the file, the Duration time readout tool shows the length of that selection. When there is no selection, the Duration time readout tool shows the length of the entire file.



Rulers

The default ruler along the top of the each view has markings in minutes, seconds and fractions of a second. Control-click or right-click the ruler to select different units. Alternately, go to **View ▶ Ruler Units** to choose a different unit of measurement. See the commands reference section on “Ruler Units” for more info.

The amplitude ruler is along the left side of the window, and it has markings in decibels. At 1x scale, the top and bottom limits of the waveform would normally be 0 dB, and are not labeled. The center between the limits is $-\infty$ dB, which corresponds to zero amplitude. There is an amplitude zoom popup directly under the amplitude ruler, and this allows you to zoom the waveform vertically in the window. This zooming does not affect how the audio sounds when played back; it only lets you see more detail in low-amplitude signals. The amplitude zoom ranges from 1x to 64x.

Playback Cue (the Green Line)

While playing back the audio, this green line shows you where the playback head currently is. Control-clicking or using the right mouse button on the main view will let you reposition the playback head. See the “Transport” section above for more info on playback.

Recording Cue (the Red Line)

When recording, this red line will show you how the recording is progressing. See the “Transport” section above for more info on recording.

Sample Rate Display (“44,100”)

In the toolbar at the top of the window is a button that shows the file’s sample rate. Clicking on this button pops up a menu of sample rates. For new, empty files, it just changes the sample rate. Any existing audio will be resampled to the new sample rate.

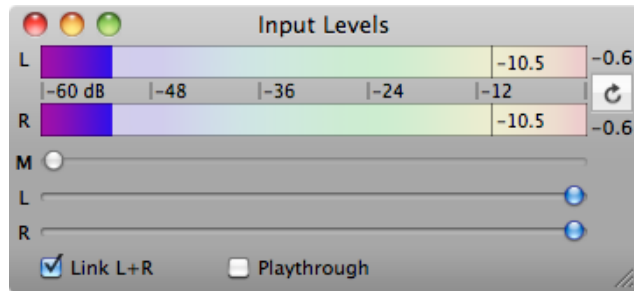
Zoom Controls

The zoom controls change the horizontal zoom, which is how many samples each pixel represents horizontally. A ratio of 1:1 means that you are seeing every sample; a ratio of 1:100 means that each pixel (dot) represents 100 samples in the audio file. The zoom can be set with the slider, the pop-up menu, the scroll wheel, or the menu commands.

Clicking on the zoom readout button reveals the pop-up menu where you can select a pre-defined zoom ratio.

Input Levels Window

The Input Levels is a normal window that shows the current levels at the sound input, and lets you control the input gain (amplification). Not all audio hardware supports input gain. You can show and hide the Input Levels by using the Window menu. The lower left-hand corner of document window also shows the input and output level meters.



Level Meter

The level meter, at the top of the palette window, shows the current levels in the left and right sound input channels. When the bar is entirely light colored, there's no signal at the inputs. The scale is in decibels. The two readouts to the right of the meters show the last peak value, with the word "Clip" displayed if clipping has occurred. Clicking on the readouts will reset them. For best results, try to keep the peaks below -6 dB, and use the Normalize and Dynamics Compressor to bring up the levels after recording.

Input Gain Control

The input gain control is directly under the level meter. There are three sliders labeled M, L, and R. M is the master gain, which controls the gain or amplification across all channels of the device. L is the left channel and R is the right. The functioning of these sliders depends on the specific audio hardware being controlled.

When any or all of these sliders are grayed out and unavailable, it means that the hardware does not allow Sound Studio to change the input gain. Some audio hardware will not let Sound Studio control the gain of one or more of the M, L, or R sliders. If all of the sliders are grayed out, it may be necessary to use an external means of controlling the audio levels.

The left and right controls are normally linked. However, if you uncheck the "Link L+R" checkbox, the left and right controls will move independently.

Playthrough Checkbox

This controls software playthrough, which lets you hear on the output device what's coming in on the input device, even if the input and output are different devices. This checkbox does the same thing as the one in the preferences window. You want to be sure to keep your microphone away from your speakers when playthrough is on, or you'll get feedback noise.

To enable hardware playthrough, check your audio hardware's manual. In some cases, you may be able to enable it by using the Audio MIDI Setup application. To do this, go to Sound Studio preferences, and click on the "Open Audio MIDI Setup" button. Then select the audio hardware, and enable the "Thru" checkbox for the appropriate channels, if this is a supported feature of the hardware. Not all hardware supports hardware playthrough.

Commands

This section describes in detail the menu commands in Sound Studio, and each command will generally have a section on how each command works, a section on how they typically are used, and a section of tips on using the command.

The commands are organized here the same way they are organized in the application, that is, grouped into categories such as File, Edit, Audio, etc.

File Menu

The File menu lets you work with files or documents as a whole.

New

The New command creates a blank audio document with the default sample rate of 44.1 kilohertz, sample size of 16 bits, and two tracks (stereo). The new document is immediately available for recording.

After creating a new document, the sample rate can be changed in the Audio menu with the “Change Sample Rate” command. The file can be converted to mono with the “Mixdown to Mono” command. The sample size can be changed with the “Sample Size” submenu.

The sample rate and the number of channels (whether it is stereo or mono) affect the amount of disk space taken up by audio sample data. A higher sample rate will increase the disk space needed for a given duration of audio data. Choosing mono instead of stereo will halve the disk space needs. To set the default sample rate to 48,000 Hz (or your desired value), use Terminal to type in **defaults write com.feltp.SoundStudio file.sampleRate 48000**, or the desired value instead of 48000.

The sample size setting only affects the file when it is saved. While the audio document is open, Sound Studio stores the audio data as 32-bit floating-point data. Please see the entry for “Sample Size” for more details.

New Mono File

This variation creates a one-channel document, perfect for spoken word recordings.

Open

The Open command allows you to select a file on the disk to be read into Sound Studio. Conceptually, this is moving a file from the file system into the Sound Studio application. For a list of file formats, please see the File Formats section of this manual.

The **Open Recent** submenu shows the last ten files opened in Sound Studio.

Save

The **Save As** and **Save Selection As** commands let you choose from different file formats and options. See the File Formats section for a list of file formats. The Save Selection As command does not change the file type or options of the open document.

Revert

The Revert command returns the current document to the state when it was last saved.

Get Info

Under Get Info, you can edit iTunes metadata and ID3 tag data, such as artist, album, genre, and other music information. You can also edit the cover artwork, lyrics, and many other tags that show up in iTunes. Sound Studio supports reading ID3 tags in version 2.2, 2.3, and 2.4 format, and writes out version 2.2 tags, in MP3, AIFF, AIFF-C, Sound Designer II, and WAV formats. It also supports the iTunes metadata format in MPEG-4 AAC and Apple Lossless (m4a) files.

Artwork that is added must be a file in the JPEG or PNG format. Images can be dragged from files in the Finder into the artwork box, dragged between Sound Studio files, and rearranged within the box.

If you are creating a podcast, make sure to fill in the Artist field with your name and the Album field with your podcast's name.

Use the “Tap for BPM” button to tap out a steady beat in rhythm with the music while it plays. It'll show you the BPM after 4 beats, and will keep getting more accurate with the more beats you put in. Hold off the beats for about 2 seconds to reset it.

Edit Menu

The edit menu lets you edit the current selection, whether it be text or audio sample data. Some commands only apply to audio data. Some commands in this menu are explained on this page, while others have their own pages.

Undo and Redo

The Undo and Redo commands allow you to walk through the history of changes to the current document, up to the point when the file was opened. Before any significant data-modification action, the application saves a snapshot of the document so that you can undo your actions. After executing an undo, the redo command is available to move forward through the history.

The downside to unlimited undo actions is that it can take up a lot of disk space. For example, if you have a 100 MB file and you apply five different filters to the entire file, the undo system will have to save 500 MB of data in addition to the original file. To clear the undo system, close and save the file, and reopen it.

Cut, Copy, and Paste

The Cut, Copy, and Paste commands allow you to move audio data between and within documents. The commands work similar to a text editor. The application keeps track of a single “pasteboard,” which can hold audio data. The Copy command puts the selected audio on the pasteboard, and the Cut command does the same thing as the Copy command except it also removes the selected audio from the document. The Paste command inserts the audio on the pasteboard into the current document. It does not overwrite the existing audio unless you have already selected some audio in the document you are pasting into. In that case, it will replace the selected audio with the pasteboard audio.

There are a few things to note about the Paste command. First is that you can copy and paste multiple tracks of audio, but you have to be careful about where you put the insertion point or selection when you paste, or you may get unexpected results. When you select the Paste command, the first thing it will do is to delete any selected audio. Then it will insert the pasteboard audio where the deleted audio was. If the pasteboard audio contains only one track, then the audio will simply be pasted into the selected track. (The selected track is the first track where the insertion point or selection is located.) If the pasteboard audio contains more than one track, it will keep inserting audio for each track, adding tracks to the document if necessary.

The Paste command only affects the minimum number of tracks needed to insert the audio, so, for example, if you paste a single-track (monophonic) of audio into track one of a two-track (stereophonic) document, only the audio in track one will be moved to accommodate the new audio, thus causing the stereo audio after the insert point to be out of sync.

When moving audio data between two files with different sample rates, the audio will be resampled so that the pitch and duration of the audio will remain the same. When there is a sample rate mismatch, this resampling can take a few seconds or more for large amounts of data.

To copy marker times and names to a text editor, just use Copy on the sound data that has markers, and paste into the text editor. The marker times and names will be pasted as a tab-delimited list.

The audio pasteboard is private to the application and not available to other applications. It cannot be exported nor can you copy and paste audio data into another application. To move data to other applications, you can save the audio file in the appropriate file format and open it in the other application.

Mix Paste

Displays a dialog that will let you mix the audio on the clipboard (from a Copy or Cut) with the audio in the current file. It mixes the audio starting from the left edge of the selection, and its length is determined by the clipboard audio. It will extend the file if necessary, if the option to paste all of the pasteboard audio is selected.

Delete

The Delete command removes the selected audio from the current document. The audio that comes after the edit will be shifted to fill in the gap made by the deleted section.

Select All

The Select All command selects all of the audio of every track of the document.

Silence

The Silence command clears the selected audio by replacing it with zero-volume audio data, making it silent, while keeping the timing of the audio around it the same.

Crop

This command trims the file down to the selected audio. It deletes the audio before and after the selection, and deletes any tracks not within the selection.

Split By Markers

This command divides the file into sections based on the markers in the file. It then saves each section as a new file to a folder which you can specify by clicking the “Change Folder” button. Each file will be named for the marker at the start of the section.

You can select the file format to save in. If you are digitizing vinyl records or tapes, you can record an entire side in one session, place markers at the beginning of each song, and use this command to save each song separately. You will want to select a compressed format such as MP3 or AAC (m4a) to save hard disk space. After splitting the file, the new files can be added to your iTunes library, where you can listen to them and burn them to an audio CD.

Create file names with track number: This option will prefix the track number to the file names. For example:

Marker Name	File without Prefix	File with Prefix
Cat	Cat	1 Cat
Dog	Dog	2 Dog
1-Fish	1-Fish	3 1-Fish
2-Fish	2-Fish	4 2-Fish

The **Info for Multiple Items** section lets you set the iTunes metadata or ID3 tags for the files to be saved. Note that the Track Number is for the first file saved, and will be automatically incremented for each new file.

Go To

The Go To command scrolls the main (lower) view to show a specified time, the selection start, or the selection end. If you choose a specific time, it will also deselect and move the insertion pointer to that time.

Set Selection

The Set Selection command allows you to numerically enter the start and end points of the selection. It will also zoom the main (lower) view to show the selection.

You can also select markers in the above two commands.

Audio Menu

The audio menu has audio-specific commands.

Loops while Playing (key shortcut: 4)

Checking this option will keep playback going in an endless loop. If there is selected audio, only the selection will play in a loop. Otherwise, the entire file will repeat. You can turn it on and off while playing.

Play (key shortcut: space)

This command starts playback from the start of the selection or the insertion point, and plays to the end of the selection or end of the file. To start playback using the mouse, click the right mouse button or control-click (by holding down the control key while clicking with the mouse button) in either the upper or lower waveform views, and playback will start where you clicked.

During playback, use these keys to control playback:

1 Rewind

2 Fast Forward

4 Loop

M Drop Marker

return Jump to 0:00 time

These keyboard shortcuts and more are listed in the File Formats section of this manual.

Stop (key shortcut: escape)

This command stops playback or recording.

Record (key shortcut: 3)

This command starts recording to a temporary storage file on disk. When you press the stop button, the recorded audio will be inserted into the file at the cursor, replacing any selected audio.

Auto Start/Stop Recording

This allows you to set up certain timers or conditions under which recording will start and stop. You would typically create a new file, and then select this command to set up the recording options. These options only apply to the current file, and are not saved with the file.

You can select a preset and then fine-tune the options for your particular setup and for the material you are recording. There are presets for timer recording and voice-activated recording.

Under “Start Recording,” when “Manually” is selected, recording always starts as soon as you press the record button in the current file's window. The manual option is the default setting when you create a new file. The “start recording at peak level” option tells the recording to wait until the peak audio level, which you can see in the “Input Levels” palette window, rises above a certain decibel level, which you can set in the box next to this option. This option can be used when recording from a tape or vinyl record, so that you can tell Sound Studio to wait until the music starts before recording. You can also use it when the computer and your audio equipment are far away from each other, and you don't want to record several seconds of silence while you walk over to your audio equipment to start playback.

Under “Pause Recording,” you can turn on the option to automatically pause recording if the levels drop below a certain decibel threshold for a certain length of time. The recording will be automatically restarted when the levels rise above the threshold. This is useful if you are recording someone speaking and there are sections of silence you want to cut out during recording.

Under “Stop Recording,” you can set up the condition under which recording will stop. The manual setting will let recording continue until the disk runs out of free space, which is not a good idea because the OS becomes unstable without disk space.

The total time option lets you specify how long the file should be, and recording will stop when the duration of the entire file is equal to the time you enter in the box next to this option. You can enter the time in hours, minutes, seconds, or samples. Sound Studio looks at the length of the entire file, and not just the current recording session, so if you set the limit to 30 minutes and the file is already over 30 minutes long, it won't record anything which would add to the file's length.

The “At the time” option lets you specify an exact time of day when recording should stop. This is useful if you are going to record or have already started recording a radio or television program that stops at an exact time.

The “Stop recording at peak level” option tells Sound Studio to stop recording when the peak levels drop below a certain decibel level for a certain length of time. You can monitor the peak levels in the “Input Levels” palette window. This option is useful in many situations. If you are recording from a tape or record, when the audio ends, you will only hear silence, so you can have it stop recording when it detects a long period of silence. When used in conjunction with the “start recording at peak level” option, you can have it start and stop recording when it hears audio above a certain level. This way, you can avoid recording a lot of silence if you go away from the computer and are not around to stop the recording manually.

Change Sample Rate

This command allows you to change the sample rate of the document, to change the total duration of the document, and to resample the audio data. When the “Resample” checkbox is not checked, the two fields are linked together since without resampling, the sample rate and the total duration are interrelated. Changing one will change the other. Note that changing the total duration will also change the pitch of the audio, regardless of the “Resample” setting.

The most common sample rate is 44.1 kHz, but you can use any sample rate in Sound Studio. A higher sample rate such as 48 kHz will give better definition, especially in the higher frequencies, but also consumes a larger amount of disk space per second of audio. At 88.2 kHz, a recording generates twice as much data as 44.1 kHz, while going to 22.05 kHz will generate half the data.

Sample Size

The Sample Size submenu allows you to change the sample size of the document as it is saved to disk. It does not have any effect on the file while it is open in Sound Studio, because audio data is stored as 32-bit floating-point data within the application. Only when the file is saved to disk is the sample size used.

The sample size will affect the sound quality and file size of the file on disk. A smaller sample size (8-bit being the smallest) will result in a smaller file size, but also lower sound quality. The 16-bit setting will produce a file twice the size of an 8-bit file, but with significantly better sound quality, while the 24-bit setting will produce a file three times the size of an 8-bit file. The 16-bit settings is usually the best compromise between file size and sound quality.

Mixdown to Stereo/Mono

These two commands allow you to convert the file between mono and stereo formats.

Fourier Spectrum Analysis

This command charts the frequency spectrum of the selected audio.

Insert Menu

Marker

This command will place a marker at the insertion point, or if there is a selection, at the beginning and end of the selection.

Silence

This command will insert a specified length of silence in the selected tracks. If any audio is selected, it will be replaced by the silence.

Noise

This command will insert a specified length of noise in the selected tracks. If any audio is selected, it will be replaced by the noise. You can control the volume of the noise, in decibels.

Tone

This command will insert a specified length of a steady audio tone in the selected tracks. If any audio is selected, it will be replaced by the tone. You can control the frequency, the wave shape, and the volume in decibels of the tone.

FM Synthesis

This command will insert a specified length of a variable-frequency audio tone in the selected tracks. If any audio is selected, it will be replaced by the tone. You can control the frequency and the volume in decibels of the tone. The carrier frequency is the base frequency, the deviation frequency is how many hertz the tone frequency should go above and below the carrier frequency, and the modulation frequency is how fast the tone frequency should go from the lowest frequency to the highest and back to the lowest again.

Add Track

This command will insert an empty track after the last track in the document.

Remove Track

This command will remove the last track in the document. Any audio data in this track is deleted. The document must have at least one track.

Filter Menu

A filter applies an effect on the selected audio. An effect works by taking the audio as input, applying some kind of transformation to it, and putting it back in the file. If nothing is selected, the entire file will be selected before applying the filter.

Some filters will let you save presets. A preset allows you to save the position of the controls of an options window so that you can recall them later. You can save several presets in the default presets folders. Any files in the default presets folder will show up in the “Presets” popup. To create a preset, select “Create New Preset” from the pop-up button, give it a name and save it. Then make changes to the options window. Any changes will be immediately saved to the presets. To make changes without affecting a preset, select “Default” from the presets pop-up.

To rename, delete, or otherwise manage presets, select the “Reveal in Finder” command from the pop-up. This will show the presets folder in the Finder, where you can manage the preset files. Presets are stored in the “~/Library/Application Support/Sound Studio/Presets/” folder, in your home directory.

Last Filter

The first item in the menu shows the name of the last filter used. Selecting it will reapply the filter to the current selection without showing its options window.

Amplify/Volume

This command will raise or lower the volume of the selected audio. The options dialog box shows the amount of the change in both decibels (dB) and percentage factor (%). To adjust the volume to a uniform level, see *Dynamics Compressor* and *Normalize*, below.

Fade In

This command applies a linear fade from zero volume to full volume.

Fade Out

This command applies a linear fade from full volume to zero volume.



Fade Special

This command allows you to draw and apply a fader envelope over the waveform to create custom fade effects. There will always be a start point and an end point, but you can add several points in between. You can drag these points up and down to modify the envelope. To delete a

point, hold down the option key and click on the point. The vertical scale is in linear amplitude. The horizontal scale is in time, as a percent of the selected time. The actual time scale will depend on what is selected.

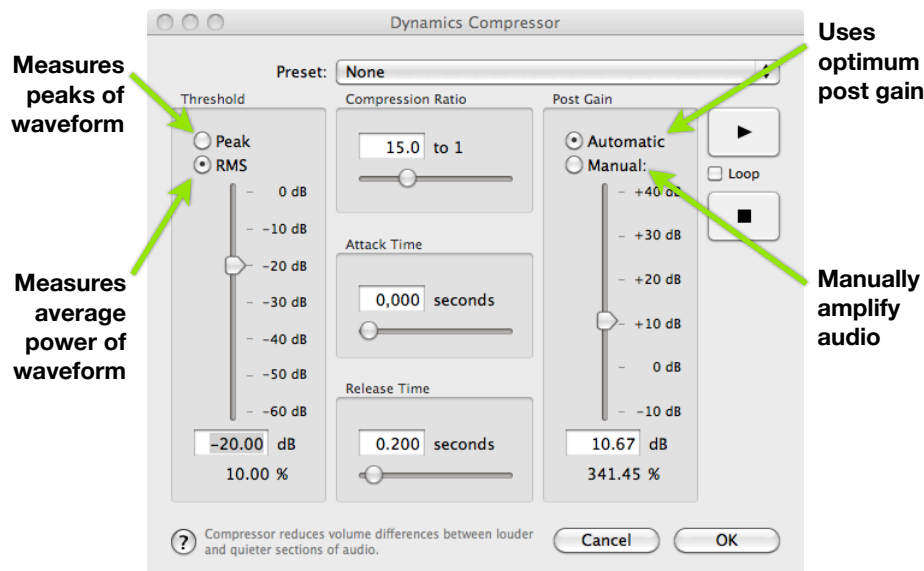
Normalize

This command normalizes the volume of the audio. If loudest peak level is chosen, then this command scans the selected audio for peaks, which are the tops of the waveform, and adjusts the volume so that the peaks coincide with the target level specified in the options dialog box. If loudest RMS level is chosen, this command adjusts the volume based on the average power of the waveform, which closely follows what our ears perceive to be the level of the audio. The target level is shown in both decibels (dB) and percentage factor (%), with 0 dB being the maximum bandwidth of the file, and lower levels being negative decibel numbers.

-  **Normalize each track independently:** This option will treat each track separately when searching for and adjusting the peak levels. When multiple tracks are selected, this option will effectively bring each track to about the same volume relative to each other.
-  **Normalize all tracks together:** This option will treat all track as a unit when searching for and adjusting the peak levels. When multiple tracks are selected, only the highest peak level of all tracks is considered. Each track is adjusted by the same amount so that the volume of each track relative to each other will stay the same.

Dynamics Compressor

A compressor reduces differences in volume between quieter and louder sections of audio. It is applied to audio which goes above the threshold level by turning the volume down by an amount determined by the compression ratio. The attack time is how fast it responds to levels above the threshold, and the release time is how fast it reacts when the levels drop below the threshold. The post gain increases the overall volume to compensate for the lowered volume in the formerly loud sections.



Another way to describe the dynamics compressor is to imagine someone watching the peak level meter while the audio is playing, and whenever the level goes above a certain threshold, he turns down the volume. This is what a compressor does. The compression ratio determines how far to turn down the volume when the level goes above the threshold, with larger ratios resulting in the volume being turned down more. The attack time determines how fast the compressor

reacts to the level going above the threshold. The release time determines how long after the level goes below the threshold the volume is returned to normal.

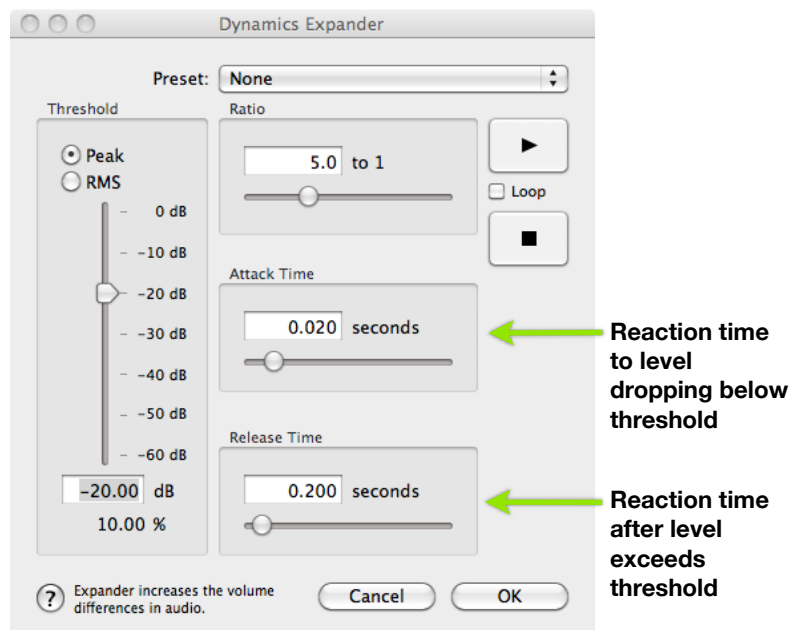
The “Post Gain” control compensates for how the compressor turns down the volume on the loudest passages. It does this by amplifying the audio after it goes through the compressor. Turning on the “Auto” option causes it to calculate the optimum post-gain, so that the volume of the loudest passages stays about the same but the quieter passages are increased in volume. The effect of this is to bring up the level of quieter passages without increasing the volume of louder passages, essentially compressing the dynamic range of the audio.

Compression is useful for taking music with a wide dynamic range, where the quiet passages are very quiet and the loud passages are very loud, and making all the passages about the same volume so that they can all be heard in a noisy environment such as in a moving car. Compression is also used when someone is speaking into a microphone, so that if that someone moves closer to or farther from the microphone, or speaks softer sometimes and louder sometimes, the compression can compensate for these differences and output a consistent level of audio.

The Peak/RMS options under Threshold determine what kind of level meter to use. A peak level meter measures the very tops of the waveform, and is triggered whenever the wave goes above the threshold. A RMS level meter measures the average power of the waveform, and closely follows what our ears perceive to be the level of the audio.

Dynamics Expander

An expander increases the differences in volume between quieter and louder sections of audio. It works by turning down the volume when the volume level stays below the threshold level, and turning the volume back up when the level goes up above the threshold. The attack time is how fast it responds to levels above the threshold, and the release time is how fast it reacts when the levels drop below the threshold.



Another way to describe the dynamics expander is that it is the reverse of a compressor. An expander works by reducing the volume of the audio when the levels drop below a certain threshold. The ratio determines how much to turn the volume down, with a larger ratio resulting

in the volume being turned down more. A high ratio of 12:1 or more is considered a noise gate. The attack and release times behave the same way as in a compressor. The attack time is how fast it reacts to the level dropping below the threshold, and the release time is how long before the volume returns to normal after the level goes above the threshold.

Expanding is useful when you want to increase the dynamic range of the audio. It is also useful if you have a noisy recording and want to reduce the volume of the quieter passages so you don't notice the noise as much. It does have the side effect of changing the way sounds decay and can end up silencing some parts that are quieter.

Noise Gate Expander

The Noise Gate is a version of the Expander which turns the volume all the way down. It generally silences any audio that falls below the threshold level. The attack time is how fast it responds to levels above the threshold, and the release time is how fast it reacts when the levels drop below the threshold.

This is most useful on percussive instruments such as drums and on spoken word recordings, where you have silent gaps between parts with audio signal, and you know you want the gaps to be completely silent.

Add Noise

Adds white noise to the current selection. You would use noise to mask out other sounds in the recording.

DC Offset

This command can be used to shift the waveform vertically in the display. In the manual mode, you would enter the amount to shift the waveform, and the entire selection will move by that amount. In automatic mode, it removes any signal below 20 Hz, resulting in a signal that is centered on the zero voltage line. This can help increase the dynamic range of a recording.

Interpolate

This command can be used to repair individual spikes in waveforms by drawing a smooth line from the start of the selection to the end of the selection. For best results, zoom in to 1:1 and select the area just around such a spike. If you select the entire file and apply this filter, you will get a straight line from the start to the end of the file, which is not how you want to use this effect.

Invert Signal Polarity

Changes the waveform so that positive voltages become negative voltages and vice-versa. In the waveform view, this has the effect of flipping the waveform on the x-axis, so that it is upside-down relative to the original. You won't hear a difference after you've applied the filter, but this filter is useful if you reversed the positive and negative terminals on your cable connections, and you ended up recording the audio with the voltage inverted. Applying this filter will fix the inversion.

If you consider waves that go above the centerline of the display to be positive and waves that go below to be negative, the Invert filter just makes the positive parts negative and vice-versa. This filter is useful when you have a stereo file and one of its channels is inverted relative to the other. The audio will sound like it's coming from the sides when you listen to it in stereo, with no audio coming from the center.

Swap Left and Right Channels

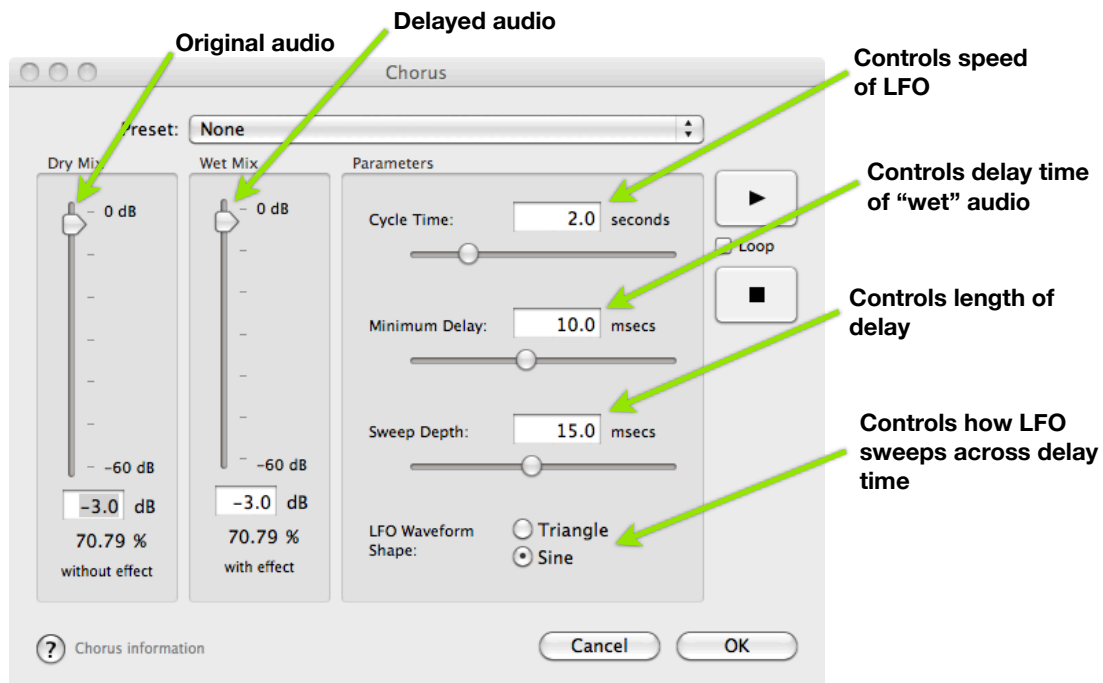
When there are two tracks selected, this command swaps the two tracks. With more than two tracks, it swaps the odd number tracks with the even number tracks.

Backwards / Reverse

This reverses the audio in the file so that it sounds like it is playing backwards.

Chorus

This effect adds the effect of having two voices or instruments playing the same tune, but since they're not perfect, they are slightly out of tune and out of time with each other. It does this by applying a variable-length delay to the audio and mixing the delayed audio with the original audio. An LFO (low-frequency oscillator) controls the amount of delay so that it oscillates between a longer delay and a shorter delay.



The “Dry Mix” is the original audio, and it adjusts the level of the original audio. The “Wet Mix” is the delayed audio, and it adjusts the level of the audio with the variable-length delay applied to it.

The “Cycle Time” controls how fast the LFO (low-frequency oscillator), which drives the variable-length delay, runs. Changing the speed also changes how far out of pitch and time the “wet” audio becomes. A higher speed causes the LFO to run faster, which means that the “wet” audio is being sped up and slowed down much faster.

The “Minimum Delay” controls how long of a delay the “wet” audio will always have. This prevents the “wet” audio from matching up exactly with the “dry” audio.

The “Sweep Depth” controls how far the delay is allowed to go. Increasing this will also cause the pitch and time to vary more.

The “LFO Waveform Shape” controls how the LFO sweeps across the delay time, either in a gentle sine wave or in a more abrupt triangle wave.

Delay and Echo

This effect takes the original audio as input, delays it, and mixes the delayed version with the original version. You can control how much of the original audio (the “dry” mix) and delayed

audio (the “wet” mix) are mixed into the resulting audio. You can also have the delayed audio take its output as its input, resulting in feedback, which is controlled by the “Feedback” checkbox.

Without the feedback, the delay is like having a tape loop which records what you say and plays it back a certain amount of time later. With feedback, it is like shouting across a valley and hearing your voice echo back to you over and over until it dies out.

The “Delay Time” determines how long the delay line is, or how much time there is between each echo.

Note that the effect will end abruptly at the end of your selection, or the end of the file. If you want the delay or echo to decay naturally, you will need to select a few seconds beyond the sound, first adding some silence to the end of the file if necessary.

Flanger

This effect creates a “flanger” sound in the audio. Technically, it is the same as the chorus effect, but with much shorter delay times. The effect causes a sweeping whoosh sound in audio with a lot of broad frequency content, such as hi-hats. You can try the effect by inserting about five seconds of noise and applying the filter with the default settings.

See the Chorus effect for an explanation of the controls.

Pitch and Tempo

This effect allows both pitch and tempo to be changed independently. Pitch is how high or low the notes sound. Tempo is how fast or slow the notes are played.

Pitch is measured in cents, which is one-hundredth of an equal tempered semi-tone. One semitone is 100 cents, and one octave is 1200 cents.

Tempo is measured in percent of the original file's tempo. Fifty percent is half speed, and 200% is double time. Tempo also affects the total duration of the file.

If you want the traditional, analog-style pitch control, turn on the “Link Pitch & Tempo” checkbox. This turns off the pitch shifting compensation and just varies the playback speed.

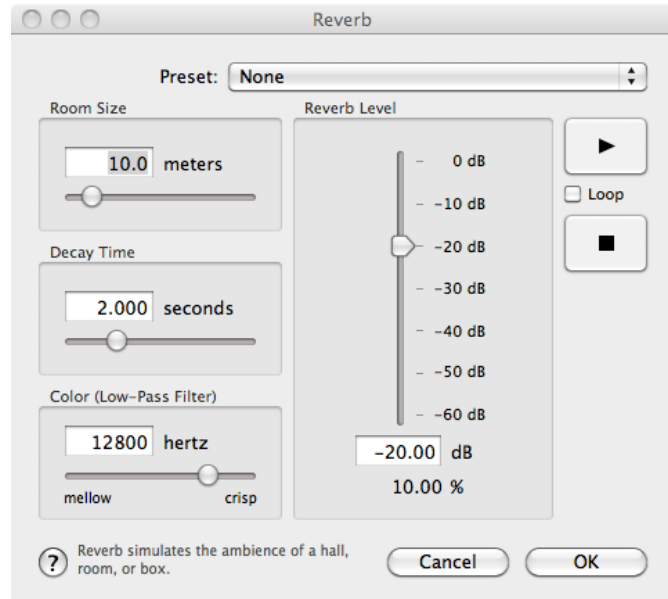
Reverb

This effect adds a natural reverberation effect to the selected audio. When you are in a room, a hall, an auditorium, a stadium, or any other kind of enclosed chamber, the sounds you hear have some kind of reverberation because of the sound waves bouncing back and forth between the walls, the floor, and the ceiling. This effect is most noticeable in a large enclosed stadium, where the announcer's voice echoes through the stadium. You first hear the announcer's voices, and then you hear several, less distinct echoes of the announcer's voice. Usually you don't notice reverb because your ears are used to hearing it, but without it, the audio sounds flat, dry, and lacking in character. Our ears use reverb to define the size and shape of the room we're in.

Audio signals in the computer are often recorded without any reverb. If you record an instrument directly, or if you use a unidirectional microphone or one close to the sound source, you will get little or no reverb in the signal. To make the audio sound grander, we add reverb.

The “Room Size” controls how long the audio takes to bounce back and forth between the walls. The shape is fixed in this effect, but you can adjust the size, so you can simulate anything from a small room to a large stadium.

The “Decay Length” controls how long the reverberations can be heard bouncing between the walls. A short decay means that the reverberations die away quickly, while a long delay means that they can be heard longer. Generally, a bare room with hard surfaces like tile and stone reflect sound well, and will allow the reverberations to keep bouncing around longer. A room with carpets, drapes, and lots of soft furniture will cause the reverb to die away very quickly because all those soft surfaces absorb the sound.



The “Low-Pass Filter” simulates the effect that sound going through the air and bouncing off of softer surfaces will tend to lose their higher frequencies. This option can be adjusted to help color the sound of the reverb.

The “Level” controls how loud the reverbs are compared to the original sound. Typically, you want to keep this low so that the reverb doesn't overwhelm the original audio and make it a mess. You can set the level high if you want the effect of a bad PA system in large stadium where you can't tell what the announcer is saying.

The reverb filter uses six comb filters with low-pass filters and a simple delay. The “Room Size” parameter controls the delay length of each comb filter. The “Decay Length” parameter controls the amplitude of each comb filter. For stereo files, the effect is applied equally to both channels.

A comb filter is a type of filter where the input audio is delayed by a constant amount of time, and then fed back into itself. It is called a comb filter because it will cancel out those frequencies that are 180° out of phase when the delayed audio is fed back upon itself. When you look at the output of such a filter on a frequency graph, it will look like a comb.

3-Band EQ

This effect allows you to change the levels of bass frequency, mid frequency, and treble frequency audio content.

10-Band EQ

You can think of audio as being made up of many different frequencies of sound all mixed together. Your ear has nerves which are sensitive to various frequencies, with the lower frequency nerves detecting a bass guitar while the higher frequency nerves detecting the violins. In reality, these instruments generate many different frequencies of sound at the same time, and the different level of each of these frequencies contributes to the tone, or color, of their distinctive sounds.

Using the sliders, you can adjust the levels of each of these frequencies to color the sound. If the highs sound too harsh or tinny, you can reduce the higher frequencies. If the bass seems to be

too weak, you can increase the lower frequencies. If you want to bring out the vocals, you can increase the frequencies around 1 kHz. Ultimately, you will need to use your ears to find the ideal equalization for your audio.

Graphic Equalizer

This filter allows you to draw in a chart to represent how to adjust the levels in each frequency band.

High Pass

This filter will allow only frequencies above a certain frequency to pass, thereby removing lower frequency audio. You can use this filter to remove lower-frequency sounds. The frequency control determines the cut-off frequency, or the point at which the curve hits the -6 dB level, and the “Steepness” control determines how many samples the filter uses as input, indirectly affecting how sharp the drop-off is at the cutoff frequency. You can use the graph in this dialog to see how the filter affects different frequencies.

One effect you can get with this filter is making audio sound as if it were being played on a cheap radio with tinny speakers. Another effect is to use this filter and the Low Pass filter to emulate audio coming through a telephone.

Low Pass

This filter will allow only frequencies below a certain frequency to pass, thereby removing higher frequency audio. This filter is useful for removing high-frequency noise due to artifacts from the digitization process. You can also use it if you know your source material doesn't have any audio content above a certain frequency, and you want to remove the noise above that frequency. You can also use this filter to create certain effects. One effect is to make the audio sound like it is being played far away and muffled by walls and other things between you and the sound source.

The frequency control determines the cut-off frequency, or the point at which the curve hits the -6 dB level, and the “Steepness” control determines how many samples the filter uses as input, indirectly affecting how sharp the drop-off is at the cutoff frequency. You can use the graph in this dialog to see how the filter affects different frequencies.

View Menu

The main view is the bottom area of the window, while the smaller waveform at the top of the window is the overview. The main view is the editing area.

In addition to this menu, you can also use the mouse scroll wheel and the Zoom In and Zoom Out icons in the window to zoom. On the keyboard, use the E key to go back to the previous zoom.

Zoom In (key shortcut: T)

Makes the waveform appear wider on the screen.

Zoom Out (key shortcut: R)

Makes the waveform appear narrower on the screen.

Fit In Window

Changes the horizontal scale of the main waveform view so that the entire file is visible in the main view.

Zoom All the Way In

Changes the horizontal scale of the main waveform view so that one pixel width represents one sample point.

Zoom to Selection

Scrolls and zooms the view so that the entire selection is visible in the main view.

Vertical Scale

Changes the vertical scale of the main waveform view. The 1x setting is the default. This applies amplification to the waveform to make it easier to see low level audio, but does not affect the audio data.

Ruler Units

Changes the time units displayed in the time rulers and in the time readouts in the main window. The 30 and 25 fps settings are generally used in video editing settings. The 24 fps settings, including Feet and Frames, are generally used in film editing settings.

Markers Shown

Selecting this item will cause the markers to be visible. Turning this off will cause the marker lines to only be shown in the ruler area. In either case, markers can still be edited in the Markers sidebar. Marker line color can be changed in the preferences.

Snap to Markers

Selecting this item will cause selections to be aligned to the closest marker when the mouse pointer is close to a marker.

Lock Markers

Selecting this item will prevent modifications to markers in the main view. You will still be able to edit them in the Markers sidebar.

Delete All Markers

Removes all markers from the file.

Grid Shown

Selecting this item will cause a grid, or lines set at a regular interval, to be visible. The position of these lines is set in the Edit Grid command. Grid line color can be changed in the preferences.

Snap to Grid

Selecting this item will cause selections to be aligned to the closest grid line when the mouse pointer is close to a it.

Edit Grid

Allows the offset to the first grid line and the spacing between grid lines to be changed. Use the BPM setting to have a tempo grid based on beats per minute.

Use the “Tap for BPM” button to tap out a steady beat in rhythm with the music while it plays. It'll show you the BPM after 4 beats, and will keep getting more accurate with the more beats you put in. Hold off the beats for about 2 seconds to reset it.

Show/Hide Sidebar

Selecting this item will show or hide the editing window's sidebar, which contains the table of markers.

Customize Toolbar

Selecting this item will show all available tools that can be added to or removed from the toolbar.

Window Menu

In addition to the commands below, the window menu has a list of all the open windows in Sound Studio, including alert windows and progress windows. Selecting the name of the window from this menu will bring it to the front. This menu also lets you open some utility windows.

Zoom

Toggles window size between maximum size and normal size.

Minimize

Moves the window into the dock.

Input Levels

Shows the Input Levels window, which shows the audio levels on the audio inputs specified in the preferences. This window also has controls for the input gain, which adjusts the levels before the audio is digitized. The “M” slider is the master slider, which controls the gain for the entire device, and the “L” and “R” sliders control the gain for individual channels on the device. The “Link L+R” option causes the “L” and “R” sliders to move in sync. The “Playthrough” option enables a preview of the audio inputs by playing it on the audio outputs; this option is the same as the playthrough checkbox in the preferences.

If you minimize the Input Levels window, it will turn into a little levels meter in the dock, so you can monitor the levels even when you're not in Sound Studio. Increase the dock icon size or turn on dock magnification for more detail.

Sound Studio is copyright © 2010 Felt Tip Inc. All rights reserved worldwide.
Sound Studio 4 Manual copyright © 2011 Felt Tip Inc. All rights reserved worldwide.

For the most up-to-date help and support information, please go to
<http://www.felttip.com/ss/>.

About Ogg Vorbis

Sound Studio includes the Ogg Vorbis library frameworks. The notice below is required for redistribution of these libraries:

Ogg Vorbis

Copyright (c) 2002 Xiph.org Foundation

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Xiph.org Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

About LAME

If installed, MP3 encoding is provided by the LAME framework.

<http://www.mp3dev.org>.

To install LAME, see the instructions at <http://www.felttip.com/ss/>.