

SoundStudio3

Record, edit, and add effects to audio

Users Manual for Sound Studio 3

Written by Lucius Kwok

Edited by Bruce Morrison





Introduction	3
About Sound Studio.....	3
About this User's Manual.....	4
Getting Started	5
Installation	5
Audio Hardware.....	5
How-To Tutorials.....	6
How to Open a CD Audio track	6
How to Turn a Record or Tape into a CD	6
How to Record an Internet Audio Stream (or From Another App)	9
How to Cross-Fade Two Files	10
How to Make a File Fit Within a Specific Time.....	11
Reference for Sound Studio.....	12
Preferences	12
Windows.....	15
Document Window	21
Input Levels Window.....	21
Commands	23
File Menu	24
Edit Menu	25
Audio Menu.....	29
Insert Menu	32
Filter Menu.....	33
View Menu	44
Window Menu	47
Index	48

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:

- ② Operating System: Mac OS X 10.3.9 "Panther" or better, including 10.4 "Tiger".
- ② CPU: PowerPC G3 500 MHz or better, including G4 and G5.
- ② Hard Drive Space: 1 GB minimum (for storing audio data). 5 GB recommended.
- ② Audio Hardware: Built-in audio or other Core Audio compatible sound hardware.
- ② QuickTime: 6.5.2 or later.

About this User's Manual

This document has three major parts:

1. **Getting Started**
2. **How-To Tutorials**
3. **Reference for Sound Studio**

The first part, "**Getting Started**," tells you what you need to know to run Sound Studio.

The second part, "**How-To Tutorials**," walks you step-by-step through some common tasks in Sound Studio.

The third part, "**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.

Words like this denote web page links and email addresses.

The words "app" and "application" refer to the Sound Studio application software. The "operating system" refers to the Mac OS system software.



This chapter contains information you need to know in order to install and set up Sound Studio for the first time.

Installation

- ② Insert the Sound Studio disc in your computer.
- ② Open the disc and double-click on the Sound Studio installer package.
- ② Follow the instructions in the installer.

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 play back, 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. 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.



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 Turn a Record or Tape into a CD

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

- ② **Sound input device.** If your computer does not have built-in audio hardware, you will need a sound card or a USB or FireWire audio capture device.
- ② **CD burner (CD-R or CD-RW).** Most computers come with a CD burner built in, but check that you have one.
- ② **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.

- ② **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 Audio or the name of your USB audio device. Also check that the audio output is properly selected.

Turn on the play-through 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. If the Input Levels window is not open, select **Window > Input Levels**. 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 the 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 **Select Folder** button and create a new folder where you want to save the tracks. Click on the **Choose** button to choose the newly created folder. Then click OK 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 turn the tracks into MP3 files by going to iTunes and selecting **Advanced > Convert to MP3** to convert the files to MP3s.

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.

To be able to hear what you're recording, you should use a stereo Y-connector to connect the two sound jacks together along with external speakers. When you do this, you will want to make sure that play-through is turned off (check in **Preferences** or in the input level window). If your computer doesn't have a sound input jack, you can use a USB audio input device.

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 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: seconds: milliseconds, so for example 3:05.00 is three minutes and five seconds. The pitch of the audio will stay the same.

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.

Preferences

General

- ② **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.
- ② **Make new document when no files are open:** This feature automatically creates a new document when there are no other files open when you click on the application icon in the dock or double-click the application icon in the Finder.
- ② **Default Zoom:** This sets the horizontal scale of newly opened and created documents.
- ② **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 the system "tmp" directory on your system disk.

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.

- ② **Play-through 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. Play-through is useful for monitoring your recording through the computer, to make sure you have the levels and connections set up properly.

- ② **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.

The very top row of this window has the transport controls (stop, play, pause, and record), four time readouts, a control for adjusting the playback volume, and the Sound Studio badge.

The box just below the buttons is the “Overview” waveform view. It always shows the waveform of the entire file. This 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. To the right is the sample rate readout, and clicking it will pop up a menu allowing a different sample rate to be chosen, which will resample the audio if necessary. 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.

Here is a list of keystrokes and what effect they have on the document window:

Key	Effect
Space bar	Starts and stops playing sound
M	Inserts marker at cue point
P	Pauses and unpauses playback and recording
Command-period Escape	Stops playback and recording

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

- ② **Scroll wheel.** 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.
- ② **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 (Stop, Play, Pause, Record)

The playback and recording transport controls are in the top left of the window. They allow you to start and stop 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.

- ② **Playback** starts when you press the play button. Playback will start from the insertion point, or the beginning of the current selection. If you have “Loops when Playing” 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 go back to the beginning of the selection or file and play it again.

You can use the right mouse button to start playback from any point in the file. Click the right mouse button 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.

- ② **Recording** starts when you press the record button. Any recording options you have set up in 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 stop after 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. The window will only show the newly recorded data after you press stop.
- ② **Stop** will always stop playback or recording. The cue will return to the original point where it started when you press stop.
- ② **Pause** will pause playback or recording at the current cue point. For playback, you can use pause to stop playback without losing the current cue position. Click on the pause button again to un-pause playback or recording.

Time Readouts (cue, start, end, length)

The time readouts along the top of the window shows the current cue position and the start, the end, and the length of the current selection. The cue represents the point where the playback and recording head would be, and updates while playing or recording. The start, end, and length readouts only update when you move the insertion point or make a new selection. When there is no selection, the length readout is the length of the entire file, and the start and end readouts show the current insertion point position.

Rulers

The default ruler along the top of the each view has markings in minutes, seconds and milliseconds. 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 (“44100 Hz”)

In the lower left corner 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 plus and minus zoom buttons, the scroll wheel, or the menu commands.

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

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 hold uncheck the "Link L+R" checkbox, the left and right controls will move independently.

Play-through checkbox

This controls software play-through, 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 play-through is on, or you'll get feedback noise.

To enable hardware play-through, 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 play-through.

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.

Sound Studio supports the following file formats:

Acronym	Description	File Extensions	Notes
AIFF	Audio Interchange File Format	aif, aiff	Native Mac OS format. Supports markers.
AIFF-C	Audio Interchange File Format - Compressed	aif, aifc	Extension of AIFF, allows compressed data formats including IMA 4:1
MP3	MPEG Layer III Audio	mp3, mpg	MP3 encoding requires the LAME.framework
MPEG-4 AAC	MPEG-4 Advanced Audio Coding	mp4, m4a, mpg	Unprotected format only; cannot open or save protected "m4p" files
SDII	Sound Designer II	sd2, L, R	Uses the resource fork for important data
WAVE	Windows Wave	wav	Native Windows PC format. Supports markers.

* Sound Studio also supports reading AU and QuickTime movie files. It also supports exporting audio sample values as a tab-delimited text file that contains each and every sample point as numbers.



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.

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.

Open

The Open command shows a panel to select one or more files to be opened.

The Open Recent menu shows the last ten files opened.

Save

Save the current file in the same format and location it was last saved as, which is AIFF by default.

Revert

Discards all changes you’ve made to the current file, and reloads the last saved version.

Edit Menu



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 blinking 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.

This behavior is different from earlier versions of Sound Studio, because earlier versions of Sound Studio were designed for only one- or two-track operation, while this version allows multiple tracks. Earlier versions would try to keep all tracks in sync with each other, so that inserting or deleting audio would affect all tracks. The current version 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.

The pasteboard is private to the application and not available to other applications. 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

Selects the entire sound file from beginning to end and on every track.

Silence

Makes the selected part of audio silent, so that it has zero amplitude. This creates a silent passage in your audio.

Crop

Trims the file down to the part that is selected. This means that the audio to the left of the selection and to the right of the selection is deleted. This is useful if you want to chop off the left and right ends of the file in one operation. You just select the part you want to keep, and crop it. It also removes 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, separate file to a folder, which you can specify by clicking the "Change Folder" button. By default, this folder is the same folder that contained the original file. Each file will take its name from the marker on the left side of each respective section. You can rename the markers in the file before using this command, and new files will be named after the markers after this command is used. The very first section before the first marker, if such a section exists, will be named after the original file.

Add Numeric Prefix to File Names: This option will prefix each file name with a number before saving the file.

The "Split By Markers" command is useful for splitting up a recording of an entire side of an album into separate songs or chapters to be burned as a CD or made into MP3 files.

Go To

This command allows you to go to a specific point in time in the file. It also moves the blinking insertion point to the new location. You can enter a time in hours, minutes, seconds, and milliseconds, and it will move the insertion point to that time. Another option is to scroll the window to the start or end of the current selection.

Set Selection

This command allows you to precisely enter the start and stop points for the current selection. You can enter it in hours, minutes, seconds, and milliseconds.

Audio Menu



The audio menu contains commands which affect the entire file and which are specifically related to audio recording and playback.

Loops when Playing

If this item is checked, the playback head will repeat the selected audio by jumping back to the beginning of the selection when it gets to the end of the selection. If there is no selection, the entire file will be repeated. This command is also known as “Repeat.”

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.

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. 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.

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 when saved to 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 setting is usually the best compromise between file size and sound quality.

Mixdown to Stereo/Mono

These two commands allow you to convert the file to a mono or stereo format.

Fourier Spectrum Analysis

This will run a frequency spectrum analysis on the selected audio, and graph the average levels of each audio frequency.



Marker

This command drops a marker at the current insertion point, at the beginning and end of the selection, or at the current playback cue point if the file is playing. You can also use the 'M' key to drop a marker.

Silence

This command inserts a specified duration of zero-voltage signal, which is also known as silence.

Noise

This command uses the built-in white noise generator to generate triangular-PDF dithered noise.

Tone

This command generates a simple one-frequency tone with a choice of three wave shapes (sine, square, and triangle), any frequency, and any amplitude. Use the keyboard button to bring up a dialog box that lets you select the frequency based a note from a piano keyboard.

FM Synthesis

This command generates a sine wave with frequency modulation. It has a carrier frequency that modulates at the modulation frequency rate and goes up and down by the deviation frequency. For example, a carrier frequency of 440 Hz means that the sound starts out at 440 Hz. The deviation frequency of 20 Hz means that the sound goes up to 460 Hz and goes down to 420 Hz, and a modulation frequency of 5 Hz means that it goes through a complete up and down cycle five times per second.

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, and 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/Audio/Presets/Sound Studio/” folder, in your home directory.

Last Filter

There is a “Last Filter” command that repeats the last selected filter using your previous setting without showing a dialog box. After using a filter, the name of the filter will replace the words “Last Filter” for this command.

Amplify Volume

This will adjust the volume of the recorded audio. You use a decibel (dB) scale to adjust the volume. Positive dB values increase the volume, or amplify, the sound, and negative dB values decrease the volume. You can use the percent readout to see how dB relates to a linear amplification scale.

If you're trying to set the peaks in the levels to an exact value, try the Normalize filter.

Fade In/ Out

Fades the selected audio, using a linear fade. If you wanted to start fading out at five seconds before the song ends, you would select the last five seconds and apply the Fade Out filter.

Fade Special

This filter allows you to specify a custom envelope for fading in or out. You can click in the envelope view to add points to the envelope. 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. It scans the selected audio for peaks and adjust the volume so that the peaks coincide with the target level specified in the options dialog box. 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.

- ② **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.

- ② **Normalize all tracks together:** This option will treat all tracks 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 that 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 that 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

This filter simply turns your waveform upside down. 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

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 “Speed/Rate” 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 / 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 similar to 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 above 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.

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 thorough 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.

Graphic EQ

There is a 3-band, a 10-band, and a custom graphic EQ (equalizer). They all work in the same way. You use each band on the EQ to control the level of the audio within a certain frequency band.

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 a graphic EQ, 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.

The changes you make to your file are permanent and integral to the audio of the file.

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 "Window Width" 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.

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 "Window Width" 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.



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 also the edit area.

In addition to this menu, you can also use the mouse scroll wheel and the zoom slider in the window to zoom.

Zoom In

Zooms the main view by a magnification factor of 2. This lets you see more detail in the waveform. When you get to zoom level 1:1, you can't zoom in any further.

Zoom Out

Zooms the main view by a reduction factor of 2. This lets you see more of the entire file in the window.

Fit All in Window

This lets you see the entire file in the main view.

Zoom All the Way In

This zooms in all the way to 1:1, where one pixel in the main view represents one sample.

Zoom to Selection

This sets the zoom and the scroll position so that the selected audio fits exactly into the main view. You select the area you want to zoom into, and then choose this command. This is useful if you find a glitch in the audio and want zoom in on it to see it better. You can also use the command-click zoom technique to zoom in on a section of audio. Hold down the option key, and then click and drag the mouse pointer over the area you want to zoom in on.

Ruler Units

This sub-menu lets you set the units of measurement for the current file. The ruler and the time readouts will be shown in the units selected.

The “Time” option is the default. This shows the time in hours, minutes, seconds, and milliseconds, in the form hh:mm:ss.xxx, where the ‘x’ is milliseconds. The hours and minutes are omitted if they are zero.

The “Seconds” option shows the time in absolute number of seconds from the beginning of the file.

The “Samples” option shows the time in number of samples from the beginning of the file. Relative to the time in seconds, this number will depend on the sample rate of the file. A 44.1 kHz file will have 44,100 samples per second.

The “Time :Frames” options show the time in hours, minutes, seconds, and frames, in the form hh:mm:ss:fr, where ‘fr’ is frames. The number of frames is determined by the fps (frames per second) setting, either 30, 25, or 24 fps. The 30 and 25 fps settings are generally used for video, while 24 fps is used in film.

The “Feet ‘ Frames” option shows the time in feet and frames at 24 fps. This is the traditional unit of measurement for film, and is where the term “footage” comes from. There are sixteen frames per foot, and at 24 fps, that means that a foot of film is two-thirds of a second long.

Markers Shown

Turning this option on will allow you to see markers in the main view. Turning it off hides the markers. To snap to markers, and to edit markers, you will need to show them. To move a marker, make sure “Lock Markers” is off, and drag the marker to the new location. To edit or delete a marker, also make sure “Lock Markers” is off and double click on the marker. To just edit the name of a marker, click on the marker label and type in a new name.

The color of the markers can be changed in the Preferences.

Snap to Markers

This option causes your selections to try to align with any markers near where you clicked. Turn this option on if you have inserted markers in your file and you want

your selections to snap to those markers.

Lock Markers

This option causes markers to be locked, meaning that you cannot be moved or edited. This is useful if you want to make a selection and you don't want to accidentally move or edit a marker. To delete a marker, this option must be unchecked.

Delete All Markers

This deletes all the markers from the current file.

Grid Shown

This option lets you see a regular grid of lines that repeat at a fixed interval, which you can set up in "Edit Grid." The color of the grid lines can be changed in the Preferences.

Snap to Grid

This causes your selection and insertion point to try to fall on the grid lines, as set up in "Edit Grid." This option is in effect even if you cannot see the grid lines.

Edit Grid

This lets you edit the spacing between the grid lines. You can use this if you want to set up a certain rhythm to your audio, or if you have to edit your audio in chunks of a specific number of samples, for example.

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.

Minimize

Moves the window into the dock.

Input Levels

This command shows the Input Levels window. Please see the “Input Levels Palette Window” section under “Windows” above for more information about this window.

Markers

This command shows the Markers window, which contains a list of all the markers in the current file. You can add, delete, and edit markers in this window.



- amplitude, 16
- automatic recording, 26
- CD burning, 6
 - reading from, 6
- clearing audio, 24
- comb filter, 37
- compression, 31
- crop, 25
- cross-fade, 10
- cue playback, 20
 - recording, 19
- cut, 25
- delay, 39
 - variable-length, 38
- deleting audio, 25
- document window, 16
- dynamics
 - compression, 35
 - expansion, 36
 - noise gating, 37
- equalizer, 42
- expansion, 35
- feedback, 13
- file, 24
 - closing, 24
 - creating new, 24
 - format, 24
 - opening, 24
 - reverting to last saved, 24
 - saving, 24
- filter, 33
 - tempo, 40
 - Add Noise, 37
 - Amplify, 33
 - Backwards, 38
 - Chorus, 38
 - Compressor, 35
 - DC Offset, 37
 - Delay/Echo, 39
 - Expander, 36
 - Fade In, 10
 - Fade In, 34
 - Fade Out, 10
 - Fade Out, 34
 - Fade Special, 34
 - Flanger, 40
 - Graphic EQ, 42
 - High Pass, 42
 - Interpolate, 37
 - Invert, 38
 - Last Filter, 33
 - Low Pass, 43
 - Noise Gate, 37
 - Normalize, 34
 - pitch, 40
 - Reverb, 40
 - Swap Channels, 38
 - Volume, 33
- flanging, 40
- footage, 45
- Fourier, 31
- frames per second, 45
- frequency modulation, 32
- generate, 32
- Go To, 28
- grid, 46
- input gain, 21
- input levels, 21
- insert, 32
 - FM synthesis, 32
 - markers, 32
 - remove track, 32
 - silence, 32
 - tone, 32
 - track, 32
- installation, 5
- iTunes, 8
- keyboard, 17
- keystrokes, 17
- level meter, 21
- loop, 29
- markers, 45
 - deleting, 46
 - inserting, 32
 - split by, 24
- menu
 - Audio, 29
 - Edit, 25
 - File, 24
 - Filter, 33
 - Insert, 32
 - View, 44
 - Window, 47
- mix, 10
- mixdown, 31
- modifier keys, 17
- MP3, 23
- noise, 32

Index (continued)

- noise gate, 37
- paste, 25
 - mixing, 10
- phonograph, 6
- playback, 12
 - cue, 18
 - looping, 18
- play-through, 13
- QuickTime, 23
- RealPlayer, 9
- recording, 18
 - automatic, 26
 - cue, 20
- repeat, 18
- reverberation, 40
- reverse audio *See filter, Backwards*, 38
- RIAA equalization curves, 6
- ruler, 19
 - units, 19
- sample rate, 12
- sample size, 12
- selection, 17
 - manual entry of, 44
 - Select All, 27
 - single channel, 18
 - zoom to, 44
- silence, 27
- sine wave, 32
- spectrum analysis, 31
- split
 - using markers, 27
- stadium, 40
- tape loop, 39
- time readouts, 19
- transport, 18
- turntable, 6
- undo, 25
- USB audio, 7
- voltage, 32
 - inverted, 38
 - offset, 37
- waveform, 12
- window, 15
 - Input Levels, 21
 - Markers, 47
 - minimize, 47
- zoom, 16
 - horizontal, 12

Lucius Kwok
Felt Tip Software

<http://www.felttip.com>
<http://www.freeverse.com>

Software copyright © 2005 Lucius Kwok. All rights reserved worldwide.

This manual copyright © 2005 Lucius Kwok. All rights reserved worldwide.

Sound Studio © 2005 FeltTip software.

Portions © Freeverse, Inc.

Audio Units™ Apple Inc.

All other copyrights are the properties of their respective owners.