

Felt Tip
Sound Studio 2

User's Manual
for version 2.0

Written by Lucius Kwok

Copyright Page

Lucius Kwok
807 Keely Place
Philadelphia, PA 19128

Email: lucius@felttip.com

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

Software copyright © 1999-2002 Lucius Kwok. All rights reserved
worldwide.

This manual copyright © 2002 Lucius Kwok. All rights reserved
worldwide.

Contents

Introduction	4
About Sound Studio	4
14-day Trial Period and Purchasing	5
About this User's Manual	6
Getting Started	7
How-To Tutorials	9
How to Open a CD Audio track	9
How to Turn a Record or Tape into a CD.....	9
How to Record an Internet Audio Stream (or From Another App)	10
How to Cross-Fade Two Files	11
How to Make a File Fit Within a Specific Time	11
Reference for Sound Studio	12
Windows in Sound Studio	12
Document Window.....	12
Input Levels Palette Window	16
Commands in Sound Studio.....	17
File Menu.....	17
Edit Menu.....	19
Audio Menu.....	20
Insert Menu.....	23
Filter Menu.....	23
View Menu.....	29
Window Menu.....	31
Purchase Menu	31
Index	33

Introduction

About Sound Studio

Welcome to Sound Studio

Sound Studio is an audio recording and editing app for the Mac OS. It allows you to take full advantage of your Mac's built-in sound recording and playback capabilities, to digitize your collection of vinyl and tapes, to do live recording of audio, to edit new and pre-existing digital audio, and to save your audio in AIFF or Sound Designer II format to be used in other audio apps.

Main Features

Audio Recording:

- CD-quality recording (44.1kHz, 16-bit stereo)
- Records from any Sound Manager input device, including USB audio input
- Auto start/stop recording options for VOX-style and timed-length recording
- Disk-based recording
- AppleScript support for automatic recording

Audio Editing:

- Waveform display window with zooming down to sample level
- Standard cut/copy/paste editing, plus mixing
- Effects filters including Chorus, Flanger, Reverb, EQ, Compressor, Expander, Noise Gate, Fade In and Out, Normalize, Volume
- Repair clicks and pops with the Interpolate filter
- Remove DC offsets and fix inverted signals
- Tone and noise generators
- Resample to different sample rates
- Adjust the pitch of the sound

File Formats:

- Apple AIFF
- CD Audio Track (requires QuickTime)
- Windows WAVE
- System 7 Sound
- Sound Designer II
- Dual-mono import/export
- QuickTime import/export

Requirements

- Mac OS 8.1 or later (Built for Mac OS X)
- QuickTime 3.0 or later for some features
- Power Mac, iMac, PowerBook, iBook (G3 or higher recommended)
- 12 MB free RAM (20 MB recommended)

What's Included?

The Sound Studio folder should include:

1. **Documentation folder:** contains help, tutorials, and other information.
2. **Sound Studio.app** the application

3. **AppleScripts**: example scripts for controlling Sound Studio.

If you are missing any of these items from your Sound Studio folder, please download the complete package. You can check for the latest version of Sound Studio and download it at:

<http://www.felttip.com/products/soundstudio/>

14-day Trial Period and Purchasing

Felt Tip Sound Studio is distributed as 14-day trial, with the save feature disabled after the trial period is over. You can use it on 14 different days before the trial period runs out.

To purchase a license for Sound Studio, find current prices, and get more information on purchasing, please use this URL:

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

The version of Sound Studio you have downloaded is a complete version, but without a license code, you won't be able to unlock it and have unlimited use of it. Once you purchase a license code, it will be sent to you via email.

To unlock your copy of Sound Studio, you can enter your license code when Sound Studio asks for it, or go to the Purchase menu and select "Enter License Code." In the dialog that appears, you must enter the name and license code exactly as they appear in the email. You can also enter your company's or organization's name in the dialog, but this field is optional.

Please print your license code email for future reference.

Where's My License Code?

Many people place an order with Kagi to purchase a Sound Studio license code, and then wonder when the license code will arrive. Typically, the license code is sent within seconds after the credit card charge is approved (though sometimes it takes one or two hours longer if the system is down for maintenance). The email is easy to miss since it will look like a generic message, since it will come from the address admin@kagi.com, have the title "Thanks for your payment," and have several lines of form data. Don't delete this message, because it contains your license code. You need to scroll down to the bottom half of this email to find your code.

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, and what system requirements are needed and recommended.

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.

Getting Started

This chapter describes how to install Sound Studio on your Mac and set it up for the first time. It also contains information on the minimum system requirements to run Sound Studio.

Installing Sound Studio

You can get the latest version of Sound Studio from the Felt Tip Software Web site:

<http://www.felttip.com/products/soundstudio/>

An alternate download URL:

<http://www.downloadSoundStudio.com/>

Sound Studio is distributed as a StuffIt 5 archive. You will need Aladdin's StuffIt Expander 5.0 or later to decompress the archive, or a compatible decompression utility. Once you have decompressed the archive, you will have the Sound Studio folder on your hard disk.

You can get Aladdin StuffIt Expander from:

<http://www.aladdinsys.com/expander/>

You can install Sound Studio by dragging the entire Sound Studio folder to any place on your hard disk.

If you received Sound Studio on a CD, you can run it from the CD, but it is recommended that you copy it to your hard disk so that you can save presets.

Requirements

Sound Studio has the following system requirements:

- Mac OS 8.1 or later, for the Appearance Manager and Navigation Services
- PowerPC CPU (from the 601 to the G4 and beyond).
- 10 MB of available RAM
- Hard disk to use as a scratch disk

Also, it is recommended that you have the following:

- Mac OS X 10.1.2 or later
- QuickTime 5 installed
- G3 or G4 processor
- 20 MB of RAM
- Built-in audio input or USB audio input
- Free hard disk space, approximately twice the size of the audio files you will be working with.

CarbonLib is not required. If you are running Mac OS 8, you may need to find the Sound Studio application in the path "Sound Studio.app: Contents: MacOSClassic: Sound Studio".

Setting Up Sound Studio

Scratch Disk

The first time you launch Sound Studio, you will want to set up your scratch disk. Choose "Preferences" from the Edit menu, select the name of the disk you want to use as your scratch disk. This is where Sound Studio stores audio data while it is

recording and editing. You will want to select a disk with sufficient available disk space and is fast.

Remember, you will need scratch disk space to open, edit, and record audio. If you are recording an hour of audio, you will need about 600 MB of disk space free just to record the audio. Then, you will need another 600 MB to save the file. Also, you will need another 600 MB if you apply any filters to the entire file. That would total to about 1.8 GB of disk space for recording and editing an hour of audio at CD-quality settings.

System Options: Virtual Memory and AppleTalk

You will want to turn off virtual memory if you are not booting into Mac OS X. You may also want to turn off file sharing and AppleTalk. Both Virtual Memory and AppleTalk can interfere with Sound Studio's ability to record and play back audio without digital glitches. If you find that you get skips and clicks when recording and during playback, try turning off virtual memory and AppleTalk.

Application Location: Where to Put Sound Studio

Sound Studio should be run from a folder on a hard disk where you have write access, because it stores a "Presets" folder in the same folder as the application package. This is where you can save presets for some effects filters. While you can save presets to any location, Sound Studio expects the presets you want to use are located in the "Presets" folder hierarchy.

You can put Sound Studio in a locked folder or volume, but you won't be able to add any presets in that situation.

Hardware that Sound Studio Uses

In order to record, Sound Studio needs some kind of audio input hardware. On most older Macs, this is already built into the CPU and you don't need any other hardware to record audio. On PowerBooks and iBooks, there is a built-in microphone which can be used for low-quality recordings. Other Macs may have a 1/8" stereo mini-jack for sound input with a microphone icon, or RCA jacks for line-level input. These Macs have built-in audio input.

On most recent Macs models, the sound input ports were removed in favor of USB audio input, which requires that you purchase an extra device for sound input. This would either be a USB audio input device such as the Griffin iMic and Roland UA-30, or an internal PCI audio card.

USB audio input devices will work with Sound Studio, since Apple provides the drivers which communicate with the USB devices. You can plug the audio input device while Sound Studio is running, and then go to "Sound Input/Output Setup" to choose your input device.

PCI audio cards may work if the card manufacturer provides Sound Manager (for Mac OS 8-9.x) or Core Audio (for Mac OS X) drivers for the card. Since Sound Studio does not support ASIO, there are some cards which will not work with Sound Studio.

As for sound output hardware, Sound Studio has a selector box in "Sound Input/Output Setup" for choosing between available output devices, so you can choose between the built-in audio system or USB speakers if you have them plugged in.

How-To Tutorials

This chapter shows you step-by-step how to do common tasks in Sound Studio. You will learn how to digitize vinyl records and analog tapes, record Internet audio streams, import CD audio tracks, cross-fade between two songs, and how to fit a file into a specific duration using pitch shifting.

The first thing to remember is that Sound Studio is designed to let you experiment with the various editing commands and filters. You can undo your last command (though it is only a single-level undo), and you can revert to your last saved version in case something really goes wrong.

The second thing is to be sure to save your file at various times when you are editing, since you might want to go back to a previous version if the edits don't turn out OK.

How to Open a CD Audio track

For Mac OS X: use **File→Import with QuickTime** to open CD tracks. The direct-open method below won't work under OS X.

For Mac OS 9: To open a track on an audio CD, select **File→Open**, navigate to your Audio CD, and open your CD track directly. The file will be converted into an AIFF file which you can edit and save. All CD audio tracks will open up as 44.1kHz, 16-bit, stereo files.

If that doesn't work, you can also use **File→Import with QuickTime** to open CD tracks. Both of these methods require QuickTime 2.0 or later be installed.

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. These are usually built into the computer, but the iBook and Cube, for example, don't have them. You will need a USB Audio input device such as the Roland UA-30 or Griffin Technology iMic for these Macs.
- CD Burner (CD-R or CD-RW). These come with software that allows you to make audio CDs. This article won't go into detail about using your burner software.
- For phonographs: If you have a turntable, you will also need an amp 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.

When you have that all set up, launch Sound Studio and go to **Audio→Sound Input/Output Setup** and make sure that **Built-In→Sound In** is selected if you have built-in sound input, or USB Audio is selected. Please see the USB Audio section if you have trouble here.

Make sure the play-through option is on, and click OK to close the dialog. If you start your analog source playing, and your Input Levels window is open, you should see the levels meter start working. 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.)

Once you've set up the levels, you can start recording. Select **File→New** and create a 44.1kHz, 16 bit, stereo file. You can start recording now. The maximum file size is 2 GB, which translates to about 3 hours and 20 minutes at the current settings. The best way to record is to do an entire side of the record or tape in one pass, and then edit later.

When one side is done, click stop. The window will start to redraw, but you can start editing right away. 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 in Window**. Then, 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 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**. This command will create a new file for each track. In the "Split By Markers" dialog, click on "Select Folder" and create a new folder where you want to save the tracks. Click on "Choose" 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 **Audio→Sound Input/Output Setup** or in the input level palette 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.

Mix (under the Edit menu) into the first file at a point five seconds before its end. The audio on the clipboard will be mixed on top of the first file at the insertion point, and the file will be extended to hold the new mix. And there you have your cross-fade.

How to Make a File Fit Within a Specific Time

To change the duration of an entire audio file, select **Audio→Adjust Pitch** and enter the duration the file should be. The duration is in the format hours: minutes: seconds: milliseconds. The pitch of the audio will change with the duration, but as long as the change is not more than 5%, it shouldn't sound noticeably odd.

When you change the duration, you will notice that the sample rate and pitch controls will change. All the controls are linked together, since the pitch is adjusted by speeding up or slowing down the sample rate, as if a tape deck were being sped up or slowed down.

After using **Adjust Pitch**, it's a good idea to **resample** the audio back to its original sample rate, typically 44.1 kHz. This way, you can record a CD of your audio, because CDs require 44.1 kHz, stereo, 16-bit audio. There is a preset for this in the Resample dialog.

Reference for Sound Studio

This chapter describes in detail each of the commands in Sound Studio, and the parts of its windows. This chapter is organized into two sections. The first section, “Windows” has information on the layout of the main document window, the input levels palette window, and some conventions used in dialog boxes.

The second section, “Commands,” goes through each menu command individually and gives some hints and tips on using each command to its fullest.

To check for the latest version of Sound Studio, please go to:

<http://www.felttip.com/products/soundstudio/>

For the latest support issues in Sound Studio, and for more information, please go to:

<http://www.felttip.com/products/soundstudio/support.shtml>

Windows in Sound Studio

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

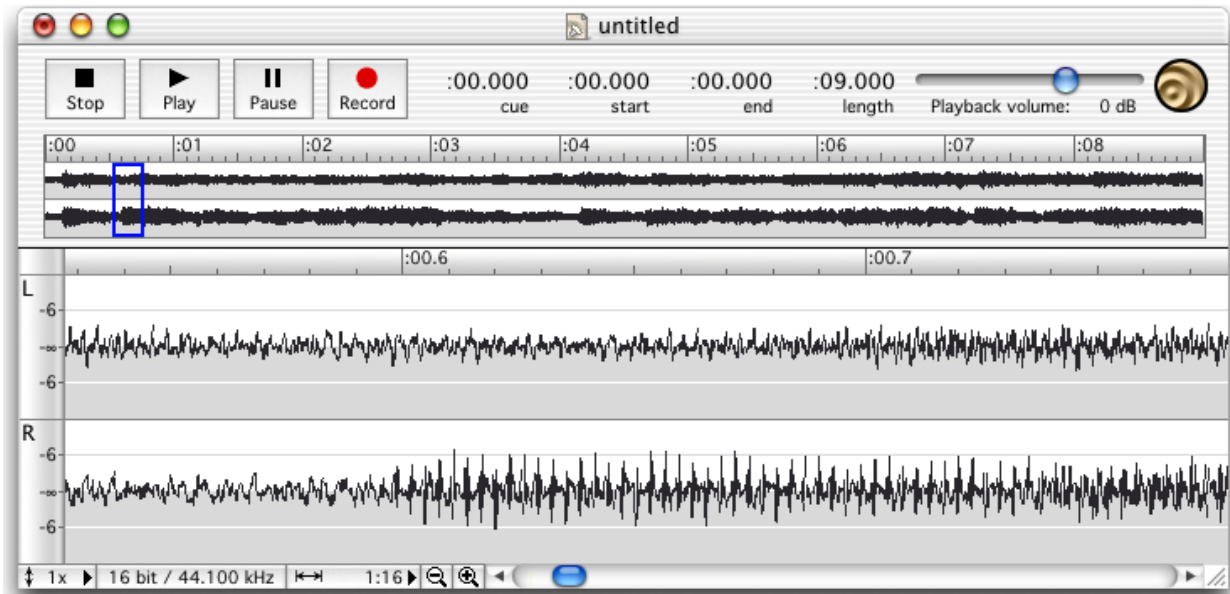
Document Window

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 which is drawn around the part of the file which 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. Just below the amplitude ruler is a popup button which controls the vertical zoom level of the editing area. To the right is a status display of the current sample bit-depth and sample rate of the file. Next to that is the horizontal zoom level, with 1:1 being the closest you can be, and higher numbers meaning



you have zoomed further out. The magnifying glasses with plus and minus can be clicked on to increase or decrease the horizontal zoom.

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	Stops playback and recording

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

Modifier Key	Effect on editing area	Effect on overview
None	Normal selection	Moves view box
Shift	Extends selection	Same as above
Option	Grabbing hand: scrolls selection	Same as above
Command	Zoom In. Click-drag to draw a zoom marquee.	Click-drag to draw a new view box.
Command + Option	Zoom Out	Same as above
Control or right mouse button	Start playback at point	Start playback at point

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.

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 “Loop When Playing” turned on, the selection or the entire file will be looped, 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.

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.

If the file is empty, recording will start from the beginning of the file and when you press stop, the length of the file will update to reflect the length of your recording. If your file already has audio in it, recording will start from the insertion point or the beginning of the selection, and overwrite any audio after that point. You will see the red cue line which represents where the recording head is writing the audio data. If you have selected some audio, recording will only write audio data within the selection. Otherwise, it will keep recording new audio data until it reaches the limit set up in “Auto Start/Stop Recording.”

The file will only update its display to reflect the newly recorded data when you press stop. Recording is not undoable.

Stop will always stop playback or recording. The cue will return to the original point where it started when you press stop.

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 ruler along the top of the view has markings in minutes, seconds and milliseconds. The units can be changed by going to View->Ruler Units and choosing 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. The top and bottom limits of the waveform would normally be 0 dB, and are not labeled. The center between the limits is - 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 32x.

Playback Cue (the Green Line)

While playing back the audio, this green line shows you where the playback head currently is. Some time drift may occur since the timer is only about 99% accurate. 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 over existing audio, this red line will show you where the record head is overwriting the existing audio. When recording to the end of the file, the line will only be visible when you just start recording, and then it will continue off the right side of the window. See the “Transport” section above for more info on recording.

Sample Size and Rate Display (“16 bit / 44.100 kHz”)


In the lower left corner of the window is a button which shows the file’s sample size (8 or 16 bit) and sample rate (in kilohertz). Clicking on this button brings up the Resample dialog, which will let you change the sample rate, sample size, and number of channels.

Zoom Pop-up Menu | |

The Zoom button, located next to the Sample Size/Rate Display, shows the zoom ratio, or how many samples each pixel (dot) represents horizontally. A ratio of 1:1 means that you are seeing every sample, while a ratio of 1:100 means that each pixel (dot) represents 100 samples in the audio file.

Clicking on this button reveals the pop-up menu where you can select a predefined zoom ratio. You can also change the zoom using the View menu, by selecting a new view box in the overview, or using modifier-key and mouse combinations as defined in the modifier key table above.

To the right of this popup are two buttons with the plus sign (+) and the minus sign (-) in magnifying glasses. These buttons let you zoom in and out.



Input Levels Palette Window

The Input Levels palette is a floating window that shows the current levels at the sound input, and lets you control the input gain (amplification). You can show and hide the Input Levels palette by using the Window menu. This window is automatically hidden when Sound Studio is put in the background.

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 there are no bars lit, 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 set the levels to just peak at -6 dB on the scale most of the time, while avoiding any signal that goes into the clipping range.

Stereo Input Gain Control

The input gain control is directly under the level meter, and controls the gain, or amplification, at the sound input, before the audio is digitized. You want increase the gain as much as possible without going into the red on the level meter and possibly causing clipping.

Double-clicking on the control will return it to the middle position.

The left and right controls are normally linked. However, if you hold down the option key while moving the controls, the left and right controls will move independently.

Play-through checkbox

This controls software play-through. This checkbox does the same thing as the one in the Sound Input/Output Setup dialog. You want to be sure to keep your microphone away from your speakers when play-through is on, or you'll get feedback noise.

Commands in Sound Studio

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

New

To create a new document which you can record into, select this command. It will present a dialog box where you need to choose your sound quality settings. In most cases, you want to use 44.1 kHz, 16-bits, and stereo, which is CD quality. The name of the document can be left as it is, since you will have a chance to name the file when you save it.

The New command creates a new document, which is where you will be doing all your work, including recording, editing, and playback. It creates on your scratch disk a scratch file which will grow as you add more audio to it. The document also keeps track of your view settings.

The sound quality settings are important because you need to have them set correctly in order to record.

The 11.127kHz and 22.254kHz sample rates are there for historical reasons. These rates were the standard sample rates for 68K Macintosh computers. The 11.025kHz and 22.050kHz rates are even multiples of the standard 44.1kHz rate, and are useful for reducing the size of files by cutting off the high-frequency end of the audio. The 32kHz and 48kHz rates are used by some USB audio input devices, some MPEG standard audio, DAT decks, and other audio equipment.

Open

Opens an existing Apple AIFF, CD Audio Track, Windows WAVE, System 7 Sound, or Sound Designer II audio file. Sound Studio can open audio files saved with an 8 or 16 bit sample size, one or two channels, and any sample rate up to 65335 Hz.

If the file is encoded or compressed, it will ask you if you want to import the file instead. This will use the Import with QuickTime command to open the file.

Close

Closes the current file, asking to save changes if the file has been modified.

Save

Save the current file in the same format it was last saved as, which is AIFF by default. Disabled if no changes have been made to the window.

Save As

Asks you where you want to save the current file, and allows you to change the file format to AIFF, Windows WAVE, System 7 Sound, or Sound Designer II. The file format selector requires Navigation Services.

Only AIFF format files will save marker information. This information is saved according to the AIFF standard, so other audio applications should be able to read them.

Save Selection As

Saves the selected audio to a new file. It works the same way as “Save As” described above.

Revert

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

Import/Export with QuickTime

Lets you use QuickTime to open and save files in any file format that QuickTime recognizes. When importing, you will see an import dialog where you can set various sound quality settings — the compression setting is ignored and the audio is imported as an uncompressed file. When exporting, it will honor the compression setting, so you can use IMA 4:1, for example, to reduce the file size.

Usually, you won’t be able to export in as many file formats as you can import. Under QuickTime 4, you can import MPEG-encoded audio files, including MP3s, but you can’t export to the MP3 format.

This feature requires QuickTime, and version 4.0 or later is recommended.

Import/Export Dual Mono

Lets you open and save stereo files as a pair of mono files. This format is commonly used by Digidesign ProTools to store stereo audio. When importing, it will ask first for the left channel to open, and then the right channel, and will combine the two into a stereo file. When exporting it will let you save the left and right channels separately.

Export as Text

Saves the current file as a tab-delimited text file, with each sample represented by an integer. For a mono file, each sample is on its own line. For stereo files, the left sample is first, and then the right channel is second, separated by a tab, and each pair of samples is on its own line. Files saved with this command can easily be four times as large as the original file

Preferences

Shows the preferences dialog, where you can select which disk to use as your scratch disk, and whether to scroll the window while playing in order to keep up with the green playback marker.

This command is in the application menu under Mac OS X.

Quit

Closes all open windows, asking to save those with unsaved changes, closes the sound input driver, and quits the application.

This command is in the application menu under Mac OS X.

Edit Menu

Undo / Redo

After you have done something that can be undone, you can use this command to undo that action. This item toggles between undo and redo, where redo repeats your action after you undo it. With this command you can switch between the before and after states of a filter or edit, and compare how they sound.

Cut

Cuts the selected audio from the current file and puts it on the clipboard.

Copy

Copies the selected audio onto the clipboard

Paste

Pastes the audio on the clipboard into the current file, replacing any selected audio. This command will automatically resample the audio on the clipboard to match the current file's sound quality settings, so you can copy from one file with a 22 kHz sample rate and paste into another file with a 44.1 kHz sample rate, for example.

Mix

Displays a dialog which 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.

Clear

Removes the selected audio and moves the audio after it to fill in the gap. If you want to remove audio without changing the timing of the audio, use the "Silence" command instead.

On a single-channel selection in a stereo file, this command behaves the same as "Silence."

Select All

Selects the entire sound file from beginning to end. If only one channel of a stereo file was selected, it only selects the entirety of that channel, and not the other channel.

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 which 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.

Split at Insertion Point

Turns the current file into two files by splitting it at the current blinking insertion point. Everything before that point stays in the current file, while everything after it goes into a new file.

This command is not available if you have made a selection or if you put the insertion pointer at the beginning or end of the file.

The preferred method of splitting is to start from the end of a file and work your way to the beginning, since it will be a little bit faster and will use up less scratch disk space. However, if you have many songs to separate into different files, it is faster to place markers between each song and use the “Split By Markers” command, described below.

Split By Markers

If you have markers inserted in your file, you can use this command to split the file into many new files. Imagine your file is one long tape with your audio on it, and your markers are the points where you will take your scissors and cut the tape. This command does essentially the same thing, with some minor differences. It only copies the data to new files, so the original file isn't changed.

It takes the clippings and has to save them somewhere. This is where the “Save To” option comes in. You can select a folder where the new files will go. By default, this is the same folder that the current file is located. You will probably want to create a new folder so that you keep all the new files in one place and you don't get the new files mixed with other files.

It also has to name the files something. Each file gets its name from the marker that marks the point where its starting point was cut. Also, you have the option of prefixing the names with a number, so that the files retain the same order they were in when they were part of one file. The “Add numeric prefix to file names” option adds these prefixes.

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 to go to in hours, minutes, seconds, and milliseconds, you can scroll the window to the start or end of the current selection, or you can go to the location of a specific marker. If you go to a time or marker, the current selection will be unselected.

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, or you can select a marker to serve as the start or end points.

Audio Menu

Loop When Playing

If this item is checked, the playback head will loop over 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 looped. Also known as “Repeat.”

Resample

Changes the sample rate, sample size (8 or 16 bits), and number of channels (mono channel or stereo) of the active window, without changing the pitch or the overall duration of the entire audio file. This command is a bit slower than the other commands in this application, so it's advised that you use it with small or empty files only.

In the Resample dialog, there is a sample rate field where you can type in a sample rate or select a pre-defined one from a pop-up list. All Macs support not all rates, so for some rates you will see an alert saying that you can't record at that rate. This does not affect the editing or saving functions of Sound Studio, only recording.

Please see the "New" command in the "File Commands" section for a discussion of sample rates.

Adjust Pitch

Changes the sample rate of the active window without touching the actual audio data, resulting in audio that runs at a different speed and has a higher or lower pitch, based on your new settings. Both pitch and duration are simultaneously changed.

The reason that both pitch and duration are changed at the same time is because this requires changing only the sample rate of the file. Changing them independently would require a complex algorithm which separates the frequency information from the time information.

Spectrum Analysis

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

Auto Start/Stop Recording

Allows you to set up the way audio recording works in the current file. 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.

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 start again 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 2 GB file size limit is an absolute limit which is applicable for all the options below, so even if you select an option which would cause the file size to go beyond 2 GB, the recording will stop before it goes beyond the 2 GB limit. The default is to keep recording until it hits this limit.

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 which 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.

Sound Input/Output Setup

Allows you to select the audio input device, its audio source, and the audio output device. These settings are application-wide, which means that they apply to everything in Sound Studio, and aren't specific to any one file.

The "Input" section lets you set up the audio input device. The device list lets you choose an audio input device, which is an abstraction which typically represents a piece of audio hardware, either an external unit which you can hold in your hands, or something internal or abstract which you can't actually hold in your hands. For example, if you have an iMic plugged in, it will show up as an "iMic USB audio adapter" in this list. If you click on a device, the "Source" list will update to show the available audio sources on the device. A source typically represents an input port, jack, or connector on the device. If your device only has one input port, it will probably have only one source. The names of these devices and sources are supplied by the operating system, so different operating systems and device drivers may report different names.

The option "Play Input Audio Through Output Device" turns on software play-through, which will let you hear the audio coming in on the audio input by playing it on the output device. Some devices have hardware play-through, which does the same thing as software play-through, but much faster. If your device has hardware play-through, you can turn the software play-through off.

The "Output" section lets you select which device you want your play-through and your audio files to playback on. The names of these devices are supplied by the operating system, so different operating systems and device drivers may report different names. Under Mac OS X, there may be a "Mac OS X Audio HAL (System)" device, which is the output device which you select in the "Sound" control panel in System Preferences. The other devices may refer to other available audio output devices.

Insert Menu

Marker

Drops a marker at the current insertion point, or at the current playback cue point if the file is playing. You can also use the 'M' key to drop a marker.

Silence

A very simple sound generator which creates a stream of zero-voltage signal (a.k.a. silence).

Noise

Uses the built-in white noise generator to generate triangular-PDF dithered noise.

Tone

This generates a simple one-frequency tone with a choice of three wave shapes (sine, square, and triangle), any frequency, and any amplitude.

FM Synthesis

This generates a sine wave which has a carrier frequency which 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.

Filter Menu

All filter commands require that some audio be selected. If you don't have a window open and some audio selected, the menu commands will be dimmed, since you haven't selected anything.

Some filters will let you save presets. A preset allows you to save the position of the controls of a filter so that you can recall them later. You can save several presets, and they will be available in the "Presets" popup if you save them to the appropriate folder. To create a preset, set up your filter controls, and then click the "Save" button in the dialog. It will ask you what to name the preset, and the location will be already set to the appropriate presets folder. Give it a name, and after you save it, it will appear in the "Presets" popup for the current filter. To delete a preset, select it from the popup and click "Delete."

The presets will be saved to a "Presets" folder located in the same folder as the Sound Studio application. Inside the "Presets" folder are several sub-folders, each one corresponding to a filter. You can open those folders to copy, move, and rename your saved presets.

Last Filter

There is a "Last Filter" command which 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.

Fade In/ Out

The fade is a logarithmic fade, which sounds more even than a linear fade to a human ear.

You'll want to select just the part that you want to fade. For example, 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 filter brings the volume of your audio as high as it can without clipping. This way, you can make sure you're using all the bandwidth (height of the channel) that you can get. Adjusting the 'Normalize to' value changes the maximum amount of bandwidth you want to use. At 50%, the sound volume will be brought up (or down) so that only half of the available bandwidth will be used.

The 'Adjust tracks individually' option controls whether you want to treat the two left and right tracks as independent audio tracks or as one stereo track. If the audio you're editing already has its balance between the left and right sides perfectly set, then you want to leave this option off. However, if one side is much louder than the other and you want them balanced, select this option.

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.

The “Blending” controls how much the edges of the selection are blended. Blending helps smooth the transition between the amplified selection and the unamplified audio around the selection. It essentially cross-fades between the amplified and unamplified audio at the edges of the selection.

If you're trying to get the sound volume to a specific peak loudness, try the Normalize filter.

Compressor

This command attempts to lower the volume on the loudest passages of audio in the selection. 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.

Expander

An expander is in some ways 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

A noise gate is essentially an expander with the ratio turned up. 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 which 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

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 connected the positive and negative terminals backwards 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 which go above the centerline of the display to be positive and waves which 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 Channels

This filter swaps the left channel with the right channel, in case you had the left and right cables mixed up when you recorded the audio. Or you can use this effect on a short selection of audio to Ping-Pong sounds from left to right and vice-versa. This filter only works on stereo files.

Backwards

I don't know if anybody really uses this filter for anything but having fun. It turns your sound around so that it plays backwards. The effect is the same as playing a tape 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. This delay is controlled by a LFO (low-frequency oscillator), which makes the delay longer and shorter.

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.

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 voice, 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 delay length of each comb filter is governed by the “Room Size,” and the amplitude of each comb is determined by the “Decay Length.” For stereo files, the effect is applied equally to both channels.

Graphic EQ

There is a 3-band, a 10-band, and a 30-band 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.

View Menu

Remember that the main view is the larger waveform in 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.

All of these options will be saved with the file. The window position, the scroll position, and the zoom level are also saved with the file.

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 in Window

This lets you see the entire file in the main view, by zooming out.

Every Sample

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.

Show Waveform

Turning this option on will allow the waveform to be drawn as it represents the audio in the file. Turning it off will show a “tape-style” display with a solid bar representing the audio. The color of the waveform can be changed in the Preferences. By default, this option is on.

Show Markers

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 can not 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.

Clear All Markers

This deletes all the markers from the current file.

Show Grid

This option lets you see a regular grid of lines which 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 can not 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.

Snap to Zero Crossings

This option causes your selection and insertion point to try to fall on places where the waveform crosses the “zero” line in the display, which is shown by the label “- ” on the amplitude ruler in the window. This is typically where the waveform lies when there is silence, and the waveform generally crosses this line many times when there is a signal. Since the waveform crosses “zero” line many times, it is often best when doing edits to have the selection start and end at these points. Also, when pasting two pieces of audio together, their waveforms must meet at a common point, or you will hear a pop or click. The “zero” line is a convenient reference point.

The reason that the “zero” line is shown as “- ” in the amplitude ruler is because the ruler is shown in decibels, while the waveform is shown in linear amplitude. Converting from the decibel scale to the amplitude scale, - dB equals 0 linear amplitude.

Window Menu

The window menu has, in addition to the commands below, 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 Window

Available only under Mac OS X, this command minimizes the current window and puts it into the dock.

Bring All to Front

Available only under Mac OS X, this command brings all Sound Studio windows to the foreground, un-minimizing any windows in the dock.

Show/Hide Input Levels

This command shows the input levels palette window when it is hidden, and hides it when it is shown. Please see the “Input Levels Palette Window” section under “Windows” above for more information about this window.

Cycle Through Windows

This command lets you go through all the open windows in Sound Studio, making each one the current window in turn.

Purchase Menu

The Purchase menu is only available if your installation of Sound Studio does not have a valid license code entered into it.

Purchase a License

This brings up a dialog which will tell you how to purchase a license for Sound Studio. It will also let you launch your web browser and open a URL to the web page to purchase a license. The URL is:

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

Enter License Code

This brings up a dialog where you can enter your license code into Sound Studio. This is how you can unlock the “Save” function and have full, unlimited use of Sound Studio.

Index

Add Noise.....	25	license code.....	5
Adjust Pitch.....	21	loop when playing.....	20
Amplify.....	24	Low Pass.....	28
audio menu.....	20	markers.....	30
Auto Start/Stop Recording.....	21	Minimize Window.....	31
Backwards.....	26	Mix.....	19
Bring All to Front.....	31	New.....	17
CD.....	9	Noise Gate.....	25
Chorus.....	26	Normalize.....	24
Clear.....	19	Open.....	17
Close.....	17	Paste.....	19
Compressor.....	24	playback.....	14
Copy.....	19	playback cue.....	15
crop.....	19	preferences.....	7, 18
cross-fade.....	11	presets.....	23
Cut.....	19	purchase.....	5
Cycle Through Windows.....	31	Purchase a License.....	32
DC Offset.....	25	Purchase Menu.....	31
Delay.....	27	Quit.....	18
document window.....	12	RealPlayer.....	10
download.....	7	recording.....	14
Echo.....	27	recording cue.....	15
edit menu.....	19	Resample.....	20
Enter License Code.....	32	Reverb.....	27
Expander.....	25	Revert.....	18
Export as Text.....	18	ruler.....	15
Fade In.....	23	Ruler Units.....	29
Fade Out.....	23	sample rates.....	17
Fade Special.....	24	Save.....	17
file menu.....	17	Save As.....	17
Filter Menu.....	23	Save Selection As.....	18
Flanger.....	27	scratch disk.....	7, 18
Go To.....	20	Select All.....	19
Graphic EQ.....	28	Set Selection.....	20
grid.....	30	Show / Hide Input Levels.....	31
High Pass.....	28	silence.....	19
Import/Export Dual Mono.....	18	Snap to Zero Crossings.....	31
Import/Export with QuickTime.....	18	Sound Input/Output Setup.....	22
input gain.....	16	Spectrum Analysis.....	21
input levels palette.....	16	Split at Insertion Point.....	19
Insert FM Synthesis.....	23	Split By Markers.....	10, 20
Insert Marker.....	23	Swap Channels.....	26
Insert Menu.....	23	time readouts.....	14
Insert Noise.....	23	transport controls.....	14
Insert Silence.....	23	trial period.....	5
Insert Tone.....	23	Undo / Redo.....	19
Interpolate.....	25	View Menu.....	29
Invert.....	26	Volume.....	24
iTunes.....	10	Window Menu.....	31
Last Filter.....	23	zoom.....	29
level meter.....	16		