

SNDSAMPLER™

USER'S GUIDE

Copyright © 1997 by Alan Glenn
All rights reserved

SndSampler™ Version 3.4
By Alan Glenn and M. Q. Edison
Copyright © 1997 by Alan Glenn
All rights reserved

SndSampler 3.x (meaning SndSampler version 3-point-whatever) is **shareware**. If you like it and want to keep it, please register by sending \$20 U.S. (cash OK, but we're not responsible if it gets stolen!) to:

Alan Glenn
4516 Cruz Drive
Midland, MI 48642
U.S.A.

Please fill out the accompanying registration form and include it along with your check. Please make checks drawn on a U.S. bank and payable to Alan Glenn. PLEASE DO NOT SEND CREDIT CARD NUMBERS! We cannot process credit cards at this time.

Special upgrade offer for registered users of SndSampler 2.x:

If you are a registered user of SndSampler 2.x you can upgrade to 3.x for a mere \$8. This is simply the difference between the price of 2.x and 3.x.

Comments and bug reports are also welcome. If you report a reproducible bug that we haven't found yet, you'll get **free registration!** When reporting bugs, try to be as specific as possible, i.e. what were you doing exactly, where were the start and end pointers, was the sound compressed, what was the sample size and number of channels, what kind of Mac do you have, what system version, etc. Please realize that we must be able to reproduce the bug before we can give you free registration. Also, it has to be a bug that we haven't already fixed.

Again, SndSampler is shareware. **You are allowed to use it on a trial basis for 30 days.** After that, you must either register or delete SndSampler. You are also allowed to freely distribute *unregistered* copies of SndSampler. As a matter of fact, we strongly encourage this. Registered copies, on the other hand, may not be distributed. And it probably goes without saying that you may not alter any of the materials which accompany SndSampler, or the application itself.

Introduction

SndSampler is the ultimate tool for recording, editing, and working with Macintosh sounds. SndSampler is simple enough for the novice but contains many powerful features which will appeal to the sound-editing professional. SndSampler will work on any Mac running System 7.0 or later. It is helpful, but not necessary, for your Mac to have a sound input device. Of course you won't be able to record sounds without a sound input device, but you can still edit existing sounds.

Please note: if your Mac doesn't have 16-bit audio hardware, you must have Sound Manager 3.0 or later installed in order to play 16-bit sounds. Otherwise your Mac will emit horrible screeching noises when you try to play a 16-bit sound. (System 7.5 and later have Sound Manager 3.0 and later built-in.) Also, some Macs don't offer stereo playback. Many models will only play the left channel of a stereo sound. Check to see what your Mac hardware offers if you are having problems with stereo sounds. If you don't have stereo output capability, Sound Manager 3.0 will mix the left and right channels during playback. If you haven't already, you really ought to get yourself a copy of Sound Manager 3.0 or later. Be advised that we have found that Sound Manager 3.2 does not work well (or at all) on some older Macs (like ours).

SndSampler Reference

The Windows

SndSampler has two main display windows: the sample window and the stat window. The **sample window** is by default the one on the top (which is where we prefer to keep it). The first thing you see in the sample window is a big rectangle, which is for viewing a graphical representation of your sound plotted as amplitude versus time, something like you might see on an oscilloscope. (When displaying a stereo sound, SndSampler will alternate between drawing the left and right channels, so what you see will be kind of an average of the stereo signals.) You can resize this window however you like. In the lower left corner of the sample window the number of free bytes of RAM is displayed. When you are performing some kind of operation, SndSampler will display messages here to keep you updated as to what is happening.

Below the sample window is the **stat window**. First, the stat window shows you how many frames are in the entire sound. A frame can be thought of as one digital sample of an analog sound, taken at a specific instance in time. For 8-bit mono sounds the sample frame size is 1 byte (1 byte = 8 bits). For 8-bit stereo the frame size is 2 bytes (the first byte for the left

channel, the second byte for the right channel). For 16-bit mono the frame size is 2 bytes (two bytes for one sample). For 16-bit stereo the frame size is 4 bytes (you get the idea). The sound length value will not change until you do something which changes the size of the entire sound. Also, if you click on this number, it will change to display the sound's **playing time**, either in milliseconds (millisecond = ms = 1/1000th of a second), or, if it is long enough, in seconds. Click again to return to the frame display. The start and end points will also convert to time when you click on the sound length.

Next, going down, the start and end positions of the sample are given in sample frames. The portion of the sound which falls between the start and end positions is called the "selected portion" of the sound. When you save, the selected portion of the sound will become the new sound. In other words, you will only be saving the selected portion of the sound. Next comes the disk size of the sound. This is the actual disk space, in bytes, which would be taken up by your sound if it were to be saved immediately. Below the disk size is the file type. This is the type of file from which the sound data originated. Note that if SndSampler doesn't recognize the file's type, it will simply display the Finder's four-character type identifier for that particular file. Sometimes, depending on the file, one or more of the four characters are non-printable. So for strange files, like invisible icon files, the file type field may look all screwy.

Going up and to the right comes the sample rate. This is how many digital samples were taken of the analog sound signal per second. kHz = kilohertz = thousand samples per second. Lowering this value makes the sound become slower and deeper and makes the playing time longer, while raising it makes it higher-pitched and makes the playing time shorter. The sampling rate can be no higher than 65 kHz. Changing the value in this box simply changes the rate at which the sound will be played back. It does not affect the sound data in any way. Also, note that when you enter any of the values 5, 7, 11, 22, or 44, SndSampler will set the sample rate to Apple's pre-defined values of 5563.6363, 7418.1818, 11,127.27273, 22,254.54545, and 44,100 Hz, respectively. (For more information on how to change the sample rate, see the Header and Resample sections below.) Note that when you change the number displayed here, it will be immediately changed in the sound. Other changes, such as the format, don't take effect until the sound is saved.

Below the sample rate is the bit resolution and number of channels. Below this is a pop-up menu which displays different compression options. Some of the compression options will be disabled depending on the format of your sound. For example, only 8-bit mono sounds can be compressed using the MACE algorithms. Only 16-bit sounds can be compressed using the μ -law algorithm.

Finally, below the compression menu, SndSampler shows you what version of the Sound Manager (SM) and QuickTime (QT) you are running.

File Menu

New

This function allows you to record a new sound. First, you will see the input device dialog, which offers you a number of options which depend on the chosen input device. SndSampler includes special code to make the Pro Audio Spectrum PAS16 sound card work like it should. However, if the input device dialog does not seem to adequately display or understand the features offered by the chosen input device, click Bypass. When this button is pressed, SndSampler will not attempt to alter the input device parameters and the input device will remain in its default state. **Normally, though, you should click OK.**

One particular option of interest is the "Record directly to file" check box. If you check this box, then you will have to specify an AIFF file to which the sound input device will write the sound data. This gives you the ability to record more sound data than you can store in RAM. When recording to the file is finished, SndSampler will attempt to open it. It may not be able to if you have recorded more data than can fit in RAM. When recording really huge sounds that won't fit in the available RAM, we recommend using the Segment and Join items to edit them (see below).

After the input device dialog is gone, you will see the sound recording dialog, with three buttons marked in tape recorder fashion. The far left button with the red circle is the record button. Press the record button to begin recording. The middle button is the pause button. Press the pause button once to pause recording, and press it once again to resume recording. The far right button is the stop button. Press this when you are finished recording. If you are recording directly to a file, there will be a 3 second pause before you can click the OK button. This is to allow Apple's Sound Manager time to finish writing all the sound data to disk before proceeding. In addition, the input level is shown directly above the buttons. Only one input level meter will be active if you are recording mono sound. Keyboard shortcuts: press the space bar to begin recording; once recording has commenced, press the space bar to pause and re-start recording; and press the escape key or return or enter to stop recording. The up and down arrow keys will increase and decrease the input gain, respectively. The left arrow key will set the input gain to its minimum value, while the right arrow key will set it to the maximum. Pressing the "1" key will set the gain to 1.00. Note that the recording dialog is semi-modeless, meaning that you can move it around and also switch applications when this dialog is up.

One more thing: be sure to turn off speech recognition before attempting to record a sound.

Open

Allows you to open a System 7 sound file or an AIFF file. If there is a problem and SndSampler cannot interpret the sound, you will have the option of converting it to a raw audio file. This may allow you to recover the sound data. Note that sometimes when you download an AIFF from the internet, the file type will not be set correctly. If the file type is not set correctly, the AIFF will not show up in the file dialog if you try to open it. What you should do in this case is simply drop the file onto SndSampler's icon. When you do this, SndSampler will know that the file is an AIFF, and will be able to open it (or play it, if you hold down the shift key). SndSampler will also set the file type to AIFF if it can. Actually, then, it is better to drop files onto SndSampler to open them, especially if you think there might be a file type problem.

Extract

This handy-dandy little gem allows you to steal the 'snd ' resources out of any file or application. When you select a file which contains at least one sound, you are presented with a dialog box that allows you to choose which resources to load. The resources are listed in the order of ascending resource ID. Sounds can be selected from the list using standard Macintosh list conventions. You click on a sound to select it, and you shift-click to select a range of sounds. Use cmd-click to select and deselect single (possibly non-contiguous) sounds. The arrow keys will move the selection up or down one position, while shift-arrow will extend the selection range up or down.

The dialog box features a Play button, which allows you to listen to a sound before loading it. To play the sound, SndSampler calls the Mac OS routine `SndStartFilePlay`. Unfortunately, this routine is a little buggy, and will not always play a 'snd ' resource from disk. If the sound is corrupted, it may even crash your Mac. Also, if your Mac doesn't support play from disk, the Play button will be disabled. Your Mac must have the Apple Sound Chip to support play from disk. Another reason the button might be disabled is if there isn't enough free memory to play the sound from disk.

Also, as a convenience (i.e. so you don't have to go to ResEdit), you are allowed to delete sounds from the file's resource fork. Be careful when deleting sounds, however, as the action cannot be undone, and the sounds will be gone forever.

Incidentally, the keyboard shortcuts work a little differently here than elsewhere in the program. Pressing the space bar is the same as clicking the Play button. Then, if you press the space bar while a sound is playing, the sound will stop. If you press the escape key, the results will be the same as if you had clicked the Cancel button. The same goes for cmd-period: it's as if you clicked the Cancel button. Any sounds which are playing will of course stop when you leave this dialog.

The Extract dialog will also come up when you drop a file that isn't a System 7 sound file or AIFF onto SndSampler's icon. Assuming that there are any 'snd ' resources in the files you have dropped onto SndSampler, the Extract dialog will come up for each file that has at least one 'snd ' resource in it.

Save (selected portion)

This function will save the selected portion of the sound in the current file (other programs might call this "Save Selection," but we don't because in SndSampler you *always* save only the selected portion of the sound). The current file is the file you last opened the sound from or saved the sound as, and you can see which file this is by choosing the Path item (see above). If the file is an AIFF, the sound will be saved as an AIFF. The Save function does not ask you to specify a filename or a folder. It just goes ahead and saves your sound.

Note that because saving the sound may change it, you are allowed to Undo the Save. When you Undo the Save, though, the contents of the file you saved in are not changed. The home file, i.e. the currently active sound file, will still hold whatever sound data you saved into it. If you Undo a Save, only the sound itself will be restored to its state immediately previous to the Save. Also, when you Redo the Save, the file contents are not altered. Only the sound resource itself is changed back to the way it was immediately after the Save.

Sound resources can be a maximum of 16 megabytes and AIFF files can be a maximum of 2 gigabytes. 2 gigs is a limit imposed by the Mac OS (it's the maximum positive value represented by a 32-bit integer, for you propeller heads out there). That should be more than enough, though: it's equivalent to over three hours of CD-quality stereo sound, which is much more than you should ever need to store in a single file.

Save As (selected portion)

Saves the selected portion of the sound as either a System 7 Sound File or an AIFF file. The specified file becomes the current file. You may Undo a Save As. Note that the previous home file becomes the new home file. So if you open file X, change it, and Save As file Y, when you choose Undo Save, the home file will revert to X. File Y will still exist, and will still hold the data you saved into it.

Save In (selected portion)

This function allows you to save the selected portion of the sound in another file's resource fork. When the Save In item is selected, you will have the option of giving the 'snd ' resource a new name and changing the resource ID. Then you choose which file or application in which you want to save the sound. If there is already a sound there with the same ID, SndSampler will let

you decide if you want to replace that sound. If you don't want SndSampler to replace a resource with the same ID, then it will choose a new ID at random which is not used by any other 'snd ' resource in the specified file.

Revert To Saved

Restores the sound to its state immediately following the last Save or Save as. Note that if you have never saved the sound, e.g. you pasted in a whole new sound from the clipboard, you won't be able to Revert because there is no original file to revert to.

Import

The Import function allows you to convert certain files (WAVE, Sun/NeXT .au, raw audio, QuickTime movies, and CD audio tracks) to AIFFs, which SndSampler will then open. SndSampler uses its own internal routines for converting WAVE, Sun, and raw audio files. QuickTime is used to convert movies and CD tracks, so you must have QuickTime 1.6 or later installed in your System Folder in order to do this. QuickTime also converts WAVE and .au files, but we have found it to be slightly flaky when doing this (especially versions of QuickTime earlier than 2.5). However, if SndSampler can't decipher a file, it lets QuickTime give it a try. This is in case Apple ever gets around to fixing QuickTime or expands its importing capabilities sometime in the future. QuickTime's importing capabilities will *not* be invoked during batch import sessions (this of course is excepted when you are importing movies).

Importing Raw Audio

Any file of any type can be imported as raw audio. One of the neat things about importing a file as raw audio is that you can sometimes recover the sound data from a corrupted file. When importing raw audio, SndSampler assumes that all the data in the file is made up of pure sound samples, without any header information. This means that files with corrupted headers can still be opened. You are required to tell SndSampler what format the sound data is in. SndSampler presents the data format options intelligently, so that choices which don't matter or are illegal under certain conditions are dimmed. For example, if you are loading 8-bit uncompressed samples, it doesn't matter if the file came from a Motorola or an Intel machine, so the little-endian and big-endian options are disabled. If you're not sure what format the data is in, try different combinations until you get something that sounds right. You'll have to guess at the sample rate of the sound data, but that shouldn't be too hard. Also, there will probably be a little blip at the beginning of the sound data, which corresponds to the header information. Just move the start pointer past the blippy stuff.

Of special interest will be the byte offset choice. This is provided for 16-bit data. Usually you will want to choose 0 here. However, with 16-bit data it might happen that the file has been corrupted in such a way that there is an odd number of bytes before the sound data begins. In this case, if you choose a 0-byte offset, the 16-bit sound data will be wrong: 8 bits of it will come from the first sample, and 8 bits will come from the next sample. So if you can't seem to import a 16-bit raw audio file correctly, try changing this to 1.

Guide to sound formats:

8-bit System 7 sound file: 8 bits, unsigned.

16-bit System 7 sound file: 16 bits, signed, little endian.

8-bit AIFF: 8 bits, signed.

16-bit AIFF: 16 bits, signed, little endian.

8-bit WAVE: 8 bits, unsigned.

16-bit WAVE: 16 bits, signed, big endian.

8-bit Sun/NeXT: 8 bits, signed.

16-bit Sun/NeXT: 16 bits, signed, little endian.

Export (selected portion)

This item allows you to take the selected portion of the active sound and export it into a foreign file format. The active sound is not altered in any way. Exporting is **not** the same as saving. For example, if the start and end pointers are not at the beginning and end of the sound, only the exported sound will be shortened. You should always save before exporting.

Exporting Raw Audio

Exporting raw audio is similar to importing raw audio in that you have a number of options regarding the data format. Some rules of thumb: in general, 8-bit data should be unsigned and 16-bit should be signed. The options are handled intelligently by SndSampler, just as when you are importing raw audio. SndSampler gives the raw audio file the type TEXT, which means you could open it in a word processor if you wanted.

Quit

Quits the application. When quitting, SndSampler will attempt to delete the undo files. It will also attempt to copy the application's clipboard sound to the Finder's clipboard. If the clipboard sound is dangerously large, you will get a warning alert which allows you to tell SndSampler not to copy the sound to the Finder's clipboard (unless you turn off the warning with the Clipboard option). Note that unless you allow the warning dialog, the clipboard sound will be copied to the

Finder's clipboard when you quit, even if you choose not to have the clipboard copied on an application switch. The only way to stop it from being copied when you quit is to allow the warning dialog, which then lets you choose not to copy to the Finder's clipboard.

Edit Menu

Undo

The Undo function allows you to restore the sound to its state just prior to the last undo-able operation (most operations are undo-able). The method behind this function involves saving the entire sound in the data fork of one of two special invisible undo files (imaginatively titled "SndSampler Undo 1" and "SndSampler Undo 2" and located in the Preferences folder of your System Folder) whenever an undo-able operation is performed. Then, when Undo is chosen, the sound will revert to its state just before the last undo-able operation was performed. After you Undo an operation, you have the ability to Redo it. The text of the Undo menu choice will change from "Undo xxx" to "Redo xxx."

If you don't like the way the Undo operation works, you can turn it off. Look at the Options menu for more details. Why would you want to turn it off? Well, if you're working with really huge files, it might take a lot of time to save the whole sound to the undo file. In this case you might want to turn Undo off. In fact, when you start working with really big sounds, the time spent saving the undo data will probably be much longer than the time taken to do whatever operation is being performed on the sound. Why does SndSampler save the whole sound, you ask? Why doesn't it just save the part which is being changed? A good question. The simple and truthful answer is that it's much easier to program it to simply save the whole sound. We may change this in the future.

What should you do if there has been some kind of catastrophic error and SndSampler can't undo? One thing you could try is to import the undo file as raw audio. Locate the undo files in the Preferences folder, then import them as raw audio using the unsigned, no compression options for 8-bit and the signed, no compression, Mac format options for 16-bit. SndSampler will attempt to convert these files to AIFFs, which you can then try to open. There will be some residual garbage at the start of the sound data (this is due to the header information in the undo file) which you can simply delete.

Cut, Copy (selected portion)

These two functions copy the selected portion of the sound to the application's clipboard. This sound data can then later be inserted or mixed at a different point in the same sound, or in another sound altogether. The clipboard sound data can also be transferred to the Finder's

clipboard, which will then make it available to other applications which deal with sound resources, such as ResEdit. Cut and copied sounds are always format 1, and are always uncompressed. So if for example you open a 3:1 compressed format 2 sound, then copy it, then delete the entire sound, and then paste the copied sound, it will be format 1 and uncompressed. To make it format 2, 3:1 compressed again, just click the appropriate buttons in the stat window and then save (since it won't get compressed and the format won't be changed until you save). Another point of interest is that the sound header of the copied sound may be different than that of the original sound. The cut or copied sound will always have the minimum required header, i.e. if you open an 8-bit mono sound which has an extended sound header, the cut or copied version of the sound will have a standard sound header.

Please note the following: using your clipboard sound in another application will require that the sound be copied to the Finder's clipboard. This will occur when you switch applications, and also when you quit SndSampler (unless you specify otherwise--see Options below). The Finder's clipboard is not intended to hold a large amount of data, however, and the Mac's system software is not prepared to deal with a large clipboard. In particular, **if you copy a large sound (or large anything) to the Finder's clipboard, small applications may crash after launch.** For example, if you copy a 500K sound to the Finder's clipboard, and then try to launch an application whose memory allocation is set to around 500K, the application may crash, or even crash the system. SndSampler will warn you when you are copying a dangerously large sound (>200K) to the Finder's clipboard. You can use the clipboard options to select if and when you want to copy the application's clipboard to the Finder's clipboard. Note that you can clear the Finder's clipboard with the appropriate choice from the Options menu. (Programmer's note: the Clear Finder's Clipboard function simply calls ZeroScrap.)

As of version 2.4, SndSampler can no longer warn you about copying large sounds when you simply switch applications. It will only warn you when you quit.

Paste (selected portion)

This function should actually be called Replace, but in the interest of compatibility with other applications it has been named Paste. The selected portion of the sound is replaced with the clipboard sound. If there is no existing sound, this function will take the clipboard sound and make it the existing sound. The clipboard sound remains unchanged.

Paste New

Takes the clipboard sound and makes it into a new, editable sound. The clipboard sound remains unchanged.

Select All

Will select the entire sound.

Make Selection New

Takes the selected portion of the currently active sound and makes it into a new sound, which will in turn become the current sound. The formerly active sound will not be affected.

Delete (selected portion)

Will delete the selected portion of the sound. It's the same as pressing the delete key.

Insert Start, Insert End

These functions will insert the clipboard sound at either the start or end points of the existing sound, respectively.

Mix Start, Mix End

These functions will mix the clipboard sound with the existing sound at either the start or end points of the existing sound, respectively. The resulting sound will be lengthened if necessary.

Play Clipboard Sound

Lets you listen to the clipboard sound, in case you've forgotten what it is. You can't do anything else while the clipboard sound is playing. Press the escape key or cmd-period to stop the clipboard sound from playing. Clicking the mouse will also stop playback. Keyboard shortcut: option-space.

Clear Clipboard

Clears the application's clipboard. This will make more memory available for other operations, because any sound in the application's clipboard is stored in RAM. If you are getting inexplicable "out of memory" problems, try clearing the clipboard. Sometimes there can be something in there you didn't expect, since SndSampler will load any sound in the Finder's clipboard when it starts up.

Sound Menu

Play

Plays the selected portion of the existing sound. You have to wait until the sound is finished until you can do anything else. You can stop the sound at any time by pressing the escape key, hitting

cmd-period, or clicking the mouse. You can also play the sound by pressing the space bar. If you press the space bar while the sound is playing, the sound will start over again from the beginning. A thin moving line will show you the progress of the sound when it is playing.

Play Unselected

Plays the unselected portion of the sound. This is handy if you're trying to delete a specific portion of a sample and you want to see what it will sound like when you've deleted it.

Play File

Plays a sound file of your choosing. The file must be an AIFF or a System 7 sound file. The sound will not be loaded into RAM, which is good for playing a sound which is too big to fit into available RAM. Currently loaded sounds will not be affected. While the file is being played, a small dialog will pop up which shows some information about the sound in the file. The dialog will be up for a minimum of 3 seconds, even if the sound is shorter than that. Note that you will not get to see all the information about a System 7 sound file unless you have Sound Manager 3.0 or later installed.

A shortcut for this menu item is cmd-space bar. You can also play a series of files by holding down the shift key and then dropping the files onto the SndSampler icon. Be sure to hold down the shift key until the files begin playing.

Amplitude (selected portion)

Allows you to multiply or divide the amplitude of the selected portion of the sound by an amount that you type in. Use whole numbers with no decimal points for the least distortion. You can use decimal points if you want, however. The amplitude function allows you to save three of your favorite parameters. The first two choices in the submenu are predefined to "times 2" (cmd-2) and "divide by 2" (cmd-1, think "1/2" to remember it). The next three are yours to define. You can change the names of these user-definable functions to make it easy to pick the one you want from the submenu. The first user-definable function has a command key (cmd-G, think "Gain" to remember it) associated with it. You'll probably want to put your most often-used amplitude factor here, so you can invoke it with a quick keypress.

Normalize (selected portion)

This function will multiply your sound's amplitude by the largest value it can without causing clipping. You can choose to only multiply by integer values (for no distortion) or fractional values (which will give some distortion but which will better increase the amplitude). The multiplication factor will be shown before the sound is redrawn. If the sound's amplitude could

not be increased, you will hear a system beep and the sound will not be altered. Note that this function only works on the selected portion of the sound.

Fade In (selected portion)

Ramps the sound's relative amplitude from 0 to 1 over the selected portion. So, if you want to fade in the first 200 samples of the sound, set the start pointer at 0, the end pointer at 200, and choose Fade In. Then you can restore the end pointer to its previous position. Note that you should **select only the portion that you want to fade**. If you leave the entire sound selected, then the entire sound will be faded. This may take some getting used to. The command key for this is "U," for "Up."

Fade Out (selected portion)

Ramps the sound's relative amplitude from 1 to 0 over the selected portion. Otherwise, works identically to Fade In. This one's command key is "D," for "Down."

Dynamic Fade (selected portion)

Allows you to change the amplitude of the selected portion of the sound in a linearly changing fashion. This function uses the same interface as Pitch Bend and Dynamic Pan (see below). You can ramp the amplitude anywhere from 0% to 200% between any two points.

Zero Endpoints (selected portion)

This function reduces to zero volume the first and last sample of the selected portion of the sound in order to help eliminate clicks and pops. Then, to help even more, it takes the next respective sample and adjusts it so the sound's first (or last) three samples form a linear ramp to zero. This function won't get rid of the click the Mac sometimes makes when turning on or off the sound circuitry. A good way to see if there is a click at the beginning or end of your sound is to play it several times in quick succession (pressing the space bar while the sound is playing will start it over from the beginning). If the click isn't always there, then you know it is not due to your sound. Another good way to get rid of clicks is to do a very small fade in at the beginning and a very small fade out at the end.

Extend (whole sound)

Allows you to add zero-amplitude samples to the end or the beginning of the sound. Note that the extra samples will be added at the *very* beginning or *very* end, *not* where you have set the start and/or end pointers. Times (and frames, although they will be truncated) can be entered with decimal points.

Header

Allows you to edit certain information in the 'snd ' resource header. In particular, it allows you to enter a more exacting sample rate than in the sample rate box on the main window. You may enter a number which includes one decimal point. Note that this will only change the rate at which the sound is played back. The sound data is not altered in any way.

Standard sound headers are allowed only for 8-bit mono sounds, while Extended and Compressed headers will support 16-bit and/or stereo sounds, although a Compressed header is required for a compressed sound. Changes to the header type will take effect immediately. Note that when a sound is copied to the clipboard, a new header is created for it. This header will be whichever is most appropriate for the sound data. In other words, if an 8-bit mono sound is copied to the clipboard, it will have a Standard sound header. If any other type of sound is copied, it will have an Extended sound header (it won't be a Compressed sound header because clipboard sound data is never compressed). Then, if this sound later replaces the whole existing sound, the new header type will reflect whatever it was in the clipboard. Other operations do not change the sound header. For example, if a 16-bit sound is downsized to 8 bits, the header will not change from Extended to Standard: it will remain Extended. Conversely, though, if you upsize an 8-bit sound with a Standard sound header, the header must become Extended, since Standard sound headers do not support 16-bit sound data. Confused? You're not alone.

Note: loop points are to be specified in **frames**, not bytes (Inside Macintosh is *wrong!*). Also, base frequencies can be MIDI note values or numbers. **SndSampler does not determine the base frequency for you.** When you make a new recording, SndSampler will insert the arbitrary value of C3 as the base frequency.

Loop Points

Choosing this function brings up a dialog which will help you to select a sample's loop points. The loop points are the sample frame numbers which a synthesizer, or Apple's Sound Manager, will loop around when playing a sample for a longer period of time than there is sound data. In other words, if you have a 1 second sample, and your synth wants to play it for 3 seconds, it will repeatedly play the sound data between the two loop points until your sample has been played for a total of 3 seconds. When you click the Test Loop Points button, SndSampler commands Apple's Sound Manager to play the sample for twice as long as its normal duration using the loop points you have specified.

SndSampler lets you zoom in and out on the sound so you can more easily see where to place the loop points. Note that the highest level of zoom corresponds to a 1:1 frame-to-pixel ratio, i.e. every frame of the sample is displayed, with one frame to one screen pixel. A general

rule of thumb is that you should set the loop points to be somewhere in the middle of the sound where the amplitude doesn't change much. Also, it is a good idea to make the width of the loop points be equal to one pitch period (pitch period in frames = sample rate / sample pitch, i.e. $44,100 \text{ Hz} / 440 \text{ Hz} = 100$ frames for a concert A note sampled at 44 kHz). If you don't know that sample's pitch, you will have to guess by looking at the sound wave itself. This is where the zoom capability can come in handy. Oh yeah: only the left frames of a stereo sound are shown.

Resource

This function allows you to alter the sound format, the sound resource ID, the sound resource attributes, and also to modify the sound channel initialization option and the sound command. For more information on initialization options and sound commands, see Inside Macintosh: Sound. Most users will only want to change the ID and won't need to adjust the other parameters. For maximum robustness, we recommend that you always set the initialization option to zero (Clear) and the sound command to a bufferCmd. The resource ID should probably be somewhere from 128 to 32767. Don't set the resource ID to zero. The Mac OS will not allow you to double-click on a sound resource file whose resource has an ID of zero.

Hex Editor

The hex editor allows you to alter any byte in the sound resource. **BE CAREFUL!** You can easily muck things up very badly if you don't know what you're doing. Use the delete key to delete a data byte, and use the tab key to insert a new data byte immediately prior to the selected byte. The first actual sound sample will be shown in bold. If you are editing a stereo sound, this will be the left half of the stereo frame. Note that for a 16-bit sound, each sample is two bytes, so the first and second bytes will be shown in bold. You can go directly to the start and end selections by pressing space and cmd-space, respectively. You can create new start and end selection points by pressing s and cmd-s, respectively. This could be useful if, say, you wanted to set the selection end just after the first zero-valued sample (80 for 8-bit sounds, 00 00 for 16-bit) near the end of the sound.

Tone Generator

This one is pretty self-explanatory. Using this function you can create a sine wave, square wave, or sawtooth wave at a given frequency, amplitude, and duration. Numbers note: the frequency and sample rate can be entered as fractional values, but the duration and amplitude must be whole integers. The amplitude must be less than or equal to 100%, and the duration must be less than the maximum duration as shown in the dialog. There are a number of other options as well, but figuring out what those do is left as an exercise for the reader.

FX Menu

Echo (selected portion)

This function gives your sound depth by adding a fading echo effect through an FIR (Finite Impulse Response) filter. This function operates only on the selected portion of the sound. Also, the Echo effect allows you to customize parameters. You can choose values for both the decay constant k and the echo delay (in milliseconds or sample frames). The decay constant k is taken from e^{-kn} , which is the amplitude factor of echo number n . The total number of echoes N is chosen by the application such that echo number N will have an amplitude factor which will reduce the sound amplitude to nearly zero. If you're not mathematically inclined, just remember this: increasing k will cause your echoes to die away faster, while decreasing k will give you more echoes and they will last longer. The delay indicates how far apart each echo will be. If the Auto Extend box is checked, then SndSampler will increase the length of your sound so that you don't lose any of the echoes. SndSampler will generally make your sound REALLY big when doing echoes, and this will require more free RAM than otherwise.

Note that the best way to figure out exactly what the user-adjustable parameters can do (for this and the other fx) is to play around with them. When you do find some effects that you like, you can save them. SndSampler gives you several slots to fill with your own favorite echo effect parameters. You can change the names of these user-definable echo effects to make them easier to select from the submenu. For example, if you find an effect which makes the sound sort of robotic, you can type "Robotic echo" in the name box. Then the Echo submenu will show the name "Robotic echo." You can do the same thing with Reverb, Chorus, and Flange (see below). The first user-definable function has a command key (this one is cmd-E, for "Echo") associated with it. You can put your most often-used effect here, so that you can invoke it with a quick keypress.

Note also that with this and the following three fx, decay, delay, and speed values can be entered as decimal numbers. With Chorus, however, the number of voices must be a whole number.

Reverb (selected portion)

The Reverb function imposes an orchestra hall type of echo on the selected portion of your sound. You can choose to use the allpass filter, which directly implements an allpass reverberation IIR (Infinite Impulse Response) filter of the equation

$$y(n) = -g * x(n) + [1 - g^2] * [x(n - D) + y(n - D)]$$

where y = output, x = input, g = gain factor (< 1), n = current sample number, and D = delay in samples. Or, you can choose the comb filter, which has the equation

$$y(n) = x(n - D) + g * y(n - D).$$

You can specify both g and D (D can be in samples or milliseconds). The gain factor g controls the ratio of direct sound to reverberated sound. Larger g = more direct sound, less reverberation. You will probably want to make g smaller for a comb filter effect, since it can be quite loud and harsh. Note that g can be negative.

If you check Auto Extend, SndSampler will make your sound longer if necessary in order to not lose any reverb, and will increase the length of the selected portion of the sound to include the delayed samples. Note that the larger the selected portion of the sound, the more free memory required to make the reverb effect. More complicated reverbs can be created by using different parameters and then mixing together the results. For example, commercial reverb generators will commonly do three or four comb filters in parallel, mix this result together, and then do an allpass reverb on the mixed data. (And they'll probably add a little bit of echo, too.)

Special note: SndSampler will *always* extend your reverb sound by at least the length of the delay, regardless of whether or not you have selected Auto Extend. However, sharp-eyed users may notice that if you choose a comb filter and don't check Auto Extend, the sound does not appear to increase in length. The reasons for this are complicated, and if you are really interested in why this happens, write to us. Rest assured, however, that SndSampler is not making a mistake when this happens. It is *not* a bug! (So don't try to get free registration for reporting it.)

The command key for the first Reverb effect is cmd-R (for "**R**everb").

Canceling FX

Some fx, like Reverb, will show you a progress report in the lower left corner of the sample window. Any procedure that does this can be canceled by pressing and holding the escape key. Beware: when you do this, it is possible that the sound may become corrupted. If this happens, you'll probably want to Undo immediately. Actually, you'll probably want to Undo immediately regardless, just in case the sound is corrupted but you can't tell right away.

Chorus (selected portion)

Chorus adds a delayed version of the selected portion of the current sound to the current sound to make it sound like there is more than one voice making the sound. You can choose the number of

voices (delayed versions) to add to the original sound. Note that as you add more sounds, the resulting sound will tend to get louder. (The process will also become really s-l-o-w as you add more voices. Super slow for more than four.) The delay value is a starting point for how much the other voices will be delayed. The speed factor determines how fast the variable delay will vary. You can type any number here, including numbers as high as 100,000 and as small as 0.0001 (although those numbers give lousy results). Feel free to experiment with the settings. Note that 100% is the default setting for the speed factor. The Auto Extend function works as above. Note also that a lot of the Chorus parameters as determined by SndSampler are random! So if you don't like what happens to your sound with Chorus, try it again. The result will be different each time.

The command key for the first Chorus effect is cmd-K (think "**K**horus"; cmd-C was already taken by Copy.)

Flange (selected portion)

This function is very similar to Chorus. Flange perpetrates a kind of phasey, in-and-out effect on the selected portion your sound. The gain factor represents the relative amplitudes of the original and delayed sounds which are adjusted so that they add up to 1. In other words, if you type 0.6 here, the delayed sound will be multiplied by 0.6, and the original sound will be multiplied by 0.4. If you check the auto-extend box, SndSampler will continue to flange past the selection end in order to make the flanged portion join up nicely with the non-flanged portion. Also, you might notice a decrease in volume after a flange, but you can fix this with a slight increase in amplitude, like 1.2 or so.

The command key for the first Flange effect is cmd-F (for "**F**lange").

Backwards (selected portion)

Reverses the selected portion of the sound so that it plays backwards. This function can produce interesting effects when combined with reverb or echo. (See the cool effects listing below.)

Stereo To Mono (whole sound)

Makes a mono sound out of a stereo sound. Operates on the entire sound.

Mono To Stereo (whole sound)

Changes a mono sound to a stereo sound. You can pan by moving the slider left and right. This function splits the sound into two parts, so there will probably be a noticeable decrease in volume. Operates on the entire sound.

Dynamic Pan (selected portion, but whole sound becomes stereo)

Allows you to dynamically change the stereo field of a monophonic sound. When choosing Dynamic Pan you will be presented with a dialog which displays the selected portion of the sound. (This is similar to the Pitch Bend effect as described below.) Below that is a box with a line drawn between two points at either end. By clicking somewhere in the box, you can make new points. Each point represents a certain position in the stereo field. Left is at the top and is represented by +1. Right is at the bottom and is represented by -1. The sound will be altered so that its position in the stereo field appears to move linearly between any two points you have made. You can click in an existing point to drag it to a new position, or you can enter numbers directly in the text boxes at the bottom of the dialog. Cmd-delete will delete the currently selected point (the one which is white). Cmd-left arrow and cmd-right arrow will move among the points. (If you just type an unmodified delete or arrow you will move the cursor in the active text box.) Note that the two end points can only be moved up or down, and you can't put any points at an equal horizontal position with the end points. Also, you can't delete the two end points.

Even if you haven't selected the entire sound, the portions of the sound that are not selected will still be changed to stereo. The portion which precedes the start pointer will have a stereo value which is identical to that of the first point. The portion of the sound which follows the end pointer will have the same stereo value as that of the last point.

For maximum speed, select only the portion of the sound you want to dynamically pan, and only use two points. Three or more points will slow it down dramatically.

Resample (whole sound)

This function allows you to change the sample rate of a sound while not changing the way it sounds when played back. The Resample function uses linear interpolation to calculate new values for each individual sample so that when the whole sound is played back, it still sounds like it did before. The new sound will take up more or less disk space, depending on whether the new sample rate is higher or lower than the original sample rate. Resampling affects the entire sound, regardless of the locations of the start and/or end pointers. Remember that the farther away you get from the original sample rate, the more distortion there will be in the resulting sound. A good rule of thumb is to only resample a sound once. For instance, don't resample from 44 kHz to 22 kHz and then decide you want to resample from 22 kHz to 34 kHz. Just go directly from 44 kHz to 34 kHz.

A note about sample rates on the Macintosh

Apple specifies sample rates in sound ('snd ') resources as fixed-point binary numbers. A fixed-point number on the Macintosh only has 32 bits of resolution. Sample rates such as 22,254.54545... Hz (Apple's standard 22 kHz rate) can only be approximated with a fixed-point number. For most applications, this won't cause a problem. However, if you are upsampling by a factor of 2 from Apple's `rate11kHz` (which is actually 11,127.27272 Hz, as represented by the fixed-point value \$2B77.45D1) you will not come up with Apple's `rate22kHz`. Multiplying `rate11kHz` by 2 results in a rate which is 1/65,536 shy of `rate22kHz` (which is actually 22,254.54546 Hz, as represented by the fixed-point value \$56EE.8BA3). `SndSampler` recognizes this, and when upsampling from `rate11kHz` by a factor of two it will force the resulting sample rate to be `rate22kHz`. Conversely, when downsampling by a factor of 2 from `rate22kHz`, `SndSampler` will force the new sample rate to be `rate11kHz`. `SndSampler` doesn't do anything special with `rate5kHz` or `rate7kHz` because these are not used very much. By far the most popular sample rates are 11 and 22 kHz.

Note that typing 11 or 22 (etc.) in the Sample Rate box in the main window will force the playback rate to be `rate11kHz` and `rate22kHz`, respectively. This will not change the sound data, but will only change the rate at which Apple's Sound Manager (part of your system software) plays back the sound.

Downsample (whole sound)

Resamples the sound to half its original sample rate, e.g. 22 kHz -> 11 kHz, or 44 kHz -> 22 kHz. This is a quick-and-dirty operation which simply discards every other sample frame. It is written in pure 68K assembly language and is lightning-fast, even on a Mac Plus. Another bonus for this routine is that you will *always* have enough RAM to use it. Guaranteed!

Pitch Bend (selected portion)

Allows you to dynamically change the pitch of the sound. When choosing Pitch Bend you will be presented with a dialog which displays the selected portion of the sound. Below that is a box with a line drawn between two points at either end. By clicking somewhere in the box, you can make new points. Each point represents a certain pitch change away from the current pitch. The pitch change values are displayed in semitones, with +12 semitones being one octave higher and -12 semitones being one octave lower. Each note on a standard musical scale can be thought of as being a certain number of semitones away from the other notes. For example: C + 1 semitone = C sharp (#) = D flat (b), C# + 1 semitone = D, D + 1 semitone = D# = Eb, D# + 1 semitone = E, E + 1 semitone = F, F + 1 semitone = F#, F# + 1 = G, G + 1 = G#, G# + 1 = A, A + 1 = A# = Bb, A# + 1 = B, and B + 1 = C. (There are no sharps between E and F or between B and C.)

The pitch of the sound will change linearly between each point you create. You can click in an existing point to drag it to a new position, or you can enter numbers directly in the text boxes at the bottom of the dialog. Cmd-delete will delete the currently selected point (the one which is white). Cmd-left arrow and cmd-right arrow will move among the points. (If you just type an unmodified delete or arrow you will move the cursor in the active text box.) Note that the two end points can only be moved up or down, and you can't put any points at an equal horizontal position with the end points. Also, the pitch change must remain between +12 and -12. What SndSampler actually does is change the sample rate to simulate a pitch shift, so that when the pitch goes higher it will sound more chirpy and when it goes lower it will sound slowed-down and play for a longer time. The Pitch Bend effect also requires lots of RAM. There may need to be as much free RAM in SndSampler's partition as 3 times the size of the selected portion of the sound. Pitch Bend can be kind of slow, too.

16-bit -> 8-bit (whole sound)

This should be fairly self-explanatory. (Yeah, we know, this isn't really an effect, but the Sound Menu was getting too crowded.)

8-bit -> 16-bit (whole sound)

Opposite of above.

Goodies Menu

Options

Sound Color

Allows you to choose a new color for the sound waveform display.

Sound Click

Allows you to choose either a double-click or an option-(or ctrl or cmd)-click to select the end pointer. If you choose either of these options, you will select the position of the start pointer with a single click. You may also choose the nearest click option, where a single click will move the pointer which is closest to where you clicked. New in version 3.1 is the ability to click and drag to select a portion of the active sound.

Undo

Allows you to turn the Undo feature on or off.

Clipboard

Allows you to tell SndSampler when to copy its own clipboard to the Finder's clipboard. Also, this option will allow you to turn off SndSampler's warnings when the sound you are copying to the Finder's clipboard is dangerously large.

Drag and Drop

Allows you to auto-quit after a batch operation is finished.

Cmd Key Shortcuts

When cmd key shortcuts are on, certain functions will skip the dialog and use their default parameters when you select them by using a menu key (like cmd-E for Echo).

Edit Movie

This item allows you to fiddle with the soundtracks of a QuickTime movie, and will only be enabled if you have QuickTime 1.6 or later installed in your System Folder. Note that you can only play with the soundtracks, not the video tracks or any other type of track. You can't paste or copy any movie-ish type data to or from the movie you are editing, either. Bummer. When you select a movie you will be presented with a dialog which displays the movie at its normal size on the right. On the left will be some statistics, which include the name of the movie, a menu listing all the movie's soundtracks, the format of the sound samples in the currently selected soundtrack, and the soundtrack's volume. Below these are some buttons which let you alter the movie's soundtracks. They function as follows:

Insert: this button will insert a new soundtrack at the movie's current position using the selected portion of the active sound. The sound data will be compressed according to the compression menu in the stat window. If there is no active sound, this button will be dimmed.

Enable/Disable: make various soundtracks active or inactive. Note that **inactive soundtracks will not be saved** to the movie file.

Move To: moves the start of the selected soundtrack to the movie's current position.

Isolate: disables all the soundtracks except the currently selected one.

Pressing the up and down arrow keys will change the volume of the currently selected soundtrack (note that you can overdrive the volume, i.e. make it greater than 1).

Holding down the cmd key while pressing the up or down arrow will change the playback rate of the movie. If the movie has a soundtrack, it will make the soundtrack's playback rate change also. It has the same effect as changing the sample rate of a sound in SndSampler simply by typing in a new value in the appropriate box in the stat window. If you are adding digital audio to a QuickTime music track, you probably will not want to change the playback rate in this manner, because the apparent pitch of the digital audio will be affected while the pitch of the notes in the music track will not change.

You can play and change the position of the movie using the normal MoviePlayer controls: space bar plays the movie, left and right arrow keys change the movie's current position, etc. Of special note is that the return key will play the movie instead of hitting the save button.

Some caveats to editing movies: first, QuickTime can be slightly (!) unstable, so watch out for crashes. Also, when the sound data is inserted, it will not be "flattened" (interleaved), so that after you have added a few soundtracks (or even just one really long soundtrack) you might get funny playback, i.e. the sound skips or the entire track doesn't play. All you can do is go ahead and save, during which procedure the movie data will be "flattened," and see what it turns out like. Usually it will turn out OK. You may want to add a little bit of sound, then save, then add a little bit more sound, then save again, etc. Finally, you should **ALWAYS WORK WITH A COPY OF THE MOVIE**, since when you choose to cancel QuickTime might leave some of the sound data in the file even if it isn't being used. This might cause the movie to bloat out with many megabytes of unused information. Not pretty sight.

Convert to Movie

Allows you to use QuickTime to convert System 7 sound files and AIFFs into movie files. The movie will be saved as a MoviePlayer document. You'll need to have QuickTime 1.6 or later installed in your System Folder in order to use this option.

Segment

This function lets you choose an AIFF file to divide into smaller AIFFs. This can be handy if you want to edit a sound which is larger than RAM (see "How Can I Edit Sounds That Are Larger Than RAM?"). SndSampler will suggest a segment size to you based on the amount of free RAM it currently has. Make the segment less than or equal to this suggestion if you want to be able to edit it later. Each segment becomes a stand-alone AIFF in its own right and can be used for whatever you might want to use an AIFF for. Segment works best, however, when used in conjunction with Join.

Join

Once you are done fooling around with your segmented AIFFs, you can use Join to put them back together. Select the segments in the order in which they were created, i.e. select #1 first, then #2, then #3, etc. When you come to the last segment, select it and then click the Join button (don't click Select). To see if all went well, use Play File to listen to the rejoined AIFF (assuming it's too big to fit in RAM).

Clear Finder Clipboard

Clears the Finder's clipboard (by calling *ZeroScrap*). A convenient way to get rid of a large sound you may have copied to the Finder's clipboard, before it causes a system crash.

Window Menu

Choosing a sound from this menu will make it the active sound. Note that you can cycle through all the open sounds by pressing the return key (assuming there are no active text boxes when you press return).

Network Notes

As of this version, *SndSampler* should not be launched over a network, i.e. the actual application should reside on the hard drive of the computer from which it is launched. At any rate, your registration is for you and you alone and does not permit multiple simultaneous users. Each new user must register his or her own copy.

Answers to Common Questions

Why Doesn't SndSampler Have Lots of Fancy Graphics? (Attention *MacWorld* reviewers!)

SndSampler is intended to be a tool, not a game. It doesn't need a lot of fancy graphics. All it needs to do is get the job done. Many applications will use lots of unnecessary graphics to make the windows and tool bars look pretty. These graphics are disk and RAM hogs. We have left out the graphics to keep the application size as small as possible and also to give you as much free RAM as possible to work with your sounds. Think about it: with *SndSampler* you may actually be able to take advantage of all that extra RAM you bought. What a concept!

Should I Use Virtual Memory To Edit Sounds That Are Larger Than Available RAM?

No. See the next question for an alternative to using virtual memory. If you absolutely feel you *must* use virtual memory, it might be a good idea to turn off Undo (see Goodies -> Options).

How Can I Edit Sounds That Are Larger Than Available RAM Without Using Virtual Memory?

SndSampler provides you with an alternative to virtual memory via Segment and Join (under the Goodies menu). Say for example you have a 40 megabyte CD audio file and you want to fade it in at the beginning and fade it out at the end. Rather than turning on virtual memory, segment the audio file into three sub-files. Make the first segment be the part at the beginning that you want

to fade. Be sure to make it small enough to fit in RAM! Then, make the second segment consist of the rest of the sound up to where you want to fade it out. Finally, the third segment will be what's left, i.e. the part you want to fade out. Now, load the first segment and fade it in to your satisfaction. Then do the same with the third segment. When that's done, join all the segments back together. Ta-da! You have a finished audio file ready to burn on your fancy-schmancy CD-R drive.

Of course, this approach will not work every time. It will only work if you can make the segments small enough to fit into RAM, and if the operation you perform on the sound doesn't depend on the sound data in any other segment. Another example of a situation where Segment and Join will work well would be if you wanted to double the amplitude of a sound file that was twice as big as available RAM. What you would do here would be to divide the file into two segments, multiply the amplitude by 2 for each, then rejoin them.

For yet another example, let's say you wanted to make some CD tracks and you wanted to cross-fade two of them, i.e. one fades up while the other fades down. What you need to do is segment both audio tracks. For track #1, make the first segment contain all the audio data right up to where you want to start fading down. Then make the second segment be the rest of the data, i.e. what you want to fade down. For track #2, you'll do the opposite: make the first segment be what you want to fade up, and make the second segment the rest of the audio data. Now comes the tricky part. What you'll do is open the second segment of track #1 and fade it down. Then you'll open the first segment of track #2 and fade it up. Now copy the faded up data and mix it with the second segment of track #1. So far so good. Now you'll have to pick some place where you want the tracks to separate, i.e. where track #1 becomes track #2 on the CD. When you pick your spot, copy the beginning of track #2 and replace the first segment of track #2 with this sound data. Then save the second segment of track #1 and the first segment of track #2. Finally, join track #1 and track #2 back together from their respective segments. Voila! Track #1 will fade right into track #2. We know this works because we've done it!

An example of where Segment and Join wouldn't work so well would be if you wanted to normalize a file that was twice as big as RAM. If you segmented it into two smaller files and then normalized them, the two segments would probably end up being multiplied by different factors. In other words, the first segment might get multiplied by 2, while the second segment might get multiplied by 3. This is probably not what you want.

On the other hand, there *is* a way to normalize a file bigger than RAM using Segment and Join: it's just a little more complicated. Here's how: make each segment of the sound small enough to fit in RAM. Then do a normalize on each segment, and write down the value it gets multiplied by. After you've done this with all segments, choose the smallest normalization multiplier from your list and then change the amplitude of each segment by this factor. Ta-da!

This is just one specific instance, however. You will have to determine whether or not Segment and Join are appropriate for your needs on a case-by-case basis.

Help! Something's Wrong With My Audio File!

Don't panic (yet). If there's something wrong with your audio file and SndSampler can't read it, you can try importing it as raw audio. When you import as raw audio, SndSampler treats the file as if were composed entirely of sound samples. You then have the responsibility of specifying the format of those sound samples. For example, say you want to import a WAVE file and you know that it is supposed to be 16-bit stereo. But SndSampler can't read it because the header information in the file has been corrupted. What you would do is import the file as raw audio, and choose 16-bit stereo, signed, no compression, byte offset 0 (try 1 if 0 doesn't work), and big-endian byte format. The WAVE will be converted to an AIFF, which will then be opened if you have enough free RAM. There will be some garbage noise at the beginning of the imported sound due to the WAVE header information, but you can easily discard this. If you don't know what format the sound file you're having trouble with is supposed to be, you'll have to try different approaches until you get one that works. There are a lot of possible combinations, however, so you've got your work cut out for you.

What About This Batch Mode Business?

Ah-ha! Glad you asked that. SndSampler has some incredibly cool features which allow you to automatically--without any required input on your part--perform common operations on sounds. Let's say, for instance, that you needed to downsample (reduce the sample rate by half) 100 sound files. Opening, downsampling, and saving each file would be extremely tedious. SndSampler, however, allows you to perform these options automatically, with very little input required from you. What you do is select all the files you want to downsample. Then hold down the cmd key and drop the files onto SndSampler's icon. (Note that you have to keep holding down the cmd key until the batch mode dialog appears.)

After you have done this, you will be presented with a dialog which allows you to choose what batch operation you want to perform. In this case you would choose Downsample. Then you specify a folder in which the modified sounds will be saved. After that, SndSampler will open each file, downsample the sound, and save it to a new file inside the folder you specified earlier. You also get to watch a progress report about which sound is being processed and how far along it is. If you want to cancel the batch mode, press and hold the escape key. If you cancel the batch mode in this way, it is probable that the sound that is currently being worked on will become corrupted. Also, if there is an error and one or more of the sounds cannot be processed (because there isn't enough free RAM, for example), SndSampler will write a short description

of the problem in a file called "Batch Errors" inside the destination folder. If there is a problem with one or more of the batch sounds, try to convert it normally without using batch mode. That way SndSampler will tell you in more detail what is wrong.

Three of the batch operations which in particular deserve discussion are Extract, Insert, and Replace. When you extract, all the 'snd ' resources in all the files which you drop onto SndSampler's icon will be converted into separate System 7 sound files. They will appear inside the folder you specify. When you insert, SndSampler will take all the 'snd ' resources from all the files dropped onto its icon and open the AIFFs which are dropped onto its icon and insert all those sounds into the resource fork of whatever file you choose. (Special note: the inserted 'snd ' resources will have different ID's than they did in their original files, courtesy of the Resource Manager routine `UniqueID`. If you are writing a program which will try to load these 'snd ' resources, you will probably want to refer to them by name, using `GetNamedResource`, rather than by ID.) Inserting an arbitrary resource type will allow you to take all the resources of a given type, such as 'DITL"s or 'LDEF"s, from all the files you drop onto SndSampler and insert them into the resource fork of a file of your choosing. Finally, replacing is nearly the same as inserting, except all the resources ('snd ' or what you have specified) will be deleted before the new resources are inserted. Be careful: SndSampler won't just delete the resources that have the same IDs as the ones you're inserting--it will delete **ALL** the existing resources of the specified type. The resources will be deleted before the batch operation begins, so they will be removed even if the batch operation fails.

If you can't seem to drop files of all types onto SndSampler's icon, you'll need to rebuild the desktop. To rebuild the desktop, restart and hold down the cmd and option keys until you get the "Rebuild desktop?" dialog.

One more thing to note about batch mode is that in most cases, sounds that are not converted will still be saved into the destination folder. This is best illustrated with an example. Say you have 200 sounds you need to normalize. It is likely that some of these sounds will already be at maximum volume and hence will not be altered by the normalization process. However, even if a sound is already at maximum volume, it will still be copied over to the "SndSampler Normalized" destination folder. The reasoning behind this is that when batch mode is finished, you will know that all the sounds in the destination folder are at maximum volume. Otherwise, you would have to check to make sure all 200 sounds were in the destination folder, which is probably not something you want to do. If there is an error and one or more of the sounds cannot be opened, SndSampler will write a short message about the problem in the file called "Batch Errors" in the destination folder.

Finally, **it is a good idea to close all sounds before initiating batch mode**. You get more memory that way, and it's also safer.

New in version 3.4: SndSampler will automatically detect if a file dropped onto its icon is an AIFF, even if the file type is set incorrectly.

How Can I Convert My System 7 Sound Files to AIFFs While Adding ".aiff" To The File Names?

It is absolutely true that the batch function "snd -> AIFF" does not allow you to add a filename extension. This is because it is assumed that, as a Macintosh user, you don't cotton to filename extensions. That's a Windoze thing, after all. However, you are allowed to add a filename extension of your choice when using the batch "Import" function. And, believe it or not, you can batch import System 7 sound files! They will simply be converted to AIFFs (with the filename extension of your choice added to them). You can also import AIFFs. So if you want to add a filename extension to a bunch of AIFFS, just import them.

How Can I Read CD Audio?

Well, there are two ways to do this. One is to simply choose New, and then select the Internal CD option from the Input Source menu of the input device dialog.

A better way, however, is to choose Import and then select one of the CD audio tracks. You will be asked to specify a file to save the audio data to, and you will get a dialog which lets you choose from several options regarding the data format. With this method, the audio data is read directly off the CD. With the other method, the audio data first must pass through your Mac's D/A converter to make analog sound and then through your Mac's A/D converter to make it back into digital sound. This will cause some (hardly noticeable) loss in fidelity.

Special note: we don't recommend that you import CD audio tracks in batch mode. This is because in batch mode you will not see the settings dialog, and the CD track will be converted using the default settings of 8-bit mono, 22 kHz.

How Can I Change the Pitch of a Sound?

As of this version, SndSampler can only simulate a pitch shift by altering the sound's sample rate. The best way to do this is to use Pitch Bend (see above). If you don't want to make the pitch vary with time, but want to keep it at a constant value over the whole sound, just make the first and last Pitch Bend points have the same pitch offset.

Hey! How Come I Can't Export/Import/Segment IMA 4:1 Sounds?

Well, IMA 4:1 compression is kind of funny. (That's funny strange, not funny ha-ha.) First of all, you can't export a sound using IMA 4:1 compression because every computer system implements it differently. That means that Apple's IMA 4:1 is different from Windoze IMA 4:1. So different

that we're not going to bother trying to do the conversion. This is also the reason you can't import IMA 4:1 sounds. Now, the reason you can't segment an IMA 4:1 AIFF has to do with the nature of the ADPCM (Adaptive Differential Pulse-Code Modulation, upon which IMA 4:1 is based) algorithm. With ADPCM it is the *difference* between succeeding samples which is stored, not the samples themselves. So, taking a slightly oversimplified look at the situation, if you segment an IMA AIFF and then go about changing the samples in one segment, the next segment will get all screwed up because the samples in the first segment have changed. Remember, it is the *differences* which are stored in the AIFF file. If you change the last sample in the first segment, then the first sample in the next sample will be wrong, because it expects the previous sample to still be the same. And then every succeeding sample that segment will be wrong. Yeah, it's kind of complicated, so you'll just have to trust us on this one.

Why Don't My Resampled Sounds Sound Good?

There could be many reasons for this. However, you should keep in mind that in most cases your Mac's hardware can only play sounds at one specific sample rate, usually 22.254 kHz. That means if you try to play a sound which has a different sample rate, the Mac OS must resample the sound to 22.254 kHz before playing it. So if you use SndSampler to downsample a sound from 22.254 kHz to 11 kHz, say, then when you play it the Mac OS will upsample it right back to 22.254 kHz. The more times you resample a sound, the more distorted it becomes, and in this example we have resampled it twice before it gets played, which is definitely going to introduce noticeable distortion. For the best sounding sounds, resample to your Mac's hardware rate. You can usually find this somewhere in the Sound control panel. Of course, other Macs out there may not have the same hardware sampling rate as yours, so a rule of thumb is to make the sample rate 22.254 kHz to sound as good as you can on all Macs.

Where Did My Windows Go?

If for some reason your windows disappear (like they might if you change your screen resolution, or if you use two monitors for a while and then get rid of one), you will need to trash the file "SndSampler Prefs" from the Preferences folder inside your System Folder. Then the windows will become itchy-bitsy and will be in the upper left corner of your main monitor (the one with the menu bar). You'll lose all the other stuff you've saved as preferences, but that's the way it goes.

Why Isn't SndSampler PowerPC Native?

Many of SndSampler's crucial byte-manipulation routines are written in highly-optimized 68K assembly language. In order to go native, we'd have to re-write these routines, which would be a

big pain; and doing this would mean having two separate bases of source code, which would be an even bigger pain; and we have an extremely low threshold of pain. Besides, if you upgrade your PowerPC to 7.5.3 or later, you get dynamic recompilation of 68K code, which should be fast enough for all but the most impatient. 'Nuf said?

Why Doesn't This Document Use Smart Quotes?

Because some word processors (like ClarisWorks, for instance) don't seem to be able to understand smart quotes contained in an RTF document created by MS Word. Go figure.

Cool Things To Try

Simple Stereo Chorusing

1. Change a mono sound into a stereo sound with the slider all the way left
2. Copy the new sound
3. Undo mono to stereo
4. Change the mono sound to stereo with the slider all the way right
5. Move the start pointer to around 100 or so
6. Mix start
7. Move start pointer back to the beginning

Shifting Echo

1. Change a mono sound into a stereo sound with the slider all the way left
2. Echo the entire sound with a delay slightly longer than twice the sound's duration (use frames)
3. Move the start pointer to just before the first echo
4. Copy
5. Revert to saved
6. Change to a stereo sound with the slider all the way right
6. Repeat step 2
7. Move the start pointer to just after the original sound
8. Mix start
9. Move the start pointer back to the beginning

Reverse Echo

1. Reverse the entire sound
2. Echo the entire sound
3. Reverse again
4. Cool, huh?

Keyboard Shortcuts

Space bar: plays the existing sound; starts recording a sound; pauses and re-starts recording.

Cmd-space bar: plays a file from disk.

Option-space bar: plays the clipboard sound.

Escape key: hits dialog Cancel button; stops a sound from playing; stops recording; halts any effect which shows percentage completed.

Return or enter key: hits default dialog button; stops recording.

Cmd-period: hits dialog Cancel button; stops a sound from playing.

Delete: deletes the selected portion of the sound if there are no active text boxes.

Left arrow / right arrow: move the start pointer; make input gain minimum / maximum.

Up arrow / down arrow: move the end pointer; increase / decrease input gain; increase / decrease volume of QuickTime soundtrack.

Cmd-left arrow: moves the start pointer to the beginning of the sound; selects next point to the left in pitch bend / dynamic pan dialogs.

Cmd-right arrow: moves the end pointer to the end of the sound; selects next point to the right in pitch bend / dynamic pan dialogs.

Cmd-delete: deletes current point in pitch bend / dynamic pan dialogs.

Tab: moves between text-edit fields in stat window; moves between text-edit fields in most dialogs.

Return: cycles through the list of open sounds.

Cmd: initiates batch mode when dropping files on the SndSampler icon.

1: sets input gain to 1.00 in record dialog.

Shift + menu choice: no dialog is shown for Amplitude, Mono to Stereo, Echo, Reverb, Chorus, and Flange. It is best to hold down the shift key until the action you have chosen starts happening.

Mouse Shortcuts

Drag the triangles to move the start and end pointers.

Single-click on the sound to move the start pointer.

Double-click (or option-click--see the Options menu) to move the end pointer.
Click on any unused portion of the stat window to deactivate an active text box.
When a sound is playing, click once to stop it.

Click on the Sound Length number and it will change from frames to time.

(We made this bold because nobody ever seems to notice it.)

Legal Notice

This product is sold as is. The authors make no claims about its correctness or suitability for a specific purpose or otherwise. The authors accept no responsibility whatsoever for damages of any kind, be they direct, incidental or otherwise, incurred by the user while using, or not using, any version of SndSampler, or any other software we may have written, regardless of what state law or local law or federal law or international law or intergalactic law or any other law may say. If you use this software you are contractually agreeing to these terms. We think that covers everything pretty well, don't you?

Rhapsody, Schmapsody! Hurrah For The Good, Old-Fashioned Mac!

SndSampler was developed mostly in Pascal (sorry folks, but C stands for "crap") on a good, old-fashioned 68030 Mac running good, old-fashioned System 7.1. We only have 8 MB of RAM, too!