AUDIO SAMPLE EDITOR

USER GUIDE

GARRICK M^CFARLANE

BSC (HONS) COMPUTING SCIENCE—FINAL YEAR

DEPARTMENT OF COMPUTER SCIENCE AND APPLIED MATHEMATICS

ASTON UNIVERSITY

TABLE OF CONTENTS

❶*Getting Started*

The Sample Editor is an application that allows you to edit and manipulate sampled sound on the Apple Macintosh. It supports sound input and output from *Audio Interchange* (*AIFF*) files, Macintosh ‘`snd` ’ resource files, and the Clipboard. The sound may be compressed by a factor of up to 6:1. You can also record direct from any connected microphone, limited only by the available memory. The waveform of a sound can be displayed in any level of detail you require by selecting the part you wish to zoom in on. You can edit it using *Cut*, *Copy*, *Paste* and *Clear*, and apply transformations such as *Fade*, *Normalise*, *Crossfade* and *Digital Delay*.

This document contains user instructions for the Sample Editor. It assumes that the reader is familiar with the general operation of Apple Macintosh applications, and has a basic knowledge of digital audio.

Background

I wrote the Sample Editor as a final year project in early 1992 for the Department of Computer Science and Applied Mathematics, Aston University. The project was supervised by Dr Peter Coxhead and moderated by Catherine Johns-Lewis.

The code was written using Symantec's *Think C* development environment, so some parts of the program are copyright © 1989, 1991 Symantec Corporation. The Sample Editor application and this documentation are copyright © 1992 Garrick McFarlane.

Minimum System Configuration

This version of the Sample Editor has been written to run on a specific Apple Macintosh system. It was written for a Macintosh IIsi, though it should work on any machine with at least a 68020 processor. It requires System Software version 7.0 or

later, and at least a 640 by 480 pixel monitor with 8-bit colour. To record directly into the application requires

a sound-input device that conforms to Apple's sound input standards, such as the built-in microphone that is supplied with the newer Macintosh models.

Starting the Sample Editor

There are two ways to start a session with the Sample Editor. You can double-click on the Sample Editor application icon or one of its files. Alternatively, you can drag a sound file to the Sample Editor's icon and drop it in. The application will start up and present you with a window as shown in figure 1. If you selected a file, that sound's waveform will be displayed, otherwise the *Main View* and *Overview* will be empty.
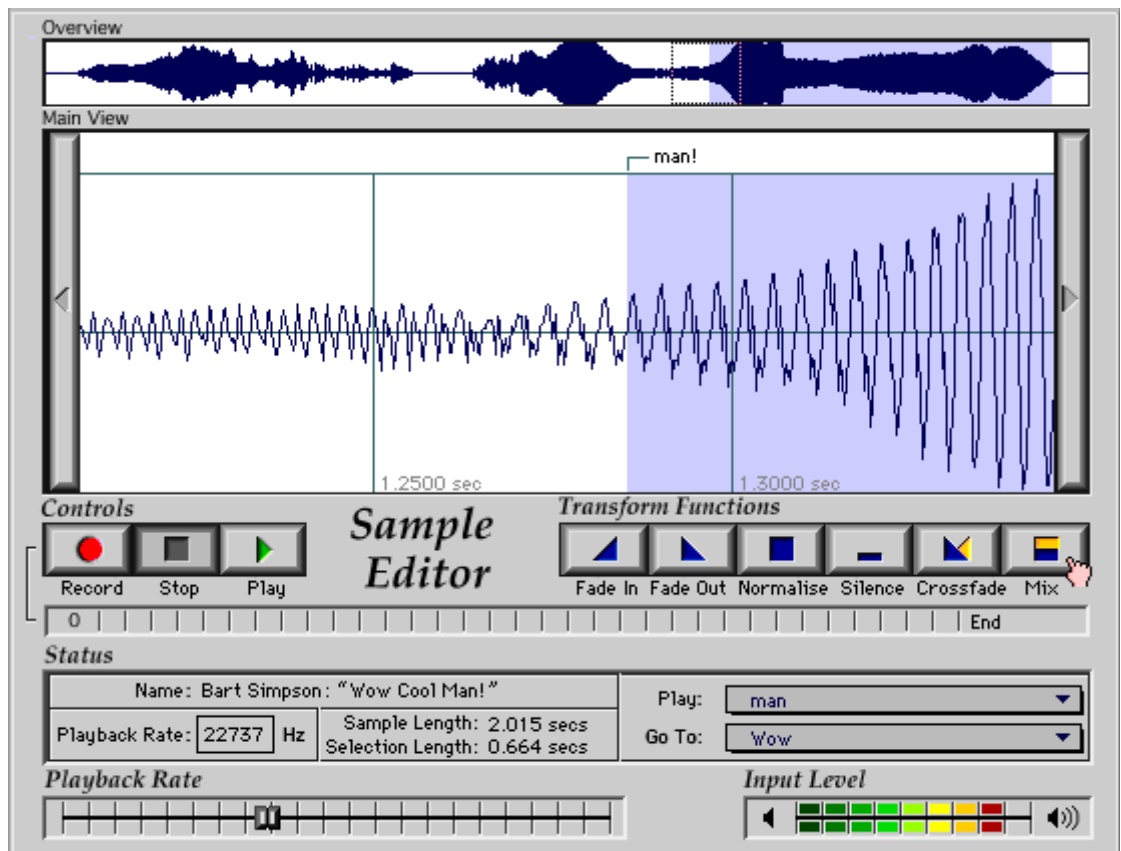
Figure 1—A Sample Editor Document Window

Finding Your Way Around

Balloon Help

The Sample Editor provides full support of Apple's Balloon Help system. To turn help on, choose *Show Balloons* from the [?] menu. To find out about any item, point to the part of the screen that you are interested in. A balloon will be displayed telling you what the item is and how to use it. An example is shown in figure 2.

Figure 2—Using Balloon Help

You can also use Balloon Help to find out what all menu items do, and what the controls in dialogue boxes are for.

Screen Layout

The main window of the Sample Editor is designed to represent the control panel of

an imaginary piece of equipment. It contains a number of separate parts that show different

information. A brief description of these items is given here, starting at the top of the window; for a full explanation of their use, read the relevant section of this documentation. Refer to figure 1 to see their positions.
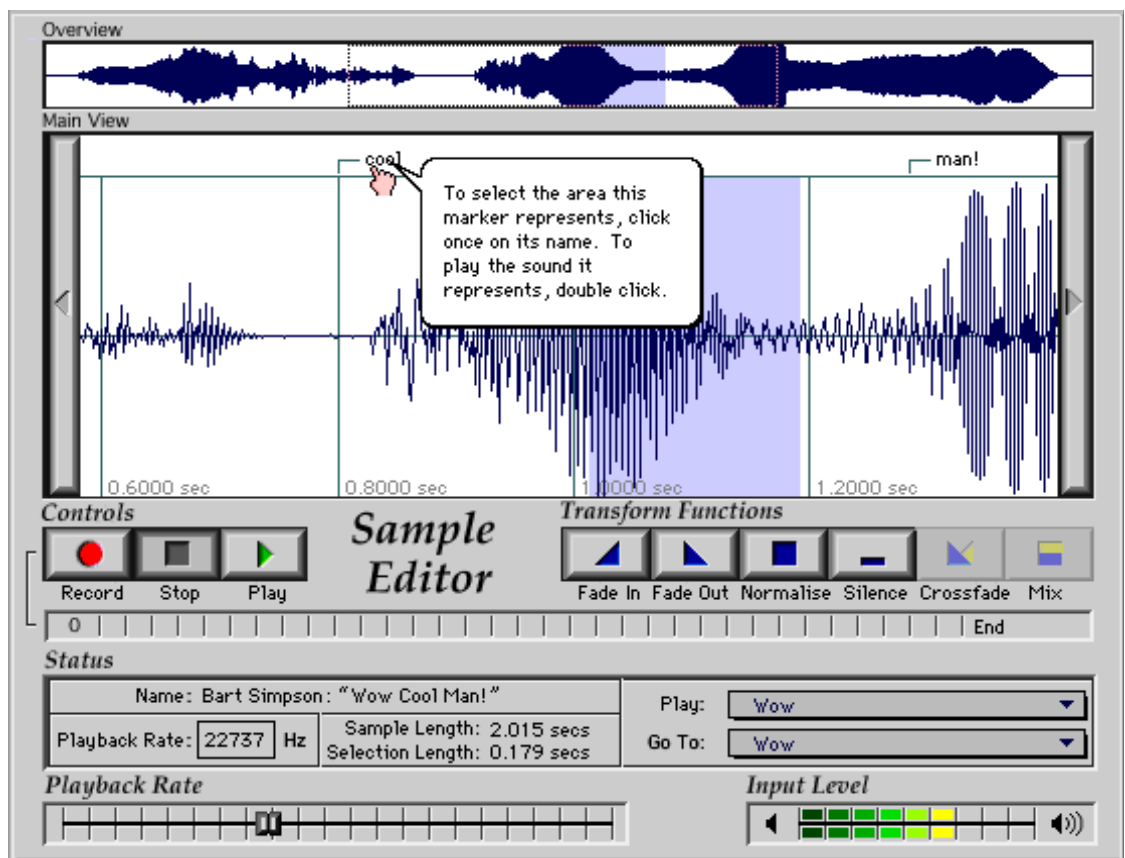
Overview

The *Overview* displays the whole waveform of the sampled sound you are currently working on. As the level of detail is not very high, it acts as a rough guideline as to the overall shape of the sound.

Main View

The *Main View* is the primary display of the sampled sound's waveform. It can show more detail than the *Overview*, not just because it is larger, but also because you can use it to look at a magnified section of the waveform. It is here that you do all editing of the sound.

Controls

In the *Controls* section are three buttons to control the Sample Editor much like the controls of a tape deck. When you begin, the *Stop* button will be pressed down. You can click on another button to operate the appropriate function.

Transform Functions

These buttons are operated in the same way as the *Controls*, but their purpose is to apply certain audio transformations to the selected part of the sampled sound.

Record Bar

Directly below the *Controls* and *Transform Functions* is a grey, empty bar called the *Record Bar*. This is only used while a recording is taking place; it is used to display the available recording time, the amount of that which has been used, and the shape of the recorded signal.

Status

The *Status* section contains various items regarding the current sample. There are displays to give you information such as the sample's name, length and playback rate. It also contains a couple of pop-up menus to help you move around the sample.

Playback Rate

This is a slider that you can use to adjust the sample's playback rate directly using the mouse.

Input Level

The section marked *Input Level* is an animated level meter, which constantly monitors the amplitude of the signal coming in through the Macintosh's sound input port.

❷*Viewing Waveforms*

As described in the previous section, there are two parts of the screen on which you can see the sample's waveform. At the top is the *Overview* which shows the whole sample's shape; below is the *Main View*.

Overview zooming

The *Overview* is not just used to display the shape of the sample. As you move the pointer into the *Overview* area, you should see the pointer change shape to become " ▶". This indicates that you can select a rectangle to zoom in on. By clicking and dragging on it with the mouse, you mark the part of the waveform that you are interested in. That part of the waveform will then be magnified and displayed in the *Main View* for editing.

By decreasing the height of the *Overview* selection, you can magnify the amplitude of the waveform displayed in the *Main View*. This might be useful, for example, to look closely at noise in a quiet section of a sample.

> *Tip:* To select the whole of the waveform for editing, just click once anywhere on the *Overview*.   The *Main View* will then display the whole sample, fully zoomed out, ready for editing.

Selecting and Scrolling

Once you have selected part of the overview as described above, it will be displayed in as much detail as possible in the *Main View* window. In order to specify which part of this you actually want to work on, you need to select a bit of it in the *Main View* display. This works in a similar way to an overview selection, but instead of selecting

a rectangle, you just need

to drag the mouse across the part of the waveform you wish to select. The chosen part will become highlighted as can be seen in figure 1, where the user has selected the right hand part of the sample.

If you need to select more than is shown on the *Main View*, you do not need to zoom out before re-selecting. Instead, you can just drag the mouse off the edge of the *Main View*. The waveform will automatically scroll in the appropriate direction, allowing you to select more than can be shown at any one time.

> *Tip*: If you hold down the *Shift* key when starting a selection, the previous selection will be automatically extended towards the place you click at. This can be useful if you want to adjust an existing selection by a small amount.

To scroll the *Main View* waveform without altering the selection, you must use the *Scroll Buttons*. These are pictured below in figure 3. Click on the appropriate button to move the waveform in that direction. These buttons will have no effect if you are fully zoomed out, or have scrolled as far as you can go in that direction.

> *Tip:* To help you keep track of the current selection when scrolling, watch the *Overview* display. This shows you a small version of the current selections, so you can see how much of the waveform you have selected.
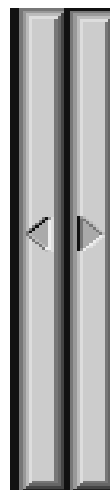

Figure 3—The Scroll Buttons

Naming Selections

A lot of the time you are working with the Sample Editor, you will find that you want to remember where specific sounds are in the waveform. To help you with this, the Sample Editor provides *Markers* that allow you to attach names to sections of the waveform.

To name a section of the waveform, you must first select the part you want to name on the *Main View*. Alternatively, make an insertion point by just clicking once at the place you want to mark; you can name this point in the same way as naming a selection. Once you have chosen the area you want to mark, pick the *Mark* command from the *Sample* menu. The dialogue box in figure 4 will be shown, in which you should enter the desired name.
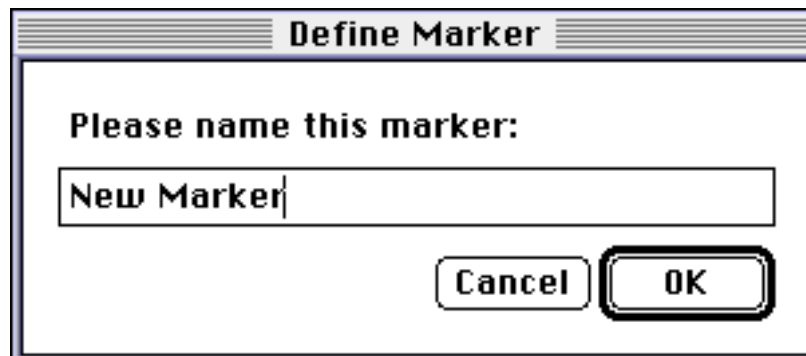

Figure 4—Defining a New Marker

Once you have named part of a sample, the named segment is indicated at the top of the *Main View* by a marker (shown in figure 5 below). The marker can thereafter be used to refer to the section you named. If you click on the marker, the piece of the waveform it is associated with will be automatically selected ready for editing.
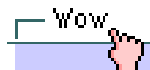
Figure 5—A Marker

Naming the important parts of a sound can help in navigation around a large sample. As well as seeing the markers on-screen, you can use the two pop-up menus (*Go To* and *Play*, figure 6) to quickly access parts of the sample. Pick an item off the *Play* menu to play it regardless of the current selection. Pick an item off the *Go To* menu and the Sample Editor will select that part of the waveform and display it in the *Main View*.
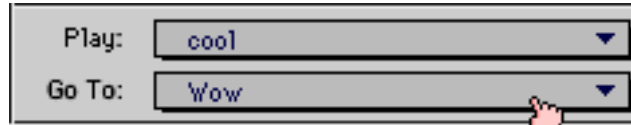
Figure 6—Navigating using Markers

When you edit a sample that you have named some sections of, the edit may affect the positions of these markers. For example, you may *Cut* out the whole of a marked section, or *Paste* some sound into the middle of a marked section. The Sample Editor will try as far as is sensible to maintain the connection between the markers and the original piece of *sound* that they referred to. In the examples given above, the markers would automatically be deleted or moved as necessary to reflect the new whereabouts of the sound.

> *Tip:* You can play a visible named selection easily—just double-click on its marker.

❸*Playing Samples*

You can use the Sample Editor to play sampled sounds through the Macintosh's internal speaker, or through connected headphones or speakers for better sound quality. To start a sample playing, choose *Play* from the *Sample* menu, or just click the *Play* button (shown in figure 7 below).

Figure 7—The Controls

If you have selected part of the waveform in the *Main View*, then only that section will be played when you press the *Play* button. If, however, there is no current selection, playback will begin at the current insertion point and continue until the end of the sample is reached. You can force the Sample Editor to play the whole sample, regardless of the current selection or insertion point, by holding down the *Option* key as you click *Play*.

While the sample is playing, you will see a pair of bars moving across the *Overview* and *Main View* areas. These bars are provided as indicators of the current playback position; they give you a rough idea of the part of the sample that is currently being played.

> *Tip:* A quick way to select and play a part of the waveform is to hold down the *Option* key while you make a selection. The selection you make will be played immediately you release the mouse button.

Playback Rate

When you first listen to a sample, it will be played at its default playback rate (usually 11 or 22 kHz). You can vary the playback rate, to speed the sample up or slow it down, in two ways.

If you do not need to set the rate exactly, it is easiest to use the *Playback Rate Slider*, shown in figure 8. Move the mouse so that the pointer is over the knob, then click and drag the mouse. As you move the knob left or right, the sample's playback rate is automatically adjusted as appropriate. If the sample is playing while you're adjusting its playback rate, you will hear the effect of your changes as you make them. When you release the mouse, the new playback rate will be displayed in the *Status* area.

Figure 8—The Playback Rate Slider

To set the playback rate more accurately, use the *Status* area itself. Move the mouse to the playback rate section of the *Status* area, and double-click to select the text displayed there. You can then type a new playback rate for the Sample Editor to use directly. Press *Return* or *Enter* to accept the new figure, or click on another part of the window.

Warning: If you alter the playback rate of a sample *while it is playing*, the playback position indicators will no longer be accurate and should be ignored.

₄*Editing Samples*

Before editing a sample in any way, you need to select the part you wish to modify, as described in *Selecting and Scrolling* in section ❷. Once you have highlighted the area to edit, choose the appropriate command. The standard Macintosh editing commands can be used from the *Edit* menu as normal, or for more complex transformations choose one of the transform functions from the buttons on the screen or the *Transform* menu.

The Sample Editor fully supports *Undo*. This means that you can try out any command to see how it affects the current sample. If you don't like the result, you can choose *Undo* from the *Edit* menu and the sample will revert to the state it was in before your last command.

This section describes the available editing features. Don't forget though that, when using the Sample Editor, you can use *Balloon Help* to remind you what does what.

Basic Editing

The Sample Editor supports the standard Apple Macintosh *Edit* commands—*Cut*, *Copy*, *Paste* and *Clear*—on sections of sampled sound.

If you *Cut* or *Copy* a section of sound, it will be placed in the Clipboard for future use. You can then *Paste* the sound back into the document at another position, or into another document. Some other Macintosh applications will also allow you to paste sound from the Sample Editor into them, such as *Microsoft Word 5.0*, *Hypercard 2.1*, and the *Sound* control panel that comes with the computer's system software.

To *Paste* a piece of sound into a Sample Editor document, you must first tell the program where you want the new bit of sound to go. Do this by selecting the part of

the sample you wish to replace, or if you just want to insert the Clipboard contents into the document, click once to make an insertion point at the place you wish the new sound to be inserted. If the Clipboard contains sound that has been compressed by another Macintosh application, the

Sample Editor will automatically expand the sound data before pasting it into the current document.

> *Warning:* Some Macintosh 'sounds' do not contain pure sampled sound, but also include a number of related items. If you try to *Paste* one of these sounds into the Sample Editor, it may not be able to successfully extract the sampled sound data, in which case you will be informed.

Transformations

The Sample Editor allows you to perform a number of more complex transformations to the sampled sound. To do one of these transformations, select the part of the sample you wish to affect, then choose the transformation from the *Transform* menu, or click the appropriate button on the *Transform Controls* area of the screen (shown in figure 9 below).

Figure 9—The Transform Control Buttons

Fade In/Out

These two transformations cause the selected part of the sample to be faded in or out respectively, to or from silence. To soften the start of an abrupt sound, just select a small area around the beginning of the sound, then choose *Fade In*. Alternatively, to do a long, gradual fade over the end of a sound, select a large part of the sample including the end of the sound, and choose *Fade Out*.

Normalise

Sometimes you may find that the whole (or part) of a sampled sound is too quiet. *Normalisation* increases the amplitude of the selected sound as far as it can go

without becoming distorted. It is often a good idea to normalise the whole sample after you make a recording, so that any discrepancies in the input level are compensated for. Normalisation is also useful when attempting to pick out sounds that appear in background noise, such as

faint voices. Normalising the 'noise' will often amplify it enough to be able to make out any sounds which are present in it.

Silence

This 'transformation' simply silences the selected block of sound. Silence generated with this function is completely pure and noise-free. To improve the perceived quality of a spoken passage, you could apply silence to the gaps between the spoken words, then add short fades in and out of each silent part to give a smooth transition.

Crossfade

The *Crossfade* transformation is used to perform a seamless edit between two different sections of sound. To be able to choose *Crossfade*, you must have a piece of sound in the Clipboard, and a section of the *Main View* waveform selected. The transformation will cause the selected sound to be faded out whilst the Clipboard contents are faded and mixed in at the same location. For example, if the document contained a sample of a cow mooing, and the Clipboard contained a sample of a dog barking, the result after *Crossfade* would be a "Moof!" with no discernible crossover between the "Moo" and the "Woof!".

Mix

This is used to blend two sounds together to become one. As for *Crossfade*, you need to have some sound in the Clipboard to be able to use the *Mix* feature. Select a part of the waveform in the *Main View*, then choose *Mix*, and the sound in the Clipboard will be merged with the current selection. This could be used to add vocals over an instrumental track, for example, or to add a background sound effect to another sample.

> *Warning:* When you mix two sounds together, you are effectively *adding* the two sounds. You should be careful not to mix two very loud sounds together, as this may lead to distortion in the resulting

sample.

The following two transformations do not have on-screen buttons to invoke them directly, but can be chosen from the *Transform* menu.

Digital Delay

The *Digital Delay* transformation can be thought of as a controllable *echo*. By varying the parameters, you can alter how many echoes are added to the sound, and how often the echoes occur. Because it is more complex than the other transformations, when you choose the *Digital Delay* command, the Sample Editor presents a dialogue box in which you must enter further parameters. The dialogue box is shown in figure 10.
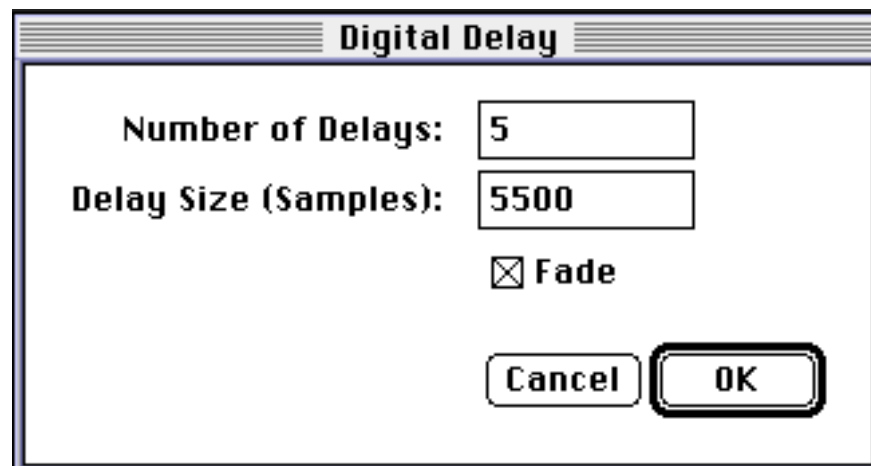

Figure 10—The Digital Delay Dialogue Box

The first field, *Number of Delays*, is where you enter the number of echoes you wish the Sample Editor to create of the selected sound. Note that the larger the number you enter here, the longer it will take for the Sample Editor to perform the transformation.

The other text field, *Delay Size*, is for you to specify how often the echoes should occur. This value is measured in samples. As an example, editing a sound to be played back at 22 kHz, you could enter 22,000 here to cause a one-second delay, or 11,000 to cause a half-second delay. Smaller delays than this tend to merge together

so that they are not audible as separate echoes. This can be put to use to apply special audio effects to the selected sound.

The *Fade* check box causes successive echoes to fade out, which makes the *Digital Delay* more like real-life echoing. The transformation is slightly slower when fading is selected.

Warning: As for *Mix*, you should be careful not to cause distortion when using *Digital Delay*; too many echoes may result in the sample being too loud. You should also be aware that complex delays using many faded echoes can take a long time to compute.

Reverse

The final transformation is *Reverse*. This simply reverses the currently selected sound in time. It is useful for percussion effects such as creating a reversed snare drum sound. It could also be used in conjunction with *Digital Delay* to create reverse echoes.

❺*Recording*

The Sample Editor can record sound directly from any connected sound input device that conforms to Apple's standards. This includes the built-in microphone on newer models such as the Macintosh IIsi and LC.

Record works in a similar way to the *Paste* function. First, select the part of sample you wish to record over, or the point at which you want to insert the new recording. You are now ready to record.

As soon as you click the *Record* button, recording will start. The pointer will change to a "🎤", and the record bar will begin to fill, indicating the amount of available space that has been filled with the new recording. During recording, you can monitor the input level using the on-screen level meter; the overall amplitude shape of the recorded signal so far is shown in the record bar. An example of a recording in progress (just after *Stop* has been clicked) is shown in figure 11, below. During recording, the word 'End' on the right of the record bar changes to show you how much time you can record for according to the current settings.

Figure 11—Recording in Progress

If you leave the recording going, the Sample Editor will automatically stop when it has filled the available memory space. To stop the recording manually, you should click on the *Stop* button with the mouse. Whilst recording is in progress, you will not be able to use any other functions of the Sample Editor.

> *Warning:* It is good practice to make recordings as short as

possible. The *Record* function requires large amounts of memory; a long

recording may cause the Sample Editor to run short of memory afterwards.

Microphone Configuration

The Sample Editor can accept audio signals from any device that conforms to Apple Computer's standards for sound input devices. Because different devices can digitise at different sampling rates and sample sizes, you may need to alter the settings before starting a recording. To do this, choose *Microphone* from the *Configure* menu. The Sample Editor will display the dialogue box shown below in figure 12.
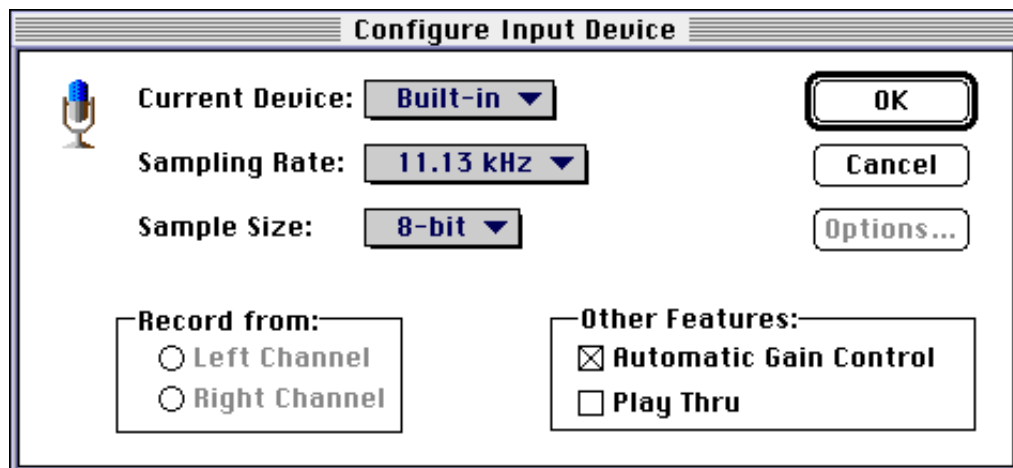
Figure 12—The Configure Input Device Dialogue Box

When the Sample Editor starts up, it assumes you wish to record using the built-in microphone if it exists. To change device, choose the one you require from the *Current Device* pop-up menu. When you choose a device from this menu, the other items in the dialogue box will be updated to reflect the settings of that device. To make it compatible with the Sample Editor, you must choose an 8-bit sample size, and if it is stereo, the channel you wish to record from. If your microphone has further options available, you can access them by clicking on the *Options* button. Note that any changes you make in the *Configure Input Device* dialogue box will not take effect until you click the *OK* button.

Automatic Gain Control

When this is turned on, the Macintosh automatically compensates for signals that are

too loud by reducing the sensitivity of the microphone. This helps prevent distortion in loud

environments. This feature is turned on by default; click the check box to turn it off, after which it is up to you to make sure the sound being recorded is not too loud.

Play Thru

Play Thru, which is only available for certain hardware configurations, allows you to listen to the incoming signal as it is being recorded. This is most useful when you are recording directly from a piece of audio equipment, such as a personal stereo, to enable you to hear the sound that it is being fed into the Macintosh. If you turn *Play Thru* on when recording from a microphone, and you are not using external headphones, you should be careful not to put the microphone too near to the Macintosh's internal speaker as feedback ('howl') may result.

Record protect

Like all other Sample Editor functions, *Record* can be undone if you aren't happy with a new recording. However, it is still a potentially dangerous feature—especially if a large recording results in a memory shortage—so the Sample Editor provides the *Record Protect* option. When this is turned on (using *Preferences*, described later), you will always be asked before making a recording to verify your decision. The alert box shown in figure 13 will be displayed, allowing you to change your mind and cancel the command. If you click the *Start Recording* button instead, recording will begin immediately.

Figure 13—The Record Protect Alert

₆*Audio Files*

When you have finished editing or recording a new sound, you will want to save your work on disk. There are two main ways the Sample Editor stores sounds—Audio Interchange files and as Macintosh 'snd ' resource files (the sound files used by System Software version 7). This section describes the two types of file, and explains the available audio compression options.

Compression

One problem with storing sampled sound is the large amount of storage space required. The Sample Editor can help alleviate this problem by allowing you to compress the sound as you store it on disk. Whichever file type you choose to use (see below), you can opt to compress the file by a ratio of up to 6:1. Files compressed in this way are still compatible with most other Macintosh audio-related applications. For example, you could use the Sample Editor to compress sounds that you want to use as System beeps to save disk space.
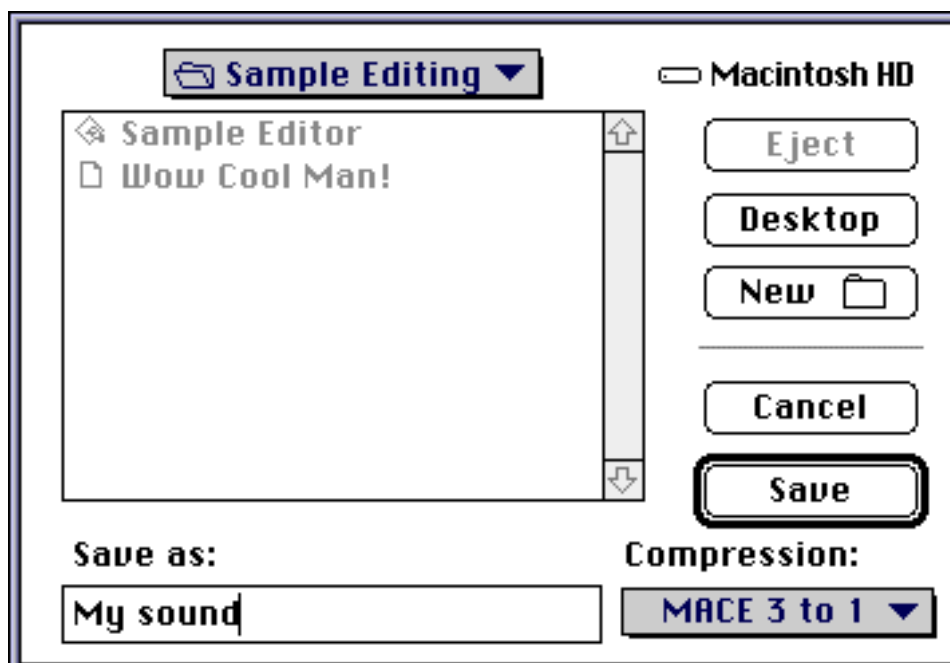
Figure 14—Choosing Audio Compression Type

You specify the type of audio compression to use as you *Save* (or *Export*) a sound file. In the dialogue box (shown in figure 14), as well as the usual items, is a pop-up menu from which you choose the audio compression type. In the example, 3:1 compression has been selected. 'MACE' refers to the method used to compress the file (Macintosh Audio Compression and Expansion)—the standard Macintosh compression method.

> *Tip:* If you are compressing sampled music, use 3:1 compression as it does not have much of an effect on the fidelity of the sound. 6:1 compression is more suited to speech or other low-fidelity sounds, where the sound quality is not important.

If you re-open a compressed sound file using the Sample Editor, it will automatically be expanded to allow you to edit it as normal.
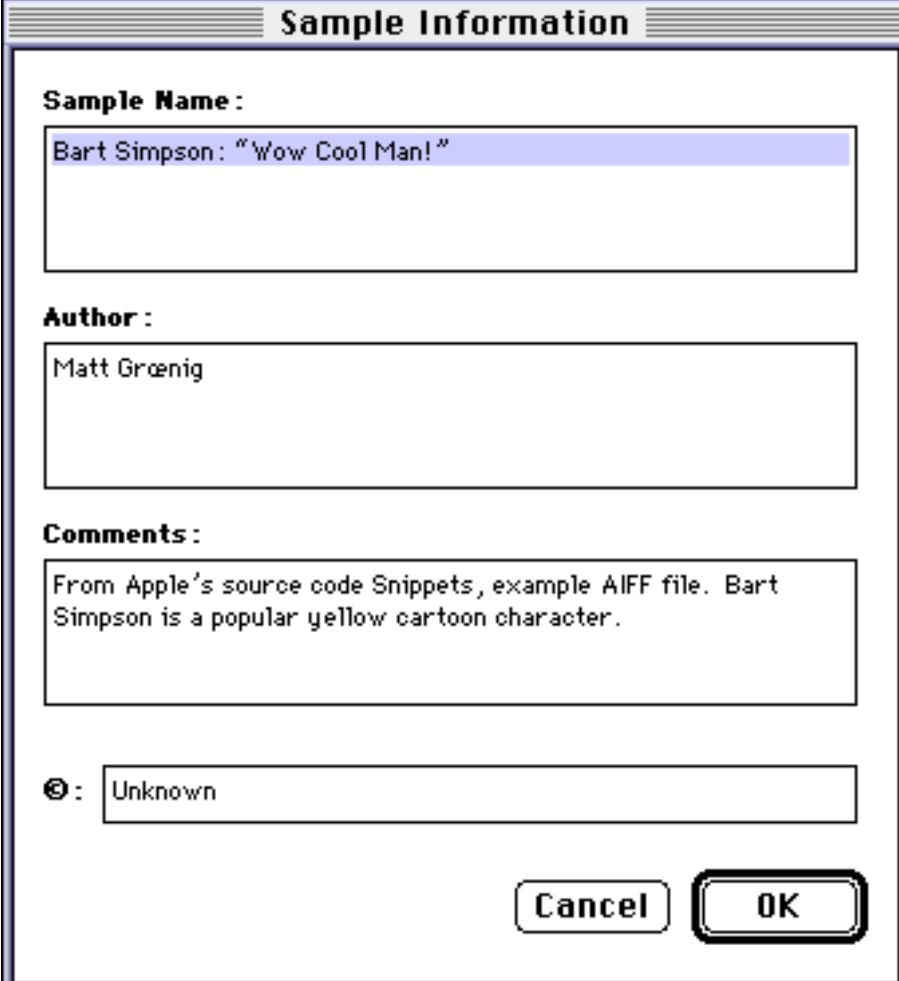
Audio Interchange Files

The Audio Interchange File Format (AIFF) defines a new standard type of file that audio applications can read. All Sample Editor files are stored in this way, and other audio editing applications should be able to read them directly without needing to perform any kind of conversion.

Sample Information

A lot of information regarding the sampled sound is stored inside an Audio Interchange file. As well as low-level data concerning playback rates and so on, you can store the details of the sample's name, author, copyright and general comments. To see or edit this information, choose *Get Information* from the *Sample* menu. A dialogue box like that shown in figure 15 will be displayed. If the sample has already been named, you can edit the information presented, or you can enter new details into the boxes.

Note that the *Sample Name* as stored here is *not* the sample document's filename. Changing the document's filename has no effect on the sample's actual name. You can use the *Sample Name* to store the full name of the sample without worrying about the constraints placed by the filing system.

Figure 15—The Sample Information Dialogue Box

Macintosh '`snd`' Files

The other method you can use to store your sounds is as a Macintosh '`snd`' file. '`snd`' files contain less information than Audio Interchange files, and are less portable; they cannot be used on computers other than Apple Macintoshes, for example. They are useful because this is the kind of file that the System Software uses to move sounds around. You can drop '`snd`' files into the System suitcase to add them to the list of System sounds available for beeps and other applications.

To access '`snd`' files using the Sample Editor, use the *Export* and *Import* commands on the *File* menu. Choose *Export* to create a '`snd`' file from the current document; use *Import* to create a Sample Editor document from an existing '`snd`' file.

₇*Miscellaneous*

This section covers some of the more general features of the Sample Editor—*Preferences* and *Auto-Play*.

Auto-play

When you are browsing through files in the *Open* or *Import* file dialogue boxes, you can choose to use the *Auto-Play* feature. When *Auto-Play* is on, the Sample Editor will automatically play any sound file as you click on it in the file list. As soon as you click on another file, the old one will stop playing and playback of the new one will begin. This makes it easy for you to scan a group of sound files quickly without the need to *Open* the file into a Sample Editor document.
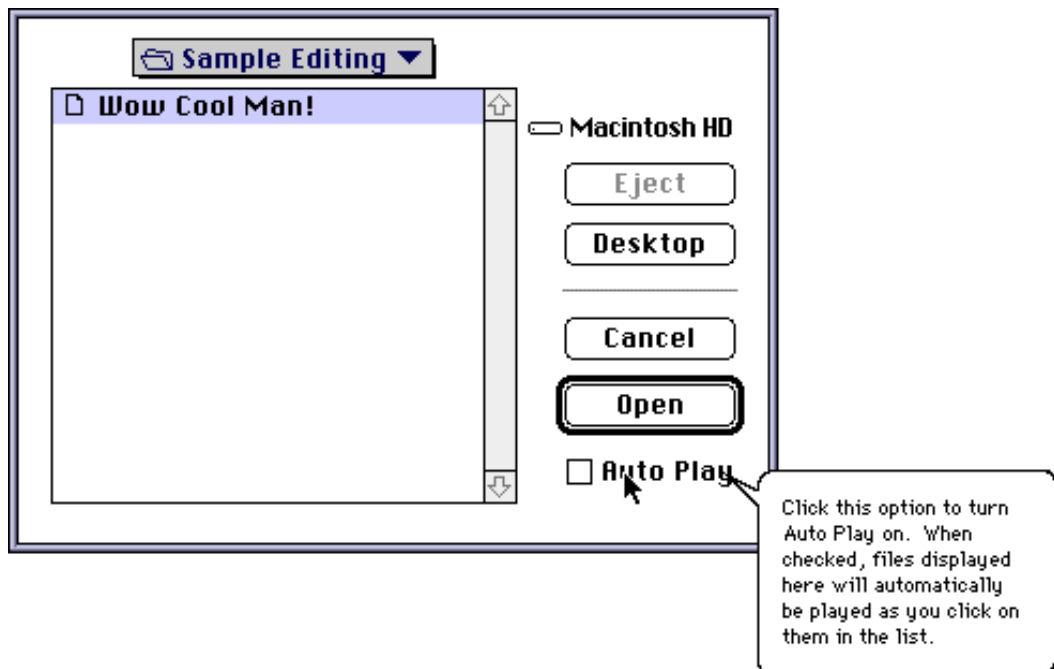
Figure 16—Auto Play in the Open File Dialogue Box

The diagram above (figure 16) shows the *Open File* dialogue box. Notice that there is an extra item, the check box marked *Auto-Play*. You can use this check-box to turn *Auto-Play* on or off temporarily. To change its state permanently, make the appropriate setting in the *Preferences* dialogue, described next.

Sample Editor Preferences

Sample Editor Preferences

Sample Colour:	**Text Colour:**
Axis Colour:	**Time Colour:**

Overview Accuracy: 100 %

☐ **Record Protect**   ☐ **Interpolation**

☒ **Auto Play in Open File Dialog**   ☐ **Drop-sample Conversion**
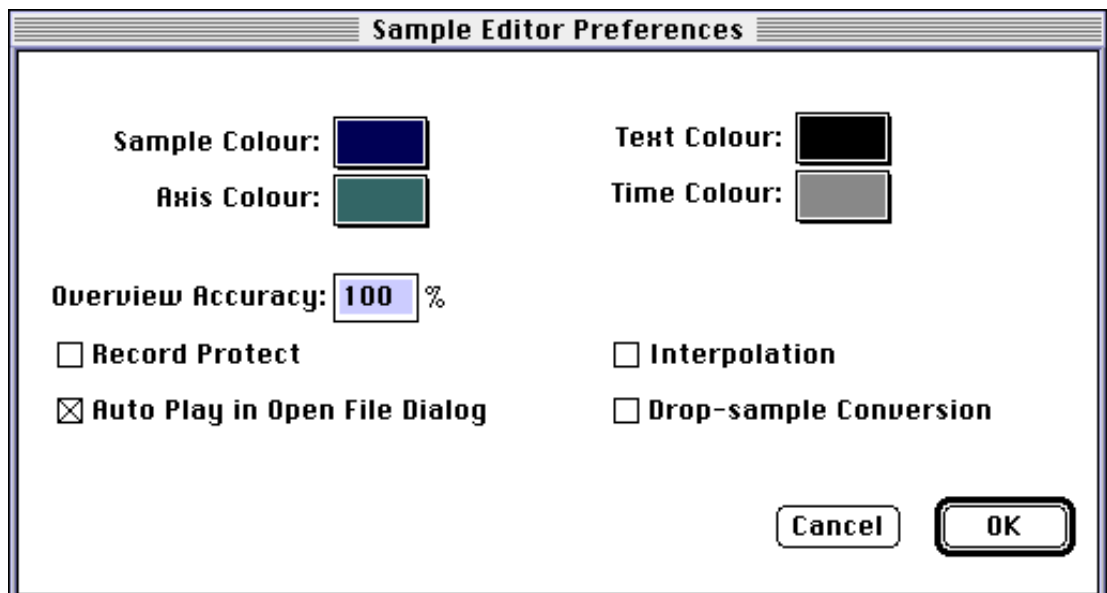
Cancel   OK

Figure 17—The Preferences Dialogue Box

Choose *Preferences* from the *Edit* menu at any time to see the dialogue box shown in figure 17, above. This lets you alter the default settings that the Sample Editor uses.

Colours

These four boxes let you change the colours used by the Sample Editor to draw various parts of the display. Click in the box you want to change, and the standard Macintosh colour wheel will be displayed from which you should choose the new colour.

The *Sample Colour* is the colour used to draw the waveforms. The *Axis* colour is used to draw the gridlines behind the waveform on the *Main View*. General text items are displayed in the *Text Colour*, whilst the *Time Colour* is the colour of the time readings given at the bottom of the *Main View*.

Overview Accuracy

Every time you make an edit to the sample, the Sample Editor must recalculate the shape of the waveform to display in the *Overview*. This calculation can take some time, especially for large complex samples on slower machines. You can speed up the calculation by telling the Sample Editor not to worry so much about its accuracy. Enter a new number in the *Overview Accuracy* field of the dialogue box, between 0 and 100, indicating the percentage accuracy. The normal value is full accuracy, i.e. one hundred per cent.

Auto-Play

This check box is used to set the default action of the *Open* and *Import* dialogue

boxes. If it is checked, *Auto-Play* will be on by default; if it is unchecked, you will need to turn *Auto-Play* on manually when you want to use it.

Interpolation

Use this option to select whether or not linear interpolation is used when playing a sound back at a different frequency from the sound's recorded frequency. Sounds most affected by the absence of linear interpolation are sinusoidal sounds. Sounds least affected are noisy sound effects like explosions and screams. Turning off linear interpolation decreases the amount of work the computer has to do to play back sounds.

Drop-Sample Conversion

Use this option to select whether or not to use drop-sample conversion when playing a sound back. This improves the playback quality of certain kinds of sound, but as for linear interpolation, increases the amount of processing the computer must do to play a sound.

If you have downloaded this user guide—

The user guide was originally written in QuarkXPress. I have converted it to Word 4.0 format as I presume not many people can read Quark format. However, much of the formatting has been lost in the translation process (as indeed have some of the pictures). So—don't take too much notice of the Table of Contents page numbers for example, and if something appears to be missing, it probably is...

If you want the User Guide in its original Quark format, let me know and I'll send it to you. See the Read Me note for contact details.