

GoldWave™ User's Manual

Copyright 1993-1994 Chris S. Craig

January 1994

Table of Contents

1 Introduction

2 Getting Started

- 2.1 System Requirements
- 2.2 Notation Used in this Manual
- 2.3 Installation

3 Using GoldWave

- 3.1 Window Overview
- 3.2 Editing Overview
 - 3.2.1 Selecting Part of a Sound
 - 3.2.2 Sample Editing with the Mouse
- 3.3 Device Controls Overview
 - 3.3.1 Playing Sounds
 - 3.3.2 Pausing Playback
 - 3.3.3 Stopping Playback and Recording
 - 3.3.4 Rewinding and Fast Forwarding
 - 3.3.5 Recording Sounds
 - 3.3.6 Volume Scroll Bar (Vol)
 - 3.3.7 Balance Scroll Bar (Bal)
 - 3.3.8 Speed Scroll Bar (Spd)
 - 3.3.9 Setup
- 3.4 Keyboard Overview
- 3.5 File Menu Commands
 - 3.5.1 New
 - 3.5.2 Open (Open Button) and Supported File Formats
 - 3.5.3 Save
 - 3.5.4 Save as
 - 3.5.5 Export
 - 3.5.6 Exit
- 3.6 Edit Menu Commands
 - 3.6.1 Undo (Undo Button)
 - 3.6.2 Cut (Cut Button)
 - 3.6.3 Copy (Copy Button)
 - 3.6.4 Paste (Paste Button)
 - 3.6.5 Mix (Mix Button)
 - 3.6.6 Delete (Delete Button)
 - 3.6.7 Trim (Trim Button)
 - 3.6.8 Insert space (Space Button)
 - 3.6.9 Select view (Sel Vw Button)
- 3.7 View Menu Commands
 - 3.7.1 Zoom block (Block Button)
 - 3.7.2 Unzoom (Out Button)

- 3.7.3 Zoom user defined (User Button)
- 3.7.4 Zoom 20:1
- 3.7.5 Zoom 5:1
- 3.7.6 Zoom 1:1 (1:1 Button)
- 3.7.7 Zoom 1:10
- 3.7.8 Zoom 1:100
- 3.8 Effects Menu Commands
 - 3.8.1 Cross-fade
 - 3.8.2 Echo
 - 3.8.3 Exchange channels
 - 3.8.4 Filter
 - 3.8.4.1 Lowpass 50%
 - 3.8.4.2 Lowpass 25%
 - 3.8.4.3 Highpass 10%
 - 3.8.4.4 Highpass 25%
 - 3.8.5 Invert
 - 3.8.6 Mechanize
 - 3.8.7 Pan
 - 3.8.8 Reverse
 - 3.8.9 Silence
 - 3.8.10 Volume
 - 3.8.10.1 Change
 - 3.8.10.2 Fade in
 - 3.8.10.3 Fade out
 - 3.8.10.4 Maximize
 - 3.8.10.5 Shape
 - 3.8.11 Playback rate
 - 3.8.12 Resample
 - 3.8.13 Expression evaluator
- 3.9 Options Menu Commands
 - 3.9.1 Input device
 - 3.9.2 Output device
 - 3.9.3 Aux volume
 - 3.9.4 Settings
 - 3.9.5 Editing (Chan Button)
- 3.10 Window Menu Commands
 - 3.10.1 Tile
 - 3.10.2 Cascade
 - 3.10.3 Arrange icons
- 3.11 Help Menu Commands
 - 3.11.1 Contents
 - 3.11.2 About

4 General Information

- 4.1 Warranty, Trademarks, and Copyrights
- 4.2 Contacting the Author

- 4.3 Reporting Problems and Bugs
- 4.4 Revision History
- 4.5 Future Versions
- 4.6 Acknowledgements

Appendix A: Audio Attributes/Terminology

- A.1 Digital Sound Attributes
- A.2 Problems with Recording
- A.3 Frequency Spectrums

Appendix B: Expression Evaluator

- B.1 Overview
- B.2 Evaluation Range, Variables, and Constants
- B.3 Valid Operations and Functions
- B.4 Signal Generation
- B.5 Filters

Appendix C: Trouble Shooting and Q&A

- C.1 Trouble Shooting
- C.2 Common Questions and Answers

1 Introduction

GoldWave™ is a digital audio editor for the Microsoft Windows™ 3.1 platform. GoldWave allows you to easily and quickly record and edit music, voices, instruments, and any other sounds. Built in to GoldWave are many features such as an expression evaluator, dual amplitude/spectral oscilloscopes, an independent device controls window, intelligent and fast editing, and numerous effects.

The expression evaluator allows you to generate everything from simple dial tones to band-pass filters. It supports over 20 common functions and operations.

The dual oscilloscopes graphically display sounds as they are being played or recorded. Data can be graphed as a series of amplitudes or as a frequency spectrum. The left and right channels of stereo sounds each have their own oscilloscope.

The device controls window contains buttons for playback, recording, volume, etc. This window is independent from GoldWave's main window and allows flexible control over audio devices.

All editing operations, such as paste and mix, are intelligent. This means that clipboard audio data is automatically converted to a compatible format before the data is used. When editing files with different attributes (sample rates, bits, and channels), intelligent editing simplifies the process.

GoldWave employs a multiple document interface (MDI) so that up to five sounds can be loaded at one time, making file-to-file editing easy. The sounds are displayed graphically and the level of detail can be changed by zooming in or out. Samples can be edited directly with the mouse when zoomed in to a sub-sample level (the level at which individual samples are easily visible).

GoldWave also includes many audio effects, including reverse, echo, volume shaping, panning, filters, etc. These allow you to enhance or distort sounds in various ways.

This manual presents and explains many features of GoldWave. The first part discusses system requirements, installation, etc. The second part covers the menu structure. Topics are covered in the order that they appear in GoldWave's main menu. The final section provides general information (file list, warranty, etc.). Please be sure to read this section.

For users who are unfamiliar with digital audio, **Appendix A** briefly introduces some of the fundamentals of digital sound. It also provides some solutions to common recording problems. This manual assumes that users are familiar with Microsoft Windows.

2 Getting Started

2.1 System Requirements

The minimum system requirements for GoldWave v2.00 are:

- 16MHz 386 (or better) IBM PC or compatible
- Microsoft Windows™ 3.1
- 4 Megabytes of RAM
- 1 Megabyte free on your hard disk
- Mouse
- VGA card and monitor
- Audio card and an MPC software driver (or PC speaker driver)

Note that GoldWave will not run on a 286. GoldWave was compiled using the 386 instructions set. These instructions are used for the spectral (FFT) oscilloscope and for faster processing of effects and editing.

A math co-processor is not required, but it is recommended if you will be using the expression evaluator.

If you need to edit large files, you will need a large amount of RAM. One minute of CD-quality sound requires 10 megabytes of memory.

2.2 Notation Used in this Manual

Bold text and a vertical bar, "|", are used to indicate menu commands. For example, **File | New** means to select the **New** command from the application's **File** menu.

This notation is also used to refer to other sections within this manual. In the above example, you can find information by looking for **New** under the **File Menu Commands** section.

2.3 Installation

GoldWave and all its files should be kept in a single directory of its own (the file list is

given in the **General Information** section). You can simply copy all the files from the GoldWave disk to a new directory or you can use the install file included on the disk.

For example, to install GoldWave from drive "A:" to the directory "gw" on drive "C":

- 1) C:> mkdir gw
- 2) C:> copy a:*. * gw

or

- 1) C:> a:install c:\gw

Before running GoldWave make sure that you have installed an appropriate Windows sound driver using the Control Panel "Drivers" Applet. The driver and installation instructions should be included with the hardware. If you do not have a sound card, then you can use the PC-Speaker driver (you may have to search BBS or ftp sites to find this, since it is not included with Windows). Note: Since the PC-Speaker driver does not return control to Windows until after the sound has finish playing, the oscilloscope feature and many of the control buttons will not work.

To add a GoldWave icon to a Program Group use **File | New** on the Program Manager's menu, select **Program Item** then click on **OK**. Type "GoldWave v2.00" as the description. Choose **Browse**, find and double click on "goldwave.exe", then click on **OK**.

To run GoldWave, either double click on the new GoldWave icon or use **File | Run** on the Program Manager's menu.

GoldWave allows you to choose which devices to use for input and output. Selecting **Options | Output device** presents you with a list of installed output devices and shows their capabilities. Choose the device you wish to use. A similar list is presented for **Options | Input device**.

Be sure to check the readme.txt file for any additional information not available at the time this manual was printed.

3 Using GoldWave

3.1 Window Overview

GoldWave is composed of three types of windows: the Main window, the Device Controls window, and Sound windows. The Main window contains the main menu and a row of buttons. It groups together and manages all the Sound windows when they are opened. At the bottom of the Main window is an information bar that gives you attributes of the active Sound window. The upper half of the bar gives details for the entire sound. The lower half gives similar detail for the currently selected part of the sound.

The Device Controls window provides an easy to use interface to your audio hardware. It contains buttons to play, record, and stop sound as well as controls for volume, balance, and playback speed (provided your audio device supports these features). The amplitude/spectral oscilloscope is displayed here whenever a sound is played or recorded. See **Device Controls Overview** in a following section.

Sound windows are created when you open a file. Up to five sounds can be opened at any one time. These windows contain a graph of the sound with a time axis near the bottom. For stereo sounds, two separate graphs are shown with the top (green) graph as the left channel

and the bottom (red) graph as the right channel. Near the bottom of the Sound window is a small "Overview" box. This box shows you what part of the sound is selected (highlighted green and/or red), what part is displayed in the above graph (black background), and what part is not visible in the above graph (dark grey background). Initially, the entire sound is visible and selected.

GoldWave makes extensive use of "bitmapped buttons". These buttons have pictures which help describe their purpose. Buttons containing no colour (only shades of grey) are disabled. They are disabled when the function they perform is invalid. Most buttons are disabled when you initially run GoldWave. After opening a sound file, most of these buttons will become enabled. Information about each button can be found by looking in the Table of Contents. Buttons are listed in parentheses as "(Open Button)", for example.

3.2 Editing Overview

The GoldWave editor is block based. A block is defined as the highlighted part of the sound graph between two vertical markers. Most editing and effects are performed on this highlighted (selected) block. As mentioned in the previous section, when a sound file is first opened, the block is the entire sound. The vertical markers are cyan lines located at the far left side (start marker) and far right side (finish marker) of the graph; they may be difficult to see.

3.2.1 Selecting Part of a Sound

There are two ways to move the block selection markers, depending on **Options | Settings** setup. GoldWave does not use Window's standard "click-and-drag" method because it does not allow accurate positioning of the markers. The methods below allow you to independently set the start and finish markers to an individual sample (when zoomed in at a 1:1 level or better).

Method 1:

The default method (check **Quick selecting** in **Options | Settings**), uses the left and right mouse buttons to set the start and finish markers respectively. Whenever the mouse arrow is inside the sound's graph, the arrow changes into a vertical bar with small arrows on either side. This means that you can change a marker's position.

To move the start marker:

Click the left mouse button on the graph at the point where you want to place the start marker.

To move the finish marker:

Click the right mouse button on the graph at the point where you want to place the finish marker.

Note that you cannot place the finisher marker before the start marker. The same is true for setting the start marker after the finish marker.

The drawback of this method is that it is easy to accidentally change the start/finish marker

when clicking on a Sound window to activate it.

Method 2:

If the **Quick selecting** check box in **Options | Settings** is not checked, then selection method 2 is used.

To move a marker, position the mouse arrow over it. You should see the arrow change into a vertical bar with small arrows on either side. You can then click and drag the marker to a new location. For example, to move the start marker from its original position:

- 1) Move the mouse arrow to the far left side of the graph. Make sure it is inside the graph area and the arrow shape has changed.
- 2) Press and hold the left mouse button.
- 3) Move the mouse to the right.
- 4) Release the mouse button when the start marker is at the desired position.

To find a certain part of the sound, you can play the sound and mark it as it plays. GoldWave indicates the part being played with a white vertical line.

In most cases, editing and effects are performed only on the active Sound window. The expression evaluator, however, can modify any part of any Sound window.

Window's clipboard (a temporary storage area that all applications can use) is used for most of the editing commands. Note that cutting or copying a large block into the clipboard can consume a lot of memory.

See **Options | Editing** for information about controlling independent channel editing.

3.2.2 Sample Editing with the Mouse

GoldWave lets you directly edit samples with the mouse. To do this, you must first zoom in so that individual samples are visible (see **View | Zoom 5:1** or **View | Zoom 20:1**).

- 1) Zoom in 5:1 or closer.
- 2) Place the mouse arrow directly over the sample you want to edit. The arrow will change into a thin horizontal line with two small arrows pointing to it.
- 3) Click and hold the left mouse button.
- 4) Move the mouse up or down to change the sample.
- 5) Move the mouse slowly left or right to change neighbouring samples (if you move too fast, you may skip a sample).
- 6) Release the mouse button.

Direct editing is great for removing pops or clicks. You can also use it to draw new sound waves of almost any shape.

3.3 Device Controls Overview

The Device Controls window is an easy to use interface to your audio hardware and/or drivers. On the left half of this window are the oscilloscopes in which sound is graphed "real-

time" during playback and recording. On the top right half of the window is a standard set of tape recorder controls, including play, stop, record, rewind, pause, and fast forward. In the middle right half of the window are controls to set the device's output volume, balance, and playback speed. In the lower right half, is a setup button for configuring the Device Controls window.

3.3.1 Playing Sounds

After opening a sound, you can use the play button (the one with a green triangle that points to the right). While a sound is playing, it is shown graphically in the oscilloscope box. The current position is shown in the graph of the Sound window by a vertical white line.

Sounds can be played only if the driver supports the sound's format and attributes. By using **Edit | Copy** and **Edit | Paste**, it is possible to convert a sound to a playable format.

3.3.2 Pausing Playback

While a sound is playing, you can pause it with the Pause button (two vertical blue bars). Remember to use either play or stop later. Pause will freeze the oscilloscope so you can see the shape of the sound.

3.3.3 Stopping Playback and Recording

Recording and playback can be stopped immediately with the stop button (red square).

3.3.4 Rewinding and Fast Forwarding

While a sound is being played, you can use the rewind button (blue triangle pointing left) or fast forward button (blue triangle pointing right) to quickly move back and forward through the sound. The current position is shown in the graph of the Sound window by a vertical white line. Playback is suspended until the play button is pressed.

You can adjust the speed of rewind and fast forward with the **Setup** dialog described in a following section.

3.3.5 Recording Sounds

The record button (red circle with 'R' in it) allows you to record your own sounds. Recording occurs inside the selected block, overwriting any data that was previously there. To make room for recording without overwriting the block, you can use **Edit | Insert space**. Recording will automatically stop when the block is full. You can stop recording at any time with the stop button.

Recording is successful only if the driver supports the active sound's attributes/format. You can try recording into a new sound (**File | New**) and use **Edit | Copy** and **Edit | Paste** to overcome this problem.

3.3.6 Volume Scroll Bar (Vol)

If your audio device supports volume control, you can use the **Vol** scroll bar to change the output volume of the device. Move the thumb right or click the right arrow to increase the volume. Move the thumb left or click the left arrow to decrease the volume.

3.3.7 Balance Scroll Bar (Bal)

If your audio device supports left and right volume control, you can use the **Bal** scroll bar to change the device's left/right balance. Move the thumb in the direction you want to shift the balance.

3.3.8 Speed Scroll Bar (Spd)

If your audio device supports multiple playback speeds, you can change the speed with the **Spd** scroll bar. Move the thumb right to increase the speed, and left to decrease it. Changing the speed also changes the pitch.

3.3.9 Setup

The Setup button presents a dialog containing many configurable options for the Device Controls window. These options are described below.

Graph type

The **Graph type** group allows you to set the type of oscilloscope to display. These include:

Amplitude: the sound is shown as a series of amplitudes.

Spectral: the sound is shown a frequency spectrum (see **Appendix A**).

None: the oscilloscope is disabled.

Controls

The **Controls** group lets you configure the speed of the fast forward and rewind buttons. By entering a value of "3.00" for the **Rewind time**, for example, the rewind button will move the current position backwards through the sound three times as fast as regular playback speed.

By entering small numbers (such as "0.01") the rewind and fast forward buttons will move the current position very slowly through the sound. This is useful for finding pops or clicks, since the oscilloscope will also move slowly through the data.

Graph length

The size of the oscilloscope windows can be changed by using the radio boxes in the **Graph length** group. If **Long** is selected, the graph will show 256 points. If **Short** is selected, the graph will display 128 points.

Looping

Looping options allows both playback and recording to be looped continuously. Record looping is useful when you want to monitor the audio inputs. This lets you to adjust volumes to prevent clipping. You can also analyze the input using the Spectral oscilloscope.

Playback

The **Playback** options define what part of the sound will be played when the play button is pressed. If **All** is checked, the entire sound is played. If **Window** is checked, only the visible segment of the sound is played. If **Selected** is checked, only the sound between the start and finish markers is played.

Other

The **Other** group contains a few of miscellaneous options. The **Spectral windowed** check box enables (checked) or disables (unchecked) windowing of sound data before spectrum processing (See **Appendix A** for more information).

If the **Grid on graph** check box is checked, the oscilloscope windows will have horizontal and vertical gridlines.

If the **Save setup** check box is checked, all the settings will be saved in GoldWave's ".ini" file after the **OK** button is pressed.

Positioning

The **Positioning** group lists several methods for determining the current position when a sound is played. This position is used to synchronize the oscilloscopes and the white "current position" line. GoldWave relies heavily on correct positioning for many of the device controls such as rewind, fast forward, and speed adjustment.

The first three methods (**Milliseconds**, **Samples**, and **Bytes**) uses the device driver to get the current position. The fourth method (**Estimate**), uses the system timer and the sampling rate to approximate the current position. In many cases, the **Estimate** method is more accurate than the others. However, with certain devices, this method may slowly loose synchronization when playing long sounds, since the device does not play at the exact sampling rate (the Sound Blaster, for example, plays 11025Hz sounds at 11111Hz).

3.4 Keyboard Overview

Table 1 summarizes keystrokes used in GoldWave.

Table 1: Keystrokes

Keystroke	Action
Shift+Del	Remove block and copy it into the clipboard.
Ctrl+Ins	Copy block into the clipboard.
Paste	Copy clipboard into the active sound at the start marker's position.
Del	Delete the block, permanently.
Shift+Ctrl+Ins	Mix clipboard with the active sound at the start marker's position.
Ctrl+T	Trim sound. Removes all data outside the block.
Left	Scrolls the active Sound window left.
Right	

Pg Up	Scrolls the active Sound window right.
Pg Dn	Scrolls the active Sound window left one screen.
Home	Scrolls the active Sound window down right screen.
End	Redraws the active Sound window at the start marker's position.
Ctrl+Home	Redraws the active Sound window at the end marker's position
Ctrl+End	Redraws the active Sound window at the beginning of the sound.
F1	Redraws the active Sound window at the end of the sound.
Space	Starts on-line help.
	Plays a sound (only when Main window is active).
	Plays/Stops a sound (when Device Controls window is active).
F4, F5, F6, F7, F8	Plays, fast forwards, pauses, rewinds, and stops respectively (only when Main window is active).
Pause	
	Pauses a sound (only when Main window is active).
Esc	
	Stops a sound (only when Main window is active).
Alt+F6	
	Switch between Main window and Device Controls window.
Ctrl+F6	
	Switch between Sound windows.

3.5 File Menu Commands

3.5.1 New

Creates a new sound file with attributes you specify. These attributes are discussed in **Appendix A**. Note that GoldWave allows you to create and edit sounds that may not be playable with your audio hardware. Use the **Options | Output device** dialog to find compatible attributes.

3.5.2 Open (Open Button) and Supported File Formats

Presents you with a list of files. By default, all files with the extension ".wav" are listed. You can change the extension to search for different sound formats. After you select a file, it is opened and given its own Sound window. Its attributes are displayed in the information bar.

Sound files are similar to document files in that different word processors save files in different formats. GoldWave supports all the sound formats listed in Table 2.

Table 2: Supported File Formats

Extension	Comments
.au	

	Sun or NeXT format. Supports 8 & 16-bit linear and law encoded files. Encoded files are converted to 8 bits or 16 bits, depending on Options Settings selections (8 bits by default). The conversion process can take a minute or so for long sound files. The header block is preserved.
.iff	Amiga 8SVX format. The blocks NAME, COPY, ANNO, AUTH, and CHAN are all preserved.
.mat	Matlab file. It should contain only one variable. The data must be normalized (i.e. -1.0 to 1.0) for double precision data. If the wavedata variable is two dimensional, the data is assumed to be stereo. GoldWave saves data using "wavedata" as the variable name.
.snd	Raw or NeXT format. NeXT format is automatically detected. Raw format presents a dialog to specify the attributes.
.wav	This includes RIFF WAVE 8-bit and 16-bit PCM mono or stereo, law encoded files (MULAW), and Microsoft ADPCM compressed files. Only files with one data chunk are supported. Only the INFO chunk is preserved.
.voc	Sound Blaster files. Only 8-bit unpacked mono files are supported.

GoldWave will present an attribute dialog if the file is corrupted, unrecognized, or contains data without header information (raw files). This dialog allows you to specify attributes. If you don't know the file's format or attributes, experiment with trial-and-error. Generally, sounds will be noisy if the format or number of bits is incorrect, in which case you will have to close and re-open the sound using a different format.

If the sampling rate is wrong, the playback speed will be too fast or too slow. You can use **Effects | Playback rate** to quickly re-adjust this. **Appendix A** has more information about sound attributes.

3.5.3 Save

Saves the active sound file to the disk. The sound is saved using its original name and format. If memory or disk space is low, the file may not be saved successfully. GoldWave will inform you if this happens. If **Save** fails, try deleting some unneeded files and close other applications. Make sure that the file is saved successfully before closing GoldWave, otherwise data may be lost.

3.5.4 Save as

Allows you to save a sound file under a different name. If no file name extension is specified, an appropriate one will be appended depending on the format of the sound (as described in the information bar). To save a file in a different format, you should use the **Export**

command described below.

3.5.5 Export

This is a flexible command that lets you save a sound in a wide variety of formats. You are presented with a dialog containing two lists. Use the left list to specify the file format and the right list to specify the attributes. Note that not all attributes are supported by some formats. Pick the format first to see what attributes are available.

After you have selected the format and attributes, choose the **OK** button. You will then be presented with the standard file dialog where you can specify the filename for the exported file.

When exporting a file in certain formats such as Sun (.au), GoldWave automatically converts the data to a compatible sampling rate.

This command is useful when files must be used on platforms other than Microsoft Windows.

3.5.6 Exit

Closes all Sound windows and quits GoldWave. All sound playback / recording is stopped. You will be asked to save any changes.

3.6 Edit Menu Commands

This section assumes that you are familiar with the terms introduced in the **Editing Overview** section.

See also: **Options | Editing**

3.6.1 Undo (Undo Button)

Reverses the most recent change made to a sound. Only one undo is possible across all sound windows. The undo feature keeps a copy of the original data in a temporary file. This file is created in the temporary directory specified by the TEMP environment variable.

Since the undo feature has to save large amounts of data, you will notice a slight delay before each edit/effect command is performed. The undo feature can be disabled via the **Options | Settings** dialog.

3.6.2 Cut (Cut Button)

Removes the block from the active sound and puts it in the clipboard. The contents of the clipboard can then be pasted into a Sound window using **Paste**.

To cut:

- 1) Move the start and finish markers to the part of the sound you want to cut.
- 2) Choose **Cut** from the **Edit** menu or click on the **Cut** button.

3.6.3 Copy (Copy Button)

Copies the block into the clipboard. The block is not removed. The contents of the clipboard can then be put into a Sound window using **Paste**.

To copy:

- 1) Move the start and finish markers to the part of the sound you want to copy.
- 2) Choose **Copy** from the **Edit** menu or click on the **Copy** button.

3.6.4 Paste (Paste Button)

Places the sound in the clipboard into the active sound. The sound is inserted at the start marker's position. The length of the active sound is increased so that clipboard sound will fit. The clipboard sound is automatically converted to match the attributes of the active sound.

To paste:

- 1) Move the start marker to the place where you want to paste the clipboard sound.
- 2) Choose **Paste** from the **Edit** menu or click on the **Paste** button.

Tips:

By copying a small block and pasting it several times, a stutter effect can be achieved.

You can use Paste to convert stereo to mono by copying from a stereo file and pasting to a mono file.

If a sound format is not supported by your audio device, use **File | New** to make a new compatible sound. **Copy** from the original, then **Paste** to the new.

Technical Note:

Clipboard data is automatically up-sampled or down-sampled to keep pitch and playback speed constant. You can use **Effects | Playback rate** to change the playback speed before copying/cutting.

3.6.5 Mix (Mix Button)

This is a special editing feature for sounds. The sound in the clipboard is blended (superimposed) with the active sound. You are asked for the mix volume. This volume is applied to the clipboard sound as it is being mixed (see **Effects | Volume**).

To mix:

- 1) Move the start marker to the place where you want to mix the clipboard sound.
- 2) Choose **Mix** from the **Edit** menu or click on the **Mix** button.
- 3) Enter the volume to mix the clipboard (100 = normal volume).
- 4) Choose **OK**.

Tip:

With this feature you could add your own vocals to music.

Technical Note:

Clipboard data is automatically up-sampled or down-sampled to keep pitch and playback speed constant. You can use **Effects | Playback rate** to change the playback speed before copying/cutting.

3.6.6 Delete (Delete Button)

Permanently removes the block from the active sound. The block is not copied to the clipboard. Note that if only the left or right channel is active in a stereo sound, then only that channel is removed. Since it is not possible for one channel to be larger than the other, the end of the deleted channel is padded with silence (this is also true for **Cut**).

To delete:

- 1) Move the start and finish markers to the part of the sound you want to delete.
- 2) Choose **Delete** from the **Edit** menu or click on the **Delete** button.

3.6.7 Trim (Trim Button)

Permanently removes everything *outside* the block. The block is not affected.

To trim:

- 1) Move the start and finish markers to the part of the sound you want to keep.
- 2) Choose **Trim** from the **Edit** menu or click on the **Trim** button.

3.6.8 Insert space (Space Button)

Inserts some blank space in the active sound at the start marker's position. You are asked how long (in seconds) the space should be. This command can be used to increase recording time. It can also be used to insert a delay.

3.6.9 Select view (Sel Vw Button)

Moves the start and finish markers to the far left and far right of the graph. It selects the part of the sound shown in the graph.

Tip:

After editing or adding an effect, a quick way to play the entire sound is to click on the **Out** button, then the **Sel Vw** button, then the play button.

3.7 View Menu Commands

This section assumes that you are familiar with the terms introduced in the **Editing Overview** section.

The following commands allow you to get more (or less) detail about the actual shape of the sound. They are similar to Paintbrush's zoom commands. The Overview box near the

bottom of each Sound window gives you useful information about what section of the sound is currently in the graph.

When you have zoomed into (magnified) a part of the sound, a scroll bar will appear at the bottom of the Sound window so you can move to different parts of the sound while still keeping the same level of magnification. The current level of magnification is displayed in the Main window's information bar next to the word "Zoom".

If you opened a long sound, the graph will not be accurate unless you zoom in. Generally, the longer the length of time that is graphed, the less accurate the graph will be.

3.7.1 Zoom block (Block Button)

Increases detail by graphing only the part of the sound in the block. It essentially magnifies the block. By selecting a new block, you can zoom in even further.

To zoom in on the block:

- 1) Move the start and finish markers to the part of the sound you want to magnify.
- 2) Choose **Zoom block** from the **View** menu or click on the **Block** button.

3.7.2 Unzoom (Out Button)

Shows the entire sound in the graph. When sound files are opened, the graph is initially unzoomed.

3.7.3 Zoom user defined (User Button)

This magnifies the graph to the level specified in the **Options | Settings** dialog. This level can be changed to any level you find useful (see **Options | Settings**).

3.7.4 Zoom 20:1

Gives a very detailed graph of the sound, where one sample is graphed across twenty pixels on the screen (represented as 1:0.05 in the information bar). This level of magnification makes direct editing of single samples easy.

3.7.5 Zoom 5:1

Gives a detailed graph of the sound, where one sample is graphed across five pixels (represented as 1:0.2 in the information bar). At this level of magnification, direct editing of a series of samples is possible.

3.7.6 Zoom 1:1 (1:1 Button)

Gives an accurate, full detail graph of the sound at the start marker's position. One sample is mapped to one pixel. This level of magnification is good for finding defects (pops and clicks) in sounds.

3.7.7 Zoom 1:10

Gives a roughly accurate level of detail. Ten samples are mapped to one pixel. This level of magnification is good for setting the start and finish markers when editing.

3.7.8 Zoom 1:100

Gives a very rough level of detail. One hundred samples are mapped to one pixel.

3.8 Effects Menu Commands

This section assumes that you are familiar with the terms introduced in the **Editing Overview** section.

With these commands, you can change a sound or create new sounds. The commands are comparable to font menu commands in word processors. For example, using font commands, you can change the size of text. In GoldWave, using the "Volume" effect changes the "size" (amplitude/volume) of a sound wave.

See also: **Options | Editing**

Technical Note:

Even though the word 'volume' is used throughout this section, the amplitude is changed linearly.

3.8.1 Cross-fade

Recalculates samples in the block such that the sample value at the start marker is "faded" to the sample value at the finish marker. This is useful for removing pops and clicks. It is usually used on small blocks.

3.8.2 Echo

Produces an echo effect with the block. The echo delay and volume can be entered after choosing this command. You may want to try the default values first.

The longer the delay, the longer it takes for the echo to "bounce" back. Try values less than 0.1 for a large room, 0.3 for a baseball stadium, above 0.3 for a canyon echo.

The lower the volume, the quieter the echo will be. Values less than 50 give good results.

If you check the **Reverb** check box, the echo will be regenerated at intervals specified by the delay. This means that if the delay was 0.1 seconds, the echo at 0.1s is regenerated at 0.2s, and this new echo is regenerated at 0.3s, etc. The volume is applied to each regenerated echo so that if the volume was 50%, the first echo volume is one half the original, the second echo volume is one quarter, etc.

To make the echo sound correct, GoldWave extends the effect slightly outside the block boundary. This may increase the length of the sound. If you do not want to modify sound outside the block or you do not want the length of the sound to change, use the **Options | Editing** dialog to disable modification outside the block.

To add an echo:

- 1) Move the start and finish markers to the part of the sound you want to add

- an echo.
- 2) Choose **Echo** from the **Effects** menu.
 - 3) Enter the delay time.
 - 4) Enter the volume.
 - 5) Check the **Reverb** check box if reverb is desired.
 - 6) Choose **OK**.

3.8.3 Exchange channels

Exchanges the left and right channels (i.e. the right channel becomes the left channel and the left channel becomes the right channel).

3.8.4 Filter

Brings up a pop-up menu listing several filter related commands. Filters are used to remove a range of frequencies from the sound. They can produce a variety of effects on the sound.

Tip:

Before down-sampling (**Effects** | **Resample**), it is a good idea to lowpass filter the sound to prevent aliasing.

3.8.4.1 Lowpass 50%

Passes the sound through a lowpass filter. Frequencies above one half of the Nyquist (or one quarter the sampling rate) are removed. For example, a sound recorded at 44100Hz will have frequencies above 11025Hz removed.

Technical Note:

This filter is a sixth order Butterworth filter with a cutoff at 0.47 of the Nyquist frequency. To get a higher order, you can use this command several times.

3.8.4.2 Lowpass 25%

Passes the sound through a lowpass filter. Frequencies above one quarter of the Nyquist (or one eighth the sampling rate) are removed. For example, a sound recorded at 44100Hz will have frequencies above 5512.5Hz removed.

Technical Note:

This filter is a third order Inverse Chebyshev filter with a cutoff at 0.25 of the Nyquist frequency and stopband 15dB down from passband.

3.8.4.3 Highpass 10%

Passes the sound through a highpass filter. Frequencies below one tenth of the Nyquist (or one twentieth the sampling rate) are removed. For example, a sound

recorded at 44100Hz will have frequencies below 2205Hz removed.

Technical Note:

This filter is a third order Inverse Chebyshev filter with a cutoff at 0.10 of the Nyquist frequency and stopband 15dB down from passband.

3.8.4.4 Highpass 25%

Passes the sound through a highpass filter. Frequencies below one quarter of the Nyquist (or one eighth the sampling rate) are removed. For example, a sound recorded at 44100Hz will have frequencies below 5512.5Hz removed.

Technical Note:

This filter is a third order Inverse Chebyshev filter with a cutoff at 0.25 of the Nyquist frequency and stopband 15dB down from passband.

3.8.5 Invert

Reflects the sound about the time axis. The sound is essentially turned up-side-down. This produces no noticeable effect in mono sounds and has a slight effect in stereo sounds. Inverting a single channel of a stereo sound produces an "in" or "out" effect.

3.8.6 Mechanize

Adds a robotic or mechanical characteristic to sounds. The percentage of quality can be entered after selecting this command. Low values produce an untuned radio effect. Higher values give a rough distorted effect.

To mechanize part of a sound:

- 1) Move the start and finish markers to the part of the sound you want to mechanize.
- 2) Choose **Mechanize** from the **Effects** menu.
- 3) Enter the quality percentage, then choose **OK**.

3.8.7 Pan

Brings up a "shape" dialog. The shape dialog consists of a graph window and several buttons. The graph window initially contains a single horizontal line with two endpoints (shown as square boxes). By clicking the left mouse button anywhere inside the graph window, you can insert new points to "bend" the line into a variety of zig-zag shapes. You can move points up and down by clicking and dragging them to a new location.

To remove a point, click the right mouse button over the point. To remove all the points (except the endpoints), choose the **Clear** button.

For panning, the shape dialog is divided into green and red regions, representing the left and right channels respectively. The line, initially located between the regions, represents the center for panning. By bending and/or moving the line, you can alter the sound's left/right balance or pan to and from each channel.

3.8.8 Reverse

Reverses a sound so that it plays backward. Now you have an easy way to listen to all those "satanic verses".

3.8.9 Silence

Erases the block. The sound in the block is muted (set to zero volume).

See also: **Effects | Volume**

3.8.10 Volume

Brings up a pop-up menu containing several volume related commands. Volumes are usually specified by a percentage of the sound's original amplitude. Unlike the volume control in the Device Controls window, these commands actually alter the sound's data to change the volume.

See also: **Effects | Pan** and **Effects | Silence**

3.8.10.1 Change

Modifies the block so that it sounds louder or softer. The percentage of change can be entered after selecting this command. Values less than 100% make the sound softer and values greater than 100% make the sound louder. A value of 25% reduces the volume by one quarter. A value of 200% doubles the volume.

By using independent left/right editing (**Options | Editing**), the internal balance of a sound can be adjusted.

To change the volume of part of a sound:

- 1) Move the start and finish markers to the part of the sound you want to change.
- 2) Choose **Change** from the **Effects | Volume** pop-up menu.
- 3) Enter the volume percentage, then choose **OK**.

3.8.10.2 Fade in

Gradually increases the volume throughout the block. The initial volume percentage can be entered after choosing this command. A value of 25% starts with one quarter volume and fades in to full volume. A value of 0% starts at silence and fades in to full volume.

To fade in part of a sound:

- 1) Move the start and finish markers to the part of the sound you want to fade in.
- 2) Choose **Fade in** from the **Effects | Volume** pop-up menu.
- 3) Enter the initial volume percentage, then choose **OK**.

3.8.10.3 Fade out

Gradually decreases the volume throughout the block. The percentage of fade can be entered after choosing this command. The fade percentage is the amount that the volume should decrease. A value of 100% fades to complete silence. A value of 50% fades to one half of the original volume.

To fade out part of a sound:

- 1) Move the start and finish markers to the part of the sound you want to fade out.
- 2) Choose **Fade out** from the **Effects | Volume** pop-up menu.
- 3) Enter the fade percentage, then choose **OK**.

3.8.10.4 Maximize

Searches the sound for the maximum percent that the volume can be increased without producing distortion or clipping errors. This volume is then applied to the sound. If the percentage is above 2000% or equal to 100%, no changes will be made.

3.8.10.5 Shape

Brings up the "shape" dialog (see **Effects | Pan**). With the shape dialog, you can define the amplitude envelope of a sound. The shape line is initially horizontal at 100%. By bending/moving the line, you can easily change the volume of different parts of the sound.

Several examples of panning shapes are given below.

3.8.11 Playback rate

Changes the playback rate of the entire sound. The sound will play faster (or slower) and its pitch will be higher (or lower). Essentially, this just changes the first number in the top information bar. Values of 11025, 22050, and 44100 are the recommended values. When recording, the new rate is used as the sampling rate.

See also: **Effects | Resample**

3.8.12 Resample

Changes the sampling rate of the entire sound. Unlike **Playback rate**, this command recalculates and optionally interpolates all the data so that the pitch and playback time are not affected. You are prompted to enter a new rate. Values of 11025, 22050, and 44100 are the recommended values. If you check the **Interpolate** check box, interpolation is used to recalculate sample values at fractional positions. This usually gives a higher quality result.

Tip:

If you have a sound recorded at 44100Hz and do not need CD quality, you can save large amounts of disk space by resampling the sound to 22050Hz or 11025Hz. This reduces the size

by 2:1 or 4:1.

Before down-sampling (converting 22050Hz to 11025Hz, for example), the data should be lowpass filtered to prevent aliasing. See **Effects | Filter**.

See also: **Effects | Playback rate**

3.8.13 Expression evaluator

Allows you to create new effects or generate signals. See **Appendix B** for more details.

3.9 Options Menu Commands

3.9.1 Input device

Presents a list of input devices that were installed with Control Panel's Drivers applet. If no input devices are installed, a warning message is displayed and recording is disabled.

The list shows you the name of the device/driver. Above the list is a table that tells you what sound attributes/formats the driver supports. All boxes with check marks in them are supported by the currently selected driver. You can select any driver to see its capabilities.

Tip:

Since you can choose the input device, you can have several devices/drivers installed at once (provided they do not interfere with each other).

3.9.2 Output device

Presents a list of output devices that were installed with Control Panel's Drivers applet. If no output devices are installed, a warning message is displayed and playback is disabled.

The list shows you the name of the device/driver. Above the list is a table that tells you what sound attributes/formats the driver supports. All boxes with check marks in them are supported by the currently selected driver. You can select any driver to see its capabilities.

Tip:

Since you can choose the output device, you can have several devices/drivers installed at once (provided they do not interfere with each other). For example, it is possible to have the PC-Speaker driver and a Sound Blaster driver installed at the same time.

3.9.3 Aux volume

Sets the output volume on an auxiliary device. This may be useful when recording from a CD-ROM. A list of installed auxiliary devices is displayed or a warning message is given if no auxiliary devices are present.

To change auxiliary output volume:

- 1) Select a device. If the device supports volume control, the "left" control will be enabled. If the device supports stereo volume control, both the "left" and

- "right" controls will be enabled.
- 2) Adjust the volume with the "left" and/or "right" controls.

3.9.4 Settings

Allows you to change and configure several features in GoldWave, as described below.

Sound path

Sound path is used by **File | Open**. You can specify the directory where you keep your sound files. If you enter "d:\sound\wave", for example, the **File | Open** dialog will automatically list the files in that directory whenever you run GoldWave.

Save window positions

If **Save window positions** is checked, the Main window and Device Controls window positions are saved when you exit GoldWave. This means you will not have to adjust window positions every time you run GoldWave.

Decode files to 8-bits

If **Decode files to 8-bits** is checked, any encoded file will be converted to 8-bit PCM format. This will cause some loss of quality, but will allow playback on a larger variety of audio devices. If the box is not checked, the data is converted to 16-bit PCM format. If your audio device supports 16-bit PCM files, uncheck this box.

Quick selecting

The **Quick selecting** check box allows you to choose which method to use when selecting a block. See **Editing Overview** for more information.

Graph averaging

If this box is checked, the graphs displayed in Sound windows will be drawn using an average of samples rather than just one sample. While averaging can give a cleaner display, the amplitude will be somewhat inaccurate. Averaging is not applied when the zoom level is 1:10 or closer (such as 1:1, 5:1, or 20:1).

Allow undo

Enables or disables the undo feature. Since the undo feature can use large amounts of disk space and slow editing, you may want to disable undo by unchecking this box.

User zoom

This is where you specify the level of zoom for the **View | Zoom user defined** command (**User** button). Values between 0.01 to 100 are valid. See the **View Menu Commands** section for more information.

3.9.5 Editing (Chan Button)

Sets the behaviour of editing and effects. These options apply only to the currently active sound. This means you can have different "behaviours" for different Sound windows.

Channel editing

The **Channel editing** group lets you control which channel (or channels) will be edited for stereo sounds. For example, if you choose "Left channel", all effects and editing will only change or use the left channel. You can then use **Edit | Copy** to extract the left channel into the clipboard or use **Edit | Mix** to mix the clipboard with the left channel and so on.

The channel editing status is displayed in the bottom information bar on the far left. The left and right Sound window graphs are highlighted appropriately.

The **Chan** button provides a quick way to control channel editing. It cycles through the channels with each click. Initially, stereo editing is enabled. After one click, left channel editing is enabled. The next click enables right channel editing. Another click enables stereo editing and the cycle starts again.

Modify outside block

The **Modify outside block** check box allows you to protect sounds outside the block from being changed during certain effects (**Effects | Echo** uses this).

3.10 Window Menu Commands

The **Window** menu contains a list of all Sound windows currently opened. The commands organize Sound windows.

3.10.1 Tile

Arranges Sound windows side-by-side so that they are all fully visible.

3.10.2 Cascade

Layers Sound windows on top of each other so that their title bars are visible.

3.10.3 Arrange icons

Arranges minimized Sound window icons in rows on the bottom of the Main window.

3.11 Help Menu Commands

3.11.1 Contents

Starts Window's on-line help utility and gives a list of contents for GoldWave help.

3.11.2 About

Displays version/registration information. Shows how much memory is available under the GoldWave icon

4 General Information

4.1 Warranty, Trademarks, and Copyrights

GoldWave ("the package") includes the following software and documentation:

GOLDWAVE.EXE	Executable file
GOLDWAVE.HLP	On-line help
GOLDWAVE.WRI	Manual
README.TXT	Extra information
EXPRESS.TXT	A collection of expressions
ORDER.TXT	Order form
BWCC.DLL	Borland's OWL DLL

The package is provided as is, without warranty of any kind. The author shall not be liable for damages of any kind. Use of this software indicates you agree to this.

The package and this documentation are copyright 1993-1994 by Chris S. Craig.
All rights reserved.

CompuServe is a registered trademark of CompuServe Incorporated.

GoldWave is a trademark of Chris S. Craig.

IBM is a registered trademark of International Business machines.

MATLAB is a trademark of The Math Works Incorporated.

Sound Blaster is a trademark of Creative Labs Incorporated.

Windows is a trademark of Microsoft Corporation.

All other trademarks/registered names acknowledged.

4.2 Contacting the Author

I would appreciate any comments, suggestions, and/or bug reports. If you have some new ideas for features or ways to improve GoldWave (and the documentation), you can contact me at the addresses below.

Please do not ask programming or development questions such as "Can I have the source code?", "How can I make an oscilloscope?", etc. These type of questions can take hours to respond to. Therefore, I will only answer questions relating to the use of GoldWave.

You can send Internet e-mail to:

chris3@garfield.cs.mun.ca

This e-mail address can be used on other networks. For example, CompuServe users can send e-mail as follows:

>INTERNET:chris3@garfield.cs.mun.ca

I prefer e-mail.

Regular mail:

Chris Craig
P.O. Box 51
St. John's, NF
CANADA A1C 5H5

Thank-you for using GoldWave.

4.3 Reporting Problems and Bugs

In the event that you encounter a problem, please try to determine the circumstances that cause it. The more information you can provide (file attributes, menu commands used, etc.), the better the chances are of finding and fixing the problem. See **Contacting the Author** for more information.

4.4 Revision History

v2.00:

New effects: cross-fade, resample, mechanize, pan, volume shaping, fade in, echo reverb, and 4 filters.

Added trim to editing.

More view commands: user, 20:1, 5:1, 1:10, 1:100.

Enhanced Device Controls window: dual oscilloscopes, setup dialog.

Setup for Device Controls: graph type, controls, looping, graph length, positioning, other.

New tool bar buttons: Trim, User, Chan, Undo.

Added second information bar for block info.

Added Spectrum analyzer (easier said than done).

Added direct editing with the mouse.

Improved keyboard support.

Added Export command.

Added file support for ADPCM, MULAW, unsigned 16-bit, byte swap, Matlab.

Added progress indicator in the information bar.

Changed the expression evaluator, added edit buttons and "f" variable.

Fixed bugs: Stereo echo, balance, PC Tools loading, refresh, etc.

v1.00:

First release. Includes: mono/stereo, 8/16 bit, uLAW, independent channel editing, up to

5 sound windows, detached Scope and Controls window, tool bar, information (attributes) bar, MDI, expression evaluator, and several effects.

4.5 Future Versions

Here's a list of possible future enhancements to GoldWave:

- As soon as I can figure out the magic behind pitch shifting & time scaling, I'll add it.
- OLE support (quite painful)
- Support for large file (this is really the job of the OS!)
- Markers, A-Law support, Files for expression evaluator
- Got any ideas?

4.6 Acknowledgements

I thank the following people for their suggestions and/or help:

Keith Peterson, SimTel Repository

Ari Hovila, Garbo Windows Archives

David Komatsu, Sound Site Newsletter and Archives

Guido van Rossum, Audio Formats FAQ

Steve Sampson & Steve Haehnichen, for their clearly written FFT code I found on the net. A great help for understanding FFTs and testing GoldWave's 386 spectrum analyzer.

Josh Templin, for graphical panning and other good ideas

and the many others who gave helpful feedback!

My thanks to the registered users, who made this version possible.

Appendix A: Audio Attributes/Terminology

A.1 Digital Sound Attributes

A digital sound (sampled sound) is composed of thousands of pieces of data, called samples. Each sample holds the loudness (amplitude) of a sound at a given instant in time. This is similar to computer graphics where each point of light (pixel) has a certain brightness and location. All these points combine to make a picture. In digital sound, all the samples combine to make a sound. There are several attributes that determine the quality and quantity of digital sound. They are: **sampling rate**, number of **bits**, and number of **channels**.

The **Sampling rate** is the number of times, per second, that a sample is recorded. It is measured in Hertz (seconds⁻¹, Hz). A high sampling rate will yield a high quality of digital sound in the same manner that high graphics resolution will show better picture quality.

Compact disks, for example, use a sampling rate of 44100Hz, whereas telephone systems use a rate of only 8000Hz.

The rate to use depends upon the type of sound and the amount of memory and disk space you have available. Higher rates consume larger quantities of memory. In the above example, the compact disk requires over 5 times the amount of storage as the telephone system for the same digital sound.

Certain types of sounds can be recorded at lower rates without loss of quality. Some standard rates are listed in Table A.1 at the end of this section.

The number of **bits** determines how accurately the amplitude of a sample is recorded. The two most common are 8-bit and 16-bit formats. In an 8-bit sample, there are 256 different levels of amplitude possible. 16-bit samples have 65536 levels. To compare the difference, let's say that you are a teacher grading tests and you can use one of two marking schemes. In scheme #1, the mark is out of 10. In scheme #2, the mark is out of 1000. All marks must be rounded off (no fractions allowed). If a student gets two thirds of the questions right, then in scheme #1, the grade will be 7 out of 10. In scheme #2, the grade will be 667 out of 1000. Obviously, scheme #2 is more accurate. In digital sound, low levels of accuracy can cause quantization errors, as discussed in the next section.

One drawback of using 16 bits is that it uses twice as much storage space. This is shown in Table A.1 at the end of this section.

The number of **channels** also affects the quality and quantity of digital sound. A single channel sound, referred to as a mono sound, contains information for only one speaker (i.e. similar to AM radio). A two channel sound (stereo sound) contains data for two speakers (i.e. similar to FM stereo). Stereo sounds require twice as much storage as mono sounds.

Table A.1: Sound attributes

Format/Attributes	Quality/Sound type	Quantity/Storage per second
11025Hz, 8-bit, mono	Fair quality. Good for speech and low pitch sounds.	11025 bytes
11025Hz, 8-bit, stereo	Fair quality stereo.	22050 bytes
11025Hz, 16-bit, mono	Less noise.	22050 bytes
11025Hz, 16-bit, stereo	Stereo, less noise.	44100 bytes
22050Hz, 8-bit, mono	Good quality. Good for music and relatively complex sounds.	22050 bytes
22050Hz, 8-bit, stereo		

	Good quality stereo.	44100 bytes
22050Hz, 16-bit, mono		
	Very good quality.	44100 bytes
22050Hz, 8-bit, stereo		
	Very good quality stereo.	88200 bytes
44100Hz, 8-bit, mono		
	High quality. Good for all sounds.	44100 bytes
44100Hz, 8-bit, stereo		
	High quality stereo.	88200 bytes
44100Hz, 16-bit, mono		
	Excellent quality.	88200 bytes
44100Hz, 16-bit, stereo		
	Excellent quality stereo. Large storage requirements.	176400 bytes

A.2 Problems with Recording

There are five potential problems when recording sound: aliasing, clipping, quantization, internal noise, and DMA/memory boundaries.

Aliasing occurs when the sampling process does not get enough data to correctly determine the shape of the sound wave. The digital sound will have high or low pitch tones that never existed in the original sound (see **Appendix B**, at the end of the **Signal Generation** section). This problem can be eliminated by using higher sampling rates (or by using anti-aliasing filters).

Clipping errors occur when the sampled amplitude is outside the range of valid values. For example, if the range is -1.0 to 1.0, and a value of 1.2 is sampled, then this value must be clipped to 1.0. A way to reduce these errors is to lower the input volume when recording.

Quantization errors occur when the sample is rounded to the nearest level of amplitude. This can be explained by using the "marking schemes" example of the previous section. The number two thirds (2/3) is represented by 7/10 in scheme #1. This gives a quantization error of:

$$| 7/10 - 2/3 | = 1/30$$

Similarly, in scheme #2, the quantization error is:

$$| 667/1000 - 2/3 | = 1/3000$$

Clearly, scheme #2 has the smallest error. Therefore, using 16 bits instead of 8 bits is a good way to reduce quantization errors.

The other two recording problems deal with computer hardware. To minimize internal noise, make sure your audio card is installed as far away from your graphics/monitor adaptor card as possible. If you use a microphone, keep it away from your monitor and computer fan. Remember to use shielded cables.

DMA and memory boundaries can cause loss of quality during recording and playback (pure sine waves rarely play without some "defects"). Some of the latest audio cards and drivers may reduce this problem. When recording at high sampling rates, use a RAM disk.

A.3 Frequency Spectrums

GoldWave features a built-in frequency spectrum analyzer in the Device Controls window. Essentially, this allows you to see what frequencies (or pitches) are present in a sound. A rainbow is an example of a frequency spectrum of visible light. The sun's light is broken down into a set of fundamental colours.

GoldWave generates the spectrum by using a radix-2 fast Fourier transform (FFT). FFTs require intensive computations, making them somewhat unsuitable for real-time applications. To speed up these computations, GoldWave makes extensive use of 32-bit 386 assembly language instructions. For accuracy, 64-bit temporary results are used. GoldWave can perform a 512 point FFT 18 times per second on a 20MHz 386. This can generate a real-time spectrum for monaural sounds sampled at about 10000Hz.

GoldWave optionally applies a windowing function to the data before performing the FFT (see **Setup** in the **Device Controls Overview** section). This reduces errors that occur when dividing the data into small chunks (such as 512 point chunks). The Hamming window is used.

To make the spectrum more realistic to human hearing, magnitudes are scaled logarithmically. This means that if one frequency "sounds" twice as loud as another, it will be graphed with twice the height.

During playback with a short spectral oscilloscope, the following operations are performed each time the oscilloscope is updated:

- 1) The current position is obtained or estimated.
- 2) The position is drawn on the Sound window's graph.
- 3) The sample data is windowed.
- 4) The FFT is performed.
- 5) The logs of the magnitudes are calculated.
- 6) The result is converted to screen coordinates.
- 7) The graph is drawn.

All this requires a significant amount of CPU time. On a 20MHz 386 without accelerated graphics, about 5 to 10 frames per second are possible. Playing a 16-bit, 44100Hz, stereo file with a long spectral oscilloscope is likely to consume nearly 100% of the CPU time. Under some circumstances, this may prevent dialogs from being displayed. If you notice that a dialog is taking an unusually long time to appear, press the pause or stop button to free up some CPU time.

Appendix B: Expression Evaluator

B.1 Overview

The Expression Evaluator is a versatile tool for manipulating and generating audio data. Before reading this section, you may want to try a few of the expressions given in the **Signal Generation** section. After you select **Expression evaluator** from the **Effects** menu, you are presented with a dialog that is similar in appearance to a calculator.

On the left side, above the number buttons, is the Destination box. The destination is defined as the Sound window where results of the evaluation will be placed. The box's drop-down list contains a list of all Sound windows in the form "wavex", where x is the number on a Sound window as given in its title bar. For example, a Sound window with the title "3 - HELLO.SND" would appear as "wave3" in the list. By default, the destination is set to the active Sound window. You can change the destination, if more than one Sound window is opened, by using the up and down keys or by selecting it with the mouse from the drop-down list.

A large Expression box is located in the middle of the dialog. This is where the expression is entered. To enter an expression, you can either type it in or click on the "calculator" buttons with the mouse (you can also use the cut, copy, paste, and delete buttons or keystrokes listed below). A list of valid operations and functions is given in a following section.

Near the bottom of the dialog are three boxes where the initial values of the variables "n", "t", and "f" may be entered.

After you have specified the destination, expression, and initial values, choose the **Start** button (or just press the "enter" key) to begin evaluation. If you entered an expression incorrectly, a message will be displayed in the **Message** box. The **Message** box also gives the progress during evaluation. Since the evaluation process takes time, you can stop it at any time with the **Cancel** button. Pressing the **Cancel** button a second time will close the Expression Evaluator dialog.

Tip:

You can copy, cut and paste expression in the Expression box using the following keystrokes:

Copy	Ctrl+Ins
Cut	Shift+Del
Paste	Shift+Ins

You can use Window's Notepad to keep a set of useful expressions and copy and paste them in to the Expression box when needed. You can also copy and paste expression from the on-line help. A co-processor can dramatically increase the speed of the evaluation, since GoldWave uses floating point mathematics.

B.2 Evaluation Range, Variables, and Constants

Expressions are evaluated over the destination Sound window's block. You must set start and finish markers before you open the evaluator dialog (Refer to **Editing Overview** for more information). Let's assume we have the following before opening the dialog:

Title bar: "1 - Untitled"
Sampling rate: 8000Hz
Start marker at 0.5 seconds
Finish marker at 1.2 seconds

After we open the dialog, we set:

Destination: "wave1"
Expression: "sin(t)"

The initial value for t is automatically set to 0.5. If you now choose OK, the expression "sin(t)" will be evaluated from t = 0.5 to t = 1.2 in steps of 1/8000 seconds. You will get:

wave1(0) = sin(0.500000)
wave1(1) = sin(0.500125)
wave1(2) = sin(0.500250)
...
wave1(3) = sin(1.200000)

Time Variable and Constant

The variable "t" is the time variable and, as demonstrated above, it changes throughout the start/finish range. The step of change is determined by the inverse of the destination's sampling rate. This inverse is assigned to a special constant, "T". In the above example, T = 0.000125 (i.e. 1/8000). Constant "T" and variable "t" are automatically reset when a new destination is selected.

User variable

The variable "f" can be set to any value you choose. In many cases, this value is used as a frequency, hence the letter "f" is used. For example, if you entered the expression:

$\sin(2*\pi*f*t)$

You can then generate any sine wave by specifying a frequency in the "f" edit field. This value does not change through evaluation, but stays at the value you assigned to it.

Sample Index Variable and Total Constant

The variable "n" is the sample index variable. It is always an integer value. During evaluation, this variable is continuously incremented in steps of 1. The sample index variable is used whenever an operation must be performed on an existing sample. A value of n = 0 (the initial default value) is the index of the sample at the start marker's position of the given signal. For example, the expression

$\text{wave2}(n) + \text{wave3}(n+100)$

takes the wave2 sample at wave2's start marker's position and adds it to the wave3 sample 100 samples after wave3's start marker's position. If a sample index is outside the wave's block, the value is 0 (i.e. $\text{wave2}(n-50) = 0$ while $n < 50$).

You must be careful when indexing signals with different sampling rates to get a desired effect. Assume wave1 is a voice recorded at 11025Hz and wave2 is music recorded at 22050Hz. If you want to sum (superimpose) the two signals with wave1 as the destination, then the expression

$$\text{wave1}(n) + \text{wave2}(n*2)$$

must be used (ideally, wave2 would have to be low-pass filtered first). Whereas, if wave2 is the destination, the expression should be

$$\text{wave1}(n/2) + \text{wave2}(n)$$

A constant "N" is set to the total number of samples in the evaluation range. It has several uses, such as reversing a sample. Assume wave2 is a new sound that has the same sampling rate and length of wave1. By setting the destination to wave2 and using

$$\text{wave1}(N-n-1)$$

wave2 will be the reverse of wave1.

Conversion Between Variables

$$n = (t - \text{initial value})/T$$

$$t = n*T + \text{initial value}$$

B.3 Valid Operations and Functions

The following table summarizes the operations and functions included in the evaluator.

Table B.1: Evaluator Operations and Functions

Label	Operation, function
(,)	Operation, function
(,)	Parenthesis
+, *, -, /	Add, multiply, subtract (negate), and divide
%	Modula operator (remainder)
^	To the power of, y^x
pi	Constant (3.14159...)
cos	Cosine
sin	Sine

	Sine
tan	Tangent
acos	Arccosine
asin	Arcsine
atan	Arctangent
cosh	Hyperbolic cosine
sinh	Hyperbolic sine
tanh	Hyperbolic tangent
sqrt	Square root
abs	Absolute value
log, ln	Log base 10, natural logarithm
exp	Exponential base e
step	Unit step (0 for $t < 0$, 1 for $t \geq 0$)
int	Integer value
rand(n)	Random number between 0 and n
wavex(n)	Signal value at n

B.4 Signal Generation

Generation expressions of several signals are listed below. Words given in italics are values that you must enter.

To try one of the following expression, perform the following steps:

- 1) Choose **File | New**.
- 2) Choose **OK**.
- 3) Choose **Effects | Expression evaluator**.
- 4) Type in the expression as given in the example. For example:
 $\sin(2*\pi*261.7*t)$
- 5) Choose **Start**.
- 6) Wait for the "Finished" message.
- 7) Choose **Cancel**.

8) Play the sound.

Type: Sine wave

General Expression:

$$\sin(2*\pi*frequency*t)$$

Examples:

Middle C:

$$\sin(2*\pi*261.7*t)$$

Telephone dial tone for "5":

$$(\sin(2*\pi*773*t) + \sin(2*\pi*1336*t)) / 2$$

Type: Square wave

General Expression:

$$\sin(\pi*\text{int}(2*frequency*t))$$

Examples:

1000Hz tone:

$$\sin(\pi*\text{int}(2*1000*t))$$

Middle C:

$$\sin(\pi*\text{int}(2*261.7*t))$$

Type: Saw (triangle) wave

General Expression:

$$1 - 2*\text{abs}(1 - 2*frequency*t\%2)$$

Example:

200Hz tone:

$$1 - 2*\text{abs}(1 - 2*200*t\%2)$$

Other examples of signals:

White noise:

$$1 - \text{rand}(2)$$

Exponential decay of 500Hz cosine wave to 50%:

$$\cos(2*\pi*500*t) * (0.5*\exp(-t) + 0.5)$$

Slow sweep (up to 20kHz):

$$\sin(2*\pi*160*(t\%5)^3)$$

The Sweep is a great way to see and hear aliasing and quantization if the destination signal has a low sampling rate. The ideal output should be a single tone increasing in pitch until it is nearly outside the range of human hearing. In reality, you will be able to hear two or more tones, some rising in pitch and some falling.

B.5 Filters

One of the easiest ways to create digital filters is to use MATLAB™ (The Student Edition of MATLAB, by The Math Works Inc., published by Prentice-Hall, ISBN 0-13-855974-0). It has many built-in commands that generate filter coefficients. The coefficients can then be used in the expression evaluator.

Example of a Low-Pass Filter

In preparation for down-sampling, you can use MATLAB to generate the coefficients of a 4th order Butterworth low pass filter that will remove noise above half the Nyquist frequency (1/4 the sampling rate). Enter:

```
[b,a] = butter(4, 0.5)
```

The result should be:

```
b = 0.0940 0.3759 0.5639 0.3759 0.0940  
a = 1.0000 0.0000 0.4680 0.0000 0.0177
```

To implement this filter in GoldWave, let's assume that the sound to be filtered is in the Sound window titled "1 - SOUND.SND".

- 1) Use **File | New** to create a new Sound window with the same sampling rate as "1 - SOUND.SND". The new Sound window will be titled "2 - Untitled".
- 2) Make sure the length of the new sound is as long (or longer) than the original. If it is not, use **Edit | Insert space** to increase its length.
- 3) Use **Effects | Expression evaluator** to open the expression evaluator dialog.
- 4) Set the destination to wave2.
- 5) Enter the following expression:
$$\text{wave1}(n)*0.0940 + \text{wave1}(n-1)*0.3759 + \text{wave1}(n-2)*0.5639 + \text{wave1}(n-3)*0.3759 + \text{wave1}(n-4)*0.0940 - \text{wave2}(n-2)*0.4860 - \text{wave2}(n-4)*0.0177$$
- 6) Choose **Start** (You might want to get a coffee while you wait).

Appendix C: Trouble Shooting and Q&A

C.1 Trouble Shooting

Problem	Cause/Solution
Cannot open large files	When GoldWave opens a sound file, it reads all the data into RAM. Therefore, the size of the file you can open depends upon how much RAM your system has.
Cannot play sounds	GoldWave or audio device/driver is incorrectly installed. Make sure a WAVE driver is installed in the Control Panel's Drivers applet. Check that Window's Sound Recorder can play sounds. If it doesn't, the driver is not installed correctly. Make sure that your audio device is capable of playing the sound. For example, an 8-bit audio device cannot play a 16-bit sound.
Cannot record sounds	See above. Make sure your audio device/driver can record sounds. The sound may be paused. Click on the stop button. Cannot use the stop or pause button
Oscilloscope does not work	The audio driver is synchronous. This means that Windows (and GoldWave) loses control until the sound has finished playing.
System freezes or crashes or a General Protection fault occurs	Make sure that you have a 386 or better system. GoldWave will not run on a 286 system. You may have encountered a problem. Contact the author for more information.
Oscilloscope is out of sync	Try using a different positioning method in the Device Controls Setup dialog. Many audio drivers return inaccurate 'current' positions. Make sure you have the most up-to-date device driver.
Periodic popping or clicking while playing a pure sine wave	Most audio devices/drivers make periodic pop/clicks between DMA transfers and/or memory boundaries. It is most noticeable when playing pure sine waves. Pops and clicks can occur at the beginning or ending of a sound if the first or last sample is not 0 (silence). Fading in/out a small block can sometimes fix this.
Sound won't play for more	

than a few seconds.

Make sure your driver is configured to play for more than few seconds. The PC-Speaker driver, for example, will play for only 4 seconds unless you configure it to play longer.

C.2 Common Questions and Answers

How can I play a 16-bit sound on my 8-bit sound card?

The following steps convert a 16-bit sound to an 8-bit sound:

- 1) Open the 16-bit sound.
- 2) Choose **Edit | Copy** or click the **Copy** button.
- 3) Choose **File | New**.
- 4) Set the sample rate and time to the values in the information bar.
- 5) Make sure **8 bit** is selected, then choose **OK**.
- 6) Choose **Edit | Mix** or click the **Mix** button.
- 7) Enter 100 as the volume, then choose **OK**.
- 8) Choose the play button.

You can also use the **File | Export** command to save a 16-bit sound as 8-bits.

How can I play a stereo sound on my mono sound card ?

You can use the same procedure above, except make sure **Mono** is selected.

You can also use the **File | Export** command to save a stereo file as mono.

How can I play/record large wave files?

GoldWave is limited by the amount of RAM and virtual memory you have. Since Windows does not allow you to use more virtual memory than RAM, there is no way to play large files in GoldWave. GoldWave is theoretically capable of opening a 64 megabyte file, provided the operating system supported (and fully *used*) a 64 megabyte swap file.

You could try using Media Player in the Accessories group.

How can I slow down the oscilloscope?

You can set the "time" of the fast forward button to a small value (i.e. 0.01) using the Device Controls Setup dialog. Play the sound, then click on the fast forward button.

You can also adjust the speed (**Spd**) scroll bar.

How are you?

Fine, thanks.

Can I edit individual samples with the mouse?

Yes, see **Editing Overview, Sample Editing with the Mouse**.

Can I customize the colours for graphs?

Yes, but you have to edit the GOLDWAVE.INI file. You can add the following lines at the end of the file to change graph backgrounds to white:

```
[Colours]
Graph Background=255, 255, 255
Start Marker=0, 0, 0
Finish Marker=0, 0, 0
Scope Background=255, 255, 255
Scope Grid=0, 0, 0
Amplitude Scope=0, 0, 0
Spectral Scope=128, 0, 0
```

Will there be any new versions?

Yes, version 3.0 is being planned. If you have any suggestions, please pass them along. It is very important that you register this version, since the cost of developing GoldWave is rather expensive. For version 3.0, I would like to upgrade my PC (a 20MHz 386 is just not suitable for developing this type of application).

Are there any plans of a multilingual version?

Not yet. The effort and resources needed to develop a multilingual version is far too much for one person to handle. Maybe in the future.

Can I play MOD files with GoldWave?

No, but you can extract the instrument samples. After you select the file with **File | Open**, the attribute dialog is displayed. Set the attributes to 8-bit, mono, signed, 16000Hz, no swap. By using the start/finish markers, you can extract the individual instruments. *Warning: do not save the MOD file within GoldWave. GoldWave does not update the MOD header.*