

Contents

GoldWave is a sound editor, player, recorder, and converter. Before using GoldWave, make sure sound drivers have been installed and configured with Window's Control Panel.

To learn how to use help, press F1.

Introduction

[Main Window, Device Controls Window, and Sound Windows](#)

[Direct-to-disk editing, RAM editing, and the Flash feature](#)

[Editing Overview](#)

[Author](#)

Commands

[File Menu Commands](#)

[Edit Menu Commands](#)

[Effects Menu Commands](#)

[View Menu Commands](#)

[Tools Menu Commands](#)

[Options Menu Commands](#)

[Keyboard Commands](#)

How To...

[Convert files from one format or type to another](#)

[Directly edit individual samples](#)

[Play or record a sound](#)

[Select part of a sound](#)

[Set reference/cue points](#)

[Select a device for playback or recording](#)

[Create a tremolo effect](#)

GoldWave Windows

There are three types of windows used by GoldWave: the **Main** window, the **Device Controls** window, and **Sound** windows. The **Main** window contains a menu bar, tool bars (sometimes called speed bars or button bars), and information/status bars. This window keeps all the **Sound** windows grouped together. It is similar to the main window of a word processor, where you can have several document windows open at once.

The **Device Controls** window contains a oscilloscopes and sound device controls. The oscilloscopes display the sound as it is being played. The controls determine the volume, left/right balance, and speed at which sounds are played. A set of standard cassette buttons for playing and recording sounds are kept in this window. Several features can be configured using the setup button.

Sound windows are created whenever sound files are opened. You can open as many sounds you like, provided there is enough memory and disk space. **Sound** windows display a graph of the sound. Mono sounds are shown with a single green graph. Stereo sounds are shown using a green graph for the left channel and a red graph for the right channel.

Just below the graph is an **Overview Bar**. It provides useful information about what part of the sound is currently selected and what part of the sound is displayed by the graph. You can zoom in on a section of a sound by using the View Menu Commands.

See Also

[File Menu Commands](#)

[View Menu Commands](#)

[Device Controls](#)

Editing Overview

Almost all commands in GoldWave operate on the currently selected part of a sound. The selected part, or **selection**, is the highlighted part of the sound graph between two vertical markers. The vertical markers are cyan lines located at the far left side (**start marker**) and far right side (**finish marker**) of the graph.

Note that GoldWave does not use Windows' standard "click-and-drag" method to make a selection because it does not allow accurate positioning of the markers. The method below allows you to independently set the start and finish markers to the nearest sample (when zoomed in at a 1:1 level or better).

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.

See Also

[Direct editing of individual samples](#)

[Edit Menu Commands](#)

[View Menu Commands](#)

[Disk, RAM, and Flash](#)

Disk, RAM, and Flash

GoldWave supports both **direct-to-disk** editing and **RAM** editing, along with a time saving **flash** feature. These features are described below.

In **direct-to-disk** editing, the entire sound is stored in a temporary file where it can be modified. This allows you to edit very large files (up to about 1 billion bytes in size) provided the required disk space is available. Only a small amount of RAM is required for each opened sound. The drawback is that editing and effects processing take more time since audio data must be transferred to and from the disk.

In **RAM** editing, the entire sound is stored in RAM. This allows you to edit and process files very quickly. It saves time and reduces the burden on your hard disk. The drawback is that the size of the files must be small enough to fit in the available RAM (not including virtual RAM).

The **flash** feature allows large files to be opened instantly. The entire file is not copied to temporary storage and only one second of the sound is graphed. This can save a great deal of time if you only want to play a file and not modify it.

A flashed file requires no disk space and only a small amount of RAM, which means that several large files can be opened at once, regardless of how much free space is available on the disk. The drawback is that you need a fast system, especially when a file has to be decompressed before it can be played.

See Also

[Options | File](#)

[Edit | Deflash](#)

Author

Chris Craig is a programmer and engineer. He graduated from Memorial University of Newfoundland in 1994 with a Bachelor of Engineering.

He has been developing software for over 10 years, working in a variety of languages including C, C++, 0x86 assembly. His recent works include ScopeTrax (a Microsoft DOS sound editor) and GoldWave (a Microsoft Windows sound editor).

He currently lives in St. John's, Newfoundland, Canada and can be reached by:

Mail:

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

E-mail:

chris3@cs.mun.ca

World Wide Web:

<http://web.cs.mun.ca/~chris/>

File Menu Commands

New

Creates a new sound file and allows you to choose the sound's attributes through the New Sound dialog. New files are saved as Waves (.WAV) by default.

Open

Opens an existing sound file. If GoldWave does not recognize the file's format, the Raw File Format dialog is presented. The settings in Options | File control how the file is to be stored for editing.

Close

Closes the current sound.

Save

Saves the current sound to the disk.

Save as

Saves the current sound to the disk under a different name and/or format. This command can be used to convert files from one format to another by selecting an appropriate type and attributes. If the current file is flushed, use Edit | Deflash before **Save as**.

Exit

Exits GoldWave.

1..5 File History

At the bottom of the File Menu is a history list of five recently opened files. These files can be re-opened by selecting them from the menu.

Edit Menu Commands

Undo

Reverses the last modification.

Cut

Copies the selection into the clipboard and deletes it from the sound.

Copy

Copies the selection into the clipboard.

Paste

Pastes the clipboard into the sound.

Paste at

Pastes the clipboard into the sound at a specified location.

Mix

Mixes the clipboard with the sound.

Delete

Deletes the selection.

Trim

Deletes the ends of the sound outside the selection.

Insert space (Space Button)

Inserts a blank space into the sound of a length you specify in seconds. The space is inserted at the start marker's position.

Select view (Sel Vw Button)

Moves the start and finish markers to the left and right sides of the sound graph. This essentially selects all the sound currently displayed in the graph.

Channel (Chan Button)

Controls which channel of a stereo sound will be used or modified by editing or effects. You can use this feature to copy a single channel from a stereo sound or apply an effect to only one channel. The currently selected channel is shown in the status bar. The channel setting has no effect when recording, using the [Expression Evaluator](#), or using pitch modifying effects such as [Doppler](#), [Transpose](#), etc.

Deflash

Usually a flashed file will be deflashed automatically. If you are trying to play a compressed file on a slow system, you can use this command to decompress the file into temporary storage. The [flash](#) feature can be configured using the [Options | File](#) command.

See Also

[Direct editing of individual samples](#)

[Editing Overview](#)

Effects Menu Commands

Distortion

Distorts or limits output levels.

Doppler

Dynamically changes pitch.

Echo

Produces an echo.

Exchange channels

Exchanges the left and right channels of a stereo sound.

Filter

Displays a pop-up menu listing several filter related commands.

Flange

Creates a number of flange related audio effects.

Interpolate

Uses linear interpolation to smooth out samples between the start and finish markers. Use this command on a tiny selection to remove a pop or click.

Invert

Turns the sound upside-down.

Mechanize

Adds a mechanical characteristic.

Offset

Changes the dc offset.

Pan

Sets internal balance and left/right channel panning.

Reverse

Reverses the selection so that it will play backwards.

Silence

Silences (mutes) the selection.

Volume

Displays a pop-up menu listing several volume related commands.

Playback rate

Changes the rate at which the entire sound is played.

Resample

Changes the entire sound so that it can be played at a different sampling rate.

Transpose

Changes the pitch of the entire sound by resampling it.

See Also

Editing Overview
Edit Menu Commands

View Menu Commands

View commands act like a zoom lens allowing you to zoom in to and zoom out of the sound graph. By using these commands, you can get a more detailed view of the selection. The **Overview Bar** at the bottom of a Sound window shows what part of the sound is currently graphed.

All (All Button)

Graphs the entire sound.

Selection (Sel Button)

Magnifies the selection, increasing the detail of the graph. You can zoom in many times (by changing the selection and zooming again) until a single sample is shown in the graph.

User (User Button)

Magnifies the sound to the level of detail specified under the Options | Window dialog. The start marker is used as the zoom location. The level can be set to any value you find convenient.

Other

This magnifies the graph to any level you specify. The level is given as 1:X, where X is the number you enter at the prompt. A value of 10 gives a 1:10 level as described below. A value of 0.10 is equivalent to a 10:1 zoom level. If the given level is not possible, the closest valid level is used.

Zoom 10:1

Magnifies the sound to give a very high level of detail. At this level, individual samples are easily visible. The start marker is used as the zoom location. Direct sample editing with the mouse is possible at this level.

Zoom 5:1

Magnifies the sound to give a high level of detail. At this level, individual samples are visible. The start marker is used as the zoom location. Direct sample editing with the mouse is possible at this level.

Zoom 1:1 (1:1 Button)

Magnifies the sound to give a true representation of the graph of the sound. The start marker is used as the zoom location.

Zoom 1:10

Shows a larger part of the sound, but with less detail. The start marker is used as the zoom location.

Zoom 1:100

Shows a much larger part of the sound, but with less detail. The start marker is used as the zoom location.

Start

Scrolls the view to the start marker.

Finish

Scrolls the view to the finish marker.

See Also

[Direct editing of individual samples](#)

Tools Menu Commands

Cue points

Presents a dialog where cue points can be created or edited.

Expression evaluator

Presents a dialog where sound can be generated from mathematical expressions.

CD player

Shows or hides the CD Player window.

Device controls

Shows or hides the Device Controls window.

View Menu Commands

Colours

Presents a dialog where the Sound window colour scheme can be changed.

File

Presents a dialog where file related features can be configured.

Window

Presents a dialog where window related features can be configured.

Keyboard Commands

Sound Windows

Keystroke	Action
Left	Scrolls the Sound window view left.
Right	Scrolls the Sound window view right.
Pg Up	Scrolls the Sound window view left one screen.
Pg Dn	Scrolls the Sound window view right one screen.
Home	Moves the Sound window view to the start marker's position.
End	Moves the Sound window view to the finish marker's position.
Ctrl+Home	Moves the Sound window view to the beginning of the sound.
Ctrl+End	Moves the Sound window view to the end of the sound.

Main Window

Keystroke	Action
F1	Starts on-line help.
Alt+F6	Switch between Main window and Device Controls window.
Ctrl+F6	Switch between Sound windows.
Space	Plays a sound.
F4, F5, F6, F7, F8	Plays, fast forwards, pauses, rewinds, and stops respectively.
Pause	Pauses a sound.
Esc	Stops a sound.

Editing

Keystroke	Action
Ctrl+X or Shift+Del	Cut selection.
Ctrl+C or Ctrl+Ins	Copy selection into the clipboard.
Ctrl+V or Shift+Ins	Paste clipboard into the sound at the start marker.
Ctrl+B	Paste clipboard into the sound at the beginning.
Ctrl+F	Paste clipboard into the sound at the finish marker.
Ctrl+A	Paste clipboard into the sound at the end (appears at the end).
Del	Delete the selection, permanently.
Ctrl+M or Shift+Ctrl+Ins	Mix clipboard with the sound at the start marker.
Ctrl+T	Trim sound. Removes all data outside the selection.

Save As Dialog

This dialog lets you save a sound under a different name and format. You can convert a file from stereo to mono, 8 bits to 16 bits, .WAV to .AU, and a wide variety of other formats.

To rename a sound:

- 1) Choose **Save as** from the **File** menu.
- 2) Type in a new name in the **File Name** box.
- 3) Choose **OK**.

To convert the sound to a different format:

- 1) Choose **Save as** from the **File** menu.
- 2) Type in a new name in the **File Name** box.
- 3) Choose the new file type from **Save File as Type** list.
- 4) Choose an appropriate set of attributes from the **File Attributes** list.
- 5) Choose **OK**.
- 6) If you want the Sound window to have the new format, choose **Yes**.
- 7) If you want the Sound Window to retain its current format, choose **No**.

If you choose **Yes** to update the Sound window, the sound is reopened from the new file on disk. Depending on what options you have selected under [Options | File](#), the file may be flashed. If no filename extension is given, an appropriate one will be appended.

Note that different file types support different attributes. Always pick the type first to see what attributes are available.

This command is useful when files must be used on platforms other than Microsoft Windows, such as audio files for the World Wide Web.

To convert a .WAV file to a standard .AU file, you should first resample the sound to 8000Hz. You can then choose the **.AU** type and the **µlaw mono** attributes for **Save as**.

To save an .AU file as a .WAV file, you should choose the **.WAV** type and the **16 bit mono signed** attributes.

Direct Editing of Samples

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 Menu Commands](#)).

Direct editing:

- 1) Zoom in 2:1 or closer.
- 2) Place the mouse arrow directly over the samples 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, down, left, or right to change the samples and reshape the sound.
- 5) Release the mouse button.

See Also

[Editing Overview](#)

[Edit Menu Commands](#)

Playing and Recording Sounds

The **Device Controls** window provides easy access to audio devices. Before playing or recording a sound, you should use the  (setup) button to make sure that input and output audio devices have been selected. Many device related features can be configured through the [Device Setup](#) dialog.

To play part of a sound:

- 1) Move the start and finish markers to the part of the sound you want play.
- 2) Choose the  button on the **Device Controls** window.

The  (user play) button can be configured to loop or play different parts of the sound.

You can pause playback with the  button and stop playback with the  button. To save screen space, you can minimize the **Device Controls** window. The play and stop buttons will still be usable.

The device **volume** and **balance** can be set using the  and  scroll bars. The **playback speed** of the device can be controlled with the  scroll bar.

The  (rewind) button and the  (fast forward) button move quickly through the file. The default speed for these buttons can be configured in the [Device Setup](#) dialog.

To record a sound:

- 1) Move the start and finish markers to the part of the sound you want to record over (you can use Edit | Insert space to create a place to record the sound).
- 2) Choose the  button on the **Device Controls** window.

Before recording, make sure that Line In or Mic In is turned on in the Mixer application included with your sound card.

GoldWave requires that the recording time be known in advance so that sufficient resources can be pre-allocated. This results in higher quality recordings with less pops and clicks because resources do not have to be allocated on-the-fly.

Use the  button to stop recording. This button appears only during recording.

See Also

[Editing Overview](#)

[Device Controls Window](#)

Cue Points

Cue points allow you to remember and describe specific positions within sounds. Cue points can be set at the start or finish marker's position. Also, the start or finish marker can be moved to a cue point. Cue points are saved only in Wave files.

To set a cue point (using the start marker):

- 1) Move the start marker to the position where you want to create a cue point.
- 2) Choose **Cue points** from the **Tools** menu.
- 3) Choose the **Start marker** radio button.
- 4) Enter a name and description for the cue point.
- 5) Choose the **Add** button.
- 6) Choose **OK**.

To delete a cue point:

- 1) Choose **Cue points** from the **Tools** menu.
- 2) Select the cue point from the list.
- 3) Choose the **Delete** button.
- 4) Choose **OK**.

To change a cue point:

- 1) Move the start marker to the position where you want to move the cue point.
- 2) Choose **Cue points** from the **Tools** menu.
- 3) Select the cue point from the list.
- 4) Choose the **Start marker** radio button.
- 5) Change the name and description.
- 6) Choose the **Revise** button.
- 7) Choose **OK**.

If you do not want to change the cue point's position, you should move the start marker to the cue point first.

To move the start marker to a cue point:

- 1) Choose **Cue points** from the **Tools** menu.
- 2) Select the cue point from the list.
- 3) Choose the **Start marker** radio button.
- 4) Choose the **Set** button.
- 5) Choose **OK**.

Note:

Currently, cue points are fixed and do not change position when a sound is modified. This should be considered when certain commands, such as delete, are used. Any cues inside the deleted selection will not be deleted and the cues outside the selection will not be adjusted to account for the new positions.

Expression Evaluator

The Expression Evaluator is a versatile tool for manipulating and generating audio data. Expressions can be entered by clicking on the buttons, using the keyboard, or selecting an expression from the Group and Expression name lists.

For example, say you want to generate a sine wave. The mathematical equation is:

$$y = \sin(2\pi ft) \quad \text{where } f \text{ is the frequency, } t \text{ is the time and } \pi = 3.14159\dots$$

To generate this function in the current sound:

- 1) Choose **Expression evaluator** from the **Tools** menu
- 2) Type in the expression in the large Expression edit box:
`sin(2*pi*f*t)`
- 3) Enter the desired frequency in the **f** edit box (500Hz, for example).
`500`
- 4) Choose the **Start** button.

Evaluation can be stopped at any time with the **Cancel** button. Choosing **Cancel** a second time closes the Expression Evaluator.

You can also use the expression evaluator to implement new effects.

To create a tremolo effect:

- 1) Choose **Expression evaluator** from the **Tools** menu
- 2) Type in the following expression:
`wave(n) * (0.6 + 0.4 * sin(2 * pi * f * t))`
- 3) Enter the tremolo rate in the "f" edit box (4 per second, for example).
`4`
- 4) Choose the **Start** button.

You can alter the depth of tremolo by changing the values `0.6` and `0.4` (to `0.8` and `0.2`, for example).

Group and Expression Names

The **Group name** and **Expression name** lists allow you to organize and store expressions in the **express.eqx** file located in your Windows directory. Similar expressions can be stored together in groups. **Group name** lists all these groups, while the **Expression name** lists all the expressions in a group.

To retrieve an expression:

- 1) Select the group from the **Group name** list.
- 2) Select the expression from the **Expression name** list.

To add an expression:

- 1) Enter the expression.
- 2) Type in the group name or select a group from the **Group name** list.
- 3) Type in an expression name.
- 4) Choose the **Add** button.

To delete an expression:

- 1) Select the group from the **Group name** list.
- 2) Select the expression from the **Expression name** list.
- 3) Choose the **Delete** button.

When a group becomes empty, it will automatically be deleted.

New Sound

GoldWave supports a wide range of sound attributes. These attributes must be specified when creating a new file. By default, the Wave file type is used when the file is saved.

Bits

The number of bits determine how precisely the audio data is stored. Use 16 bits for high quality and low noise processing and recording. Use 8 bits if your system does not support 16 bit sounds or you have limited RAM or disk space. 16 bit sounds use twice as much storage as 8 bit sounds.

Channels

Mono sounds have one channel of audio. Stereo sounds have two channels. Stereo sounds use twice as much storage as mono sounds and can take twice as long to process. You can edit individual channels of a stereo sound by using the Edit | Channel command.

Sampling Rate

This value affects the quality and quantity of the sound. It is the number of times per second data is sent to or received from the audio device. The drop down list box provides many standard rates. Microsoft recommends that this value should be one of 11025, 22050, or 44100.

11025	Voice and low pitch sounds, relatively small storage usage.
22050	Medium quality music, medium storage usage.
44100	High quality music, potentially CD quality, high storage usage.

Time

This value specifies the length of the sound in seconds.

CD quality sound is produced with 16 bit, stereo, 44100Hz. Note that this uses 176,400 bytes for every second of sound or about 10 million bytes per minute.

You can use 8 bit, mono, 11025Hz for vocal sounds. This give good quality and requires only 11025 bytes per second.

If you plan to do some recording, make sure you specify attributes that are compatible with your audio device.

Raw File Format

The Raw File Format dialog lets you open a file that GoldWave could not open automatically. If the file is not compressed or corrupted, GoldWave should still be able to interpret the audio data, provided the correct attributes are specified with this dialog.

The Raw File Format dialog is presented if any of the following occur:

- 1) The file has no header information.
- 2) The file does have a header, but the structure is incorrect for the detected file type.
- 3) The file has a new format that GoldWave does not recognize.
- 4) The file type or format is not supported by GoldWave.

Format

Specifies how each sample of the audio data should be interpreted. For uncompressed files **8 bit** or **16 bit** should work. Files from the World Wide Web are usually **μ-Law**.

Channels

Specifies that the file is either mono or stereo. Stereo data must be interleaved (i.e. left channel, right channel, left channel, ...). If you don't know, choose **Mono**.

Sampling rate

Specifies the rate at which the sound was recorded. This value does not affect how the data is translated. If you select the wrong rate, the sound will either play too slow or too fast. You can fix this with the [Effects | Playback rate](#) command.

Signed

Specifies that 8 bit, 12 bit, or 16 bit samples are signed. Amiga and Apple systems use signed 8 bit (-128 to 127) or signed 16 bit (-32768 to 32767). Wave and Sound Blaster files for PCs are usually unsigned 8 bit (0 to 255) or signed 16 bit (-32768 to 32767). Generally, all 12 bit and 16 bit samples are signed.

Byte swap

When more than one byte is required for each sample, the order in which the bytes are stored can vary from system to system. Systems with Intel processors (0x86 PCs) store bytes in a certain order (1, 2, 3, 4, 5, 6...). Systems with Motorola processors (Macs) store the data in a different order (2, 1, 4, 3, 6, 5, ...) where each pair of numbers has been swapped. If the file came from a non-Intel system, you should check this box.

If you have no idea how the data should be interpreted, experiment with trial and error. If the attributes are wrong, the sound will play distorted. You will have to close the sound window and re-open it using different attributes.

If you save the file later, use a different file name and type using [File | Save as](#).

Undo

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 sound in a temporary file. This file is created in the temporary directory specified in [Options | File](#).

Since the undo feature has to save large amounts of data, you will notice a slight delay before each edit or effect command is performed. The undo feature can be disabled in [Options | File](#).

Cut

Removes the selection from the sound and puts it in the clipboard. The contents of the clipboard can then be Pasted or Mixed into another sound. Note that if only the left or right channel is selected 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 cut channel is padded with silence.

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.

See Also

Delete

Trim

Copy

Copy

Copies the selection into the clipboard. The selection is not removed. The contents of the clipboard can then be Pasted or Mixed into another sound.

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.

See Also

Cut

Paste

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

See Also

[Mix](#)

[Copy](#)

[Paste at](#)

Paste At

Puts the sound in the clipboard into the current sound. The clipboard sound is inserted at the location selected from the pop-up menu. The length of the current sound is increased so that the clipboard sound will fit. The clipboard sound is automatically converted to match the attributes of the current sound.

To append the clipboard to the current sound:

- 1) Choose **End** from the **Edit | Paste at** pop-up menu.

See Also

[Mix](#)

[Copy](#)

[Paste](#)

Mix

This is a special editing feature for sounds. The sound in the clipboard is blended (superimposed or layered) with the current sound. You are asked for the mix volume. This volume is applied to the clipboard sound as it is being mixed.

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 = full volume).
- 4) Choose **OK**.

See Also

[Copy](#)

[Paste](#)

[Volume](#)

Delete

Permanently removes the selection from a sound. The selection is not copied to the clipboard. Note that if only the left or right channel is selected 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.

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.

See Also

[Cut](#)

[Trim](#)

Trim

Trim permanently removes the ends of the sound that are outside the selection. The ends are not copied to the clipboard. Note that if only the left or right channel is selected 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 ends of the trimmed channel are padded with silence.

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.

See Also

[Cut](#)

[Delete](#)

Distortion

Distortion alters the amplitude mapping of the selection. It can limit, compress, or expand a range of amplitudes. The amplitude mapping is set using Shape Controls, where x-axis and y-axis both have a range of -1 to 1. When the line stretches diagonally from the lower left corner to the upper right corner, the input amplitude (x) and output amplitude (y) are the same for every point on the line. By changing the line, the output will be different from the input.

Doppler

Doppler dynamically alters or bends the pitch of the selection. Shape Controls are presented where the pitch can be varied over the length of the selection from half to 1.5 times normal. A doppler effect can be heard at race tracks when a fast car passes in front of you. The pitch of the engine appears to drop slowly as the car speeds away. You can use Effects | Volume | Shape to dynamically alter the volume as well.

The Doppler effect alters both channels of a stereo sound, even if only one channel is selected.

Echo

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

Delay

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.

Volume

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

Reverb

The Reverb check box enables or disables regeneration of the echo. This makes the echo deeper.

To make the echo sound correct, the effect extends slightly outside the end of the selection. This may increase the length of the sound.

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 **Reverb** if desired.
- 6) Choose **OK**.

Filter Menu

Presents a pop-up menu containing filter related commands. Filters are used to remove a range of frequencies from a sound and can produce a variety of effects.

Noise gate

Removes noise from quiet sections of the selection.

User defined

Applies a filter to the selection using coefficients you specify.

User Defined Filter

The **User Defined Filter** dialog allows you to specify coefficients to use for filtering. Up to 15 coefficients can be given. Almost any kind of linear filter can be created with this command because it uses a general digital filter equation:

$$b(0)y(n) + b(1)y(n-1) + \dots + b(14)y(n-14) = a(0)x(n) + a(1)x(n-1) + \dots + a(14)x(n-14)$$

For speed, fixed point calculations are used. Coefficients can be floating point values with magnitudes from 0.0001 to 500. The number of coefficients entered for **a** and **b** must be the same. For FIR filters, you would usually enter a one followed by a number of zeros for **b**.

You can use Ctrl+C and Ctrl+V to copy and paste coefficients. The **Clear** button clears the coefficient edit boxes. GoldWave includes several useful coefficient presets.

Coefficient Sets

The number following lowpass filter presets indicates what percentage of frequencies are kept.

Lowpass 25, for example, keeps the lower 25% of frequencies. The number following highpass filter presets indicates the percentage discarded. **Highpass 10**, for example, discards the lower 10% of frequencies. The actual frequencies kept or discarded depends on the sampling rate of the sound.

Lowpass 25 on a 22050 Hz sound will remove frequencies from about 2700 Hz to 11025 Hz.

To fully use this command requires a working knowledge of digital filter theory. Several of the filters in the presets were generated by Matlab (info@mathworks.com).

Noise Gate

Noise gate removes background hiss from quiet parts of the selection. You can use this after recording to clean up some of the noise created by the audio device when it converted the sound to digital data. Noise gate does not remove background hiss from louder parts of the selection.

Decay time

This is the amount of time (in milliseconds) that it takes for the noise gate to fully close. When the gate is closed, no sound can pass and this leaves only silence. Values less than 200 milliseconds are common.

Threshold

This is the amplitude level at which the gate will start to open and let sound pass. If you specify a value of 0.05, for example, all samples with levels from 0.05 to 1.0 will be allowed to pass. Samples with levels 0 to 0.05 are blocked. If you still notice a hiss in quiet sections, increase this value.

Flange

Flange uses variable delays and mixing to create some unusual audio effects. Parameters can be entered or read from [presets](#).

Input volume

Specifies the volume of the unmodified sound to send to the output.

Mix volume

Specifies the volume of the delayed sound to mix with the output.

Feedback

Specifies the volume of the feedback to mix with the output.

Invert

The input, mix, or feedback can be inverted by checking the appropriate check box. Normally, the mixing adds sounds together. Checking **Invert** subtracts the sound instead.

Depth

Specifies in milliseconds how much the delay should vary. A value of 40 will let the delay vary from 0 to 40 milliseconds.

Frequency

Specifies how fast to vary the delay. A value of 2 will vary the delay over its depth twice a second. For a value of 0.2, the full delay depth is reached every 5 seconds.

Fixed delay

The fixed delay is added to the depth to change the minimum delay. If the depth is 40 and the fixed delay is 10, the delay will vary from 10 to 50 milliseconds.

Invert

Reflects the selection about the x (time) axis. This produces no noticeable effect in mono sounds and has a slight effect in stereo sounds. Inverting a single channel of a stereo sound will produce a surround-sound effect.

To invert part of a sound:

- 1) Move the start and finish markers to the part of the sound you want to invert.
- 2) Choose **Invert** from the **Effects** menu.

Mechanize

Adds a robotic or mechanical characteristic to the selection. The percentage of quality can be entered after selecting this command. Low values produce an untuned two-way 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**.

Offset

Adjusts or removes a dc offset in the selection by shifting it up or down. If you notice that silent sections of a sound are not at zero in the graph, you can use this command to adjust them to zero.

When this command is selected, it first scans the sound for any existing offset. The existing offset is then displayed in a dialog where it may be changed. A positive value shifts the graph up and a negative value shifts it down.

To adjust the offset of part of a sound:

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

Pan

You are presented with Shape Controls. For panning, the shape graph 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 dynamically alter the sound's left/right balance or pan to and from each channel.

Volume Menu

Presents a pop-up menu containing volume related commands. Volumes are usually specified by a percentage of the sound's original amplitude.

The output volume of the audio device can be controlled in the Device Controls window (if your audio device supports it). Volume commands actually change the internal volume of a sound.

Change

Changes the volume of the selection.

Fade in

Gradually increases the volume throughout the selection.

Fade out

Gradually decreases the volume throughout the selection.

Maximize

Increases volume as high as possible without distortion (i.e. normalizes the volume)

Shape

Allows envelope shaping.

See Also

Pan

Change Volume

Modifies the selection so that it sounds louder or quieter. The percentage of volume can be entered after selecting this command. Values less than 100 make the sound quieter and values greater than 100 make it louder. A value of 100 is normal volume and has no effect. Enter 50 to halve the volume and 200 to double the volume.

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 the volume.
- 2) Choose **Change** from the **Effects | Volume** pop-up menu.
- 3) Enter the new volume percentage, then choose **OK**.

Fade In

Gradually increases the volume throughout the selection. 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**.

See Also

[Fade out](#)

Fade Out

Gradually decreases the volume throughout the selection. 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 half 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.
- 2) Choose **Fade out** from the **Effects | Volume** pop-up menu.
- 3) Enter the fade percentage, then choose **OK**.

See Also

[Fade in](#)

Maximize

Searches the selection for the maximum percent that the volume can be increased without producing distortion or clipping errors. This volume is then applied to the selection. If the percentage is above 5000% or equal to 100%, no changes will be made.

To maximize the volume of part of a sound:

- 1) Move the start and finish markers to the part of the sound you want to increase the volume.
- 2) Choose **Maximize** from the **Effects | Volume** pop-up menu.

Shape Volume

You are presented with Shape Controls. With these controls, you can define the amplitude envelope of a sound. The shape line is initially horizontal at 100. By bending/moving the line, you can dynamically change the volume of the selection.

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 information bar. Values of 11025, 22050, and 44100 are recommended.

To change the playback rate of the entire sound:

- 1) Choose **Playback rate** from the **Effects** menu.
- 2) Enter the new rate, then choose **OK**.

The playback rate of the audio device can be controlled in the [Device Controls](#) window.

See Also

[Resample](#)

[Transpose](#)

Resample

Changes the sampling rate of the entire sound. Unlike **Playback rate**, this command re-calculates and 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 recommended.

To change the sampling rate of the entire sound:

- 1) Choose **Resample** from the **Effects** menu.
- 2) Enter the new rate, then choose **OK**.

See Also

[Playback rate](#)

[Transpose](#)

Transpose

Changes the sampling rate of the entire sound so that it plays back at a different pitch. This is useful for converting instrument samples from one note to another. Transpose combines the **Resample** and **Playback rate** effects. You are prompted to select the **original pitch** and the **new pitch** from drop down lists. These lists specify notes.

By specifying "C" as the **original pitch** and "C2" as the **new pitch**, you can shift the pitch up one octave.

Fine tune lets you make a slight pitch adjustment.

See Also

[Playback rate](#)

[Resample](#)

CD Player

The **CD Player** provides easy access to your CD-ROM device for playing audio CDs. Before using the **CD Player**, you should carefully follow the CD-ROM installation instructions and make sure that all the necessary CD-ROM drivers have been installed (in DOS and Windows).

Some CDs contain both computer data and audio. Attempting to play a track that does not contain audio will cause an error.

A bug in the Microsoft Windows 3.1 MCICDA driver can cause problems on some system when a CD contains more than 15 tracks.

The buttons from left to right are play, pause, stop, track up, track down, and eject. The track can be changed by using the track up/down buttons or by using the scroll bar.

The status gauge shows how much of the current track has been played. You can jump to any point in the track by clicking inside the gauge.

Like the Device Controls window, the CD Player has working buttons when minimized.

Colour Options

Use the **Sound Window Colours** dialog to set the colours for Sound windows. The **Preview** box shows the Sound's current colour scheme. The **Colour Settings** group allow you to change the scheme. To set the colour of a particular item, select the item from the drop-down list or click on the item in the **Preview** box. Once the item is selected, you can change its colour.

File Options

Use the **File Options** dialog to setup directories and file storage options.

Sound files

Specifies the directory where you keep your sound files. The **File | Open** command will automatically list files in this directory whenever you start GoldWave. If you prefer to use Program Manager's **Properties** feature to specify a working directory, enter a period, ".", for this directory.

Temporary

Specifies the directory to use when creating temporary files. This directory should be located on a large disk with plenty of free space. Using a compressed drive is not recommended. It will slow processing and give poor results when recording if **hard disk** storage is enabled. Changing this directory does not affect opened files already in temporary storage.

Undo

Specifies the directory to use for storing undo data. In most cases, this should be the same as the temporary directory. Changing the undo directory does not affect the current session of GoldWave since the undo file has been created already. **Undo** can be enable or disabled by checking or unchecking the check box.

Flash open

These radio buttons let you control the flash feature. If you usually edit small files or have a slow system, set this to **Never**. If you always play files and rarely modify them, choose **Always**. If you have a fast system and often modify files, choose **Limit** and specify the minimum size (in units of 1000 samples) for a file to be flash. Any file larger than this will be flashed.

Temporary storage

Specifies where files should be stored for processing. RAM storage is very fast, but limits the size of files. Hard disk is slower, but allow huge files to be processed. Changing this options does not affect files currently opened.

Window Options

Use the **Window Options** dialog to configure the positions of windows and specify the zoom value for View | User.

Main window size

Controls the Main window's position and size every time GoldWave is started. **Normal** gives control to Windows. **Maximize** makes the Main window occupy the entire screen. **Save position** saves the Main window's position and size when GoldWave is closed so that it will appear in the same location next time.

Main window size

Controls the position and size of Sound windows. **Normal** gives control to Windows. **Maximize** makes a Sound window occupy the entire Main window. **Auto-tile** resizes all Sound windows whenever a new sound is opened or closed so that every one will be visible.

User zoom

This is where you specify the level of zoom for View | User and **User** button. Values between 0.01 to 1000 are valid.

Device Setup

The **Device Setup** dialog lets you select input and output audio devices and configure device controls and realtime oscilloscopes.

To select a new output device, choose the **Output** button. To select a new input device, choose the **Input** button.

Graph type

This group allows you to set the type of oscilloscope to display. Clicking on the oscilloscope changes the type as well. Types include:

Amplitude	The sound is shown as a series of amplitudes.
Spectrum	The sound is shown as a frequency spectrum.
Spectrum bar	The sound is shown as a frequency bar graph commonly found on stereo systems.
Spectrogram	The sound is shown as a coloured frequency spectrum, with frequency on the y-axis and time on the x-axis. The colours represent the magnitude of the frequency. The colours, in increasing magnitude, are black, purple, blue, cyan, green, yellow, red, and white.

The frequency range of the three spectrum graph types depends on the sampling rate and the current device speed. For a speed of 1.00, the frequency range is 0 to half the sampling rate. A sound sampled at 22050 Hz would have frequencies from 0 to 11025 Hz shown in the graph.

Controls

This group lets you configure the speed of the fast forward and rewind buttons. By entering a value of 3.00 for **Rewind speed**, for example, the rewind button will play the sound backwards three times faster than regular playback speed. By entering small numbers (such as 0.1) the rewind and fast forward buttons will play very slowly. This is useful for finding pops or clicks, since the oscilloscope will move slowly through the data.

Save

If **Setup** is checked, all the values in the **Device Setup** dialog will be saved. If **Window position** is checked, the position and size of the **Device Controls** window will be saved.

User play button

Configures the **user play** button to play and/or loop sections of the sound. If **All** is chosen, the entire sound will be played. If **Selection** is chosen, the sound between the start and finish markers will be played. If **Unselected** is chosen, the regions outside the start and finish markers will be played. This lets you quickly test how a cut or delete will sound without actually changing the sound. **Intro/loop/end** is a special playback feature that plays the sound in three section. The beginning of the sound, outside the selection, is played first. Then the selection is play and looped. Finally the end of the sound, outside the selection, is played. This is useful for musical accompaniment or looped instrument samples. **Loop** specifies the number of times playback should be repeated. A zero value loops forever.

Positioning

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

Options

Monitor input connects the input to the oscilloscopes and LED meters so you can adjust input levels before recording. The **Triple buffering** check box enables triple buffering (checked) or double buffering (unchecked) for playback and recording. Unless you experience problems during recording or playback, you should keep triple buffering enabled. The **Input buffer** box controls the size of the buffer (in

seconds) to use for recording. Three of these buffers are created for triple buffering. The **Output buffer** box controls the size of the buffer (in seconds) to use for playback. Three of these buffers are created for triple buffering.

Increasing buffer sizes may reduce the number of pops and clicks, but will require more memory. A two second playback buffer for a CD quality sound requires a total of 1MB of memory. Larger buffers also take longer to fill when using rewind and fast forward.

Device Select

The drop down list shows you the names of audio device drivers. Below the list is a table that tells you what sound attributes 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. To play or record CD quality, the driver must support 16 bit, stereo, 44.100 kHz.

Presets

Presets store parameters and shapes in the **gwpreset.ini** file so they can be recalled again the next time the effect is used. Controls for presets consist of a drop down list box, a **[+]** button, and a **[-]** button.

To add a new preset:

- 1) Enter in all the new parameters and/or draw the new shape.
- 2) Type in a new name for the preset in the drop down list. This name cannot be the same as one currently in the list.
- 3) Choose the **[+]** button.

To delete a preset:

- 1) Select the preset from the drop down list.
- 2) Choose the **[-]** button.

To change a preset:

- 1) Delete the preset, as above.
- 2) Add the preset, as above.

When you delete a preset, the current parameters and name remain on the screen so they can be changed.

Shape Controls

Shape controls usually consists of a graph window, a **Clear** button and shape presets.

Graph window

The graph window initially contains a single line with two endpoints (shown as large dots). By clicking the left mouse button anywhere inside the graph window, you can **add** new points to bend the line into a variety of zigzag shapes.

To **move** a point, click on it and drag it to a new location.

To **remove** a point, click the **right** mouse button over the point. Endpoints cannot be removed.

To remove all the points (except the endpoints), choose the **Clear** button.

GoldWave Order Form

Please type in the following information:

Name (first & last): _____

[Company]: _____

Address: _____

Country: _____

[Telephone (with Area Code)]: _____

Version of GoldWave you are using: _____

Registration:

Standard (\$30 US)

 Password to remove shareware message and unlock
 all future versions.

Deluxe (\$50 US)

 Password (as above). Diskette and manual for
 new version, when available.
 Please include \$5.00 US extra for deluxe orders
 outside USA and Canada.

Please mail this form along with your check or money
order. Be sure to make the check or money order payable
to Chris Craig. Thank you for registering!

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

Registering

GoldWave is a shareware program. To register and encourage further development, please follow the directions in the ORDER.WRI file included with this program. If you require more than one registration, please send an order form each person who will be using GoldWave.

Note: Passwords cannot be sent by e-mail.

Standard Registration (\$30 US)

You will receive a personalized password that will remove the startup shareware message and unlock all versions of GoldWave. The password will make all future versions of GoldWave fully functional. This is equivalent to a lifetime of free updates! All you have to do is download the latest version from the GoldWave home page:

<http://web.cs.mun.ca/~chris3/goldwave/>

Deluxe Registration (\$50 US)

You will receive a personalized password as above. You will also receive a diskette and a spiral bound manual for the next version of GoldWave (when completed). The diskette will contain a new version of GoldWave, some wave files, and other software.

GoldWave Shareware Version

This application is a fully functional shareware version of the GoldWave digital audio editor. Shareware allows you to try a program before you buy it. If you find GoldWave useful, you are required to [register](#) it.

As shareware, you can give copies of GoldWave to anyone you think might find it useful. You can also upload it to BBS or ftp sites and post it to appropriate forums or newsgroups. Before copying, uploading, or posting, please click [here](#).

Please check the [Contents](#) for an introduction to direct-to-disk/RAM/flash and other changes.

The **readme.txt** file contains helpful information and important installation instructions. The **whatsnew.txt** file lists many of the changes made to this version. For recent information and software updates, please refer to the GoldWave WWW Home Page:

<http://web.cs.mun.ca/~chris3/goldwave/>

Toll Limit

The **toll** in the status bar gives you a rough idea of the benefit you have received from using GoldWave. This shareware version is limited to 150 commands (or a \$1.50 toll) per session. When the limit is reached, the Device Controls window becomes disabled, but all other commands will continue to work. You can enable the Device Controls window again by exiting and restarting GoldWave. By registering, you will receive a password that removes the toll limit.

Only the original GoldWave zip file may be copied, uploaded, or posted. This ensures that everyone will get a complete and working copy. Distributing modified or incomplete copies is a violation of the copyright.

Bugs

Before reporting bugs, please check the GoldWave home page to see if the bug has been reported already:

<http://web.cs.mun.ca/~chris3/goldwave/>

If you encountered an error that caused GoldWave to terminate, you should delete any temporary files that were created (**gu?????. \$\$\$** and **gw?????. \$\$\$**).

Try to remember the command(s) that led to the error. The best report to send would be a sequence of commands that is guaranteed to recreate the error.

Unless you have an unusual system configuration, you do not need to send information about your system.

Constructive criticism is welcome.

Please send bug reports and comments to:
chris3@cs.mun.ca

WARNING:

This is a pre-release (beta) version of GoldWave. This means that it is still in the testing stage and may contain bugs that can cause a loss of data. Use at your own risk.

This beta version of GoldWave is provided through the GoldWave WWW home page only and may **not** be distributed in any other form. Under no circumstances should it be uploaded to BBS or FTP sites, posted to newsgroups or forums, or placed on CD-ROMs.

A full shareware release will be available shortly that can be redistributed.

This version has undergone extensive changes, so there are bound to be some rough edges. Click [here](#) if you have discovered a bug, an error in the help/text files, or would like to send some feedback.
