

Contents

GoldWave is a sound editor, player, recorder, and converter. It can create entertaining sound files for Web pages, answering machines, or Windows sounds. A rich set of effects and editing features are included for professional sound production. High quality audio CDs can be created by using GoldWave in conjunction with CD writer software.

Before using GoldWave, make sure sound drivers have been installed and configured with the Windows Device Manager.

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 a sound](#)

[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. In GoldWave, you can have several sounds open at once to copy and paste between them.

The **Device Controls** window contains 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 as 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 the standard "click-and-drag" method to make a selection because it does not allow accurate positioning of the markers. The method, described 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](#)

[Cue points](#)

[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 the first five seconds of the sound is graphed initially. 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. Note that some compressed files cannot be flash opened due to limitations of the compression algorithm.

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 degree.

He has been developing software for over 11 years, working in a variety of languages including C, C++, 0x86 assembly. His recent works include Multiquence (a real time Windows multitrack audio editor), GoldWave (a Windows digital audio editor), and ScopeTrax (a simple MS-DOS 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/~chris3/>
<http://www.goldwave.com>

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 File Format dialog is presented. For unrecognized formats, you can associate the file name extension to a format using Options | File types. 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 appropriate type and attributes.

Exit

Exits GoldWave.

1..7 File History

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

Edit Menu Commands

Undo

Reverses the last change, edit, or effect.

Cut

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

Copy

Copies the selection into the clipboard.

Copy to

Copies the selection directly to a new file.

Paste

Pastes the clipboard into the sound.

Paste new

Pastes the clipboard to a new Sound window.

Paste at

Pastes the clipboard into the sound at a specified location.

Mix

Mixes the clipboard with the sound.

Replace

Replaces the selection with the sound in the clipboard.

Delete

Deletes the selection.

Trim

Deletes the ends of the sound outside the selection (trims off the ends).

Insert silence (Space Button)

Inserts silence into the sound of a length you specify in seconds. The silence 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 rate modifying effects such as Resample and Playback rate.

Marker

Shows a submenu of marker related commands. These commands let you set markers to an exact time or sample position and move the markers to the current playback position.

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

Doppler

Dynamically changes pitch.

Dynamics

Distorts, compresses or limits output levels.

Echo

Produces an echo.

Filter

Displays a submenu listing several filter related commands.

Flange

Creates a number of flanging 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.

Pitch

Changes or transposes the pitch (frequency).

Reverse

Reverses the selection so that it plays backwards.

Silence

Silences (mutes) the selection.

Stereo

Displays a submenu showing stereo related effects.

Time warp

Changes the speed and stretches or compresses the length of the selection.

Volume

Displays a submenu 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.

See Also

Editing Overview
Edit Menu Commands

View Menu Commands

View commands act like a zoom lens allowing you to zoom in and/or 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.

Other

This magnifies the graph to any level you specify.

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.

Zoom in

Magnifies the sound to give twice as much detail. The middle of the window is used as the zoom location.

Zoom out

Reduces magnification to show half as much detail. This shows more of the sound. The middle of the window is used as the zoom location.

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 and Zoom 1:1000

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

Vertical zoom all

Vertically zooms all the way out so that the entire vertical/amplitude range of the sound is shown.

Vertical zoom in

Magnifies the graph vertically to show twice as much amplitude detail. Zooming is centered on the middle of the graph.

Vertical zoom out

Reduces vertical magnification to show half as much amplitude detail. This show a larger range of the amplitude. Zooming is centered on the middle of the graph.

Start

Scrolls the graph to the start marker.

Finish

Scrolls the graph to the finish marker.

See Also

[Direct editing of individual samples](#)

View User

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 (medium detail) and a value of **0.10** is equivalent to a 10:1 zoom level (high detail). If the given level is not possible, the closest valid level is used.

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 the Windows CD Player accessory. If this option is not enabled, you will have to install the CD Player. See Windows help under "CD Player, Installing" for more information.

Volume control

Shows the Windows Volume Control accessory. By default, output volumes are displayed. To set recording volumes, select **Options | Properties** from the Volume menu, choose the **Recording** radio button, and then choose **OK**. For more information, see Recording Sounds.

Device controls

Shows or hides the Device Controls window.

Audio extraction

Displays a submenu with several tools for extracting audio from different sources. Audio can be copied from an AVI movie file using the AVI audio command. You can make exact digital copies of music on audio CDs using the CD audio command.

Options 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.

File types

Presents a dialog where you can associate a file extension to an audio format.

Tool bar

Presents a dialog where you can add/remove buttons to/from the tool bars.

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.
Shift+Right	Move the start marker right.
Shift+Left	Move the start marker left.
Ctrl+Shift+Right	Move the finish marker right.
Ctrl+Shift+Left	Move the finish marker left.
Shift+Up	Horizontally zooms in.
Shift+Down	Horizontally zooms out.
Ctrl+Up	Vertically zooms in.
Ctrl+Down	Vertically zooms out.

Main Window

Keystroke	Action
F1	Starts on-line help.
Alt+F6	Switches between Main window and Device Controls window.
Ctrl+F6	Switches between Sound windows.
Space	Plays a sound.
F4, F5, F6, F7, F8	Plays, fast forwards, pauses, rewinds, and stops respectively.
Shift+F4	Plays the sound using the User play button settings.
Ctrl+F9, Ctrl+F8	Records, and stops recording 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
Ctrl+B	Paste clipboard into the sound at the begin
Ctrl+F	Paste clipboard into the sound at the finish

Ctrl+A

Ctrl+N

Del

Ctrl+M or Shift+Ctrl+Ins

Ctrl+T

Ctrl+R

[

]

Paste clipboard into the sound at the end (clipboard).

Paste clipboard into a new Sound window.

Delete the selection.

Mix clipboard with the sound at the start of the selection.

Trim sound. Removes all data outside the selection.

Replaces the selection with the sound in the clipboard.

Moves the start marker to the current playhead.

Moves the finish marker to the current playhead.

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 compressed 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 flash opened.

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, choose the **Sun (*.au)** type and the **Java/Web** attribute.

To save an .AU file as a .WAV file, you should choose the **Wave (*.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](#)). If the sound was [flash opened](#), you will need to use the [Edit | Deflash](#) command to prepare the sound for editing.

Direct editing:

- 1) Zoom in 2:1 or closer (such as Zoom 5:1).
- 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](#)

Device Controls



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 playback and recording devices have been selected. Many device related features can be configured through the [Device Controls Setup](#) dialog. To save screen space or enlarge the oscilloscopes, you can resize the **Device Controls** window.

Playback controls




The  (play) button always plays the selection. The  (user play) button can be configured to loop or play different parts of the sound. You can pause playback with the  (pause) button and stop playback with the  (stop) button. For more information, see [Playback](#).

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 Controls Setup](#) dialog. These buttons are enabled only during playback.

Recording controls

The  (record) button starts recording. Use the  (stop record) button to stop recording. This button appears only during recording. For more detailed information, see [Recording](#).

Other controls

The device playback **volume** and **balance** can be set using the  and  scroll bars. The **playback speed** of the device can be controlled with the  scroll bar. Note these controls affect only the device. To change the volume, balance (pan), or speed within a file, use the corresponding command in the [Effects](#) menu. To change recording volumes, see [Recording](#).

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.

See Also

[Editing Overview](#)


[Device Controls Window](#)

[Device Controls Setup](#)


[Tools Menu Commands](#)

Playing Sounds

To play a sound:



- 1) Choose **Open** from the **File** menu.
- 2) Find and select the sound you want to open.
- 3) Choose the **Open** button. The Sound window is now shown.
- 4) Choose the  button on the **Device Controls** window.

To play part of a sound:

- 1) Open the sound, as above.
- 2) Move the start and finish markers to the part of the sound you want to play. See [Editing Overview](#) for details.
- 3) Choose the  button on the **Device Controls** window.

To quickly select the entire sound again, click the **All button** then the **Sel Vw** button on GoldWave's tool bar.

To loop part of a sound:

- 1) Open the sound, as above.
- 2) Move the start and finish markers to the part of the sound you want to loop. See [Editing Overview](#) for details.
- 3) Choose the  (setup) button on the **Device Controls** window, then choose the **Play** tab.
- 4) In the **User play button** group, choose the **Selection** option, check the **Loop** box, set the number of loops, and choose **OK**.
- 5) Choose the  (user play) button on the **Device Controls** window.

The user play button can also be configured to always play the entire sound, the unselected part, or the part currently shown in the Sound window.

See Also


[Device Controls Setup](#)

[Device Controls](#)

[Recording](#)

Recording Sounds

To record a sound:

- 1) Choose **New** from the **File** menu.
- 2) Choose attributes for recording and specify the length of time you want to record, then choose **OK**. For an existing sound, you can move the start and finish markers to the part of the sound you want to record over or use the **Edit | Insert silence** command to create room for recording.
- 3) Select the correct input, as described below, such as line-in, microphone, or CD audio.
- 4) Choose the  button on the **Device Controls** window.

Note that there are several options available for recording under the **Record** tab in the **Device Controls** setup window, including a countdown timer and level activated recording. See the [Device Controls Setup](#) for more information.

To select a different input source, such as the line-in, microphone, or CD audio:

- 1) Select **Volume control** from the **Tools** menu to start the **Volume Control** accessory.
- 2) Select **Properties** from the Volume Control **Options** menu.
- 3) In the Properties dialog, select the **Recording** option.
- 4) Make sure that all the needed inputs are checked in the checklist.
- 5) Choose **OK**. You will now see the recording volume controls.
- 6) Choose the input you want by checking the appropriate **Select** check box and make sure the volume is not too low.

The **Monitor** recording option can help you adjust the volume level for recording.

See Also

[Device Controls Setup](#)

[Device Controls](#)

[Playback](#)

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 more detailed information, refer to the **goldwave.txt** file.

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

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

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 folder. 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 number of sound attributes, as explained below. These attributes must be specified when creating a new sound. You can choose one of the three **Quick settings** buttons to automatically set appropriate attributes based on the class of audio, then just specify a length.

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.

8000	Telephone quality.
11025	Low end radio quality. Good for voice.
22050	Radio quality. Good for music and voice.
44100	CD quality. Excellent for all sounds.

Time

This value specifies the length of the sound in minutes and seconds. You can enter any time up to 3 hours for the length. For example, enter **1:00** for one minute or **1:00:00** for one hour.

For CD quality sound use stereo, 44100Hz attributes. Note that this uses 176,400 bytes for every second of sound or about 10 million bytes per minute.

You can use mono, 11025Hz for vocal sounds. This give good quality and requires only 22050 bytes per second.

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

File Format

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

- 1) The file has no header information and has no file association.
- 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 the format classification of the audio data. The **PCM** format means the samples are binary data of 8, 12, 16, or 32 bits in size. The **Telephony** format means that the samples are encoded Dialogic ADPCM (VOX), μ -Law, A-Law, or ISDN A-Law. The **Text** format means the samples are plain text numbers ranging from -1.0 to 1.0 or -32768 to 32767.

Attributes

Specifies the actual nature and organization of the samples. If you copied audio from a CD, for example, the format would be PCM and the attributes would be "16-bit, stereo, signed". Dialogic VOX files often use Telephony, with "4-bit VOX ADPCM, mono" attributes, but can also use μ -Law or A-Law. Attributes are as follows:

Mono versus stereo

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**.

Signed versus unsigned

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 & Pentium 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.

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. A CD audio recording has a sampling rate of 44100Hz. A Dialogic VOX file can have a rate of 6000Hz or 8000Hz.

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 filename and type using [File | Save as](#).

If you want GoldWave to automatically assume a format for a file, you can use [Options | File types](#) to associate a format with the filename extension.

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 folder 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

Copy To

Copies the selection to a new file. This is ideal for dividing a large file into smaller sections.

To copy:

- 1) Move the start and finish markers to the part of the sound you want to copy.
- 2) Choose **Copy to** from the **Edit** menu or click on the **Cpy2** button.
- 3) Enter a filename and choose a file type and the attributes.
- 4) Choose **OK**.

See Also

[Save as](#)

[Copy](#)

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](#)

[Replace](#)

[Paste at](#)

[Paste new](#)

Paste New

Creates a new Sound window and copies the sound in the clipboard into the new sound. The new sound will have the attributes and length of the clipboard sound. This command is useful when you need to save part of a sound to a new file.

To paste part of a sound into a new sound:

- 1) Move the start marker 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) Choose **Paste new** from the **Edit** menu or click on the **PNew** button.

See Also

[Copy](#)

[Replace](#)

[Paste](#)

[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 submenu. 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** submenu.

See Also

[Mix](#)

[Copy](#)

[Replace](#)

[Paste](#)

[Paste new](#)

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](#)

[Replace](#)

[Paste](#)

[Volume](#)

Replace

Replaces the selection with the sound in the clipboard. The length of the current sound is changed 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 and finish marker to the section of sound you want to replace.
- 2) Choose **Replace** from the **Edit** menu.

See Also

[Mix](#)

[Copy](#)

[Paste](#)

[Paste at](#)

[Paste new](#)

Delete

Removes the selection from the 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 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

Marker Menu

Presents a submenu containing marker related commands. For more information about markers, see [Editing Overview](#).

Set

Sets the start and finish markers to an exact time or sample position.

Drop start

Drops the start marker at the current playback position. This command is enabled only during playback.

Drop finish

Drops the finish marker at the current playback position. This command is enabled only during playback.

See Also

[Editing Overview](#)

[Cue points](#)

Set Marker

Sets the start and finish marker to an exact time or sample position. To specify a time, choose the **Time** option and enter the time in hours, minutes, seconds, and thousandths of a second. For example, you could enter **1:04:27.873**. To specify a sample position, choose the **Sample** option and enter the position.

If you want the length of the selection to be aligned to a **CD sector** or **1 kilobyte**, select the appropriate option. When the **OK** button is pressed, the finish marker will be adjusted to align the selection length.

See Also

[Editing Overview](#)

[Cue points](#)

Dynamics

Dynamics alters the amplitude mapping (dynamic range) 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.

Echo

Create an echo or reverb in 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 submenu 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.

Noise reduction

Reduces noise in the selection.

Low/Highpass

Applies a lowpass or bandpass filter to the selection at the specified cutoff frequency.

Bandpass/stop

Applies a bandpass or bandstop filter to the selection with the specified frequency range.

Equalizer

Boosts or reduces fixed frequency bands using a 7-band equalizer.

Parametric EQ

Boosts or reduces frequencies. Up to 30 bands can be created with precise control over the center frequency, gain, and bandwidth.

User defined

Applies a filter to the selection using coefficients you specify.

User Defined Filter

The **User Defined Filter** command is for advanced users who want to implement their own filter. Up to 30 coefficients can be given. Almost any kind of linear filter can be created with this command because it uses a general digital filter equation:

$$b(0)y(n) + b(1)y(n-1) + \dots + b(29)y(n-29) = a(0)x(n) + a(1)x(n-1) + \dots + a(29)x(n-29)$$

Coefficients can be floating point values with magnitudes up to 10000. 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**. Note that **b(0)** must not be zero.

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. If, for example, you recorded a voice over a noisy phone line, the noise gate can remove the noise any place where the speaker paused or stopped talking. Noise gates do not remove background hiss from louder parts of the selection.

Attack 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. Start with a value of about 200 milliseconds.

Release time

This is the amount of time (in milliseconds) that it takes for the noise gate to fully open. A value of 50 or lower usually gives good results.

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 and decrease the attack time.

Anticipation

This settings allows the noise gate to predict when the gate should open or close. This significantly reduces the problem associated with analog noise gates where the beginning of the sound is chopped off during the time it takes to release the gate. A value of 10 makes the noise gate look 10ms head. In general, it is best to set this value equal to the release time.

Noise Reduction

Noise reduction helps eliminate unwanted noise within a sound, such as a background hiss, a power hum, or random interference. It cannot be used to separate or remove complex sounds, such as removing vocals from music.

You are presented with a spectral analysis window, with a shape line, and several other controls. The **Coordinates** group show the **x** and **y** coordinates when you click-and-drag a shape point. The **x** coordinate is the frequency in Hertz and the **y** coordinate is the magnitude in decibels. The time of the spectral analysis is given as **T**, in seconds. If you move the time scroll bar, located below the analysis window, the time can be changed to show the spectral analysis of a different part of the sound.

A reduction envelope is used to remove noise. The envelope can be created in three different ways, depending on the **Reduction envelope** setting:

Use shape

Lets you manually create an envelope shape or select a preset shape. See [Shape Controls](#) for information about creating shapes. By creating a horizontal line at about 75 dB, you can remove a hiss from a sound.

Use current spectrum

Creates an envelope based on the shape of the graph shown in the spectral analysis window. This is particularly useful for removing a complex buzz or hum. Before using this effect, find a time in the sound where only the noise can be heard. Use that time for the analysis window by adjusting the time scroll bar. Select this option, then choose **OK**. If the noise is consistent through the sound, this is usually the best setting to use.

Use average

Applies an averaging envelope throughout noise reduction processing. The envelope is continuously updated. Use this setting if the noise changes frequently through the sound.

Settings

The **FFT size** determines the detail of the spectral analysis and the noise reduction envelope. Usually values between 9 to 11 give the best results. The **Overlap** value specifies the percentage of the FFT size to process at once. A value of 75 is best. The **Scale** value lets you alter the reduction envelope scale. A value of 100 uses the envelope as it is. A value of 200 doubles the envelope, which double the amount audio removed from the sound. A value of 50 halves the envelope, which halve the amount removed. Normally it should be set to 100.

Low/Highpass

Lowpass filters block high pitched frequencies (treble), but allow low pitched frequencies (bass) to pass. They can be used to reduce high end hiss noise or remove unwanted sounds **above** the given cutoff frequency. If you were to apply a lowpass filter with a cutoff frequency of 1000Hz on speech, it would make it sound mumbled and deep. Lowpass filters can also be used to eliminate aliasing when used before downsampling.

Highpass filters block low pitch frequencies, but allow high pitched frequencies to pass. They can remove deep rumbling noise or remove unwanted sounds **below** the given cutoff frequency. If you were to apply a highpass filter with a cutoff frequency of 1000Hz on speech, it would make it sound thin and hollow.

Cutoff frequency

The **Initial** box specifies the constant cutoff frequency for static filtering. If the **Dynamic** option is selected (see below), then a final cutoff frequency can be given in the **Final** box.

Filter options

Select **Lowpass** if you want to keep only the frequencies below the cutoff frequency. Select **Highpass** if you want to keep only the frequencies above the cutoff frequency.

If you want the cutoff frequency to remain constant throughout the selection during processing, select the **Static** option. If you want the cutoff frequency to change from the initial value to the final value, select the **Dynamic** option. Note that dynamic filtering will take more processing time.

The **Steepness** value specifies how sharply the filter cuts off frequencies outside the cutoff frequency. A higher steepness makes the filter sharper, but it also increases processing time. In technical terms, the steepness specifies the number of second order cascade filters used.

Examples

To make speech gradually become more hollow and thin:

- 1) Enter 60 in the **Initial** box.
- 2) Choose **Dynamic**.
- 3) Enter 1000 in the **Final** box.
- 4) Choose **Highpass**.
- 5) Choose **OK**.

Filtering before downsampling from 44100Hz to 22050Hz:

- 1) Enter 11025 in the **Initial** box.
- 2) Choose **Lowpass**.
- 3) Choose **Static**.
- 4) Enter 20 in the **Steepness** box.
- 5) Choose **OK**.

Bandpass/stop

Bandpass filters block all frequencies outside the specified range, keeping only frequencies within the range.

Bandstop filters block all frequencies within the specified range, keeping all other frequencies outside the range.

Frequency range

The **From** and **To** boxes specify the frequency range of the filter. If the **Dynamic** option is selected, then a final frequency range can be given in the other **From** and **To** boxes.

Filter options

Select **Bandpass** if you want to keep only the frequencies within the range. Select **Bandstop** if you want to keep only the frequencies outside the range.

The remaining options are explained under the [Low/Highpass](#) command.

Equalizer

The **Equalizer**, which is commonly found on stereo systems, boosts or reduces certain ranges of frequencies.

Center frequencies for each of the 7-bands are given at the top of each scroll bar. The scroll bars can be adjusted to boost or reduce a band by +12dB to -24dB.

Presets

Predefined settings can be read from [presets](#).

Parametric EQ

The **Parametric Equalizer** is a flexible tool for reducing or enhancing ranges of frequencies. GoldWave presents an easy to use interface where all the parameters for up to 30 band can be quickly configured.

Graph window

The graph shows frequency on the x-axis in Hertz and the gain on the y-axis in decibels. Each enabled band is displayed in the graph as a diamond shaped box located at its center frequency and gain. The width of the box shows the bandwidth. The currently selected band is shown in blue and its exact settings are shown by the scroll bars.

Controls

A band can be configured by selecting its number from the **Select band** box and adjusting the scroll bars. A quicker way is to **drag-and-drop** the band to a new location on the graph. Note that because of the logarithmic frequency scale, the width of a diamond changes as you move it left or right. The bandwidth, however, remains constant.

Any bands that are not needed can be disabled by unchecking the **Enabled** box. Disabling unused bands will speed up processing.

Presets

Predefined parameters can be read from [presets](#).

Flange

Flange uses variable delays and mixing to create some unusual audio effects. Parameters can be saved to 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.

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 simulated 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 the waveform 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**.

Pitch

Changes the pitch of the selection. This is useful for converting instrument samples from one note to another. The pitch can be changed by using a scale factor or by specifying semitone and fine tune values.

Scale

This option scales the pitch by the value you specify. If you set the scale to **0.5**, that will be equivalent to a downward shift by one octave. A value of **2.0** is the same as an upward shift of one octave and would make a voice sound like a chipmunk. A value of **0.75** would make a woman's voice sound like a man's.

Semitone

This option changes the pitch by semitones (notes). If your sound is a note at middle C and the semitone value is **2**, the note will be changed to D. A value of **-1** changes the note to B. A value of **12** make the note one octave above middle C. The **Fine tune** value lets you make a slight pitch adjustment in hundredths of a semitone. For example, a value of **50** would let you change a note from C to halfway between C and C#.

Preserve length

If this option is checked, a complex algorithm will be used to keep the length of the original note the same as the new note. In other words, the tempo will not be changed. In terms of a voice, this changes the pitch without speaking faster or slower. This option requires a substantial amount of processing time. In general, the **FFT size** should be set from 9 to 11 and the **Overlap** should be at least 88. Values of 90 and 95 will give better results, but require more processing time.

See Also

[Playback rate](#)

[Doppler](#)

[Time warp](#)

Stereo Menu

Presents a submenu containing stereo related commands. These commands work only with stereo sounds.

Exchange channels

Exchanges the left and right channels of a stereo sound.

Pan

Dynamically sets internal balance and left/right channel panning.

Remove vocals

Removes vocals from music. In some stereo recordings, it is possible to remove vocals by subtracting the left and right channels. In order for this to work, vocals must be located in the exact center of the stereo image. Unfortunately, any instruments located in the center will be removed as well. Another side effect is that the stereo image is lost and the final sound will be monaural.

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 over time or pan to and from each channel.

Time Warp

Changes the playback speed or stretches/compresses the length of the selection. Three different techniques are provided, each with certain advantages and disadvantages. All three let you specify the change either by a speed factor or by a new length. The **Speed** factor lets you specify a relative change. A value of **0.5** make the selection play twice as slow. A value of **2.0** make the selection play twice as fast. The **Time** options lets you specify a new length for the selection. This is useful if you need to make a sound fit a certain time, such as squeezing a 35 second commercial into a 30 second spot.

Speed

Changes the sampling rate of the entire sound so that it plays back at a different speed, similar to spinning a record faster or slower. It works the same way as the speed scroll bar in the Device Controls window, but in this case, the sound itself is changed. This technique is very fast and produces excellent quality, however, the pitch of the sound is changed as well.

Similarity

Uses correlation to add or overlap small, similar sections of the sound. This technique preserves the pitch and generally produces high quality for voice and fair quality for music, when using small changes. A fair amount of time may be require for processing, depending on the **Search range** value. For voice, the **Window size** should be set between 20 and 30 and the **Search range** set to between 5 and 10. For music, a larger **Window size** and **Search range** gives better results, such as 50 and 25.

FFT

Uses Fourier transforms and interpolates or decimates the spectral analysis to change the length. This technique preserves the pitch, but can introduce some artifacts into the sound. Best quality is obtained by using the **Oscillator synthesis** option, but that requires significant processing time. The **FFT size** should be set from 9 to 11 and the **Overlap** should be at least 75, but can be set to 88, 90, or 95 for slightly better quality.

Note: If you changed the Device Controls playback speed, remember to set it back to 1.00 so that the sound plays at the correct speed.

See Also

[Playback rate](#)

[Resample](#)

[Pitch](#)

[Device Controls Window](#)

Volume Menu

Presents a submenu 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). Shows rms and maximum volume levels.

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** submenu.
- 3) Enter the new volume percentage, then choose **OK**.

Fade In

Gradually increases the volume throughout the selection. An initial volume 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** submenu.
- 3) Enter the initial volume percentage, then choose **OK**.

See Also

[Fade out](#)

Fade Out

Gradually decreases the volume throughout the selection. The fade value is the amount that the volume will 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** submenu.
- 3) Enter the fade percentage, then choose **OK**.

See Also

[Fade in](#)

Maximize

Searches the selection for the maximum volume level, then displays the rms (root-mean-square) level, the maximum level and the time at which the maximum occurs. You can then specify a new maximum for the sound. A value of **1.0** normalizes the volume so that it covers the full dynamic range. A value of **0.5** sets the volume to half the dynamic range. Values above **1.0** should not be used if clipping is to be avoided, since the volume will be increased beyond the allowable dynamic range.

To normalize 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** submenu.
- 3) Set **Volume** to **1.0**.
- 4) Choose **OK**.

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 over time.

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](#)
[Time warp](#)
[Pitch](#)

Resample

Changes the sampling rate of the entire sound. Unlike **Playback rate**, this command recalculates 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](#)

[Time warp](#)

[Pitch](#)

Extract AVI Audio

The **Extract AVI Audio** tool lets you copy audio from an AVI file into a wave file, where it can be edited. You select the source AVI file by clicking the folder button. If the file contains more than one audio stream, you can select the stream number to extract in the **Audio stream** box. The destination wave filename is automatically set to the same name as the source AVI file. If you need to change the name, click the folder button.

Windows 95 Note: A bug in the Microsoft AVISave function prevents the wave file from being closed properly. You will have to close GoldWave before you can open the new file.

Extract CD Audio

The **Extract CD Audio** tool lets you digitally copy audio directly from an audio CD to a file. Note that only some SCSI CD-ROM devices are supported and you must have an ASPI driver installed. ATAPI and EIDE devices are not supported.

If your system is configured correctly, you will see a dialog that allows you to select a CD device and specify the region of audio to copy. You can preview the start and finish times of the region by using the Start and End buttons. Choosing the **Save** button will prompt you for a filename and format and then begin extraction.

The **Read technique** group lets you manually specify what SCSI command to use to read from your CD-ROM drive. In some cases, GoldWave will automatically select an appropriate technique. In general, you will have to try each command to see which one works on your system.

The **Correct jitter** option helps to eliminate pops and clicks caused by slight reading errors. When enabled, the CD is read more than once to detect and correct any errors. If your CD-ROM drive automatically fixes these errors (such as Plextor or CD-R devices), you can disable this option to speed up extraction.

The **Byte swap** option changes the order of bytes (endian) extracted from the CD-ROM. If the audio you extracted is badly distorted or a loud hiss, you will need to check this box.

Troubleshooting

If you see an error message indicating a read problem, make sure that the CD is free of dust and finger prints. If the CD appears clean, try using a different read technique. If you have tried all of the techniques and still get an error message, then CD audio extraction may not be possible on your system.

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 choose folders and file storage options.

Sound files

Specifies the folder where you keep your sound files. The **File | Open** command will automatically list files in this folder whenever you start GoldWave. If you have already configured a working folder using short cut properties, enter a period, ".", for this folder.

Temporary

Specifies the folder to use when creating temporary files. This folder 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 folder does not affect opened files already in temporary storage.

Undo

Specifies the folder to use for storing undo data. In most cases, this should be the same as the temporary folder. Changing the undo folder 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 opened. Any file larger than this will be flash opened.

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.

Clipboard

Due to size limitations in the standard Windows clipboard, copying large sections of audio can cause problems. If you encounter any unusual behaviour or frequently work with large files, choose **GoldWave** to use GoldWave's special clipboard. If you work with small files or need to copy and paste audio to other applications, choose **Windows** to use the standard clipboard.

File Types

Associates a filename extension (such as **.snd** or **.vox**) with an audio format. This is useful for automatically opening files that do not contain any information describing their format. For example, if you work with Dialogic files, you can associate the **.vox** extension with one of the Telephony formats. Whenever you open a **.vox** file, GoldWave will assume the format you specified and open the file without prompting you for the format.

To associate **.vox** with **A-law** encoding:

- 1) Enter the extension **vox** in the box. Do not enter the leading period.
- 2) Select **8000** from the Rate drop down list.
- 3) Select **Telephony** from the Format drop down list.
- 4) Select **A-law, mono** from the Attributes drop down list.
- 5) Choose the **[+]** button.

To remove an association:

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

To change an association:

- 1) Remove the association, as above.
- 2) Add the new association, as above.

Tool Bar Options

Configures the Main tool bar and Effects tool bar. You can control what buttons to show or hide. To add or remove a button, simply find the button and click on it. Buttons that are grayed will not appear in a tool bar.

Window Options

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

Main window size

Controls the Main window's position and size every time GoldWave is started. Choosing **Normal** lets Windows determine the size and position. **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.

Sound window size

Controls the position and size of Sound windows. The **Normal** option lets Windows determine the size and position. **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 they all will be visible.

Amplitude axis

Displays or sets the units of the vertical axis in Sound windows. Selecting **Off** hides the axis completely. **Normalized** shows an axis with a range of -1.0 to 1.0. **Signed 16-bit** shows an axis ranging from -32768 to 32767, which is the range of a 16-bit sample. **Unsigned 8-bit** shows an axis with a range of 0 to 255.

Time axis

Sets the format for displaying the horizontal time axis in Sound windows. **Seconds** gives the time as a floating point number, such as 1234.567. **Minutes** gives the time as minutes:seconds, such as 12:34.567.

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 Controls Setup

The **Device Controls Setup** property sheet lets you configure the playback and recording features, the oscilloscope graphs, and the recording and playback devices:

[Play Properties](#)

[Record Properties](#)

[Graph Properties](#)

[Device Properties](#)

Device Controls Play Properties

The **Device Controls Play Properties** page lets you configure the user play button and set the speed of the fast forward and rewind buttons.

User play button

Configures the **user play** button to play and/or loop sections of the sound as follows:

- | | |
|-----------------------|--|
| All | Play entire sound. |
| Selection | Play region between start and finish markers. |
| Unselected | Play regions outside the start and finish markers. This lets you quickly test how a cut or delete will sound without actually changing the sound. When possible, playback is confined to the region shown in the Sound window view so that the entire file does not have to be played. |
| View | Play all of the sound shown in the Sound window view. |
| Intro/loop/end | This is a special playback feature that plays the sound in three sections. The beginning of the sound, outside the selection, is played first. Then the selection is played 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.

Fast/Rewind speed

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.

Device Controls Record Properties

The **Device Controls Record Properties** page lets you configure recording features and options.

Recording options

Monitor input connects the input device to the oscilloscopes and LED meters so you can adjust input levels before recording. See [Recording Sounds](#) for information about selecting a different input and setting volumes.

Loop keeps recording over and over until the stop button is pressed. This is useful for capturing quotes from a TV show or Freudian slips in a speech. By loop recording a 1 minute sound, you will always have the last minute of audio stored for recall.

Ctrl key safety prevents you from accidentally recording over a sound. To record, you must hold down the **Ctrl** key. If you do not, a safety message appears.

Allow undo saves the entire selection so that you can undo recording after. If you are recording a large file, you should uncheck this setting.

Mix with selection mixes the newly recorded sound with the current sound. This lets you layer recordings. You can do this by using the Copy and Mix [Edit Menu Commands](#) as well.

Disable write cache immediately writes the newly recorded sound directly to the hard drive. Normally, Windows will temporarily save the sound in memory before writing it to the hard drive (caches it). However, when recording large files, the memory quickly becomes full and Windows is forced to write a large amount of data. This can cause skipping and gaps in the recording. In some cases, disabling write cache can solve this problem.

Delayed recording

Countdown timer delays recording for the number of seconds specified in the **Time** box. This gives you some time to get ready after you start recording or lets you automatically record something at a later time.

The **Level activated** feature pauses recording until a sound is detected. You can think of it as an automatic recording system. The **Threshold** specifies how loud a sound should be before recording begins. Normally, this value should be small (less than 0.2). The **Duration** specifies how long to record after the sound becomes quiet again. Entering a zero value allow recording to continue without stopping. Level activated recording is quite useful for automatically synchronizing a recording.

Device Controls Graph Properties

The **Device Controls Graph Properties** page lets you configure oscilloscopes and LED meters

Graph

This group allows you to set the type of oscilloscopes to display. Clicking on an oscilloscope in the Device Controls window 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 analysis.
- 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.

Show axis

Draws x and y axes on the oscilloscopes. Note that you can resize the Device Controls window to make the oscilloscopes more detailed. The frequency range of the three frequency graph types depends on the sampling rate of the sound. The 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.

Window function

Lists several common window functions used for FFT frequency analysis to minimize artifacts caused by a finite analysis window (and all that stuff).

Refresh

This group controls the frames-per-second rate and the LED hold time. **Frames/s** sets the number of times per second that oscilloscopes and LED meters are updated. **LED hold time** is the amount of time, in milliseconds, that the peak LED level is retained.

Device Controls Device Properties

The **Device Controls Device Properties** page lets you select and configure the devices for recording and playback.

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.

Buffer size

The **Playback** box controls the internal size of the buffer (in seconds) to use for playback. The **Record** box controls the internal size of the buffer (in seconds) to use for recording. In general, the buffer sizes should be set to one second or less.

Playback device and Record device

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. Under each sampling rate (11025, 22050, and 44100) are two boxes indicating mono and stereo support. If only one box is green, the device supports mono only. If both boxes are green, then stereo is supported. You can select any driver to see its capabilities. To play or record CD quality, the driver must support 16 bit, 44100 Hz, stereo.

Important: The **Microsoft Sound Mapper** is not really a device. It converts audio so that it can be played or recorded by your sound card. For best quality, you should avoid using the **Microsoft Sound Mapper**. This is especially true if you are trying to play sounds at rates other than the ones given above.

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 [+] (add) button, and a [-] (remove) 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 [/] button.

To save a shape, use the presets controls.

GoldWave Order Form

Please type in the following information:

Full Name (first & last): _____

Address: _____

Country: _____

[Telephone (with Area Code)]: _____

E-mail address: _____

Version you are using: v4.00

Registration:

- Standard (US\$30, Can\$40, UK£20)
 Password to remove shareware message and unlock
 all future versions.
- Deluxe (US\$55, Can\$77, UK£37)
 Password (as above). Diskette and manual for
 new version, when available.

Payment:

- Check/money order enclosed (payable to Chris Craig)
- VISA credit card (charged in Canadian currency)

VISA number: _____

Expires: / /

Cardholder name: _____

Signature: _____

Please mail this form to the address below. Thank you for registering!

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

Registering

GoldWave is a shareware program. To register and encourage further development, please fill in the Order.doc file. Registration details are given below. If you require more than one registration, please send the full name of each person who will be using GoldWave.

Note: Passwords cannot be sent by e-mail.

Standard Registration (US\$30, Can\$40, UK£20)

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://www.goldwave.com>

Deluxe Registration (US\$55, Can\$77, UK£37)

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 the newest version of GoldWave.

Sending Payment

Checks from banks in the **United States, Canada,** and the **United Kingdom** are welcome. VISA credit card orders are accepted. Other credit cards cannot be accepted at this time.

Customers in **Germany** can register through the KDL-Registration Service. Please contact them and they will forward your order:

KELLY Data GmbH
Lena-Christ-Str. 50
82152 Martinsried
Tel: 089/89 56 26-50
Fax: 089/89 56 26-99
norbert.burkhard@kellydata.de
CIS: 100347,2627
AOL: Kelly Data

Customers in **France** can register through WinShare. Please contact them and they will forward your order:

WindowShare, SARL
46-48 Route de Thionville
57140 WOIPPY
Tel: (+33) 87.30.85.57
Fax: (+33) 87.32.37.75
Web: <http://www.wska.com>

International postal money orders or **international postal bank drafts** are also accepted, but payment must be made in Canadian currency.

[Register](#)

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.

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.

Registered users please note that this warning is always displayed for beta versions.

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 Web, BBS, or FTP sites, posted to newsgroups or forums, or placed on CD-ROMs. A non-beta shareware release will be available shortly that can be redistributed.

Click [here](#) if you have discovered a bug, an error in the help/text files, or would like to send some feedback.

Please check the [Contents](#) for an introduction.

The **readme.txt** file contains helpful information and important installation instructions. The **goldwave.txt** file contains the manual, with instructions for using GoldWave. The **whatsnew.txt** file lists many of the changes made to this version and past versions. For recent information and software updates, please refer to the GoldWave WWW Home Page:
<http://www.goldwave.com>

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) each session. When the limit is reached, a toll message will appear whenever you use a control in the Device Controls window. Exiting and restarting GoldWave will let you use another 150 commands without interruption. By [registering](#), you will receive a password that removes this toll completely.

Bugs

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

<http://www.goldwave.com>

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 commands that led to the error. The best report to send would be a sequence of commands that is guaranteed to recreate the error. Please do not send sound files to me through e-mail.

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

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

