

Contents

GoldWave is a sound editor, player, recorder, and converter. It can create sound files for CDs, websites, answering machines, or Windows sounds. A full 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 Recorder 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](#)

Getting Started:

[Playing Sounds](#)

[Recording Sounds](#)

[Selecting Part of a Sound](#)

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 the waveform with the mouse](#)

[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, and information/status bars. This window keeps all the **Sound** windows 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 real-time graphs and sound device controls. The graphs 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 properties 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.

Unlike the standard Windows "click-and-drag" selection method, which forces you to change both the start and end points each time, GoldWave uses an innovative selection method that lets you change the start and finish markers independently. You simply use the **left** and **right** mouse buttons.

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 the Waveform](#)

[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 few seconds of the sound are 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 14 years, working in a variety of languages including C, C++, and 0x86 assembly. His recent works include Multiquence (a real time Windows multitrack audio editor) and GoldWave (a Windows digital audio 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

Website:

<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. Note that to open MPEG Layer 3 files (.mp3) you must have the new Windows Media Player installed.

Close

Closes the current sound.

Information

Sets the file's title, author, description, and copyright information. This information is saved in .wav, .aiff, and .mp3 files only.

Batch conversion

Converts a group of files to a different format.

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.

Save selection as

Saves the selected part of the current sound to the disk under a different name and/or format. See 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 the Waveform](#)

[Editing Overview](#)

Effects Menu Commands

Doppler

Dynamically changes pitch.

Dynamics

Distorts, compresses or limits output levels.

Echo

Produces an echo.

Expand/Compress

Compresses, limits, or expands dynamic range.

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.

Previous zoom

Returns the view to the previous zoom level. Use this to switch back and forward between two different zoom levels.

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 waveform 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 waveform 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 the Waveform](#)

View Other

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.

CD audio extraction

Makes exact digital copies of music on audio CDs.

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	Moves the start marker right.
Shift+Left	Moves the start marker left.
Ctrl+Shift+Right	Moves the finish marker right.
Ctrl+Shift+Left	Moves the finish marker left.
Tab	Moves the start and finish markers to the stored locations.
Shift+Tab	Stores the locations of the start and finish markers.
Shift+Up	Horizontally zooms in.
Shift+Down	Horizontally zooms out.
Shift+A	Horizontally zooms all the way out.
Shift+P	Zooms to previous horizontal zoom.
Shift+S	Horizontally zooms in on the selection.
Shift+U	Horizontally zooms to the user defined level.
Shift+1	Zooms 1:1 horizontally.
Ctrl+Up	Vertically zooms in.
Ctrl+Down	Vertically zooms out.
Shift+V	Vertically zooms all the way out.
Scroll Lock	When turned on, the Sound window view automatically scrolls to follow the playback position.

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.
Ctrl+N	Creates a new file.
Ctrl+O	Opens a file.
Space	Plays a sound.
F4, F5, F6, F7, F8	Plays, rewinds, pauses, fast forwards, 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.
F11	Displays Device Controls Properties window.

Editing

Keystroke	Action
Ctrl+Z or Alt+Backspace	Undoes last change.
Ctrl+X or Shift+Del	Cuts the selection.
Ctrl+C or Ctrl+Ins	Copies the selection.
Ctrl+V or Shift+Ins	Pastes the clipboard into the sound at the start marker.
Ctrl+B	Pastes the clipboard into the sound at the beginning.
Ctrl+F	Pastes the clipboard into the sound at the finish marker.
Ctrl+E	Pastes the clipboard into the sound at the end.
Ctrl+P	Pastes the clipboard into a new Sound window.
Del	Deletes the selection.
Ctrl+M or Shift+Ctrl+Ins	Mixes the clipboard with the sound at the start marker.
Ctrl+T	Trims the sound. Removes all audio outside the selection.
Ctrl+R	Replaces the selection with the sound in the clipboard.
Ctrl+A	Selects the entire sound.
Ctrl+W	Sets the selection to the view.
Ctrl+Q	Drops a new cue point at the current playback position.
[Moves the start marker to the current playback position.
]	Moves the finish marker to the current playback position.

File Information

Use this dialog to set the title, author, description, and copyright notice. These appear under Windows Explorer when you right click a sound file, choose **Properties**, then view the **Details** tab. Note that information is saved only in **.wav**, **.aiff**, and ID3v2 **.mp3** files. Information is discarded for other file types.

Batch Conversion

Use this window to convert a set of files from one format to another. You can choose files by using the **Add files** button or by using drag-and-drop with Windows Explorer.

You can remove files from the list by selecting one or more files and choosing the **Remove** button. The **Clear** button removes all files from the list.

The destination file types and attributes are the same as given under the Save as command. You must select the file type before selecting attributes since attributes vary from one type to the next.

The **Destination folder** specifies where all the converted files will be stored. If the destination folder is blank, the converted files are stored in the same folder as the original file.

The **Overwrite existing files** option allows you to replace the original files with the converted format. Overwriting only occurs if the destination folder and file type match the original.

The **Author** and **Copyright** text is stored in file types that support such information as explained under the File | Information command.

Press the **Begin** button to start converting all the files. A status window will appear to show the progress of the conversion and list any errors.

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 **.mp3**, and a wide variety of compressed formats.

The **Set custom** button saves the current file type and attribute settings. The **Custom** button retrieves those settings. This is useful if you frequently use certain attributes when saving files and would like to quickly select them.

For more details about attributes, see **Attributes** under [File Format](#).

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 websites.

To create an **.mp3** (MPEG Layer 3) compressed file, choose the **MPEG Audio (*.mp3)** type and select one of the listed attributes. You must have an MPEG encoder installed, otherwise an error message will appear. See the GoldWave [website](#) for details.

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 the Waveform

GoldWave lets you directly edit the waveform 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) Choose **Deflash** from the **Edit** menu, if enabled.
- 2) Zoom in 2:1 or closer (such as Zoom 5:1).
- 3) Place the mouse arrow directly over the waveform. The arrow will change into a thin horizontal line with two small arrows pointing to it.
- 4) Click and hold the left mouse button.
- 5) Move the mouse up, down, left, or right to change the samples and reshape the sound.
- 6) 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  (properties) button to make sure that playback and recording devices have been selected. Many device related features can be configured through the [Device Controls Properties](#) window. To save screen space or enlarge the graphs, 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 Properties](#) 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  faders. The **playback speed** of the device can be controlled with the  fader. Note these controls affect only the playback 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 Properties](#)


[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 play.
- 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) Use the **left** and **right** mouse buttons to move the **start** and **finish** markers to select 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, use **Edit | Select All**.

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  (properties) 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 Properties](#)

[Device Controls](#)

[Recording](#)

Recording Sounds


To record a sound:

- 1) Choose **New** from the **File** menu.
- 2) Choose attributes (or choose a **Quick setting** button) 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 recording source, 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 Properties** window, including a countdown timer and level activated recording. See the [Device Controls Properties](#) for more information.

To select a different recording source, such as the line-in, microphone, or CD audio, use one of the following methods:

Method 1: Use Device Controls Properties

- 1) Select the  (properties) button.
- 2) Choose the **Volume** tab.
- 3) Check the appropriate **Select** box and adjust the volume.
- 4) If the device has a master recording control, make sure the **Mute all** box is not checked and that the master volume is not zero.

Method 2: Use the Volume Control Accessory

- 1) Select **Volume control** from the **Tools** menu to start the **Volume Control** accessory from the Windows **Start** menu.
- 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 sources 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.

GoldWave's **Monitor** recording option can help you adjust the volume level before recording.

See Also

[Device Controls Properties](#)

[Device Controls](#)

[Playback](#)

Cue Points

The Cue points tool (**Tools | Cue points**) allows 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. Right-click on a cue point in the Sound window to see a context menu. Cue points are saved only in Wave files.

The **Move marker** button moves the start or finish marker to the currently selected cue point.

The **Copy all** button copies all the cue point information to the clipboard. You can then paste it into a text editor, such as the Notepad accessory.

The **Split file** button divides the file onto smaller files using the cue points as split points. If you've recorded one side of an album and need to divide it into individual songs, for example, you would set a cue point near the start of each song, then use this feature to automatically create a file for each song. Each file can then be written to a CD-R disc as a track.

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**.

Once a cue point has been added, you can right-click on the cue point flag in the Sound window to reveal a context menu. A cue point can be dragged-and-dropped to a new location using the left mouse button.

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 **Move marker** 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.

Split Options

Use this feature to divide a large file into smaller files. The file is split at each cue point. Smaller files are stored in the specified **Folder**. The **Base name** is used together with a 3 digit number to create the name of each file. If you entered **Track**, for example, the files would be named Track000, Track001, Track002, etc.

The **Overwrite existing files** option automatically replaces files with the same name that already exist in the folder.

The **Use CD format and alignment** ensures that each file is stored in CD compatible format and that the length of the audio is exactly aligned to a CD sector boundary.

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.htm** file.

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 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 is automatically 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. Higher values allow a wider range of frequencies to be recorded and give better quality. Smaller values require less space. 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.

If you see this dialog for **.mp3** files, you need to install the new Windows Media Player available from Microsoft. GoldWave cannot decode **.mp3** files without the Windows Media modules.

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 select this attribute.

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.

Recall positions

Moves the start and finish markers to the previously stored locations. This acts like the "Recall" button on a calculator.

Store positions

Stores the start and finish marker positions. This acts like the "Store" button on a calculator.

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.

Snap to zero-crossing

You can check or uncheck this option. Whenever you move a marker, its position is changed so that it coincides with a place at which the waveform approaches zero. This helps to reduce any clicks caused by a sudden change in the waveform after editing. Note that the marker may be snapped to a place that is several hundred samples away from the position where you released it. The best zero-crossing point is used, which is not necessarily the first point.

When editing a stereo file, both the left and right channels are scanned for the best overall position for the marker. You can use the [Edit | Channel](#) command to limit the zero-crossing search to a single channel.

When zoomed in 1:2 or closer, snapping is automatically disabled so you can place the marker at any position.

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.

See Also

[Expand/Compress](#)

[Noise gate](#)

Expand/Compress

The Expand/Compress effect is a general purpose dynamics processor that includes compressors, limiters, expanders and gates. See the manual for more detailed information.

The **Ratio** specifies the compression or expansion ratio. For compression, this value should be less than 100%, typically between 25% and 75%. For expansion, the value should be less than 100% as well.

The **Threshold** specifies the envelope level to activate the expander or compressor. Compressors change the volume level of all sounds above that level. Expanders change the volume level of all sounds below that level. Depending on the **Smoothness** setting, the threshold may have to be set much lower than expected.

The **Smoothness** specifies how quickly the compressor/expander changes from one volume level to the next and how quickly it activates. Using 0% means that volumes will change instantly, which can cause a rough distortion in sections of audio that border on the threshold level. A value of 100% means that volumes will change gradually over 100ms. With a high smoothness setting, the threshold will have to be reduced. The higher setting makes the envelope detector respond more slowly to changes in the sound, resulting in a lower envelope range.

Use the **Expander** and **Compressor** options to specify what processing is required.

See Also

[Dynamics](#)

[Noise gate](#)

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.

Pop/Click

Finds and repairs pops and clicks.

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.

See Also

[Expand/Compress](#)

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 four different ways, depending on the **Reduction envelope** setting as explained below. The **Use clipboard** option often gives the best results.

Use shape

Lets you manually create an envelope shape or select a preset shape. See [Shape Controls](#) for information about creating shapes. By setting a horizontal line at about 75 dB, you can remove a hiss from a sound. The **Preset shapes** group allow you to store or recall a shape. Preset shapes do not apply to any of the other envelope settings.

Use current spectrum

Creates an envelope based on the actual shape of the graph shown in the spectral analysis window. This is 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 the easiest 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.

Use clipboard

Analyzes the audio in the clipboard and creates an envelope based on it. To use this setting, you would select the part of the sound where the noise exists by itself, use the Copy command, select the all or part of the sound you want to remove the noise from, then use Noise reduction with this option. Like the **Use current spectrum** setting, this can remove a buzz, hum, static, and other common noises. However, with this setting the noise can be based on another file or outside the current selection. In general, this is the best setting to use.

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 overlap from one calculation to the next. 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](#).

Pop/Click

A pop/click filter is a specially designed filter that searches for abrupt changes in the sound and eliminates them.

The **Tolerance** defines how abrupt a change can be before it is considered a click. It is best to start with a value near 1000%. Lower values will detect more clicks. Values less than 500% should be used on short selections only.

The filter requires a minimum selection of 4000 samples to establish a base line.

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 given in edit box controls. A short time frequency analysis graph is drawn with the left channel in green and the right channel in red. The time of the analysis can be changed using the scroll bar located at the bottom of the window.

Controls

A band can be configured by selecting its number from the **Select band** box and adjusting the Gain, Center, and Width controls. A quicker way is to **drag-and-drop** its diamond box 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 command 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 or average) level, the maximum level and the time at which the maximum occurs. You can then specify a new maximum or a new RMS value for the sound. A maximum of **1.0** normalizes the volume so that it covers the full dynamic range. A maximum of **0.5** sets the volume to half the dynamic range. Maximums above **1.0** should not be used if clipping is to be avoided, since the volume will be increased beyond the allowable dynamic range.

Setting a new RMS value is useful for matching volumes across several files. Files that have a similar RMS value usually appear to have the same volume. If you are creating a CD with several songs that have widely varying volumes, adjusting their RMS values may help to even them out. It is best to check the RMS values of all the songs to determine an appropriate (average or minimum) value.

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 **New maximum** 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 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](#)

[Time warp](#)

[Pitch](#)

CD Audio Extraction

The **Extract CD Audio** tool digitally copies audio directly from an audio CD to a file on your hard drive, without using your sound card. This features has several advantages over normal recording:

- There is no need to create and initialize a new file.
- Audio can be compressed while extracting, saving lots of hard drive space.
- Recording volume controls do not have to be selected or adjusted.
- Synchronizing recording with CD playback is unnecessary.
- Some CD-ROM drives allow extraction several times faster than recording.

Note that only recent SCSI and ATAPI (IDE) CD-ROM devices are supported. Devices that are MMC compliant are supported using the **MMC** read technique. For Windows 95 and 98 you must have an ASPI driver installed. Due to the wide variety of interfaces, inconsistent device standards, and problems with the ASPI driver, incompatibilities may arise that will require a system reset. It is recommended that you close all other programs before proceeding.

Before using this command, insert an audio CD in the drive.

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. If you have only one CD-ROM drive, only one device will be listed in the drop down list. If you already have an audio CD in the drive, the track times will be shown in the **From** and **To** time lists. Otherwise you will need to insert an audio CD and select the device from the list.

The **Preview from** button plays the first five seconds beginning at the **From** time. The **Preview to** button plays the last five seconds up to the **To** time.

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 **Fix defects** option helps to eliminate pops and clicks caused by slight reading errors (also know as jitter). 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 get synchronization errors at the beginning, try continuing several times before giving up. Some drives are very inaccurate when spinning up and may require a dozen continues before synchronization can be reliably established.
- 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.

The **Rate** is the sampling rate of the audio.

The **Skip bytes** value is the number of header bytes to skip at the beginning of the file, if a header is present. Note that the header is lost when you save the file.

The **Format** and **Attributes** lists specify how the audio data is to be interpreted. The format must be selected first since the attributes vary from one format to the next.

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.

The **OS Association** button associates GoldWave with the file type in Windows so that when you double-click on the file in Windows Explorer, GoldWave starts and opens the file. **OS Association** adds an Open and a Play command to the context menu when you right-click on the file.

To associate GoldWave with the **.wav** file type in Windows:

- 1) Enter the extension **wav** in the box. Do not enter the leading period.
- 2) Choose the **OS Association** button.

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.

Initial zoom

When you open a new file, the waveform is drawn using this zoom level. If you choose **All**, then the entire waveform is shown using whatever zoom level is required. If you choose **Flash**, the entire file is drawn only if the file is initially deflashed.

Device Controls Properties

Use the Device Controls Properties  button to display the **Device Controls Properties**. These properties lets you configure the playback and recording features, the real-time graphs and LEDs, and the recording and playback devices:

[Play Properties](#)

[Record Properties](#)

[Volume 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

All

Plays entire sound.

Selection

Plays region between start and finish markers.

Unselected

Plays 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

Plays all of the sound shown in the Sound window view.

Finish

Plays three seconds just before the finish marker, so you can determine if the marker is in the right place without listening to the entire selection.

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 value of 1 loops playback once, so the selection is played twice. 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 graphs will move slowly through the data.

Device Controls Record Properties

The **Device Controls Record Properties** page lets you configure recording features and options. To set recording volumes, choose the [Volume](#) tab.

Recording options

Monitor

Connects the recording device to the graphs and LED meters so you can adjust recording levels before recording. See [Recording Sounds](#) for information about selecting a different recording source 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. This setting is ignored if the option Unbounded is checked.

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.

Unbounded

Continues recording until all storage is exhausted or until you press the record stop button. The file size is increased automatically to hold the new audio, which can save time when recording for a long period of time. You **must** enable [Hard disk](#) storage under [Options | File](#) when using this option. If you are using RAM storage, this option is ignored.

Delayed recording

Timer

Delays recording until the specified **Time** and **Day**. Use this feature to automatically record something at a later time. The time is given in 24 hour time. A time of **06:00:00** is 6:00 AM and a time of **18:00:00** is 6:00 PM. **00:30:00** is 12:30 AM or 30 minutes past midnight. If entering the time directly, remember to include the seconds. Entering **18:00** means 00:18:00. You must press the record button to activate the timer.

Level activated

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 Volume Properties

The **Device Controls Volume Properties** page lets you adjust recording volumes and select or unselect recording sources.

A volume fader and checkbox is shown for each source. To select a source, check the appropriate checkbox. If your sound card has a master control, make sure that the **Mute all** option is turned off and that the master volume is not zero.

You can use the Monitor option under the Record Properties tab to activate the LED meters and real-time graphs without recording. By moving the Properties window, you can see the levels as you adjust the volume faders.

Note that volumes are changed regardless of whether you choose **OK** or **Cancel**. Using **Cancel** will prevent the left and right graphs from being set the same.

To select a different recording device, use the Device tab.

Device Controls Graph Properties

The **Device Controls Graph Properties** page lets you configure real-time graphs and LED meters.

Graph

This list allows you to set the type of real-time graphs to display. Note that clicking on a graph in the Device Controls window changes the type. You can also use the right mouse button on a graph to select a type from a menu.

Amplitude

The sound is shown as a series of amplitudes.

Spectrum

The sound is shown as a frequency analysis.

Log bar spectrum

The sound is shown as a logarithmic frequency band 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, and red.

Colour spectrum

The spectrum of the sound is shown as a detailed coloured bar graph, similar to the spectrogram colours given above.

Fire spectrum

The spectrum of the sound is shown as a detailed fire coloured bar graph.

X-Y mode

The sound is plotted with the left channel against the right channel to generate a Lissajous pattern. This is often used to see the phase difference between two equal frequency signals. If the left and right channels are in phase, the pattern is a diagonal line running from the lower left to the upper right. If the channels are 90 degrees out of phase, the pattern is a circle. For general stereo sounds, it looks like a crazy scribble. The larger the scribble, the larger the difference between the channels. Monaural sounds always show a diagonal line since the left and right data are the same.

Colour X-Y mode

This is the same as X-Y mode graph above, but coloured points are used, ranging from low amplitude green to high amplitude red.

Colour Amplitude

This is the same as the Amplitude graph above, but coloured points are used instead of lines. Colours range from low amplitude green to high amplitude red.

Show axis

Draws x and y axes on the graphs. Note that you can resize the Device Controls window to make the graphs 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

Frames/s sets the number of times per second that graphs and LED meters are updated. **LED hold time** is the amount of time, in milliseconds, that the peak LED level is retained. Note that the top red overload LED remains lit until the sound is stopped. **Fade time** is the amount of time, in milliseconds, that it takes for the LED meters and the colour and fire spectrum graphs to drop from maximum to zero.

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 real-time graphs 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 playback 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. If you have an 8 bit sound card, you must select the Microsoft Sound Mapper as the Recording device.

Joystick control

Allows playback and recording to be controlled using a joystick. The first detected joystick is used. The directions control playback. Left is rewind, right is fast forward, down is pause, and up unpauses. The first button (button 1 or A) starts or stops playback. The second button (button 2 or B) starts or stops recording.

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. Some dialogs contain a clear button **[/]** that resets everything to default values.

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.

When performing FFT processing, the sound is processed one block at a time. To smooth out transitions from one block to the next, it is necessary to overlap blocks. The Overlap value controls what percentage the FFT analysis of one block overlaps the next. A high percentage makes the transition between each block more smooth. It also requires more processing time since overlapping samples are recalculated several times. A low percentage results in rougher transitions, but processes much faster. A value of 88% or higher is recommended.

When performing FFT processing, the sound is processed one block at a time. The FFT size value controls the size of these blocks. The number of samples to process is calculated by taking the value as a power of 2. A value of 10 gives 2 to the 10th power, or 1024.

Additional Installation

GoldWave requires additional packages to be installed for enhanced file support. They are not necessary to run GoldWave. However, if you are working with MP3 files or would like to extract audio from movie files, you need to install the following packages:

- **Media Player** Microsoft's new Media Player installs a number of codecs that GoldWave needs to read MP3 files and audio in movie files (AVI, MOV, MPG, etc.). If this is not installed, GoldWave will not be able to open or edit MP3 files.
- **LAME** To create and save MP3 files, you will need to install the LAME encoder.

For the download locations of these packages, please see the GoldWave website:

<http://www.goldwave.com/faq.html#install>

If MP3 and video files appear to be working in GoldWave, please disregard this message.

MP3 Files

GoldWave requires additional packages to be installed for **mp3** file support. To open and play files, **Media Player** must be installed.

When saving an **mp3** file, an MPEG audio encoder must be installed. GoldWave currently supports BladeEnc and LAME encoders. One of these must be installed before you can save **mp3** files. Note that you can use File | Save as to save the **mp3** file using a different file type.

For the download locations of these packages, please see the GoldWave website:

<http://www.goldwave.com/faq.html#install>

Ogg Files

To work with Ogg Vorbis files, GoldWave requires a copy of the CALOGG00.DLL module in the **Windows\System** or **WinNT\System32** folder.

For the download location of that module, please see the GoldWave website:

<http://www.goldwave.com/faq.html#install>

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: v4.21

GoldWave license/password: US\$40, Can\$55, UK£25

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

Website Links

Main site:

<http://www.goldwave.com>

Ordering information:

<http://www.goldwave.com/register.html>

Frequently asked questions:

<http://www.goldwave.com/faq.html>

Registering

GoldWave is a shareware program. It is not free software. However, after trying GoldWave, the modest registration fee should seem very reasonable considering all of the great features. To register and encourage further development, please see the [website](#). You can order securely online or you can fill in the [Order.doc](#) file and send it by regular mail. Registration details are given below. If you require more than one license, please send the full name of each person who will be using GoldWave.

Registration License (US\$40, Can\$55, UK£25)

You will receive a personalized password that removes the startup and toll messages and unlocks 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 website:

<http://www.goldwave.com>

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 (except online).

For information about the German version of GoldWave, please contact KELLY Media:

KELLY Media AG
Lilienthalstrasse 3
D-82178 Puchheim
Tel: +49-89/89 56 26-50
Fax: +49-89/89 56 26-99
<http://www.kellymedia.de>
<http://www.goldwave.de>

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

Order Now

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.

Getting Started:

- [Playing Sounds](#)
- [Recording Sounds](#)
- [Selecting Part of a Sound](#)

Please check the [Contents](#) for additional information.

The **readme.txt** file contains helpful information and important installation instructions. The **goldwave.htm** 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 website:

<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.

As shareware, you can give copies of GoldWave to anyone you think might find it useful. You can upload it to Web, BBS or ftp sites and post it to appropriate forums or newsgroups. Magazine and book publishers can include GoldWave on cover and companion CD-ROMs. Before copying, uploading, or posting, please click [here](#).

Only the original GoldWave zip file or self-extracting exe 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. Note that GoldWave is not free software and should not be described as such.

