

INTRODUCTION

INTRODUCING SOUND FORGE 4.5

Thank you for purchasing Sound Forge 4.5 and for your continued support of Sonic Foundry products. Sound Forge 4.5 brings new, powerful features to the Sound Forge product line such as built-in Batch Conversion, Spectrum Analysis and full support for creating Loops for Sonic Foundry's ACID.

If you are an avid Sound Forge user, you will find 4.5 to be a great leap forward. If you are new to Sound Forge, we are sure that you will find out what thousands of Sound Forge users already know – Sound Forge is the only solution for all of your audio editing needs.

READING THE MANUAL

The manual is separated into five major sections: Introduction, Using Sound Forge, Sampling, MIDI/SMPTE, Sound Forge Reference, Troubleshooting, and Appendices.

INTRODUCTION

This section includes installation and registration procedures, and a general overview of digital sound.

USING SOUND FORGE

This section demonstrates Sound Forge through pictures and examples, and provides suggestions for how to best use it.

SAMPLING

This section introduces samplers, the Sound Forge Sampler Tool, and provides tips for sampling.

MIDI/SMPTE

This section addresses MIDI, the Sonic Foundry Virtual MIDI Router, MIDI Triggers, and using MIDI/MTC synchronization.

SOUND FORGE REFERENCE

This section describes all menus and dialogs and can be quickly referred to when you have a specific question about a menu or dialog option.

TROUBLESHOOTING

This section contains the most common questions we have received from our users and includes some excellent tips.

APPENDICES

This section covers a variety of topics and can be used for general information, as well as reference.

GETTING STARTED

The first thing you should do is to register your copy of Sound Forge by filling out the enclosed registration card and returning it to Sonic Foundry within 30 days of purchase. Alternatively, you will be given the option to register online via our web site at the end of the installation process.

You must be registered for us to provide you with technical support and upgrade information. Also, we want to be able to let you know about additional products and updates as they are made available.

SYSTEM REQUIREMENTS

In order to use Sound Forge you will need:

- Intel Pentium or Alpha AXP microprocessor
- Microsoft Windows 95 or Windows NT 4.0 or later
- Windows-compatible sound card
- VGA display

- CD-ROM drive
- 16 MB RAM
- 5 MB hard-disk space for program installation

INSTALLATION

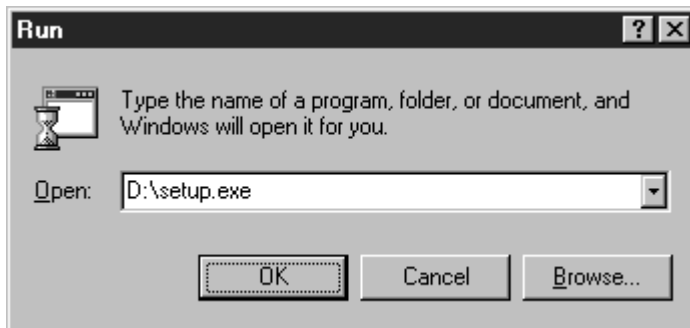
The install utility `SETUP.EXE` on the setup CD creates any necessary folders and copies all files required by Sound Forge.

NOTE

Sound Forge requires Microsoft DirectX Media 5.2a or higher to be installed on the computer. The setup program will notify you if it does not detect DirectX Media 5.2a or higher on your computer and will prompt you to install it from the Sound Forge CD-ROM.

To install Sound Forge, do the following:

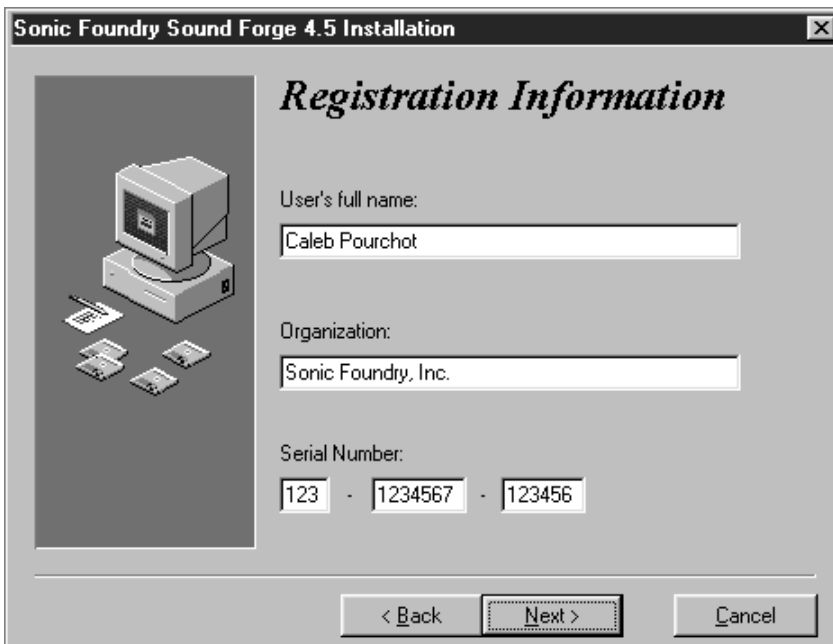
1. After placing the Sound Forge CD-ROM in the CD-ROM drive, AutoPlay will launch an installation menu. Press *Install* and follow the instructions.
2. If you have disabled the CD-ROM AutoPlay feature, run *Setup* from *Run* in the Windows Start menu by typing `E:\SETUP`, where *E* is the letter of your CD-ROM drive. Press *Install* and follow the instructions.



ACTIVATION

After you have selected a destination to install Sound Forge, you will be asked to enter your serial number. This number can be found on the inside of this manual on the registration card insert. You must have this serial number whenever you wish to run the setup program. Please do not tear out the entire insert if you are going to hand mail the registration card.

Sound Forge does not come with copy protection other than requiring your serial number whenever you run the setup program. We feel that it is in your best interest that the programs do not come with heavy-duty copy protection. We hope you will allow us to continue this policy by abiding by the license agreement and giving your friends our phone number rather than a copy of the software.



The screenshot shows a Windows-style dialog box titled "Sonic Foundry Sound Forge 4.5 Installation". The dialog box has a close button (X) in the top right corner. On the left side, there is a graphic of a computer monitor and keyboard. The main title of the dialog is "Registration Information". Below the title, there are three input fields: "User's full name:" with the text "Caleb Pourchot", "Organization:" with the text "Sonic Foundry, Inc.", and "Serial Number:" with three separate input boxes containing "123", "1234567", and "123456". At the bottom of the dialog, there are three buttons: "< Back", "Next >" (which is highlighted with a dashed border), and "Cancel".

ONLINE REGISTRATION

During the install process, you will be given the option to register online. If you choose to register online, the setup program will attempt to connect you to your Internet Service Provider and open the Online Registration page on our web site, www.sonicfoundry.com. Enter the user and survey information and press the Submit button at the bottom of the page.

If you choose not to register at the end of the setup process, you may register at any time using the Online Registration page on our web site, www.sonicfoundry.com/register.

TECHNICAL SUPPORT

If you experience problems while using Sound Forge, our technical support department is always ready to help you.

BEFORE YOU CALL

Since we have many users we ask that you do the following before calling:

- Make sure you have registered online or sent in your registration card.
- Check to see that the information you need is not already contained in the manual or available in the Help file.
- Have your serial number and the type of hardware and software you are using available to you.
- Use a phone that is near your computer so our support people can step you through any problems.

REACHING TECHNICAL SUPPORT

For technical support you can reach us via any of the methods below.

PHONE

(608) 256-5555 8–7 (CST) Monday through Friday

Call our technical support line between the hours of 8 AM and 7 PM Central Standard Time Monday through Friday.

FAX

(608) 256-7300

You can fax us your questions 24 hours a day.

INTERNET

support@sonicfoundry.com

You may also send us email via the Internet.

WEB

<http://www.sonicfoundry.com>

Our web page at <http://www.sonicfoundry.com> is another resource for program updates, and technical information and support.

COMPUERVE

GO SONIC

We maintain a forum on CompuServe that provides support as well as tips and updates. To reach our forum, simply access CompuServe and type GO SONIC at any prompt.

OVERVIEW OF DIGITAL SOUND


If you are new to digital sound editing, it will be well worth your time to become familiar with some of the basic concepts. In this section we cover the most important fundamentals. However, we strongly recommend that you page through a book on digital audio and sound recording if you want to get the most out of Sound Forge's editing and digital signal processing features.

SOUND WAVES

You can think of air pressure as the density of air molecules. When an object vibrates or moves, it displaces air molecules causing a pressure change. This in turn, causes other air molecules to move. We don't hear air pressure changes caused by the weather. Instead, we hear air pressure differences that vary rapidly over time.

When you hear a sound, you are sensing changes in the air pressure around your eardrum. These vibrations are then picked up by your ears and converted to electrical signals that your brain interprets as sound. If we were to graph the

air pressure at your eardrum as a function of time while you were listening to a short sound, it might look like the waveform that follows.




LOUDNESS AND PITCH

When there is no sound wave, the air pressure is constant. This is perceived as silence. As the sound wave reaches your eardrum, the air pressure changes above and below the normal atmospheric pressure. The amount of change is perceived as the loudness of the sound. The loudness of a sound, called its amplitude, is usually measured as a fraction of a standard level, often in decibels (dB).

The rate at which the air pressure changes is perceived as the pitch. In scientific terms, this term corresponds to the frequency of the wave. The frequency is usually measured in Hertz (Hz), or cycles per second.

Sounds in nature are not as simple as the sine wave we graphed above. In reality, a sound would look something like the one drawn below. This irregular waveform does not have a periodic amplitude or frequency.



TIMBRE

Complex waveforms like the one shown above are constructed by combining a number of simple waveforms (like the one in the first drawing) of different amplitudes and frequencies. This is why we perceive both high and low pitched sounds at once when we hear most natural sounds.

The characteristic sound of a waveform (be it produced by a grand piano or a violin) is called its timbre. Timbre, also referred to as tone color, is said to be rich or full when there are many different frequencies in a sound. A sound from a sine wave is considered dull by most people since it has only one frequency.

The different frequencies in a sound, combined with the varying amplitudes of each frequency, make up the spectral content of a waveform. The spectral content, which you might say is the more scientific term for timbre, usually varies over time. Otherwise, the sound remains static and again sounds dull. The spectral characteristic of a waveform over time is the signature of a tone that allows you to describe it as string-like or horn-like.

ANALOG RECORDING AND PLAYBACK

Let's say you're recording with a microphone. As you hold the microphone up in the air and scream, the microphone converts the changes in air pressure into changes in electrical voltage. This is called an analog signal.

If you were to graph the changing voltage inside a microphone cord, it would look exactly like the graph of the air pressure going up and down. To record your scream, you would send the signal to a medium such as magnetic tape which can store a replicate of the analog signal.

To playback your recording, you need something to create the differences in air pressure that our ear interprets as sound, i.e. an audio speaker. Speakers operate by moving a cone from one position to another in a consistent manner. In order to move the cone either forward or backward the speaker must be driven by an electrical current. During playback, a tape or record player generates a current that is then fed to an amplifier. When connected to a speaker, the current moves the speaker in a way that reproduces the pressure changes sensed by the microphone during recording.

Until recently, sound was always recorded as an analog signal on magnetic tape or vinyl grooves. One problem with storing a signal in this form is that it is hard to accurately record the analog signal without adding noise. When you

make copies of your recording you have to convert it to an electrical signal and then re-record it, adding even more noise. Listen to a third generation cassette recording and you'll know what we're talking about. Also, editing with tape is a not an easy task, since you must always be fast-forwarding or rewinding to a section, splicing, etc. Tape-based editing is called linear editing.

DIGITAL RECORDING

With recent advances in computer technology, it has become efficient and economical to record sound waves using a process called digital sampling. In digital sampling, the analog signal of the sound wave is divided and stored as numbers that represent the amplitude of the wave over very small segments of time.

For a moment, let's take a look at another process that is very similar to the way our computer makes sound – the making of movies and television. Given a scene with a person walking, we can slow down the speed at which the pictures are shown and see that each movement is captured by a different picture. As we speed up the rate of the pictures passing by, the motion becomes more fluid and eventually we stop noticing each individual picture. If we keep speeding up the movie the person appears to walk faster and faster and eventually ends up looking quite humorous. The important point is that a movie is just a collection of individual pictures.

So how do people record a movie? Obviously they pick up a movie camera and film the action. The movie camera takes a series of pictures at a fast rate and saves them on the film. The movies we see are simply a collection of pictures that are played back in rapid succession.

When we record a sound through the sound card in our computer (also called digitizing or sampling) we do much the same thing as the movie camera. The computer rapidly checks what position the microphone is in and saves it in the computer. When done recording, the computer has a collection of individual positions (normally called samples) which it can use to recreate the sound we have recorded. Individually, the samples are almost meaningless, much like an individual picture in a movie, but together they make up the recorded sound. This method of recording and playing sound is known as digitized sound.

THE PC SPEAKER

A simple model for a digital system is one where a speaker cone can be in one of two positions, either in or out, corresponding to the numbers 1 and 0 stored in a computer. The normal position, in, is when the speaker is sitting at rest with no current applied. The speaker isn't moving and we don't hear anything. When current is applied the speaker cone moves to the other position, out. As the speaker moves to the out position it forces the air around it to move and we hear a small click or pop. If we leave the speaker sitting in the out position we again hear nothing since the speaker only produces sound when it is moving, not when it is stationary.

Now if we move back and forth between the in and out positions we will hear a tone. As we move it faster and faster between the two positions (increasing the frequency) we will hear the tone increase in pitch. This model of a simple digital speaker system is exactly how the speaker inside your PC operates. Programs can put the speaker in either the in or out position to make sounds. Although there are more complicated methods to allow this type of system to produce sounds other than tones, they are beyond the scope of this introductory text.

EXPANDING THE MODEL

Instead of using just the two-position model with the in and out positions, let's suppose we had a system that allowed us to have 100 positions or even more. If this were the case then we would be able to make much more complex sounds. We could move it just a little bit or we could move it all the way out. This would allow us to have more precise control of the amplitude of the waveform. The more positions we have, the more flexibility we have in producing sound. For example, if we were to represent amplitude as a number from one to four, any values that fell between would be rounded to the closest value. This rounding error is called quantization noise. When more positions are available, rounding errors become smaller.

You will often see a sound card referred to as 8-bit, or 16-bit. We can directly relate this to the number of positions in which we can place the speaker. With an 8-bit card we can place it in 256 different positions and with a 16-bit card we can place it in any of 65,536 positions. Although you might think that a 16-bit card should have twice as many positions as an 8-bit card, this is not

the case. It actually has 256 times as many positions. Even though 16-bit samples take up twice as much space as 8-bit samples, it is recommended that when at all possible you use 16-bit samples to minimize quantization noise.

SAMPLING RATE

The number of times a sound waveform is checked for position each second is the sampling rate. The sampling rate is similar to the frame rate in movies. As you can imagine, with higher sampling rates you store more information about the sound's changing amplitude. This gives you more fidelity.

As a matter of fact, it is impossible to accurately record frequencies above one-half of the sampling rate. This threshold frequency is called the Nyquist Frequency, and should be considered when selecting a sampling rate. Frequencies higher than the Nyquist Frequency show up as alias noise.

The downside to very high sampling rates is that since each sample takes up space in memory (1 byte for 8-bit samples, 2 bytes for 16-bit samples), higher sampling rates will fill up your hard drive faster than lower sampling rates. For instance, a stereo digitized sound of 44,100 Hz 16-bit data (approximately what your CD player uses) lasting 10 seconds takes up almost 2 megabytes of space! This means if you have a 40 megabyte hard drive you couldn't even store 4 minutes of sound data, and that's without having any programs or other data on your system.

ADVANTAGES OF DIGITAL EDITING

The advantages of digital editing far outweigh the enormous storage requirements. Once you've recorded a sound as a digital sample on your hard drive, you have the ability to perform edits like copying, cutting, and pasting without losing any fidelity and, as some people like to brag, with accuracies of up to 0.000023 seconds (single sample spacing at 44,100 Hz sampling rate). With a visual editor such as Sound Forge, you can actually see a representation of the waveform to navigate through the sound file quickly and accurately.

Another advantage of storing sound digitally is the availability of digital signal processing (DSP) techniques. Digital signal processing techniques can be used for filtering, simulating room acoustics, and other special effects to restore or enhance the original recorded sound.

Finally, with a tool like Sound Forge, you can open and save your sound files to and from a number of different computer platforms, sound cards, and external samplers.

DIGITAL LEVELS

When recording to an analog medium such as magnetic tape, recording engineers always try to keep their meters as close to 0 VU (stands for Volume Unit, which is based on electrical currents) as possible. This ensures a high signal-to-noise ratio while preserving enough headroom to keep the tape from saturating and distorting. Recording a few peaks that go above 0 usually doesn't cause any problems since the tape saturation point is not an absolute.

In the digital realm, where amplitudes are stored as discrete numbers instead of continuous variables, things are quite different. Instead of having a flexible and forgiving recording ceiling, we have absolute maximum amplitudes, -32,768 and 32,767, in 16-bit audio. No stored signal can ever have a value above these numbers. Everything beyond gets clamped to these values, literally clipping off the wave peaks. This chopping effect can add large amounts of audible distortion. If the clipping is very short and infrequent such as during a very loud snare hit, it can go unnoticed. But in general, it is safe to say that digital audio has absolutely no headroom.

At what level, then, should a signal be recorded digitally? The standard method for digital metering is to use the maximum possible sample amplitude as a reference point. This value (32,768) is referred to as 0 decibels, or 0 dB. Decibels are used to represent fractions logarithmically. In this case, the fraction is: sample amplitude divided by the maximum possible amplitude. The actual equation used to convert to decibels is: $\text{dB} = 20 \log (\text{amplitude}/32,768)$

Say you have a sine wave with a peak amplitude of 50% of full scale. Plugging the numbers in gives you $20 \log (0.50) = -6.0 \text{ dB}$. In fact, every time you divide a signal's amplitude by two, you subtract its dB value by 6 dB. Likewise, doubling the amplitude of a signal increases its dB value by 6 dB. If you kept dividing your sine wave until its peak amplitude was equal to 1, you'd get the very lowest peak dB possible, -90.3 dB.

Why do we use dBs? We'll for one, it's easier to say -90 dB than 0.000030 (1/32,768). Decibels have been used for a very long time when dealing with sound pressure levels because of the huge range (about 120 dB) that the human ear can perceive. One confusing thing about using decibels is that 0% is referred to as minus infinity (-Inf. throughout this manual and in Sound Forge dialogs).

How do we measure the levels of a digital signal? Digital meters usually show the maximum instantaneous amplitudes in dB. This is called a peak meter. Peak meters are excellent for making sure that a recorded signal is never clipped. However, peak meters aren't as precise as using RMS (Root mean square...another mathematical formula) power readings when trying to measure loudness. This can be appreciated by generating a sine wave and a square wave with the same peak amplitudes and noting the square wave is much louder. When using RMS power, a maximum-amplitude square wave will be 0 dB (by definition), while a maximum amplitude sine wave reaches only -3 dB.

Now, let's get back to the real question – at what level should audio be digitized? If you know what the very loudest section of the audio is in advance, you can set your record levels so that the peak is as close to 0 dB as possible and you'll have maximized the dynamic range of the digital medium. However, in most cases you don't know in advance what the loudest level will be, so you should give yourself at least 3 to 6 dB of headroom for unexpected peaks (more when recording your easily over-excited drummer friend).

Now get in there and have some fun with Sound Forge.

This page was intentionally left blank.

USING SOUND FORGE

This part of the manual will help you quickly find your way around Sound Forge by giving you step-by-step instructions on how to use the most important features of the program.

Before you start your first edit session, make sure you have installed the sound drivers that Windows requires for your sound card. For more information on sound drivers refer to the Windows User's Guide or your sound card documentation.

Sound Forge was developed with input from musicians, sound editors, multimedia designers, game designers, studio engineers, and various other personnel from music and sound industries. You will find we have included features that are specific to these different applications. Just because a feature is intended for a specific use doesn't mean it can't be used in some other creative way. For instance, although the loop-tuning features of Sound Forge are intended for use by people creating sustaining loops for samplers, many game developers use this feature for creating background ambiance for their games.

Sound Forge is a deep program. There are many features that you may not discover immediately and others that may seem overwhelming. Don't let this scare you. Remember, there's always undo! The best advice for using Sound Forge is to jump in and start using it. Pick up the manual now and then, read through a section to learn a few new tricks, and then go back and try what you've learned.

MAIN SCREEN BASICS

This section will introduce you to the Sound Forge screen and cover the most basic operations including:

- Using the Mouse
- The Main Screen
- Toolbars
- The Data Window
- ToolTips
- Getting Online Help

USING THE MOUSE

Although using the mouse in Sound Forge is not required, it will make your editing sessions easier. Once you become familiar with Sound Forge you will probably want to use some of the built-in shortcuts provided by the mouse and the keyboard. The following list of terms will help you when reading the manual.

Pointing

Moving the mouse pointer over an item is called pointing.

Clicking

Pointing to an item and quickly pressing and releasing the left or right mouse button is called clicking. The manual will let you know whether you need to right-click or left-click on an item to execute a specific function. When there is no left or right specification you can assume we mean left-clicking. Right-clicking is often used to reach shortcut menus.

Double-clicking

Double-clicking is the same as clicking except instead of pressing and releasing the mouse button once, you do it twice in quick succession. Double-clicking always indicates clicking twice with the left mouse button.

Triple-clicking

Triple-clicking is the same as clicking except instead of pressing and releasing the mouse button once you do it three times in quick succession. Triple-clicking always indicates clicking three times with the left mouse button.

Toggle-clicking


Toggle-clicking involves clicking the right mouse button while holding the left mouse button to toggle through various options. This is a great shortcut for procedures such as drag and drop and changing the *Magnify Tool* mode of operation.

- Shift-clicking** Shift-clicking refers to holding down the Shift key on the keyboard while clicking the mouse. Shift-clicking is used mainly to skip dialogs so that you can quickly repeat operations.
- Control-clicking** Control-clicking is holding down the Control key on the keyboard while clicking the mouse. Control-clicking is used to modify the operation of a normal click.
- Dragging** Holding down the mouse button while you move the mouse pointer is known as dragging. Dragging is used to quickly move sections of data between separate windows and to move trackbars, scrollbars, and faders.
- Slow-dragging** Holding down both the right and left mouse buttons while moving trackbars and faders increases the resolution of the movement. This is especially useful when making fractional adjustments to parameters in a dialog box.
- Dropping** After dragging an item, releasing the mouse button on top of another area is known as dropping. Dragging and dropping is used to speed up operations like mixing or moving regions within the Playlist.

THE MAIN SCREEN

When you start Sound Forge, you see the main screen, or workspace, where you will do all of your editing. The very first time you open Sound Forge, no data windows are open and you will need to either open an existing sound file or create a new one. There are a variety of ways to do this, all of which are explained later in the manual.

The important parts of the Sound Forge screen are labeled for you in the diagram below. The following list briefly describes how each one functions.



Program Title Bar

Shows the program name and the name of the currently active data window (if the data window is maximized within the workspace).

Data Window

Each opened sound file has its own data window. Data windows can be arranged, resized, or minimized.

Menu Bar

Shows the menu headings for the available functions. When no data windows are open, the Process, Effects, and Tools menus are not listed because these contain functions that require an open data window.

Status Bar

On the left, help and processing information is displayed. The fields on the right show the playback sample rate, sample size, mono/stereo, total length of the active data window, and the total free storage space. These fields can be edited by double-clicking or right-clicking on them (except for the free storage space, of course). When no data windows are open, the fields are blank.

Sound Forge Workspace

This is the background area behind the data windows. You can drag sections here to create new data windows.

TOOLBARS

The first time you run Sound Forge, two toolbars also appear on the screen: the *Standard* and *Transport* toolbars. The toolbars contain buttons that are used to quickly execute commands.

Standard

Provides quick access to many commonly used Sound Forge File and Edit menus.

Transport

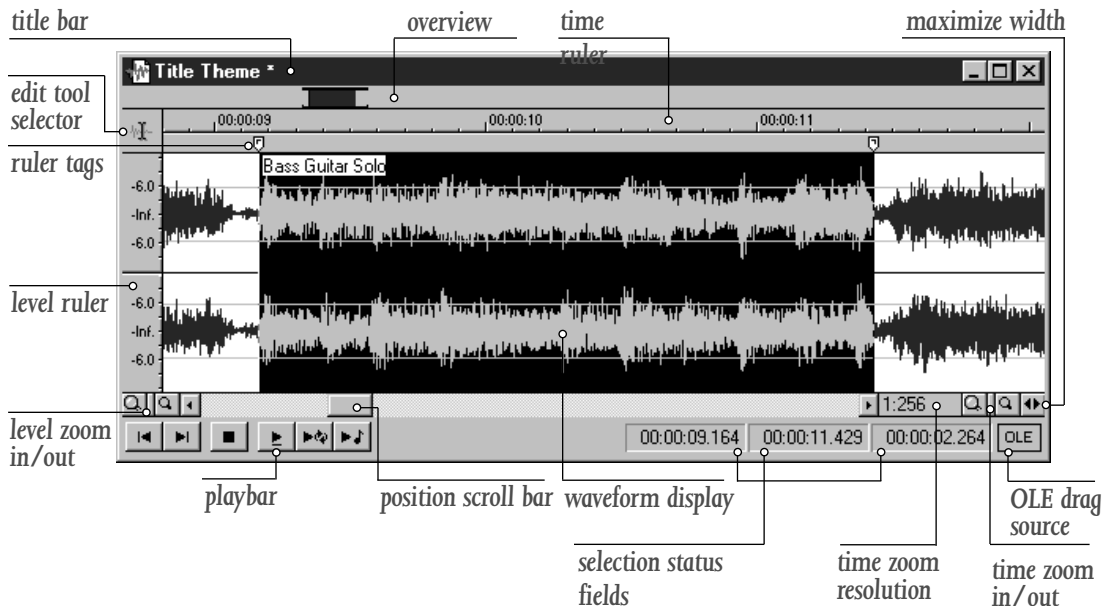
Provides the audio transport buttons: *Record*, *Play All*, *Play*, *Pause*, *Stop*, *Go to Start*, *Rewind*, *Forward*, and *Go to End*.

These toolbars, like any other toolbar in Sound Forge, can be dragged and dropped to anywhere on the screen. Toolbars are also resizable and removable. When a toolbar is dragged to any side of the main window, it docks, or attaches to the side. If you drag a toolbar away from a side, it becomes a floating toolbar. To hide a floating toolbar, just click on its *Close* button.

A list of available toolbars can be displayed by selecting the *Preferences* command under the *Options* menu and looking in the *Toolbars* folder (or *Toolbars* in the *View* menu). To show a toolbar, just check the box next to the toolbar you wish to use and then select *OK*. The number of open toolbars and their position on the screen is entirely up to you.

THE DATA WINDOW

Data windows contain sound file data. These windows contain a number of sub-windows and controls that you will use in editing and viewing your sound data. Many parts of the data window lead to shortcut menus, dialogs, and other operations.



The important components of the data window are labeled for you in the diagram above. The following list briefly describes each part and possible mouse operations:

Title Bar

Shows the file name or, if present, the title stored in the **Summary** information of .WAV files. Double-click to maximize and restore the window.

Level Ruler

Shows the amplitude of the waveform. Right-click to reach the **Level Ruler** shortcut menu. Drag to shift the view up or down when zoomed in vertically.

Time Ruler	Shows the current location in the data window as well as ruler tags. Right-click to reach the Time Ruler shortcut menu. Drag to scroll the data window.
Ruler Tags	Used to indicate the position of region end points, loop end points, and markers. Right-click to reach the Ruler Tag shortcut menu. Drag to change their locations. Double-click anywhere in the region to select.
Edit Tool Selector	Use this button to toggle between the Edit , Magnify , and Pencil Tools .


The Pencil Tool requires that the Zoom Ratio be between 1:1 and 1:16.

NOTE

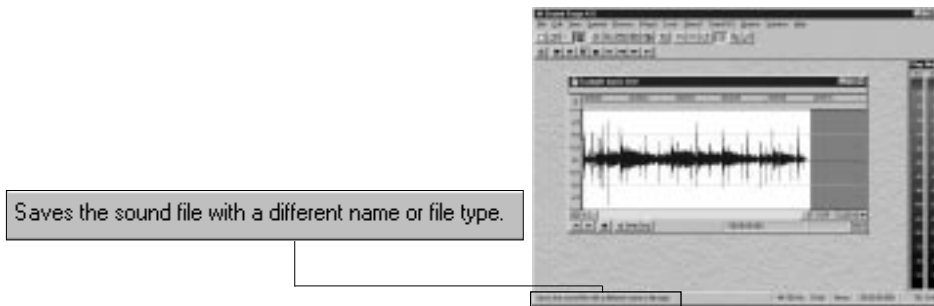
Playbar	On the left side of the playbar are the following audio transport buttons: Go to Start , Go to End , Stop , Play Normal , Play Looped , and Play as Sample . To the right are the playbar's selection status fields.
Selection Status Fields	Show the beginning, end, and length of a selection. If no selection has been made, only the cursor position (also known as the insertion point) is displayed. Double-click on the left-most field to reach the Go To dialog. Double-click on either of the other two fields to reach the Set Selection dialog. Right-click to reach the Status Format shortcut menu.
Waveform Display	Shows a graphical representation of the sound file. The horizontal axis represents time (marked in the Time Ruler) and the vertical axis represents amplitude (marked in the Level Ruler). Right-click anywhere on the waveform display to reach the Waveform display shortcut menu.

Position Scroll Bar	Use this to scroll the sound file forward and backward in time to see parts of the file not currently visible in the waveform display.
Overview	Allows for quick navigation and playback of any part of the file. Overview also shows the fraction of the waveform being shown on the waveform display, as well as the selected region. Left-click to move the cursor. Double-click to center the cursor in the waveform display. Right-click to start playback or pause. Left-click and drag to activate the <i>Audio Event Locator</i> .
Time Zoom Resolution	(Also called <i>Zoom Ratio</i>) Specifies the number of samples of data represented by each point on the screen horizontally. This determines the length of time which is shown in the waveform display. With a small resolution value (1:1, 1:2, 1:4, ...), a shorter length of time is displayed.
Time Zoom In/Out	Use this to change the zoom resolution for the time axis. Between these buttons is a spinner that allows you to continuously change the time zoom resolution.
Level Zoom In/Out	Use this to change the zoom resolution for the level (vertical) axis. Between these buttons is a spinner that allows you to continuously change the level zoom resolution.
Maximize Width	Pressing this button stretches the width of the data window to fit within the Sound Forge workspace.
OLE Drag Source	Drag from this section of the data window into another application such as Microsoft <i>Word</i> , to link the sound file to other documents.

TOOLTIPS



If you keep the mouse over a toolbar button or status bar field for more than one second, you will notice a small box appear next to the mouse pointer containing some text. This text is a quick description of the function of the button or field. Using *ToolTips* is an easy way to learn your way around the Sound Forge screen. If you wish, you can disable this function by selecting *Toolbars* under the *View* menu and un-checking *Show ToolTips*.



When you press and hold a button in a toolbar or select a menu item, you will see a one line description of the command in the left side of the status bar. This gives you a little more descriptive information about the function of a button. If you release the mouse button outside of a toolbar or menu item, the toolbar command will not be performed.

GETTING ONLINE HELP

Online help is available in a variety of ways. You can:

- Go straight to the *Help* table of contents by selecting *Contents* under the *Help* menu, or pressing *F1*. Once there, you can search for your help topic and also get information on using the online help.
- Press the *Help* button or the *F1* function key in a dialog. This will take you straight to the help information on the dialog.

- Press the F1 function key while selecting a menu item to get information on the command.

You will find that online help is extremely useful as a quick reference guide. Since many topics are cross linked, it's easy to navigate through all of the help material available on a topic. Once you have become familiar with using Sound Forge, you'll probably never have to refer to this manual to find information. All you'll need is the online help.

DATA WINDOW BASICS

With Sound Forge, you can have multiple sound files open simultaneously on the screen. Each file has its own corresponding data window that shows you a graphical representation of the waveform and other information about the file. This section covers the following topics:

- Opening an Existing File
- Playing a File
- Playing a Section
- Using the Transport Toolbar vs. the Data Window Playbar
- Creating a New Window
- Active Windows vs. Inactive Windows
- Copying Data to a New File
- Saving a File

OPENING AN EXISTING SOUND FILE

To open an existing sound file, do the following:

1. Select the **Open** command from the File menu. Sound Forge displays the **Open** dialog.
2. Use the **Open** dialog to select and open a file. You can preview .WAV files by pressing the **Play** button in the **Open** dialog. Pressing the **Play** button once more stops the playback. You can also have the files play when selected by checking the **Auto play** box.

You can use the **Files of type** list to select different types of data files. If Sound Forge recognizes the type of a file when it is highlighted in the files list, it will display its information in the bottom half of the dialog.

Open the file TUTOR1.WAV now. This file is located in the directory where you installed Sound Forge.

You will now see the data window containing TUTOR1.WAV. This file is a recording of someone saying "Wow, sound editing just gets easier and easier."

Notice how even though the file name is TUTOR1.WAV, the name of the data window is "Wow, sound editing.." This is because we have stored a title for the file within the .WAV file itself. Sound Forge allows you to embed descriptive titles as well as copyright information and other text fields for any .WAV file.

NOTE

Sound files in Sound Forge can have a title that is different from the file name. To edit the title, use the Summary page of the Properties command dialog (File menu). You can quickly get to the Summary dialog by pressing Alt+Enter.

PLAYING A FILE

Once you have opened a file, you can hear it by clicking on the **Play All** button on the **Transport** toolbar.

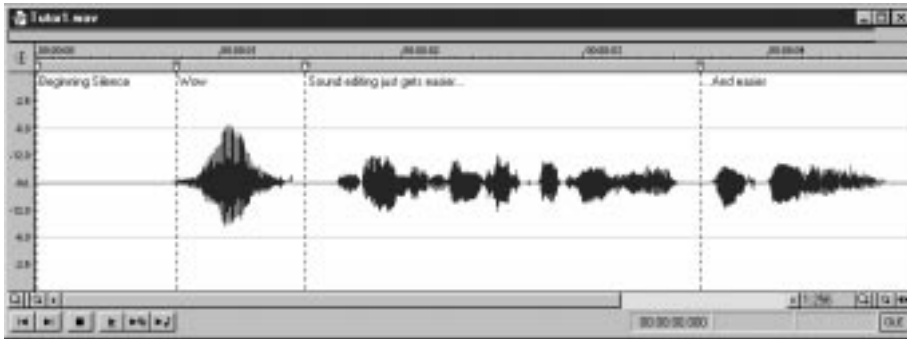


While the file is playing, a pointer will move along the data window showing the current position. You will also see the current position in the first selection status field in the playbar.

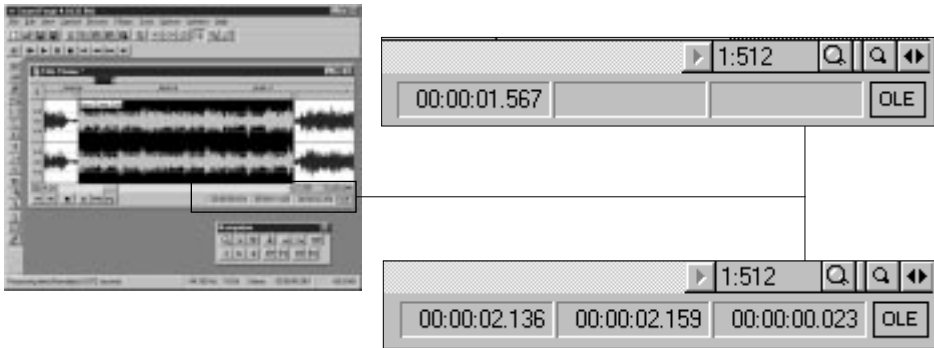
Play the TUTOR1 file now. You will hear the words "Wow, sound editing just gets easier and easier."

You can play from any point in the window by simply moving the cursor. Move your mouse pointer to the silence section right after the "Wow" portion of the file. Left-click (don't drag, we don't want a selection yet), and you will see the flashing cursor in the silence area after the word "Wow." Click the **Play** button on the **Transport** toolbar and now you will only hear the "sound editing just gets easier and easier" part.





If you don't hear anything you may have made a small selection by accidentally dragging the mouse. To tell if you've made a selection just look at the three right-hand status fields in the data window status bar. If only the first one is filled, then you don't have a selection. If these fields all have values then you need simply click in the data window again to clear the selection.



PLAYING A SECTION

You can also play portions of the sound data by making a selection on the waveform display with the mouse. To do this, click and drag the left mouse button starting at the section of silence prior to "Wow." Drag the mouse until you are in the portion after the "Wow." Notice that as you drag the data, the background appears in another color. Let the mouse button up and you will now have a highlighted section of data. Again click the **Play** button. You should now hear just the "Wow" portion of the file.

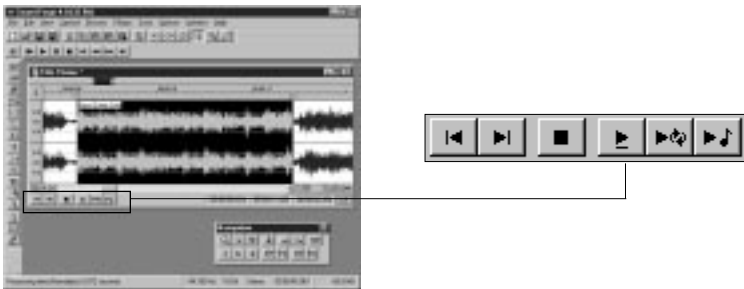
While making a selection, and after selecting, you will notice that the three selection status fields on the right side of the playbar have values. These values show you the start (*Selection Left*), end (*Selection Right*), and length (*Selection Length*) of the selection you have made. When no selection has been made, the cursor position (also called the insertion point) is shown.

Let's clear the selection and put the cursor back at the start of the file. To do this click on the *Go to Start* button in the *Transport* toolbar (the fourth button from the left).



USING THE TRANSPORT TOOLBAR VS. THE DATA WINDOW PLAYBAR

There are a number of ways to play your sound files within Sound Forge. The most common method is to use one of the *Play* buttons located either on the *Transport* toolbar or the playbar in the data window. Let's take a look at both of them.



THE PLAYBAR BUTTONS

On the left side of the playbar at the bottom of every data window are six buttons. These buttons allow you to play the sound data in a variety of ways. The first two buttons allow you to set the cursor to the start or end of the current file. These are the *Go to Start* and *Go to End* buttons. The next button is the *Stop* button which you can use at any time to stop playback.



The last three buttons play the sound in three different ways and also set the default play mode at the same time. The first button is the **Play Normal** button. This button plays the currently selected section of data. If there is no selection, pressing the **Play Normal** button plays from the current cursor position to the end of the file.



The next button is the **Play Looped** button. This button plays the selected section in a continuous loop or, if no selection is made, the entire file in a continuous loop. In addition, while the **Play Looped** button is selected, playback will follow any new cursor position or selection that is made.



The final button is the **Play as Sample** button. This button is used to play a sound file as if it was present in a sampler. This means the file will play from the start until it reaches the sustaining loop defined for the file and then will loop in the sustaining loop for the number of loops defined. If there are no loops defined the file plays straight through. This is known as a one shot sample.

Whenever you play a file from the playbar, you will notice a small line that appears underneath the play arrow in the corresponding **Play** button. This is the current play mode which will be used whenever you select the **Play** button in the **Transport** toolbar (or hit the **Spacebar**). **Play Normal** is the current play mode in the playbar graphic shown above.








To change the current play mode, press the corresponding playbar button or use **Control+Spacebar**.

THE TRANSPORT TOOLBAR

The **Transport** toolbar has nine buttons.



The **Transport** toolbar buttons are described below from left to right:

Record	Brings up the Record dialog which allows you to record to either a new or existing window. <i>Shortcut: Control+R</i>	
Play All	Plays the entire sound file from beginning to end. This allows you to hear the entire sound even if you have a selection or the cursor positioned somewhere other than the start of the data. <i>Shortcut: Shift+Spacebar</i>	
Play	Plays the file using the current play mode. To set the play mode select one of the three play buttons on the playbar as described above. The last button used becomes the active play mode. <i>Shortcut: Spacebar</i>	
Pause	Works like the Stop button in that it stops play of the current file. However when selecting Pause the cursor is placed at the current play position rather than back to where it was when play was selected. <i>Shortcut: Enter during playback</i>	
Stop	Stops play and positions the cursor back to where it was prior to selecting play. <i>Shortcut: Spacebar during playback</i>	
Go to Start	Places the cursor at the beginning of the sound file. <i>Shortcut: Control+Home</i>	
Rewind	Shuttles the cursor backward in the sound file. <i>Shortcut: Control+Left Arrow</i>	

**Forward**

Shuttles the cursor forward in the sound file.
Shortcut: Control+Right Arrow


**Go to End**

Places the cursor at the end of the sound file.
Shortcut: Control+End

CREATING A NEW WINDOW

To create a new data window do the following:

1. Select the *New* command from the *File* menu. Sound Forge displays the *New Window* dialog.
2. Select the new data format in the dialog.



Since we are already working with the TUTOR1 file let's create a new window that has the same data format as TUTOR1.

In the *New Window* dialog set the *Sample size* to 16-bit, the *Channels* to *Mono*, and the *Sample rate* to 44,100. When you click the *OK* button a new window will appear titled *Sound2*. This is an empty data window into which we are going to place data from TUTOR1.

ACTIVE WINDOWS VS. INACTIVE WINDOWS

When you have multiple windows on the screen, only one window is considered the active window. This is the window on which you are currently working. Any operations you perform will only affect this window.

To make a window the active window click anywhere on the window with the left mouse button. The window's title bar will change to the color you have defined as the active window color in **Control Panel** in Windows.

COPYING DATA TO A NEW FILE

Make “Wow, sound editing..” the active window by clicking on its title bar. If you don't still have the word “Wow” selected, select it again as we did above. From the Edit menu, select the **Copy** command. This will copy the sound data for “Wow” onto the clipboard.

Now make Sound2 the active window by clicking on its title bar (if the title bar is covered by another window, you can always use the Window menu to activate a data window). Select the **Paste** command from the Edit menu and the “Wow” data will appear in the Sound2 window. If you press the **Play** button, you can hear how our new file sounds with just the word “Wow.”

There are faster ways to copy data from one window to another, or from an existing window to a new window, but we will cover those in a later section.

SAVING A FILE

To save a data window, you first need to make it the active window. Make sure the Sound2 window is active.

1. Select the **Save** command from the File menu. Since the Sound2 window is a new file, Sound Forge displays the **Save As** dialog. If the file had been opened or previously saved by Sound Forge, the file would be saved immediately.
2. Type a file name into the **File name** field in the **Save As** dialog and select the **OK** button. Let's save the new sound file, Sound2, as MYFIRST – Sound Forge will add the file extension, **.WAV**, automatically.

The **Save As** dialog also allows you to change file types, data format, and set summary information fields. This is covered in a later section.

SIMPLE EDITING AND NAVIGATION

This section will show you the basic editing operations of Sound Forge. You will learn about:

- Common Edit Operations
- Status Formats
- Magnification and Zooming

COMMON EDIT OPERATIONS

The edit operations used most often include: *Cut*, *Copy*, *Clear/Delete*, *Trim/Crop*, *Paste*, and *Paste Special (Mix, Crossfade)*. Most of these make use of the clipboard, which is a temporary storage area. The clipboard can also be used to move data from one window to another. The following list provides a brief description of each operation:

Cut	Deletes a selected portion of data and copies it onto the clipboard. <i>Shortcut: Control+X</i>
Copy	Copies a selected portion of data onto the clipboard. <i>Shortcut: Control+C</i>
Clear/Delete	Deletes a selected portion of data but doesn't copy it onto the clipboard. <i>Shortcut: Delete</i>
Trim/Crop	Deletes all data in a window except the selected section. <i>Shortcut: Control+T</i>
Paste	Inserts the contents of the clipboard into a data window at the current cursor position or replaces the current selection. <i>Shortcut: Control+V</i>
Mix	Mixes the contents of the clipboard with the current data in a window starting at the current cursor position or at the start of selection. <i>Shortcut: Control+M</i>
Crossfade	Crossfades the contents of the clipboard with the current data in a window starting at the current cursor position. <i>Shortcut: Control+F</i>

To show how these operations are used, we will once again be using the TUTOR1.WAV file. If it is not currently open please open the file now as described in the previous sections.

MAKING A SELECTION

If TUTOR1 is not the active window, activate it by clicking on its title bar. To make a selection, you must first make sure that you are using the *Edit Tool*. To get to the *Edit Tool*, look under the Edit menu for the *Tool* item (it should already be checked) and select *Edit*. Next, select the word "Wow" with the mouse by clicking and dragging like we did in the previous section. You can verify you have the right section of data by pressing the *Play* button to hear it.

COPYING DATA ONTO THE CLIPBOARD

Once you have a selection, you can use the *Copy* command from the Edit menu. This will copy the selected data onto the clipboard. You will see no change since the *Copy* command does not change the data, it only copies it to the clipboard.

PASTING DATA FROM THE CLIPBOARD

Now move the cursor to the beginning of the file by selecting the *Go to Start* button on the playbar. To insert the contents of the clipboard into the file, use the *Paste* command from the Edit menu. You should now see the data for the word "Wow" appear at the beginning of the window. Press the *Play* button just to make sure. You should hear "Wow Wow Sound editing just gets easier and easier."



You have just made your first edit in Sound Forge! There are much quicker ways to do what we have just done with Sound Forge, but we'll get to that in the Advanced Editing section.

PASTING DATA TO ANOTHER WINDOW

Data on the clipboard remains there until it is replaced by another operation that places data onto the clipboard. Therefore, you can continue pasting the data anywhere you want.

To demonstrate this, create a new window as described in the previous section by selecting the **New** option from the **File** menu. Now, select the **Paste** command from the **Edit** menu once more. You should now have a new data window with the data for the word “Wow.” You can also do this by selecting the **Paste to New** function in the **Paste Special** item of the **Edit** menu. This creates a new window and fills it with the clipboard contents in one easy step.

CUTTING DATA

To cut data, first you need to select a section of data you want to cut. Select one of the “Wow” words (you should have two if you have been following the examples) from TUTOR1.

Now select the **Cut** command from the **Edit** menu. This will remove the selected data and place it onto the clipboard.

UNDOING AN EDIT OPERATION

After any edit operation you can cancel it by selecting the **Undo** command from the **Edit** menu.

Let's undo the cut we just made. Select **Undo Cut** from the **Edit** menu and you will see the original two “Wow” words in the TUTOR1 window.

REDOING AN EDIT OPERATION

If you change your mind again and decide you really liked an edit you undid, simply select **Redo** from the **Edit** menu. Doing this re-performs the edit operation you last undid. In this case selecting **Redo** would re-cut the second “Wow” from the TUTOR1 window.

Let's redo the cut. Select **Redo Cut** from the **Edit** menu and you will see the second “Wow” has been cut from the file again.

NOTE

The **Undo/Redo History** window (**View** menu) also allows you to restore your sound file to previous states by undoing or redoing operations. For more information, see the *Reference* section.

TRIMMING/CROPPING

The **Trim/Crop** command in the Edit menu allows you to single out a section of data and cut everything else out of the window except that section. This is a handy feature since you can keep using the **Play** button to hear selections until you have just the right amount and then get rid of everything else.

By now you should be getting used to selecting data on the screen, so select the "Wow, sound editing just gets easier" section in the TUTOR1 window, but don't select the second "easier." Remember you can use the **Play** button to hear how the selection sounds at any time.

Once you have the selection, select the **Trim/Crop** item found in the Edit menu. After cropping you will have only "Wow, sound editing just gets easier" left in the window.

At this point let's close the "Wow, sound editing.." window. Either click on the close box of the TUTOR1 window, or select the **Close** command from the File menu. You will be asked whether you want to save the changes you have made to TUTOR1. Select the **No** button since we don't want to keep the changes we made to TUTOR1.

Also close any other windows you may have open, like MYFIRST.WAV.

MIXING

Mixing is a powerful and useful edit operation that you will use often. Mixing allows you to combine two sounds together into one window so you can create complex sound effects.

This time we are going to open two files, TUTOR1.WAV and TUTOR2.WAV. Open them now as we've shown you in the previous sections.

TUTOR2 is a file that contains the sound of a snare drum and crash cymbal. We are going to mix this sound with TUTOR1; the "Wow, sound editing.." window.

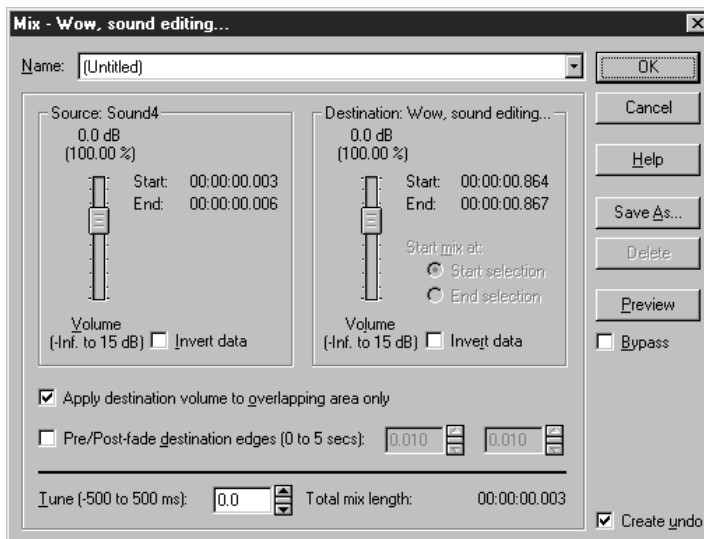
Before we begin mixing, let's make sure the **Status Format** is set to **Time** to make finding the mixing points a little easier. To do this select the item **Time** from **Status Format** under the Options menu. You will need to do this for both data windows since Sound Forge keeps track of the format type for each individual window.

To make the windows easier to view during the mix operations you may want to maximize Sound Forge by pushing the *Maximize Window* button in the upper right-hand corner and selecting the *Tile Vertically* item from the Window menu. This will arrange the windows vertically – fully utilizing the Sound Forge workspace – and make things easier to see.

If you activate each of the TUTOR windows you will notice that the length of TUTOR1 is 5.0 seconds long and TUTOR2 is 3.0 seconds long. For this example, we want the drum hit to occur just before the "Wow." We could just copy the drum hit sound on to the clipboard and then paste it before the "Wow", but this would increase the length of TUTOR1 to 8.0 seconds. So instead we'll use the *Mix* command.

Select all the data in *Drum Hit* by making it the active window and then double-clicking in the waveform display with your left mouse button (you could also use the *Select All* command under the Edit menu). Now copy the data onto the clipboard by selecting the *Copy* command from the Edit menu.

Make TUTOR1 the active window and then select the *Go to Start* button on the playbar. This will put the cursor at the beginning of TUTOR1. Select the *Mix* command from the *Paste Special* option in the Edit menu.



The *Mix* dialog now appears. Keep both levels at 0 dB and select *OK*.

You will see that the drum hit sound has been mixed into the TUTOR1 window and the length of TUTOR1 is still 5.0 seconds. Press the **Play** button to hear the results.

Select the **Undo Mix** command from the Edit menu to put TUTOR1 back to its original state.

Now let's mix the drum hit sound closer to the “Wow” portion of TUTOR1. The “Wow” occurs at about 0.8 seconds into TUTOR1, so move your cursor in the TUTOR1 window to approximately 0.8 seconds. You do this by clicking with the left mouse button in the data window of TUTOR1 and watching the cursor position status field on the status bar (the left-most field). You don't have to be exact. Once you've positioned the cursor, select the **Mix** command again and then hit **Play** for the result. Notice how the drum hit sound and the “Wow” sound overlap each other.

STATUS FORMATS

By right-clicking on the **Time Ruler** and selection status fields on the playbar, files can be set to different formats. You can coordinate sound files with other events, or edit to a timing base you feel most comfortable using.

Lengths and positions can be displayed in a variety of formats including: **Samples**, **Time**, **Seconds**, **Frames**, **Measures and Beats**, and **SMPTE**.

SELECTING A STATUS FORMAT

To select a format choose the **Status Format** item from the Options menu. This shows the ten different formats available. Choosing one of these options sets the status format for the current data window.

The available formats are:

Samples	Number of samples
Time	Hours, minutes, seconds and milliseconds
Seconds	Seconds and fractions thereof
Time & Frames	Hours, minutes, seconds and frames as defined by the Edit Frame Rate (Special Menu) command
Absolute Frames	Frames and fractions thereof
Measures & Beats	Measures, beats and 1/4's of a beat

SMPTE Non-Drop	SMPTE at 30 or 29.97 frames per second (fps) non-drop
SMPTE Drop	SMPTE at 30 fps with drop frames
SMPTE EBU	SMPTE at 25 fps
SMPTE Film Sync	SMPTE at 24 fps

Using the current file, TUTOR1, let's take a look at how the status formats affect values in the status display fields. First, select the *Samples* format from the *Status Format* menu under the Options menu.

Now select all of the data in the TUTOR1 window. To do this, choose the *Select All* option from the Edit menu. This will select all data in the window. You can also select all of the data in a window by triple-clicking the left mouse button anywhere in the waveform display of the data window.

In the selection status fields on the playbar (at the lower right-hand side of the data window) you should see the values of 0, 220,506, and 220,507. This means that the first selected sample (*Selection Left*) is sample 0, the last selected sample (*Selection Right*) is 220,506, and the total number of samples in the selection (*Selection Length*) is 220,507.

Now select the *Time* option from *Status Format* under the Options menu. You will see that all these values change to values specified in hours, minutes, and seconds rather than samples. You can see that a sound containing 220,507 samples with a sample rate of 44,100 Hz will play for 5.000 seconds. You can experiment with each of the status formats to see how each format is displayed.

CONFIGURING FRAMES AND MEASURES & BEATS

When setting the status format to *Frames* or *Measures & Beats*, there is additional information you can provide to Sound Forge to customize how these values are displayed. The *Edit Frame Rate* dialog in the Special menu allows you to change the frames per second. In the *Edit Tempo* dialog, also in the Special menu, you can specify the *Beats per minute* and *Beats per measure* values used to calculate measures and beats. The default values for *Frames* and *Beats* are set on the *Status* page in the *Preferences* folder (Options menu).

MAGNIFICATION AND ZOOMING

When sound data is drawn on the screen, it is necessary to represent many samples of data for each horizontal point on the screen. There are almost always many more samples in a sound file than there are horizontal points (pixels) on the screen. Depending on the editing operations you need to use, you may want to view the entire file at once, or you may want to look only at a small portion in greater detail. This is where the magnification, or zoom ratio, comes in.

TIME RULER ZOOMING

In the bottom right corner of the data window, just above the data window playbar, is the current time magnification ratio. The *Zoom Ratio* defines how many samples of sound data are squeezed into each horizontal point on the screen. This setting is shown as a value of 1:XX , where XX is the number of samples represented by each horizontal point on the screen.

For example, if the setting is 1:1 then each point on the screen is one sample. With this setting, a very short length of time is shown on the waveform display. You might not be able to see the cursor or selection because you are seeing it at such close range. If the setting is 1:1,024, then there are 1,024 samples represented by each point on the screen and a longer length of time can be seen. By dragging the *Time Ruler* right or left, you can easily scroll the data window. The further you drag in either direction, the faster the window will scroll.

If you right-click on the waveform display you can quickly select various Zoom commands from the Data Window shortcut menu.

NOTE

When Sound Forge opens a file it sets the horizontal magnification to the value specified by the *Normal zoom ratio* set on the *Display* page in the *Preferences* folder, unless the file is too small.



To increase or decrease the *Zoom Ratio*, use the *Zoom In/Out* spinner control found on each data window next to the ratio. Clicking on either of the magnifying glasses increases or decreases the *Zoom Ratio* by single steps. If you drag the small rectangular button between the two magnifying glasses, you can spin the *Zoom Ratio* by large amounts.

LEVEL RULER ZOOMING

In the bottom left corner of the data window, just above the *Go To Start* button, is the spinner control for zooming in and out vertically in the data window. Zooming in along the level axis allows for more precise editing at the lower levels (where many edits start and end).

The labels for this axis can be either in decibels or in percent and are switched by right-clicking on the ruler. When zoomed in, only the lower level samples can be seen; the peaks of the waveform will have moved beyond the scope of the data window. To view these peaks, drag the *Level Ruler* up or down.

NOTE

To optimize both the time and level scaling of a selection, double-click on the *Level Ruler*. Double-click again to zoom out to the normal level.

USING ZOOM SELECTION AND ZOOM OUT

If you would like to quickly magnify a section of a data window horizontally, the easiest way is to use the **Zoom Time | Selection** command found in the View menu or in the *Data Window* shortcut menu. To use the **Zoom Time | Selection** command you must first select an area which you want to magnify. After selecting **Zoom Time | Selection** Sound Forge will calculate the minimum zoom ratio which will allow the full selection to fit on the screen and then center the selection within the waveform display.

This will not affect the zoom ratio of the Level Axis. Only the Zoom Selection command in the Level Ruler shortcut menu will zoom the selection both vertically and horizontally.

NOTE

To display all data in a window at one time, use the **Zoom Time | Out Full** command. This command sets the **Zoom Ratio** to the lowest value needed to display all of the data at once in the waveform display.

If you wish to set the zoom factor to its maximum magnification (1:1), use the **Zoom Time | In Full** command also found in the View menu or the *Data Window* shortcut menu (**Zoom In Full**). You will then be able to see the actual waveform oscillations that correspond to physical sound waves when played through your speakers.

MAGNIFY TOOL

Another way to zoom in to a particular section of the sound file is to use the **Magnify Tool**. You can select this tool from the Edit menu (**Tool | Magnify**) or by clicking on the *Edit Tool* selector (upper left-hand corner of data window) until the magnifying glass appears.



When the *Magnify Tool* is enabled, instead of making a selection, the mouse cursor will become a magnifying glass and a dotted square will appear. The zoom will take place within this dotted square. The three different modes of operation for the *Magnify Tool* can be selected by toggle-clicking (See *Using the Mouse* above). These allow you to time zoom, level zoom or both simultaneously. You must then return to the *Edit Tool* mode to do edit operations.

You can also enter this *Magnify* mode temporarily by holding down the *Control* key when making a selection. If you drag the mouse while still holding the *Control* key, you can select a region that will then be magnified to a best fit ratio within the waveform display.

ADVANCED EDITING AND NAVIGATION

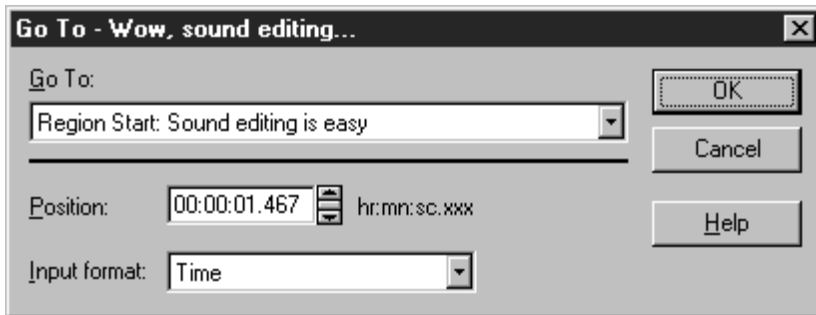
This section will show you how to use the advanced editing and navigation operations of Sound Forge. You will learn about:

- Go To
- Making a Selection Using the Set Selection Dialog
- Adjusting the Selection Points
- Using the Overview
- Using Views
- Drag and Drop Operations
- Editing Stereo Files

GO TO

To quickly move the cursor to a specific point and center it in the waveform display, use the *Go To* command. You can reach the *Go To* dialog in a number of different ways:

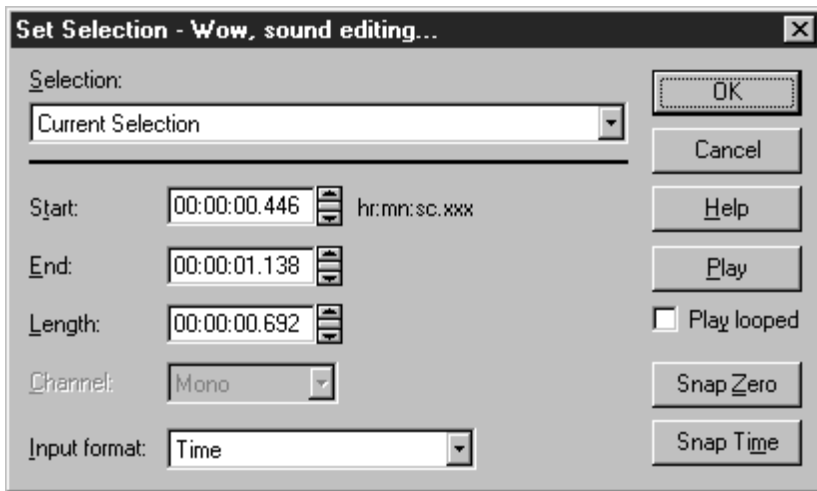
- Use the *Go To* item in the Edit menu
- Right-click on the waveform display and select *Cursor | Go To* from the shortcut menu
- Double-click on the left-most selection status field
- Press **Control+G**



The *Go To* dialog is much like the *Set Selection* dialog described on the following page. You can type in the location where you want to go or use one of the predetermined locations from the list. To create your own location marker, press the “M” key at any cursor position. You can also drop a marker at any time during playback and recording by pressing “M”. Refer to the *Regions List* and *Playlist* sections for more information on using markers.

MAKING A SELECTION USING THE SET SELECTION DIALOG

If you need to select data at specific points, the *Set Selection* dialog allows you to type in selection points or choose a selection from a list of regions. To reach the *Set Selection* dialog, choose the *Selection | Set* command from the Edit menu.



In the *Set Selection* dialog, you can modify the *Start*, *End*, *Length*, and *Channel* of the selection. Pre-determined regions can also be chosen from the *Selection* drop-down list.

Also, in this dialog you will find *Snap Zero* and *Snap Time* buttons. Pressing one of these buttons will force the current start and end selection points to their nearest zero-crossings or time increments, accordingly. A more thorough explanation of these functions can be found in the *Reference* chapter under *Selection* in the *Edit* menu section.

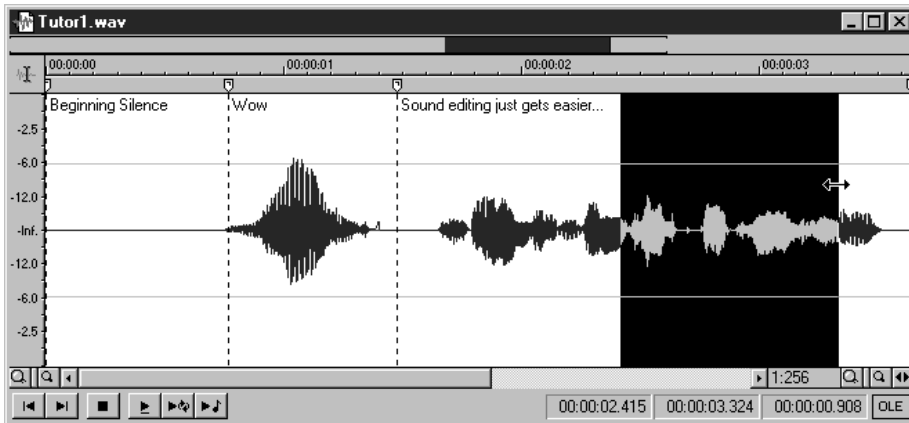
NOTE

Many commonly used functions in *Sound Forge* can be performed in a number of different ways. For example, you can reach the *Set Selection* dialog by double-clicking on the playbar field to the left of *OLE*, or you can pick the *Selection* item from the *Waveform Display* shortcut menu. *Sound Forge* has been designed to allow you to choose how you want to work. It doesn't force you to work in a certain way. To avoid confusion in our demonstrations, we won't always tell you about every method of accomplishing the same task. Refer to the *Sound Forge Reference* chapter and the *Appendix* section on shortcut keys for more information on each feature.

ADJUSTING THE SELECTION POINTS

EXTENDING A SELECTION WITH THE MOUSE

After selecting a section of sound data, you will sometimes find that the start or end points are not exactly where you want them to be. You could just re-select the data, but it is often difficult to get the start or end points just right.



Sound Forge allows you to update the selection by grabbing a selection edge and moving it. Once you have established a selection, place the mouse over the selection start or end point. You will notice that the cursor changes to a bi-directional arrow. Once this arrow appears, you can press the left mouse button and move the selection to a new position. The new position will be updated once you let up on the mouse button.

SNAPPING A SELECTION TO ZERO-CROSSINGS

If, after making a selection, the end points jump somewhere else automatically, it's probably because you have *Auto Snap to Zero* on. You may want to enable *Auto Snap to Zero* (under the Options menu) so that any drag selections start and end points will always reside on a zero-crossing.

To make a current selection snap to zero-crossings, select **Edit | Selection | Snap to Zero** (or press "Z" on the keyboard).

If you'd like *Auto Snap to Zero* disabled only when zoomed in (1:1 or 2:1), select the *Disable auto-snapping below 1:4 zoom ratios* option in the *Editing* page in the *Preferences* folder.

SELECTING WITH THE KEYBOARD

The keyboard selection controls allow you to quickly select data or update a selection accurately. For example, say you want to extend the selection end by a small amount. First make sure that the cursor is at the end of the selection. You can move the cursor from one end of the selection to another by pressing the *Home* (moves to the start) and *End* (moves to the end) keys. Use the “*C*” key to center the cursor in the waveform display. You can then use the *Shift+Right arrow* and *Shift+Left arrow* combinations to extend the selection end point by small increments.

The following list shows the available keyboard commands for selecting data.

Press	To select from cursor to...
<i>Shift+Right arrow</i>	the next screen pixel.
<i>Shift+Left arrow</i>	the previous screen pixel.
<i>Shift+End</i>	the last sample visible in the waveform display.
<i>Shift+Home</i>	the first sample visible in the waveform display.
Press	To select ...
<i>Shift+Page Up</i>	10% of the current view past the cursor position.
<i>Shift+Page Down</i>	10% of the current view previous to the cursor position.
<i>Shift+Control+Right arrow</i>	2% of the current view past the cursor position (see below).
<i>Shift+Control+Left Arrow</i>	2% of the current view previous to the cursor position (see below).
<i>Shift+Numpad +</i>	the next sample.
<i>Shift+Numpad -</i>	the previous sample.

Shift+Control+Numpad + 10 samples past the current position.

Shift+Control+Numpad - 10 samples previous to the current position.

In *Windows 95*, **Shift Numpad +** will not work unless the **Numlock indicator on the keyboard** is lit.

NOTE

Also, if there are regions, markers, or loops present in the data window:

Press **To select from cursor to...**

Shift+Control+Right arrow the next region start/end or marker.

Shift+Control+Left arrow the last region start/end or marker.

As you can see from the above lists, *Sound Forge* has extensive keyboard support for selecting data. You will find that almost any operation has an equivalent keyboard shortcut which advanced users find invaluable. For a complete listing see *Appendix A*.

Let's take a quick look at how to use some of these keys. Open the file *TUTOR1.WAV* as we've done in previous sections. First, make sure that you're using the *Edit Tool* by selecting it in *Tool* under the *Edit* menu. Now make a selection that encompasses the word "*Wow*" with a generous amount of space on each side of the word. Pressing the *Play* button you should hear a little bit of silence, "*Wow*", and a little more silence.

Choose the *Zoom Selection* command from either *Zoom Level* or *Zoom Time* in the *View* menu. This will fit the selection on the screen with the best possible resolution. Now let's adjust the right side of the selection.

Grab the end of the selection with the mouse and drag to the right. Adjust the selection to be close to the end of the word "*Wow*" and let up on the mouse button. Now let's do the same thing with the keyboard. Hold down the *Shift* key and press the *right arrow* key. The selection will extend to the right by one screen pixel. If you hold down the *Shift* key and press the *left arrow* key, the selection will decrease by one pixel. One screen pixel is equal to the number of samples shown in the *Zoom Ratio* field. You can move the end of the selection in this manor until you've got the selection just right.

Now let's adjust the start of the selection. Since the cursor is at the end of the selection, any keys we press for selecting will adjust the end. So we need to put the cursor at the front of the selection. To do this, press the **Home** key. You will see the cursor jump to the start point of the selection. Now, use the selection keys to adjust the start of the selection.

A selection can also be updated on a per sample basis rather than a pixel basis. If you look in the above list you will see that the **Shift + Numpad + (plus)** and **- (minus)** keys are used for these operations. This allows you to get the exact selection you need without having to change the **Zoom Ratio**.

RESTORING A SELECTION

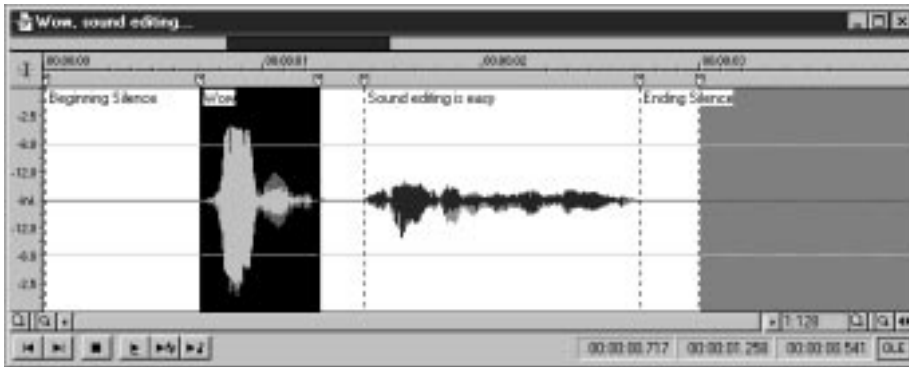
When you change the cursor position (with the mouse or keyboard) without holding the **Shift** key down, you lose your selection and only the cursor is displayed. If you want to return to the last selection, select the **Toggle Selection** item under the **Special** menu or press the "S" key.

THE OVERVIEW

While making selections and navigating through a sound file, you probably noticed the overview (the thin bar directly underneath the title bar) change. The overview represents the length of the entire sound file, as if you were zoomed out all the way. From the overview, you can determine what section of the entire sound file is being displayed, the selection made, and the cursor location.

GETTING THE WHOLE PICTURE

Zoom out all the way on the TUTOR1 window (**Zoom Out Full** in the **Data Window** shortcut menu) and make a selection over the word "Wow." Notice that the entire overview has brackets above and below it, since the entire sound file is being displayed. Also notice how the selection and the cursor are displayed on the overview.



Now press the *Zoom In* button a few times (the large magnifying glass in the lower right-hand corner of the data window) and notice how the brackets become smaller. This corresponds to the smaller fraction of the entire sound file that you can see when you zoom in on the waveform display. The selection size, however, remains the same and does not move. Even when you can't see them on the waveform display, you always know where the cursor and the selection are in the sound file by referring to the overview.

You can also change the cursor position by left-clicking anywhere on the overview. As always, when you change the cursor, you lose your selection. Remember you can use *Toggle Selection* if you want to restore it.

FAST NAVIGATION AND PLAYBACK

When you left-click outside the bracketed region in the overview, you will not be able to see the cursor on the waveform display. However, double-clicking anywhere on the overview will move the cursor and center the waveform display to the selected position in the sound file.

You can also play back the sound file starting from the current cursor position by right-clicking anywhere on the overview. Right-clicking again pauses playback. Left-clicking in the overview also moves the cursor position. Also note that left-clicking on the overview during playback moves the cursor to the point where you clicked and continues playback at that point. These navigation tools make it very easy to find sections in large files.


As an example, say you wanted to move the cursor to the beginning of a phrase in a speech. You can right-click on the overview bar to begin playback and then left-click at different positions within the overview until you find the right start point. Once you find the start point, you can press the **Stop** button to stop playback. The cursor should now be positioned at the last spot on which you left-clicked in the overview; select **Center Cursor** from the **Special** menu and you are ready to edit.

AUDIO EVENT LOCATOR

Holding down the left mouse button and dragging it in the overview bar will begin playback of small looped portions of audio around the cursor position. While this is not technically a scrub function, it serves a similar purpose i.e., it allows the user to hear brief segments of audio material in order to quickly scan or locate events within the sound file. The length of the **Audio Event Locator** loop segment can be adjusted in the **Preferences** dialog.

USING VIEWS

If you wish to save the selection points you've set, the zoom ratio, or the position of the waveform display, you can do so with **Views** Sound Forge has the ability to store up to eight views for each data window. Each view stores the selection, cursor position, magnification, and position scroll bar placement. A saved view can later be instantly restored.



You can store and recall up to eight different views (1-8) in this manner. There is also a **Views** toolbar you can use (look in the **Toolbars** folder of the **Preferences** dialog) to set and restore views. When using the **Views** toolbar, all stored views are indicated by an underscore in each **View** button.


Follow these steps that demonstrate how to use views:

1. Open the file TUTOR1.WAV and select “Wow.”
2. To store the current view, press **Control+1**.
3. Now, zoom in and make a different selection anywhere you wish. Store this new view by pressing **Control+2**.
4. To return to the first view you stored, press **1**. Notice that the selection and zoom factor are restored.
5. To go to the second view you created, press **2**.

DRAG AND DROP OPERATIONS

You can perform a mix, paste, or crossfade by dragging a selection from one data window and dropping it on another.

To drag and drop:

1. Open two sound files and select *Tile Vertically* from the Window menu just to make them easier to work with. One will be called the source window and the other, the destination window.
2. Press and hold the left mouse button in the selected area of the source data window. While holding the mouse button down, move the cursor until it changes to an arrow dragging a little box. 
3. Drag the cursor onto the destination window. You will see a shaded block that represents the range of the source material. You will also see that the little box accompanying the cursor now contains an “M”(Mix), “C”(Crossfade), or “P”(Paste). Position the block at the place in the destination window where you would like the mix, paste, or crossfade to occur.
4. Let go of the left mouse button. This will bring up the Mix dialog. To perform a crossfade instead of a mix, hold the **Control** key down while letting go of the mouse button. Holding down the **Alt** key when dropping performs a paste.

NOTE

An easier way to select whether you will perform a mix, paste, or crossfade is to simply drag the block to the destination window and then while holding the left mouse button, click the right mouse button to toggle between the Mix, Paste, or Crossfade functions. You will see the block change according to the method; a mix will look like a solid block, a crossfade will look like an “X” (or perhaps, a butterfly), and a paste will be represented by a segmented vertical line rather than a block.

As we did in the previous section on mixing, open the two files, TUTOR1.WAV and TUTOR2.WAV. For this example, TUTOR1.WAV will be the destination window and TUTOR2.WAV the source window.

To make the windows easier to view while doing the mix operation you may want to select the *Tile Vertically* option from the Window menu.

Highlight all the data in TUTOR2 by making it the active window and triple-clicking in the waveform display. Now drag the selected section by pressing and holding the left mouse button in the data window of the TUTOR2 window. Notice that the cursor changes to a pointer with a small box when you begin dragging. Drag the block to the beginning TUTOR1 data window and let go of the mouse button to drop the section.

When you let go of the mouse in the TUTOR1 window, you will see a dialog that shows the TUTOR2.WAV window as the source and TUTOR1.WAV as the destination. Leave the levels at 0 dB, and select *OK*. You will see that again the drum hit sound has been mixed into the TUTOR1 window. Press the *Play* button to hear the results.

DRAG AND DROP CROSSFADING

Undo the last operation. Now select all of TUTOR2 and drag it toward the end of TUTOR1. Press and hold the *Control* key, you will see the crossfade block and by moving the block back and forth horizontally across the end of TUTOR1, you can adjust the crossfade time and see the crossfade expand and contract over the end of the destination file. Once you have determined where you would like the crossfade to begin, let go of the mouse button. This takes you to the *Crossfade* dialog. Select the *Normal* crossfade preset from the preset

list at the top of the dialog and press **OK**. You should hear the end of TUTOR1 fade out as the beginning of TUTOR2 fades in.

DRAG AND DROP PASTING

Again, undo the last operation. Select all of TUTOR2 and drag it over to TUTOR1. Hold the **Alt** key down while you position the paste point at the beginning of TUTOR1. Now, let go of the mouse button. If everything was done correctly, you should have pasted the data from TUTOR2 at the insert position in TUTOR1.

CREATING A NEW WINDOW WITH DRAG AND DROP

A very useful feature of drag and drop is the ability to quickly create a new window from a selection. Make a selection in a data window and drop it on an empty area of the Sound Forge workspace. A new data window is automatically created containing the data from the selection.

SNAP TO EVENT IN DRAG AND DROP OPERATIONS

Another useful feature of drag and drop is the ability to snap to markers, regions, time increments and other events in the destination window. All of the drag and drop operations mentioned above will snap (or gravitate) to points in the destination file depending on what events have been established in the file.

The following list shows possible **Snap** to points in the destination file.

Cursor	Start of block snaps to cursor position
Selection	Start of block snaps to start or end points of a selection
Start	Start of block snaps to start of file
End	Start of block snaps to end of file
Markers	Start of block snaps to marker
Regions Start and End Markers	Start of block snaps to region start or end
Time, Measures, Etc.	Start of block snaps to labeled divisions on Time Ruler

Video Frames

Start of block snaps to the start of video frames appearing in the video strip

ADVANCED MIXING OPTIONS

In the *Mix* dialog, there are several options you can use to perform complex mixes. For example, when mixing vocals over music, you can have the music fade out and then fade back in when the vocal part is mixed. This is sometimes referred to as ducking. Follow these steps to perform voice-over ducking:

1. Open the file TUTOR1.WAV and select the word “Wow.”
2. Drag the “Wow” selection to an open area in the Sound Forge workspace to create a new data window.
3. In the *Summary* dialog within *Properties* (File menu), name this sound file “Wow” in the *Title* field and press *OK*.
4. Now, open the file TUTMUSIC.WAV. This file contains a short music clip to which we will mix the “Wow” sound.
5. Place the cursor in TUTMUSIC at about two seconds from the start of the file.
6. Select all of “Wow” and drag and drop it to the TUTMUSIC window.
7. In the *Mix* dialog, select the *Slow duck* preset. Notice that the *Destination Volume* is set to -6 dB and the times in the *Pre/Post-fade destination edges* are set to 0.5 seconds. This means that before and after the mix, the music will be faded to -6 over 0.5 seconds. Press *OK*.
8. Listen to the file, then undo the operation.
9. Once again, drag “Wow” to TUTMUSIC and select the *Slow duck* preset. This time, set the *Destination Volume* to -20 dB and press *OK*. The music will now fade out even lower. Both the source and destination volumes can be used to get the right mix of voice and music. The *Pre/Post-fade destination edges* control how fast the destination sound is faded out and back in.

EDITING STEREO FILES

When editing stereo files you have two channels of data on which to work. The upper channel is the left channel and the lower channel is the right channel. We will refer to them in both ways left (upper), and right (lower).

SELECTING DATA IN STEREO FILES

When selecting data in stereo files, Sound Forge allows you to select the left channel, right channel or both channels for playing, editing, and effects processing.

When editing a stereo file, the waveform display showing the two channels is split into three logical sections for selection with the mouse. The upper quarter of the waveform display is the left channel hit section, the lower quarter is the right channel hit section, and the middle half selects both channels. When selecting data with the mouse, the cursor location determines what channel(s) will be selected.

Open the file TUTOR1.WAV and convert it to stereo by editing the *Format* under *Properties* in the File menu or by right-clicking on the *Channels* playbar field and selecting *Stereo*. Set the *Channels* to *Stereo* and press *OK*. Select *Both Channels* for the destination in the *Mono to Stereo* dialog. You should now have a stereo version of TUTOR1.WAV.

Move the mouse pointer near the top of the left channel and select the word "Wow." Notice the change to the left channel selection cursor, and that only the left channel of the data becomes highlighted. Now do the same thing but in the middle half of the window near the top of the right channel and the bottom of the left channel. This time you should see both channels being selected. Do this one more time near the bottom of the waveform display and you should see only the right channel being selected.

Cursor when selecting both channels of a stereo file or when in a mono file.



Cursor when selecting the left channel of a stereo file.



Cursor when selecting the right channel of a stereo file.



Cursor when extending a selection.



TOGGLING CHANNEL SELECTIONS

Once you have made a selection in a stereo file you can switch between channel selections by pressing the **Tab** key. The **Tab** key will cycle between selections of *Left Channel*, *Right Channel*, and *Both Channels*. You can also set the channel selection by using the *Channel* drop-down list in the *Set Selection* dialog.

PREVIEWING CHANNELS

Selecting a single channel allows you to hear a preview of a single channel in the stereo file. For example double-click (triple-click if it contains regions or markers) in the TUTOR1 window to select all the data (or use *Select All* from the *Edit* menu). Press the **Play** button and listen to the clip. Next press the **Tab** key to toggle the channel selection into a single channel and press the **Play** button again. Do this one more time to hear the other channel.

SINGLE CHANNEL EDITING

Stereo data files are tied together by the nature of their format. In other words, they always play together. This means that there are some edit operations, such as *Cut* and *Paste*, which you can't use on a single channel. It would leave one channel shorter or longer than the other. This is usually not a problem in real world editing situations. To shift a single channel in time by small amounts, you can use the *Delay/Echo* function.

You can copy a selection from a single channel to the clipboard by selecting the data in either the left or right channel and using the *Copy* command. This will place a mono clip on the clipboard. You can then paste the mono clip to a mono file, both channels of a stereo file, or you can mix it into a single channel or both channels of a stereo file. When mixing mono clipboard data to a stereo file you will be asked with a dialog whether you wish to mix to a single channel or both channels.

WORKING WITH .AVI FILES

Sound Forge allows you to open and save *Microsoft Audio and Video Interleave* (.AVI) files. This allows you to edit the audio tracks of a video with single frame accuracy. This section covers:

- Opening an .AVI File
- .AVI Navigation
- Attaching a Video to a Sound File
- Saving an .AVI file
- .AVI Video Compression
- Editing Audio in .AVI Files

OPENING AN .AVI FILE

To open an .AVI file, choose *Open* from the File menu and select *Video for Windows (*.avi)* from the *Files of type* list. Then, double-click on the file you wish to open. If the .AVI file being opened in Sound Forge contains exactly one video stream and one audio stream, it will open these streams automatically. However, if no audio is present an audio stream will be created containing only silence for the entire length of the video.

HANDLING MULTIPLE STREAMS

.AVI files can contain more than one video or audio stream. A stream is not the same as a track in a multi-track video or audio editor. Multiple streams are most commonly used to contain different versions of a video or audio track. For example, you could have an .AVI file with one video stream and multiple audio streams in different languages. An .AVI player, such as Microsoft's *Media Player*, can detect which language version of Windows your computer is using and automatically play the correct audio stream for that language.

If more than one stream exists you have the option of choosing which one will be opened in Sound Forge. This is done in the *Video Stream* dialog. In this dialog, you must select a video and audio stream to open by using the *Stream Selectors* (black diamonds) to the left of each stream. Once you've selected one of each, press *OK*.

.AVI NAVIGATION

VIDEO STRIP

When you open an .AVI file, you will notice the video strip on top of the waveform display. When you're zoomed out, not all of the frames in an .AVI will be displayed. The left edge of the displayed frames are aligned with the audio in time. If you'd like to see the frame numbers on the frames, right-click on the video strip and select *Number Frames*.

NOTE

To resize the video strip, drag its lower edge up or down.

To enable frame animation in the video strip, select *Animate frames* from the *Video Strip* shortcut menu. When this option is enabled, moving the cursor will cause the nearest frame to be displayed above it. During playback, you will see the active frame window play the video as the cursor moves past it.

VIDEO PREVIEW

The *Video Preview* window is opened by selecting *Video Preview* from the *View* menu. The frame closest to the cursor (during editing or playback) will always be displayed in this window.

At the top of the window, the original frame size and current display size (in parenthesis) are displayed, along with the frame rate and current frame number. An asterisk (*) by a frame number indicates that the frame is a key frame (see the section on compression below).

ATTACHING A VIDEO TO A SOUND FILE

To attach a video to an opened sound file, select *Properties* from the *File Menu* and go to the *Video* folder. Pressing the *Attach* button will take you to the *Attach Video* dialog, where you can select an .AVI file to attach a video from. Press *OK* after selecting an .AVI with a video stream.

The *Video Properties* will now display the audio and video streams in the .AVI. An .AVI file can contain multiple streams, but in most cases you'll have one audio stream and one video stream.

The square check box to the left of each stream indicates which streams will be stored when you go to save your file. The black diamond next to it is the **Stream Selector**, which indicates which streams are currently being used in Sound Forge.

Press **OK** again and you should now see the video strip above the waveform display. Once you're done editing, you must save the file as an **.AVI** if you want to store the video.

SAVING AN .AVI FILE

To save an opened file in the **.AVI** file format, follow these steps:

1. In the **Save** or **Save As** dialog, select **Video for Windows (.avi)**. After naming the file, etc., press **Save**.
2. The **Video Save Options** dialog appears next. In here, you can select which streams to store in the **.AVI** file by checking the square check box to the left of each stream. In this dialog, it is also possible to edit the titles of the streams.
3. Pressing **OK** will take you to the **Compression Options** dialog. Here, you can select a compression scheme for each audio and video stream. To change the options for a stream, select it and press the **Options** button. By checking the **Interleave every _ Frames** option, you can specify how frequently the audio chunks are interwoven with the video.
4. If you pressed the **Options** button, yet another dialog, **Video Compression**, allows you to select from different compression algorithms, and even go to another dialog to set specific compression configuration parameters. For more information on compression, see the next section. Selecting **No Recompression** saves the stream in its original format, which is the Sound Forge default.
5. Once you're done setting compression options, hitting **OK** several times will complete the process. If the video is long and/or the compression scheme selected is slow, saving the file might take a while. However, if you didn't change any video compression formats, the save will be much faster, since no recompression takes place.

.AVI VIDEO COMPRESSION

If you thought audio files chewed up your hard-disk space quickly, wait until you start using video! CD quality audio takes up about 10 MB of hard-drive space per minute:

(60 seconds x 44,100 samples per second x 2 tracks for stereo x 2 bytes per 16-bit sample). On the other hand, a typical video for multimedia use will easily contain about 200 MB per minute (uncompressed), and this for a small window of video:

(60 seconds x 15 frames per second x 320 x 240 x 3 bytes for 24 bit pixels).

Using uncompressed video as a final distribution format is out of the question for most practical purposes. Video compression is a necessary evil.

COMPRESSION ALGORITHMS

Many different compression algorithms exist and many more will continue to be created in the next several years. With all algorithms, trade-offs will have to be made between video quality, size reduction, and compression/decompression processing time. Compressing by large amounts will quickly generally produce visual artifacts such as jumpy or grainy video.

When saving an .AVI in Sound Forge, you can select from all of the installed compressors installed in your system. Included with Windows are *Cinepak by Radius (cvid)*, *Intel Indeo (IV32)*, and *Microsoft Video 1 (CRAM)*. These can be selected from the Video Compression dialog (see section above).

COMPRESSION SETTINGS

Many compressors use key frames during compression and decompression. A key frame is a frame in the video stream that is not inter-frame compressed. The following compressed frames are derived from the closest key frame. Having fewer key frames will make the video size smaller. However, it will take a lot more time to display the frames, especially when jumping around in the video.

NOTE

When editing .AVI files in Sound Forge, the display speed will suffer greatly if there is only one key frame in the entire file. In such cases, Sound Forge needs to scan all preceding data when drawing every single frame.

Also in the *Video Compression* dialog, you can specify the *Data Rate* of the video, which relates directly to the final size of the file. Different compressors do different things with this value, since it's only an expected value. In most cases you should leave this unchecked. If you have very limited playback data rates, such as in a single-spin CD-ROM, you might then want to set this value. However, note that the *Compression Quality* control setting will also affect the output size.

NOTE

Use the Preview button to get an idea of the video's output quality and compression amount.

EDITING AUDIO IN .AVI FILES

When doing any operations on the soundtrack of an .AVI file, the video stream is always left untouched. For example, doing a cut, paste, time compression, etc. on a sound file will not change the duration of the video stream.

In Sound Forge, the video stream and audio stream will always start together, and you will see the video frames only where audio exists. Therefore, if you want to view the entire video stream, you should use *Insert Silence* to add enough audio to run the entire length of the video. You can then later mix in any needed sounds into the file.

USING THE PLAYLIST WITH .AVI FILES

Once you've created and arranged regions in the *Playlist* (see a previous section), playing from the *Playlist* will keep the original synchronization between the audio and video. In other words, if you have a *Playlist* that jumps all over the file, both the audio and video will jump around.

In a sense, you can use the *Playlist* to perform basic video edits. For example, to remove unnecessary material from the beginning and end of the video, simply make a region containing only the video that you need in the *Playlist*. Playback from the *Playlist* will play only the selected regions in the specified order.

CREATING A NEW .AVI FILE FROM THE PLAYLIST

After arranging the *Playlist* entries, you often want to create a new file containing the final arrangement. To do so, select *Playlist|Convert to New* from the *Special* menu. This will create a new file, which when played back will contain the exact sequence of audio and video that the current *Playlist* specifies.

When creating a new file, each region's start and end point must line up with a frame boundary. If the regions do not line up, they will be automatically adjusted. Region start points are moved backwards and end points moved forward, with silence being used to fill the frame. This can cause the output file to be slightly longer than the *Playlist* playback time.

To ensure that *Convert to New* works exactly like playing from the *Playlist*, make sure that all the region endpoints are frame aligned. This can be done by enabling *Auto Snap to Time* (*Options* menu) and using the *Time & Frames* or *Absolute Frames* status format before creating regions.

EDITING SOUND FORMATS

This section will show you how to deal with the many different sound formats. You will learn how to:

- Change Sound Formats
- Convert Sound Files from 16-Bit to 8-Bit
- Convert Files
- Embed Summary Information in .WAV and .AVI Files

CHANGE SOUND FORMATS

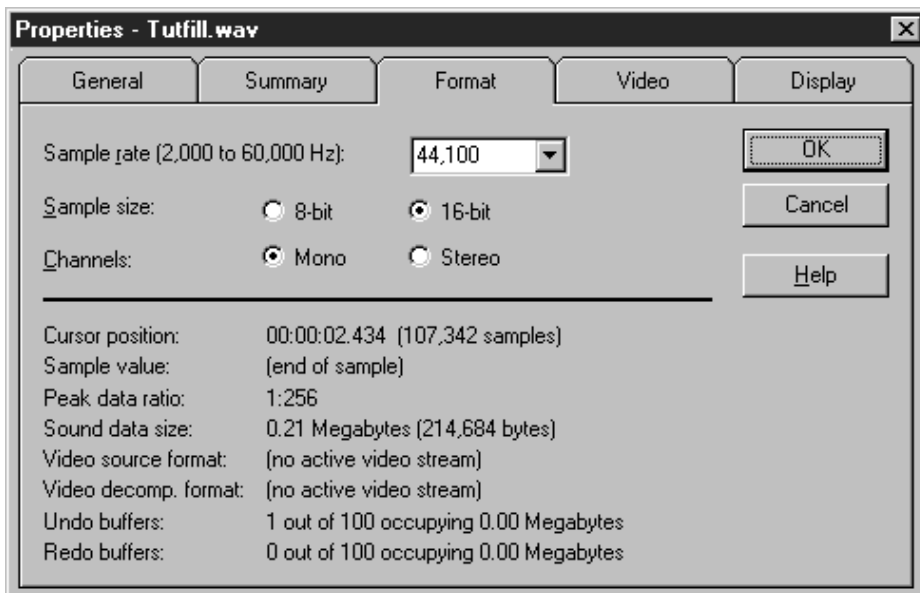
When you open a file or create a new window, the format for the window is shown on the first three status fields of the main status bar at the bottom of the program window. The format parameters consist of the **Sample Rate**, **Sample Size**, and number of **Channels (Mono or Stereo)**. The last status fields display the **Total Length** of the sound data, and amount of **Free Storage** available.



NOTE

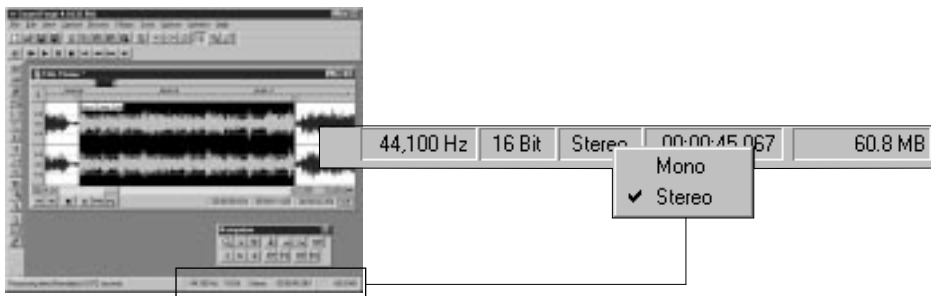
Changing the sample rate will make your sound play faster or slower and with a higher or lower pitch. If you want the file to sound the same at a different sample rate, use the Resample function under the Process menu.

When you right-click on any of the three format fields, a shortcut menu will pop up at the current mouse location allowing you to change that parameter. You can also change all three parameters at one time by selecting **Properties | Format** under the File menu. The **Properties** dialog shown on the following page can also be invoked by double-clicking on any of the format fields, or by right-clicking on the waveform display of the data window and selecting **Properties** in the shortcut menu.



MONO TO STEREO AND STEREO TO MONO CONVERSIONS

For demonstration purposes, open the file TUTOR1.WAV located in the directory where Sound Forge 4.0 was installed (in most cases this will be C:\Program Files\Sound Forge). To convert the file from mono to stereo, right-click on the status field which has the word *Mono* and select *Stereo* from the shortcut menu.



When you do conversions from mono to stereo, you will be presented with a dialog which asks you where you want to put the data. Your choices are as follows:

- Left Channel** The mono data is placed in the left channel. The right channel is set to silence.
- Right Channel** The mono data is placed in the right channel. The left channel is set to silence.
- Both Channels** The mono data is copied into both channels.

When performing stereo to mono conversions, the following options are available:

- Left Channel** Mono data is taken only from the left channel of the stereo file.
- Right Channel** Mono data is taken only from the right channel of the stereo file.
- Mix Channels** Mono data is created by mixing both channels of the stereo file.

For now let's put the data in the left channel so we can get a feel for how the function works. Click the radio button next to **Left Channel** and select the **OK** button.

You will now have data in the upper half of the data window (the left channel) and silence in the right channel. Press the **Play** button and you will hear "Wow, sound editing just gets easier and easier" in only the left channel.

If your card supports only mono data you will still be able to play stereo files if you set your playback device to Sound Mapper in the Wave section of Preferences in the Options menu.

NOTE

For now let's go back to a mono sound file, but we'll do it in a different manner. Select the option **Properties** in the shortcut menu that appears when you right-click on the waveform display. In the **Properties** dialog select the **Format** tab. This page allows you to change all the parameters at once. Set the radio button for **Channels** to **Mono** and select the **OK** button. Select the **Left Channel** for the source channel in the **Stereo to Mono** dialog and we're back where we started.

NOTE

The above steps were used for demonstration purposes. Simply using Undo would have been much more efficient.

CHANNEL CONVERTER

The *Channel Converter* (Process menu) is also used to convert between mono and stereo formats. It allows independent level settings for each channel and can be used to inter-mix the left and right channels of a stereo file to create interesting pan effects. For simple conversions, use the included presets. See Channel Converter in the Process menu section of the Reference chapter for more detailed information.

CONVERT SOUND FILES FROM 16-BIT TO 8-BIT

To save storage space, 16-bit sound files are often converted to 8-bit files. However, when you represent a sound file with only 8-bits, an audible distortion referred to as quantization error becomes very prominent. To minimize quantization error, there are several things you can do to the 16-bit sound file before converting it to 8-bit. You can:

- Apply a *Noise Gate* (Effects menu) to completely mute out the silent parts in a sound file. Often, a low level signal in a 16-bit sound file will become noise after the 16- to 8-bit conversion, so it's best to have complete silence between the sound parts.
- Select *Dynamics/Graphic* (Effects menu) to apply dynamic compression to the sound file. A small amount of compression (2:1 or less – the presets are usually sufficient) will lower the dynamic range of a sound, making it easier to represent using 8-bit samples.
- *Normalize* (Process menu) the sound to 0 dB (peak). This ensures that the entire dynamic range available in 8-bit samples is used and lowers the signal to noise ratio.

NOTE

You can apply both dynamic compression and normalization at once by using the Normalize function with the Average RMS power option. Set Apply dynamic compression in the If clipping occurs drop-down list.

Once you have performed the above operations, you should use *Convert to 8-Bit* in the Process menu to do the 16 to 8-bit conversion. Sound Forge provides three different 16 to 8-bit conversion options; *Truncate*, *Round*, and *Dither*. *Truncate* simply removes the lower 8 bits and they are lost forever. *Round* is similar to *Truncate*, only some attempt is made to represent the highest of the low bits by rounding up if possible. *Dithering* is used to mask the quantization noise with less obtrusive noise. In the *Convert to 8-Bit* dialog, set the *Dither* bit depth to a low value (like 0.1) and increase the value slowly until you reach an acceptable sound.

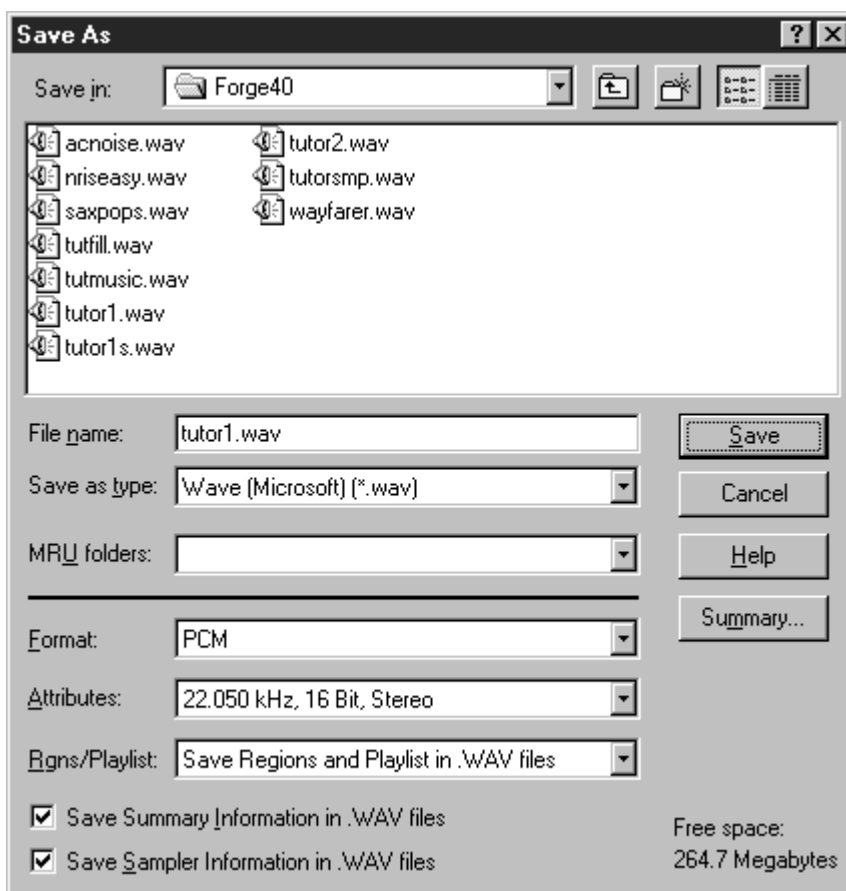
Remember, an 8-bit sound file will always sound noisier than a 16-bit sound file, so whenever possible, stick with 16-bit. When trying to save space, it is possible to get better results from lowering the sample rate (see *Resample* under the Process menu covered in the Reference chapter) instead of using 8-bit samples.

CONVERT FILES

When producing audio files you may find that you need to provide the file to a client in a variety of formats. This could be .WAV files with data formats such as 8-bit mono, 16-bit stereo, or even ADPCM. You may need the files converted to Raw data, .VOC, or some other PC format. You may even need to convert to a different platform like Macintosh .AIF files or Amiga .SVX.

We have seen in the previous sections that you can quickly convert between mono/stereo, 8/16-bit, and change the sample rate in Sound Forge. You can then use the *Save As* command to save the file to a new name, or you can save in a new format within the *Save As* dialog.

If you don't have TUTOR1.WAV open, open it up now. Select the *Save As* command from the File menu and the dialog for saving will be displayed.



When the *Save As* dialog is displayed, there are three controls that are used for file conversion; the *Save as type*, *Format*, and *Attributes*. The *Save as type* is the list of all available output file types. The *Format* field sets the format of the data that is saved in the file. The *Attributes* field controls sample size, sample rate, and the number of channels.

When you bring up the *Save As* dialog, the *Save as type* will be set to the type from which the file was opened. The default for new files is *Microsoft Wave (.WAV)*. You can change to the type of file you want to save by selecting the new type from the *Save as type* list. Notice that the extension of the file name will be changed to reflect the type of file you are saving.

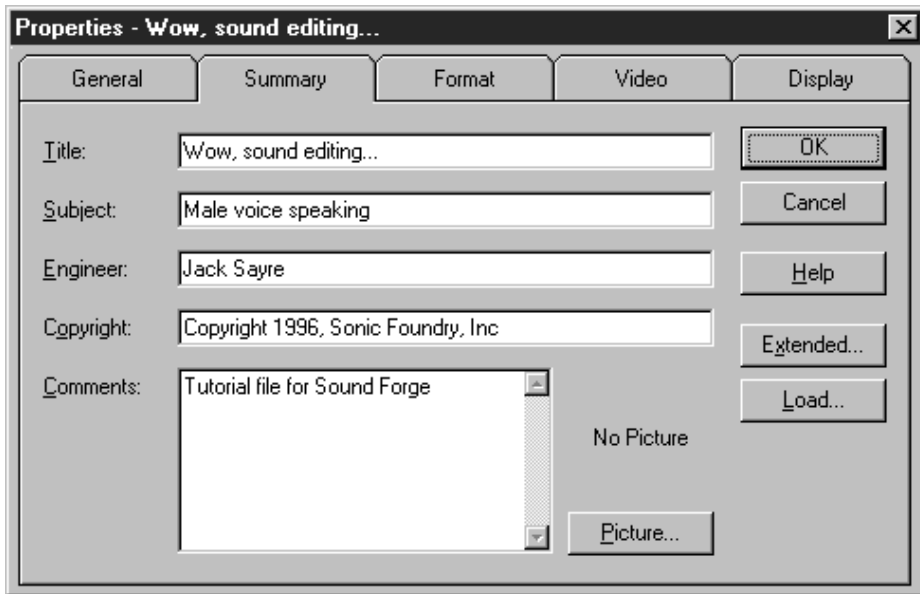
The *Format* drop-down list shows the format of the data that will be saved in the file. Usually this will be PCM. PCM (*Pulse Code Modulation*) is a standard format for normal uncompressed audio. If you are saving a file to a file type that supports other formats, they are displayed in the drop-down list. Other formats are typically used when saving audio data in a compressed form.

The *Attributes* drop-down list shows the sample size, channels, and sample rate that will be stored when saving this file. Examples include: 11,025 Hz, 8-bit Mono or 44,100 Hz 16-bit Stereo. You can quickly change the attributes of the data by selecting one of the options available in the drop-down list. When you change from a mono file to a stereo file using the *Attributes* field, that data is copied to both channels. If you change a stereo file to a mono file, the data is mixed to one channel.

When you change the *Save as type* you may notice that the *Format* and *Attributes* fields will also change. This is because not all file types support the same types of sound data. For example the *Dialogic .VOX* format only supports 16-bit mono *Dialogic ADPCM* files. This means you can't store stereo 8-bit data in this file. Don't worry, in most cases Sound Forge automatically converts any file you save to an acceptable format.

EMBED SUMMARY INFORMATION IN .WAV AND .AVI FILES

Microsoft .WAV and .AVI files allow you to store over 30 text fields including: *Creation Date*, *Copyright*, *Keywords*, and a variety of other informational text data. Sound Forge supports viewing and editing any of these fields.



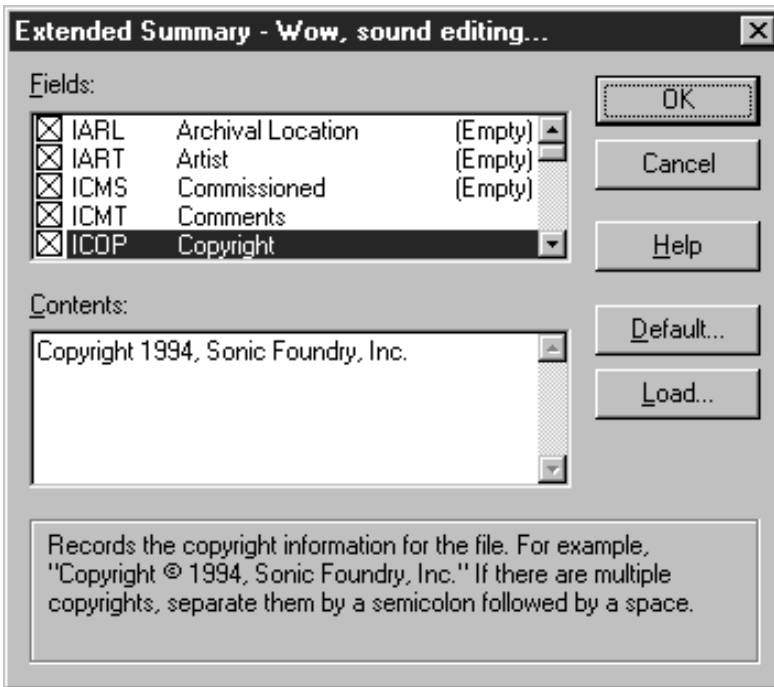
SUMMARY INFORMATION

When editing .WAV and .AVI files, the *Summary* tab is available in the *Properties* section of the File menu. This allows you to view and/or change basic summary information currently stored in the file.

EXTENDED SUMMARY INFORMATION

The *Extended Summary* dialog consists of a list of available fields, each of which may or may not have attached text. The *Fields* list shows the abbreviation of the field type, a short description, and if the field is currently empty, the word (*Empty*) appears after the short description. Below the list of fields is the *Contents* which shows the contents of that field. At the bottom of the dialog is a longer description of the current field.

To the left of each field in the list is a check box that is used to enable or disable fields of this type when saving .WAV and .AVI files. The field is enabled or disabled by clicking on the check box. If a field is empty, i.e. has no text associated, enabling the check box has no effect on a saved file. Field information is only saved if text information exists for that field.



The *Save Summary Information in .WAV/.AVI files* check box in the *Save As* dialog is used to quickly enable/disable saving of all summary fields to a file during a *Save As* operation. This allows you to quickly strip all summary information from a file.

For a description of the fields available refer to the *Properties* within the *File* menu section in the Reference chapter later in the manual.

DEFAULT

If the *Default* button is selected, the text in the summary fields is saved with the default fields that are automatically filled when creating a new .WAV file. The *Creation date* field (ICRD) is always filled with the current date for new files. Saving a custom default setup is handy for saving copyright and engineer information for new files created at your site.

Open the file TUTOR1.WAV and select the **Summary** tab from **Properties** in the **File** menu. This will show you the **Summary** dialog. Press the **Extended** button and notice that the default for saving summary information is to have all fields enabled. The fields that have already been filled in for TUTOR1.WAV are **Comments**, **Copyright**, **Creation Date**, **Engineer**, **Keywords**, **Name/Title**, **Product**, **Subject**, **Software**, and **Source**. Scroll to and select each field. You will see the embedded text for each of these fields.

ADDITIONAL EMBEDDED INFORMATION

The Microsoft .WAV and .AVI file formats allows non-text data to be embedded in the sound file. If you open and edit a file which has additional information created by other software, Sound Forge will keep track of this data and place it back in the file when you save to the original format.

If you wish to add or remove additional data (such as embedded bitmaps, metafiles, etc.), you can use the **Save As** item from the **File** menu. When you save data using the **Save As** option, Sound Forge will ask you if you wish to copy the additional non-audio information to the new file. At this point you can answer either **Yes** to place the additional information in the new file or **No** to save only audio data and any additional data types you have selected.

NOTE

Sound Forge will only prompt you to save this type of data if such data exists in the file.

CREATING REALAUDIO, REALVIDEO AND NETSHOW FILES

Sound Forge supports two different formats for creating files that are optimized for streamed playback over (usually) slow networks such as the Internet. RealNetwork's RealAudio and RealVideo (.RM,.RA), and Microsoft's NetShow Advanced Streaming Format (.ASF) files can be created in Sound Forge. To play these files, you must have the appropriate player application installed in your computer.

USING SOUND FORGE TO CREATE REALAUDIO AND REALVIDEO FILES

Sound Forge gives you the ability to save sound and video files in RealAudio and RealVideo formats along with editing tools that help maximize the audio's fidelity. This section first explains how to save files in RealAudio format, then covers saving files in RealVideo format. All of the tips for converting files to RealAudio also apply for converting AVI files to RealVideo. For more information on authoring RealAudio and RealVideo content, refer to RealNetworks' RealMedia home page (www.real.com).

PREPARING A SOUND FILE FOR REALAUDIO ENCODING

When you have finished recording a new file or opening pre-existing files, it is recommended that you perform the following processes on the file before saving the file in a RealAudio format.

Sound Forge allows you to save preset settings for each of these processes. So once you set them up the first time, you can always use the same settings.

NOTE

NORMALIZE

This process ensures that the loudest part of the file is at the maximum level.

To begin the normalization process press **Ctrl+A** to select the entire file. Select **Normalize** from the **Process** menu and the **Normalize** window will appear.

Next, select the **Peak Level** option and move the **Normalize to** slider to **-.50 dB (94.41%)**. Press the **OK** button and Sound Forge will normalize the file.

COMPRESS

If you have not already audio compressed the clip, you can do so now. RealAudio files benefit from compression because it increases the overall volume and intelligibility of the sound. The suggested method is to use very light compression.

Start the compression process by pressing **Ctrl+A** to select the entire file. Select **Dynamics** from the **Effects** menu and then select **Graphic**. The **Graphic Dynamics** window will open.

Next, check the *Sync Stereo Gain* and the *Auto Gain Compensate* checkboxes. Move the *Threshold* slider to -10.00 dB and the *Ratio* slider to 2.0:1.

Finally set the *Attack* control to 1.0 milliseconds, the *Release* control to 300.0 milliseconds and press the OK button.

SET THE REALAUDIO AND REALVIDEO TEXT FIELDS

RealAudio and RealVideo files can store information about the title, author and copyright of the clip. To set the RealAudio and RealVideo text fields:

Choose *Properties* from the File menu. The Properties window for the current file will now open. Click the *Summary* tab and enter the Title, Engineer, and Copyright information. This information will appear in the Title, Author, and Copyright text boxes on the RealPlayer.

Press the OK button and the fields will be transferred in to the RealAudio or Video file when it is saved.

SAVING AN AUDIO FILE AS REALAUDIO

To save an audio file as RealAudio select *Save As* from the File menu. Enter the file name of the file you are saving and select *RealMedia (*.rm;*.ra)* as the *Save as type*. The default extension for this type of file is .rm. To save using the .ra extension, type .ra at the end of the file name.

Next, press the *Save* button and the *RealMedia Save Options* dialog will open. The options in this dialog are covered in the reference section under *Save As*.

Select the *Audio compression type* from the drop-down list. A brief description of the highlighted compression type is displayed below this list. If the file you recorded was stereo and you selected a mono compression type, select the *Stereo to mono conversion*. It is recommend that you use the *Both Channels* option.

Finally, press OK. After the file has been saved, a message appears asking whether you want to listen to the saved file through the RealPlayer. This message verifies that you have saved the file in RealAudio or RealMedia format. You can listen to the file by clicking the *Yes* button if you have the RealPlayer installed.

SAVING AN AVI FILE AS REALVIDEO

To save an AVI file as RealVideo select *Save As* from the File menu. The *Save As* dialog will open. Enter the file name of the file you are saving and select *RealMedia (*.rm;*.ra)* as the *Save as type*.

Next, press the *Save* button and the *RealMedia Save Options* dialog will open. The options in this dialog are covered in the reference section under *Save As*.

If you wish to include the audio in the new *RealVideo* file, make sure the *Save Audio Stream* option is checked. Then, select the *Audio compression* type from the drop-down list. A brief description of the highlighted compression type is displayed below this list.

If you wish to include the video in the new *RealVideo* file, make sure the *Save Video Stream* option is checked. Then, select the *Video compression* type from the drop-down list. For a brief description of the two compression types, see the reference section under *Save As*.

Finally, specify the *Desired maximum data rate* in total kilobits per second (Kbps). This option is only available if you are saving a video stream. Refer to the documentation for *Desired Maximum Data Rate* in the *RealMedia Save Options* section for more information.

Press the *OK* button. After the file has been saved, a message appears asking whether you want to listen to the saved file through the *RealPlayer*. This message verifies that you have saved the file in *RealVideo* format. You can view the file by clicking the *Yes* button if you have the *RealPlayer* installed.

USING SOUND FORGE TO CREATE MICROSOFT NETSHOW FILES

This tutorial shows how to create *Advanced Streaming Format (ASF)* files using *Sound Forge*. ASF is a multimedia file format used for streaming information across the Internet using *Microsoft NetShow*. With ASF files, you can deliver multimedia content, audio or video, open Web pages, or deliver scripting commands to the client computer.

Sound Forge includes support for saving WAV and AVI files and their associated markers and script commands in an ASF file. ASF files can be played off of your local drive, streamed off of an HTTP server, or optimally stored and streamed off of a *NetShow* server.

NETSHOW REQUIREMENTS

In order to play ASF files you must install the *Microsoft NetShow Player and Tools*. These components install all of the necessary audio and video codecs for saving ASF files. Both of these components are available on the Sound Forge CD-ROM and can also be downloaded free of charge from the NetShow web page, www.microsoft.com/netshow.

To install the NetShow Player select *Microsoft NetShow Player* from the list in the main setup program on the Sound Forge CD-ROM. This will run Microsoft's setup program for the player.

To install the ASF authoring and conversion tools, select *Microsoft NetShow Tools* from the list in the main setup program on the Sound Forge CD-ROM. This will run Microsoft's setup program for the NetShow Tools.

Further documentation for this software is available on the NetShow web page. At this time, the components shipped on the Sound Forge CD-ROM are for the x86 and Pentium platforms only. See Microsoft's web page for availability of components for other platforms.

USING THE BUILT-IN TEMPLATES

Sound Forge contains a number of built-in templates for saving AVI or WAV files as ASF files. These templates have been designed by Microsoft to be optimized for certain bandwidths.

Open *Sample1.avi* from the Tutorial Files\NetShow folder on the Sound Forge CD-ROM. This file is a short video clip that uses the Intel Indeo Interactive VCM video codec installed with the NetShow Player. The Intel codec is well suited for mid to high bit-rate video. The audio track in this file is PCM.

Choose *Save As* from the File menu. Select *Advanced Streaming Format* from the *Save as type* drop-down list. Now, press *Save* and the *NetShow Settings* dialog will appear.

A high bit rate ASF file can be created from this file by simply choosing a built-in template. The numbers preceding the names of the templates indicate the target bandwidth for the ASF file. Higher numbers require wider bandwidth connections.

For this file, choose *100 Presentation Video*. This template is designed for a Dual ISDN connection and a "talking head" type video like *Sample1.avi*. Press the *Create* button and the AVI will be saved.

CREATING CUSTOM TEMPLATES

Custom templates can also be created in the *NetShow Settings* dialog and saved for future use. To create a custom template, press the *Custom* button. The *NetShow Encoder Template* wizard will start and guide you through the template creation process.

The first step is naming the template. Type the name of your custom template in the space provided and press the *Next* button to continue.

The second step is determining the desired bandwidth for the output file. The bandwidth selected in this step will determine what compression options are available in the next step. You can enter a custom bandwidth in bits per second by selecting the *Custom* option.

A bandwidth selection of 14.4 will not allow you to save a video stream due to the amount of bandwidth required for video.

NOTE

The final step is to choose the audio and video codecs you wish to use and press the *Finish* button. Each audio codec will have a set of *Formats* that are available depending on the bandwidth that was selected in the previous dialog.

Microsoft provides a free Content Creation Guide on their web site, www.microsoft.com/netshow. This guide is a great resource for understanding what codecs work best in certain situations.

NOTE

ADDING MARKERS AND SCRIPT COMMANDS

Sound Forge can embed markers in ASF files. These markers can be used for navigation in the file or for triggering events or custom script commands.

ADDING NAVIGATION MARKERS

Open *Sample1.avi* from the *Tutorial Files\NetShow* folder on the Sound Forge CD-ROM if it isn't already open. If you want to viewers to have the ability to jump to certain parts of the ASF file, you will want to add markers in the source file.

Find the part of the AVI file where John says, “No, not gonna do that.” Place the cursor right before that point and press the M key on the keyboard. The *Edit Marker/Region* dialog should now be displayed. Type, “No, not gonna do that” in the *Name* field and press OK. You now have a named marker in the file.

Choose *Save As* from the *File* menu and select *Advanced Streaming Format* from the *Save as type* drop-down list. Press the *Save* button to get to the *NetShow Settings* dialog. Choose the *100 Presentation Video* template.

Before you press the *Create* button, look directly above it at the *Use markers in source file* checkbox. Make sure that this box is checked and then press *Create*.

When the file is done saving choose *Yes* to open the ASF file. If you look just above the copyright information part of the *NetShow Player* window, there is a drop-down list that will contain all of the markers that have been saved in the file. The marker you just added should now appear in that list. Choose it and press the *Play* button. Playback should now start from that point in the ASF file.

ADDING SCRIPT COMMANDS

Markers are also used to store script commands that are embedded in the ASF file. The most popular of these is the URL flip. This command will cause the opened web browser to open specified pages as the ASF file reaches certain points in the file.

Sound Forge has a special function to automatically add certain script commands but regular markers will work just as well. Using the built in function will simply save you from remembering the syntax for basic commands.

Open *Sample1.avi* again. Put the cursor close to the beginning of the file and choose *ASF Marker/Command | Add* from the *Special* menu. Notice that the name of the new marker has already been filled in with the correct syntax for a URL flip. After the *http://* type, *www.sonicfoundry.com/SoundForge* and press OK. A new marker has been inserted into the file.

Now save the file as ASF as described above, but before you press the *Create* button in the *NetShow Settings* dialog, make sure that *Use script commands in source file* is checked.

After the file is saved, choose *Yes* to listen to the file in the *NetShow Player*. As the file is played, you will notice that your default web browser has been opened and the *Sound Forge* page of the *Sonic Foundry* web site has been loaded.

If you aren't already connected to the internet, the time it takes to connect will, of course, throw off the timing of the URL flip. Ideally the viewer of the ASF file will already be online and the URL flip will be more accurate.

NOTE

SAMPLE HTML USING AN ASF FILE

On the *Sound Forge* CD-ROM in the `\SAMPLES\NETSHOW` folder there is HTML file containing an example of how use the *NetShow* ActiveX control to add an ASF file to a web page.

The HTML sample requires Microsoft Internet Explorer v3.0 or better.

NOTE

Copy `SAMPLE.HTM` to the folder on your hard drive that contains the ASF file you wish to use. Next, modify the `FILENAME` parameter within the `NSOPlay` object to reference your local ASF file (a fully qualified path is required). Then, open the edited `SAMPLE.HTM` file within your browser. There are additional instructions on the sample page.

RECORDING

Sound Forge has a variety of sound recording modes and options each of which come in handy for different situations. These include recording new data, recording over existing data, commonly known as **Punch-In** mode, and **Remote Recording**. In this section we will cover:

- Recording to a New Window
- Selecting an Alternate Record Window
- Sync While Recording
- Available Recording Time and Recording Levels
- Record Meters
- DC Adjust
- Previewing Recorded Sounds
- Using the Prepare button
- Recording Modes



To start a record session you can either select **Transport** then **Record** from the **Special** menu or press the **Record** button on the **Transport** toolbar. The **Record** button is the first button on the **Transport** toolbar with the red circle on it.

NOTE

If you experience problems recording, refer to the Troubleshooting chapter for information on common recording problems.

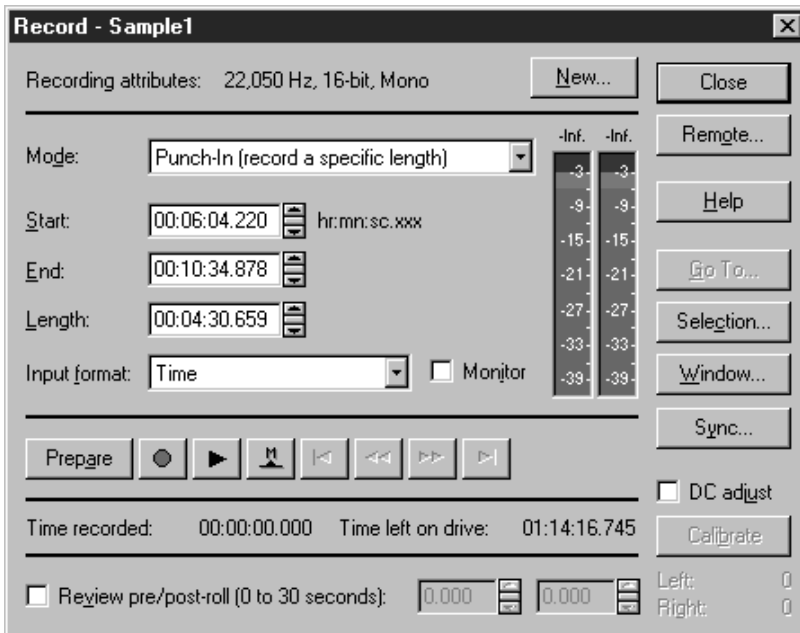
After pressing the **Record** button or selecting **Transport | Record** from the **Special** menu, you will be presented with the **Record** dialog. Notice that the window into which you will be recording has its title displayed in the dialog title.

NOTE

Warning: *The Record and Record Remote dialogs are always destructive and contain no Undo capabilities.*

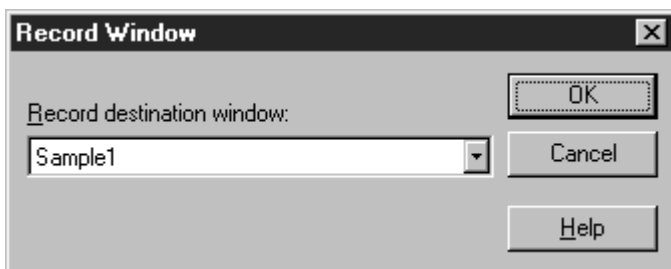
To prevent accidentally recording over sound data, record into a new or scratch data window and paste the takes you want to keep into the desired sound data window. For extra safety Sound Forge automatically defaults to record in a new window when the current sound file is opened in direct edit mode.

In the upper left corner of the dialog are the *Recording attributes*. These are the record sample rate, sample size, and number of channels that will be used when recording. These attributes are applied to the data window into which you will be recording. If you want to change these attributes, exit the *Record* dialog and change them in the data window, record to a new window, or record to another window.



RECORDING TO A NEW WINDOW

If you want to record to a new window rather than the currently selected record window, press the *New* button found at the upper right of the *Record* dialog. This brings up the *New Window* dialog where you can specify the *Sample rate*, *Sample size*, and *Channels* for the new data window. These attributes will be applied while recording.



SELECTING AN ALTERNATE RECORD WINDOW

If you wish to record to a window other than the one currently displayed in the Record dialog title, you can do so by selecting the *Window* button. Pick the window you want to record to from the drop-down list in the *Record Window* dialog and select *OK*. The title of the window you select will now appear in the Record dialog title.

SYNC WHILE RECORDING

Use this dialog to set sync options for recording in conjunction with other devices or applications that respond to MIDI/SMPTE time code.

When *MTC/SMPTE Output Synchronization* is enabled, Sound Forge generates SMPTE/MTC to the output device selected in the *MIDI/Sync Preferences*. When you start recording, SMPTE output begins at the *Start* value. If *Pre-roll* is selected, SMPTE output will begin before recording by the specified amount. For example, if the *Start* value is 00:00:30:00 and the *Pre-roll* is 00:00:10:00, SMPTE time code will be generated starting at 00:00:20:00. After 10 seconds (at 00:00:30:00) have elapsed, recording will commence.

When *MTC/SMPTE Input Synchronization* is enabled, recording begins (is triggered) when the *Start* time is received from the current *MIDI/Sync Input* device (see the *MIDI/Sync* page in the *Preferences* folder). Recording will then stop when the end time is reached.

NOTE

The *Prepare* button must be activated for record to be triggered by incoming MTC/SMPTE.

Output Synchronization Start	Start time which will be generated when record begins
Output Synchronization Pre-Roll	Pre-roll time from the start time at which SMPTE generation will begin
Input Synchronization Start	Time at which recording will begin
Input Synchronization End	Time at which recording will end

Check to make sure that the input and output devices are specified correctly in MIDI/Sync under Preferences (Options menu). The SMPTE format here is determined by Status Format in the Options menu.

NOTE

When recording tracks in from a tape deck that can perform a chase lock, the **Synchronization Pre-roll** can be set to compensate for the time it takes the tape deck to rewind and begin chasing prior to the synchronization start time.

AVAILABLE RECORDING TIME

Near the bottom of the dialog you will see the **Time recorded** and **Time left on drive** fields. These two boxes show how much time you have recorded and how much time is available on your hard drive for additional recording. If your **Time left on drive** field is displaying a limited amount of available time you may want to free up some space on your hard drive or pick an alternate drive where Sound Forge stores its temporary files. You can get more information on temporary file usage by referring to the Reference chapter on **Temporary Storage**.

Sound Forge can only record a single file up to 2 Gigabytes in size regardless of the capacity of your hard drive. This should be sufficient for most users, however.

NOTE

RECORD METERS

Like the Sound Forge *Play Meters*, the *Record Meters* can be scaled to view differing dynamic ranges by right-clicking on the meters and selecting the appropriate range from the shortcut menu. For most recording situations it is probably best to select -42 to 0 dB, as this is the most practical range to view a good record signal. To view very low levels, select -90 to 0 dB. This is a good way to measure the amount of noise you have in your system due to noisy sound sources and/or poor equipment (such as an inexpensive sound card).

Also in this shortcut menu, you can set the meters to *Hold Peaks* and *Hold Valleys*. It is recommended that *Hold Peaks* be checked, as this is a good way to view the peak levels while setting the record level (see below). To reset the current peak, single-click on the meter's text output (just above the meters).

NOTE

For Windows 95 and NT 4.0 keyboard shortcut lovers, the Play Meter shortcut menu can be reached by the keyboard shortcut Alt+Spacebar.

CHECKING RECORD LEVELS

Sound Forge allows you to check the level of your input source before recording begins. To view your levels engage the *Monitor* check box. The meters will light up in relation to the level of the incoming signal. For best results, the level should be somewhere in the yellow range with an occasional red. Once your levels are checked you can immediately begin recording by selecting the *Record* button. If you do not see the meters light up, you may have your mixer levels or input source set incorrectly. You can refer to the Troubleshooting chapter for more information on these problems.

ADJUSTING LEVELS USING THE PEAK

The peak values displayed above the level meters are useful for maximizing your input level without clipping. When recording you generally want your input signal to be as hot as possible without clipping. By this we mean you want your input levels to be as high as possible without exceeding the range of values which can be stored digitally when recording. When you clip, the peaks of your waveform become clipped off resulting in distortion. The peak values show you (in decibels) the highest peak that you have reached since hitting the *Monitor* button.

To adjust your levels, select the **Monitor** check box so that Sound Forge begins to listen to your recording device. This is just like recording except that Sound Forge doesn't store any of the data it receives. Apply an input signal by speaking into your microphone, playing your CD, or whatever it is you're trying to record. If the peak value stays at a low value, increase the levels of sound you are supplying so that the peak value is somewhere in the -6 dB range. If the peak reaches 0 dB then you have clipped and will see the word **Clip** above each meter. Once you lower your input levels, right-click on the meters and select **Reset Clip** to clear the current peak value. Sound Forge always keeps the maximum peak displayed above the meters or via **Hold Peaks**.

Once you have adjusted your levels you can immediately begin recording by selecting the **Record** button or end monitoring the levels by un-checking **Monitor** below the meters.

It is particularly important to record sounds with the hottest levels possible when you plan to later convert 16-bit data to 8-bit. This assures that you will use the greatest dynamic range possible in an 8-bit file (which uses fewer values to represent the waveform).

NOTE

DC ADJUST

When the **DC Adjust** option (underneath the **Sync** button) is turned on, Sound Forge will attempt to automatically compensate for any DC offsets generated by the sound card during the record process. If you plan on using the **DC adjust**, it will be necessary to calibrate it at least once for the currently selected record device. To calibrate the DC adjust amount, simply press the **Calibrate** button. The two left and right values shown below are the calculated DC offset amplitudes.

A DC offset of around 30 is not uncommon even for sound cards with very low noise floors. This will show up in the record meter as -60 dB! Don't worry, this doesn't mean you're losing 30 dB of resolution. However, for the meters to display down to -90 dB you need to calibrate for this offset. You can recalibrate at any time, even while recording. However, it's best to calibrate with silence at the record inputs.

NOTE

If you have more than one sound card installed in your system, it is a very good idea to calibrate the DC adjust each time you switch devices. Some sound cards will also need to be recalibrated each time you change sample rates and/or between mono and stereo. Experiment with calibration in different modes to find out what works best for your card.

PREVIEWING RECORDED SOUNDS

After recording your material, you can listen to what you have recorded by selecting the **Play** button. You can also listen to the section over which you plan to record in **Punch-In** mode. To stop playing select the **Stop** button at any time.

USING THE PREPARE BUTTON

The **Prepare** button is used when you need Sound Forge to begin recording as soon as possible after selecting the **Record** button. The **Prepare** button opens the wave device and loads all recording buffers in order to minimize the time between selecting the **Record** button and when the sound actually begins recording.

NOTE

*The **Prepare** button is optional. It is not necessary to select this button prior to recording, however it does allow for more accurate takes in the **Punch-In** mode.*

RECORDING MODES

Sound Forge has five different modes of recording. These are **Automatic retake** (automatically rewind), **Multiple takes creating Regions**, **Multiple takes (no Regions)**, **Create a new window for each take**, and **Punch-In** (record a specific length). Each mode is described below.

AUTOMATIC RETAKE (AUTOMATICALLY REWIND)

The *Automatic retake* mode is the easiest method of recording. Recording starts at the position shown in the *Start* field when you select the *Record* button and continues until you select the *Stop* button. Any data that is currently after the position in the *Start* field will be replaced. When recording is stopped, the start position is reset to the beginning of the take allowing an immediate review and retake if desired.

Automatic Retake is the default mode when recording into an empty data window or when you select the *Record* button with no data selected in the current data window.

MULTIPLE TAKES CREATING REGIONS

The *Multiple takes creating Regions* mode allows multiple takes to be recorded with each take automatically defining a region in the *Regions List*. Recording starts at the position shown in the *Start* field when you select the *Record* button and continues until you select the *Stop* button. Any data that is currently after the position in the *Start* field will be replaced. When recording is stopped, the start position remains at the end of the take, allowing the next take to be recorded immediately.

MULTIPLE TAKES (NO REGIONS)

The *Multiple takes (no Regions)* mode allows multiple takes to be recorded, but no regions are defined in the *Regions List*. Recording starts at the position shown in the *Start* field when you select the *Record* button, and continues until you select the *Stop* button. Any data that is currently after the position in the *Start* field will be replaced. When recording is stopped, the start position remains at the end of the take allowing the next take to be recorded immediately.

CREATE A NEW WINDOW FOR EACH TAKE

The *Create a new window for each take* is similar to *Multiple takes creating Regions* except rather than creating new regions with each take, it will actually create a new data window for each take. This will prove helpful in recording data where each take is intended to be saved into its own file.

PUNCH-IN (RECORD A SPECIFIC LENGTH)

Punch-In mode is used when you want to record over a region of data in an existing data window. Recording starts at the position shown in the **Start** field when you select the **Record** button and continues until you select the **Stop** button, or the length of the data recorded is equal to the length in the **Length** field. This makes it easy to record over a section of audio without effecting the rest of the file. You can use the **Play** button to hear the selected punch-in region at any time.

You may adjust the punch-in region by changing the values in the **Start**, **End**, or **Length** edit fields. You may also adjust the format of these fields to a variety of different display status formats by selecting a format from the **Input format** drop-down list.

Punch-In mode is the default mode when you select the **Record** option while you have a region of data selected in the current data window.

USING PRE/POST-ROLL WITH PUNCH-IN MODE

At the bottom of the **Record** dialog are two edit fields which contain the pre-roll and post-roll times. These can be used when listening to a region in **Punch-In** mode. These times define the amount of audio you will hear prior to (pre-roll), and just after (post-roll) playing the selected region. This allows you to hear the transitions between the punch-in region and sound before and after the region. Click on the check box at the bottom of the dialog if you want to use the **Review pre/post-roll** option. To disable pre/post-roll un-check this box.

CHANGING THE START POSITION

The **Start** position determines where playback or recording begins from the time you press the **Play** or **Record** button. Four button controls – **Go to Start**, **Rewind**, **Forward**, and **Go to End** – allow you to quickly change the **Start** position.

Although potentially confusing at first, these controls are really handy when you're doing multiple recordings. The **Go to Start** and **Go to End** buttons move the **Start** position to the beginning or end of the file to which you are recording. After you've recorded a take, the **Rewind** and **Forward** buttons move the **Start** position to the start or end of the take.

For example, to listen to the entire file, press *Go to Start* before pressing the *Play* button. Likewise, to start recording at the end of the file, press *Go to End* before selecting the *Record* button. If you wish to start playback or recording at a specific point, press the *Go To* button and enter the desired *Start* position.

DROPPING MARKERS

The button just to the right of the *Play* button (marked with an “M”) is for dropping markers during recording. For example, if you were recording a series of voice translation prompts, the person reading the prompts can simply deliver a prompt, pause for a moment, and then deliver the next prompt. During the pauses, a marker can be dropped to easily tag each take without interrupting the recording process. The *Markers to Regions* function could then be used to convert all markers into regions in the *Regions List*.

RECORDING STATUS

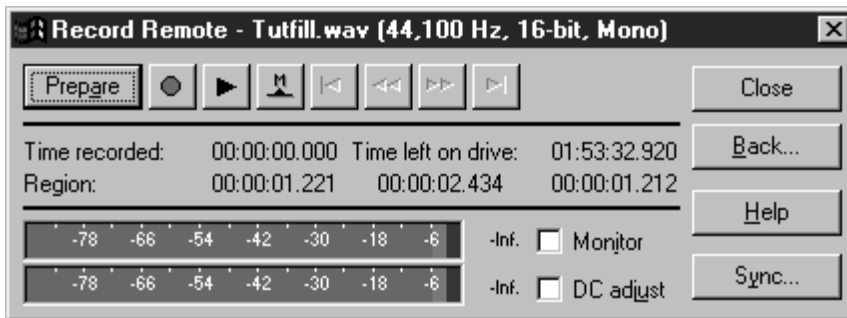
While you are recording, the amount of *Time Recorded* will increase and the *Time left on drive* will decrease. Make sure to keep an eye on your *Time left on drive* if your available record time is limited. It's never fun running out of recording time!

FINISHING RECORDING

When you have finished recording, select the *Close* button to exit the *Record* dialog and return to normal editing mode.

REMOTE RECORDING

By selecting the *Remote* button you enter into *Remote Recording* mode. *Remote Recording* mode hides the main *Sound Forge* window and puts up the small *Record Remote* dialog which always stays as the top-most window.



The Remote dialog is a fully functional but condensed version of the regular Record dialog. It's particularly useful when using applications that control your input sources, sound levels, CD audio, or MIDI sequencing. To return to normal record mode from the Remote dialog select the Back button.

NOTE

You can access all of the same Record dialog features from the Remote Record dialog (right-click on the title bar).

APPLYING SOUND PROCESSING FUNCTIONS

The Effects, Process, and Tools menus contain a variety of functions you can use to apply an effect to the sound data, display information, or even synthesize new sounds. Each function is documented in detail in the Reference chapter later in the manual. In this section, we will run through a few of the functions to help you get a feel for using them.

You will learn how about:

- Applying Simple Processes and Effects
- Applying Functions to Stereo Files
- Using Controls
- Presets
- Previewing an Operation
- Processing Shortcuts

APPLYING SIMPLE PROCESSES AND EFFECTS

APPLYING AN EFFECT TO THE ENTIRE SOUND FILE

To apply an effect, first select a section of data on which you wish to operate. To select the entire file, double-click (triple-click if regions or markers are present) on the waveform display.

If you don't have a selection when you perform an effect that requires one, Sound Forge will apply the effect to the entire file.

NOTE

To show how effects are used we will once again be using the TUTOR1.WAV file. If it's not currently open, please open the file now.

Select the *Reverse* item from the Process menu. The *Reverse* operation reverses the data selection, making it sound like it's playing backwards. Once the operation has finished, press the *Play* button to hear how TUTOR1 sounds when played backwards.

APPLYING AN EFFECT TO A SECTION OF THE SOUND FILE

Select the *Undo Reverse* item from the Edit menu to put TUTOR1 back to its original state. Now select only the word "Wow" in the window and again select the *Reverse* item in the Process menu. Play the file again and notice how only the word "Wow" is played backwards. After you have listened to the effect, undo the reverse again so we can try another effect.

Again, let's select only the word "Wow" in the TUTOR1 window. This time after selecting the data, select the *Volume* item in the Process menu. Use the scale factor fader to set the scaling to -6 dB (50%). Clicking on the fader and using the up/down arrows will give you more precision. Now select *OK*. If you press *Play* you will notice that the volume is now at 1/2 of its original volume. You can undo the last operation by selecting *Undo* from the Edit menu.

The previous two examples show you how easy it is to apply effects to files. Sound Forge has a large number of different processes and effects, ranging from simple volume changes to complex multi-tap reverb/delays. To learn about a particular effect, please refer to the Effects, Process, or Tools menus that are covered in the Sound Forge Reference chapter.

CANCELING A FUNCTION IN PROGRESS

While applying an effect, the progress meter at the lower left-hand side of the main screen shows what percentage of a selection has been processed while running a function. You can select the *Cancel* button to stop processing the file. When you cancel an operation in progress, the affected data remains in the sound file. Undoing the operation will return the file to its original state.

APPLYING FUNCTIONS TO STEREO FILES

In a stereo file, only the selected region in the channel that you've selected is processed. Most functions can be applied to the left, right, or both channels. The only functions that cannot be applied to separate left and right channels are functions that affect the length of the data, since each track in a stereo file must be of the same length. These include: *Insert Silence*, *Resample*, *Time Compress/Expand*, *Gapper/Snipper*, *Pitch Bend*, and *Pitch Shift (without preserving duration)*.

If you need to run a process like *Insert Silence* in only one channel, you should divide the stereo file into two mono files first, insert the silence, and then join them back together into a stereo file. Separating stereo files into two mono files is a good idea if you are going to be changing the duration of each track by cutting or inserting space to synchronize different events.

APPLYING A FUNCTION TO A SINGLE CHANNEL IN A STEREO FILE

Previously, we reversed TUTOR1 so that we could hear how the data sounded when played backwards. Now let's try this on a single channel of a stereo file. First convert TUTOR1 to stereo by right-clicking on the status field which has the word *Mono* and select *Stereo* from the shortcut menu as we did in a previous section. Make sure to select the *Both Channels* option in the dialog so we have data in both channels of the new stereo file.

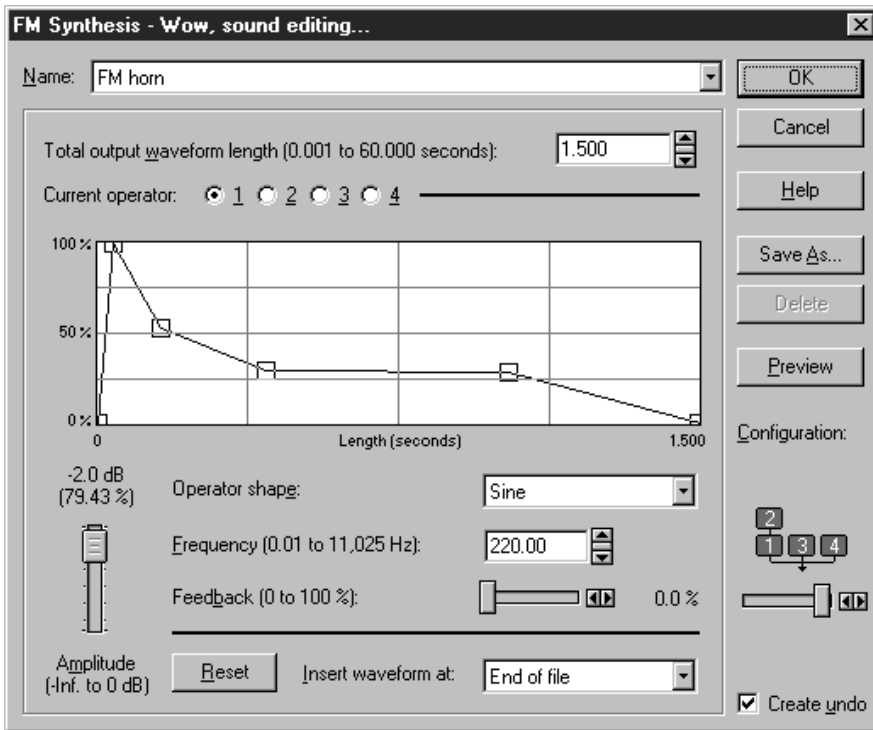
Select the entire left channel by triple-clicking with the mouse in the upper 1/4 of the TUTOR1 waveform display. If you select both channels, just use the *Tab* key to toggle the selection to only one channel. Now select the *Reverse* item from the *Process* menu. You will notice that only the left channel data is reversed. Press the *Go to Start* button on the playbar to clear the selection and place the cursor at the start of the file. Press the *Play* button and you will hear "Wow, Sound editing is easy" backwards in the left channel and forward in the right channel. If you only hear one channel then you are in single channel preview mode. To hear both channels use the *Tab* key to toggle the cursor or selection to encompass both channels and again press the *Play* button.

NOTE

If you want more information on how to use a function, press the F1 function key while selecting the function title in the menus. Also, you can press the F1 function key or the Help button in all function dialogs.


USING CONTROLS

Before going any further, let's become familiar with all of Sound Forge's different controls used to enter function parameters. To do this, open a new file and select the *Synthesis | FM* item from the Tools menu. Don't worry about understanding what all of the parameters mean. Let's just play with the knobs.



VERTICAL FADER AND HORIZONTAL TRACKBARS

In the *FM Synthesis* dialog, the *Amplitude* is controlled by a fader and the *Feedback* is controlled by a trackbar.



There is nothing complicated about these controls. To change the parameter values, just left-click and hold on top of the trackbar thumb and drag left and right or up and down.


NOTE

You can fine tune a control's value by holding down both the left and right mouse buttons at the same time (or holding down Control on the keyboard).

There are many keyboard shortcuts when using faders and trackbars. You can use the *Up*, *Down*, *Left* and *Right* arrow keys to change the value in small increments or the *Page Up* and *Page Down* keys to change the value in larger increments. The *Home* and *End* keys change the parameter value to its maximum or minimum.

If you double-click on a fader or trackbar thumb, it returns to its reset value (usually 0%, 50%, or 100%). Left-clicking on the hash marks in a fader also changes the value by very small increments.


EDIT BOX SPINNER CONTROL



In the FM Synthesis dialog the *Total* output waveform length and *Frequency* parameters are modified with edit box spinner controls. The edit box is the box containing text, while the spinner is the small control to the right. With this type of control, you have several options for changing a parameter value. You can:

- Type in the number by left-clicking on the edit box and then typing in the value.
- Increment the value by clicking on the two small up/down buttons.
- Use the spinner to change the value. This is done by left-clicking on the center button (between the two up/down buttons) and dragging the mouse up or down. Again, holding down both the left and right mouse buttons at the same time will cause the values to increment in finer steps. You can also use the *Up/Down* arrows and *Page Up/Page Down* keys to alter the value.
- Adjust the spinner while holding down both mouse buttons to see the finer increment resolution (or hold down the *Control* key).

DROP-DOWN LIST



To select an item from a drop-down list, click on the drop-down list and select the item. If you have to scroll through a large list, click on the scroll buttons or use the arrow keys. The FM Synthesis dialog uses a drop-down list to change presets, select waveforms, or change operators.

NOTE

If you scroll down the drop-down list for a preset, you can see all of the parameters change for each preset. This is useful for getting a feel for which parameters are used to create different effects.


PUSH BUTTON

To use a push button, left-click on it or press the Spacebar while it is selected. In the FM Synthesis dialog, the OK, Cancel, Help, Save As, Delete, Selection, Preview, and Reset controls are push buttons.

RADIO BUTTON


Radio buttons always come in groups of two or more. They are used like the radio station selector in a car radio; turn one on and the rest turn off. For example, the Current operator control in the FM Synthesis dialog is a group of radio buttons.

CHECK BOX



A check box, as its name implies, is a square box you can check or leave blank. They are used to turn a feature on or off. *Create undo* in the *FM Synthesis* dialog is a check box.

ENVELOPE GRAPH



Envelope graphs are used to draw the shape of a frequency or amplitude envelope that will be applied. The horizontal axis represents time, with the left-most point specifying the start of the selection and the right-most point specifying the end of the selection. The vertical axis represents either amplitude or frequency.

- To create a new knob on the graph, left-click anywhere on the line connecting the envelope points and drag the knob in any direction, sort of like pulling a rubber band.
- You can delete a knob by double-clicking or right-clicking within it.
- To reset the graph (delete all knobs) click on the *Reset* button.

MOVING MULTIPLE ENVELOPE POINTS

It's possible to select and move more than one envelope point at once. Just try:

1. Clicking on an unused section of the *Envelope Graph*.
2. While holding the mouse button down, drag the mouse. A dotted-line selection square will appear.
3. Move the mouse to surround the needed points with the selection square.
4. Release the mouse button. The selected points will now be colored red.
5. To move all of the selected points, drag any one of the points.

You can also select all points by pressing **Control+A** when the *Envelope Graph* is active.

DISPLAYING THE WAVEFORM ON AN ENVELOPE GRAPH

Some envelope graphs (such as in the *Fade | Graphic* or the *Pitch | Bend* dialogs) allow you to view the waveform of the selected data region on the graph. If the selection is small, the waveform is automatically displayed. Otherwise, you must press the *Show wave* check box. If you have a stereo file, you can choose between displaying the *Left channel only*, *Right channel only*, or *Mix channels* in the drop-down list.

PRESETS

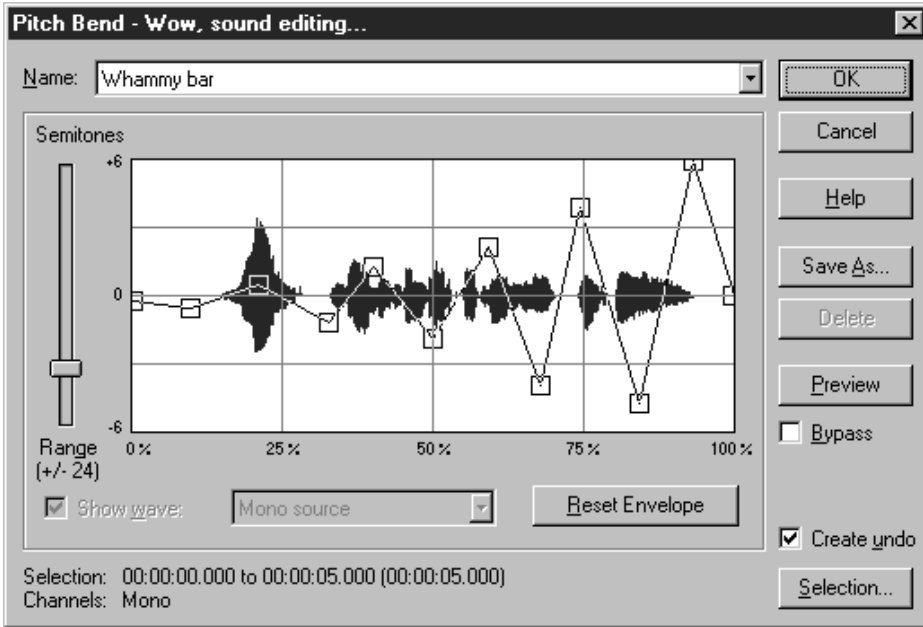
USING A PRESET

Most processing functions allow you to store the current settings for later use. There are also system presets for each effect that can be used without having to learn how to use the knobs and tweakers. These presets are a good place to start when you're creating your own effects.

Open the file *TUTOR1.WAV* again. Under the *Effects* menu, select *Noise Gate*. The *Noise Gate* function mutes any sounds below a threshold level. Using the drop-down list labeled *Name*, scroll down the presets list. Notice that the all of the controls change automatically. Now choose one of the presets and select **OK**.

NOTE

If you didn't have a selection when you ran the effect, the entire file was processed.



CREATING AND STORING YOUR OWN PRESET

Now open the file TUTOR2.WAV. Select a large portion of the drum hit and then select the item **Pitch Bend** from the Effects menu.

The Pitch Bend function changes the pitch of the wave variably over time by using a pitch bend envelope. Use the following steps to create your own pitch bend effect:

1. Press the **Reset Envelope** button. This restores the pitch bend envelope to its reset state.
2. Now raise the **Range** trackbar to 24 Semitones. This affects the range of the pitch change. In this setting, the maximum pitch bend is 24 semitones (two octaves).
3. To change the pitch bend envelope, left-click on a point on the envelope. You can move a point around by dragging it or you can remove it by double-clicking on it. Create an envelope similar to the one shown here.
4. Now, select **Save As** to store this preset. You can give this preset whatever name you'd like and then select **OK**.
5. Press **OK** again to run the function. You should be able to hear the drum hit change in pitch relative to the envelope you drew.

6. From now on, you have the option of using the effect you just created by selecting it from the drop-down list. If you decide you don't have use for this effect in the future, you can always delete the preset by selecting it in the drop-down list and selecting the *Delete* button.

NOTE

To reset all parameters in a dialog to their default setting, select *Reset All* from the shortcut menu which comes up when you right-click on any unused part of the dialog.

PREVIEWING AN OPERATION

You can preview how a function will affect a sound file by using the *Preview* button found in most sound processing dialogs. This is useful for fine-tuning parameters without having to leave the dialog.

To A/B test the sound (with and without processing), use the *Bypass* check box. When the *Bypass* box is checked, selecting the *Preview* button will play the unprocessed data. Otherwise, you will hear the processed sound.

NOTE

You can change the state of the *Bypass* check box while the preview is playing to quickly switch between A and B.

PREVIEW CONFIGURATIONS

You can temporarily select the length of data you wish to preview by right-clicking anywhere in the dialog and selecting *Configuration* from the shortcut menu. In the *Preview Configuration* dialog, enable the *Limit previews to* option and specify the number of seconds from 1 to 30.

You can set a permanent preview configuration by using the *Preview* folder in the *Sound Forge Preferences* section (or by selecting the *Default all previews to the current settings* in the *Preview Configuration* dialog).

You'll also see the *Reactive previewing* check box in the *Preview Configuration* dialog. When activated, this item automatically generates a preview that allows you to hear any change in dialog parameters. This function becomes very handy when tweaking the center frequency in the *Paragraphic EQ*, for instance.

NOTE

When Reactive Preview is active, holding down the Shift key while changing parameters will suspend the automatic preview generation until the Shift key is released.

Pre-roll and *Post-roll* times can be set in the *Preview Configuration* dialog. These determine the amount of material that will play before and/or after the selected area in a preview. The *Pre-roll* and *Post-roll* can be quickly turned on and off by right-clicking anywhere within a dialog and selecting from the shortcut dialog.

PREVIEWING A CUT

To preview a *Cut* operation, make a selection you wish to cut and select *Preview Cut/Cursor* from the *Edit* menu. *Sound Forge* will pre-roll and post-roll audio as specified in the *Other* page of the *Preferences* folder, skipping the selection.

If no selection is made, you can use this feature to listen to sound immediately surrounding the cursor.

PROCESSING SHORTCUTS

REPEATING AN OPERATION

To execute a function with the same parameters as the last time the function was used, hold the *Shift* key down while selecting the function. This is known as shift-clicking and works for most operations in *Sound Forge*, not just the effects. Shift-clicking works for toolbar buttons, as well as menu items.

You can also use the *Repeat* item (*Control+Y*) under the *Edit* menu to quickly repeat the last operation used. This works for all *Process*, *Effects*, and *Tools* functions in *Sound Forge*. Holding down the *Shift* key while selecting *Repeat* takes you to the last function's dialog.

TOOLBARS

If you like toolbars, there are toolbars for functions under the Process, Effects, and Tools menus. To use these toolbars, you must go in the *Preferences* folder (Options menu) to the *Toolbars* tab and turn on those you wish to appear on your screen.

CHANGING THE SELECTION AFFECTED BY A FUNCTION

To apply a function to a region other than the current selection, choose the *Selection* button in the function's dialog (you cannot do this with functions that do not have a dialog, such as the *Reverse* function). This causes the *Set Selection* dialog to appear.

In the *Set Selection* dialog, you can change the *Start*, *End*, *Length*, and *Channels* selections which will be affected by the function. You also have a choice of the format in which you are going to enter these values or you can even use a preset selection.

SOUND PROCESSING TECHNIQUES

In this section, we will cover some basic concepts relating to several of the Sound Forge processing functions used for modifying (or creating) sound. Detailed information about each effect is also available in the Reference chapter.

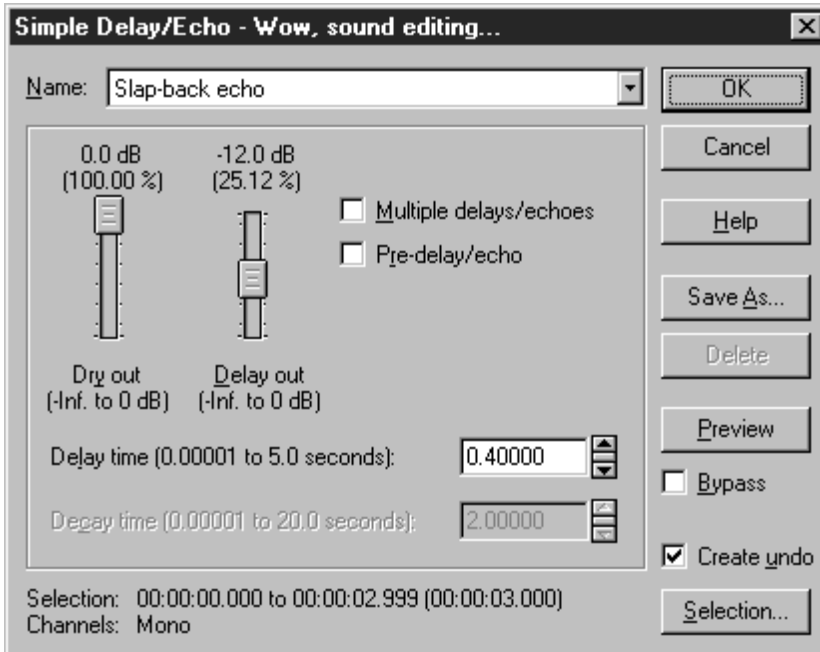
You will learn about:

- Delay/Echo
- Reverb
- Chorus
- Flange/Wah-Wah
- Dynamics
- Changing Pitch and Time Duration
- Synthesis
- Filtering

DELAY/ECHO

SIMPLE

Simple Delay/Echo creates copies of the original sound which are then mixed with the sound file to create simple echo effects.



Open TUTOR1.WAV and select *Simple* from *Delay/Echo* in the Effects menu. Choose the *Slap-back echo* preset and select *Preview*. A single copy of the sound is heard 0.4 seconds after the original. Changing the *Delay time* determines the time between the original and echoed sound.

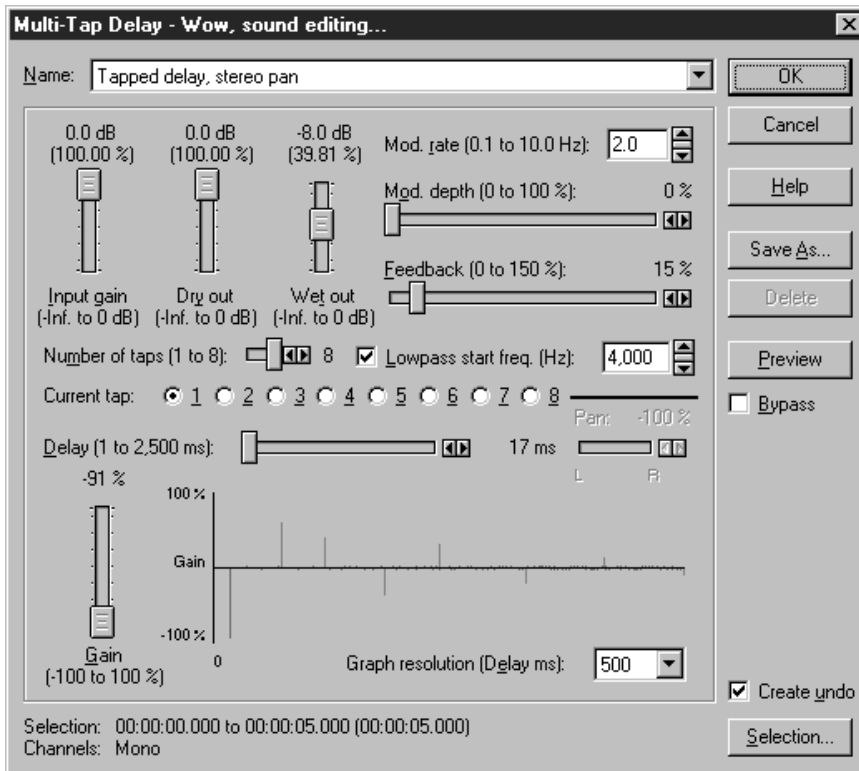
Now, select the *Grand Canyon* preset. Notice that the *Multiple delays/echoes* check box is marked. This means that instead of just one echo, you will hear multiple decaying echoes of the original sound. The *Decay time* determines how long it takes for these echoes to fade out.

The *Pre-delay/echo* function creates echoes heard before the original sound. Use this effect while the spaceship in your next sci-fi flick travels back in time by means of a wormhole.

MULTI-TAP DELAY

The **Multi-Tap Delay** function allows you to specify up to eight delay-taps spaced anywhere within 2.5 seconds (0.5 seconds in 16-bit Sound Forge) of the original sound.

Open the file TUTOR1.WAV and select **Multi-Tap** from **Delay/Echo** in the **Effects** menu.



Now, select **Tapped delay, stereo pan** from the end of the preset list. To hear the effect, select **Preview**. Selecting **Bypass** allows you to listen to the unprocessed file. Listen to how echoes bounce between left and right channels.

This preset uses all eight delay taps. To switch between taps, use the **Current Tap** function. There is a delay time and an amplitude associated with each delay tap which determine when and how loud the delay tap is heard after the original sound. In stereo files, each tap can also be panned left or right, to simulate sound coming from all directions.

The selected tap is represented as a red vertical line on the echogram. The echogram represents an approximation of the impulse response of the current settings and can be used to estimate the decay time. Blue lines correspond to other active taps and black lines are feedback echoes.

The horizontal axis of the echogram is time, with the right-most point equal to the **Graph resolution** scale setting. The vertical axis represents the amplitude of the echo. You can change the scale factor between 500, 1,000, or 5,000 milliseconds for viewing the impulse response over different time ranges.

By changing the **Feedback**, you can change the decay time of the delay.

If you set the Feedback too high, the echoes keep adding up until they are clipped (use the 5000 ms Graph resolution setting). Too much feedback also creates unnatural ringing of some frequencies.

NOTE

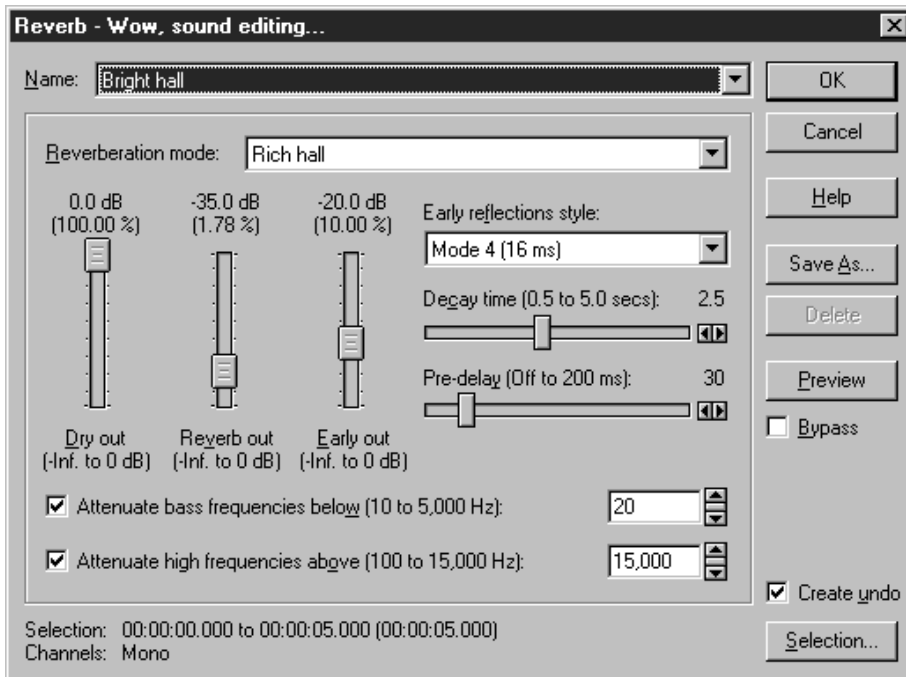
Select the **Default all parameters** preset. Here, there is only one tap active, with no feedback. This corresponds to a simple delay. Set the **Dry out** to zero and turn on the **Low-pass start frequency**. Low-pass filtering removes high frequencies, in the same way that many building materials absorb high-frequency sounds. Adding modulation creates slight pitch changes, which can make a delay sound fuller.

Now, switch to the **Chorus-reverse verb 1**. Here, the delay taps increase in magnitude over time, thereby creating a swelling effect. Many other sound effects are possible if you experiment a little with tap arrangements.

REVERB

The **Reverb** function is used to simulate the acoustics of different environments. When a sound is generated inside a room, you first hear the original sound, followed by a quick succession of echoes. These echoes are the early reflections. Then, as the echoes accumulate, you hear of thousands of echoes combined into a smoothly decaying reverberation.

Open TUTOR1.WAV and select **Reverb** from the Effects menu. Select the **Bright hall** preset and hit **Preview**. Notice how the imagined space where the recording took place changes from a small studio to a large, rich hall. You can adjust the loudness of the reverberation and early reflection by using the **Reverb out** and **Early out** faders.



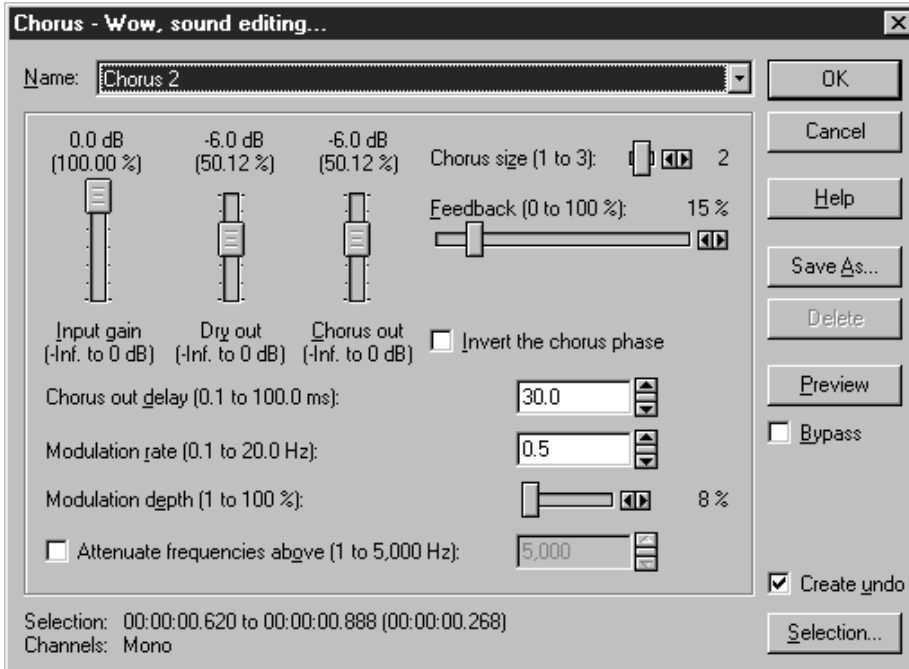
To increase the reverberation length, try changing the **Decay time** from 2.5 to 5.0 seconds. The **Pre-delay** parameter specifies the time between the original sound and the start of the reverberation.

Now, change to the **Sewer** preset. This preset uses a different **Reverberation mode**. Notice how its tone quality is very different – very metallic. The reverberation modes provided give you many choices.

CHORUS

The **Chorus** function is used to simulate multiple sound sources from a single sound. This is achieved by mixing a delayed, pitch-modulated copy of itself to the sound source.

With TUTOR1.WAV open, select the **Chorus** item from the Effects menu.



Preview the **Chorus 2** effect to hear an example of chorusing. Increase the depth to 50% to have a slowly de-tuning effect. The rate changes how fast the de-tuning occurs.

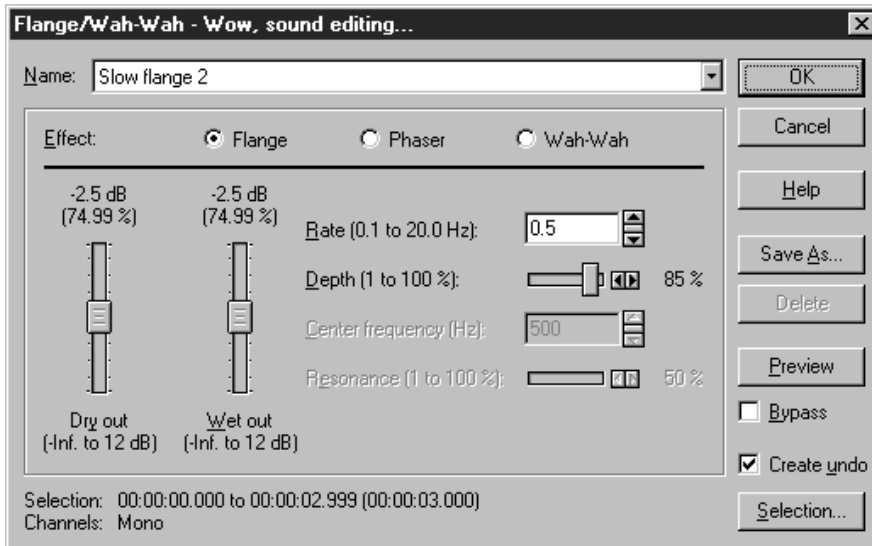
Changing the **Chorus out delay** time affects how much later the modulated sound is heard. Very small values usually create interesting filtering effects, such as flanging. With larger values, chorused echoes are created.

The **Feedback** and **Chorus size** parameters change the strength of the chorus. Large **Feedback** or **Chorus size** values intensify the chorus. You can use the **Low-pass filter** in to remove unwanted high frequencies.

FLANGE/WAH-WAH

The *Flange/Wah-Wah* function is used to create the sweeping effects often heard in 60's guitar recordings and techno-sounds of today. With this function, you can recreate the sounds that flanger, wah-wah, and phaser pedals are famous for.

To hear what flanging sounds like, open TUTOR1.WAV and select *Flange/Wah-Wah* from the Effects menu.



Preview the *Slow flange 2* preset. The *Rate* and *Depth* parameters are similar to the ones in the *Chorus* effect and can be used to create faster, slower, or stronger flanging.

DYNAMICS

Dynamics processing is used to change volume levels within a sound file. Applying dynamics specifically adjusts the lower levels and higher levels in the file to create more of a dynamic range, less of a dynamic range, or to keep the amplitude of a file within a certain range.

NOISE GATING

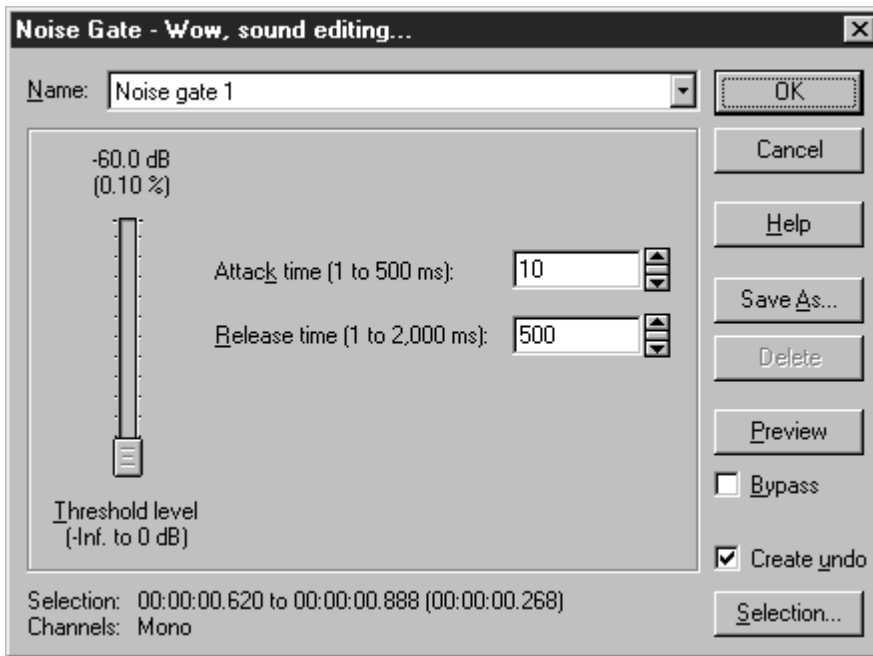
When recording a sound, there is often an audible noise floor during silent breaks. Noise is generated by many different things, including electrical equipment, machinery, and traffic outside your window. When your sound source is much louder than this background noise, it is simple to remove the noise during silent breaks, where the noise is most noticeable, with a noise gate.

For example, open TUTOR1.WAV. Move the cursor to a location between “Wow” and “Sound” and select **Insert Silence** (Process menu). Specify 1 second of silence to be inserted at the cursor. Next, create a new file. In the new file, use **Synthesis | Simple** (Tools menu) to generate four seconds of noise at -60 dB amplitude (keep the frequency around 2,000 Hz). Next, copy all of the noise on to the clipboard and mix it to the beginning of TUTOR1. Use 100% for the noise (*Source*) and 100% for the voice (*Destination*). You should now have a file with voice and a relatively low hissing sound in the background.

To remove the noise during silent breaks, first you should analyze the amplitude of the noise. In TUTOR1, select a region with only noise (no voice), and run the **Statistics** function under the Tools menu. Take note of the **Maximum sample value percentage**. If you followed the preceding steps, it should be close to -60 dB.

Next, select **Noise Gate** from the Effects menu.

In the **Noise Gate** dialog, select the **Noise gate 1** preset. Slide the **Threshold level** to a bit over the **Maximum sample value** (in percentage). A value of -50 dB should be high enough to differentiate the noise from the vocals. Press **OK**. Notice that the noise during the silent regions should have disappeared (total attenuation). If it didn't, try raising the **Threshold level** until it does.



The other two parameters in the Noise Gate determine how fast the gate will open and close. This will affect how the beginning and end of the sound will fade in and out. If the *Attack* time is too slow, the beginning of the vocals might get cut off. Likewise, if the *Release* time is too fast, the end of the vocals will be cut. Experiment with these times until you can remove as much noise as possible without removing any of the voice.

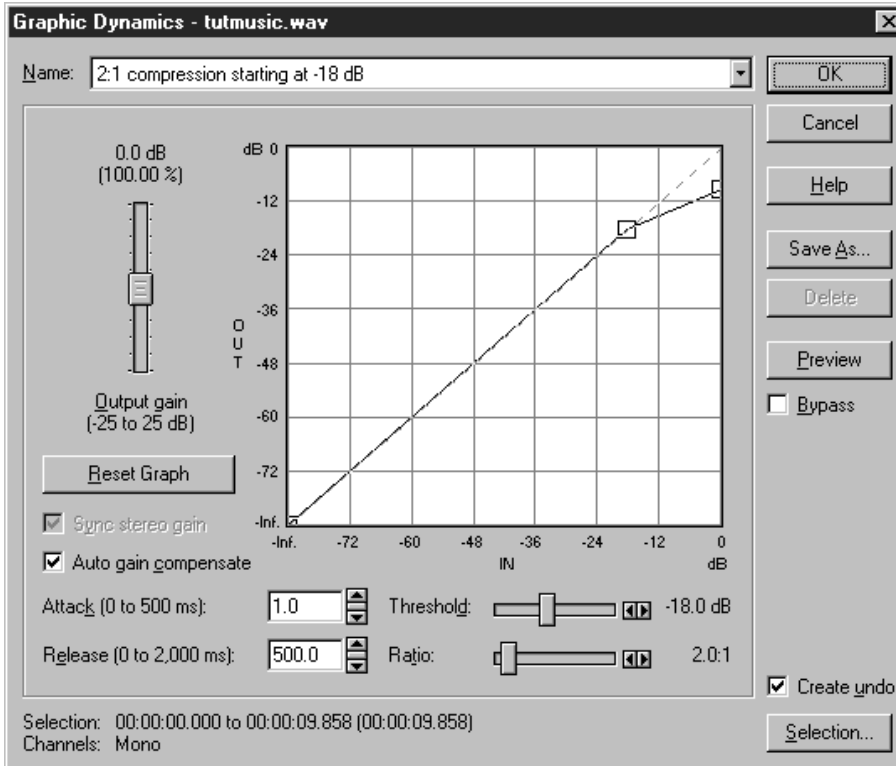
COMPRESSION AND LIMITING

Compressing and limiting are terms used to indicate effects that lower the dynamic range of a sound. When you compress a sound, you lower the volume of loud sections and then raise the overall volume to compensate. This is done to keep the volume level from fluctuating too much over time. Limiting works exactly like compression, but to a higher degree.

NOTE

Compression and Normalization can be used together to increase the apparent loudness of a file. This is currently a popular mastering technique.

Open the file TUTMUSIC.WAV. Make a copy of it by selecting the whole region and dragging it to an open space in the Sound Forge workspace. Now, with the new file active, select **Dynamics** from the Effects menu and choose **Graphic**.



The **Graphic Dynamics** dialog allows you to precisely tailor the gain at all input levels of a signal. The **Input vs. Output level** graph shows you the gain that will be applied to the input signal, depending on its level over time. The diagonal line, referred to as the **No Gain** line, indicates where the input and output levels are equal. When an envelope point is below the line, gain reduction occurs.

If you understand the principle of the graph, you can edit its points directly to create your own compression, expansion, limiting, or noise gating curves (see the Reference chapter for more details). However, you don't need to use the graph at all. Using presets and the threshold and ratio trackbars is sufficient for most cases.

Choose the **2:1 compression starting at -18 dB** preset, turn *Auto gain compensate* off, and select **OK**. If you can see both files at once and are zoomed out, you'll be able to notice how the levels of the new file are more constant than the original. Listening should reveal the same thing.

In this preset, sounds above the **Threshold** (-18 dB) are attenuated at an input to output ratio of 2:1. Sounds below the **Threshold** are not affected by the compressor. The *Auto gain compensate* check box raises the total volume of the file after processing to bring the overall output level back to an acceptable level. Turn it on once and you'll see how the overall level is much higher now. This can also be done manually with the **Output gain** fader.

Undo the previous operation and select **Dynamics** again. Now, increase the **Ratio** to 7:1 and select **OK**. The sound will be even more compressed, meaning that more attenuation and boosting will be applied to compressed signals.

NOTE

Too much compression starting at too low of a threshold level will usually lead to a distorted sound.

Again, undo the operation and select the **Limit levels to -6 dB (hard limiter)** preset. Now, signals above the **Threshold** (-6 dB) will still be attenuated, and with a huge ratio of **Infinity:1**. Essentially, after processing, there will be no levels above -6 dB included.

Limiting is often used to keep signals from going above a certain level, but can also be applied to create heavily-compressed effects. Limiting should only be done on peaks; if the **Threshold** level is set too low, heavy distortions will occur.

EXPANSION

Dynamic expansion is the opposite of compression and limiting. Any signal below the threshold will be attenuated by a certain ratio. The most common effect of expansion is to attenuate low-level noise, such as a noise gate. However, you can also use expansion to add more dynamic range to a sound.

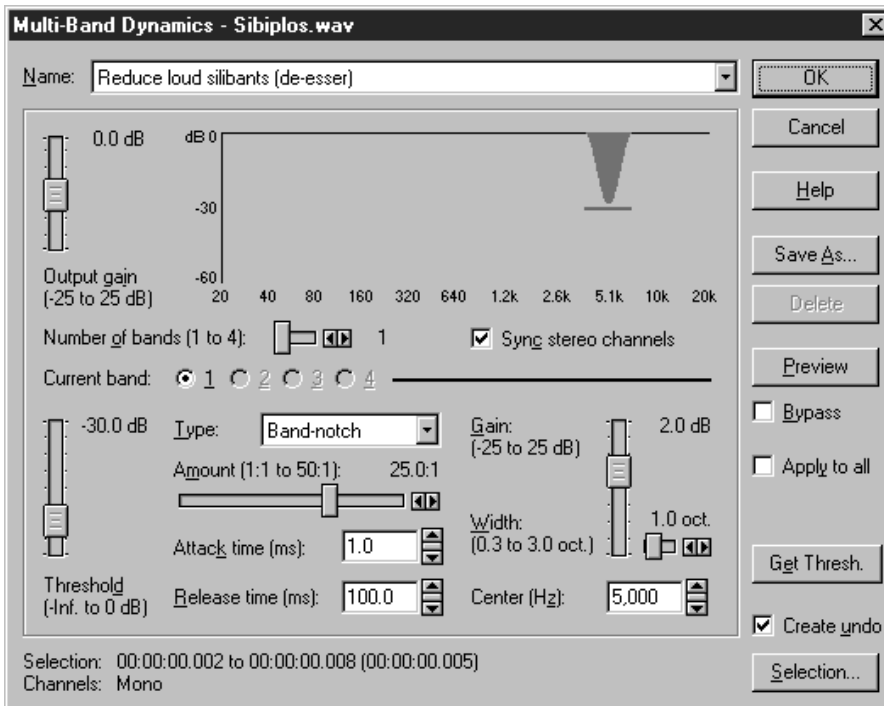
With the **TUTOR1.WAV** file open, select **Dynamics | Graphic** from the **Effects** menu. Choose the **Soft noise gate below -36 dB** preset and select **OK**. Notice how very low-level sounds have been attenuated. This is useful for removing noise from silent regions in a similar fashion as a noise gate.

MULTI-BAND DYNAMICS

The other type of dynamics processing that Sound Forge allows for is **Multi-Band Dynamics**. Essentially this is a compressor that only affects certain frequency bands. The possibilities with this type of compressor are endless, but the most common use is for de-essing and de-popping.

NOTE

You can also use this function as an advanced EQ for such things as tightening up a bass line or drums.



Open the file SIBIPLOS.WAV. This file contains both loud sibilants (S's) and plosives (P's). Let's first remove the sibilants. Double-click on the region named *...and sibilants* to select it. Next, select **Dynamics | Multi-Band** from the Effects menu and choose the *Reduce loud sibilants (de-esser)* preset. This preset acts upon frequencies around 5,000 Hz. Previewing the effect should make the "S" sounds less painful.

Next, make a selection on the region named plosives. This can be done by pressing the **Selection** button and choosing **Region: ...and sibilants**. In the **Multi-Band Dynamics** dialog, choose the **Reduce loud plosives** and hit **Preview**. The loud “P” becomes less obtrusive and unlike EQ, this reduction doesn’t affect the entire sound file. It only acts upon the plosives.

NOTE

There is also a Multi-Band preset called Reduce plosives and sibilants which you can use to remove both S’s and P’s at once. This is basically the same effect as running one first and then the other.

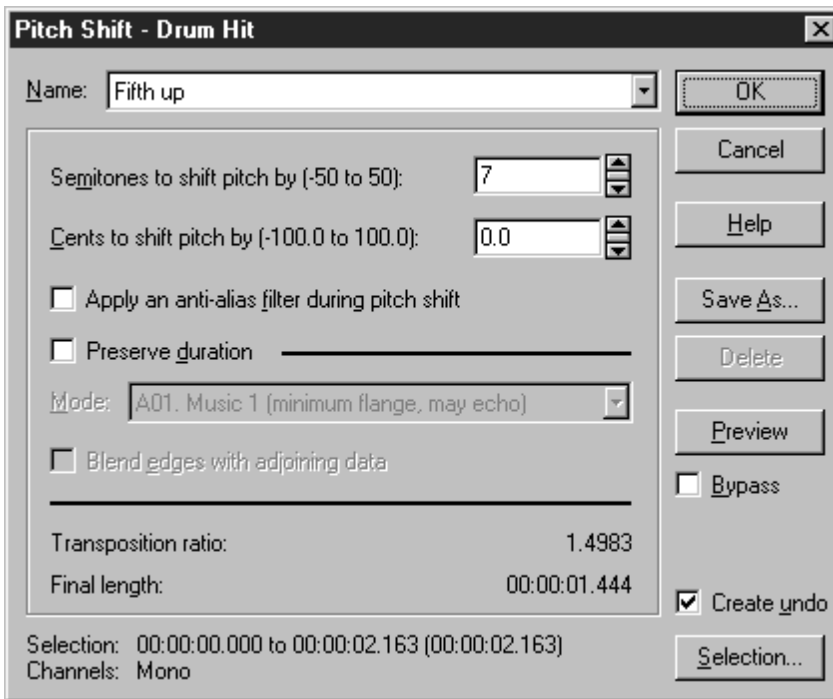
When applying these presets to your own files, you will most likely need to change the **Threshold** setting until only the unwanted sounds are affected. If the **Threshold** is too low, too much filtering will occur. Setting the filter to **-Inf.** forces filtering of all the data. Also, you can also boost a band after compression by using the **Gain** control to compensate for any loss during processing.

CHANGING PITCH AND TIME DURATION

ALTERING PITCH BY CHANGING PLAYBACK RATE

You are probably familiar with how changing the playback rate of a recording affects its pitch. For example, playing a 33 1/3 RPM vinyl record at 78 RPM makes the Beatles sound like the Chipmunks. Likewise, playing a 78 RPM record at 33 1/3 makes a trumpet sound like a tuba. This concept is used by most sampler units to make a single sound file achieve different pitches.

Open the file **SAXRIFF.WAV**. Under the **Effects** menu, select **Pitch | Shift**.



Next, select the *Fifth up* preset and select **OK**. Notice that the length of the file is less than its original size. Listen to the sound to hear both a higher pitch and a faster rate.

Undo the last operation and return to the *Pitch Shift* dialog. Now, select *Fifth down* and select **OK**. The file is now larger, slower, and lower in pitch.

ALTERING PITCH WITHOUT CHANGING THE DURATION OF THE SOUND

Undo the last operation and return to the *Pitch Shift* dialog. In the last two examples, the *Preserve duration* check box was left unchecked. When you check this box, a completely different algorithm is used to change the pitch of the sound.

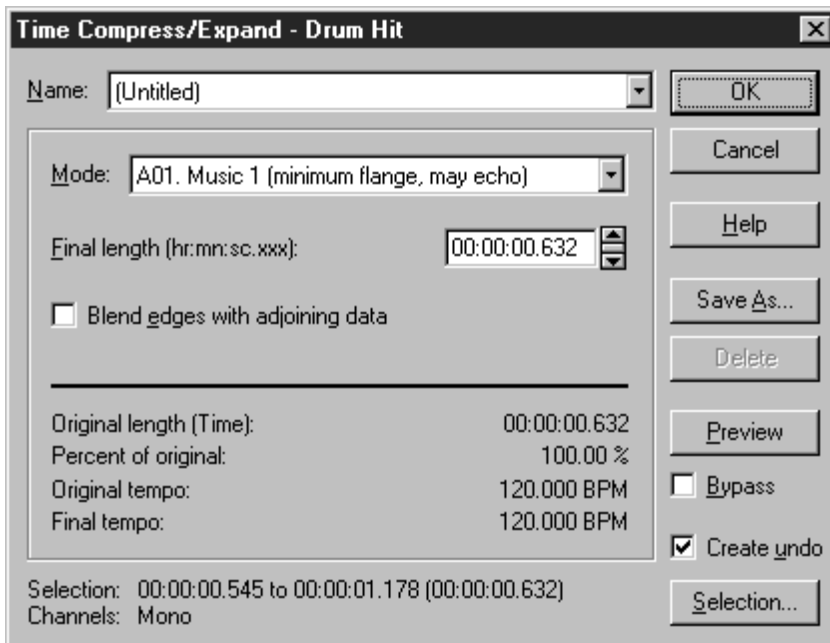
Select the **Third down** preset and check the **Preserve duration** check box. Next, you must select a mode. Different algorithms are included for different kinds of sounds. These modes are only suggestions, since many will work with different types of sound. It's best to just try a few different ones and choose the one that sounds best to you. For this saxophone riff, the mode named **A13. Solo instruments 4** does an excellent job.

Now, preview the sound. The sound file will now have a higher pitch, but it will not sound as if it's being played back faster.

CHANGING TIME DURATION WITHOUT CHANGING THE PITCH OF A SOUND

Another Sound Forge process unique to digital processing is the ability to stretch or compress the time duration of a sound without altering the pitch. This is useful for lengthening or shortening sounds to meet a specific time length.

Open **SAXRIFF.WAV** and select **Time Compress/Expand** from the Process menu.



This process is very simple to use. In the dialog, you have the ability to specify the final duration of a sound, from 50% to 500% of the original length.

Like in the *Pitch Shift* dialog, you must choose a mode from a list depending on the type of material you're processing. The mode named *A13. Solo instruments 4 (less echo)* does a good job with this sound file. Change the *Final length* from 2.5 to 2.0 seconds (80% of original) and select *OK*. Press the *Play* button and you should hear the riff played at a faster tempo.

The A and B modes used in both *Time Compression/Expand* and *Pitch Shift* without change of duration work in a similar fashion. Modes labeled C are from Sound Forge 3.0. Basically, the reason why there are so many options is that compromises must always be made between having excellent time localization (no echoes) and accurate frequency reproduction (no low-frequency distortion or flanging).

If you have sounds that have few loud, sustaining low-bass frequencies, you can use modes that will result in no echoes. These include the *Drum* modes and some of the *Solo Instrument* modes. For complex music, you must sacrifice some time localization for the ability to reproduce long bass notes.

Since classifying sounds into strict categories is not possible, different settings are included for each sound type. For example, if you have a drum recording with bass drum, snare, and cymbals, you find the *A17. Drums, unpitched (minimum echo)* does an impeccable job. However, once you start adding ringing low toms and congas, the *A18. Drums (better for toms)* preset will work better. Again, experimentation is the key.

SMALL AND ACCURATE CHANGES IN TIME DURATION

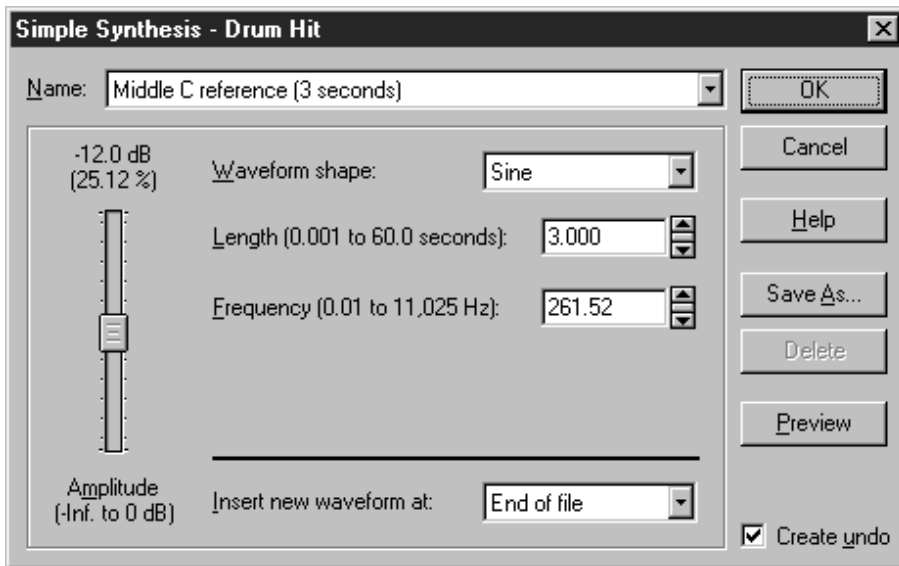
If you have a sound file that is one minute and three seconds long, but you need it to be exactly one minute, there are two ways to go about it. If a very small change in pitch is acceptable, you might try using the *Pitch Shift* with no *Preserve duration*. Keep the *Semitones* at 0 and adjust the *Cents* until the required *Final length* is reached. It is sometimes possible for this method to give better sounding results than when using the *Time Compression* algorithm.

SYNTHESIS

SIMPLE SYNTHESIS

The Simple Synthesis tool can be used to generate a simple waveform of a given shape, pitch, and length. More complex waveforms can be generated with the *FM Synthesis* tool.

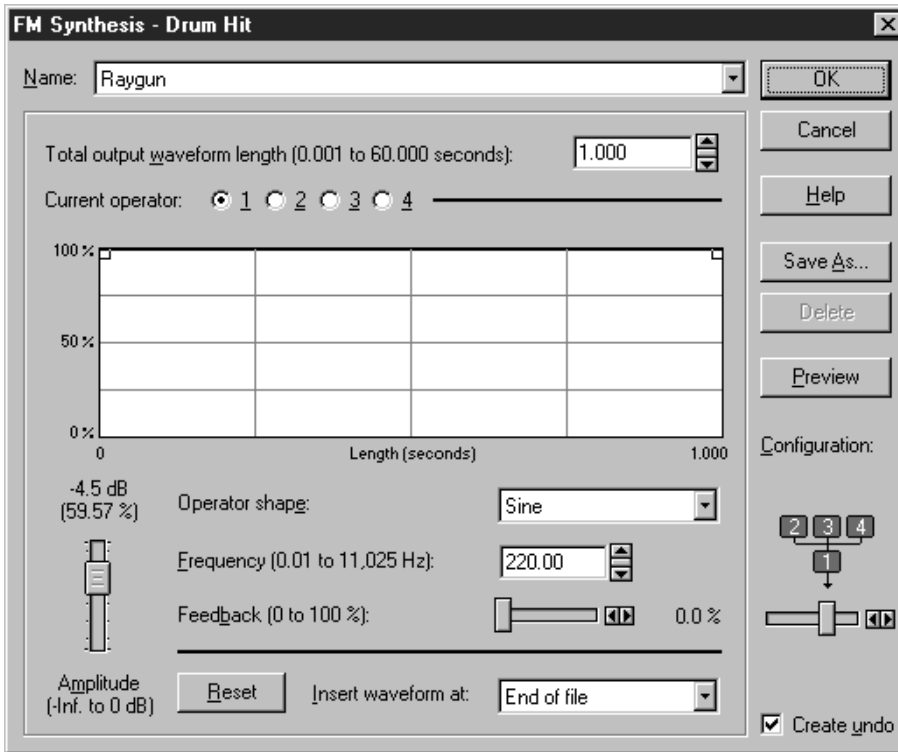
Create a new window by selecting the *New* item from the *File* menu and set the data format for the window to be *16-bit, 22,050 Hz, Mono*. Next select *Synthesis | Simple* from the *Tools* menu.



Select the *Middle C Reference (3 seconds)* preset and press *OK* to generate a sine wave in the window that is 3 seconds long and has a pitch of 261.52 Hz (Middle C). Press the *Play* button to hear your reference tone.

FM SYNTHESIS

Sound Forge's *FM Synthesis* tool allows you to use frequency modulation (FM) and additive synthesis to create complex sounds from simple waveforms.



In frequency modulation, the frequency of a waveform (the carrier) is modulated by the output of another waveform (the modulator) to create a new waveform. If the frequency of the modulator is low, the carrier will be slowly de-tuned over time. However, if the frequency of the modulator is high, the carrier will be modulated so fast that many additional frequencies, or sidebands, are created.

Sound Forge allows up to four waveforms (operators) to be generated in a variety of configurations. Depending on the configuration, an operator can be a carrier, a modulator, or a simple unmodulated waveform. Waveforms can also be added together (additive synthesis) to add more complexity to the sound.

The best way to become familiar with this module is to start with a single operator configuration and try out all the different waveforms by themselves. Experiment using *Feedback* on each waveform to hear the different effects of self-modulation. *Feedback* adds overtones to the sound by modulating the sound with itself. Zero feedback indicates simple synthesis.

In the second configuration, two unmodulated operators are mixed together (horizontal connection) and heard simultaneously (additive synthesis).

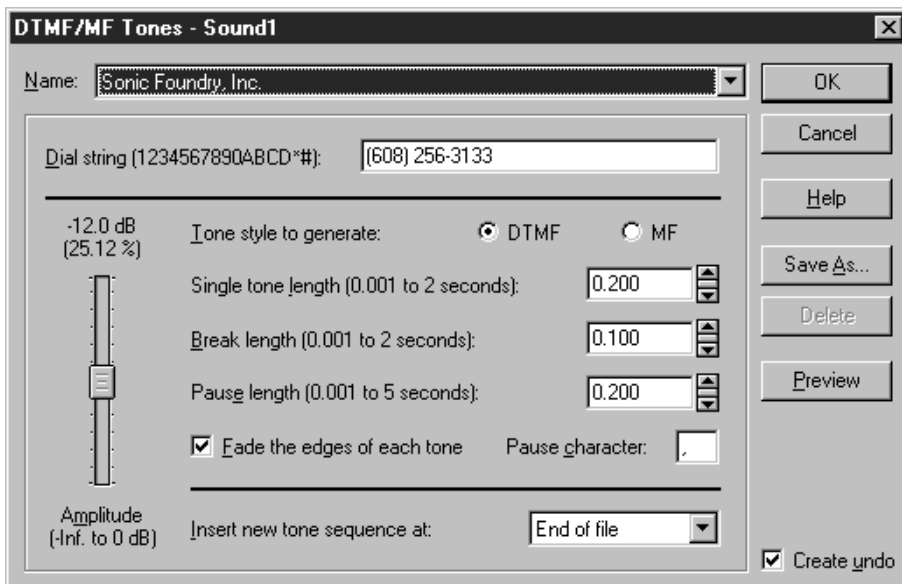
Now change the *Configuration* to two stacked operators (vertical connection) and again experiment using different waveforms and frequencies on the carrier (bottom operator) and the modulator (top operator). With the modulator *Frequency* set low (1-5 Hz), lower the modulator *Amplitude* to create slight de-tuning. Raising it creates big pitch bends. If the modulator *Frequency* is set high, many unusual FM sounds can be achieved. Make sure that the modulator *Amplitude* is not too high, otherwise the result will be harsh, noise-like sounds.

Setting the modulator *Frequency* to 0.00 allows absolute control of the carrier's pitch by using the modulator's envelope, which you can use for de-tuning over time.

Adding more operators increases the complexity of the waveform. When both FM and additive synthesis are combined, an almost endless variety of sounds can be generated.

DTMF SYNTHESIS

The *DTMF Synthesis* tool is used to generate the dial tones used by telephone companies. These dial tones correspond to the numbers in a telephone unit, along with other special codes.



For example, select the *Sonic Foundry, Inc.* preset. The dial string shows Sonic Foundry's phone number. The *Amplitude* control determines how loud the tones will be. Press *OK* to generate the tones. Now when you play back the sound file, you'll hear the tones which, when played over a phone line, would trigger a call to Sonic Foundry.

FILTERING

What can you do when you have a sound that is not quite perfect, but not bad enough to throw away? Sometimes, judicious use of filtering can keep a favorite sound bite from losing favor. Filtering is not a cure-all for bad sound, but slight alterations of a sound can bring it back to life.

Sounds are composed of varying amounts of one or more frequencies. For example, a sound with a rich timbre will contain many different frequencies. A sine wave sound has only one frequency in it, like 60 or 440 Hz.

Filters allow the tailoring of the frequency spectrum of a sound. A spectrum is nothing more than a representation of how much of each frequency component (from 20 Hz up to one-half the sample rate) is present in a signal. Less of any given component means that the given frequency component is not as prevalent aurally as another component. Filters pass or reduce (attenuate) frequency components.

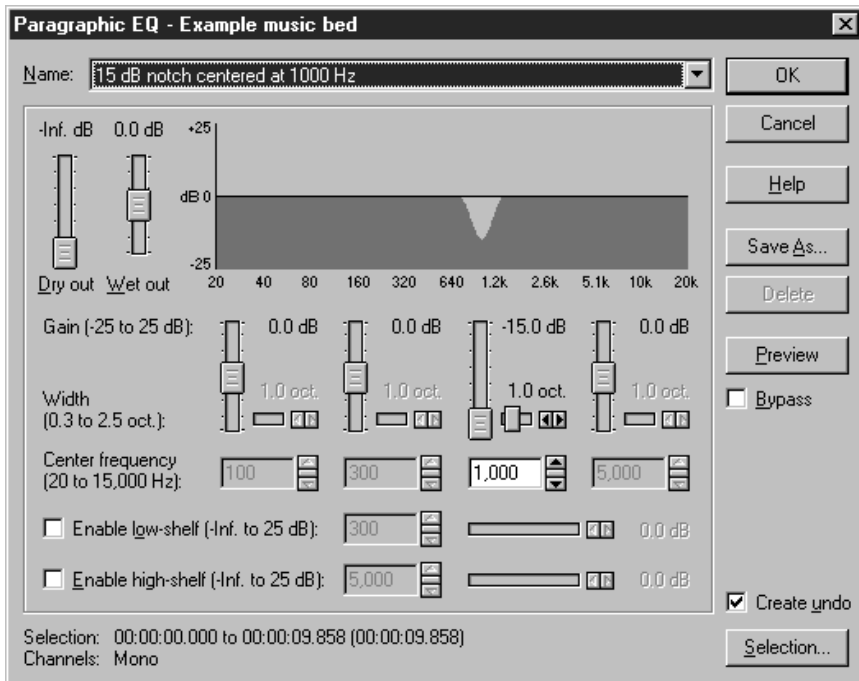
Many factors, such as the frequency response characteristic of your speakers and sound card, can also affect which frequencies are more prominent during the playback of a sound. In general, the goal is to make a recording sound as close to the original as possible. However filtering is also commonly used to remove unwanted sounds such as noise, or make individual instruments in a recording sound louder.

PARAGRAPHIC EQ

The *Paragraphic EQ* is a set of six very flexible parametric filters. Four independent band filters allow you to boost or attenuate specific frequency ranges. Also, two shelving filters let you control the amount of low and high frequencies in your recordings. A *Gain vs. Frequency* graph shows the overall effect of the combined filters, making it easier to visualize the final sound.

Open TUTMUSIC.WAV and select EQ | Paragraphic from the Process menu. Next, select the 15 dB notch centered at 1000 Hz preset. Notice how the graph goes below the centerline at 1000 Hz.

Hit Preview and listen to the very strange, hollow sound caused by the missing mid-range frequencies. Moving the Gain fader of the third filter further down from -15 dB to -25 dB will increase the effect. Next, move the fader - 25 dB to +9 dB and listen to the drastic tone difference.

**NOTE**

Double-clicking on a band's Gain fader resets its value to 0 dB, which disables the band. To disable all bands, right-click on an unused portion of the dialog and select Reset All from the shortcut menu.

Each of the four bands has a *Center frequency*, *Gain*, and *Width* control. The *Width* control determines how many octaves above and below the center frequency will be affected by the filtering. For the most part, try to keep the width around one octave. If the width is small, fewer frequencies will be affected.

Select the *Boost bass frequencies below 250 Hz by 4 dB* preset and hit *Preview*. Here, the low-shelf filter is used to make low frequencies louder. Change the gain from +4 to -15 dB and check out how most of the bass has been drastically reduced, as if being played through a small radio.

Open the file CREEK.WAV. This file contains a recording of a small stream. If you look at the waveform, you'll notice that there are some very high level, low frequency waves in this recording, probably caused by the wind. They are inaudible, yet if you look at your speaker's woofers during playback, you'll see them bouncing up and down. This large excursion can damage some speakers and forces both the amplifier and woofer to work overtime. *Woof!*

To remove these frequencies, go to the *EQ | Paragraphic* dialog and select the preset titled *Remove very low and inaudible frequencies below 20 Hz*. Now hit *Preview*. Turning the *Bypass* check box on and off, you should notice no audible difference. Hit *OK*. Notice how the large waves are gone and your woofers are no longer crying for help.

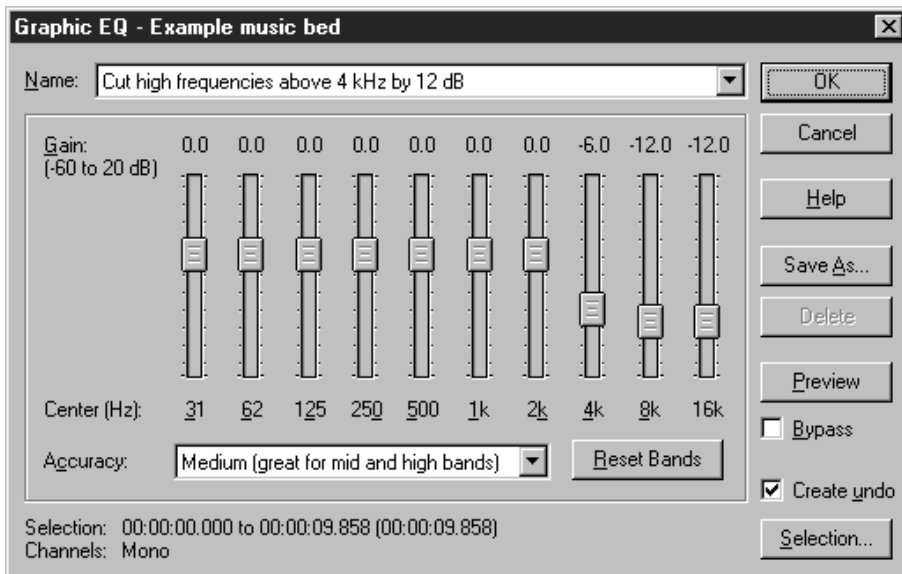
Removing very low and inaudible frequencies after recording is a very good idea. For one, your waveform will look more natural and the level meters more accurate. Also, edits and effects will behave more predictably, giving you extra headroom.

N O T E

GRAPHIC EQ

The *Graphic EQ* divides all the possible frequencies into ten bands which you can boost or attenuate (cut). Each band has a *Center* frequency related to it. The *125 Hz band*, for example, affects frequencies between 90 and 190 Hz. The *Gain* value above each band indicates the amount of cut or boost applied to the band. When a band is set to zero, it means that the frequencies in the band will not be modified. Positive gain values indicate a boost and negative values indicate a cut in amplitude.

Open *TUTMUSIC.WAV* and select *EQ | Graphic* from the *Process* menu.



Select the *Cut high frequencies above 4 kHz by 12 dB* preset and press *OK* to remove the high frequencies from the music bed file. Press the *Play* button to hear how the music is much duller.

PARAMETRIC EQ

Sound Forge has four basic types of *Parametric Filter* styles: *High-pass*, *Low-pass*, *Band-pass (peak)*, and *Band-reject (notch)*. A low-pass filter alters frequencies above a certain value, and a high-pass filter alters frequencies below a certain value. A band-pass filter alters frequencies outside of a specified range. A band-reject (or notch filter) alters frequencies inside a specified range. Generally, each filter attenuates frequencies in its stopband range(s) and leaves frequencies in its passband range alone.

The rate at which the passband transitions into the stopband is the transition band (or rolloff). When you specify the characteristics of a filter in Sound Forge, you will no doubt be tempted to specify that a low-pass filter will start to transition at, say, 1,000 Hz and finish transitioning by 1,001 Hz. The problem with doing so is that there are mathematical limitations to creating such a filter. When you specify a sharp transition band, the software will try to design a filter that will take a long time to apply and will unnecessarily filter the signal to overly exacting requirements. Razor-sharp filtering can often create unwanted side-effects. The more subtly you filter a signal, the more natural it sounds.

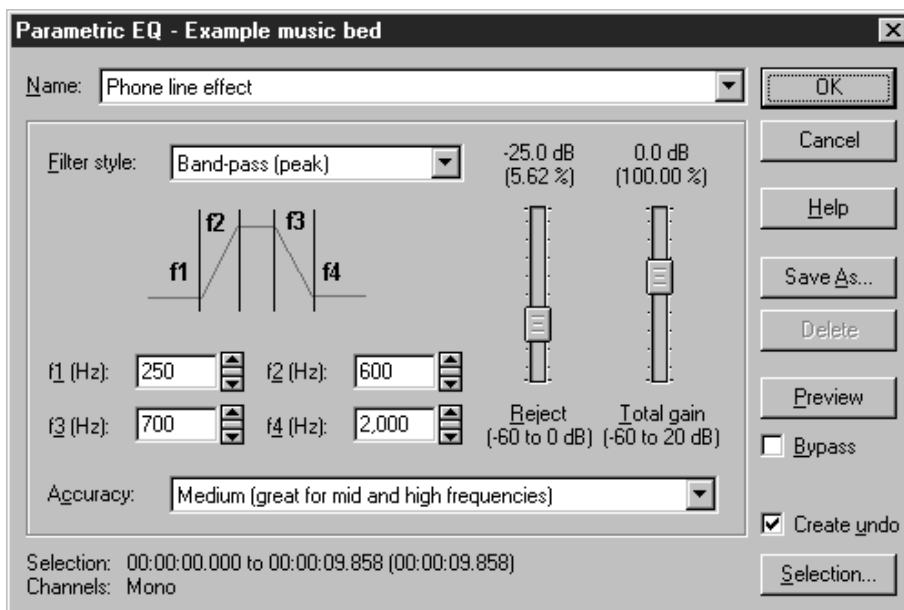
When specifying the attenuation amount for your filter, you are selecting how much to de-emphasize the frequency components that you don't want to hear. Numbers that are more negative suppress the frequency components more heavily. In other words, -6 dB of stopband rejection is going to let more unwanted frequencies through than -15 dB will.

So, let's say you just recorded a great hamster sneeze, but there is something not quite right about it when you play it back later. What can you do to restore the hamster sneeze to its rightful place as your favorite sound in your Windows environment?

Maybe you used a microphone preamp with bad common-mode rejection characteristics, or you set the hamster right next to a lamp when you recorded the sneeze? Well, don't panic, because Sound Forge can fix that sound. Try applying a 60 Hz notch filter to the sound. A notch filter will notch-out the frequencies in a range about 60 Hz so that they are not as readily heard relative to other frequencies.

Maybe the hamster sneeze has too much hiss in it? Try applying a low-pass filter and see if that makes it better. Remember, a small amount of filtering goes a long way.

Open TUTMUSIC.WAV and select EQ | Parametric from the Process menu.



Select the *Phone-line effect* preset and press OK to run the filter. Press the Play button to hear how the music is much thinner sounding, quite like hearing it over the phone. The *Band-pass* filter has attenuated frequencies below 250 and above 7000 Hz.

THE PLAYLIST AND REGIONS LIST

In this section, we will go over Sound Forge's ability to create regions and markers in a data window and how to use the *Playlist* to arrange them. Each data window can have its own *Regions List* and *Playlist* stored along with the sound data when saved in the .WAV file format or saved to an external playlist file.

We will cover:

- Why to Use the Regions List and Playlist
- Creating Regions
- Editing a Region
- Adding Regions to the Playlist
- Arranging the Playlist Regions
- Treating the Playlist as a Cutlist
- Using Markers

WHY TO USE THE REGIONS LIST AND PLAYLIST

FAST NAVIGATION

The most basic use of the *Regions List* is for dividing a sound file into separate, named regions. These regions can then be quickly played or highlighted in the data window. You can also tag important time positions with markers so that you can get around large files faster.

DESTRUCTIVE VS. NON-DESTRUCTIVE EDITING

When you perform edits on sound files using cut and paste operations, Sound Forge must shift the sound data on your hard drive. When working with large files, this process can take longer than your patience may be able to handle. However, when you create regions in the *Regions List* and then rearrange them in the *Playlist*, only the playback order changes (as opposed to the hard-disk storage). This is called non-destructive editing since no large chunks of data are moved or deleted in your original file.

With non-destructive edits, as opposed to hard-disk intensive destructive edits, you can listen to your changes right away. Once you're done arranging the regions, it's also possible with Sound Forge to create a new file containing your final arrangement.

MIDI SYNCHRONIZATION AND TRIGGERING

The regions you create can also be triggered using MIDI or SMPTE time code. This is used for synchronizing sound files to your sequencer, MIDI controller, or any other time-based media. For example, you can assign a MIDI trigger to a voice sound file in Sound Forge and then have your sequencer trigger it along with other MIDI instruments. Another use is to assign SMPTE times to special effects sound files to match the visual action on a soundtrack. Refer to the MIDI/SMPTE chapter for more information on synchronization.

CREATING REGIONS

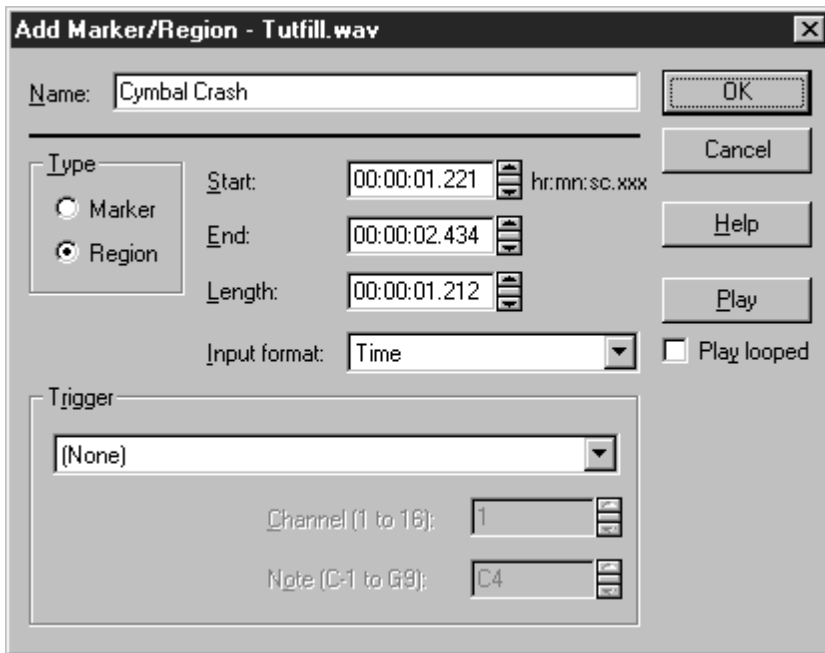
This section will describe how to create regions within a sound file.

CREATING A REGION USING MENU COMMANDS

As with most features in Sound Forge, you can create a region in your sound file in a number of different ways. One of the simplest ways is to drag a selection from a data window into the *Regions List*.

Let's create a region in the file *TUTFILL.WAV*. Open the file as explained earlier in the manual and then open the *Regions List* window. This window is opened by selecting *Regions List* from the View menu or by pressing the *Regions List* button on the *Regions/Playlist* toolbar. You can also open this window by pressing *Alt+1*.

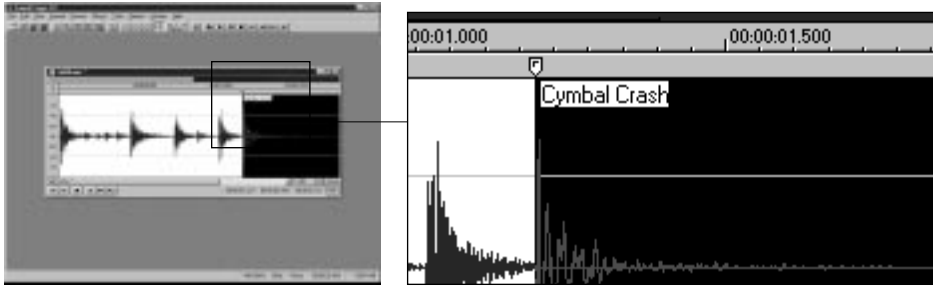
To create a region, listen to the file by pressing *Play* and then select the cymbal crash at the end of the file. Next, select the *Add* command from the *Regions List* shortcut menu in the Special menu. The *Add Region/Marker* dialog will appear.



Name the region *Cymbal Crash* and then select *OK*. A region will be created and placed in the *Regions List*. Each region in the *Regions List* has a *Play* button that you can use to play that particular region. In addition, *Region Tags* that show the location of the region are created in the *Time Ruler*.

OTHER WAYS OF CREATING REGIONS

There is also a shortcut menu to let you quickly create regions. Select the beginning part of the sound file with the drum roll and then right-click on the ruler. From the *Ruler* shortcut menu, select *Add Region/Marker* and the *Add Region/Marker* dialog will appear. Name this region *Drum Roll*. Notice that if you click on *Drum Roll* or *Cymbal Crash*, in the *Regions List* window, the region is selected in the data window.



The easiest way of creating regions is to drag them from a data window directly into the *Regions List* window. To do this, select the portion of data between the *Drum Roll* and *Cymbal Crash* regions and then drag it to the *Regions List* window and release the mouse button. The *Add Region/Marker* dialog will immediately appear upon releasing the mouse button. Name this region *Toms* and press *OK*.

Pressing “R” after a selection is made will create a region as well. This is a very fast method of creating regions, however, the *Add Regions* dialog won’t pop up. This means the region will be named automatically by its time boundaries. You can change the region name later (see below).

AUTOMATIC CREATION OF REGIONS

Sound Forge can create regions automatically for you in three different ways. One is to select the *Multiple takes creating Regions* mode while recording. This will create a regions every time you hit *Record* and *Stop*. Another way is to use the *Auto Region* function (*Tools* menu). This function analyzes the audio data and creates regions based on peaks, or it can create regions based on your *Beats per minute* settings. The other way of automatically creating regions is to use the *Regions List | Markers to Regions* function in the *Special Menu*, which converts all adjacent markers into regions. See the *Reference* chapter for more information on these functions.

EDITING A REGION

This section describes how to edit an existing region in a sound file.

EDITING REGIONS WITH THE RULER REGION TAGS

Changing the end points of a region is as simple as sliding the region tags on the data window ruler (above the waveform display). When you left-click on a region tag, the corresponding region end point tag is also highlighted. Also, the name of the region is displayed in the left side of the status bar. Double-clicking on a region tag changes the selection points to the region area.

Region edit operations can be performed by right-clicking on the region tags. This brings up the *Marker/Region* shortcut menu. You can then choose from the following options:

Select	Modifies the current selection points to be equal to the region end points.
Delete	Deletes the region, not the sound data, from the <i>Regions List</i> .
Split	With the cursor between the region boundaries, this command splits the region in two using the cursor as the new region boundary.
Edit	Takes you to the <i>Edit Marker/Region</i> dialog for manually editing region attributes.
Update	Modifies the region end points to be equal to the current selection.


EDITING REGIONS IN THE REGIONS LIST

Another way to edit region points is to double-click on a region name in the *Regions List*. This brings up the *Edit Marker/Region* dialog, which allows you to type in region parameters or select triggers (see the MIDI/SMPTE chapter). Also, right-clicking on the *Regions List* allows you to choose from various commands in the *Regions List* shortcut menu.

ADDING REGIONS TO THE PLAYLIST

Let's now arrange the regions you've just created in the *Playlist*. To add regions to the *Playlist* you can either use the *Special* menu option (*Playlist/Cutlist | Add*) or you can use drag and drop. Dragging a selection from the waveform display or a region from the *Regions List* into the *Playlist* is the quickest way to create *Playlist* entries.

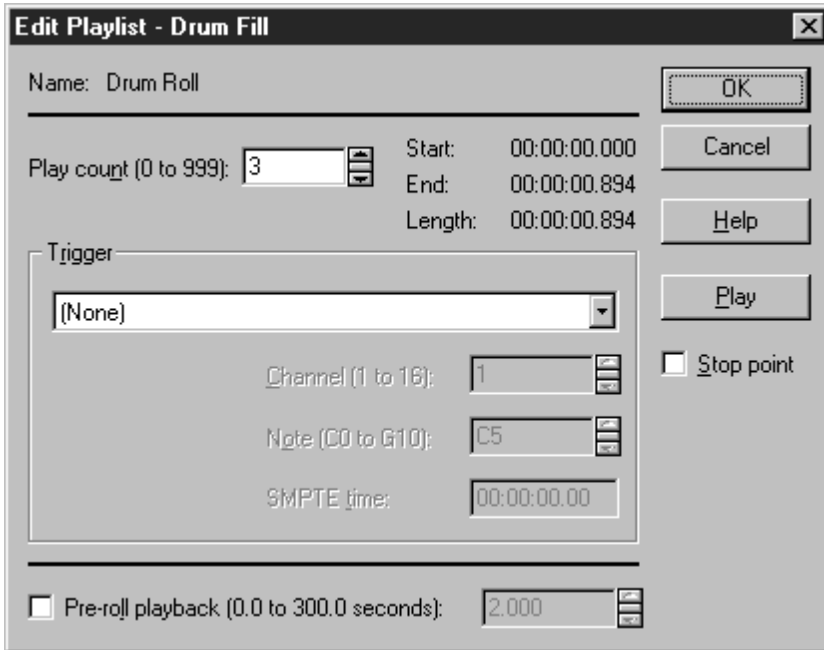
Since we already have some regions defined, we can just drag the regions from the *Regions List* into the *Playlist*. First, open the *Playlist* window by selecting the *Playlist* item from the *Window* menu or pressing *Alt+2*. Next, drag the *Drum Roll* region into the *Playlist*. Repeat this procedure with the other two regions. Selecting the *Play* button in front of any region in the *Playlist* plays the regions sequentially beginning with the selected region.



The *Playlist* shows the sequential order that regions will be played back. In front of the *Name* of the region, there is a number which indicates the number of times the region will be played before the next region is played (*Cnt*). After the *Name* is the *Start* time of the region. Several display options are available on the *Regions List* and *Playlist* pages in *Preferences*.

ARRANGING THE PLAYLIST REGIONS

Let's assume that you want the *Toms* region before the *Drum Roll* region. Select *Drum Roll* from the *Playlist* and drag and drop it to below *Toms*. When you start playback from the beginning region (press the *Play* button in front of the top region) you can hear the new arrangement. To place a region between two other regions, drag and drop the region between the two other regions.



CHANGING THE REGION PLAYBACK COUNT

Each region in the *Playlist* can be played (looped) a number of times. To do this, double-click on a region in the *Playlist* and in the *Edit Playlist* dialog set the *Play count* to any number between 0 and 999. You can quickly change the count for any selected region by using the + or - keys rather than bringing up the dialog. If you set the *Play count* for the *Drum Roll* region to 3 you can hear the drum roll played 3 times in a row.

PLAYING A REGION AT DIFFERENT POINTS IN THE PLAYLIST

One other advantage of using a playlist is the ability to repeat the same region at different points in time without having to copy the sound data in the file. If you want to create another region with the same data, select a region and then select *Replicate* from *Playlist* in the *Special* menu or from the *Playlist* shortcut menu. You can also replicate a region by dragging the region and, while holding the *Control* key, dropping it at another position in the list.

REMOVING A REGION FROM THE PLAYLIST

To remove a region from the *Playlist*, select the region and select *Delete* from *Playlist* in the *Special* menu or from the *Playlist* shortcut menu. You can also select a region and press the *Delete* key. This operation does not affect the sound file data.


CREATING A NEW FILE FROM THE PLAYLIST

Once you've arranged the *Playlist* regions, *Sound Forge* allows you to create a completely new file containing the sound data in the *Playlist* regions. For example, say you've created regions in a file and then re-ordered them in the *Playlist*. To make a new file containing the re-ordered regions, select *Playlist | Convert to New* from the *Special* menu or the *Playlist* shortcut menu.

TREATING THE PLAYLIST AS A CUTLIST

If most of your editing consists of trimming long recordings, using the *Cutlist* will greatly speed up your editing time. To enable the *Cutlist*, you must select *Treat as Cutlist* from the *Special | Playlist/Cutlist* menu. The *Cutlist* is opened by selecting *Playlist/Cutlist* from the *View* menu. Any regions in the *Cutlist* will not be played by *Sound Forge* when playing in *Cutlist Mode*. *Cutlist* playback mode is selected by choosing *Transport | Play as Sample/Cutlist* from the *Special Menu*.

For example, say you've recorded a voice-over into one large file and now wish to delete the bad takes, silent regions, coughs, etc. To "cut" sections of audio, you simply highlight the regions you wish to skip during playback and drag them to the *Cutlist* window (or press the *Delete* key). This operation is instantaneous; there's no waiting around for your hard drive to keep up with you. The regions contained in the *Cutlist* are displayed in the waveform as grayed out regions (a line pattern is drawn across the region).

Pressing the **Play as Sample/Cutlist** button  (the third play button in the Data Window) switches you to Cutlist playback mode. From then on, pressing the Spacebar will also play in Cutlist mode. Regions in the Cutlist will simply be skipped during playback.

When in Cutlist mode and a selection exists, pressing the Delete key adds the current selection to the Cutlist. This is a handy shortcut key when you need to perform multiple non-destructive deletes. If you want to make a destructive cut while in Cutlist mode, press Shift+Delete.

NOTE

Once you've created all of the regions you wish to delete, you can either create a new file or delete all the regions in the opened file. To create a new file without all of the cut regions, select **Special | Playlist/Cutlist | Convert to New**. Otherwise, selecting **Special | Playlist/Cutlist | Delete Cut Regions** will delete all regions in one pass.

If you want to use the Playlist instead of the Cutlist, uncheck the **Treat as Cutlist** option.

USING MARKERS

While regions are used to specify a section in a sound file, a marker is a location in the sound file you wish to use as a reference point. Markers make it easier to get around the sound file, as they are often used to tag important locations.

CREATING A MARKER WITH THE RULER SHORTCUT MENU

A marker can be created in the same ways discussed earlier for making regions. The quickest way is to place the cursor where you want the marker and then right-click anywhere on the ruler. This will bring up the Ruler shortcut menu, where you then select the **Marker/Region** option. In the **Add Region/Marker** dialog, type in a name for the marker and select **OK**. If you had no selection in the data window, Sound Forge will automatically create a marker at the cursor position. If you have a selection, Sound Forge defaults the **Type** to **Region** in the **Add Region/Marker** dialog.

EDITING MARKERS

Markers can be edited in the same way as regions. The marker tags in the ruler can be moved back and forth, and shortcut menus appear when you right-click on the tag.

STARTING PLAYBACK FROM A MARKER LOCATION

To play the sound file beginning at a marker location, press the small **Play** button to the left of the marker name in the *Regions List* window.

DROPPING MARKERS WHILE PLAYING OR RECORDING


Another useful feature in Sound Forge is the ability to create markers as you listen to the sound file. To do so, first start playback on any sound file. A marker will be created each time you press the “M” key or the *Mark-In* toolbar button. When finished, you can then rename each marker to whatever you wish. The most accurate method for creating markers during playback is to use the “M” key.

While recording, you can press the *Drop Marker* button (to the right of the *Play* button) in the *Record* dialog. Once focus is on this button, pressing the *Spacebar* will create other markers.

USING THE BATCH CONVERTER

OPEN THE BATCH CONVERTER WINDOW

The Batch Converter window is opened by selecting Batch Converter from the Tools menu in Sound Forge. It contains a file list (initially empty) showing sound files to convert and their status. The window also displays the output path (at the top) and brief information about the currently selected output file format and processing parameters in the bottom status bar.

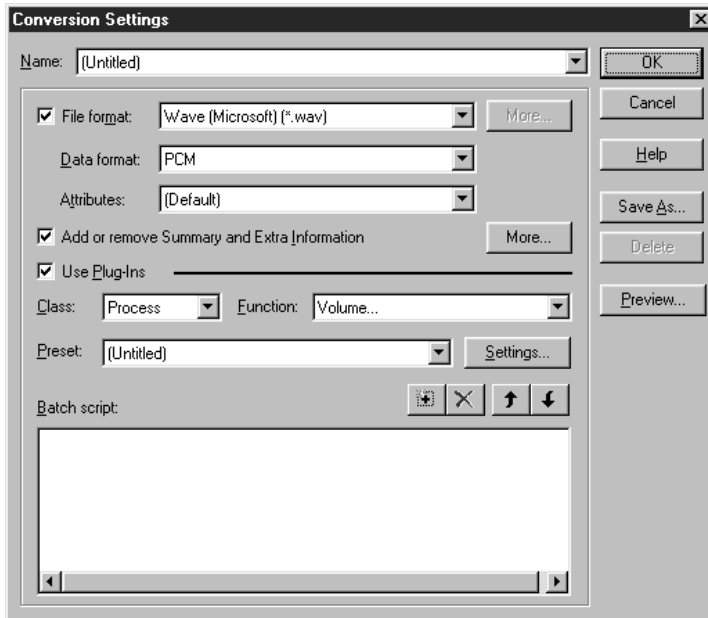


ADD THE SOUND FILES YOU WANT TO CONVERT

Select *Add* from the File menu of the Batch Converter window. In the *File Add* dialog, you can select one or more sound files you wish to convert and add them to the batch list. Once you've selected the files, press the *Add* button. You can also add files by dragging sound file icons from outside of Sound Forge (from any Windows Explorer window) into the Batch Converter window.

SELECT THE DESTINATION FILE FORMAT

Select *Conversion Settings* from the *Options* menu.



The *Conversion Settings* dialog is used to determine both the destination file format and audio processing that will be performed on the selected files. If you want the files to be converted to a specific file format, the *File format* checkbox must be turned on.

ADD PROCESSING FUNCTIONS

Select the audio processing functions you wish to apply to your batch of files. This is done by selecting functions, setting their parameters (or using presets), and pressing the *Add* button (the button with the plus sign). Once you're done adding functions, press *OK*.

SELECT DESTINATION FOR CONVERTED FILES

Select the *Destination Folder* item in the Options menu to set the location that the converted files will be saved. If you enter a folder that does not exist, the Batch Converter will create it for you.

CONVERT

Select *Run Current Script* from the Convert menu to begin the conversion. A progress meter will be displayed that shows the status of each function performed on each sound file until the batch processing is complete. You can cancel a batch at any time by clicking the *Cancel* button.

SAVING CONVERSION SETTINGS

All of the settings in the *Conversion Settings* dialog can be saved for convenient reuse by selecting the *Save As* button. The Batch Converter comes with several presets you can use as good starting points.

PROCESSING SPECIFIC SOUND FILES IN THE LIST

To process a single sound file from the file list, first select the file by clicking on the file in the list. Then, select *Convert Selected Only* from the Convert menu. You can select multiple files by holding down the mouse button while dragging the mouse over the files. Also, the Edit menu contains several options for managing the sound files to be processed. To re-process an already processed file, you must first reset it by using *Reset Selected* in the Edit menu.

USING BATCH SCRIPTS

Batch Scripts are essentially presets made of the *Conversion Settings* that you specify. However, they add to the flexibility of the Batch Converter greatly. By creating Batch Scripts, you can unlock the power of the *Run Batch Script* tool, the *Automatic Batching* function, and the *Run Multiple Scripts* function.

To create a Batch Script, specify the *Conversion Settings* you wish to use from the Options menu. While in the *Conversion Settings* dialog, press the *Save As* button. You will be prompted with a dialog asking for a name. Type a name for the script and press the *OK* button. That script is now available to the functions listed above.

To save a set of Batch Scripts into a single file that may be passed to a different computer, select *Batch Script/Save As* from the File menu. To import a Preset Package (.SFZ) that has been saved, choose *Batch Script/Add* from the File menu.

RUNNING MULTIPLE SCRIPTS

The Batch Converter allows you to run multiple scripts on a single group of files. This is useful for creating different bandwidth versions of Internet streaming files from a single source, for example.

To run a batch of multiple scripts, add the files you wish to batch as explained earlier in this manual. Select *Run Multiple Scripts* from the Convert menu. Place an 'X' in each script you wish to run.

The result of each script can be output to a different destination. To change the destination for any script, choose the script and press the *Destination* button.

When you have set up the parameters, press the *Run Scripts* button.

AUTOMATIC BATCHING

The Batch Converter can monitor a directory in the background and automatically batch any new files that are added. This makes it convenient to run on a network computer where users can dump finished files into the directory and have them processed and converted automatically.

To set up this feature, choose *Automatic Batching* from the Convert menu. Set the directory to monitor by typing in a path or press the *Browse* button to manually find the directory.

Press the *Scripts* button to choose the Batch Scripts you wish to use on the files. This dialog is exactly like the *Multiple Scripts* dialog mentioned above. When you have decided on settings, press the *Run Scripts* button. The screen will switch back to the *Automatic Batching* window.

To initiate *Automatic Batching* press the *Start* button. To stop *Automatic Batching* press the *Stop* button. The Batch Converter cannot be used on the same machine while in *Automatic Batching* mode.

PREVIEWING A CONVERSION

Before performing all of the conversions, you can preview the results on smaller audio sections in the *Preview Conversion* dialog. To open the *Preview Conversion* dialog, click the *Preview* button in the *Conversion Settings* dialog.

The *Preview Conversion* dialog allows you to build up to five different preview clips. To build a preview clip, select a file from the *Active File* list in the *Preview Conversion* dialog (one or more sound files must have previously been added to the batch list using the *Add* command in the *File* menu).

Once you've selected a sound file to preview, pressing **B** (**Build**) will create and play a preview clip of the selected file using the current parameters in the *Conversion Settings* dialog. You can re-play the preview clip by pressing the **>** (**Play**) button, while **S** (which appears on the play button) stops playback. To play the unmodified clip from the active file, press the larger **Play** button on the right side of the window.

The section of a sound file used for creating the preview clip is determined by the *Preview Start* and *Length* controls. Also, comments can be typed in the edit fields to the right of the **B** (**Build**) and **>** (**Play**) buttons.

Since the *Conversion Settings* dialog and the *Preview Conversion* dialog can be open and accessed simultaneously, you can conveniently change parameters and preview the results.

The effects of saving to a compressed audio format will not be heard when using the Preview dialog. You must open the processed file to do so.


N O T E

USING SPECTRUM ANALYSIS


This chapter introduces modes of displaying spectral data, the concepts behind performing a *Fast Fourier Transform* (FFT), and the basic steps for using the Spectrum Analysis Plug-In.

ANALYZING SPECTRAL DATA

FREQUENCY DOMAIN



Time Domain Waveform



Frequency Domain


The waveform display in a Sound Forge data window represents audio in the time domain (amplitude vs. time), while the Spectrum Analysis Plug-In allows you to examine a sound recording in the frequency domain.

Data displayed in the frequency domain (whether in the form of a *Spectrum Graph* or *Sonogram*) shows the amplitudes and frequencies of sine waves that, if mixed together, would sound very much like the original sound. Since it's relatively easy to remember what a sine wave sounds like at different frequencies, it's possible to visualize what simple waveforms sound like by looking at the spectrum of the sound.


Viewing the frequency components of a sound and their corresponding amplitudes, make it simple to determine the fundamental frequency and the other overtones present in a recording. Similarly, unwanted noise can be analyzed for determining the required filtering steps needed to remove the noise.

Two different methods are used to display spectral data in the Spectrum Analysis Plug-In – *Spectrum Graph* and *Sonogram*.


SPECTRUM GRAPH DISPLAY



The horizontal axis represents frequency in Hertz (Hz). The vertical axis represents amplitude in decibels (dB). It is also possible to view multiple spectrums corresponding to different points in time in the sample, in which case the display becomes three dimensional.



SONOGRAM DISPLAY



The Sonogram is another way of displaying spectral data variations over time. In this graph, the horizontal axis represents time and the vertical axis represents frequency. The amplitude of each frequency component is represented by the color intensity of each point in the graph. This method of displaying spectral information is useful for identifying distinctive spectral patterns created from sounds such as speech, musical instruments, bird calls, etc.

PERFORMING A FAST FOURIER TRANSFORM

The mathematical method used to convert a waveform from the Time Domain to the Frequency Domain is called a Fourier transform, named after the French mathematician Jean Baptiste Fourier. Fourier was one of the first mathematicians to claim that any periodic signal could be recreated by adding a series of harmonic sine waves together. Since his work in the early 1800's, Fourier analysis has been applied to many types of signals in hopes of gaining a better understanding of their composition.

Since the Fourier transform is very computationally intensive, it is common to use a technique called a *Fast Fourier Transform* (FFT) to perform spectral analysis. The FFT uses some tricky mathematical shortcuts to lower the processing time at the expense of putting limitations on the analysis size (such as requiring it to be a power of 2, which is the case in the Spectrum Analysis Plug-In).

The analysis size, also referred to as the *FFT size*, indicates the number of samples from the sound signal used in the analysis and also determines the number of discrete frequency bands. When a high number of frequency bands are used, the bands have a smaller bandwidth, which allows for more accurate frequency readings.

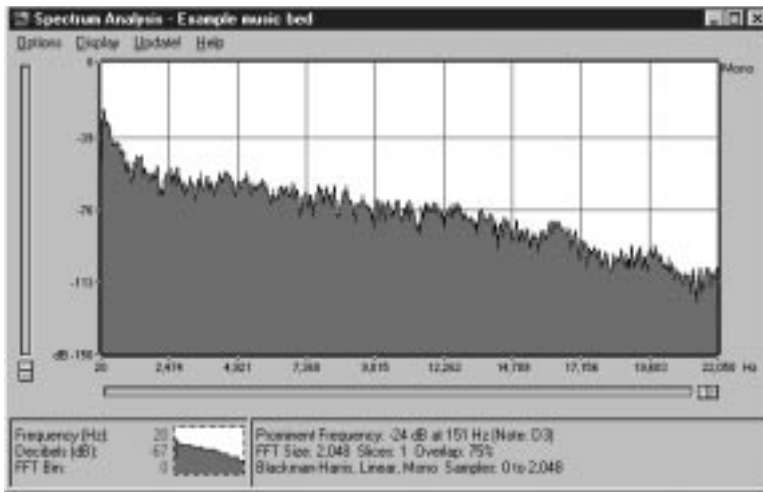
However, since most interesting sounds have a rapidly changing spectrum, a large analysis size can smear the time-changing frequencies of a sound together. For example, when doing an FFT analysis of a sound file sampled at 44,100 Hz using an analysis size of 4096, almost 100 milliseconds ($44,100/4096$) of sound are analyzed. If the sound is not constant for those 100 milliseconds, you will not be able to focus on the instantaneous spectrum at smaller time intervals. This is the ever-present trade-off between time resolution and frequency resolution often encountered when analyzing signals.

There is a lot of information to digest about these and other topics concerning audio frequency analysis. Outside reading is recommended if you want to more fully understand the subject.

SPECTRUM GRAPH

DISPLAYING THE SPECTRUM OF A SOUND FILE

Open a sound file and select the portion of the waveform that you wish to analyze. Select **Spectrum Analysis** from the **Tools** menu in the **Sound Forge** menu bar.



In the **Spectrum Analysis** window, the **Spectrum Graph** displays the amplitude (in dB) of each frequency component from 0 Hz (DC) to the Nyquist frequency (half of the sample rate). As you move the cursor in the **Spectrum Graph**, the frequency and amplitude values at the current position are displayed next to the cursor.

If the **Sonogram Graph** is displayed, select the **Normal Graph** option from the **Display** menu to view the **Spectrum Graph**. If the currently selected file is stereo, see the section titled **Stereo Files** below.

NOTE

You can continue to make selections in the sound file Waveform Display with the Spectrum window still open (i.e. move the cursor or make selections as you normally would in Sound Forge). Selecting Update! from the Spectrum Analysis main menu bar will update the spectrum display by analyzing the current selection in the sound file. If no selection is made, analysis is performed on the samples immediately following the cursor position.

To view multiple spectrums at a time, go to the Spectrum Settings dialog (Options menu) and increase the value in the Slices displayed control. Select the number of spectrums you wish to see and hit OK. Each graph is a spectral analysis performed starting at different times (samples) in the sound file. To scroll between the different slices, use the smaller horizontal trackbar under the graph.

The Spectrum Settings dialog contains advanced parameters used in the analysis and display of the spectrum. Different settings can be stored to best display the type of audio being analyzed. For a detailed description on the use of each parameter, refer to Spectrum Settings in the Reference section of this manual.

SPECTRUM GRAPH NAVIGATION

The Spectrum Graph has some powerful tools that allow you to zoom in on a particular frequency and/or amplitude range. The frequency and amplitude range determine the maximum and minimum frequencies and amplitudes displayed on both the Spectrum and Sonogram graphs.

SELECTING A FREQUENCY RANGE

1. Click and hold the left mouse button on the Spectrum Graph at one end of the frequency range you wish to zoom into.
2. Drag the mouse left or right until you've selected the frequency range that you're interested in and then release the button.
3. The Spectrum Graph frequency axis now ranges from the start and end frequencies you selected. To slide the selected range across the frequency spectrum, use the longer of the two horizontal trackbars below the graph or use the right and left arrow keys.

4. Selecting *Zoom Out Full* from the *Display* menu or *Spectrum* shortcut menu allows you to view the entire graph again. You can also zoom out completely by double-clicking on the front-most slice of the *Spectrum Graph*. *Zoom to Range* sets the range to the *Display Range* values specified in the *Spectrum Settings* dialog.

The Spectrum shortcut menu can be reached by right-clicking on the Spectrum display or with the keyboard shortcut Control+Shift+F10.

NOTE

SELECTING AN AMPLITUDE RANGE

1. Click and hold the left mouse button on the *Spectrum Graph* at one end of the amplitude range you wish to zoom into.
2. While holding down the left mouse button, click and release the right mouse button. The magnify range rectangle changes from vertical selection to horizontal selection.
3. Drag the mouse up or down until you've selected the amplitude range that you're interested in and then release the button. Only the selected amplitude range is then displayed.

You can also select the amplitude and frequency range by using the Display Range parameters in the Spectrum Settings dialog (Options menu).

NOTE

SELECTING BOTH AN AMPLITUDE AND A FREQUENCY RANGE

1. Click and hold the left mouse button on the *Spectrum Graph*.
2. While holding down the left mouse button, click and double-click with the right mouse button. The magnify range rectangle can now be changed vertically and horizontally to specify both a frequency and horizontal range.

3. Release the left mouse button to zoom into the currently selected region.

NOTE

Note: You can also change the frequency and amplitude display ranges by using the Spectrum Settings dialog.

GRAB/PAN MODE

Once you are zoomed in, switching to *Grab/Pan* mode allows you to scroll vertically and horizontally around the *Spectrum Graph*. To turn *Grab/Pan* mode on and off, select it from the *Display* menu or the *Spectrum* shortcut menu. The *Escape* key can also be used to turn it off. When in *Grab/Pan* mode, dragging the mouse (click and hold the left mouse button) slides the zoomed view of the *Spectrum Graph*, bringing temporarily hidden sections of the spectrum into view.

OVERVIEW WINDOW

The *Overview* window, located on the bottom left side of the window, indicates where you are in the file in relation to the entire FFT. The rectangle with dashed lines represents the region you see in the larger graph.

STEREO FILES

When viewing the spectrum of a stereo file, a spectral graph is displayed for each channel. Each graph's settings are set independently. For example, if you click on the top graph (notice how the graph's border becomes highlighted, meaning that it's selected) and then select *Normalize dB* from the *Display* menu, only the top graph's vertical scaling changes. The same is true for most of the other settings; you must click on the graph you wish to edit before making changes.

Cursor readings only appear on the selected graph. To switch graphs, click on the graph or use the Tab key.

NOTE

However, if the Sync Graphs option is checked (Display menu), zooming and other settings are synchronized between the two graphs.

SONOGRAM

DISPLAYING A SONOGRAM

1. Open a sound file and select the portion of the waveform that you wish to analyze.

Analyzing very long sections of audio can take a long time and lowers the time resolution, so it is suggested that the section be relatively short. Also, if the audio has a very low amplitude level, you might want to boost it by using the Volume or Normalize functions in Sound Forge.

NOTE

1. If the Spectrum Analysis window is not open already, select Spectrum Analysis from the Tools menu in the Sound Forge menu bar.
2. In the Spectrum Graph or the Spectrum Settings dialog, select the frequency and amplitude range you are interested in viewing (see section above). In most cases you'll want to display the entire ranges (fully zoomed out), but if you're interested in only human speech, for example, you might want to set the frequency range from 100 to 5,000 Hz to show only the important details.
3. Select either Sonogram (Color) or Sonogram (B&W) from the Display menu or the Spectrum or Sonogram shortcut menus. The sonogram is calculated while a progress meter is displayed and then the sonogram is revealed.

NOTE

If no selection is present in the waveform display window, the sonogram will analyze the sound data from the cursor position to the end of the file.

You will often have to experiment with different parameters in the *Spectrum Settings* dialog to get the best graph possible. Try to narrow the frequency and amplitude ranges as much as possible to achieve large contrasts. If the graph is too chunky, raise the *Sonogram resolution* setting to around 200 samplings. For more frequency resolution, increase the *FFT Size*. To lower processing times, lower the number of samplings or the *FFT Size*.

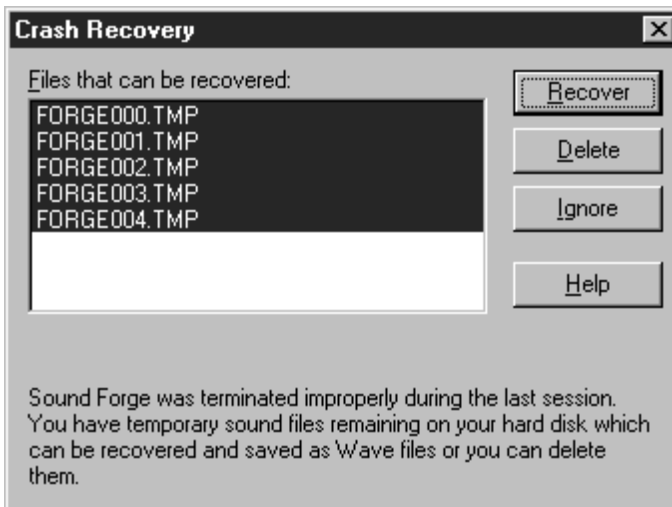
The color intensity can be adjusted by sliding the trackbar beneath the sonogram. The bottom right corner of the window will show the color scale in dB. This function may be fairly slow if you do not have a palletized driver and *Video for Windows* is not installed.

If you make a different selection in the sound file, selecting *Update!* will cause a new sonogram to be generated using the newly selected sound data. To return to the *Spectrum Graph*, select *Normal Display* from the *Display* menu.

Also, if the *Set sonogram resolution* checkbox in the *Spectrum Settings* dialog is turned off, the *FFT overlap* setting is used to determine how many samplings will take place. For very long files, this can be a very large number and will be very time consuming.

CRASH RECOVERY

If for some reason Sound Forge is terminated improperly – when Sound Forge or, most likely, another program crashes – all of the opened and unsaved sound files can be recovered excluding those opened in direct and read-only modes. Unless a file is opened in direct or read-only modes, a temporary file is created and any edits made are stored in this file. When an improper termination of the program occurs, these temporary files remain on your hard drive and can be reopened to recover any work done to the sound files before crashing. Also, the original sound files will remain unchanged until you save your work.



CRASH RECOVERY

When starting the program, any temporary files detected in your temporary file directory will indicate that something went wrong. You have the option to delete these files if no important work has been done, or to just ignore these files. Ignoring the files leaves them intact on your hard drive.

RECOVER

Use this to change the ending of all file names from FORGExxx.TMP to FORGExxx.WAV. You can then open these .WAV files using the *File Open* dialog. The files will be placed in the directory specified on the *Storage* page of the *Preferences* folder.

DELETE

This command deletes all the temporary files regardless of selection in the drop-down list. Use this only when you are sure that no important work is in the file.

IGNORE

If you don't want to deal with the temporary files left over from an improper termination at this moment, use this command. However, it is recommended that you either rename the files and check their contents or delete them, since they are taking up valuable disk space.

OPTIMIZING SOUND FORGE

This chapter contains information on configuring your system to maximize Sound Forge's performance.

HARD DRIVE USE

Sound Forge is a disk-based sound editor, which means that all editing operations are performed on your hard drive rather than in your computer's memory. This allows Sound Forge to edit sound files of almost unlimited size. Sound Forge also stores undo/redo information and clipboard data on the hard drive.

The drive you specify for *Temporary storage* (on the *Performance* page of the *Preference* folder) must have enough free space to store the data for each opened sound file, as well as undo/redo information, and the clipboard. Although this can use a lot of disk space, it allows you to edit long sound files that could never fit into your computer's memory (RAM). And hard drives are typically much cheaper than memory.

Since Sound Forge is hard-drive intensive, the faster your disk access, the faster Sound Forge will be. The easiest thing to do to improve your performance without buying new hardware is to defragment your hard drive. Over time, and especially with heavy use, a hard drive can become fragmented. This means the data stored on the hard drive gets scattered across the surface of the disk making your files discontinuous. When files are discontinuous, the hard drive must jump around to retrieve the data. This jumping around can slow down access to your files and can be very noticeable on older hard drives. DOS 6.0 and above come with a program called *Defrag* that will defragment your hard drive. Windows 9x comes with a *Disk Defragmenter* utility that can be found in the *System Tools* folder in your *Accessories* folder.

The following guidelines will help you optimize your hard drive depending on which operating system you are running.

WINDOWS 9X

WINDOWS 9X DEFAULT CONFIGURATION

Choosing system configuration settings other than the default Windows 9x settings can result in reduced performance from Sound Forge, as well as Windows 9x. To check these settings go to *System* under *Control Panel* and click the *Performance* tab. *File System* and *Virtual Memory* should both be set to 32-bit. In the advanced settings area of the window, click on *File System*. In the *File System Properties* dialog, the *Hard Disk* tab setting for *Read-ahead optimization* should be set to *Full* and the *Troubleshooting* tab should have all of the options unchecked.

Now, click on *Graphics* in the *Advanced Graphics Settings* dialog. In most cases this setting should be set to *Full*. If you are experiencing unexplained glitching in your audio playback, try different settings. However, if changing these settings doesn't clear up the problem, change them back. Finally, click on the *Virtual Memory* button under the *Advanced* settings of the *Performance* page. In most cases, you should choose to let Windows handle your virtual memory.

WINDOWS NT

WINDOWS NT DEFAULT CONFIGURATION

There are no special considerations for Windows NT.

EDITING IN DIRECT MODE

Whenever you open a file in Sound Forge, a backup copy is created so that the original file is not affected until you select to save your work. When you open a file as read-only or direct, Sound Forge will not create a backup copy of the file, which makes opening files quicker and uses less disk space. In fact, editing in direct mode can make opening a file almost instantaneous. Direct mode can only be used on files that are 16-bit and stored in the Microsoft Wave format.

The first time you open a file in direct mode Sound Forge will scan the file to create a peak file. This allows Sound Forge to display the data in the file extremely fast as well. This peak file is stored in the same location as the wave file with the same name as the wave file but with the extension *.sfk*. After a peak file is created Sound Forge automatically keeps it updated when you edit

the file. The only time the peak file may need to be regenerated is if you modify the file with an application other than Sound Forge. This allows Sound Forge to instantaneously load the file next time you open it.

The only drawback to editing a file in direct mode is that you no longer have the security of a backup file. In most cases this is not a problem, since all changes can be undone (with ample disk space). But if, for some reason, Sound Forge terminates improperly, the file would remain in its edited state.

FAST FILE SAVING

When creating a new file Sound Forge creates a temporary file that contains all of your new sound data. When you are ready to save the file, Sound Forge must copy this data into the file you specify in the *Save As* dialog. If you are saving the sound data as a 16-bit Microsoft Wave file, Sound Forge can make this process almost instantaneous by quickly moving the file rather than copying the data to a new file. In order for this to work, the new file into which you are saving your data must be on the same drive as your temporary storage folder. Thus it is recommended that you keep your temporary storage folder on the same drive where you keep most of your sound files. You can set your *Temporary storage folder* on the *Performance* page in the *Preferences* folder.

SETTING UP WAVE DEVICES

THE SOUND MAPPER

Sound Forge allows you to choose any installed Windows compatible sound card for playback and recording. It also provides one more option for playback and record called the *Sound Mapper*. The *Sound Mapper* is a special device which attempts to select the most appropriate sound card (map) on which to play a sound, or it will translate the sound into a format that can be played on your sound card.

If you are having problems with breakups this may be caused by the overhead introduced when the *Sound Mapper* does format conversions. To check this, make sure that your sound card supports the data format you are playing or recording. The easiest way to do this is to choose a playback and record device that is not the *Sound Mapper*. If Sound Forge is able to play and record the sound

using the actual sound card's Wave driver, then the *Sound Mapper* is not causing the breakups.

However, if the *Sound Mapper* must be selected to play or record the format that you are having trouble with, then you should convert your sound data to a format that is directly supported by your sound card. This will remove all overhead required to translate the sound data for your sound card.

There is nothing wrong with using the *Sound Mapper* when the sound data format is directly supported by your sound card. The *Sound Mapper* will simply pass the data through to your sound card which requires negligible overhead. For more information on the *Sound Mapper*, refer to the Appendix called *Sound Forge and The Microsoft Audio Compression Manager*. The *Sound Mapper* is a component of the *Audio Compression Manager (ACM)*.

TOTAL BUFFER SIZE

The *Total buffer size* option, found in the *Wave* folder of the *Preferences* dialog, specifies the total amount of RAM that should be used for buffering when recording to or playing from the hard drive. For most systems 512 kb is a good number. If you are still experiencing gaps during playback you may want to use more buffering or record at a lower sample rate.

There are some drawbacks to increasing the *Total buffer size*. It requires more of your computer's memory and, if you have set a large *Preload size*, there can be small delays when starting and stopping playback.

PRELOAD SIZE

The *Preload size* option in the *Preferences | Wave* folder tells *Sound Forge* how much sound data it should prepare and load into your sound card driver before starting playback. This preloading is performed between the time you press the *Play* button and the time you hear the sound play.

Turning this option on may help skipping (audible gaps) on initial playback of a sound. These gaps are normally observed by people with slow or fragmented hard drives. The trade-off for enabling this option is a slight delay before play begins; the delay time is greater for larger preload sizes. This option does not affect operations other than playing.

Some Windows sound drivers do not support this option correctly. If you are experiencing noise or dropouts at the very start of playback, try disabling this option.

FORCED WRITE-THROUGH DURING RECORD

Due to a poorly implemented disk cache in Windows 95, Sound Forge implements a forced write-through during recording. This method will make recording less prone to drop outs and gapping under Windows 95. However, this does tend to make the hard drive more active and possibly noisy.

If this is a problem, you may uncheck the option, *Enable forced write-through on Record* in the *General* page of *Preferences*. This is not recommend, however, as gapping during record may occur.

PLAYBACK CURSOR AND RECORD COUNTER

The playback cursor and record counter options in the *Preferences | Performance* folder allow you to configure whether or not the numeric counters (position information displayed in status fields) will be updated during playing and recording. If you are recording or playing at high sample rates (44,100 or 48,000 Hz) and are experiencing gapping or skipping problems in the data, try disabling the playback cursor and/or record counters to decrease overhead.

METERS

The *Play* and *Record Meters* in Sound Forge take a small amount of processing overhead during playback and record. If you are experiencing gapping during playback and you have the *Play Meters* showing, try turning off the *Play Meters* by selecting the *Play Meters* option in the *View* menu. If you are experiencing gapping during record and you have the *Monitor* option enabled in the *Record* dialog, try turning off the *Record Meters* by un-checking the *Monitor* option.

PASSIVE UPDATING FOR VIDEO AND TIME DISPLAYS

Selecting *Passive Update* for *Video* and *Time Display* from the *Options* menu will lower the priority of the repainting for these displays during playback. If your computer is relatively slow and you are experiencing gapping during playback or your computer just seems to be bogging down during playback, try turning on *Passive Update*. This will force the repaints for the *Video* and *Time Display* to only update if there is time to do so. In most cases you won't even be able to

tell if it is missing some updates, so turning these options on can decrease the amount of overhead during playback with little inconvenience.

DISABLE SCAN FOR EXTERNAL MONITOR

Sound Forge has the ability to display its AVI preview on an external video monitor. By default, Sound Forge will scan for an external video monitor everytime you open an AVI file. If you do not have an external video monitor, you will want to disable this feature, as it can slow down the open process for AVI files.

To disable this feature, check the *Disable scan for External Monitor support* checkbox in the *Video* page of *Preferences*.

BACKGROUND PROCESSING

When Sound Forge is performing operations that take a long time, such as opening a file or performing audio processing, a *Progress Meter* is displayed on the status bar. While this *Progress Meter* is visible, Sound Forge is doing its work. You can cancel the operation at any time by pressing the *Cancel* button to the left of the *Progress Meter*, or you can press the *Escape* key.

All processing performed by Sound Forge is done in the background. What this means is you do not have to wait for Sound Forge to finish its work before you can use other applications. If an operation is taking a long time in Sound Forge, simply switch to another application and work on something else. Sound Forge will continue processing while you work in another application.

If you minimize Sound Forge while it is processing, the percentage processed will update in the title of the Sound Forge icon. This allows you to keep an eye on Sound Forge's progress while using other applications.

You may have noticed that Sound Forge displays the amount of processing time an operation required in the status bar when it is completed. This processing time is intended for entertainment purposes only—this value is not intended for accurate profiling. The accuracy of this processing time value is heavily dependent on what other activity is being performed by your computer while Sound Forge is working. Using other applications while Sound Forge is working can cause the processing time to increase substantially.

AUDIO AND VIDEO SYNCHRONIZATION

If you are having trouble synchronizing the audio with your video, there are a few points to keep in mind. If your video has been opened from a slow device, such as a CD-ROM or network drive, then Sound Forge may have trouble accurately playing back the audio and video in sync. This happens because video information is streamed at all times from the original .AVI file. You should always copy the videos you are working with to a fast hard drive.

After you open an .AVI file you may be notified that your video has very few key frames. This is generally representative of tutorial videos captured with a program that saves the video as a *Run Length Encoded (RLE)* stream with a single key frame. Due to non-linear access requirements of Sound Forge this may result in difficulty with audio synchronization. It is recommended that you save your file with more key frames and re-open it for better performance during digital editing.

Here are a few other tips that can help when trying to synchronize the audio and video:

- After assembling or editing the audio you wish to use with your video, place markers during video playback to correspond to any major synchronization points. You can locate a particular frame by dragging the cursor along the audio – provided the *Video Preview* window is open or the *Animate Video Strip* option is enabled. After primary locations have been located, double-click and drag your audio to these markers bringing up corresponding dialogs for mixes, pastes, and crossfades.
- Features like *Insert Silence*, *Clear/Delete*, and *Time Compress/Expand* are commonly used to tighten synchronization. Another useful trick is to create a region representing the delta between a video frame and audio event. Then you can enable the *Lock Loop/Region Length* and drag the delta region to a preceding silent section. Use the region as a template for adjusting the audio stream length — either copying and pasting to insert time or deleting to remove time.
- When cropping beginning and end silence for your .AVI file, drag your finalized selection to the *Playlist*. Choosing *Convert to New* will generate a new file depicting your selected region.

This page intentionally left blank.

SAMPLING

SAMPLERS

Without the convenience of a large computer monitor, an advanced operating system like Windows, and an intuitive interface, working with samples can be a nightmare. With the *Sampler Tool*, you can use Sound Forge's powerful and easy-to-use sample editing capabilities to create and edit samples and then transfer them to your external and internal samplers.

This chapter provides all the information you need to transfer (dump) samples between your computer and sampler using the *Sound Forge Sampler Tool*. In addition, we have included a section titled *Tips and Ideas* that may stir up some new, and perhaps unconventional, ways of using your sampler.

The documentation that came with your sampler undoubtedly has insights on sampling and loop tuning as well, so you might want to find it, dust it off, and browse through it. Many of these manuals can be very useful when you read them.

EXTERNAL SAMPLERS

External samplers are devices that can play samples at different pitches for use as a musical instrument. Most external samplers allow you to record samples or transfer pre-recorded samples into their memory. There are two methods of transferring (dumping) samples to external samplers supported by the *Sampler Tool* – MIDI Sample Dump Standard (SDS) and SCSI MIDI Device Interface (SMDI).

MIDI SAMPLE DUMP STANDARD (SDS)

The MIDI SDS is used to send and receive digital samples with your normal MIDI hardware and cable connections. Because of the limited bandwidth of the MIDI protocol and the large amounts of data needed to store digital samples, performing a sample dump using the normal MIDI SDS can be slow, often taking minutes to transfer short samples. SDS is also limited to mono-only samples, though some samplers allow you to combine two mono samples to create a single stereo sample.

SCSI MIDI DEVICE INTERFACE (SMDI)

The SCSI MIDI Device Interface (SMDI) is a relatively new standard for music hardware and software to communicate using SCSI hardware and cables. Because SCSI hardware has a much larger bandwidth than MIDI, SMDI sample dumps are many times faster than SDS dumps. Another advantage to using SMDI is that it supports mono and stereo sample dumps (unlike SDS).

INTERNAL SAMPLERS

Internal samplers are cards that are installed inside of your computer. Unlike most sound cards, internal samplers allow you to download your own sounds into the card's memory. These sounds can then be played at different pitches to simulate a musical instrument just like an external sampler.

The Sampler Tool currently supports only one internal sampler, the Digidesign SampleCell II. Once installed, use the procedures outlined in the Appendix section on Sample Transfer Procedures for Digidesign SampleCell II.

Internal samplers that are not directly supported by the Sampler Tool could use the MIDI Sample Dump Standard (SDS). Though no Windows-compatible samplers support SDS (that we know of), it would be very simple for a manufacturer to implement this feature. If nothing else, SDS could be supported using an open loop transfer which would eliminate the complexity of supplying an input device. Open loop would also make it possible for an internal sampler to accept data at far greater speeds than MIDI normally allows.

If you own a Windows-compatible internal sampler, you may want to contact the manufacturer about supporting SDS in their Windows drivers. SDS may not be the best solution, but it is far better than no solution.

SAMPLER TOOL DIALOG OPTIONS

There are two dialogs available for the Sampler Tool. The first dialog is called the Sampler dialog and is used to send and receive samples to pre-configured samplers. The second dialog is the Sampler Configuration dialog and allows you to configure one or more sampler setups for use in the Sampler dialog.

SAMPLER DIALOG

The *Sampler* dialog (Tools menu) allows you to send and receive data to and from a sampler. To set up sampler/computer communications, press the *Configure* button.

CONFIGURATION

This drop-down list contains preset sampler setups (saved in the *Sampler Configuration* dialog). Once you save different configurations, the settings can be conveniently selected from this list.

LOGICAL SEND/RECEIVE SAMPLE NUMBER (0 – 16,383)

This is the number your sampler uses as its location reference for samples sent or received. This number can be biased for specific samplers with the *Sampler bias* option in the *Sampler Configuration* dialog.

ACTUAL SEND SAMPLE NUMBER

This text field contains the actual sample number to be used by your sampler to store the sample. This number is based on the *Logical send/receive sample number* combined with the *Sample bias* for sending, set in the *Sampler Configuration* dialog.

ACTUAL RECEIVE SAMPLE NUMBER

This text field contains the actual sample number to be retrieved from your sampler into the current data window. This number is based on the *Logical send/receive sample number* combined with the *Sample bias* for receiving set in the *Sampler Configuration* dialog.

SEND SAMPLE

Press this button to begin transmitting the active data window's sample to the sampler.

GET SAMPLE

Press this button to begin receiving a sample from the sampler. The sample will always be placed in a new data window.

N O T E

A message box will be displayed requesting confirmation if there is already data in the active window.

CONFIGURE

Opens the *Sampler Configuration* dialog where you can determine the settings for communication with your samplers.

SAMPLER CONFIGURATION DIALOG

The *Sampler Configuration* dialog is used to create and save sampler communication configurations.

NAME

The name of the sampler configuration that is saved as a preset.

SEND TO

Lists supported samplers and methods of communication for sending samples. If your sampler is not listed here, try *Generic SCSI/SMDI* or *Generic MIDI/SDS*.

MIDI OUTPUT (MIDI/SDS)

Select the MIDI interface connected to the input of the sampler to which you wish to send samples.

MIDI INPUT (MIDI/SDS)

Select the MIDI interface connected to the output of the sampler to which you wish to send samples.

SCSI HOST (SCSI/SMDI)

Select the SCSI card used to send data between the computer and sampler.

SCSI SAMPLER (SCSI/SMDI)

Select the SCSI-compatible sampler which will receive the data.

SAMPLE BIAS**(-16,383 - 16,383)**

This is a numbering offset that allows you to add or subtract a value to the logical sample number to compensate for differing sampler storage schemes. It is often easiest to set the *Sample bias* so a logical sample number of zero corresponds to the first available sample storage number in your sampler.

MIDI CHANNEL (MIDI/SDS)**(1-16)**

Specifies the MIDI channel on which the sample data will be sent when using SDS.

WAIT FOR REQUEST (MIDI/SDS)

Selecting this option tells the *Sampler Tool* to wait for a request (handshake) from the sampler before beginning an SDS sample transfer.

OPEN LOOP (MIDI/SDS)

Selecting this option will cause the *Sampler Tool* to send SDS sample data immediately upon clicking the *Send Sample* button. This is an unconditional transfer of sample data (no handshake). When using *Open loop* for send, the sampler's MIDI out does not have to be connected to your computer's MIDI in. However, using *Open loop* decreases sample dump performance.

RECEIVE FROM

Lists supported samplers and methods of communication for receiving samples. If your sampler is not listed here, try *Generic SCSI/SMDI* or *Generic MIDI/SDS*.

MIDI OUTPUT (MIDI/SDS)

Select the MIDI interface connected to the input of the sampler from which you wish to receive samples.

MIDI INPUT (MIDI/SDS)

Select the MIDI interface connected to the output of the sampler from which you wish to receive samples.

SCSI HOST (SCSI/SMDI)

Select the SCSI card used to receive data from your sampler.

SCSI SAMPLER (SCSI/SMDI)

Select the SCSI-compatible sampler which will send the data.

SAMPLE BIAS**(-16,383 - 16,383)**

This is a numbering offset that allows you to add or subtract a value to the logical sample number to compensate for differing sampler storage schemes. It is often easiest to set the *Sample bias* so a logical sample number of zero corresponds to the first available sample storage number in your sampler.

MIDI CHANNEL (MIDI/SDS)**(1 - 16)**

Select the MIDI channel on which the sample data will be received when using SDS.

SEND REQUEST (MIDI/SDS)

Selecting this option tells the *Sampler Tool* to send a request to your sampler asking for a sample when you press the *Get Sample* button. If you turn this option off, you will always have to initiate transfers to the *Sampler Tool* from your sampler after pressing the *Get Sample* button. In other words, you will have to press a button on your sampler to start the actual transfer.

OPEN LOOP (MIDI/SDS)

Selecting this option will cause the *Sampler Tool* to wait for sample data after pressing the *Get Sample* button. You must initiate all sample transfers from your sampler. This is an unconditional transfer of sample data (no handshake). When using *Open loop* for receive, the sampler's MIDI in does not have to be connected to your computer's MIDI out. However, *Open loop* decreases sample dump performance.

CONFIGURING THE SAMPLER TOOL

After you have your sampler hardware installed and functioning correctly, you must configure the *Sampler Tool* for sending and receiving samples from your sampler.

NOTE

If you do not have your sampler and MIDI or SCSI card set up for transferring samples, you should refer to the *MIDI/SDS Setup and Troubleshooting* or *SCSI/SMDI Setup and Troubleshooting* later in this chapter. This information will help you install the necessary components for transferring samples using the Sampler Tool. There are also suggestions for troubleshooting problematic hardware configurations.

Select **Sampler** from the Tools menu and click the **Configure** button. Note that you must have at least one data window open to access the Tools menu. The **Sampler Configuration** dialog lets you select your sampler and sample transfer mode for both sending and receiving. The most common configuration is to have both the send and receive sampler and transfer mode the same. This is not a requirement, however, and there are situations where differing send and receive setups can be useful. One example of this might be if you want to transfer samples from one sampler to another with editing in between. Multiple configurations can be saved as presets to match different scenarios.

CREATING A SAMPLER CONFIGURATION

The first step to configuring your sampler setup is to choose the sampler and transfer mode (MIDI/SDS or SCSI/SMDI) from the *Send to* and *Receive from* drop-down lists. Several samplers and transfer modes are listed by name in these lists. For example, *Peavey SP SCSI/SMDI* specifies a *Peavey SP (SCSI/SMDI)* transfer protocol. If your sampler is listed, select it from the list.

If your sampler is not listed in the *Send to* and *Receive from* drop-down lists, try using the *Generic (SCSI/SMDI)* or *Generic (MIDI/SDS)* options. As long as your sampler properly supports the transfer mode (protocol) you select, the **Sampler Tool** should be able to work with your sampler. If you have trouble and you are certain your sampler is properly set to send or receive samples, contact your sampler's manufacturer. They may have an updated ROM available to correct the problem.

INTERNAL SAMPLER CONFIGURATION OPTIONS

Currently, only the Digidesign SampleCell II internal sampler is directly supported by the Sampler Tool. If you have a SampleCell II and you have selected it as the *Send to sampler*, then there are no further configuration options. A complete discussion on the SampleCell II can be found in the Appendix section on Sample Transfer Procedures for Digidesign SampleCell II.

NOTE

The SampleCell II is only available for the Send to option; you cannot receive samples from the SampleCell II.

MIDI/SDS CONFIGURATION OPTIONS

If you choose a sampler that uses the MIDI/SDS transfer protocol in the *Send to* or *Receive from* drop-down lists, then you need to specify a few more options. These options include the *MIDI input* and *MIDI output* ports, the *MIDI channel* to transfer on, and whether handshaking should be used. You can also configure how a transfer will be initiated if handshaking is enabled. These options are discussed in detail below.

MIDI INPUT AND OUTPUT PORTS

Select the *MIDI input* and *MIDI output* ports that are connected to your sampler from the drop-down lists. If *Open loop* (discussed later) is being used for sending samples, then the *MIDI input* port can optionally be set to *(None)*. If *Open loop* is being used for receiving samples, then the *MIDI output* port can optionally be set to *(None)*. If *Open loop* is not being used, then you must supply both a *MIDI input* and *MIDI output* port.

THE MIDI CHANNEL

Now select the *MIDI channel* to use when transferring samples. This channel must match the channel configured in your sampler for SDS transfers.

NOTE

A channel must be specified even if the sampler is set to Omni (any channel) mode.

HANDSHAKING

The last few options determine how the sampler and the *Sampler Tool* interact while transferring samples. These options are *Wait for request (send)* and *Open loop* or *Send request (receive)* and *Open loop*.

Open loop is a term used to describe a unidirectional communication protocol – much like a one-sided conversation. When the *Open Loop* option is enabled, the data source will send all data to the destination without listening for directions from the destination. The destination has no control over how the data is sent and cannot ask for information to be repeated. This lack of communication makes open loop transfers prone to errors.

If the *Open Loop* option is not enabled, we refer to the communication protocol as closed loop. A closed loop allows information to flow in both directions; the two communicators are said to be handshaking because they both agree on how and when data is transferred. When using closed loop transfers, the data is sent from the source in small packets (one packet at a time). The destination receives the packet and keeps the data if it is valid or discards the packet and asks for the data to be re-sent. The source will not send the next packet until the destination asks for it. This form of communication makes closed loop transfers much more reliable than open loop transfers.

It should be noted that open loop transfers are not only more prone to errors, but they can be slower than closed loop transfers, especially when sending samples with the *Sampler Tool*. The reason for this is due to intentional delays placed between data packets. These small delays allow slow samplers to keep up with the data being sent, and since these delays are not fine-tuned to exact requirements for the sampler, idle time is bound to occur. A closed loop transfer guarantees the most efficient timing between packets making transfers faster than open loop transfers.

Beware of using the Open loop option to receive samples from a sampler. Since the Sampler Tool cannot control the flow of data packets being sent, there is a high probability that data will be missed. Using closed loop transfers guarantees that no packets will be missed.

NOTE

You may be asking why *Open loop* is an option if it is so unreliable. Good question. The answer is that there are cases when using *Open loop* is either necessary or more convenient than using closed loop. If a sampler does not

support closed loop (handshaking), then you must use *Open loop* for transfers. You may also want to enable *Open loop* if you are out of MIDI cables or are having problems with a MIDI port. See the section titled *MIDI/SDS Requirements and Troubleshooting* for further details on open loop transfers.

If you choose to use the recommended closed loop transfer, then you also have the option of determining how a sample transfer is initiated. For sending samples, you can specify that the *Sampler Tool* should wait for the sampler to request the sample transfer before sending the sample by enabling the *Wait for request* option. This means that the *Sampler Tool* will not send the sample (after you press *Send Sample*) until the sampler requests the sample to be sent. This usually requires you to press a button on the sampler to start the transfer.

For receiving samples, you can specify that the *Sampler Tool* should send a request to the sampler asking for a sample to be sent by enabling the *Send request* option. This means that the *Sampler Tool* will send a request for the specified sample to the sampler as soon as you press *Get Sample*. If you do not enable the *Send request* option, you must initiate the sample transfer from the sampler after you press *Get Sample*.

TYPICAL SETUP FOR MIDI/SDS TRANSFERS

Send to Configuration (typical)

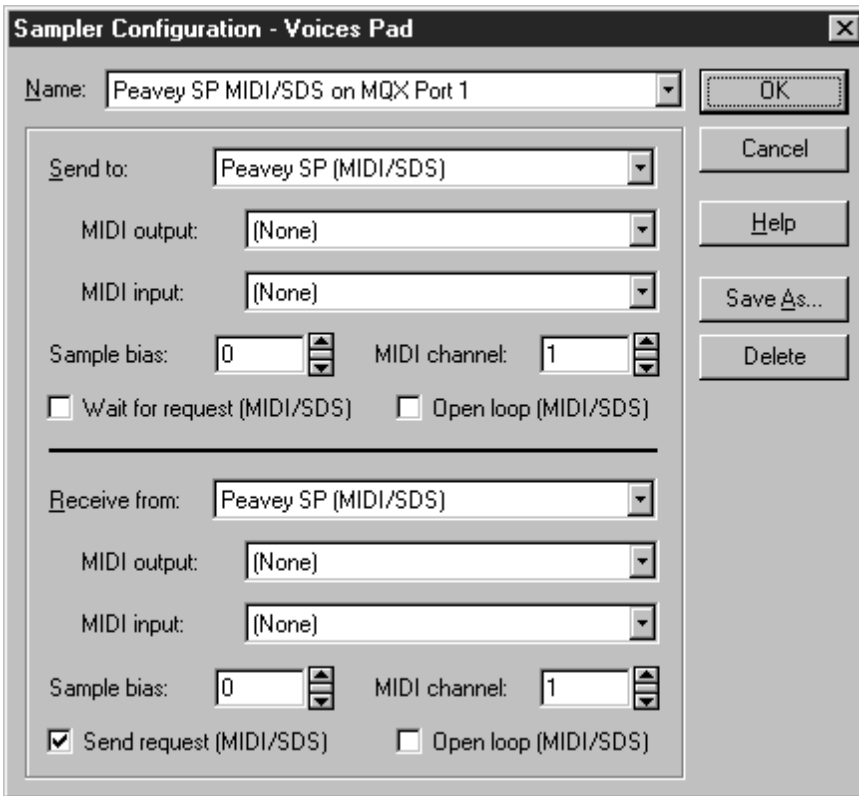
Connect and select a *MIDI input* and *MIDI output* device. Use your preferred *MIDI channel*. Disable *Open loop*. Disable *Wait for request*.

Receive from Configuration (typical)

Connect and select a *MIDI input* and *MIDI output* device. Use your preferred *MIDI channel*. Disable *Open loop*. Enable *Send request*.

The following picture shows a typical configuration for the Peavey SP (MIDI/SDS) transfers. The MIDI card being used in this example is a Music Quest MQX-32(M).

The *Sample bias* fields will be discussed in the *Selecting a Sample Bias for your Sampler* section later in this chapter.



SCSI/SMDI CONFIGURATION OPTIONS

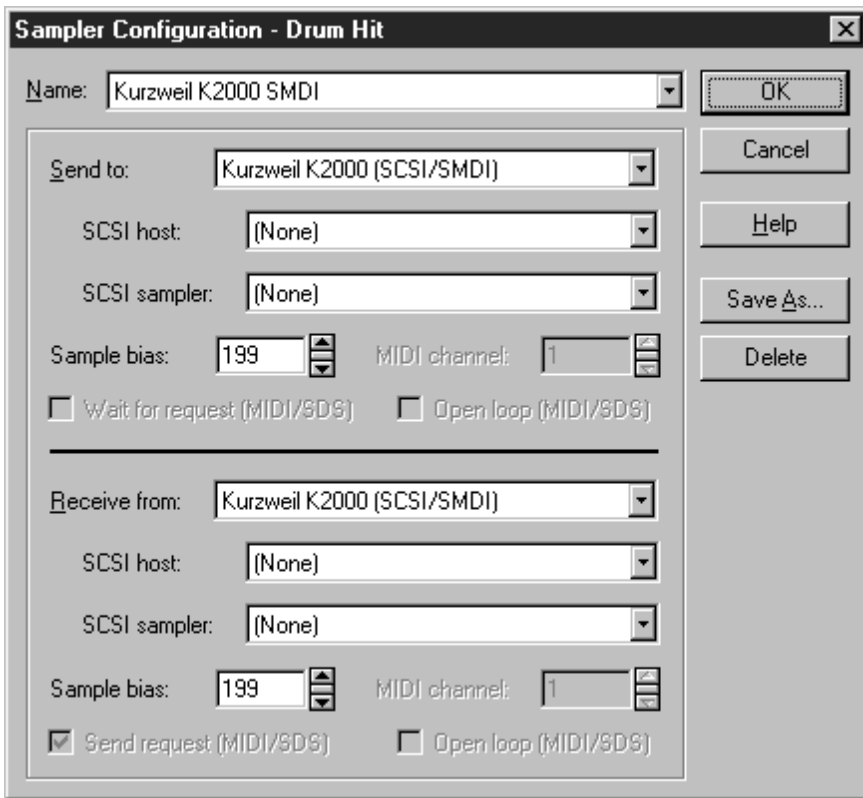
If you choose a sampler that uses the SCSI/SMDI transfer protocol in the *Send to* or *Receive from* drop-down lists, then you need to specify two other options. These options are the *SCSI host* and *SCSI sampler*.

The *SCSI host* drop-down list contains all *Sampler Tool* compatible SCSI cards found on your computer. You need to select the SCSI card that is connected to your sampler from this list. Note that most computers will have only one SCSI card listed.

Once you have selected the *SCSI host*, the *SCSI sampler* list will contain all SCSI devices connected to that host. Select your sampler from this list (the number

preceding the device name is its SCSI ID). Depending on your SCSI setup, you may notice that hard drives, CD-ROM's and other SCSI devices are listed along with your sampler. This is to help troubleshoot SCSI ID conflicts should one occur. If you accidentally select a device other than a sampler, no harm is done. However, a message box will be displayed stating the problem if you try to send or receive samples from any device other than a sampler.

If your SCSI card or sampler does not show up in the SCSI host or SCSI sampler drop-down lists, refer to the section titled SCSI/SMDI Requirements and Troubleshooting to determine the cause. The problem may be as simple as turning on your computer and sampler in the wrong order. It may also be that your SCSI card is not compatible with the *Sampler Tool* or your SCSI card and sampler are installed incorrectly.



The preceding picture shows a typical configuration for the Kurzweil K2000 (SCSI/SMDI) transfers. The fields that are grayed are for MIDI/SDS transfers only.

SELECTING A SAMPLE BIAS FOR YOUR SAMPLER

Those of you that have worked with samplers from different manufacturers know that the sample numbering scheme can vary dramatically between samplers. In an effort to simplify the confusion caused by working with different samplers, the *Sampler Tool* lets you specify a sample bias. A sample bias is simply a value that gets added to the logical sample number to derive the actual sample number used for sending or receiving. Don't panic, we'll explain everything.

When you want to send or receive a sample from a sampler, you must specify a sample number which is the location that the sample is stored within the sampler. This is pretty simple. However, not all samplers allow you to use any sample number for storing and retrieving samples.

Samplers that have factory-defined samples burned into their ROMs usually define a range of sample numbers for these samples. This range cannot be used to store and retrieve your own samples. These samplers will define one or more sample number ranges that are available for your own samples. For example, the Kurzweil K2000 uses sample numbers 1 through 199 as factory-defined samples and does not allow access to sample number zero. Sample numbers 200 through 999 are available for user-defined samples. Also note that a sample number of 200 specified in the *Sampler Tool* actually addresses sample number 201 on the Kurzweil K2000. (Confusing, isn't it?)

The Akai S1000, on the other hand, uses sample numbers 0 through 3 to store its factory-defined samples. Sample numbers above 3 are available for user-defined samples. And the list of inconsistencies goes on and on.

Wouldn't it be nice if you could simply use sample numbers 0 through n for your samples no matter what sampler you are using? This is what the *Sample bias* allows you to do. Once you figure out what sample number range is available for your custom samples, and determine whether the sampler counts in the same manner as the *Sampler Tool*, you can enter the sample bias and forget about it.

For the above examples, a sample bias of 199 works perfectly for the Kurzweil K2000 (remember it adds one, so sample number 199 actually addresses

sample number 200), and a sample bias of 4 makes the Akai S1000 a bit easier to manage. Once you have entered the *Sample bias* for your sampler, you can use sample numbers 0 through n as the *Logical send/receive sample number* and the *Sampler Tool* will convert it to the *Actual send/receive sample number* automatically.


You can also use the *Sample bias* to define different bases for multiple projects on which you are working. For example, if you are composing two different pieces that use different samples, you might want to create a sampler configuration for each of these projects (see the *Saving Sampler Configurations* section below). Within each configuration, you can define a different sample bias to keep your samples separated in the same sampler.

If you are still a little fuzzy on how the *Sample bias* option works, don't sweat it. As soon as you use it once or twice, it'll become old hat. Try different bias values to see what happens (do this before you start working on a critical project). The worst that can happen is you will overwrite a sample that you didn't mean to, and all writeable samples should be restorable from your disks. The section titled *Sending and Receiving Samples* also contains more information on using the *Sample bias*.

SAVING SAMPLER CONFIGURATIONS

After filling in all options for a sampler configuration, you can save it as a preset by selecting the *Save As* button in the *Sampler Configuration* dialog. Type a descriptive name into the *Save Preset* dialog and press *OK*. Once a sampler configuration preset is saved, you can choose the configuration by name from the *Sampler* dialog.

The following picture shows a configuration being saved for the Kurzweil K2000 (SCSI/SMDI). Note that a preset name can be anything you want, so use whatever names with which you are most comfortable.



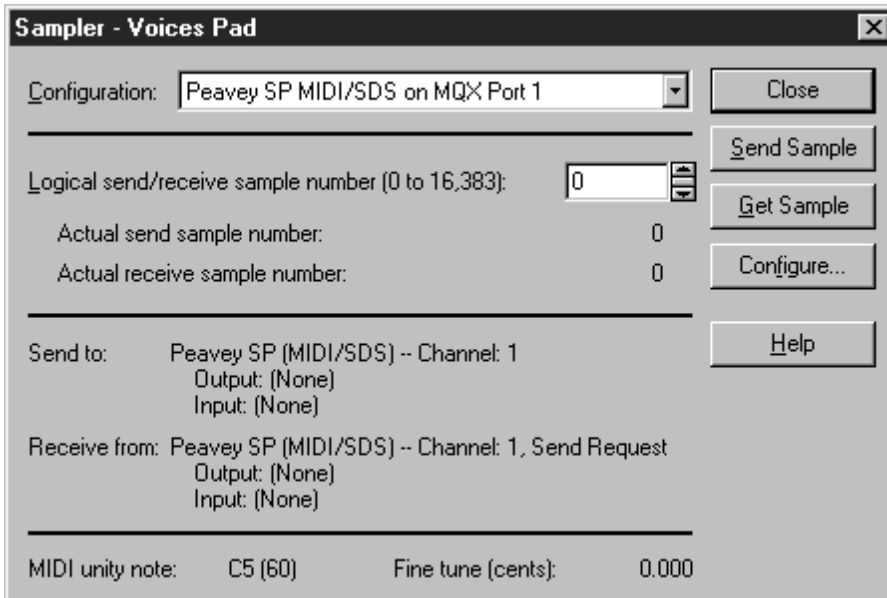
You can create and save as many sampler configuration presets as you want. You might want to create configurations for all of your samplers and presets for specific projects you are working on. Having presets for each of your samplers and projects allows you to quickly switch back and forth without reconfiguring the *Sampler Tool*.

If you want to get rid of a preset, select it from the *Name* drop-down list at the top of the *Sampler Configuration* dialog and press the *Delete* button.

SENDING AND RECEIVING SAMPLES

Once your sampler setup is configured, you should be able to send and receive samples by using the *Send Sample* and *Get Sample* buttons in the *Sampler* dialog. Select the sampler you want to *Send to* or *Receive from* in the *Configuration* drop-down list, enter the sample number you want to transfer in the *Logical send/receive sample number* field, and press the *Send Sample* or *Get Sample* button.

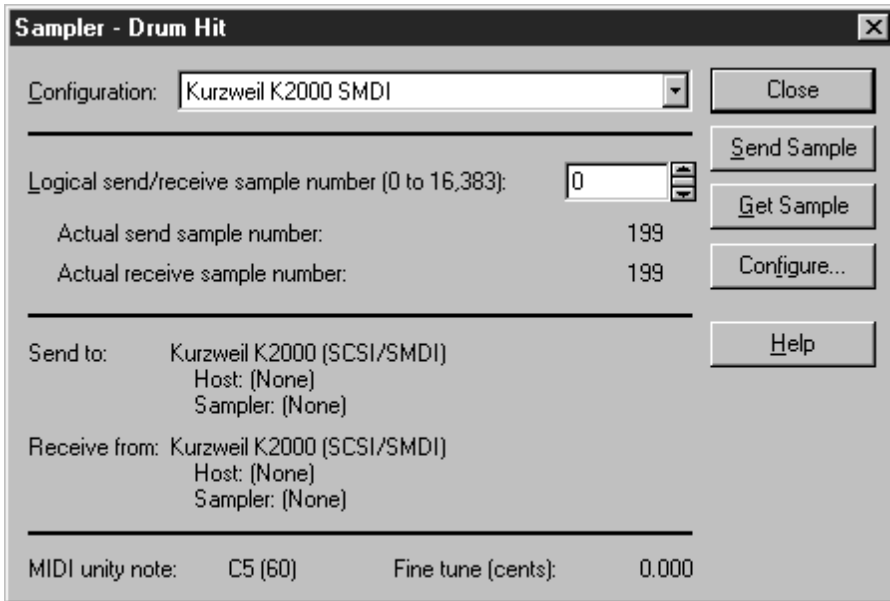
The following picture shows the *Sampler* dialog ready to send or receive sample number zero from a *Peavey SP* using *MIDI/SDS* on *Channel 1* across a *Music Quest MQX-32(M)* MIDI card.



Note that the *Actual send sample number* and *Actual receive sample number* are equal to the *Logical send/receive sample number*. This is because no sample bias is being used for this particular configuration.

Other areas of interest are the *Send to* and *Receive from* descriptions where all information relevant to the sampler configuration is printed. The fields on the bottom called *MIDI unity note* and *Fine tune* display the current values for the active sample. The C5 (60) displayed for the *MIDI unity note* represents that the sample is tuned for Middle C. The *Fine tune* value is set to 0.0 cents, so no minor tuning differences exist. These values are set using the *Edit Sample* command under the *Special* menu and may be used by samplers that support tuning information for samples. It is important to note that *Sound Forge* does not use this information, it is for samplers only.

As another example, the following picture shows a setup ready to send or receive sample number 199 (which is really 200) from a *Kurzweil K2000* using *SCSI/SMDI* through an *Adaptec AHA SCSI* card.



The most notable difference between this example and the previous one is the use of a *Sample bias* of 199 for the send and receive sample numbers.


To send or receive a sample after selecting the sampler setup from the *Configuration* drop-down list, you need only press the *Send Sample* or *Get Sample* button. The transfer will begin and a *Progress Meter* will be displayed in the status bar of Sound Forge. You can cancel a transfer at any time by pressing the *Cancel* button to the left of the *Progress Meter* or by pressing the *Escape* key.

Sending a sample will always send the entire contents of the active data window. Any selection contained in the waveform display is ignored. When receiving a sample, the sample will always be placed in a new data window.

If you're having trouble, verify that the sampler is set up correctly for sending or receiving samples by referring to the sampler's documentation on sample dumps. For further information on using your sampler with the *Sampler Tool*, refer to the Appendix at the end of this manual. These sections contain detailed information about several samplers that have been used successfully with the *Sampler Tool*. You may want to read some or all of these sections even if your sampler is not mentioned. There are tips you may find helpful when working with your sampler.

MIDI KEYBOARD SAMPLER TOOL

With the Sound Forge *MIDI Keyboard*, you can control internal or external synthesizers and samplers from Sound Forge. For example, you might want to test a sound in a sampler after sending (downloading) it from Sound Forge. You can also use the *MIDI Keyboard* to listen to the different sounds on a synthesizer, in the synthesis section of your sound card, or to practice your mouse keyboard chops.



OPENING THE MIDI KEYBOARD

First, open the *MIDI Keyboard* window by selecting *Keyboard* from the *View* menu. Like any other toolbar or shortcut menu window in *Sound Forge*, you can resize it, move it around, or dock it to any side of the *Sound Forge* window. Depending on how you resize the *MIDI Keyboard*, the keys will be arranged horizontally or vertically.

When you left-click on any of its keys, the *Keyboard* automatically turns on if it's not already on. If you don't hear anything, it's probably because its output has not been connected to a MIDI output device.

CONFIGURING THE MIDI KEYBOARD OUTPUT PORT AND CHANNEL

To select a MIDI device, left-click on the *MIDI Out* button. The *MIDI Out* shortcut menu will appear listing all the devices available for output, like your sound card or other MIDI port. Select the desired device. At the bottom of the list is an option to have the *MIDI Keyboard Send Program Changes* or not. Check it if you want to be able to select instruments in the device from the *Keyboard*.

Next, you must select the MIDI channel to be used on the chosen device. Most MIDI devices are configurable to accept MIDI commands on any channel.

STILL DON'T HEAR ANYTHING?

If, after selecting the correct device, you still can't hear anything when you play on the *Keyboard*, check the following things:

- Set the MIDI output velocity of the *Keyboard* to a high value (above 100).
- Set the MIDI input channel in your sound module to the same channel as the *Keyboard* channel. Also, make sure the device is set to receive MIDI input.
- Check the device output volume level. (Is the mixer level for your sound card or the output volume of the sampler and speakers at a high enough setting?)
- For external devices, check your MIDI cables.

SELECTING INSTRUMENTS

To change to a different voice in the synthesizer, select the instrument name from the instrument list and press a key. Patch names are arranged as specified in the General MIDI Standard (with synthesizers not using the General MIDI convention, use the patch number instead of the instrument name). Make sure that you have the *Send Program Changes* option under the *MIDI Out* shortcut menu checked, otherwise you won't be able to switch between patches. In addition, program changes are only sent when a key is pressed.

You can also have Sound Forge generate chords instead of single notes by selecting a chord structure from the chord structure drop-down list (default: *Note*). A chord is then generated using the pressed note as the root of the selected chord or interval.

CREATING AND EDITING LOOPS

A loop is a region in your sound file which is repeated on playback. Since samples are not infinite in length, in most samplers loops are used to create sounds which sustain for a long time. For example, when a sampler is told to sustain a flute sound for a long period of time, it loops a short section of the sample over and over. Loops are also used to repeat sections of music, although for these purposes it is recommended that you use the *Playlist*.

WAYS TO CREATE LOOPS

To create a loop, first select a section of the sound file. You can hear how this section will sound in looped mode by selecting the *Play Looped* button on the playbar. Next, select the *Edit Sample* item from the *Special* menu. In the *Edit Sample* dialog, set the *Sample type* to *Sustaining* and switch the *Loop count* from *Infinite loop* to 3 and select *OK*. You could have also created the loop by right-clicking on the ruler and selecting *Edit Sample* or hitting "L" on the keyboard.

The data window now has tags in the ruler specifying the loop end points. You can play the sound file with the loop counts by selecting the *Play as Sample* button on the playbar. This plays the entire file, with the looped section repeated three times.

EDITING A LOOP WITHOUT THE LOOP TUNER

Once you've created a loop, you can change its length, start, or end points by sliding the loop markers in the ruler. Left-click on the loop start marker. While holding the left button down, slide the marker to a new location and then release the button. It's as simple as that. You can accomplish the same thing by right-clicking on the marker and selecting *Edit* from the shortcut menu. In the *Edit Sample* dialog, you can then manually enter *Start* and *End* points for the sample.

Other options available when right-clicking on a loop point marker are:

SELECT

Modifies the current selection to be equal to the loop end points.

DELETE

Deletes the loop markers.

UPDATE

Modifies the loop end points to be equal to the current selection.

LOOP TUNER SAMPLER TOOL

Have you ever tried to create loops on an external sampler using the sampler's microscopic LCD display? Only then will you fully appreciate the ease of use of *Sound Forge's Loop Tuner*. With the Loop Tuner, you can create loop points visually and avoid obvious sound wave glitches which result in annoying pops.

LOOPING TECHNIQUES

Creating a natural sounding loop is not something that comes naturally. There are many factors that can cause annoying pops and glitches. While experimentation is always required, here are a few basic guidelines to follow:

MATCH END POINT AMPLITUDES

It's most common to select the end points with amplitude of zero (zero-crossings).

MATCH END POINT WAVEFORM SLOPE

If the waveform changes slope drastically, a pop will be heard.

MATCH END POINT SOUND LEVEL

What this means is that the overall loudness near (~ 100 ms) the end points should be equal. It is often very difficult to avoid this with quickly decaying sounds, since the start point will always be greater than the end point sound level.

AVOID VERY SHORT LOOPS

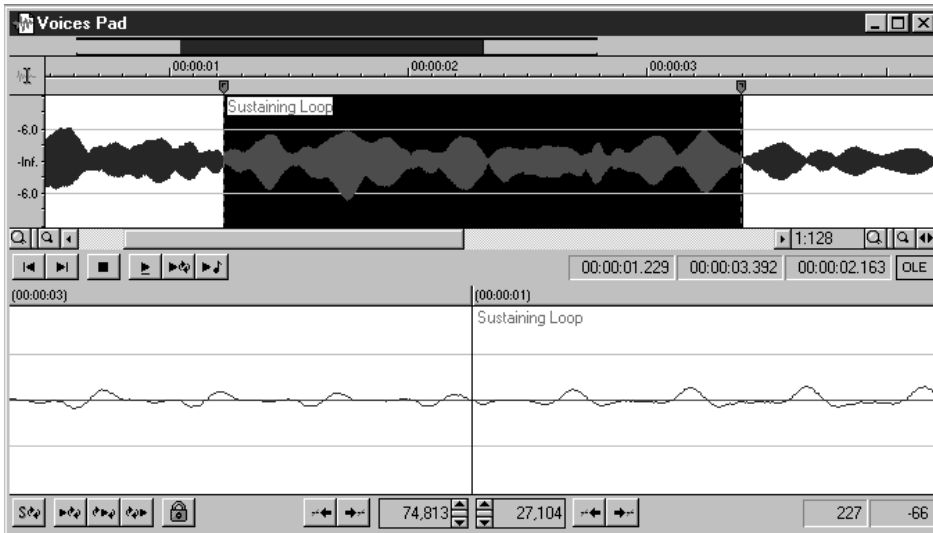
If the loop is smaller than ~ 50 ms ($1/20$ Hz), you will most likely create a loop with a pitch not equal to the sample pitch. You can experiment with pitch-tuning a loop by creating short loops with a length equal to $1/\text{frequency}$. For example, if you had a sample of pitch 440 Hz which corresponds to A5 on the Keyboard, you can try making the loop 2.27 ms. Pitched loops, however, don't sound like the original sample.

TRY THE CROSSFADE LOOP TOOL

If you still can't get the loop to sound the way you want, you can try using the *Crossfade Loop* tool in the Tools menu. This tool allows the end of the loop to be crossfaded with the end of the pre-loop material to eliminate some of the whoomping that may occur when the loop repeats itself. Some experimentation is required and not everything can be effectively looped, but this tool may help to find that perfect loop.

OPENING THE LOOP TUNER

Before opening the *Loop Tuner* window, you must have created a loop in the sound file. For this demonstration, open the file TUTORSM.P.WAV. This file has a *Sustaining Loop* already specified in it.



Under the View menu, select *Loop Tuner*.

The *Loop Tuner* allows you to view both loop end points at once. The left-hand side of the *Loop Tuner* window shows the end of the loop while the right-hand side shows the start of the loop. A seamless loop will not have any discontinuities; in other words, the waveform should look smooth as it goes from the end of the loop back to the start.

PLAYING FROM THE LOOP TUNER

There are three playback buttons in the *Loop Tuner*: the *Play Pre-Loop*, *Play Loop*, and *Play Post-Loop*. You can also still use the *Play as Sample* button in the data window to hear the entire sample with repeated loops.

SWITCHING BETWEEN SUSTAIN AND RELEASE LOOPS

When your *Sample type* is *Sustaining with Release* you can switch between *Sustaining* and *Release* loops in the *Loop Tuner*. Press the left-most button in the *Loop Tuner* playbar. An "S" indicates *Sustaining*, while an "R" indicates *Release*. Keep this on *Sustaining* for now (the "S" should be showing).

FINE-TUNING LOOP END POINTS WITH THE LOOP TUNER

If you want to change the loop end points by a small amount, you can use the *Start point* and *End point* edit box spinners. To increment the loop start point by one sample, click on the start point edit box spinner up arrow.

If you left-click on the center of the spinner, you can increment or decrement the *Start point* by larger amounts. Slide the mouse up and down while holding the left-button down. Notice how the start point waveform (right side of the *Loop Tuner* display) and the *Loop Marker* in the data window slide along horizontally in time.

On the lower right-hand side of the *Loop Tuner* display, the sample amplitude at the loop end points is displayed. Although it is highly dependent on the waveform, a good rule of thumb is that the closer these two values are to each other the better.

You can also click and drag either of the rulers at the top of the *Loop Start* and *Loop End* displays to move the corresponding loop points.

NOTE

ZERO-CROSSING FINDER

As we mentioned earlier, loop end points should match amplitude and slope. The *Zero-crossing Finder* accomplishes this by finding the next zero-crossing of the waveform with a positive slope (you can select the slope in the *Editing* folder in the *Preferences* dialog).

There are two *Zero-crossing Finders* for each loop end. One for finding the next zero-crossing and one for finding the previous zero-crossing. To use the finders, just left-click on one of the buttons. By trying out different locations, moving the start and end points forward and backward you should be able to quickly find loop ends that meet smoothly.

Press on the *Play Loop* button in the *Loop Tuner* to hear the current loop. If the loop still clicks when played, try moving to different zero-crossings. Remember, you can also use the loop point spinners to fine tune the end point location.

LOCK LOOP LENGTH

Let's assume that you want to shift a loop's position without changing its length. Pressing the *Lock Loop Length* button allows you to do so. When on, any operation you use to change a loop point affects both the start and end of the loop, therefore keeping the loop length constant.

For example, with the *Lock Loop Length* button down, drag the *Loop Start Marker* in the *Data Window Ruler* to the left. The start and end markers will slide together. Turning the *Lock Loop Length* off allows you to shift each end point individually. The *Lock Loop Length* button works the same as the *Lock Loop/Region Length* item in the *Options* menu.

REAL TIME LOOP TUNING

The tools mentioned above for changing the loop end points can be used while the loop is being played. This allows for instant feedback on how the loop sounds. To do so, start playback of the loop by selecting the *Play Loop* button and proceed to use any method to change the loop length. *Sound Forge* will play the modified loop points.

SAVING LOOP POINTS

Loop points are automatically saved when you save your sound file in .WAV, .SDS, .SMP, .DIG, .PAT, and .AIF file formats. You are allowed to disable the saving of loop points when saving to the .WAV format.

MIDI/SDS SETUP AND TROUBLESHOOTING

MIDI/SDS HARDWARE REQUIREMENTS AND SETUP

To use the MIDI/SDS protocol with an external sampler that supports the MIDI Sample Dump Standard (SDS), you need a MIDI card (with MIDI input and output ports) installed in your computer.

Using MIDI cables, connect the MIDI out port of your sampler to the MIDI in port of the card and the MIDI in port of the sampler to the MIDI out port of the card. This is the same configuration you would use to connect a MIDI keyboard to your computer for sequencing.

Internal samplers do not require a MIDI card and MIDI cables. However, you may need to use open loop protocol when sending samples to an internal sampler. The documentation that comes with your sampler should specify the requirements for performing SDS transfers if it is supported.

MIDI/SDS TROUBLESHOOTING

When using MIDI/SDS to send or receive samples, it is usually in your best interest to not use the *Open loop* option. The *Open loop* transfer is slower and has a much greater chance of sending or receiving incorrect data. Reasons for using *Open loop* are:

- Your sampler does not support closed loop transfers (handshaking).
- You don't have enough cables to connect both the MIDI in and MIDI out ports.

Open loop can also help with troubleshooting SDS hardware setup problems. If you can't get the *Sampler Tool* to transfer data to (or from) your sampler, it may be helpful to try turning on the *Open loop* option to see if single cable transfers work. If open loop works but closed loop does not, you probably have one of the following problems:

- Your sampler does not support closed loop transfers (handshaking).
- Your MIDI cable or connection for MIDI input (send) or MIDI output (receive) to your computer is bad.
- Your MIDI card is not receiving MIDI input (send) or sending MIDI output (receive). Interrupt conflicts are a common problem for MIDI input.

SCSI/SMDI SETUP AND TROUBLESHOOTING

SCSI/SMDI HARDWARE REQUIREMENTS AND SETUP

To use the SCSI/SMDI protocol with an external sampler that supports the SCSI MIDI Device Interface (SMDI) under Windows 3.x, you must have an ASPI compatible SCSI card for your machine as well as the *WINASPI* software provided by *Adaptec*. The best way to acquire this setup is to buy an *Adaptec*

SCSI controller with the EZ-SCSI software. The EZ-SCSI software is required by the Sampler Tool for accessing your Adaptec SCSI card.

NOTE

Windows 95 includes ASPI drivers, therefore you will only need a compatible SCSI adapter.

If you own an older Adaptec card and it did not come with EZ-SCSI, you will need to contact your Adaptec dealer for information on purchasing EZ-SCSI. You may also contact Adaptec Pre-Sales directly at 1 (800) 959-7274.

If you do not own a SCSI card and are looking for an inexpensive solution for transferring samples to your SMDI compatible external sampler, the Adaptec AVA-1505 TotalCD kit is an excellent choice. This kit comes with EZ-SCSI Lite which provides everything you need for using the Sampler Tool with an external sampler (except a SCSI cable).

CABLING

To connect your SCSI card to your external sampler, you will need to use a SCSI cable. There are several types of SCSI cables, so make sure you get the right one for your SCSI card and sampler.

There are three types of connectors found commonly on external SCSI cables. One type is the Centronics 50 pin. This connector is the beast of connectors and if you own an Adaptec card other than an AVA-1505 or AVA-1515, this is probably the connector you will need to connect to your SCSI card. If you own an AVA-1505 or AVA-1515, the connector on your SCSI card is a DB-25 connector. This connector is a 25-pin connector and looks like the connector found on serial and parallel cables. The DB-25 connector is most likely used by your sampler as well. The third type of connector is the SCSI-2 type connector that is used with the popular Adaptec 2940 series of PCI SCSI adapters. This connector looks similar to a DB-25, but has smaller pins and many more of them.

If you are buying a cable to connect your sampler to your system, it is important that you get the terminology right when ordering. You will need to know the type of connector found on your SCSI adapter, as well as the connector on the sampler itself. Most samplers have DB-25 connectors and most SCSI adapters these days use Centronics or SCSI-2 connectors. Determine what your system uses and order appropriately.

NOTE

Warning: Do not use a serial, modem, or printer cables to connect your SCSI devices. You may damage your equipment by using these cables.

Once you have a SCSI card installed (refer to the card's manual for installation instructions), connect the card and the sampler's SCSI ports with a SCSI cable. Both your computer and your sampler must have the power turned off before connecting or disconnecting SCSI cables! Failing to turn the power off before connecting and disconnecting SCSI cables can damage your computer and/or hardware. SCSI connections are much more fragile than MIDI connections.

SCSI/SMDI TROUBLESHOOTING

Here are some of the common problems encountered with SCSI and samplers.

CONFLICTING SCSI ID'S

When connecting devices on a SCSI chain, each device must have a unique device identifier (ID). SCSI allows you to set up to eight unique ID values from 0 to 7. Typically device ID 7 is used for the internal SCSI controller card which leaves ID's 0 through 6 available for other devices. Also, if you have a bootable SCSI hard drive its ID must be set to 0.

Typical SCSI configuration:

ID	Device
0	Hard Drive
1	CD-ROM Drive
2-6	Available for use by samplers
7	SCSI Controller Card

OCCASIONAL TRANSFER FAILURES

If you occasionally receive messages such as *The SCSI Device is not responding* or *A problem was encountered while transferring the sample*, you may just have an upset SCSI bus. Try bringing up the *Sampler Configuration* dialog and reselecting your SCSI host from the drop-down list. Doing this causes a series of SCSI commands to be executed that in many cases will settle an upset SCSI bus. If

this fails to work, power down and restart all equipment (computer, samplers, etc.).

If the above mentioned messages are encountered frequently, you may have a more serious problem. The following section discusses several possibilities.

THE SAMPLER IS RECOGNIZED BUT DOESN'T TRANSFER RELIABLY

The following is a list of causes for unreliable SCSI transfers.

SYNCHRONOUS TRANSFER MODE

Some samplers (like the Kurzweil K2000) do not operate properly if there is a SCSI device set to synchronous transfer mode on the same SCSI chain. SCSI hard drives and CD-ROMs often have the option of using a synchronous transfer mode. If there is a host vs. device synchronous transfer option, make sure you select host. Refer to your SCSI device manuals for more information.

SCSI TERMINATION

If your SCSI chain is not properly terminated, you may experience unreliable SCSI transfers. Refer to your SCSI card and SCSI device manuals for more information.

LONG OR CHEAP SCSI CABLES

SCSI cables that are very long or not properly shielded can cause unreliable operation. Keep in mind that SCSI cables are not as forgiving as MIDI cables and overly long cables will not work reliably. Also beware of using cables that are not certified SCSI cables (such as tantalizingly cheap parallel printer cables); the extra money spent on a true shielded SCSI cable can save you endless hours of frustration.

ADAPTEC 1540/1542CF DOESN'T RECOGNIZE A SAMPLER

If you own the *Adaptec 1540/1542CF* (not the C version, the newer CF version) and it won't recognize your sampler you may need to make a change in the configuration of your *Adaptec* controller.

Some samplers do not want to operate when the *Reset SCSI Bus at Power-On* option of the *Adaptec* controller is enabled. This is the default operation for the *1540/1542CF* and you will have to disable this option to allow your system to work with your sampler.

DISABLING THE RESET SCSI BUS OPTION

You should review your 1540/1542CF owners manual before proceeding with this change.

N O T E

1. Reset your machine by pressing the hardware reset button or turning the machine off and then back on.
2. While your system is booting it will display the message, *Press Control+ A for the SCSISelect Utility*. Press *Control+A* when this message appears.
3. The screen will display the main screen of the *SCSISelect* utility and show the current IO port in use by the SCSI card. Press *Enter* to proceed to the *Options* menu.
4. In the *Options* menu press *Enter* on the *Configure/View Host Adapter Settings* option to proceed to the *Settings* screen.
5. Move the cursor using the arrow keys until the *Advanced Configurations Options* item is selected. Press the *Enter* key to proceed to the *Advanced Options* screen.
6. Move the cursor using the arrow keys until the *Reset SCSI Bus at Power-On* option is selected. Press *Enter* to change this option.
7. Select the *Disabled* setting by moving the selection with the arrow keys, and then press *Enter*.
8. Press the *Escape* key twice to return to the previous screens.
9. Answer *Yes* to the *Save Changes Made?* prompt.
10. Press the *Escape* key again to exit the *SCSISelect* utility and answer *Yes* when the *Exit Utility* prompt is displayed.
11. *SCSISelect* will reboot your machine after you have made the changes.

N O T E

Disabling the Reset SCSI Bus at Power-On option may keep other devices on the SCSI chain from resetting correctly when using the Soft Boot feature of your machine (i.e. pressing Control+Alt+Del). On some systems this will cause the system to hang when pressing Control+Alt+Del. To guarantee that devices are reset when rebooting with this option disabled, you should use the Reset button on your computer or turn the power off and back on to reset the machine.

MY SCSI/SMDI-COMPATIBLE MENU DOESN'T SHOW UP UNDER WINDOWS 95

The first thing to do is to check for correct SCSI termination and multiple devices on the SCSI chain using the same SCSI ID. If you've already done this and have an Adaptec SCSI card it is possible that you need to update the mini-port drivers in your system.

Adaptec has a series of updated mini-port drivers available for Windows 95 on both CompuServe and on the world wide web.

TO FIND THE DRIVERS ON COMPUSERVE

1. Log on to CompuServe and type in GO ADAPTEC. (With WinCIM press the GO button and type in ADAPTEC.)
2. In the library section labeled Windows 95 download the file WIN95MPD.EXE

TO FIND THE DRIVERS ON THE WEB

1. Go to the Adaptec site at <http://www.adaptec.com>
2. Navigate to the support section which contains updated drivers or you can go directly to the files page at the web address http://www.adaptec.com/techsupp/BBS_95.html
3. Click on the file name WIN95MPD.EXE to download the updated driver package.

UPDATING YOUR DRIVERS

After you have downloaded the file WIN95MPD.EXE you need to create a temporary directory or folder on your system and run the WIN95MPD.EXE program.

Follow the directions in the `readme.txt` file to update your drivers for your Adaptec SCSI card.

TIPS AND IDEAS

This chapter contains several tips and ideas for using the *Sampler Tool* with your sampler.

USING FM SYNTHESIS WITH YOUR SAMPLER

One nice use for the *FM Synthesis* tool in *Sound Forge* is to turn your sampler into an FM synth! Well, not quite. This procedure will not turn your sampler into a true FM synthesizer, or at least the types that have become most familiar. While *Sound Forge* will produce four operator sounds, this is still sampling. To get different pitches across the keyboard, the sampler will vary the sample rate of the original sample rather than produce a unique timbre based on the frequencies dictated by that pitch (key).

Also, any dynamic envelope and/or velocity control of the FM is not variable like that of true FM synthesizers. Nonetheless, sending samples to your sampler that were produced using the *FM Synthesis* tool of *Sound Forge* can yield some very cool results. If you have a sampler with synthesis capabilities, this technique becomes quite powerful.

1. First, create a sample using the *FM Synthesis* tool found in the *Tools* menu. The sample can be rather short (0.2 - 1.0 sec) and it is probably a good idea to make sure that some portion of the envelope is static so that a good loop is easy to attain.
2. Once you have created the sample, create a *Sustaining Loop* and edit it with the *Loop Tuner*.
3. When you have your sample the way you want it, follow the usual *Sampler Tool* procedure to transfer the sample to your sampler.
4. From there any further manipulation will need to take place inside your sampler. That's about it! Simple but effective.

CHANGE THE GAME BUT NOT THE RULES

It's easy to see that you can expand the processing power of a sampler by using *Get Sample* to retrieve a sample from the sampler, processing it with Sound Forge, and then using *Send Sample* to send the sample back to the same location in the sampler. This in itself is no mind-blowing discovery. However, there may be some inspiring side effects waiting to be uncovered.

Most samplers will replace samples already in memory locations with new samples received from Sound Forge if the *Logical send/receive sample number* matches the original location. Other samplers will automatically put a received sample in an empty location regardless of the sampler number. The following idea is only possible with a sampler that replaces a sample.

When the Kurzweil K2000 receives a sample via SCSI/SMDI, it will replace samples that have the same incoming send sample number (this is not true if you are using MIDI/SDS with the K2000). The interesting part of this is all *Keymaps* that reference that sample location will remain intact even if the sample has been replaced with an entirely different one. This can lead to some very unique and sometimes spine-tingling sounds.

BACK UP AND CATALOG SAMPLES ON YOUR COMPUTER

You can even store and catalog samples with your computer using the *Sampler Tool*. Saving your samples as .WAV files on your computer's hard drive or normal floppy disks allows you to store detailed summary information with each sample. Summary information can include descriptive titles, subject details, general comments, and more.

Unfortunately, *Keymaps* and other sampler specific information are not transferred from the sampler to the *Sampler Tool*, so none of it gets saved with .WAV files. If this is the case with your sampler and you still want to save your samples as .WAV files, you might be able to save *Keymaps* and other information without the sample data on a floppy disk. By saving this special information without the sample data (or a very small sample if the sampler requires it), you could store quite a few on a single floppy disk. Later, the *Keymaps* can be recombined with the sample data stored on your computer.

MIDI/SMPTE

In this section, we will cover how you can synchronize Sound Forge with devices generating SMPTE time code or use a MIDI sequencer or MIDI controller to trigger audio playback.

This chapter covers:

- What is MIDI?
- The Sonic Foundry Virtual MIDI Router
- Sound Forge and MIDI Triggers
- Sound Forge and MIDI Time Code Synchronization

WHAT IS MIDI?

The Musical Instrument Digital Interface (MIDI) is a set of commands that many pieces of music software and hardware use to speak to each other. MIDI is most often used for sending commands such as, “Play Middle C Now,” but can also be used to send information like, “Current Time is: 00:00:01:23 SMPTE,” or even digital sound data (see Sampling chapter).

The most common way to use MIDI is to have a master device (such as a MIDI sequencer) to generate MIDI commands to a slave device (such as a synthesizer, which plays a note when instructed). If both were in separate hardware boxes, you would run a MIDI cable from the sequencer’s MIDI out port to the synthesizer’s MIDI in port.

THE SONIC FOUNDRY VIRTUAL MIDI ROUTER

The *Sonic Foundry Virtual MIDI Router (VMR)* is a software-only driver that allows you to control Sound Forge from other MIDI-capable applications and vice versa under Windows 9x and Windows NT without using MIDI hardware.

A MIDI Router simply transfers MIDI data from one port to another. The *Sonic Foundry VMR* driver does exactly this but requires no hardware. The driver provides up to four devices for MIDI output and input. Each of the output devices sends all MIDI data to its corresponding input device.

If one application sends MIDI data through the #3 Sonic Foundry MIDI Router output device, another application can receive this MIDI data as input from the #3 Sonic Foundry MIDI Router input device.

NOTE

The first character of the device name is a device number from 1 to 4. The device number is placed at the beginning of the name for compatibility with sequencer software that tries to display device names in very small places.

The Sonic Foundry VMR can be very useful for synchronizing two MIDI capable applications. You can trigger Sound Forge from your MIDI sequencer, or you can drive your MIDI sequencer from Sound Forge using SMPTE/MTC. This and more can be accomplished without using MIDI hardware for routing.

INSTALLING THE SONIC FOUNDRY VIRTUAL MIDI ROUTER

The Sonic Foundry Virtual MIDI Router is included on the Sound Forge CD. To install the Sonic Foundry Virtual MIDI Router, follow the instructions below for the operating system you are running.

WINDOWS 9X

1. Open Control Panel from the Settings menu in the Start menu.
2. Open the Add New Hardware Item in the Control Panel.
3. Click the Next> button in the first page of the Add New Hardware Wizard.
4. In the Detect New Hardware page of the Wizard set the Detect New Hardware option to No.
5. Click the Next> button.
6. In the Select type of hardware page select Sound, video, and game controllers.
7. Click the Next> button.
8. In the Manufacturer and model page click the Have Disk... button.
9. Insert the Sound Forge CD into your CD-ROM drive.
10. In the path list box, type in E:\Extras\Virtual MIDI Router\Win9x (where "E" is your CD-ROM drive) and click OK.

11. Select the *Sonic Foundry Virtual MIDI Router* from the *Models* list and click *OK*.
12. Click the *Finish* button in the final *Add New Hardware* page.
13. Once you have restarted Windows, all programs that support MIDI (including *Sound Forge*) can use the *Virtual MIDI Router*.

WINDOWS NT 4.0

1. Open *Control Panel* and double-click the *Multimedia* icon.
2. Click on the *Devices* tab and then press the *Add* button.
3. Select *Unlisted/Updated driver* and click *OK*.
4. You can now insert the *Sound Forge* CD. Click the *Browse* button and select the drive in which you inserted the CD. Double click the *Extras* directory, then the *Virtual MIDI Router* directory, then the *WinNT* directory. Click *OK*.
5. The driver should now be listed as *Sonic Foundry Virtual MIDI Router*. Click once on the driver to select it and click *OK*.
6. Windows will now install the driver. When it is finished you will be prompted to restart Windows to complete the driver installation. There are no further setup procedures.
7. Once you have restarted Windows, all programs that support MIDI (including *Sound Forge*) can use the *Virtual MIDI Router*.

CONFIGURING THE SONIC FOUNDRY VIRTUAL MIDI ROUTER

After installing the *Sonic Foundry Virtual MIDI Router (VMR)*, a new MIDI device is created with the name *Sonic Foundry MIDI Router*. If you want to install additional devices for MIDI routing (the *VMR* supports up to four) you need to configure the *VMR* .

*When only one routing device is used, the name of the driver is displayed as *Sonic Foundry MIDI Router*. When multiple ports are used each port is labeled starting with a number (1 - the number of ports) to identify each connection.*

NOTE

The following procedures describe configuring the number of VMR devices for Windows 9x and Windows NT 4.0.

WINDOWS 9X

NOTE

Warning: If you have a large number of MIDI devices in your system, or a multi-port MIDI board such as an eight port MIDI interface, you may have problems installing additional VMR MIDI interface ports. Windows 9x will crash on startup if more than 11 MIDI devices are installed. If you have trouble with MIDI devices causing system errors after installing the VMR, you should reduce the number of ports in use by the VMR, remove the VMR, or reduce the number of MIDI devices in your system.

1. Double-click on My Computer.
2. Open the Control Panel folder.
3. Double click on the Multimedia icon.
4. Select the Advanced Properties page.
5. Expand the MIDI Devices and Instruments list by selecting it and double-clicking.
6. Select the Sonic Foundry MIDI Router device and click the Properties button.
7. In the General page of the Properties for the Sonic Foundry MIDI Router click the Settings button.
8. In the Configure Sonic Foundry VMR dialog select the number of ports you would like to enable from the Ports combo list.
9. Click OK in the dialog.
10. Click OK in the Sonic Foundry MIDI Router Properties dialog.
11. Click OK in the Multimedia Properties dialog.
12. Restart Windows for the new number of devices to take effect

WINDOWS NT 4.0

1. From Control Panel, click on the Multimedia icon.
2. Click on the Devices tab.
3. Double-click on the MIDI Devices and Instruments icon.
4. Select the Sonic Foundry Virtual MIDI Router from the list of drivers and press the Properties button.
5. Click the Setup Settings button in the Properties window.
6. In the Configure Sonic Foundry VMR dialog, select the number of ports you would like to enable from the Ports combo list.
7. Click OK in all of the the dialogs.
8. You will need to restart Windows and to have the settings take effect.

SOUND FORGE AND MIDI TRIGGERS

There are many different ways to generate MIDI commands, including MIDI sequencers, MIDI keyboards, MIDI guitars, and there's probably MIDI maracas out there somewhere. Whatever their source may be, MIDI commands can be used to trigger audio playback in Sound Forge.

TRIGGERING FILE PLAYBACK FROM THE SONIC FOUNDRY MIDI KEYBOARD

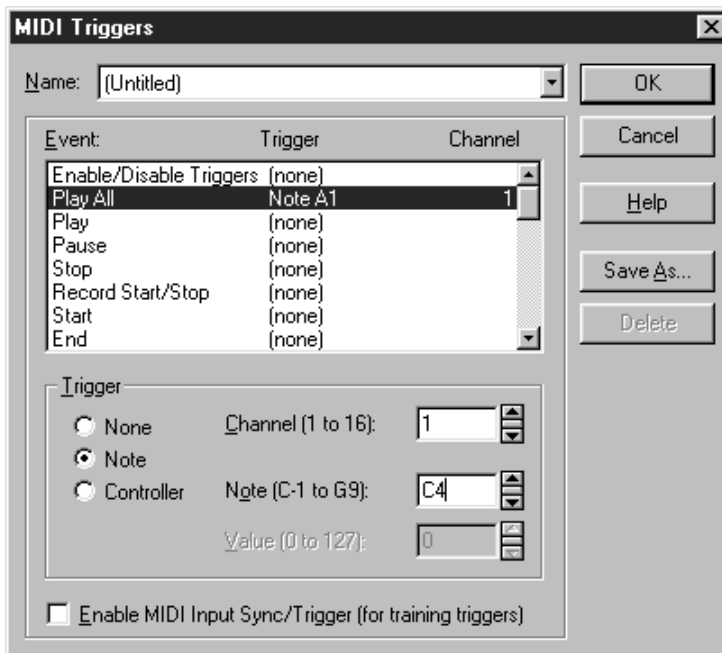
To demonstrate how to use MIDI Triggers, we will first show you how to control Sound Forge using the internal MIDI Keyboard. This might not be the most useful way of using MIDI triggers, but since everybody using Sound Forge will have access to the Sonic Foundry MIDI Keyboard, this will be a good demonstration of using triggers and the VMR. We are assuming that you have the driver installed.

First, open the Keyboard window by selecting Keyboard under the Window menu or by pressing Alt+3. Even though the Keyboard window is intended for controlling MIDI devices like sound cards or samplers, it can also be used to trigger Sound Forge itself. Next, open a sound file that you wish to trigger.

The first step in setting up MIDI triggering is to make sure that the output of the master goes to the input of the slave. In this case, the Sound Forge MIDI Keyboard is the master and Sound Forge is the slave. Left-clicking on the MIDI Out button in the Keyboard allows you to set where the Keyboard will send its MIDI commands. For this example, set it to 1 Sonic Foundry MIDI Router. Also, make sure that the output channel is set to Channel 1. You won't be able to hear anything at this point when you press on the Keyboard because nothing is connected to the other end of the VMR.

Now, go into the MIDI/ Sync folder under Preferences in the Options menu. For Sound Forge's MIDI In, also select 1 Sonic Foundry MIDI Router. At this point, we've connected the two devices, much like running a MIDI cable between them.

Next, you must select MIDI Triggers under the Options menu to determine what you want to trigger in Sound Forge. For this example, select Play All from the Event list box. In the Trigger section, select Note. Type in the note C4 and set the channel to 1. What this does is tells Sound Forge to play the current sound file when it receives a C4 Note-On command on MIDI Channel 1.



Finally, you must tell Sound Forge to begin waiting for MIDI commands. Select the *MIDI Input Sync/Trigger* item in the Options menu to enable synchronization to MIDI input. Or in the *Regions/Playlist* toolbar, select the *MIDI Input Sync/Trigger* button. (If the toolbar is not open, go to the *Toolbars* folder under *Preferences* in the Options menu and check the *Regions/Playlist* option). The *MIDI Input Sync/Trigger* button will remain pressed down and *Waiting...* will appear in the left-hand status box. Sound Forge is now ready to accept MIDI commands.

To test triggering, press the *C4* key in the *MIDI Keyboard* (the key with the number 4 on it). The sound file should start playback and play all the way through. Notice that when you press *C4*, the status box displays that it has received a *C4 Note-On* message on *Channel 1*.

It's important to understand the difference between triggering playback with the *MIDI Keyboard* and selecting the *Play* button in Sound Forge. When you press a note in the *MIDI Keyboard*, MIDI commands are sent from Sound Forge to an external device (in this case the *VMR*) and then back to Sound Forge's *MIDI Input* port. The MIDI commands which arrive at the Sound Forge *MIDI Input* port could be coming from any software or hardware device that generates MIDI.

USING AN EXTERNAL MIDI CONTROLLER WITH SOUND FORGE

There are a large number of MIDI trigger commands in the *MIDI Triggers* dialog box. These commands include transport, marker/region creation, and navigation commands (for a complete list of commands, refer to the Reference chapter.) This enables external control of Sound Forge not only for performance but also for everyday editing.

There are a variety of MIDI controller boxes on the market which include transport controls (like play, stop, rewind,), function keys, and jog wheels. To use one of these devices with Sound Forge all you need do is enable *MIDI In* synchronization and then configure the *MIDI Triggers* dialog to respond to the corresponding control on your external device.

Sound Forge comes with a system preset to use the sequencer, cursor, dial, and function buttons on the Roland MCR-8 Multi-Controller. This allows you to move the cursor in Sound Forge by rotating the dial and drop markers by clicking the dial. You can zoom in and out of the wave form by using the cursor keys (vertically and horizontally) and play and stop audio with the sequencer controls. If you have an MCR-8 you don't need to use our default configuration, but can configure the buttons to be whatever you find most comfortable for editing. Just select a function and press the key on your controller to configure it.

If you are using Sound Forge with an external controller for dropping markers or creating regions you may want to disable the *Always ask for Region, Marker, and Loop names* option in the *General Preferences* folder. This will keep Sound Forge from popping the creation dialog when adding a new marker or region from your MIDI controller. When this option is disabled Sound Forge picks a default name for the new region.

TRIGGERING REGION PLAYBACK FROM THE SONIC FOUNDRY MIDI KEYBOARD

Let's now trigger regions from the *Regions List*. First, go to the *MIDI Triggers* dialog (Options menu) and set all the triggers to none by selecting the *All triggers off* preset from the *Name* drop-down list. If you have triggers in both the *MIDI Triggers* dialog and the *Regions List*, make sure that they correspond to different MIDI notes if you don't want unexpected results to occur. Next, open a file that has regions specified in the *Regions List*. Double-click on the first region in the *Regions List* and select *MIDI: Note - On Play*. Also, set the note to *C4*.

Make sure that the output of the *MIDI Keyboard* and the *MIDI Input* of Sound Forge are set to *Virtual MIDI 1* and all channels set to 1 as described in the previous section. Press down the *MIDI Input Sync/Trigger* toolbar button to synchronize *MIDI Input* (you can also do this in the Options menu).

If everything has been done correctly, when you press *C4* in the *Keyboard* the first region will be played.

By repeating the previous steps, you can give each region a MIDI trigger. Set the MIDI notes to different values, like *C3* and *C5* so notes on the *Keyboard* correspond to different regions.

Sound Forge can play only one region at a time. Overlapping causes the earlier region to cut off and the new region to be played back.

TRIGGERING FROM INTERNAL SEQUENCERS

The basic concepts of MIDI in and MIDI out routing used in triggering from the internal MIDI Keyboard apply to any hardware or software device that can generate MIDI commands. You should be able to follow the instructions for triggering with the MIDI Keyboard and substitute your device's output port and channel number.

Suppose you have a MIDI sequence in your sequencer you want to add a vocal track to starting at a set time. The easiest way to do this is to attach a MIDI trigger note to the sound file in Sound Forge and make the sequencer play the trigger note at the correct starting time. You can also accomplish this using SMPTE times as shown in the next section.

TRIGGERING FROM EXTERNAL DEVICES

The same procedure for internal devices applies for triggering with external devices. You will have to set Sound Forge's MIDI Input port (under MIDI\Sync Preferences) to the external device output port.

SOUND FORGE AND MIDI TIME CODE SYNCHRONIZATION

MIDI Time Code (MTC) is a way to transmit SMPTE timing signals between devices for synchronization. For an overview of SMPTE, check out the Appendix which explains the differences in 29.97 FPS, 30 FPS, Drop and Non-Drop, and other confusions. Although MIDI Time Code is most often used to sync audio tracks to visual action on video, it can also be used to synchronize other playback devices, like a MIDI sequencer, to Sound Forge.

Sound Forge has the ability to synchronize to external MTC or generate MTC for other devices to follow.

PLAYING REGIONS USING MIDI TIME CODE FROM A SOFTWARE MIDI SEQUENCER

It is possible to specify start times for each region in your *Regions List* for synchronizing digital sounds with other timed events. For example, if your MIDI sequencer or other device generates SMPTE time code, Sound Forge can make its SMPTE time accurately follow the device's SMPTE time and play the regions at the specified times.

Synchronizing to MTC is very similar to synchronizing to other MIDI events. However, if you just want to trigger a few sounds in Sound Forge from your sequencer, it might be easier to use Note-On MIDI Triggering and not worry about MTC or SMPTE synchronization (refer to the section on Triggering from Internal Sequencers). MTC synchronization uses up more of your computer's processing power, yet it's the only way to go if you don't want to (or can't) generate MIDI triggers. Also, MTC synchronization allows you to specify accurate SMPTE start times in Sound Forge rather than other software or hardware devices.

Here are the basic steps to follow when triggering a sound region in Sound Forge from a sequencer using MTC:

1. In Sound Forge, open a file and create a region in both the *Regions List* and the *Playlist*.
2. Double-click on the region in the *Playlist*. Select *SMPTE: Play at Time* from the *Trigger* drop-down list.
3. Set the SMPTE start time for the region. Sound Forge displays SMPTE time as hours:minutes:seconds:frames. Make sure that you are using the same SMPTE time format and frames per second in both Sound Forge and the sequencer. Refer to *Status Format and Status Preferences*.
4. Now, in the *MIDI/ Sync* folder under *Preferences*, make sure that the Sound Forge MIDI Input port is set to the same port as the Output port of the sequencer. For example, set both ports to 1 Sonic Foundry MIDI Router.
5. Verify that in the *MIDI Triggers* dialog (Options menu) all the triggers are set to (none). This prevents other MIDI commands from also creating triggers.

6. Enable MIDI Input Sync/Trigger mode by either pressing the MIDI Input Sync/Trigger toolbar button or selecting MIDI Input Sync/Trigger from the Options menu. Now, Sound Forge is ready and waiting for synchronization to MTC.
7. Switch to your sequencer. Set the MIDI output port to be the same as Sound Forge's MIDI Input port. Refer to the sequencer's manual for information on setting MIDI ports.
8. Enable MTC output in your sequencer. (Refer to the sequencer's manual for information on how to do this.) If your sequencer supports generating MTC you should now be ready to sync.
9. Press play in the sequencer. In Sound Forge, you should see the MIDI In status box display the same SMPTE time as the sequencer's time. At the set SMPTE time, Sound Forge should play the region without interrupting your sequencer's output to any other devices.

PLAYING REGIONS USING MIDI TIME CODE FROM AN EXTERNAL DEVICE

If you have a hardware device (i.e. SMPTE time code to MTC card) that generates MTC, the procedure is basically the same as outlined above. In Sound Forge's MIDI/Sync Preferences folder, you should be able to directly select the device's MTC output driver as Sound Forge's MIDI Input port.

USING SOUND FORGE TO GENERATE MTC FOR A MIDI SEQUENCER

Follow these steps if you want Sound Forge to generate MTC for other devices to follow.

1. Open a large sound file. Sound Forge only generates MTC while playing a sound file or playlist.
2. Now, in the MIDI/Sync folder under Preferences, make sure that the Sound Forge MIDI Output port is set to the same port as the output of the sequencer. For example, set both ports to 1 Sonic Foundry MIDI Router.
3. Verify that in the MIDI Triggers dialog (Options menu) all triggers are set to (none). This prevents other MIDI commands from also creating triggers, which might create unexpected results.

4. Enable MIDI Output Sync mode by either selecting the MIDI Output Sync toolbar button or selecting MIDI Output Sync from the Options menu. Now, Sound Forge will output MTC whenever you start playback.
5. Switch to your sequencer. Set the sequencer's MIDI input port to be the same as Sound Forge's MIDI Output port. (Refer to the sequencer's manual.)
6. Set the sequencer's SMPTE offset time as required. Sound Forge always uses 00:00:00:00 as its output start point. For some sequencers, it is suggested that the SMPTE offset time be set to at least four seconds to ensure correct synchronization.
7. Enable MTC input (again, check the sequencer's manual). This usually entails switching the timing generator in the sequencer from internal to MTC. You should now be ready to sync.
8. Press play in the sequencer (some sequencers don't require you to do this). The sequencer should not start playback. Instead it should switch to a *Waiting for MTC* mode.
9. In Sound Forge, start playback in the data window or Playlist. The sequencer should start following Sound Forge's SMPTE time when the SMPTE offset time is reached.

USING SOUND FORGE TO GENERATE MTC FOR AN EXTERNAL DEVICE

If you want Sound Forge to send MTC to an external device, follow the previous instructions for sending MTC to a sequencer with a few changes. Instead of sending to the VMR, have Sound Forge's MIDI Output port send directly to the device's MIDI driver. Check the device's manual for more detailed information on using MTC.

SOUND FORGE REFERENCE

This part of the manual contains a detailed explanation of every operation available to the user. It's intended for the user who is familiar with using Sound Forge. If you aren't familiar with Sound Forge's windows and basic controls, we suggest you read through the Using Sound Forge chapter first. This section is organized in the same order as the on-screen menus to allow for quick reference to information about all the Sound Forge menu and dialog options.

FILE MENU

NEW

The *New* command creates a new sound data window.



SAMPLE RATE

A *Sample rate* can be chosen from the drop-down list or typed in as a custom rate.

SAMPLE SIZE

Sample size selects whether the new data will be 8-bit or 16-bit.

CHANNELS

Channels selects whether the new data will be mono or stereo.

MAXIMUM EDITING TIME

Maximum editing time shows the maximum time currently available on the hard drive for editing a file of the specified type. Note that files open in Sound Forge are edited in 16-bit format and are only converted to 8-bit or compressed audio when closing the file.

NOTE

Using the *New* toolbar button (or holding the *Shift* key down while selecting *New* from the menu) will skip the *New* dialog and use the previously selected *Mono/Stereo*, *8-bit/16-bit*, and *Sample Rate* settings.

Shortcut: *Control+N* or Drag and drop a selection to the *Sound Forge* desktop to create a new file from the selection or select a portion of the *Sound Forge* desktop will create a new file with the most recently used attributes.

OPEN



Use *Open* to load a sound file into *Sound Forge*.

LOOK IN

Look in lists available drives and directories.

FILE NAME

Type the *File name* you want to open or a ***. and an extension to see a list of files with the corresponding extension in the current directory.

FILES OF TYPE

Files of type shows the type of files displayed. Choose a type from the drop-down list. If you select *All types (*.*)*, *Sound Forge* will try to auto-detect the file format.

If you know that the file format is unsupported, select *Raw File (*.*)*. Hitting *OK*, will take you to the *Open raw data* file dialog where you'll be able to specify format parameters.

MRU FOLDERS

Use *MRU folders* to speed the opening of regularly used files in different folders.

OPEN AS READ-ONLY

When *Open as read-only* is checked, the sound file will be opened but you will not be able to alter the data in the file. This feature is useful if you only need to play the file or copy sections from the file. You can still change the *Regions List*, *Playlist*, and *Summary* information for the file, but these changes must be saved to a new file.

When you open a sound file, Sound Forge defaults to whichever edit mode you last selected (either Open as read-only or Operate directly on the sound file). You can change modes in the Open dialog.

NOTE

OPERATE DIRECTLY ON THE SOUND FILE

When *Operate directly on the sound file* is checked, the sound file will be opened for editing but with no temporary file created. This greatly speeds the opening process but doesn't give you the security of a backup file. In most cases this is not a problem since all changes can be undone (if you have enough disk space), but if, for some reason, Sound Forge terminated improperly the file would remain in its edited state.

Whenever you open a file in Sound Forge, a backup copy is created so that the original file is not affected until you elect to save your work. When you open a file as read-only or direct, Sound Forge will not create a backup copy of the file, which makes opening files quicker and uses less disk space.

NOTE

AUTO PLAY

When *Auto play* is checked, a file selected in the window will automatically begin playback without opening it in Sound Forge.

MERGE L/R TO STEREO

Merge L/R to stereo will merge two mono files to left and right channels of a stereo file when opening.

PLAY

Play will audition the sound file without opening it in Sound Forge. If playing a file from a slow network or CD-ROM, some skipping might occur. This feature is only available for .WAV files.

Shortcut: Control+O or Control+F12 or Control+Alt+F2

OPEN RAW DATA FILE

This dialog is reached when you select Raw from the list of *Files of Type* in the Open dialog and attempt to open a file. Use this function to open a file that is not stored in one of the standard sound file formats supported by Sound Forge.

NAME

Name displays a list of preset formats available. To create your own format from the selected parameters, press *Save As*.

SAMPLE RATE

Sample rate selects the sample playback rate that will be used by Sound Forge when playing the file.

SAMPLE TYPE

Sample type selects the format used to store each sample value.

8-BIT PCM

8-Bit PCM is an uncompressed linear format.

16-BIT PCM

16-Bit PCM is an uncompressed linear format.

G.711 U-LAW

G.711 u-Law is a compressed format commonly used for telecommunications in the United States.

G.711 A-LAW

G.711 A-Law is a compressed format commonly used for telecommunications in Europe.

FORMAT

Format shows the sample format of data storage. Sound files on the PC are most often saved using the *signed* format.

CHANNELS

Use *Channels* to select the number of tracks stored in the file.

BYTE ORDER

The *Byte order* specifies the order in which the high and low bytes of a 16-bit sample are stored. Little Endian is used by Intel microprocessors, while Big Endian is used by Motorola microprocessors.

HEADER BYTES

Header Bytes specifies the number of bytes stored in the file before the sound data.

TRAILER BYTES

Trailer Bytes specifies the number of bytes stored in the file after the sound data.

OPEN VIDEO STREAM

This dialog is reached when you select an .AVI file that has multiple audio or video streams. It allows you to select a particular video stream and an audio stream to open in Sound Forge. Most .AVI files have just one video stream and one audio stream, in which case this dialog doesn't appear.

STREAM

The black diamond on the left side of a stream indicates that it has been selected to be open. If you have more than one audio stream or video stream, click on the left of the stream to select it for opening. Additional information about each stream is available by clicking on the + area to expand the list.

CLOSE

Use this command to close the current sound file. If the data has been modified, you will be asked if you wish to save your changes.

WORKSPACE

Workspace files (.sfw) store the names of opened files along with their position, size, and views.

OPEN

Opening a workspace opens all of the files in the workspace and restores their positions on the screen.

LOOK IN

Look in lists available drives and directories.

FILE NAME

Type the name of the file that you want to open or a *. and an extension to see a list of files with the corresponding extension in the current directory.

FILES OF TYPE

Files of Type shows the type of files displayed. Choose a type from the drop-down list.

SAVE AS

Saving a workspace saves the names of all of the opened files along with their position, size, and views.

SAVE IN

Save in lists available drives and directories.

FILE NAME

Type the name of the file that you want to save.

SAVE AS TYPE

Save as type shows the type of files displayed. Choose a type from the drop-down list.

SAVE

Use this command to save the current sound data. If the data is new data which was not retrieved from a previous file or was loaded from a format which is not supported for saving, you will be prompted with the *Save As* dialog. The Macintosh Resource (SND) format is not supported for saving. If you need to export sound files to a Macintosh, use the Macintosh AIFF format.



Shortcut: *Control+S* or *Shift+F12* or *Alt+Shift+F2*

SAVE AS

Use this command to save the current sound data in a different file and/or file format. The Macintosh Resource (SND) format is not supported for saving. If you need to export sound files to a Macintosh, use the Macintosh AIFF format.



SAVE IN

Save in lists available drives and directories.

FILE NAME

Type the file name you want to save.

SAVE AS TYPE

Save as type shows the type of files displayed. Choose a type from the drop-down list.

MRU FOLDERS

Use *MRU folders* to speed the opening of regularly used files.

FORMAT

Format shows the sample format in which the data will be saved. Sound files are most often saved in the PCM format. However, other formats are often used for compression purposes. If you have the *Microsoft Audio Compression Manager (ACM)*

installed, you can save files in a variety of compressed formats including the Microsoft ADPCM format.

When saving as a Raw file, other options are available after selecting OK.

ATTRIBUTES

Attributes controls the sample rate, sample size, and whether the file is stereo/mono. Here, you can change the sample size (8 or 16-bit) that the file will be saved as. You can also convert a file between mono and stereo formats. When you convert from mono to stereo, the data will be stored in both channels. When converting from stereo to mono, the data from both channels will be mixed to a single channel.

REGIONS/PLAYLIST

Regions/Playlist lists options for saving *Regions/Playlist* information. Options include saving in the .WAV file, saving to a Sound Forge external playlist file, saving to a Session 8-compatible regions file or not saving at all.

SAVE SUMMARY INFORMATION IN .AVI/.WAV FILES

Check *Save summary information in .AVI/.WAV files* if you want summary information (such as author and copyright) stored in the .AVI/.WAV file.

SAVE SAMPLER INFORMATION IN .WAV FILES

Check *Save sampler information in .WAV files* if you want sampler information (such as note assignment and loops) saved in the .WAV file.

SAVE AS RAW DATA FILE

The *Save Raw File Type* dialog is reached when you select *Raw* from the *Format* list in the *Save* dialog. Use this to save an opened file into a file format that is not one of the standard sound file formats supported by Sound Forge.

NAME

This list of file formats contains preset versions of raw file storage formats commonly used by other applications. You can create a new file format by using the *Save As* button to save the current storage parameters.

SAMPLE TYPE

The *Sample type* controls the format used to store each sample.

FORMAT

Format specifies the binary format of each sample value.

CHANNELS

Use *Channels* to select the number of tracks stored in the file.

BYTE ORDER

Byte order specifies the order in which the high and low bytes of a 16-bit sample are stored. Little Endian (low, high) is used by Intel microprocessors, while Big Endian (high, low) is used by Motorola microprocessors.

VIDEO SAVE OPTIONS

This dialog is reached when you select Video for Windows (.avi) from the *Save as type* list in the *Save As* dialog. In this dialog, you can select which video streams you wish to save and edit the title, language, and priority of the saved streams.

STREAMS

Check the streams that you want to save. The black diamond next to each check box indicates currently opened streams.

Additional information is available by clicking on the + area to expand the list. This information can be edited by right-clicking on a field and selecting *Edit*.

SAVE AUDIO TO NEW STREAM

If *Save audio to new stream* is checked, the audio track currently open in Sound Forge will be saved to a newly created stream.

COMPRESSION OPTIONS

This dialog is reached when you select Microsoft .AVI from the *Format* list in the *Save As* dialog and then press OK in the *Video Save Options* dialog.

CHOOSE A STREAM

Check the audio or video stream for which you want to set the compression formats.

INTERLEAVE FRAMES

When *Interleave frames* is checked, the audio and video streams will be interleaved. Interleaving the audio and video streams improves playback performance when using storage devices that are more efficient at providing data sequentially (such as CD-ROM disks). The *Frames* value indicates the interval between video and audio segments. The default is 1 for CD-ROMs, but can be increased for faster media.

OPTIONS

Press the *Options* button to view and/or edit the audio and video compression options for the selected stream.

VIDEO COMPRESSION

This dialog is reached when you select Microsoft .AVI from the *Format* list in the *Save As* dialog, press OK in the *Video Save Options* dialog, and after selecting a stream pressing the *Options* button in the *Video Compression* dialog.

COMPRESSOR

Select a compression format from any of the installed .AVI compression schemes. Compression formats are available from different vendors (not from Sonic Foundry), and are also included with Windows.

If the file is already compressed, you should select *No Recompression*. This will be faster and will also avoid the loss of quality inherent in most recompression.

COMPRESSION QUALITY

Compression quality determines the final output quality. Usually, higher quality (close to 100) means less compression and fewer visual artifacts.

KEY FRAMES EVERY

Some compression formats are based on key frames (sometimes called temporal compression). A key frame is a frame that is usually less compressed. Compressed frames that follow a key frame are smaller but slower to draw. This option is only relevant to video streams.

A high number of key frames increases quality, but decreases compression. Also, if the number of key frames is very low, it will take a long time to perform decompression when drawing the video strip or moving randomly through the file.

DATA RATE

Data rate determines the data rate which will be required to play the compressed stream in real time. Lower data rates are less demanding on computer systems than higher rates. If the data rate is too high for a certain computer (usually because of a bottleneck such as a slow CD-ROM, network, or hard drive) the quality begins to suffer and glitching can occur.

CONFIGURE

Press this button to set more detailed configuration parameters specific to the compression scheme being used. Not all compressors provide this option.

PREVIEW

Press this button to view a preview of the final output.

SAVE .RA FILE

This dialog is reached when you select RealMedia (*.rm;*.ra) from the Save as type list in the Save as dialog. This feature is only available for Windows 95 and NT running on x86 processors.

ENABLE SELECTIVE RECORD CAPABILITIES

When checked, RealMedia-compatible players can allow users to save (record) RealMedia streams to their hard drives. Use this option only if it is acceptable for users to create copies of the RealMedia file. The RealPlayer Plus from RealNetworks supports Selective Record capabilities.

ENABLE PERFECTPLAY CAPABILITIES

When this option is enabled, users connected with standard modems can select to download a higher quality version or the audio and/or video that is normally only available over ISDN or LAN connections. The higher bandwidth RealMedia stream is then played back from the user's hard disk after the download process is complete. The RealPlayer Plus from RealNetworks supports PerfectPlay capabilities.

CREATE EVENT FILE (RAE) FROM MARKER AND REGION LABELS

When this option is enabled, events can be embedded from the Regions List (View menu) into the RealMedia stream.

To embed an event, create a marker or region with a label starting with one of the below prefixes followed by a colon.

RM URL: Specifies a URL (web address) to open.

RM Title: Sets the Title field in the RealPlayer.

RM Author: Sets the Author field in the RealPlayer

RM Copyright: Sets the Copyright field in the RealPlayer

For example, if you create a marker with the label RM URL:<http://www.sonicfoundry.com>, the RealPlayer would open the Sonic Foundry home page at the marker's location in the file during playback. You can just as easily open specific web pages. For example, RM URL:<http://www.sonicfoundry.com/realaudio.html> would open the [realaudio.html](http://www.sonicfoundry.com/realaudio.html) page on Sonic Foundry's web site.

RealNetworks publishes extensive information on RealMedia events on their web site (www.real.com). However, Sound Forge is capable of generating all event files automatically rather than requiring several tools as described in RealNetworks' documentation.

SAVE AUDIO STREAM

When checked, the file's audio stream will be converted and saved into the RealMedia file.

AUDIO COMPRESSION

This option selects the RealAudio compressor that is used when saving the file. If you need to save RealAudio files with Bandwidth Negotiation capabilities, you can automate the process using Sonic Foundry's Batch Converter Plug-In (available for Sound Forge 4.0 only).

When a compression type is selected, a brief description of it is displayed below the list box. The frequency response parameter of the description indicates how wide, in Hertz, the bandwidth of the audio in the RealMedia file will be. The bandwidth of a normal human's hearing is about 20,000 Hz or 20 kHz.

STEREO TO MONO CONVERSION

If the source audio file is stereo and you choose to convert to a mono RealAudio format, you will need to decide what channel(s) to use for the mono signal. Selecting *Both channels* will evenly mix the signal together. If you are unsure of this option, always choose *Both Channels*.

SAVE VIDEO STREAM

When checked, the file's video stream will be converted and saved into the RealMedia file.

VIDEO COMPRESSION

Currently there are two types of RealVideo compression: RealVideo (Standard) and RealVideo (Fractal). Each type is optimized for different types of video and transfer rates. Detailed information on the two types is available on the RealMedia home page (www.real.com)

VIDEO COMPRESSION QUALITY

This option allows you to compromise between quality of compression and lowest possible data transfer rate. Setting this level at 100 provides the best quality video with the greatest required bandwidth to stream. A setting of 1 gives you the lowest quality video but requires the least bandwidth to stream effectively.

OPTIMIZE FRAME RATE

Enabling this option allows the RealMedia compressor to determine an optimized frame rate for the required data rate. It is highly suggested that this feature be enabled, since it allows the compressor to vary the frame rate over time depending on the content itself.

FRAMES PER SECOND

This setting allows you to change the frame rate of the video that is being compressed. A lower frame rate, such as three or five frames per second, will greatly lower the bandwidth requirements for streaming the video. Just type the desired frames per second into the edit box, or choose one of the presets from the drop-down list.

RE-SIZE VIDEO FRAMES TO

Use this option to change the frame dimensions of the video. Several common sizes can be selected from the drop-down menu. If you'd like to specify dimensions not on the list, select (Custom) and type in the required proportions.

DESIRED MAXIMUM DATA RATE

This setting is available only if you are saving a video stream. The data rate specifies the desired total kilobits per second (Kbps) for the RealMedia file (including audio and video). To increase the bandwidth available for streaming the video, you can either lower the data rate for the audio, or increase the desired maximum data rate. Note that you must allow some room for inconsistent transfer rates. For example, a 28.8 modem cannot sustain 28.8 Kbps. Instead, you should target your maximum data rate for between 20.0 and 22.0 Kbps. The allowed data rates are displayed next to the Audio/Video

compression options. After the file has been saved, statistics about the actual data rate will be displayed. Use this information to adjust parameters as necessary.

SAVE .ASF FILE

This dialog is reached when you select *Active Streaming Format (asf)* from the *Save as type* list in the *Save* dialog.

TEMPLATE

The *Template* section displays preset saving options as well as any custom templates that have been created. Templates contain the settings for encoding ASF files. Choose a template and then press the *Create* button to encode a file to ASF.

DESCRIPTION

The *Description* section displays a short summary of what type of encoding the selected template is good for.

DETAILS

The *Details* section displays the technical details of the selected template. Here you will find the name of the audio and/or video codec being used and details of the audio and/or video format.

PROPERTIES

The *Properties* section allows you to enter specific information to embed in the encoded ASF file. Information such as title and copyright will be displayed in the player of the ASF file as the file is being played.

GENERATE REDIRECTOR FILES

When this option is checked, ASX redirector files will be generated. Press the *ASX Settings* button for more options.

ASF SERVER PATH

This option specifies a path for the server that will be accessed to play the NetShow file.

USE ALTERNATE DESTINATION FOLDER FOR ASX FILES

When this option is checked, the ASX files will be created in the specified directory rather than the destination directory for the encoded ASF files.

PLAY START POSITION

This option allows you to specify an alternate play start position for the encoded file.

PLAY LENGTH

This option allows you to specify an alternate play length for the encoded file.

MARKERS AND SCRIPT COMMANDS

The *Markers and Script Command* options allow you to use or not use any markers that are defined in the source file. If the checkboxes are unchecked, the encoder will ignore the markers.

CUSTOM

The *Custom* button launches the Custom Template wizard. Using this wizard you can create your own custom templates.

TEMPLATE NAME

A unique template name must be provided.

BANDWIDTH**Internet 14.4(10 kbps actual).**

This bandwidth selection supports streaming of both audio and video content to clients using a network connection that operates at a minimum of 14.4 Kbps.

Internet 28.8(22 kbps actual).

This bandwidth selection supports streaming of both audio and video content to clients using a network connection that operates at a minimum of 28.8 Kbps.

Internet 56 Modem (37 kbps actual).

This bandwidth selection support streaming of both audio and video content to clients using a modem network connection that operates at 56 Kbps.

Single Channel ISDN (50 kbps actual).

This bandwidth selection supports network connections at Integrated Services Digital Network (ISDN) rates.

Intranet (100 kbps actual).

This bandwidth selection supports clients on an intranet or local area network (LAN).

Custom (bps).

This bandwidth selection is for users whose system requires a customized bandwidth setting, such as those using a dedicated network that does not use one of the preceding bandwidths. When specifying the bit rate, consider other traffic that must pass over the network. Unlike the preconfigured bandwidth settings, Microsoft NetShow Encoder uses the entire bandwidth you specify for streaming ASF information. Note that this value is entered in bits per second, rather than kilobits per second.

CODEC OPTIONS**Audio Panel**

Because some of the audio codecs have multiple formats, the Audio panel has options for both the codec and the format. Use the Codec list to select the audio codec. If the codec has multiple formats, use the Format list to select the amount or type of compression. The format of a codec is the sampling rate (in kilohertz [KHz] or hertz [Hz]) and sample depth (in bits). The larger the sampling rate and sample depth, the more information is recorded. When more information is recorded, the audio content usually sounds better; however, the audio portion of the total bandwidth allocated also is larger. Remember that the audio file must share the bandwidth with any video or script commands.

Advanced

Click *Advanced* to select the audio loss concealment setting.

Audio Loss Concealment

Audio Loss Concealment blends the audio so that if information is lost during transmission, there is no audible gap. There are two options that provide different types of concealment. If the audio track has a low dynamic range, such as when someone is speaking, select *Voice*. If the audio track has a large dynamic range, such as when a band is playing, select *Music*.

Video Panel

Choose a video codec from the list next to Codec. Be sure to select a codec that can share the bandwidth with the audio codec you have chosen.

Advanced.

To adjust some advanced video settings, such as the size of the window for Microsoft Media Player, the number of frames per second, and the number of I-frames, click *Advanced*.

Image Size Panel

To determine the height and width of the Advanced Streaming Format (ASF) content when it is viewed use the *Image Size* panel. Quarter Common Interface format (QCIF) and Common Interface format (CIF) are predefined image standards: notice that CIF height and width are exactly double those of QCIF. These standards are provided because their image size (both height and width) is a standard that is supported by many codecs and video editing tools. The 160x130 and 320x240 options provide video windows that are slightly smaller than QCIF and CIF, respectively. The slight decrease in image size allows for more bandwidth to be used for video frames, and can improve content quality in some scenarios. Use these options if video sharpness is more important than video size for your content.

Use *Custom* to define your own image size. When resizing an ASF stream with the *Custom* setting, be aware that many codecs require image size settings in multiples of 16.

Frames / Second

To set the number of frames per second for the ASF stream, use the *Frames/Second* dialog box. The *Frames/Second* value can be set from 1 through 15. How much you adjust this value depends on the speed of the computer's coprocessor, and whether the computer is encoding high-movement or low-movement video. If more frames per second are used, there are smaller changes between frames, and the video looks smoother. If fewer frames per second are used, the video has a slower motion transition rate, and thus video with a lot of motion can appear disjointed. However, the more frames per second your computer is required to encode, the harder the coprocessor must work. It is possible to overwork your coprocessor by requesting it to encode more action than its capability allows. If this happens, NetShow Encoder stops encoding, and displays a message.

Second / IFrame

The *Second/ I-frames* dialog box displays the number of seconds that separate I-frames. All frames in between I-frames are called delta frames. Delta frames contain only the information used to display change, and are considerably smaller than I-frames; however, the visual quality of delta frames is poor. As I-frames get farther apart, the size of the video (in bytes) gets smaller. As I-frames get closer together, the size of the video (in bytes) gets larger, and the computer's coprocessor must work harder.

Pixel Format

The *Pixel Format* displays the color depth standard that has been selected. The video card that the computer is using to receive the video source has certain color standards that it can display best, for example, RGBH = 64k thousand colors and RGBT = 16 million colors. For information about which pixel format your card works best with, see the video card's documentation. The pixel format is set automatically for the preferred video cards when you select one of the template stream formats.

Delay Buffer

The *Delay Buffer* is the amount of video coming in from the video card (measured in seconds) that NetShow Encoder stores before it starts transmitting ASF information. Five seconds is the default; it allows enough information to build up so that the client doesn't have to wait, (that is, buffer information) and the NetShow Encoder computer's memory isn't overloaded with delayed ASF information. Increasing the Delay Buffer can overload the NetShow Encoder computer's memory, and affect the computer's ability to encode information. Decreasing the Delay Buffer, in some circumstances, can cause clients to wait on NetShow Encoder.

Scalable Video

Scalable Video creates multiple content streams within one encoded .asf file. The streams are encoded at different target bandwidth to allow a player to render content when network bandwidth is reduced. Using scalable video require an encoding platform that has at least a Pentium II processor running at 266 megahertz(MHz). Scalable video content can only be accurately decoded by NetShow Player 3.0 and NetShow Server 3.0.

Image Quality

The *Image Quality* slider lets you fine-tune either the image quality. You can emphasize either video smoothness or image crispness. Video content is smooth when movement shown on the screen has easy transition from one position to another and the edges of objects are not jagged. Video content is crisp when images and motion are well defined and clearly delineated. As you increase the image crispness, the video becomes less smooth. In some cases decreasing the image crispness makes the video look better.

EDIT

The *Edit* button launches the *Custom Template* wizard so that you can edit a custom template. Factory defined templates can not be edited.

REMOVE

The *Remove* button clears the selected custom template from the template list. You can not remove factory defined templates

SAVE

The *Save* button will encode the source file using the currently selected template.

CANCEL

The *Cancel* button will cancel the *Save As* command.

SAVE ALL

Use this command to save all the currently opened files. You will be prompted for each unsaved file. If you hold the *Shift* key down when selecting this command, no prompts will be displayed.

PROPERTIES

The *Properties* folders list attributes and information that is embedded in the currently active sound file.

GENERAL

This *General* folder contains technical information regarding the current sound file, including size, location, file type and attributes of the file.

Microsoft .WAV files are composed of sections called RIFF chunks. These chunks contain the sound data plus other embedded information. If there are any non-standard chunks found in the current file, the junk, padding, and additional chunks fields will be checked. These chunks might be used by other sound utilities, or might just be garbage which can cause problems like glitching in some less intelligent programs.

SUMMARY

This *Summary* folder contains extra information that can be embedded in the sound file, about the sound file. It is only available in Microsoft .WAV files.

TITLE, SUBJECT, ENGINEER, COPYRIGHT, COMMENTS

Commonly used fields can be quickly edited. To view and edit all other fields, press the *Extended* button.

PICTURE

Picture allows you to select an icon, bitmap or cursor to be attached to the current file.

EXTENDED SUMMARY

Extended summary allows you to view and edit all available summary fields.

Fields

To edit a field, you must first select it from the *Fields* list. The check appearing to the left of the type tells whether this field is currently enabled for saving to the .WAV file. If the check is not shown, the field will not be written to saved files. To change the state of the check, click with the left mouse button in the area immediately to the left of the type or double-click on the entry.

If the field for a type currently has no text associated with it, the word (*Empty*) will appear to the right of the type in the drop-down list. Empty fields are not saved.

IARL Archival Location

IARL Archival Location indicates where the subject of the file is archived.

IART Artist

IART Artist lists the artist of the original subject of the file. For example, Bach.

ICMS Commissioned

ICMS Commissioned lists the name of the person or organization that commissioned the subject of the file. For example, XZZZY Records.

ICMT Comments

ICMT Comments provides general comments about the file or the subject of the file. If the comment is several sentences long, end each sentence with a period. Do not include new-line characters.

ICOP Copyright

ICOP Copyright records the copyright information for the file. For example, "© Copyright Sonic Foundry, Inc. 1999." If there are multiple copyrights, separate them by a semicolon followed by a space.

ICRD Creation Date

ICRD Creation Date specifies the date the subject of the file was created. List dates in year-month-day format, padding one-digit months and days with a zero on the left. For example, "1964-03-02" for March 2, 1964.

ICRP Cropped

ICRP Cropped describes whether an image or sound has been cropped and, if so, how it was cropped. For example, "third movement, first through fourth bars."

IDIM Dimensions

IDIM Dimensions specifies the size of the original subject of the file. For example, "8.5 in h, 11 in w."

IDPI Dots Per Inch

IDPI Dots Per Inch stores dots per inch setting of the digitizer used to produce the file, such as "300."

IENG Engineer

IENG Engineer stores the name of the engineer who worked on the file. If there are multiple engineers, separate the names by a semicolon and a blank. For example, "Sayre, Jack; Feith, John."

IGNR Genre

IGNR *Genre* describes the original work, such as, “landscape,” “portrait,” “still life,” etc.

IKEY Keywords

IKEY *Keywords* provides a list of keywords that refer to the file or subject of the file. Separate multiple keywords with a semicolon and a blank. For example, “Madison; aerial view; scenery.”

ILGT Lightness

ILGT *Lightness* describes the changes in lightness settings on the digitizer required to produce the file. Note that the format of this information depends on hardware used.

IMED Medium

IMED *Medium* describes the original subject of the file, such as, “computer image,” “drawing,” “lithograph,” and so forth.

INAM Name/Title

INAM *Name/Title* stores the title of the subject of the file, such as, “Madison From Above.”

IPLT Palette Setting

IPLT *Palette Setting* specifies the number of colors requested when digitizing an image, such as “256.”

IPRD Product

IPRD *Product* specifies the name of the title the file was originally intended for, such as “Encyclopedia of Midwest Geography.”

ISBJ Subject

ISBJ *Subject* describes the contents of the file, such as “Aerial view of Madison.”

ISFT Software

ISFT *Software* identifies the name of the software package used to create the file, such as “Sound Forge.”

ISHP Sharpness

ISHP *Sharpness* identifies the changes in sharpness for the digitizer required to produce the file (the format depends on the hardware used).

ISRC Source

ISRC Source identifies the name of the person or organization who supplied the original subject of the file. For example, “Sonic Foundry, Inc.”

ISRF Source Form

ISRF Source Form identifies the original form of the material that was digitized, such as “slide,” “paper,” “map,” and so forth. This is not necessarily the same as IMED.

ITCH Technician

ITCH Technician identifies the technician who digitized the file. For example, “Jones, Henry.”

DISP Sound Scheme Title

DISP Sound Scheme Title sets the title that is displayed for MS Sound Systems.

Contents

Contents shows the text associated with the selected field. The text can be changed by typing in this box.

Default

Default sets the default fields to the current file’s field information. Use this to save information which will be used in a number of different files, like copyright and engineer information at your recording site.

Load

Load sets the current file’s summary information fields to the default settings.

LOAD

Load sets all of the fields to the default values. The default values can be set in the *Extended Summary* dialog.

FORMAT

Use the *Format* folder to change the data format of the current data window. This includes *Sample rate*, *Sample size (bits per sample)*, and *Stereo/Mono*.

SAMPLE RATE

The playback rate can be set from 2,000 Hz to 96,000 Hz by entering the rate you wish to use. You may also select the most common sample rates by using the drop-down list. Note that this will not resample the sound file. If the

playback rate is different from the originally recorded rate, the pitch will vary unless resampling is done.

SAMPLE SIZE

Select 8-or 16-bit for the sample size.

CHANNELS

Select *Mono* or *Stereo* for the number of channels.

CURSOR POSITION

Cursor position displays the cursor position in the sound file using the current units.

SAMPLE VALUE

Sample value displays the sample value at the current cursor position in the sound file. For 16-bit audio, this value ranges from -32768 to 32767. For 8-bit audio, it ranges from -128 to 127.

PEAK DATA RATIO

Peak data ratio displays the ratio at which the peak data is stored. This means that when zoomed out beyond 1:512, the peak file, instead of the entire file is scanned when drawing the waveform.

SOUND DATA SIZE

Sound data size displays the amount of storage space being used by the sound data.

VIDEO SOURCE FORMAT

Video source format displays the frame size (height, width, bits per pixel) and compression algorithm of the original .AVI file.

VIDEO DECOMPRESSION FORMAT

Video decompression format displays the intermediate frame size (height, width, bits per pixel) and compression algorithm used by Sound Forge before displaying the video frames. The video decompression mode can be changed from the *Performance* page in the *Preferences* folder.

UNDO BUFFERS

Undo buffers displays the number of undo buffers currently existing for the sound file and the amount of hard-drive space they consume.

REDO BUFFERS

Redo buffers displays the number of redo buffers currently existing for the sound file and the amount of hard-drive space they consume.

Shortcut: Alt+Enter

VIDEO

The *Video* folder contains information specific to Microsoft .AVI files. If your sound file does not currently contain any video, you can attach a video from an .AVI file by pressing the *Attach* button.

An .AVI file can contain multiple video and/or audio streams (or tracks). For example, you can include multiple audio tracks in different languages to go along with a video. When playing back, Microsoft's Media player detects what language the operating system is set to and plays the corresponding video stream.

STREAM

Stream displays information about the video and/or audio streams present in the .AVI file, such as name, compression format, frame rate, and language.

The black diamond to the left of the streams indicate the active streams. You can view and edit additional information (stream name, language, and priority of the active audio stream) by clicking on the + sign to expand the list. Right-clicking on a field allows you to edit it.

Streams that are checked will be preserved when the file is next saved. Streams that are not checked will be deleted when the file is next saved.

DISPLAY

The *Display* folder allows you to change how the data window of the sound file appears on the Sound Forge workspace. Checking a box will display that item in the current sound file's data window. An unchecked item will not be displayed.

To set the current settings as the default display for all sound files, simply check the *Save as the default for all new windows* check box.

SEND

Use this *Send* command to send the current sound file via electronic mail.

Sends a mail message with an attached sound file. Note that you must have electronic mail (email) capabilities to use this feature. When you select the **Send** command, your email program opens with a new message and embeds the active sound file from Sound Forge in the message. You then have the option of sending the new mail message to the desired recipient.

When a sound file is embedded into a mail message, the entire contents of the sound file are inserted into the message. To do this requires that any modifications to a sound file be saved before the message can be sent. Sound Forge will prompt you to save the file if this is necessary.

NOTE

Please note that sound files are generally very large. An email message containing an embedded sound file can take a considerable amount of time to send and receive, especially if you are using a modem and not a high-speed network.

EXIT

Use the **Exit** command to exit Sound Forge. You will be asked whether you would like to save any sound files you have modified during the editing session.

THE EDIT MENU

UNDO

Use this command to undo the last edit operation. For instance, if you deleted a section of data by accident, simply choose the **Undo** option from the **Edit** menu to put the data back.



You can also undo operations from the **Undo/Redo History Window**.

NOTE

When working on large selections, it takes Sound Forge a longer period of time to create an undo buffer. To save time and file space, you can disable the undo. This is done by selecting the *Disable Undo/Redo* item under the *Edit* menu or from within most processing dialogs.

Shortcut: Control+Z or Alt+Backspace\

REDO



Use this command to re-perform an undone event. For instance, if you decide that you really did want to delete that selection of data, select *Redo* and the undo will be undone.

You can also redo operations from the *Undo/Redo History Window*.

Shortcut: Control+Shift+Z

REPEAT



Use this command to repeat the last performed operation. This can be used with most processing functions. The last dialog settings will be used unless you hold down the *Shift* key, which allows you to change the operation's dialog parameters.

Shortcut: Control+Y to repeat last operation with the last used dialog options or Control+Shift+Y allows you to change the dialog settings before repeating

CUT



Use this command to remove selected sound data and put it onto the clipboard. This command has no effect if there is no selected data. Cutting sound data replaces the previous contents of the clipboard.

Cutting data from the beginning of a very large file will take a long time to process. This is because large amounts of data need to be moved. You can preview a cut by using the *Preview Cut/Cursor* function (Edit menu). Also, the *Playlist* can be used to make instantaneous non-destructive cuts.

*You cannot cut from a single channel in a stereo file since the two channels in a stereo file must always be equal in length. To shift a single channel in time, use the *Delay/Echo | Simple* function in the *Effects* menu.*

NOTE

Shortcut: Control+X or Control+Del

COPY

Use this command to copy selected sound data onto the clipboard. This command has no effect if there is no selected data.



Shortcut: Control+C or Control+Insert

COPY OBJECT LINK

Use this command to copy an object link to the sound file onto the clipboard.

Use this command to copy an object link for the active sound file onto the clipboard. You can then *Paste* (or *Paste Special*) the *Link* into another application (such as Microsoft Word or Write) that supports *Object Linking and Embedding (OLE)*. This feature is useful for embedding sound files in documents.

Object links are simply references to existing sound files on your hard disk or network server. If you embed an object link to a sound file in a document, the sound file is not actually inserted into the document, just a reference. Therefore, if you copy the document to another computer, any sound file(s) that are embedded will not be available on the other computer unless the sound file is stored on a common network server.

If you intend to share documents between computers that contain embedded sound files, it is suggested that you store all sound files on a network server and embed them using Universal Naming Convention (UNC) path names. For example, if your server and share name for sound files is \\myserver\sounds, then open and embed all sound files using the UNC path (i.e. \\myserver\sounds\

If you are running Windows 95, UNC path names are extremely simple to use. Instead of using a drive letter when browsing in the *File Open* dialog, select *Network Neighborhood* in the *Look In* field and browse to the appropriate network server and share containing your sound files. This method of browsing uses UNC names exclusively.

For more information on *Copy Object Link*, and OLE support in general, refer to the *Object Linking and Embedding (OLE)* chapter in the Appendix.

PASTE



Use this command to insert a copy of the clipboard contents at the current insertion point. If there is a selection made, the *Paste* command deletes the selected data before inserting.

This command has no effect if the clipboard is empty.

Pasting into a stereo file will insert data to both channels since the channels in a stereo file must always be equal in length.

Pasting data of different sample rates will cause the data in the clipboard to play at the same rate as the rate of the window in which the data is pasted.

If any regions, markers, or loops are present along with the original sound data, they will also be pasted into the destination sound file. To disable this feature, uncheck *Paste Events* in the *Options* menu.

Shortcut: Control+V or Shift+Insert or Drag and drop a selection to another data window + Alt

Paste

Specifies that dropping the data here will perform a paste operation. You can paste regions by dragging any region onto an open data window. The paste will occur at the current insertion point within the data window on which the data is dropped. The paste operation is used rather than mix when the *Alt* key is held down while dragging. A small “L” or “R” will appear next to the cursor when pasting only to a left or right channel respectively.



Instead of using the *Alt* key while dropping, you can also hit the right mouse button while dragging. This toggles between *Mix*, *Crossfade*, and *Paste* drag and drop.

PASTE SPECIAL

CROSSFADE

Use this command to crossfade a copy of the clipboard contents with the sound data. The crossfade start point is either the cursor point or the start or end of the selection in the destination file.

The *Crossfade* operation is similar to a mix, except that you can apply a linear fade to the source and destination data. It is commonly used as a transition between one sound segment to another.

When applying a crossfade, you often want to fade out (100 to 0%) the destination data and fade in the source (0 to 100%) data starting near the end of the destination file. Linear fades are applied only to the area where the destination and source files overlap following the crossfade start point. Any extra source data is appended at the end of the destination file.

SOURCE START/END LEVEL

Source Start/End Level determines the linear fade applied to the source data before mixing with the destination. A gain equal to the *Source End* level is applied to the source data when appending to the end the source file.

DESTINATION START/END LEVEL

Destination Start/End Level determines the linear fade applied to the destination data. Any destination data following the crossfade end point is not affected.

START CROSSFADE AT

If there is no selection in the destination data window, the linear fade starts at the insertion point. When there is a selection in the destination data window, the linear fade can begin at the start or end of the destination selection.

TUNE

This function adjusts the crossfade start in the *Destination* file. Use to do fine tweaking of the crossfade without having to leave the *Crossfade* dialog.

Crossfading data of different sample rates will cause the data on the clipboard to play at the same rate as that in the current sound data with which the data is crossfaded.

This command has no effect if the clipboard is empty.

If any regions, markers, or loops are present along with the original sound data, they will also be pasted into the destination sound file. To disable this feature, uncheck *Paste Events* in the *Options* menu.

When using drag and drop crossfading, make sure that the *Control* key is down when releasing the mouse button. When dealing with mono and stereo data, crossfading will occur on the channel which the selection is dropped to (*Left*, *Right*, or *Both*).

Shortcut: *Control+F*

Drag and drop a selection from the source to the destination data window while holding the *Control* key. Make sure that you don't start the drag with the *Control* key down, otherwise you will switch to *Magnify* mode. The *Control* key must be down when you release (drop) the selection.

Crossfade



Specifies that dropping the data here will perform a *Crossfade* operation. You can crossfade regions by dragging any region onto an open data window. The crossfade will occur at the current insertion point within the data window on which the data is dropped. The *Crossfade* operation is used rather than *Mix* when the *Control* key is held down while dragging. When dragging a single channel from a stereo file or a mono file onto a stereo file you can select whether the data is placed in the left or right channel by the vertical position within the data window. A small "L" or "R" will appear next to the cursor when crossfading only to a left or right channel respectively.

MIX

Use this command to mix a copy of the clipboard contents with the sound data. The mix start point is either the cursor point or the start or end of the selection in the destination data window.



Mix

Specifies that dropping the data here will perform a Mix operation. You can mix regions by dragging any region onto an open data window. The mix will occur at the current insertion point within the data window on which the data is dropped. When dragging a single channel from a stereo file, or a mono file onto a stereo file, you can select whether the data is placed in the left or right channel by the vertical position within the data window. A small “L” or “R” will appear next to the cursor when mixing only to a left or right channel respectively.



SOURCE VOLUME

Source Volume controls the gain applied to the source data before mixing with destination data.

DESTINATION VOLUME

Destination Volume controls the gain applied to the destination data before mixing with source data.

INVERT DATA

Use *Invert data* to reverse the polarity of the source or destination data. This is useful to A/B the phase relationship of the two sound files being mixed.

START MIX AT

When there is a selection in the destination data window, you have the option of performing the mix starting at the beginning or end of the selection. When no selection is present, the data is always mixed starting at the cursor position.

APPLY DESTINATION VOLUME TO OVERLAPPING AREA ONLY

When this is off, the *Destination Volume* gain is applied to the entire destination sound file. Otherwise, the *Destination Volume* gain is applied to only where the destination and source data are mixed.

PRE-POST-FADE DESTINATION EDGES

Use this to apply a fade to the destination data before and after the mixing region. The *Pre-fade* time determines the fade time before the mix start, while the *Post-fade* time determines the fade-time after the mix end.

For example, when mixing vocals over a music track, you can have the music track fade before and after the region where vocals are mixed, so that the vocals are more pronounced. This technique is sometimes referred to as ducking.

TUNE

This function adjusts the mix start position in the *Destination* file. Use this to do fine tweaking of the mix without having to leave the *Mix* dialog.

When mixing two sound files together, phase cancellation of some frequencies often occurs. Very small delays can cause very audible differences in tone quality. The *Tune* control allows you to minimize (or maximize, if you so desire) this effect.

This command has no effect if the clipboard is empty.

Mixing data of different sample rates will cause the data on the clipboard to play at the same rate as that in the current sound file with which the data is mixed.

Shortcut: Control+M or Drag and drop a selection from the source to the destination data window.

OVERWRITE

Use this command to replace an area of sound data with the contents of the clipboard.

The *Overwrite* command replaces data using the following rules:

If an area is selected which is greater than the length of the clipboard, the data from the beginning of the selection is overwritten for the length of the clipboard.

If an area is selected which is less than or equal to the length of the clipboard, the data within the selection is overwritten with the clipboard for the length of the selection.

The *overwrite* command is useful when you are trying to replace silent sections of audio with background noise.

Shortcut: Control+H

REPLICATE

Use this command to copy multiple copies of the clipboard to the current data window.

The *Replicate* command will copy as many copies of the clipboard as will fit to a selected area. If no selection is currently made in the data window, the function will have no effect.

COPY PARTIALS

Selecting the *Copy partials* button will fill the selected region completely, using a partial copy of the clipboard if needed.

WHOLE COPIES

Selecting this option will use only whole copies of the clipboard, and will not fill the selected region completely if its length is not an exact multiple of the clipboard length.

Example:

Copy a 1 second section of sound data onto the clipboard.

Insert 5.5 seconds of silence in the data waveform.

Select the 5.5 second area and select *Replicate*.

If you choose *Whole copies*, 5 copies of the clipboard will be placed in the area the silence previously occupied and 0.5 seconds of silence will remain at the end.

If you choose *Copy partials*, 5.5 copies of the clipboard will be placed in the data rather than just 5. No silence will remain.

The replicate command is useful when you are trying to replace silent sections of audio with background noise. It can also be used to introduce stuttering effects or echoes.

PASTE TO NEW

Use this command to create a new data window which contains the contents of the clipboard.

Shortcut: *Control+E* or Drag and drop a selection to the *Sound Forge* workspace to create a new file from the selection

TRIM/CROP



Use this command to remove all data from the sound file except the selected data. This command has no effect if there is no selected data. Trim/Crop sound data does not copy data onto the clipboard.

CLEAR/DELETE

Use this command to remove selected sound data without copying it onto the clipboard.

NOTE

You cannot clear data from a single channel in a stereo file since the two channels in a stereo file must always be equal in length. This command has no effect if there is no selected data.

Shortcut: Delete

SELECT ALL

Use this command to select all data in the current data window.

Shortcut: Control+A or Control+Numpad 5 or Double-clicking the left mouse button in the waveform display will select all data.

PREVIEW CUT/CURSOR

Use this command to playback around a selected region. This lets you preview the result of a Cut or Clear operation without altering the file. Go to the Previews page in the Options folder (File menu) to set the amount of pre- and post-roll around the selection.

If there is no selection, the playback will pre- and post-roll around the cursor position.

Shortcut: Control-click on the Play Normal button or Control+K

PRE-ROLL TO CURSOR

Pre-roll to Cursor will play the data before the cursor position for the specified pre-roll length in the preview *Preferences* folder.

If a selection is made the pre-roll will occur for whichever end of the selection is active. To toggle between ends of the selection use the *Home* and *End* keys. This function is invaluable when tweaking the end of a long selection and not needing to hear the entire thing.

Shortcut: *Control+Shift*-click on the *Play Normal* button or *Control+Shift+K*

TOOL

EDIT

Use this command to use the *Edit Tool* of *Sound Forge*. This is the typical user tool to do all of your selecting, editing, and processing.



MAGNIFY

Use this tool to zoom in to a particular region without losing your selection. While using the *Magnify Tool*, when you make a selection *Sound Forge* will draw a dotted rectangle around an area and will magnify the area when the mouse button is released.



There are three modes of operation for the *Magnify Tool*, and these can be chosen by toggle-clicking (refer to the *Using Sound Forge* section explanation of toggle-clicking) while the *Magnify Tool* is selected. These modes allow for *Zoom Time* only, *Zoom Level* only or a *Time/ Zoom Level*.

To return to the normal editing, choose the *Edit Tool* from the *Edit* menu.

Shortcut: While using any tool, if you hold the *Control* key down before making a selection with the mouse, you will switch to the *Magnify Tool*. This is handy for quickly magnifying a particular section of your sound file.

Data Magnify Mode



When you are in the data window this cursor shows that you can select an area to magnify. Dragging within the data window will draw a dotted rectangle which specifies the region of data to magnify.

PENCIL



Use this tool to edit the waveform by drawing on it. For example, if you have a glitch in the sound data, zoom in to the glitch and smoothly re-draw the waveform using the *Pencil Tool*.

While holding down the left mouse button, you can also use the arrow keys to fine tune the pencil position.

The *Pencil Tool* is only available when operating at magnification levels between 1:1 and 1:16. To return to the normal editing mode, choose the *Edit Tool* from the *Edit* menu.

Data Pencil Mode



When you are in the data window this cursor shows that you can draw new audio data. Dragging within the data window will draw a new line which represents the new audio data. This is an easy way to remove glitches or pops within the audio data.

GO TO



Use *Go To* to set the cursor to a particular location in the sound file.

GO TO

Go To provides preset locations to pick from.

POSITION

Position determines where the cursor will be moved.

INPUT FORMAT

Input Format determines the input format that will be used to enter the position.

If you have a selection, you can switch the cursor between selection points with the ` key (reverse quote, normally above Tab key). You can also switch the cursor between selection points by holding the Control key while selecting the Go to Start or Go to End buttons in the Playbar or Transport toolbar.

NOTE

Shortcut: Control+G or F5 or Waveform Display shortcut menu or Double-click on the selection start field in the playbar

SELECTION

SNAP TO TIME

Snap to Time forces both edges of a selection to a whole time division as designated by the marks on the Time Ruler above the data window.

Shortcut: T

SNAP EDGE TO TIME

Snap Edge to Time forces the active edge of a selection to a whole time division as designated by the marks on the Time Ruler above the data window. The active edge of a selection is defined by the blinking cursor.

Shortcut: Shift+T

SNAP TO ZERO

Snap to Zero command forces both edges of a selection to the next zero-crossing of the waveform. The Editing tab in the Preferences folder allows you to choose whether this is a negative, positive or any zero-crossing.

Shortcut: Z

SNAP EDGE TO ZERO

Snap Edge to Zero forces the active edge of a selection to the next zero-crossing of the waveform. The active edge of a selection is defined by the blinking cursor.

Shortcut: Shift+Z

SET

The *Set* dialog allows you to manually define the selection by start and end points or length.

SELECTION

Selection provides preset regions to pick from.

START

Start specifies the start point of the selection.

END

End specifies the end point of the selection.

LENGTH

Length specifies overall length of the selection.

CHANNEL

If the sound file is stereo, you can specify which channel of the data will be selected.

INPUT FORMAT

Input Format determines the input format that will be used to enter the *Start*, *End*, and *Length*.

SNAP TO ZERO

Pushing the *Snap to Zero* button snaps the current selection to the nearest zero-crossing.

SNAP TO TIME

Pushing the *Snap to Time* button snaps the current selection to the nearest whole time, as defined by the major ticks on the *Time Ruler*.

Shortcut: Control+D or *Waveform Display* shortcut menu or Double-click on the *selection length* field in the *playbar*

UNDO ALL

Use this command to undo all operations performed on a file. Unless *Undo* was disabled at some point, this will return your file to its original state when opened. However, if you disable undo at any time, you might not be able to return to the files original state.

When not working in direct mode, you can simply close the file without saving changes to revert to the original file. However, when working in direct mode if you decide to not save the current changes, Sound Forge does an *Undo All* operation automatically before saving.

DISABLE UNDO/REDO

Use this command to prevent Sound Forge from creating temporary files used for undoing operations. Since creating undo files takes time and disk space, it is sometimes more convenient to have the *Disable Undo* option on, especially when working with large files.

Remember to turn the *Disable Undo* feature off when you wish to be able to undo operations.

NOTE

Since Sound Forge doesn't modify the original file unless the file is opened direct or read-only, you can always recover it by re-opening it. However, saved changes are permanent.

Shortcut: Control+U

THE VIEW MENU

MAXIMIZE WIDTH

MaximizeWidth stretches the active data window to fit between the extents of the Sound Forge window.



FULL SCREEN

Full Screen will super maximize the Sound Forge workspace in your display. To return to the regular window display, choose the *Full Screen* command again.

TOOLBARS

Toolbars will bring up the *Toolbars* tab in the *Preferences* folder and allows you to quickly bring up any toolbars you would like on the Sound Forge workspace.

CLIPBOARD

Clipboard is used to display information on the current contents of the clipboard.

CONTENTS

Contents displays information includes *Format*, *Attributes*, *Length*, *Samples*, and *Total bytes*. If no data is in the clipboard you will be notified with a message box.

PLAY

Play will play the audio sitting on the clipboard.

ZOOM LEVEL

OUT FULL

Use this command to zoom out vertically to allow viewing of the entire amplitude range. This command also centers the centerline.

WINDOW

Use this command to zoom in or out vertically to the maximum zoom ratio that allows you to view the entire waveform on the waveform display. This command also moves the centerline if the maximum positive and negative peaks are different.

SELECTION

Use this command to maximize a selected area both vertically and horizontally.

To zoom in and out vertically by small increments, you can use Shift+Up or Shift+Down arrow keys or the Zoom In/Out buttons on the lower left-hand side of the data window.

NOTE

Shortcut: Control+Up arrow (when a selection is active) or Level Ruler shortcut menu or Double-click on the Level Ruler

ZOOM TIME

IN FULL



Use this command to maximize the horizontal magnification to one sample per pixel on the screen. When you zoom in full, the entire waveform display length will span a very small amount of time, usually around 20 milliseconds (depending on the sample rate and window size).

To zoom in and out by small increments, you can use the up and down arrow keys.

Shortcut: Control+Up arrow (when no selection is active) or Waveform Display shortcut menu

NORMAL

Use this command to zoom out to the default zoom ratio preference setting. If the entire sound file can fit in the window at a closer magnification, zoom normal will only zoom out to that ratio.



To change the default for the Normal zoom ratio, go to the Display page in the Preferences folder.

OUT FULL

Use this command to zoom out to the lowest magnification so that the entire sound file will fit on the screen.

SELECTION



Use this command to maximize a selected area in the data window.

The Zoom Selection command will calculate the maximum zoom factor for the size of a selection area, draw the data at this zoom factor, and center the selection in the data window.

NOTE

To zoom in and out in time by small increments, you can use the up and down arrow keys or the Zoom In/Out buttons on the lower right-hand side of the data window.

Shortcut: Control+Up arrow (when a selection is active) or Waveform Display shortcut menu

FOCUS TO DATA WINDOW

Use this function to return the screen focus to the sound file's data window. This is does the same operation as left-clicking anywhere on the data window.

Shortcut: Alt+0

REGIONS LIST



Use this command to open or close the Regions List window. The Regions List contains all of the regions and markers that exist in the currently selected data window. To play a region, press on the Play button to the left of each region. Selecting a marker or region moves the cursor or selection in the data window.

Shortcut: Alt+1 to turn on or make active.

PLAYLIST

Use this command to open or close the *Playlist* window. The *Playlist* is used to arrange regions for playback. To re-arrange regions, drag and drop them within the *Playlist*. Pressing the *Play* button to the left of each region plays from that region onwards down the list.



Shortcut: *Alt+2* to turn on or make active.

KEYBOARD

Use this command to open or close the *Keyboard* window. The *Keyboard* allows you to send *MIDI note on* and *off* commands to your sound card or external sampler.

Shortcut: *Alt+3* to turn on or make active.

MIXER

Use this command to open or close the *Mixer* window. The *Mixer* controls volume and record levels of sound cards that support this feature.

Shortcut: *Alt+4* to turn on or make active.

TIME DISPLAY

Use this command to open or close the *Time Display* window. The *Time Display* window displays the current cursor or play position.

Shortcut: *Alt+5* to turn on or make active.

VIDEO PREVIEW

Use this command to open or close the *Video Preview* window. This window shows the video frame at the current cursor or play position.

Shortcut: *Alt+6* to turn on or make active.

PLAY METERS

Use this command to open or close the *Play Meters* window. The *Play Meters* display peak levels during playback.

Shortcut: Alt+7 to turn on or make active.

UNDO/REDO HISTORY

Use this command to open or close the *Undo/Redo History* window. This window will allow you to see all of your edit operations. The most recently performed operations will be at the top of each list.

When you perform an edit operation on your sound file, such as when doing a *Cut*, the operation name appears at the top of the *Undo* list. When you undo the *Cut*, the operation disappears from the *Undo* list and appears in the *Redo* list.

You can audition various versions of undos and redos by pressing the *Play* button next to each. However, some *Undo* or *Redo* operations cannot be auditioned.

To undo or redo up to a certain operation, double-click on the operation in the *Undo* or *Redo* list. If you *Undo* an operation at the bottom of the list, all operations that followed (which are above in the list) will also be undone (this can take some time when working on large files). *Redo* works the same way.

Shortcut: Alt+8 to turn on or make active.

LOOP TUNER

Use this command to turn on and off the *Loop Tuner* window. The *Loop Tuner* shows the point where the end and beginning of the loop meet.

Shortcut: Control+L

THE SPECIAL MENU

TRANSPORT

Transport is used to perform playback and record operations.

RECORD

The record function records data to an existing or new data window. To record to a direct-opened file simply use the *Window* button (explained later).



If the current window is opened direct, opening the Record dialog will default recording to a new window.

NOTE

MODE

Mode selects the current recording mode. Modes are as follows:

Automatic Retake (automatically rewind)

In *Automatic retake* mode, recording begins at the specified *Start* position and replaces any sound data that already exists after that position. When recording is stopped, the next record start position is automatically rewound to the original starting position, allowing you to immediately do a retake which replaces the last take if necessary.

Multiple Takes Creating Regions

The *Multiple takes creating Regions* mode will record successive takes of data and define them as regions. A new region will be defined every time you hit stop.

Multiple Takes (no Regions)

The *Multiple takes* mode will record successive takes of data but will not create a region for each take.

Create a New Window for Each Take

In the *Create a new window for each take* mode, every time you press *Record* and then *Stop*, a new window is created containing only the last take.

Punch-In (record a specific length)

The **Punch-In** mode is used to record over the top of an existing region of sound. This allows you to record a portion of a previous take without recording a whole new file.

START, END, LENGTH, INPUT FORMAT

Use *Start*, *End*, *Length*, and *Input Format* to change the selection. Sound Forge automatically fills in the values depending on your selection or cursor position.

MONITOR

Select the **Monitor** box to monitor the level of the input signal. The meters light up as a function of the volume of the recording input. For best results, the level should be somewhere in the yellow range with an occasional red. This button also activates the peak value. The peak value shows the maximum input level recorded since the monitor input box has been checked. If the peak value reads clip, clipping has occurred and the audio quality has been compromised. Clicking on the peak value display will reset it.

REVIEW PRE/POST-ROLL

The *Review pre/post-roll* boxes specify the amount of time that will be played prior to and after regions when reviewing a take.

RECORD METERS

The following commands and options affect the *Record Meters* display. The *Record Meters* options can be activated by right-clicking anywhere in this dialog.

Reset Clip

This function resets the clip indication on the *Record Meters*. After a clip occurs, the clip indication will remain lit until the *Reset Clip* function is invoked. You can also click on the clip light or levels to clear it.

Resolution

The resolution settings display the minimum and maximum readings for the *Record Meters*. Changing the resolution of the meters is helpful for metering sound files of varying dynamism.

Show Labels

This option turns on or off the markings on the *Record Meters*.

Hold Peaks

This option turns on or off the markings to indicate the highest peaks in the sound file.

Hold Valleys

This option turns on or off the markings to indicate the lowest peaks in the sound file.

GAP DETECTION

Sound Forge detects most gaps by noticing when it is not able to save wave data as fast as the record device produces it. However, it is still possible for very short gaps to occur because of conflicts between drivers. This sort of gap is usually quite short and either happens frequently or not at all. If your application has zero tolerance for gaps, we recommend that you monitor your first several recording sessions closely to determine if your system has this problem.

The *Gap Detection* switch can be reached by right-clicking anywhere in this dialog or in the *System* menu. It allows you to specify what action Sound Forge should take if it detects a high probability of a gap having occurred while recording.

BLINKING STATUS

This option enables or disables blinking of the status indicator. It is reached by right-clicking anywhere in the *Record* dialog.

DC ADJUST

Check *DC adjust* before recording to enable automatic adjusting for any *DC offset* produced by your audio hardware during the record process.

CALIBRATE

When you press the *Calibrate* button, Sound Forge will calculate the *DC offset* while monitoring incoming sound data. The numbers below show the value, in sample amplitude, of the adjustment. If you change sound cards or are recording from different digital sources or at different sample rates, it is suggested that you re-calibrate the *DC Offset* before recording.

SYNC

Pressing the *Sync* button brings up options for syncing Sound Forge and external devices or software.

Enable MTC/SMPTE Output Synchronization

Turn this on to allow Sound Forge to generate MTC/SMPTE synchronization while recording.

Start

Turn this on to specify an exact time when Sound Forge will kick into record.

Pre-roll

Turn this on to begin the SMPTE output at a time before recording.

Enable MTC/SMPTE Input Synchronization On/Off

Turn this on to trigger *Record* of Sound Forge from incoming MTC/SMPTE synchronization.

Start

Turn this on to specify an exact time when Sound Forge will kick into *Record*.

End

Turn this on to specify an exact time when Sound Forge will kick out of *Record*.

Back

Pressing the *Back* button to switch the focus back to the main *Record* dialog.

PREPARE

Pressing the *Prepare* button prepares your system for recording. This means that the record device is opened and all preparation that can be done prior to recording is completed. This allows the system to begin recording as quickly as possible after the *Record* button is pressed. You do not need to use the *Prepare* button unless you are trying to begin recording immediately after pressing the *Record* button.

RECORD

Pressing the *Record* button starts recording.

Shortcut: Control + R

PLAY

Pressing the *Play* button allows you to hear a section of data over which you are going to record or which you have just recorded. This allows you to keep recording until you get the correct take. When pressed the *Play* button turns into the *Stop* button; press again to stop playback.

Shortcut: *Alt + P*

DROP MARKER

Pressing this button while recording drops a marker at that point. This useful feature allows for easy location of mistakes or points of aural interest.

Shortcut: *Alt + M*

GO TO BEGINNING

Pressing this button returns the cursor to the beginning of the current sound file. Use this to review takes or to begin recording from the beginning.

Shortcut: *Alt + Z*

REWIND

Pressing this button moves the cursor backwards, take by take, in the current sound file. Use this to review previous takes, or simply record over them.

FORWARD

Pressing this button moves the cursor forward, take by take in the current sound file. Use this to navigate through several takes either to review, or record over them.

GO TO END

Pressing this button puts the cursor at the end of the current sound file. Use this to quickly continue recording from the end of the current sound file if the cursor was moved to review the previously recorded material.

TIME RECORDED

The *Time recorded* field displays the total time recorded for the current sound file. The display is in the units displayed in the *Input Format* field above.

TIME LEFT ON DRIVE

The *Time left on drive* field displays the total time left on the drive for recording. The display is in the units displayed in the *Input format* field above.

REMOTE

Press the *Remote* button to enter *Remote Recording* mode. This mode will hide the main *Sound Forge* window and show only the *Record Remote* dialog. This allows you to easily record while using other components of your system like a CD player, mixer, or sequencer.

Many options described in the *Record* dialog section are available by right clicking anywhere in this dialog.

PLAY ALL



Use this function to play the entire sound file from beginning to end, regardless of cursor position, selection, or playlist.

Shortcut: Shift+Spacebar

PLAY



Play will playback the file in current playback mode (*Normal*, *Looped*, or *Sample mode*)

Shortcut: Alt+P or Spacebar, right-click on overview bar

PAUSE



Pause will pause playback and keep the cursor at its current position.

Shortcut: Enter, Shift+Spacebar, right-click on overview bar (during playback)

STOP



Stop will stop playback and return the cursor to its position prior to playback.

Shortcut: Spacebar (during playback)

GO TO START

Go to Start moves the cursor to the start of file.

Shortcut: *Control+Home*



REWIND

Pressing the *Rewind* button moves the cursor backwards, take by take, in the current sound file.



FORWARD

Pressing the *Forward* button moves the cursor forward, take by take, in the current sound file.



GO TO END

Go to End moves the cursor to the end of file.

Shortcut: *Control+End*



PLAY NORMAL MODE

Play Normal Mode sets playback to *Normal* mode. When *Play* is selected while in this mode:

If there is no selection, playback occurs from the cursor to end of file.

If there is a selection, playback occurs from the beginning of the selection to the end of the selection.

Shortcut: To switch between playback modes, use *Control+Spacebar*.



PLAY LOOPED MODE



Play Looped Mode sets playback to *Looped* mode. When *Play* is selected while in this mode: If there is no selection, the entire sound file is played in an endless loop.

If there is a selection, the selection is played in an endless loop.

Shortcut: To switch between playback modes, use *Control+Spacebar*

PLAY AS SAMPLE MODE



Play as Sample Mode sets playback mode to *Sample* mode. When *Play* is selected while in this mode, the file is played with the *Sustaining and Release Loops* repeating a set number of times. Use this to listen to a sound file as it would sound when played by a sampler.

Shortcut: To switch between playback modes, use *Control+Spacebar*

REGIONS LIST

The *Regions List* is used to store regions and markers positions in a file.

ADD MARKER/REGION

Use this function to add a selection to the *Regions List*.

TYPE

A *Marker* is a single point, a *Region* is a start and end point, and has an associated length.

START, END, LENGTH, INPUT FORMAT

Use *Start*, *End*, *Length*, and *Input Format* to change the selection. *Sound Forge* automatically fills in the values depending on your selection or cursor position.

TRIGGER

Use *Trigger* if you want to trigger a region or marker using MIDI commands.

MIDI: Note On - Play

The region will be played when the specified note on message is received and will play for the full length of the region.

MIDI: Note On - Play / Note Off - Stop

The region will be played when the specified note on message is received and will stop when either the full region is played or the specified note off message is received.

MIDI: Note On - Queue / Note Off - Play

The region will be queued for play when the specified note on message is received and will play when the corresponding note off message is received. This is used to reduce the time between receiving a trigger and playing a region.

CHANNEL

Channel determines the MIDI input channel for triggering.

NOTE

Note Determines the MIDI note that will trigger region playback. This value can be entered as either a MIDI note value such as C4 or as a MIDI note number such as 60.

If you have the MIDI Input Sync/Trigger mode enabled while using this dialog, you can auto fill the Channel and Note values by pressing a key on your MIDI Keyboard.

NOTE

Triggers in the Regions List are different from triggers specified in the MIDI Triggers dialog (Options menu) and the Playlist. When using triggers in the Playlist, Regions List, or MIDI Triggers dialog, be aware that they can interact with each other to create unexpected results. Sound Forge first looks at the MIDI Triggers, then the Regions List, and then the Playlist when determining what to do when a MIDI command is detected. If you only want to use the triggers in the Regions List, turn off all the triggers in the MIDI Triggers dialog and the Playlist.

Shortcut: R to create a regions or M to create a marker or Time Ruler shortcut menu or Region Tag shortcut menu or Region List shortcut menu or Press Enter when a region in the Regions List is selected or Double-click on a region in the Regions List or Drag and drop selection to Regions List

DELETE REGION

Use this function to delete a region from the *Regions List*. (You can undo!)

Shortcut: In the Regions List, select a region and press Delete or Region Tag shortcut menu

EDIT MARKER/REGION

Use this function to modify a region in the *Regions List*.

TYPE

A *Marker* is a single point, a *Region* is a start and end point, and has an associated length.

START, END, LENGTH, INPUT FORMAT

Use *Start*, *End*, *Length*, and *Input Format* to change the selection. Sound Forge automatically fills in the values depending on your selection or cursor position.

TRIGGER

Use *Trigger* if you want to trigger a region or marker using MIDI commands.

MIDI: Note On - Play

The region will be played when the specified note on message is received and will play for the full length of the region.

MIDI: Note On - Play / Note Off - Stop

The region will be played when the specified note on message is received and will stop when either the full region is played or the specified note off message is received.

MIDI: Note On - Queue / Note Off - Play

The region will be queued for play when the specified note on message is received and will play when the corresponding note off message is received. This is used to reduce the time between receiving a trigger and playing a region.

CHANNEL

Channel determines the MIDI input channel for triggering.

NOTE

Note determines the MIDI note that will trigger region playback. This value can be entered as either a MIDI note value such as C4 or as a MIDI note number such as 60.

If you have the *MIDI Input Sync/Trigger* mode enabled while using this dialog, you can auto fill the *Channel* and *Note* values by pressing a key on your *MIDI Keyboard*.

Shortcut: Region Tag shortcut menu or Region List shortcut menu or Press Enter when a region in the Regions List is selected or Double-click on a region in the Regions List

REPLICATE

Use this function to create a new copy of the marker/region.

Shortcut: Region List shortcut menu

SPLIT

Use this function to split a region into two separate regions. This split will occur at the cursor position and will split the region in which the cursor lies.

Shortcut: Region List shortcut menu

UPDATE

Use this function to modify the current marker/region location in the *Regions List* to correspond with the current cursor selection in the wave file.

Shortcut: Region Tag shortcut menu or Region List shortcut menu

CLEAR

Use this function to remove all markers/regions from the *Regions List*. This action will also clear the *Playlist*.

MARKERS TO REGIONS

This function will take all of the markers in the *Region List* and convert them into regions using the data between each consecutive marker as the region boundary.

OPEN

This function allows the importing of another region/playlist file into the current sound file. This handy option offers the flexibility of having multiple playlists for the same sound file.

NOTE

Opening a new regions/playlist will clear the present one. Make sure you have saved the current regions/playlist before continuing.

SAVE AS

This function saves the current regions/playlist to an external regions/playlist file. This allows the ability to create multiple regions/playlists for a single sound file.

COPY ONTO CLIPBOARD

Use this function to copy the text of the Regions List onto the clipboard for use with a text editor. Excellent for making a hard copy of the Regions List.

Shortcut: Region List shortcut menu

PLAYLIST/CUTLIST

The Playlist is used to arrange regions from the Regions List. You can open the first playlist in a sound file by selecting Playlist in the View menu.

ADD

Use this function to add a region to an open playlist.

Shortcut: Drag and drop a region from the Regions List to the Playlist or Playlist shortcut menu

DELETE

Use this function to delete a region from the Playlist.

Shortcut: Press the Delete key when a region is selected in the Playlist or Playlist shortcut menu

EDIT

Use this function to add or modify a region in the Playlist.

PLAY COUNT

Play count determines number of times the playlist region will repeat before continuing on to the next region in the Playlist. Setting this value to 0 will cause the region to be skipped during playback.

TRIGGER

Trigger determines the method with which the playlist region will be invoked (triggered).

CHANNEL

If the playlist region is being triggered by a MIDI note, this will be the receive channel for that region.

NOTE

If the playlist region is being triggered by a MIDI note, Note is the MIDI note number that will trigger this playlist region. This value can be entered as either a MIDI note value such as C4 or as a MIDI note number such as 60.

NOTE

If you have the MIDI Input Sync/Trigger mode enabled while using this dialog, you can auto fill the Channel and Note values by pressing a key on your MIDI Keyboard.

SMPTE TIME

If the playlist region is being triggered from a SMPTE time cue, this is the point in SMPTE time that will trigger the region.

PRE-ROLL PLAYBACK

When *Pre-roll playback* is checked, this sets the number of seconds of pre-roll that will be heard from the end of the previous region when starting the playlist sequence from the current region.

For example, say you have two regions, *Region 1* and *Region 2* in the Playlist. If you have the *Pre-roll* time set to 0.5 seconds, when you start playback of the Playlist from *Region 2*, the last 0.5 seconds of *Region 1* will be played back before starting with *Region 2*. This allows you to hear the transition between the two regions.

Shortcut: Playlist shortcut menu or / (forward slash) when a playlist region is selected

REPLICATE

Use this command to create a new copy of the region.

Shortcut: *Playlist shortcut menu*

CLEAR

Use this command to remove all regions from the *Playlist*.

STOP POINT

Designates the current region as a stop point in the *Playlist*. When the *Playlist* is invoked it will continue to play through the regions until it hits a stop point. The *Playlist* will continue to play only after it is invoked again. This is useful when triggering play from incoming MIDI or time code and you only want certain sections of the playlist to be played at a time.

Shortcut: *Playlist shortcut menu* or * (asterisk) when a *playlist region* is selected

PRE-ROLL PLAYBACK

When *Pre-roll Playback* is on, playback from the *Playlist* begins at the specified time prior to the selected start point in the *Playlist*. The amount of *Pre-roll* time can also be set in the *Playlist* page in the *Preferences* folder or in the *Edit Playlist* dialog in the *Special* menu.

For example, say you have two regions, *Region 1* and *Region 2* in the *Playlist*. If you have the *Pre-roll* time set to 0.5 seconds, when you start playback of the *Playlist* from *Region 2*, the last 0.5 seconds of *Region 1* will be played back before starting with *Region 2*. This allows you to hear the transition between the two regions.

Shortcut: *Playlist shortcut menu* or / (forward slash) when a *playlist region* is selected

OPEN

This function allows the importing of another region/playlist file into the current sound file. This handy option offers the flexibility of having multiple playlists for the same sound file.

NOTE

Opening a new regions/playlist will clear the present one. Make sure you have saved the current regions/playlist before continuing.

SAVE AS

This function saves the current regions/playlist to an external regions/playlist file. This allows the ability to create multiple regions/playlists for a single sound file.

COPY ONTO CLIPBOARD

Use *Copy onto Clipboard* to copy the text of the *Playlist* onto the clipboard for use with a text editor.

Shortcut: Playlist shortcut menu

CONVERT TO NEW

Use this command to create a new sound file which contains the regions arranged sequentially as specified in the *Playlist*. The new file will contain the exact sequence of regions you hear when playing from the *Playlist*.

Shortcut: Playlist shortcut menu

CONVERT TO NEW

Use this command to create a new sound file which contains the regions arranged sequentially as specified in the *Playlist*. The new file will contain the exact sequence of regions you hear when playing from the *Playlist*.

TREAT AS CUTLIST

When *Treat as Cutlist* is enabled, the *Cutlist* window will appear. Regions in the *Cutlist* will be skipped during playback when you use the *Cutlist* playback button.

DELETE CUT REGIONS

When *Treat as Cutlist* is enabled, *Delete Cut Regions* will remove cutlist regions from the file.

ASF MARKERS/COMMANDS

The *ASF Markers/Commands* functions are used to add, import and export ASF specific markers and commands to the current sound file.

SAVE AS

This command saves the current set of file markers into an external text file. This text file can then be used to apply the same set of markers and commands to any file.

IMPORT

This command imports the contents of ASF marker/command text files into the current file in Sound Forge. Any markers or regions that are in the current file will be replaced by the ones specified in the text file.

ADD COMMAND

This command will insert a new marker into the file that is pre-formatted for an ASF command. The default command text is for a URL flip. Any other command can be inserted into the text field.

ACID LOOPING TOOLS

These tools are specially designed to make it easy to edit audio into Loops for Sonic Foundry's ACID. However, these tools can easily be used for other loop-based audio applications as well.

HALVE SELECTION

This command cuts the current selection in half. This is useful for subdividing a loop into smaller parts.

DOUBLE SELECTION

This command increases the size of the current selection to twice its size. This is useful for quickly selecting bigger chunks of loops when you already have a measure or beat selected.

SHIFT SELECTION LEFT

This command moves the current selection to the left by the length of the selection. This is useful when you have a subdivision of the loop selected and you want to quickly move to the next equal subdivision (i.e., from bar 2 to bar 1).

SHIFT SELECTION RIGHT

This command moves the current selection to the right by the length of the selection. This is useful when you have a subdivision of the loop selected and you want to quickly move to the next equal subdivision (i.e., from bar 1 to bar 2).

ROTATE AUDIO

This command does one of two things depending on the status of the current selection. If there is no selection, the first quarter of the entire file will be taken off of the beginning and pasted at the end.

If a selection is made at the very beginning of the file, that audio will be cut and pasted at the end of the file. If a selection is made at the very end of the file, that audio will be cut and pasted at the beginning of the file.

If the selection does not touch the very end or very beginning of the file, the function will not work.

EDIT ACID PROPERTIES

Use this command to change ACID specific information in the wave file.

ONE-SHOT

One-Shots do not stretch with tempo or change pitch with the ACID project's key. Sounds like cymbal crashes and vocal samples are usually played as One-Shots.

LOOP

Loops stretch with tempo and may also change pitch with the ACID project's key. Loops are the most common types of samples used in ACID.

NUMBER OF BEATS

This control determines the length, in beats, of the Loop. If a 16 beat Loop is specified as an eight beat Loop, it will play twice as fast as it should in ACID.

ROOT NOTE FOR TRANSPOSING

This control determines the root note that ACID will use to pitch the Loop to the project key. If the root note is A and the project key is C, the loop will be transposed up three semitones.

DISK-BASED

Disk-Based files can stretch with tempo but will not change pitch with the ACID project's key. Long files like vocal tracks are usually considered Disk-Based.

TEMPO

Disk-Based tracks can stretch to the tempo of the project, but the original tempo must be specified. If a tempo is not specified, the Disk-Based track will not stretch.

EDIT SAMPLE

Use this command to change sampler parameters.

SAMPLE TYPE

Selecting a *Sample type* will determine how the sound file is played back.

One Shot

One shot will cause the sound file to play normally (as it appears) with no loops.

Sustaining

Sustaining will cause the file to repeat the *Sustaining Loop* region the specified number of times.

Sustaining with Release

Sustaining with release will cause the sound file to repeat the *Sustaining Loop* region, play the region between the sustaining and release loop and then repeat the *Release Loop* region.

LOOP TO EDIT

Use *Loop to edit* to select the loop currently being edited.

INFINITE LOOP

By specifying *Infinite loop* it is easier to hear how transitions between loops and normal data will sound.

LOOP COUNT

Loop Count specifies the number of times that the Sustaining Loop or Release Loop will be repeated before continuing playback. This information may or may not affect how your sampler plays the data.

START, END, LENGTH, INPUT FORMAT

Use Start, End, Length, and Input Format to change the position and length of the loop.

MIDI UNITY NOTE OF SAMPLE

The MIDI unity note of sample is the MIDI note that will cause a sampler to play the sound file at the pitch (sample rate) it was originally recorded.

FINE TUNE

Some samplers will utilize the Fine tune information to adjust the pitch (sample rate) of the sound file. Sound Forge does not fine tune the sound file itself when this option is utilized. This option is an information setting that will be transmitted to a sampler via a sample transfer procedure. A sampler such as the K2000 can use this information for its playback of the sample. The K2000 should accurately display this information on Master/Sample/Misc. page, it appears as Pitch Adjust, just below Root Key Number.

SMPTE FORMAT

Some sampler editors store a SMPTE offset value in the sample. Sound Forge ignores this offset value. If you want to add SMPTE offset information to your sample, choose a SMPTE format from the SMPTE Format list and type in an offset value. Otherwise leave the format option to (No Offset).

Shortcut: L or Time Ruler shortcut menu or Loop Tag shortcut menu

EDIT FRAME RATE

Use this command to change the *Frames per second* value used by Sound Forge for status information when using *Absolute Frames* or *Time & Frames* format. The included presets are there for your convenience. You can manually enter the frame rate below.

SET FRAME RATE TO:

HALF NTSC VIDEO (14.985 FPS)

Half NTSC video is one half of the standard rate for video broadcasts.

NTSC VIDEO (29.970 FPS)

NTSC video is the standard rate for video broadcasts.

RED BOOK DIGITAL AUDIO (75.000 FPS)

When recorded to CD, the audio is divided into frames which occur at about every 8 ms. Thus, when editing a sound file that will end up in an audio CD, this format is useful for guaranteeing that the start and end of selected audio will exactly match with the final product.

SMPTE NON-DROP (29.970 FPS)

SMPTE Non-drop is the same as *NTSC video*.

VIDEO FOR WINDOWS DEFAULT (15.000 FPS)

Video for Windows Default is the default rate used by *Video for Windows .AVI files*.

FRAMES PER SECOND

Frames per Second determines the number of frames per second used. You can select a pre-set from the list above or manually enter a value.

EDIT TEMPO

Use this command to automatically calculate the musical tempo (*Beats per minute*) from the current selection.

To use the *Edit Tempo* function, first make a selection in the waveform display. The selection made should be equal in size to the length of a beat or measure. The easiest way to tune a selection to be exactly one measure long is to play the selection looped and change the selection points until a constant downbeat is heard.

Sound Forge uses the current tempo to display measures in the *Time Ruler* and in the *Auto Region* tool. Use the process, *Time Compress/Expand* to actually change the tempo of a sound.

START, END, LENGTH, INPUT FORMAT

Determines the selection used to calculate the tempo. These values will be filled in for you based on the selection you have made in the sound file. If no selection is made, the entire file will be used.

SELECTION LENGTH IN BEATS

This control specifies the number of beats in the current selection. Changing this value will automatically change the *Tempo in beats per minute* control.

TEMPO IN BEATS PER MINUTE

This control specifies the tempo of specified selection. Changing this value will automatically change the *Selection length in beats* control.

NUMBER OF BEATS IN A MEASURE

Use to specify the number of beats in a measure when selecting a measure (e.g. 4 for most pop music out there).

CENTER CURSOR

Center Cursor will center the display so that the cursor appears in the center of the data window. This item does not actually move the cursor to a new position in the data window, it simply redraws the display so that you will now see the areas of the sound file equally on either side of the cursor. If the cursor does not seem to center when selecting this, it is because the cursor is either very close to the beginning or end of the sound file, and therefore cannot be centered.



If you have a selection, you can switch the cursor between selection points using *Home* and *End*. You can also switch the cursor between selection points by holding the *Control* key while selecting the *Go to Start* or *Go to End* buttons in the *Playbar* or *Transport toolbar*.

NOTE

Shortcut: C

DROP MARKER



The *Drop Marker* command is a way to mark positions in a sound file. You can drop a marker at the current cursor position while editing or on the fly during file playback.

It's much easier to create a marker from the toolbar or by pressing 'M' on the keyboard rather than the menu (as this would be difficult to coordinate on the fly). Markers can be quickly selected from the list in the *Go To* dialog located in Edit menu. Also, markers are placed in the *Regions List* for quick playback.

NOTE

There is usually a very small delay between the time when you select Drop Marker and when the marker is actually created. Using the keyboard shortcut 'M' gives the fastest response, and thus the highest accuracy.

Depending on the sound card you are using, there might be an offset error between what you hear and where the cursor is located. To correct this, go to the *Wave* page in the *Preferences* folder (Options menu) and use the *Play position bias* to correct for the sound card driver's offset error.

Some sound cards do not return position information to Sound Forge very efficiently. If the cursor moves in large chunks during playback, check the *Interpolate play position for inaccurate devices* to allow Sound Forge to estimate where the sound card playback is occurring.

Shortcut: M

MARK IN



Mark In is similar to the *Drop Marker* function. It will allow you to mark the in point of a new selection while the sound file is playing back. You can also use this as a method of marking the in point of a new selection while editing by moving the cursor to where you want the selection to start. This is a good way to accurately set the start of a selection. For greatest accuracy you should use the keyboard shortcut key '[' when *Marking In* during playback.

If there is a discrepancy between what you hear and the cursor position on screen, use the Play position bias or Interpolate play position for inaccurate devices between fields options on the Wave page in the Preferences folder (Options menu).

NOTE

Shortcut: [(left square bracket)

MARK OUT

Use this function to create a mark out point of a new selection. Mark In and Mark Out are used to determine a selection, and can be used during playback. For greatest accuracy you should use the keyboard shortcut key ‘]’ when marking out during playback.



If there is a discrepancy between what you hear and the cursor position on screen, use the Play position bias or Interpolate play position for inaccuracies between fields options on the Wave in the Preferences folder (Options menu).

NOTE

Shortcut:] (right square bracket)

TOGGLE SELECTION

By using the “S” key (or this menu item), you can toggle back and forth between the last selection and last cursor position. For instance, you may be setting the cursor in various places while navigating throughout a sound file. If you then hit “S”, the last selected region will immediately return to the display.

Shortcut: S

UNDO/REDO HISTORY

CLEAR

Use this function to erase the temporary files created in the undo history for the active file. These files do take up disk space, so you might want to periodically clear them when you are sure that you don't want to undo any of your previous changes. This function only affects the current sound file.

CLEAR ALL

Use this function to clear the undo histories for all open sound files. *Clear all* affects all of the opened sound files. After doing this, you will be unable to undo changes you have made to all modified sound files.

REBUILD PEAK DATA

This function recalculates the peak data file for the active data window. A peak data file is used in displaying a file's waveform without scanning the entire length of the file for the sample amplitudes. This makes screen redraws much faster.

If for some reason you believe that the waveform display is not accurately representing the sound file, you can use this function to force Sound Forge to re-scan the entire file. However, under normal circumstances you should never have to do so since Sound Forge will always update the peak file after every edit operation and when loading a file.

THE PROCESS MENU

The Process menu contains commands that perform various standard audio processes.

AUTO TRIM/CROP



The *Auto Trim/Crop* function is useful for removing unnecessary silence in a sound file and automatically fading in and out the end-points of a phrase.

FUNCTION

The *Function* allows you to select the mode in which the preset will run.

KEEP EDGES OUTSIDE OF THE SELECTION

Silence between the selection points and the sound material is deleted. Use this to remove silence from the edges of a selection. Data outside of the selection is not deleted.

REMOVE EDGES OUTSIDE OF THE SELECTION

Silence between the selection points and the sound material is deleted. Anything outside of the selection is also deleted. Use this to trim a sound from a larger file.

REMOVE SILENCE BETWEEN PHRASES (CREATES REGIONS)

All silence between phrases inside of the selection is deleted. Regions are created between phrases. Use this to delete the silence between sound effects and automatically create regions in the *Regions List*.

REMOVE DATA BEYOND LOOP POINTS

All data after the selected loop is deleted. Since most samplers never play audio beyond a loop, data beyond the loop is unnecessary.

ATTACK THRESHOLD

Instantaneous threshold level used for the detection of the trim/crop start point. -Inf. is complete silence, 0 dB is a maximum amplitude level.

RELEASE THRESHOLD

Instantaneous threshold level used for the detection of the trim/crop end point. -Inf. is complete silence, 0 dB is a maximum amplitude level.

FADE IN

After the trim/crop start point is detected, this is the length of fade in applied to the sound file. This is useful for preventing glitches at end points.

FADE OUT

Fade out controls the length of the fade out applied to the sound file after the detected trim/crop end point.

MINIMUM INTER-PHRASE SILENCE

Minimum amount of silence which must exist between phrases before a new region is created.

MINIMUM LENGTH FOLLOWING LOOP END

When using the Remove data beyond loop points option, this number specifies the number of samples which must exist after the loop. This is used to ensure that a sampler can loop properly, since some samplers don't use the exact loop points supplied.

CHANNEL CONVERTER



The *Channel Converter* is used to convert between mono and stereo formats. It can also be used to inter-mix the left and right channels of a stereo file to create interesting pan effects.

OUTPUT CHANNELS

Output channels determines the number of channels that the output sound file will have.

NEW LEFT CHANNEL

FROM LEFT

From left determines the amount of the original left channel that will be mixed onto the new left channel.

FROM RIGHT

From right determines the amount of the original right channel that will be mixed onto the new left channel.

INVERT LEFT CHANNEL MIX

When *Invert left channel mix* is checked, the new left channel mix will be inverted (polarity reversed).

NEW RIGHT CHANNEL

FROM LEFT

From left determines the amount of the original left channel that will be mixed onto the new right channel.

FROM RIGHT

From right determines the amount of the original right channel that will be mixed onto the new right channel.

INVERT RIGHT CHANNEL MIX

When *Invert right channel mix* is checked, the new right channel mix will be inverted (polarity reversed).

CONVERT TO SPECIFIED OUTPUT CHANNELS ONLY (NO CUSTOM MIXING)

When *Convert to specified output channels only* is checked, the faders and invert checkboxes are disabled. When converting from mono to stereo, each new stereo channel will be an exact copy of the original mono channel. When converting from stereo to mono, the new mono channel will consist of each of the original stereo channels mixed in at 50% volume.

CONVERT TO 8-BIT

Use this function to convert 16-bit sound files to an 8-bit format. Because the signal-to-noise ratio of 8-bit audio is so small (42 instead of 96 dB), it is suggested that the volume of the sound file be maximized by using the *Volume* or *Normalize* functions before performing the conversion.



TRUNCATE

This *Truncate* method of 16 to 8-bit conversion simply removes the least significant bits of a 16-bit sample. It is the least desirable of the three.

ROUND

When rounding, 8-bit values are mapped to the nearest 16-bit value, thereby slightly reducing the quantization noise.

DITHER

Dithering is used to reduce the audible effects of quantization noise which are most audible in 8-bit samples. This is accomplished by adding a small amount of Gaussian noise to the signal to mask the more obtrusive 8-bit quantization distortion.

The dither depth is the amplitude of the added noise in terms of the least significant bits of the sound file. For most cases, a value between 0.5 and 1.2 is sufficient.

NOTE

Use very small amounts of dithering and listen to whether to quantization noise is less obtrusive. Increase the dither amount until the dither noise is louder than the quantization noise.

DC OFFSET

A small square icon with the letters "DC" inside, representing DC offset.

Use DC Offset to change the baseline of a sound file. A recorded wave that is not centered around the zero baseline in the waveform display is said to have a DC offset. To correct for offsets, a constant value is added to each sample.

DC offsets are usually caused by electrical mismatches between your sound card and microphone. Glitches and other unexpected results can occur when sound effects are applied to files which contain DC offsets.

AUTOMATICALLY DETECT AND REMOVE

When *Automatically detect and remove* is checked, the DC offset is calculated for each channel individually and then automatically corrected.

ADJUST DC OFFSET BY

Adjust DC offset by directly controls the offset value. The offset can be from -32768 to 32,767 for 16-bit data and from -128 to 127 for 8-bit data.

NOTE

An easy way to spot DC offset is to zoom close up to a spot in the sound file containing only silence and see if the silence waveform matches the centerline in the waveform display.

An approximation of the DC offset of the waveform can be obtained by running the *Statistics* function under the *Tools* menu. For example, if a DC offset of 100 exists, then you should apply a -100 DC offset to cancel out the existing offset.

COMPUTE DC OFFSET FROM FIRST 5 SECONDS ONLY

When *Compute DC offset from first 5 seconds only* option is enabled, only the first 5 seconds of a sound file are analyzed when measuring the DC offset. The only time that 5 seconds is not sufficient is if a long fade-in or mute has been applied at the beginning of the file.

EQ

GRAPHIC

The *Graphic Equalizer* can boost or attenuate selected frequency bands to alter a signal's frequency spectrum.



BAND ATTENUATION/GAIN FADERS

Each of the ten faders corresponds to the gain or attenuation factor that will be applied to the specified frequency band. The frequency on the bottom of the fader is the center frequency of the frequency band affected by the fader.

The value of the fader affects the frequency band as follows:

- +20 dB Maximum boost
- 0 dB No effect
- 60 dB Maximum attenuation

ACCURACY

Accuracy determines a compromise between filter precision and processing speed. Low precision is not recommended for performing very sharp filtering, when filtering very low frequencies, or when using a high sample rate.

RESET BANDS

The *Reset Bands* button resets each of the ten faders to 0 dB.

PARAGRAPHIC

The Paraphrastic Equalizer is a parametric EQ with a graphical display. It consists of four fully parametric peak filters, a high-shelf and a low-shelf filter.

DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output.

WET OUT

Wet out controls the amount of processed signal mixed into the output.

EQ GRAPH

The *EQ graph* displays the amplitude vs. frequency curve generated from the current settings.

BAND FILTERS**Gain**

Gain determines the amount of boost or cut for the band. To quickly disable a band, set the *Gain* to 0.0 dB by double-clicking on the fader handle.

Width

The bandwidth determines the number of octaves (centered on the selected frequency) which will be affected by the filtering, i.e. sets the size of the band. Use a high value to affect a greater range of frequencies and a low value for a more selective filter.

Center Frequency

Center frequency determines the center frequency for the band.

ENABLE LOW-SHELF

When *Enable low-shelf* is on, frequencies below the low-shelf cutoff frequency will be attenuated or boosted by the amount specified.

ENABLE HIGH-SHELF

When *Enable high-shelf* is on, frequencies above the high-shelf cutoff frequency will be attenuated or boosted by the amount specified.

PARAMETRIC

The *Parametric Equalizer* is a set of four frequency selective filters which allow for very precise changes in the frequency content of a sound signal.

**FILTER STYLE**

The *Filter style* determines the frequency response characteristic of the filter.

Low-pass

Frequencies above f_2 are attenuated by the reject ratio. A low-pass filter is used to restrict the high frequencies of a channel. Possible uses are electrical noise or hiss reduction. Low-pass filters are also used for anti-aliasing a sound file before downsampling.

High-pass

Frequencies below f_1 are attenuated by the reject ratio. A high-pass filter removes the low frequencies of a channel and can be used to remove low-frequency rumbles such as wind, electrical hum, or traffic noise.

This *High-pass filter* works best with frequencies above 150 Hz.

Band-pass (peak)

Frequencies below f_1 and above f_4 are attenuated by the reject ratio. A band-pass filter restricts frequencies higher and lower than a specified band and is useful when trying to isolate or boost a particular frequency range of, for example, voice.

The *Band pass filter* works best when the frequencies are very close together. When they are far apart, ringing can occur.

Band-reject (notch)

Frequencies between f_1 and f_4 are attenuated by the reject ratio. A band-reject filter attenuates a selected frequency band and is often used to remove narrow-bandwidth noise such as amplifier/microphone feedback or 60 Hz electrical hum.

The Band-reject filter works best with frequencies above 150 Hz.

F1 TO F4

Cutoff frequencies for each filter type corresponding to the graph. These frequencies cannot overlap, and have a minimum separation gap which depends on the filter type.

GRAPH

A visual representation of the frequency response curve of the filter is displayed. The virtual horizontal and vertical axis correspond to frequency and gain.

REJECT

Reject determines the amount of attenuation applied to the region specified by the filter style and cutoff frequencies. (0 dB is no attenuation, -60 dB is maximum attenuation)

TOTAL GAIN

Total gain determines the amount of gain applied to the signal after processing. (-60 dB is maximum attenuation, 20 dB is maximum boost)

ACCURACY

Accuracy determines a compromise between filter precision and processing speed. Low precision is not recommended for performing very sharp filtering, when filtering very low frequencies, or when using a high sample rate.

NOTE

By running a filter several times, you can create sharper cutoff curves.

NOTE

A zero-phase filter can be created by using the following procedure:

1. Run a filter on a file.
2. Reverse the file by using the Reverse function.
3. Run the same filter on the reversed file.
4. Reverse the file once again.

By running a filter in both time directions, filter phase shifts are canceled. The filters designed are FIR Chebychev linear phase filters with 15 to 512 coefficients (the number of coefficients determines both the accuracy and processing time of the filter). In some circumstances, the specified filter parameters will be too stringent for the filter design algorithm. This occurs when: filter frequencies are set too close together or too far apart; reject attenuation is set too low; or the total gain is set too high or low.

FADE

Use a fade to vary the volume of a sound file over time.

GRAPHIC FADE

The Graphic Fade dialog allows you to draw a fade envelope which will be applied across the current data selection. Up to 16 envelope points are allowed.



Draw the envelope by pulling the small square boxes (drag points) up or down. You can create a new drag point by left-clicking on any point of the fade envelope. To delete a drag point, single-click on it with the right mouse button, or double-click on it with the left mouse button. Once you have finished with the envelope, press the OK button to apply the fade.

DITHER AND NOISE-SHAPE TO MINIMIZE 16-BIT QUANTIZATION NOISE

When Dither and noise-shape to minimize 16 bit quantization noise is turned on, noise-shaping is applied to the sound file during processing to minimize quantization noise. This effect is most noticeable when the signal is at very low levels, such as near the tails of a fade-in or fade-out.

A fade-out performed without dithering, sinusoidal waves become very blocky and eventually become square waves. You really have to crank up the volume to hear this, but if you do a fade-out on a 440 Hz sine wave, you'll eventually hear all the overtones associated with a square wave of varying pulse widths. This is most apparent when the sine wave drops from -60 to below -80 dB (try it on a 15 second sine wave that drops from -50 to -Inf. dB).

Noise-shaping works by adding a very small dither signal to the audio. The signal is combined with the existing audio in such a way that the noise that is normally generated from quantization is greatly shifted towards the Nyquist frequency. In a 44,100 Hz sampled file, noise is transferred to around 22,000 Hz. Unless you have canine hearing, audibility of such frequencies is naturally attenuated by a tremendous amount. In a way then, the noise floor of 16-bit audio is lowered.

Since quantization noise is transferred to near the Nyquist Frequency, it is strongly advised that noise-shaping not be used with sample rates below 32,000 Hz. Also, it is not recommended that you perform noise-shaping more than once on a given selection, i.e. save the dithering until the very end. However, noise-shaping is not nearly as effective when performed after a fade instead of during a fade, which is why the option to do so is in this dialog.

SHOW WAVE

Press the *Show wave* check box to draw the current selection's waveform on the envelope graph. This is done automatically if the selection is small. For stereo files, you also have the option of displaying the left, right, or a mix of the two channels.

RESET ENVELOPE

The *Reset Envelope* button clears all of the envelope points except for two. These two points are set so that no fading will occur.

FADE IN



Use *Fade In* to linearly fade a selection from a volume of -Inf. to a volume of 0 dB.

FADE OUT

Use *Fade Out* to linearly fade a selection from a volume of 0 dB to a volume of -Inf.



INSERT SILENCE

Use *Insert Silence* to insert a section of silence into a sound file.



INSERT

Insert determines the length of the silence block in seconds.

Cursor

Cursor inserts the silence at the current cursor position.

Start of File

Start of File inserts the silence at the beginning of the file.

End of File

End of File inserts the silence at the end of the file.

You cannot insert silence to a single channel of a stereo file since the length of each channel must always remain equal.

NOTE

*If you wish to shift in time a single channel of a stereo file, you can use the *Delay/Echo* function (Effects menu) to delay one channel backward or forward in time.*

INVERT/FLIP

Use this command to invert the sound file at the base line (reverse polarity). Although this does not make an audible difference, it can be useful for matching a sample transition when executing certain pastes, mixes, or loops.



MUTE



Use this command to set a selection to a volume of -Inf. dB (silence).

NORMALIZE



Use this command to maximize the volume of a selection without clipping. The *Normalize* function scans the audio and applies a gain to raise its level to a specified (often very high) value.

NORMALIZE USING PEAK LEVEL

Selecting the *Peak level* option will normalize the sound file using the maximum (instantaneous) sample values detected. A constant gain is then applied to the sound.

NORMALIZE USING AVERAGE RMS POWER (LOUDNESS)

Checking the *Average RMS power* option will normalize the sound file using the detected average RMS value of the sound file. This is helpful for matching the apparent loudness of different recordings.

NORMALIZE TO

The *Normalize to* value indicates the level to which the sound file will be normalized. For example, when using peaks, if the peak level is -10 dB, and the *Normalize to* setting is set to -3 dB, a constant boost of 7 dB will be applied to the entire file.

When using RMS levels, normalizing to 0 dB means boosting the signal so that it has the same apparent loudness as a 0 dB square wave, meaning incredibly loud. If you were to do so, all of the dynamic range of the signal would be squashed and all the peaks would either be clipped or seriously compressed (see below). The lesson is, normalizing a peak to 0 dB is OK, but normalizing RMS to anything above -6 dB is playing with fire.

SCAN SETTINGS

The *Scan settings* determine the response of the RMS power scan to the dynamics in the sound file.

IGNORE BELOW

Using the *Ignore below* control, you can determine the level of material you want to include in the RMS calculation. Any sound material below the threshold will be ignored in the calculation. This is useful to eliminate any silent sections from the RMS calculation. You should set this parameter a few dB above what you consider to be silence.

If you set this value to minus infinity, all sound data will be used. If the value is set too high (above -10 dB), there is a good chance that the RMS value is always below the threshold. In this case, no normalization will occur. Therefore, it is good to test the threshold by using the *Scan Levels* button.

ATTACK TIME

Attack time determines how fast the scan should respond to transient peaks in the sound file. A slower attack time will tend to ignore fast peaking material.

RELEASE TIME

Release time determines how fast the scan should stop using transient peak material after it has begun to drop in level. A slower release time will increase the amount of material included in the RMS calculation.

USE EQUAL LOUDNESS CONTOUR

Checking the *Use equal loudness contour* option causes the scan to take into effect the *Fletcher-Munson Equal Loudness Contours*. Essentially, very low and high frequency material is less audible than mid-range audio. This option forces the scan to weigh this factor into the RMS calculation.

SCAN LEVELS

The *Scan Levels* button runs the *Peak* and *RMS* scan function on the selection and displays an RMS level as well as the highest peak level found. Use this information to set the *Normalize to level*.

An asterisk (*) following a level indicates that the values displayed are not current. This occurs when you first open the dialog or if you change the selection. If the values have never been calculated, two lines appear where the values will normally be displayed. To update the values, press the *Scan Levels* button.

NOTE

When previewing a *Normalize*, the entire file must be scanned even when previewing a small selection. Pressing the *Scan Levels* button stores the current *Peak* and *RMS* values. This allows you to preview different *Normalize to level* settings without re-scanning the entire file.

If the *RMS* level never reaches the *Ignore below threshold*, a value of *-96 dB* will be displayed. You must then lower the *Ignore below threshold* level before processing.

USE CURRENT SCAN LEVELS (DO NOT SCAN SELECTION)

When *Use current scan levels* is turned on, no scan is performed. Instead, the current scan levels shown are used (which persist when you close the dialog). This is useful if you want to use the scan levels from a different selection or file on the current selection. This allows the same gain to be applied to different files (the gain used is based on the *Scan Level* and the *Normalize to setting*).

One possible use is to scan a small selection of a track that has the loudest material or most constant levels and then using that scan to *Normalize* the entire file or different tracks.

IF CLIPPING OCCURS

These options allow you to choose how you want the *Normalize* function to handle any clipping that could occur as the sound file is processed when using the *RMS* option.

Apply Dynamic Compression

When *Apply dynamic compression* is selected, any peaks that would clip are limited to below *0 dB* using non-zero attack and release times to minimize distortion. In other words, a time-varying gain is used to ensure that no hard clipping occurs.

This option is useful for getting very loud (it is common knowledge that louder is better, right?), yet clear sound during the mastering process.

Also, you achieve best results when converting 16-bit audio to 8-bit or other compressed formats if the audio has been maximized before the conversions occur.

Normalize Peak Value to 0 dB

When *Normalize peak value to 0 dB* is selected, the selection's peak amplitude level is normalized to 0 dB. This applies the maximum possible constant gain which doesn't clip to the selection. Less gain is applied than would be necessary to achieve the *Normalize to RMS* level.

Ignore (saturate)

When *Ignore* is selected any sound data that would clip is allowed to clip. This may not be a problem if the clipping samples are very short and infrequent, but will sound horrendous if much clipping occurs.

Stop Processing

When *Stop processing* is selected any sound data that would clip cause the *Normalize* function to stop processing and display a message to that effect.

*When using *Normalize* on stereo data, if the selection includes both channels, normalization is computed on the loudest sample value found in either channel and the same gain is applied to both. If a single channel is selected in a stereo file then Normalization will effect only that channel.*

NOTE

PAN/EXPAND

The *Pan/Expand* window allows you to draw a pan, stereo expansion, or MS mix envelope which will be applied across the current data selection. You can create up to 16 envelope points.



PROCESS MODE

PAN (PRESERVE STEREO SEPARATION)

Pan (preserve stereo separation) allows you to perform right and left channel panning effects without mixing the channels together. This is used to simulate left/right positioning of stereo recordings.

PAN (MIX CHANNELS BEFORE PANNING)

Pan (mix channels before panning) allows you to perform right and left channel panning effects by first mixing the two channels and then changing the volume between channels.

STEREO EXPAND

In *Stereo expand* mode, you can contract or expand the stereo image of a stereo recording, from dead center (mono), to completely panned wide (no center channel). For realistic effects, you should not venture too far from *Normal*, but do experiment.

To move both endpoints at once, you can select all points by pressing *Control+A* while focus is on the envelope. Dragging one of the points then moves all of the points.

This effect works best with true stereo recordings and will not work with mono material.

MIX MID-SIDE (MS) RECORDING TO LEFT AND RIGHT CHANNELS

Mid-Side recording (MS) is a microphone technique where one mic is used to record the center, or mid, channel (pointed directly towards the source) and the other mic records the stereo image (pointed 90 degrees away from the source). For proper playback on most systems, MS recordings must be converted to your standard left/right (also called AB) track.

In *Sound Forge*, to convert an MS-recorded track to a left/right track, you must first ensure that the *Mid* track is in the left track and the *Side* track on the right. The MS mix function then allows you to set the width of the stereo image for the converted track.

OUTPUT GAIN

Output gain is the gain applied to the signal after processing.

SHOW WAVE

Press the *Show wave* check box to draw the current selection's waveform on the envelope graph. This is done automatically if the selection is small. For stereo files, you also have the option of displaying the left, right, or a mix of the two channels.

RESET ENVELOPE

This *Reset Envelope* button clears all of the envelope points except for two. For the two pan modes, these points are set so that no panning will occur. For the Stereo Expand and Mix Mid-Side modes, these points are set so that normal panning will occur.

RESAMPLE

Resample is used to change the sampling rate of an existing sound file.



NEW SAMPLE RATE

New sample rate determines the sample rate to which the sound file will be converted.

INTERPOLATION ACCURACY

The *Interpolation accuracy* parameter determines the complexity of the method used during the resampling process. The audible difference between the different values can be subtle without the use of test tones and is most apparent in high frequencies. In general, a setting of 1 is more than adequate for general purpose audio. Settings of 2 and 3 are good for high-end audio. Setting this parameter to 4 forces some heavy duty number crunching (i.e. slow processing) but produces very near perfect results, good for audiophiles and audio researchers.

When downsampling by an even multiple (such as when going from 44 kHz to 22 kHz), processing is always fast and perfectly accurate.

NOTE

APPLY AN ANTI-ALIAS FILTER DURING RESAMPLE

Apply an anti-aliasing low-pass filter to the sound file before resampling when changing to a lower sampling rate. Since the maximum frequency that can be represented by a particular sample rate is one-half of the sampling rate (the Nyquist frequency), high frequencies in a sound file cannot be represented if the sound file is resampled to a lower sampling rate. Therefore, when changing to a lower sampling rate, if the sound file has a strong high-frequency content,

anti-aliasing should be used to prevent these high frequencies from becoming low-frequency distortion.

SET THE SAMPLE RATE ONLY (DO NOT RESAMPLE)

If *Set the sample rate only* is checked the playback rate is changed without resampling the data. This means that the original pitch of the file is not preserved. When converting between two very close sample rates and speed is of the essence, this option can be turned on.

N O T E

Resampling data consists of changing the number of recorded samples per second. When changing to a higher sample rate, extra samples are interpolated and the file size increases. When changing to a lower sample rate, some samples are removed and the file size decreases.

REVERSE



Use *Reverse* to reverse a selection.

SMOOTH/ENHANCE



Use *Smooth/Enhance* to add or remove high frequency content in a sound file selection.

The *Smooth* function will smooth out fast-changing transients in a sound and is useful for removing glitches.

The *Enhance* function boosts the very highest frequencies (close to the Nyquist rate) in the sound file, making the sound file sound more vivid. It is useful for compensating for the effects of downsampling or for bringing out very fast transients in a sound file

OPERATION

The *Operation* determines the amount of smoothing (high frequency attenuation) or enhancement (high frequency boost) is applied to the sound file.

NOTE

For a very drastic high frequency boost or cut, repeat the effect 2 or 3 times with the amount set to either limit.

Note that since the enhance function boosts frequencies near the Nyquist rate, the current sample rate will determine what frequencies will be affected. At 44,100 Hz, the effect can be very subtle (unless you repeat it a few times until your ears ring).

TIME COMPRESS/EXPAND

The *Time Compress/Expand* function changes the duration of a sound file without altering the pitch.



MODE

Mode selects the algorithm used to process the sound. The modes have been labeled for suggested uses, but do not assume that one mode will not work well with a certain type of sound. The A, B, and C modes use different algorithms, with the numbers indicating different internal parameters. Depending on the source sound, different modes can vary the resulting quality drastically, so experimenting with all modes is suggested.

You should be able to achieve excellent results for ratios between 75% and 115%. Beyond this range, you will start to hear artifacts such as echoes, flanging, or drop-outs. The different modes are supplied so that you can determine what works best for your case.

In many cases, many of the modes will work well, but in some instances (such as complex, multi-instrument music or pitched drums) you'll have to try out all of the modes. In general, the settings designed for music will work with the largest variety of material.

Also, running the process a number of times using small increments (like 105%) will create different effects that processing all at once with a large time change.

FINAL LENGTH

Final length controls the desired length of the selection. Note that only the A modes allow for expansion above 150%.

To change the *Original tempo* displayed at the bottom of the dialog, you must go to the *Edit Tempo* dialog (Special menu).

NOTE

You cannot perform this function to a single channel of a stereo file since the length of the channels must always remain equal.

BLEND EDGES WITH ADJOINING DATA

When *Blend edges with adjoining data* is turned on, the beginning and end of the selections will be crossfaded with the original sound data. This is used to minimize glitching.

VOLUME



Use *Volume* to change the volume of a selection.

GAIN

Negative decibel values lower the volume (attenuation). Positive values are used to raise the volume (boost). A value of negative infinity corresponds to a mute (0 %)

THE EFFECTS MENU

The *Effects* functions allow you to apply different special effects functions to the sound data. Most of the functions can be applied to single channels within stereo files as well as areas selected within the file.

AMPLITUDE MODULATION

The *Amplitude Modulation* effect applies a sinusoidal or square shaped periodic gain to the input signal. The frequency of the gain waveform can be specified to create effects varying from a slow tremolo to unusual sound distortions.

A small square icon with the letters 'AM' inside, representing the Amplitude Modulation effect.

DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output.

MODULATED OUT

Modulated out controls the amount of processed signal mixed into the output.

MODULATOR SHAPE

Modulator shape controls the shape of the gain waveform. Sine wave amplitude modulation is smooth, while square wave modulation creates drastic on/off changes.

MODULATOR FREQUENCY

Modulator frequency controls the frequency of the gain waveform to be applied to the input signal. To achieve a slow tremolo, use a low frequency (Period = $1/\text{Frequency}$). With faster frequencies, the amplitude modulation is audible not as a change in amplitude, but as additional frequency side bands.

MODULATE CHANNELS OUT OF PHASE (PAN)

When *Modulate channels out of phase* is on, the two channels' amplitude gain envelopes are out of phase. This creates the effect of panning back and forth between the two channels. Panning is only available for stereo files.

CHORUS

The *Chorus* effect creates the illusion of two or more sound sources playing together. This is done by adding a pitch-modulated and delayed version of the input signal to the unprocessed input signal. The effect simulates the variances



in pitch and timing that occur naturally when two or more people try to play or sing the same thing at the same time.

INPUT GAIN

Input gain controls the gain applied to the signal before processing.

DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output.

CHORUS OUT

Chorus out controls the amount of processed signal mixed into the output.

CHORUS SIZE

Chorus size controls the number of times the selection is processed with the chorus algorithm. Larger numbers create a fuller sound.

FEEDBACK

Feedback determines the percentage of the processed signal which is fed back for re-processing.

INVERT THE CHORUS PHASE

When *Invert the chorus phase* is checked, the processed output will be inverted (180 degree phase) before mixing with the *Dry out*. This can sometimes have a large effect on the final sound.

CHORUS OUT DELAY

Chorus out delay determines the time, in milliseconds, that the chorused signal is delayed from the input signal.

MODULATION RATE

Modulation rate controls the frequency, in Hertz, of the modulating signal.

MODULATION DEPTH

Modulation depth controls the amount of frequency modulation applied to the signal.

ATTENUATE FREQUENCIES ABOVE

Attenuate frequencies above determines the cut-off frequency, in Hertz, of the low-pass filter used on the processed signal. Setting to off means no low-pass filtering occurs.

NOTE

A wide variety of non-chorus-like effects can be created with this function. If the modulation depth is high, a vibrato effect will occur. If the delay is small, flanging occurs.

DELAY/ECHO

SIMPLE

The Delay/Echo effect adds a delayed copy of the sound signal to the file. A single delay or decaying multiple delays can be added before or after the sound signal.



DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output.

DELAY OUT

Delay out controls the amount of delayed signal mixed into the output.

MULTIPLE DELAYS/ECHOES

When Multiple delays/echoes is off, a single delayed copy is added to the file.

When on, multiple delayed copies are added by using feedback (delaying the delayed signal, and so on...), which creates a decaying echo effect.

PRE-DELAY/ECHO

When Pre-delay/echo is off, the delayed copy is added after the sound, as heard in everyday life echoes.

When on, the delayed copy is added before the sound, as heard in the twilight zone.

DELAY TIME

Delay time determines the amount of time between the delayed copy and the original sound. Also, if *Multiple delays/echoes* is checked, subsequent echoes occur at intervals specified by the *Delay time*.

DECAY TIME

Decay time determines the time it takes the multiple echoes to become nearly inaudible. Technically speaking, this is the time it takes an echo to decrease in amplitude by a factor of 30 dB. If the *Decay time* is less than the *Delay time*, no echo will be created.

You can also use the *Delay/Echo* function to shift one channel of a stereo file backward or forward in time. For example, if you have a file with music in the right channel and voice in the left, make a selection in only the left channel around the voice. Set the following parameters: *Dry out level* at 0%, *Delay out level* at 100%, no *Multiple delays/echoes*, no *Pre-delay/echo*. You can then use the *Delay time* parameter to determine the time shift. If you want to shift the vocals backwards in time, check the *Pre-delay* check box.

To simulate a stereo sound from a mono file, copy the sound to both left and right channels of a stereo file and then do a very small time delay of a single channel (*Dry out level* at 0%, *Delay out level* at 100%, no *Multiple delays/echoes*, no *Pre-delay/echo*). This simple technique often works in making the sound have a more realistic stereo image. When a sound wave comes from your side, your right and left ears receive sounds at different times, which is what this effect simulates.

MULTI-TAP

The *Multi-Tap Delay* effect is used to simulate multiple echoes, reverberation and other delay-related effects. It accomplishes this with multiple delay taps, pitch modulation, and filtering.

You can use up to eight independent delay taps. Each delay tap consists of a delay time and amplitude specification for a particular echo of the original sound. These are used to simulate the early reflection echoes which are the result of the first sound reflections in a room. Feedback, in combination with filtering and modulation, is then used to simulate the thousands of other reflections that combine to form what we hear as reverberation.

ECHOGRAM

The *Echogram* represents the amplitude of sound reflections over time (impulse response) as determined by the current settings. Each vertical line represents an echo of the original sound. The length of each line corresponds to its amplitude (as a percentage of the original sound), while its horizontal distance from the left edge represents the time elapsed after the original sound. The red line is the currently selected tap, black lines are the other active taps, and blue lines are echoes resulting from feedback.

The *Echogram* can be used to get an estimate of the overall decay time. The right-most horizontal point corresponds to a delay time equal to the graph resolution. This value can be changed between 500, 1,000, 3,000, and 5,000 milliseconds.

INPUT GAIN

Input gain controls the gain applied to the signal before processing.

DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output.

WET OUT

Wet out controls the amount of processed signal mixed into the output.

MODULATOR RATE

Modulator rate determines the frequency, in hertz, of the modulating signal. 0.1 Hz is slow while 10 Hz is very fast modulation.

MODULATOR DEPTH

Modulator depth controls the amount of frequency modulation applied to the reverberation. Zero signifies no modulation, and 100% is a large amount of modulation. Modulation is used to create a fuller sound or chorusing effects.

FEEDBACK

Feedback determines the percentage of the processed signal which is fed back for re-processing and thus creating more echoes. Large feedback values allow for longer decay times but can cause ringing.

LOW-PASS START FREQUENCY

Low-pass start frequency determines the cut-off frequency, in Hertz, of the low-pass filter (high frequency damping) used on the processed signal. Setting to Off means no low-pass filtering occurs. A high frequency setting (or Off) will result in live room reverberation, while lower settings will deaden the sound by simulating the high-frequency damping which many materials impart on sound.

Most building materials and air itself absorb more high frequencies than low frequencies, so using a low-pass filter is often necessary to prevent the reverberation from sounding too live.

NUMBER OF TAPS

Number of taps controls the number of delay taps used to create the reverberation.

CURRENT TAP

Current tap is used to select the delay tap to be modified.

For each active tap there is a corresponding:

Delay

Delay controls the delay time of the tap. Since sound travels close to 1 foot per millisecond, you can simulate a given room size by setting delays to match reflections from all walls in the room. In a large room, sound must travel further to reach the reflection surfaces. Therefore, setting the echoes further apart creates a sense of being in a larger space.

Gain

Gain controls the loudness level of the echo. Negative gain will cause the echo to be opposite in polarity from the original sound.

Pan

Pan controls the channel placement of the echo, with -100% being hard left and 100% being hard right. Panning echoes to different channels creates the illusion of reverberation coming from all directions in stereo files.

GRAPH RESOLUTION (DELAY MS)

Graph resolution determines the maximum time in milliseconds displayed on the graph.

DISTORTION

The Distortion effect can simulate the overloading of an amplifier or mangle a waveform into submission.



DRY OUT

Amount of unprocessed signal mixed into the output.

DISTORTED OUT

Dry out controls the amount of distorted signal mixed into the output.

THRESHOLD LEVEL

Threshold level signals above this level will be distorted.

CLAMP LEVEL

Signals over the Threshold level are set or clamped to this level. For normal limiting, set the Clamp level to be equal to the Threshold level. Setting this level anywhere else will make your speakers do things they've never done before.

CURVATURE

Curvature determines the curvature of the limiting above the Threshold level. Minus infinity corresponds to hard clipping (no curvature) while zero decibels corresponds to a very smooth curve, with little distortion.

HARD CLIP FOR MAXIMUM SPEAKER THRASHING

When Hard clip for maximum speaker thrashing is on, signals over the Threshold level will be set to the Clamp level. When Off, signals over the Threshold level will be soft clipped according to the Curvature.

DYNAMICS

GRAPHIC



Compression, expansion, and limiting affect the dynamic range of the sound file by varying the gain as a function of the input sound level.

A compressor lowers the dynamic range of a signal by reducing the level of high volume signals and applying an overall gain to raise the level up again. Compression is often used in vocals and music to get even volume levels which often results in a smoother, fuller sound.

A limiter is a compressor with a greater ratio of dynamic range reduction. Limiters are used in broadcasting for making sure that the volume level does not go over a certain threshold which overdrives electrical equipment. Limiters are also used by guitarists to create a distorted, sustaining tone.

An expander increases the dynamic range of a signal by decreasing the level of low volume signals. Expanders are commonly used for noise reduction or for emphasizing the dynamics of a piece.

DYNAMICS INPUT/OUTPUT GRAPH

This graph maps each possible input level to an output level. The *No Gain* diagonal line shows where the input and output are equal. When the envelope is below the *No Gain* line, signals of the corresponding input level are attenuated (gain is less than 0 dB). On the other hand, when the envelope is above the *No Gain* line, the signal is boosted (gain is greater than 0 dB).

With this in mind, it is possible to draw graphs to generate compression, limiting, noise gating, and expansion effects by moving the envelope points manually or by using the *Threshold* and *Ratio* controls. You can draw the graph by pulling the small square boxes (drag points) up or down. You can create a new drag point by left-clicking on any point of the envelope, to a maximum of 16 points. To delete a drag point, single-click on it with the right mouse button, or double-click on it with the left mouse button.

OUTPUT GAIN

The gain applied to the signal after processing. If *Auto Gain compensate* is turned on, the output gain is still applied afterwards.

RESET GRAPH

Press the *Reset Graph* button to reset all envelope points.

SYNC STEREO GAIN

When *Sync stereo gain* is on, the gain of the left and right channel will always be the same. This prevents a loss of stereo imaging that can occur if both channels were to be processed independently.

AUTO GAIN COMPENSATE

When *Auto gain compensate* is on, a gain is applied during processing to try to keep the maximum input and output levels constant. In general, this gain will be equal to the highest envelope point in the graph. You can still use the *Output gain* control to fine tune the overall gain.

ATTACK

Attack controls the time, in milliseconds, required for the gain to change from a pre-threshold level to a post-threshold level.

RELEASE

Release controls the time, in milliseconds, required for the gain to change from a post-threshold level to a pre-threshold level.

THRESHOLD

Threshold determines the signal level which differentiates the *Attack* and *Release* modes. In a simple compressor or gate, the threshold is the level at which the gain begins to deviate from 1.

When changing the threshold, all points slide diagonal up and down the *No Gain* line.

RATIO

Ratio of input to output levels above the *Threshold*. This *Ratio* determines how much a signal above the *Threshold* will be boosted or faded as a function of its level.

When changing the *Ratio*, all points above the *Threshold* move closer or further away from the *No Gain* line.

NOTE

For best results, very low frequencies (below 20 Hz) present in a sound file should be removed before applying the Graphic or Multi-Band Dynamics functions. This makes it easier to monitor audible frequencies while compressing.

MULTI-BAND

The Multi-Band Dynamics dialog allows you to apply compression and limiting to up to four different frequency bands. Band limiting is used to keep a set of frequencies from being too loud. The most common use for band limiting is the de-esser. In this case, the frequencies associated with sibilant sounds (S's) are limited to a set level, without affecting other frequencies.

Because there are so many different parameters available, getting the best results requires a bit of experimentation. For most situations, you will only need to use a single band. Each band can either be a low-shelf, band-notch, or high-shelf. After you've selected the frequency range that will be affected, you must set a compression Amount. A value less than 10:1 is considered compression while greater values will create harsher limiting.

The Threshold level determines at what point will compression begin. You can find your required Threshold level by first setting it a high level (such as -6 dB) and slowly lowering it while previewing until you can hear the band-limiting effect.

You can also get an estimate of the Threshold level necessary by first selecting the portion of sound which contains the offending frequencies, bring up the dialog, and after selecting the frequency range pressing the Get Threshold button. The Threshold setting will change to 6 dB below the detected peak of the current frequency band. You'll probably then need to lower it slightly to increase the effect.

NOTE

To reset all bands, right-click anywhere on the dialog and select Reset All. This can be used in most function dialogs.

GRAPH

The *Multi-Band Dynamics Graph* gives you an idea of where the active bands are located in the frequency spectrum and the relative threshold and amount levels. The horizontal lines indicate the threshold settings of the active bands, the current band being red. The amount of compression is indicated by how far down the dark gray curve goes down. Therefore, the amount of dark gray is representative of the amount of gain reduction that can occur.

It is recommended that the bands not overlap. When the bands overlap the results are not easily visualized.

NOTE

OUTPUT GAIN

The *Output gain* controls the gain applied to the signal after processing.

SYNC STEREO CHANNELS

When *Sync stereo channels* is on, the gain of the left and right channel will always be the same. This prevents a loss of stereo imaging that can occur if both channels were to be processed independently.

NUMBER OF BANDS

Number of bands determines the number of active, independent bands. For most cases, only a single band is necessary.

CURRENT BAND

Since each band has an independent mode, threshold, amount, attack and release times, you must use this control to switch between all active bands.

THRESHOLD

When a frequency band level is above the threshold, it will be compressed. A lower threshold will cause more of the signal to be compressed. If the threshold is too low, too much of the sound data will be filtered. If it's too high, not enough processing will occur.

GET THRESHOLD

When you press the *Get Threshold* button, the current selection is analyzed to get an estimate of the necessary threshold. After analyzing, the *Threshold* will be set to 6 dB below the peak level found for the current band.

AMOUNT

Amount determines the compression ratio applied to a band when it is above the threshold level. A high compression ratio means more gain reduction will be applied.

ATTACK TIME

Attack time determines the time, in milliseconds, required for the gain to change from a pre-threshold level to a post-threshold level.

RELEASE TIME

Release time determines the time, in milliseconds, required for the gain to change from a post-threshold level to a pre-threshold level.

TYPE

Type determines the frequency band type.

Low-shelf

Frequencies below the cut-off frequency will be compressed.

Band-notch

Frequencies surrounding the center frequency will be compressed.

High-shelf

Frequencies above the cut-off frequency will be compressed.

GAIN

Gain controls the gain applied to a band after it has been compressed. This is used for compensating for the gain lost from the compression.

WIDTH

Width determines the frequency range that will be compressed when using *Band-notch* mode.

FREQUENCY (BELOW, CENTER, ABOVE)

If the band mode is a low- or high-shelf, *Frequency* determines the cut-off frequency for which the filter starts processing. Otherwise, it is the center frequency of the band.

ENVELOPE

With the *Envelope* effect, the amplitude envelope of a waveform can be forced to match a specified envelope shape. Unlike the *Graphic Fade*, which simply fades a waveform by a specific amount over time, the gain at each point is dynamically calculated to achieve the exact specified envelope.



ENVELOPE GRAPH

The *Envelope Graph* specifies the amplitude envelope over time. Click the left mouse button to define a point and drag to change its position. Click the right mouse button or double-click the left mouse button over an existing point to remove it. You can create up to 16 envelope points.

RESET ENVELOPE

Press the *Reset Envelope* button to reset the graph to a simple ADSR (Attack, Decay, Sustain, Release) curve.

SMOOTH GAIN TO MINIMIZE DISTORTION ON STEEP SLOPES

When *Smooth gain* is enabled, prevents the gain from changing too fast, which might result in unwanted distortion. Also, when this option is on the gain will always begin at 0%.

SHOW WAVE

Select the *Show wave* check box to draw the waveform of the current selection on the envelope graph. This is done automatically if the selection is small. For stereo files, you also have the option of displaying the left, right, or a mix of the two channels.

FLANGE/WAH-WAH



FLANGE

The flanging effect is heard in many 60s and 70s recordings. It is the result of mixing a modulated delay signal with the original signal to create a sweeping sound much like an airplane taking off.

PHASER

A phase is similar to a flange, except that instead of using a *Simple Delay*, frequencies are phase-shifted over time. When used on stereo files, strange effects on the stereo image are created.

WAH-WAH

The wah-wah effect was popularized by many a rock/blues guitarists favorite toy, the wah-wah pedal. It consists of a band-pass filter which attenuates low and high frequencies in varying amounts over time.

DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output. A negative percentage means that the signal is also inverted.

WET OUT

Wet out controls the amount of processed signal mixed into the output. A negative percentage means that the signal is also inverted.

RATE

Rate determines the frequency, in Hertz, of the modulating signal, which determines the period of the flange. Low frequencies create a slow, sweeping flange while higher frequencies create fast flanging.

DEPTH

In the *Flange* effect, *Depth* determines the amount of frequency modulation applied to the signal.

In the *Phaser* effect, *Depth* determines the range of frequencies which experience phase shifting.

In the *Wah-Wah* effect, *Depth* determines the sweep range of the band-pass filter.

CENTER FREQUENCY

Center frequency controls the frequency around which the sweep of the *Phaser* and *Wah-Wah* is centered.

RESONANCE

Resonance controls the amount of band-pass filtering (in *Wah-Wah*) or phase shifting (in *Phaser*) which takes place. When the resonance is high, the range of frequencies processed is smaller, yet more pronounced.

Flange, Phaser, and Wah-Wah are best heard with long, sustaining sounds so that the entire sweep can be heard.

NOTE

GAPPER/SNIPPER

The *Gapper/Snipper* cuts chunks from the sound file or inserts silence in the sound file periodically at a set frequency to create unusual (if not useful) effects.



FUNCTION

GAPPER (INSERT)

In *Insert* mode, chunks of silence are inserted into the sound file.

SNIPPER (CUT)

In *Cut* mode, chunks are removed from the sound file.

FREQUENCY TO GAP/SNIP EVENTS

Frequency to gap/snip events controls the number of times per second a chunk is inserted or cut from the sound file.

LENGTH OF ONE EVENT

Length of one event controls the size of the chunk to be inserted or cut.

FADE EDGES OF EACH EVENT BY

Fade edges of each event by determines the percentage of the chunk not cut or between inserts that is faded in and out. Zero percent corresponds to no fade, while 50% is maximum fade. Use fading to prevent glitches.

NOTE

This process changes the size of the original file. Be aware of the final size indicated, since it is easy to set the parameters to create a very large output.

NOISE GATE



A noise gate removes signals below a set threshold. It's used to remove noise from silent breaks in a sound file.

THRESHOLD LEVEL

Signals below the *Threshold level* will be removed from the sound file. Zero dB is a very high level, while -30 dB is very low. Noise levels are usually around -20 dB.

To get an estimate of the noise level, select a region of silence (which also contains noise) and run the *Statistics* tool. The maximum value (positive or negative) shown is a good estimate of the threshold noise level required. For example, if the maximum positive value is -50 dB and the maximum negative value is -45 dB, use a number slightly higher, maybe -40 dB, as the noise threshold level.

ATTACK TIME

Attack time controls the time it takes the gain of the gate to change once the level rises above the threshold. The attack time is best kept to a minimum if percussive attacks are to be preserved. Larger attack times make sounds slowly swell up in volume.

RELEASE TIME

Release time controls the time it takes the gain of the gate to change once the level falls below the threshold. The release time is usually kept large to allow for natural sounding decays, otherwise long decays will be cut off.

PITCH

PITCH BEND

Pitch Bend is used to draw an envelope that corresponds to increasing or decreasing the pitch of a sound file over time.



ENVELOPE

The *Envelope* specifies *Pitch Bend* over time. The vertical axis represents pitch change, with zero being no change. The horizontal axis represents the length of the selected region.

Click the left mouse button to define a point and drag to change its position. Click the right mouse button over an existing point to remove the point. You can create up to 16 envelope points.

RESET ENVELOPE

The *Reset Envelope* button removes all points except the outermost two.

RANGE

Range determines the maximum and minimum pitch change in *Semitones* (half-steps). Example: A range of twelve half-steps (one octave) allows an increase or decrease of the pitch by an octave.

SHOW WAVE

Select the *Show wave* check box to draw the waveform of the current selection on the envelope graph. This is done automatically if the selection is small. For stereo files, you also have the option of displaying the left, right, or a mix of the two channels.

Pitch Bend is accomplished in the same manner as changing the playback speed on a tape deck. Therefore, the length of the file will be changed.

NOTE

PITCH SHIFT



This Pitch Shift effect changes the pitch of a selection with or without preserving the duration of the sound file.

SEMITONES TO SHIFT PITCH BY

Semitones to shift pitch by determines the number of semitones (half-steps) to raise or lower the pitch. When using *Preserve duration*, the range is limited to -12 to 12 semitones.

CENTS TO SHIFT PITCH BY

Cents to shift pitch by is used to fine tune above and below a semitone. One hundred cents is equal to one semitone.

APPLY AN ANTI-ALIAS FILTER DURING PITCH SHIFT

Changing the pitch of a file requires a form of resampling. When shifting by large amounts, an anti-alias filter will minimize distortions due to aliasing.

PRESERVE DURATION

If *Preserve duration* is enabled, the sound file length will remain the same; otherwise, the length of the selection will change by $1/\text{transposition ratio}$.

MODE

The *Mode* determines what algorithm is used to do the pitch shifting. There are three different algorithms, A, B, and C, each with different modes tuned for different sound material. It is best to experiment with different modes since the names for these modes are only suggestions.

BLEND EDGES WITH ADJOINING DATA

When this option is turned on, the beginning and end of the selections will be crossfaded with the original sound data. This is used to minimize glitching.

TRANSPPOSITION RATIO

Transposition ratio determines the fraction by which all frequencies will be multiplied. If the *Preserve duration* check box is off, the length of the file will be multiplied by $1/\text{transposition ratio}$.

NOTE

Pitch change without preserving duration is accomplished in the same manner as changing the playback speed on a tape deck. Therefore, the length of the file will be changed. A pitch change can be used to accurately change the time duration of a selection if the change in pitch is not audible or important. This is useful if a very small change in time duration is necessary, for example, to correct for playback/record speed discrepancies on a tape deck.

REVERB

Reverb is an effect to simulate various acoustic spaces. It consists of early reflections, which are the first reflections that arrive back to your ear, and the reverb itself.



REVERBERATION MODE

The *Reverberation mode* affects the quality of the reverberation tail. Each mode has slightly a different frequency response, ringiness, echo density, and delay time to simulate different environments.

DRY OUT

Dry out controls the amount of unprocessed signal mixed into the output.

REVERB OUT

Reverb out controls the amount of processed reverberated signal mixed into the output.

EARLY OUT

Early out controls the amount of processed early reflections mixed into the output.

EARLY REFLECTIONS STYLE

The *Early reflections style* affects the first few echoes heard immediately following the dry source and have a strong influence on the size of the perceived space.

DECAY TIME

Decay time determines the time it takes for the reverb to decrease in level by 60 dB. Or, in other words, this is how long your reverb will last.

PRE-DELAY

Pre-delay is the time it takes from the initial sound until the onset of reverb. This setting helps control the perceived size of the room. A longer *Pre-delay* time will imply a larger room.

ATTENUATE BASS FREQUENCIES BELOW

Use the *Attenuate bass frequencies below* option to remove low frequencies in the reverberation. Too much low bass in a reverb will result in a muddy sound.

ATTENUATE HIGH FREQUENCIES ABOVE

Selecting high frequencies filters the reverb above the set frequency. This setting has a lot to do with the perceived liveliness of the room. In a very lively room, high frequencies are not attenuated when they bounce off of the walls.

VIBRATO



Vibrato is the periodic pitch modulation often used by vocalists and instrumentalists. Very fast pitch modulation creates wacky FM effects unlike any vibrato.

MODULATION RATE

Modulation rate controls the frequency, in Hertz, of the modulating signal. Determines the period of the vibrato.

Low frequencies (0.1 to 2 Hz) create a slow vibrato while higher frequencies (2 Hz to 15 Hz) create fast pitch bends. Even higher frequencies (15 Hz and up) modulate the selection so fast that instead of hearing pitch changes, new sideband frequencies are heard.

MODULATION DEPTH

Modulation depth determines the amount of pitch change applied to the selection.

TRIM THE END IF A CHANGE OF DURATION OCCURS

If the output selection length is shorter than the input length, the selection end is trimmed. This is used to prevent glitches when using a slow vibrato inside a selection.

THE TOOLS MENU

The Tools menu contains an assortment of utilities for various purposes.

AUTO REGION

The *Auto Region* function automatically creates regions in a sound file for the *Regions List* and *Playlist*. These regions are defined according to fast sound attacks (like drum beats or words) or a selected musical time interval.



ATTACK SENSITIVITY

Attack sensitivity determines how sensitive the attack-detection algorithm is to fast increases in volume. With a high setting, regions are created when the sound level increases by very small amounts and more regions are created. With a low setting, the sound level must increase by a large amount before a new region is created and fewer regions are created.

RELEASE SENSITIVITY

Release sensitivity determines the minimum decrease in sound level which must occur before a region end is created. High values indicate that the sound level does not have to decrease very much before a region end is created, and

therefore more regions will be found. When set low, the sound level must decrease to a greater extent before a region end is found. A low setting is useful if you only want a region to be created after very silent breaks.

MINIMUM LEVEL

Minimum level determines the threshold sound level which must be found before a new region is created. With a high setting, only high-level sounds will trigger the creation of a new region. This is useful if you want only the loud instrument attacks in a song (like the bass drum) to be at the beginning of a region, since they often correspond to the beginning of a measure or beat. Low threshold settings will also allow low-level sound attacks to create new regions.

MINIMUM BEAT DURATION

Minimum beat duration specifies the minimum length in seconds that must elapse before a new region can be created. A low setting will allow very short regions to be created if sound attacks occur in fast succession. This is useful for fast tempo music. A higher setting will prevent quick sound attacks from being separated into different regions.

USE RELEASE POINT FOR END OF REGION

When *Use release point for end of region* is enabled, a region will end when the sound level drops by a factor determined by the *Release sensitivity*. This is useful if you don't want the silence in between sounds or phrases to be included in the regions.

When off, region ends are only created when attacks are detected.

BUILD REGIONS USING THE CURRENT TEMPO

When *Build regions using the current tempo* is enabled, regions are created according to the current *Beats per minute* (see *Edit Tempo* function under the Special menu to change this value). Otherwise, regions are created by using an algorithm that finds sound level changes to determine where a sound attack or release occurs.

MEASURES, BEATS

Measures and *Beats* determine the musical time interval between regions according to the specified number of beats in a measure. For example, if you want a region to be created at every beat, set *Beats* to 1 and *Measures* to 0. To have a region created at every measure, set *Measures* to 1 and *Beats* to 0. If you wanted to have a section at every 7 beats, you could set *Beats* to 7 and *Measures* to 0, or set *Beats* to 3 and *Measures* to 1, if you have defined a measure to be 4 beats. You can set the number of beats per measure using the *Edit Tempo* command.

When using the attack-detection algorithm, more beats will be detected if you raise the Attack sensitivity, raise the Release sensitivity, lower the Minimum level, or lower the Minimum beat duration. The opposite is true for detecting fewer beats.

NOTE

BATCH CONVERTER

See the *Batch Converter Reference* section for details.

CROSSFADE LOOP

The *Crossfade Loop* function is designed to help create loops in material that otherwise might not be loopable. It does so by mixing sound occurring before the loop start point into the end of the loop. This can make the transition from the end to the beginning of the loop much smoother.



You may also wish to crossfade the region following the loop (*Crossfade Post-Loop*) with the beginning of the loop to make the transition from looping to playing the post loop smoother. Often times this is done to compensate for a crossfade inside the loop.

*Although it depends on your application and personal preferences, most musicians agree that the *Crossfade Loop* function is a last resort. It is suggested that the *Loop Tuner* be used before applying the *Crossfade Loop* function to get the loop ends to match as well as possible.*

NOTE

GRAPH

The *Crossfade Loop* graph shows you where in the selection the crossfade will occur. The three areas (from left to right) which represent sound before, inside, and following the loop points, are shown with different sloped lines. For example, when doing a *Sustaining loop crossfade* with no *Post-Loop*, you can see the loop data fade out as data from the *Pre-Loop* is mixed in to form a crossfade.

LOOP TO CROSSFADE

Loop to crossfade selects the sample loop that will be affected by the crossfade. Most samples only have one loop which is the *Sustaining loop*. If a sample contains two loops, the second loop is called the *Release loop*.

LOOP

Loop determines how much of the selected sample loop will be crossfaded. When checked, the end of the sample loop is crossfaded with the end of the region before the loop start. This is used for making the end of the loop blend smoothly with the beginning of the loop.

POST-LOOP

When *Post-Loop* is checked, the region following the loop end is crossfaded with the beginning of the sample loop. This is used for making the transition from the end of the loop to the rest of the sample following the loop smoother (i.e. when done looping and playback continues).

Post-Loop crossfading is only necessary if the sampler is every going to play beyond the loop points.

PREVIEW MODE

Preview mode determine how the *Preview* button will work.

LOOP THE LOOP

Loop the Loop plays the loop area only.

PLAY LOOP THROUGH POST-LOOP

Play Loop through Post-Loop plays the loop once and then the *Post-Loop*.

PLAY AS ONE SHOT

Play as One Shot plays from the beginning of the file and uses the loop settings to determine how many times to loop each loop.

EXTRACT REGIONS

The Extract Regions tool is used to create individual files from regions in the Regions List of a sound file. This feature makes it easier for you to record long files, create regions in the Regions List, and then save those regions as individual sound files.

REGIONS TO EXTRACT

This list contains all of the regions present in the active sound file. To select the regions you wish to create individual files from, click on them. You can hold the Control key down when clicking to select multiple regions. If you hold the Shift key down while clicking, all of the regions between the last and current selection will be selected. The Select All button provides a convenient way to select all regions in the list.

DESTINATION FOLDER

This option determines the location where the extracted regions will be saved. If you enter a folder that does not exist, then the Extract Regions tool will create it. Selecting the Browse button provides a convenient way to select an existing destination folder.

BROWSE

This command displays a dialog allowing you to browse the available disk drives and directories for selecting an existing destination folder.

FILE NAME PREFIX

This option provides a text field for editing the file name prefix. If the Use Long File Names for destination file names option is not checked, then this prefix is truncated to the first 5 characters.

USE LONG FILE NAMES FOR DESTINATION FILE NAMES

When this option is turned on, the file names created can be up to 128 characters in length and may contain spaces. The files names will be the concatenation of the File name prefix and the region name. For example, if your prefix is set to *This is my file_* and there is a region named *Region 004* then the name used for extracting will be *This is my file_Region 004.WAV*.

When this option is turned off, the file names will conform to the 8.3 naming convention. These names consist of the first 5 characters from the File name prefix and a unique three digit number. For example, if you have 4 regions selected for extraction, and your prefix is set to *PREFIX*, then the names used will be *PREFI000.WAV*, *PREFI001.WAV*, *PREFI002.WAV*, *PREFI003.WAV*.

This option is only available in the 32-bit version. The 16-bit version will always use the 8.3 naming convention.

EXTRACT

Selecting the Extract button will start the region extraction process. All regions that are selected in the Regions List will be extracted to separate files.

PLAY

This command plays the region that has the focus in the list.

SELECT ALL

This command selects all of the regions in the list.

FIND



FIND

The Find tool searches for clicks and pops, volume levels, or silence breaks in a sound signal.

GLITCH

Clicks and pops are unwanted sounds often found in vinyl recordings or from bad splice-editing, and are seen as sharp glitches in the normal curve of a waveform.

The *Glitch* search algorithm looks at a waveform, starting at the cursor position, moving forward in time, and determines when a glitch occurs depending on the threshold and sensitivity settings. Then, the cursor position is moved to the detected glitch to allow for further editing. In *Pencil Tool* mode (Edit menu), you can draw a smooth curve over the glitch to remove it.

Threshold Slope

Threshold slope determines the minimum slope (steepness) of the glitch. A high setting will only detect glitches with a high slope, while a low setting will detect both low and high-sloped glitches.

Sensitivity

With a high *Sensitivity* setting, the glitch finder will detect as a glitch any part of the waveform with a slope above the *Threshold slope*.

A lower setting forces the *Glitch* finder to further verify that the high-sloped waveform is in fact a glitch and not part of the original, smooth waveform.

NOTE

*If no glitches are detected, but you can hear clicks or pops, try lowering the *Threshold slope* and/or raising the *Sensitivity*.*

*If the *Glitch* search algorithm is falsely detecting glitches, lower the *Sensitivity* and/or raise the *Threshold slope*.*

*To find the next glitch, you can use the *Control+Y* shortcut key to repeat the last operation.*

LEVEL EQUAL TO OR ABOVE

Use *Level Equal To or Above* to find the next point in time after the cursor position when the waveform reaches a particular level.

Threshold Level

Threshold level determines the sound level to search for.

END OF SILENT REGION

Use *End of Silent Region* to move the cursor to the end of the next silence break.

Threshold Level

Threshold level determines the silence sound level. The *Silence search* will move the cursor to the next position where the level goes below and then back above the silence threshold level.

LARGEST PEAK

Use *Largest Peak* to move the cursor to the next largest peak amplitude.

NOTE

Searches progress from the current cursor position forward, not the entire file. To find the largest peak in the entire file, set the cursor to the beginning of the sound file or select all before processing.

PRESET MANAGER

The Preset Manager is used to backup and transfer user-defined presets from all of the effects, processes, tools and Plug-Ins that are installed in Sound Forge.

PRESETS

This section contains a tree view explorer for all of the presets that exist in your Sound Forge installation.

Presets will be organized by menu and then by function. For example, if you want to access a preset in your DirectX Fungi Filter plug-in, click the plus sign next to DirectX, then Fungi Filter and then select the preset you wish to add.

PRESET PACKAGE CONTENTS

This section contains a list of all of the presets that you have added to the current preset package.

To add a preset, select it and press the *Add* button. You can add as many presets as you like. To remove a preset, select it in the *Preset package contents* list and press the *Delete* button.

OPEN

Pressing the *Open* button allows you to import the presets from an existing preset package file into your installation of Sound Forge.

SAVE AS

Pressing the *Save As* button allows you to save all of the presets in the *Preset package contents* list into a preset package file. This file can then be transported to another location or backed up for future use.

SETTINGS

Pressing the *Settings* button will open a temporary dialog for the selected function. From here, presets can be created on the fly to add to the preset package.

RUN BATCH SCRIPT

The *Run Batch Script* tool allows you to apply batch processing to opened sound files in *Sound Forge*. Scripts that are created with the *Batch Converter* are selected and applied to the current selection.

NAME

This list contains all of the scripts created using the *Batch Converter*, along with the default scripts included with the *Batch Converter Plug-In*.

BATCH SCRIPT

Displays the list of functions that will be applied, in their perspective order.

SAVE THE RESULT TO A NEW MEDIA FILE USING THE SPECIFIED FORMAT

Enabling this option saves the converted data in the format you have chosen in a new file. This option must be checked if you want the file format conversion step in the select *Batch Script* to be run. If you do not check this option, every part of the *Batch Script* except the format conversion will be run.

SAVE CURRENT SELECTION ONLY

Enabling this option saves only the current selection to a new file. In this way, files can be cropped without having to create a separate file of the selection first.

AUTOMATICALLY UNDO CHANGES TO SOURCE

Enabling this option will force the *Run Batch Scripts* function to undo the changes to the source after saving the new file.

SAMPLER



The *Sampler* dialog allows you to send and receive data to and from a sampler. To set up sampler/computer communications, press the *Configure* button.

CONFIGURATION

This *Configuration* drop-down list contains preset sampler setups (saved in the *Sampler Configuration* dialog). Once you save different configurations, the settings can be conveniently selected from this list.

LOGICAL SEND/RECEIVE SAMPLE NUMBER

The *Logical send/receive sample number* is the number your sampler uses as its location reference for samples sent or received. This number can be biased for specific samplers with the *Sample bias* option in the *Sampler Configuration* dialog.

ACTUAL SEND SAMPLE NUMBER

The *Actual send sample number* text field contains the actual sample number to be used by your sampler to store the sample. This number is based on the *Logical send/receive sample number* combined with the *Sample bias* for sending, set in the *Sampler Configuration* dialog.

ACTUAL RECEIVE SAMPLE NUMBER

The *Actual receive sample number* text field contains the actual sample number to be retrieved from your sampler into the current data window. This number is based on the *Logical send/receive sample number* combined with the *Sample bias* for receiving set in the *Sampler Configuration* dialog.

SEND SAMPLE

Press the *Send Sample* button to begin transmitting the active data windows sample to the sampler.

GET SAMPLE

Press the *Get Sample* button to begin receiving a sample from the sampler into the active data window. This will replace all data in the active window. A message box will be displayed requesting confirmation if there is already data in the active window.

CONFIGURE

Configure opens the Sampler Configuration dialog where you can determine the settings for communications with your samplers.

SAMPLER CONFIGURATION DIALOG

The *Sampler Configuration* dialog is used to create and save sampler communications configurations.

Name

The name of the sampler configuration that is saved as a preset.

Send To

Send to lists supported samplers and methods of communication for sending samples. If your sampler is not listed here, try *Generic MIDI/SDS* or *Generic SCSI/SMDI*.

MIDI Output (MIDI/SDS)

MIDI Output select the MIDI interface connected to the input of the sampler to which you wish to send samples.

MIDI Input (MIDI/SDS)

MIDI Input select the MIDI interface connected to the output of the sampler to which you wish to send samples.

Sample Bias

Sample bias is a numbering offset that allows you to add or subtract a value to the *Logical send/receive sample number* to compensate for differing sampler storage schemes. It is often easiest to set the *Sample bias* so that a *Logical send/receive sample number* of zero corresponds to the first available sample storage number in your sampler.

MIDI Channel (MIDI/SDS)

MIDI channel determines the MIDI channel the sample data will be sent on when using SDS.

Wait for Request (MIDI/SDS)

Selecting the *Wait for request* option tells the *Sampler Tool* to wait for a request (handshake) from the sampler before beginning an SDS sample transfer.

Open Loop (MIDI/SDS)

Selecting the *Open loop* option will cause the *Sampler Tool* to send SDS sample data immediately upon clicking the *Send Sample* button. This is an unconditional transfer of sample data (no handshake). When using the *Open loop* option to send, the samplers MIDI Out does not have to be connected to your computer's MIDI In. However, *Open loop* decreases sample dump performance.

Receive From

Receive from lists supported samplers and methods of communication for receiving samples. If your sampler is not listed here, try *Generic MIDI/SDS* or *Generic SCSI/SMDI*.

MIDI Output (MIDI/SDS)

MIDI output selects the MIDI interface connected to the input of the sampler from which you wish to receive samples.

MIDI Input (MIDI/SDS)

MIDI input select the MIDI interface connected to the output of the sampler from which you wish to receive samples.

Sample Bias

Sample bias is a numbering offset that allows you to add or subtract a value to the *Logical send/receive sample number* to compensate for differing sampler storage schemes. It is often easiest to set the *Sample bias* so that a *Logical send/receive sample number* of zero corresponds to the first available sample storage number in your sampler.

MIDI Channel (MIDI/SDS)

MIDI channel selects the MIDI channel on which the sample data will be received when using SDS.

Send Request (MIDI/SDS)

Selecting the *Send request* option tells the *Sampler Tool* to send a request to your sampler asking for a sample when you press the *Get Sample* button. If you turn this option off, you will always have to initiate transfers to the *Sampler Tool* from your sampler after pressing the *Get Sample* button. In other words, you will have to press a button on your sampler to start the actual transfer.

Open Loop (MIDI/SDS)

Selecting the *Open loop* option will cause the *Sampler Tool* to wait for sample data after pressing the *Get Sample* button. You must initiate all sample transfers from your sampler. This is an unconditional transfer of sample data (no handshake). When using the *Open loop* option to receive, the samplers MIDI In does not have to be connected to your computer's MIDI Out. However, *Open loop* decreases sample dump speed and is more prone to transfer errors.

SPECTRUM ANALYSIS

See the *Spectrum Analysis Reference* section for details.

STATISTICS

The *Statistics* tool displays statistical information about the selected sound file region.



CURSOR POSITION

Cursor position indicates the cursor position in samples from the start of the sound file.

SAMPLE VALUE AT CURSOR

The *Sample value at cursor* represents the instantaneous amplitude at this position.

For 16-bit sound files, the range of values is -32,768 to 32,767.

For 8-bit sound files, the range of values is -128 to 127.

NOTE

MAXIMUM/MINIMUM POSITION AND SAMPLE VALUE

Maximum/Minimum position and sample value is the maximum and minimum sample values and the location, in samples, where they occur.

This is useful for determining if any clipping occurs in the sound file. It can also be used to determine the noise level of a signal for use on the *Noise Gate* effect. For example, to find the noise amplitude, run the *Statistics* function on a region of noisy silence.

RMS POWER

RMS Power is the *Root Mean Square* of the sample values relative to the RMS value of a maximum-amplitude square wave (the loudest possible recording).

When used on short intervals this value relates to the volume level of the sound file. However, if used on a large selection with large volume variation, this value becomes less meaningful. For another way to measure loudness, use the *Scan Levels* button in the *Normalize* dialog.

The calculation used in *Sound Forge 4.5* is slightly different than in *Sound Forge 3.0*. The value of 0 dB corresponds to a maximum amplitude square, not sine wave. A maximum value sine wave will read at -3.0 dB.

AVERAGE VALUE (DC OFFSET)

The *Average* value is a sum of all of the sample values in the selected region divided by the number of samples. If non-zero, this usually indicates that there is a DC offset in the recording process.

ZERO CROSSINGS

Zero crossings displays the times per second that the waveform changes from a negative value to a positive value. This value can be used as a rough estimate of the frequency of the sound data for very simple waveforms.

SYNTHESIS

DTMF/MF TONES



Use *DTMF/MF Tones* to generate the standard dial tones used by the telephone companies.

DIAL STRING

Use the *Dial string* to enter the phone number to be generated along with any necessary pause characters. Unknown characters will be ignored.

AMPLITUDE

Amplitude controls the peak level of the waveform.

tone style to generate**DTMF**

Dual Tone Multi-Frequency signals are used by the standard push-button phones. It uses combinations of 679, 770, 852, 941, 1209, 1336, 1477, and 1633 Hz sine waves.

MF

CCITT R1 signals are used internally by the telephone networks. They are generated by a combination of 700, 900, 1100, 1300, 1500, and 1700 Hz sine waves.

SINGLE TONE LENGTH

Single tone length determines the output length of individual tones.

BREAK LENGTH

Break length determines the length of the silence between tones.

PAUSE LENGTH

Pause length determines the length of silence to insert when the pause character is found.

FADE THE EDGES OF EACH TONE

When *Fade the edges of each tone* is enabled, the tones are faded in and out a few milliseconds to prevent glitching.

PAUSE CHARACTER

Pause character determines the character which when present in the dial string will cause a pause between successive tones.

INSERT NEW TONE SEQUENCE AT

Insert new tone sequence at determines the point where the generated waveform will be placed within the existing sound file.

FREQUENCY MODULATION



FM Synthesis uses frequency modulation (FM) to create complex sounds from simple waveforms. Frequency modulation is the varying of the frequency of a waveform (the carrier) as a function of the output amplitude of another waveform (the modulator). If the frequency of the modulator is low, the carrier will be slowly detuned over time. However, if the frequency of the modulator is high, the carrier will be modulated so fast that many additional frequencies, or sidebands, are created.

In Sound Forge, up to four waveforms (operators) can be used in a variety of configurations. Depending on the configuration, an operator can be a carrier, a modulator, or a simple, unmodulated waveform.

TOTAL OUTPUT WAVEFORM LENGTH

Total output waveform length determines output length of the generated waveform.

CURRENT OPERATOR

The *Current operator* determines the active operator being modified.

ENVELOPE GRAPH

The *Envelope graph* displays an amplitude envelope of the active operator over time. Click the left mouse button to define a point and drag to change its position. Click the right mouse button or double-click the left mouse button over an existing point to remove the point. You can create up to 8 envelope points for each operator.

RESET

The *Reset envelope* button removes all points except the outermost two.

AMPLITUDE

Amplitude determines the output gain of operator after amplitude envelope. If the operator is a modulator, this control, along with the envelope, determines the amount of frequency modulation applied to the carrier. If the amplitude of a modulator is high, very harsh sounds can result.

OPERATOR SHAPE

Operator shape determines the shape of a single period of the waveform.

FREQUENCY

Frequency determines the frequency of the waveform. This parameter does not affect the noise waveform.

If the frequency is set to 0.00, a DC (zero frequency) waveform is produced regardless of the waveform specified.

When using *Noise*, the frequency determines the high frequency content of the noise. For full bandwidth (white) noise, set the frequency to its highest possible settings.

FEEDBACK

Feedback determines the amount of the output of the operator which is used to modulate itself. If the operator is also being modulated by another waveform, the feedback path and the modulator output are mixed together to modulate the carrier.

INSERT WAVEFORM AT

Insert waveform at determines the point where the generated waveform will be placed within the existing sound file.

CONFIGURATION

Configuration displays the arrangement and number of operators to be used.

The output of operators joined horizontally are simply mixed together. The output of the bottom operators are mixed together to form the final output. Mixing together different simple waveforms is called additive synthesis.

Two operators joined vertically are FM carrier-modulator pairs. Operators which have another operator stacked above them are carriers, with the operator directly above being the modulator.

Operators which have no other operator directly above are simple waveform generators.

When three or more operators are stacked, this means that the highest operator modulates the operator below it, which then in turn modulates the following operator below, and so on.

Start out with the simplest, single operator configuration and try out all the waveforms at different frequencies, with the feedback factor set at zero. Experiment using feedback on each waveform to hear the different effects of self-modulation. Move to the next configuration, with two unmodulated operators mixed together (horizontal connection) to hear two operators simultaneously (additive synthesis).

Change the configuration to two stacked operators (vertical connection) and again experiment using different waveforms and frequencies on the carrier (bottom operator) and the modulator (top operator). With the modulator frequency set low (1 to 5 Hz), lower the modulator amplitude to create slight detuning. Raising it creates big pitch bends.

NOTE

If the modulator frequency is set high, many unusual sounds can be achieved.

Setting the modulator frequency to 0.00 (DC) allows for direct control of the carrier's pitch by using the modulator's amplitude envelope.

Adding more operators increases the complexity of the waveform. Both FM and additive synthesis can be combined to generate an endless variety of sounds.

SIMPLE SYNTHESIS



Simple Synthesis can be used to generate a simple waveform of a given shape, pitch, and length.

AMPLITUDE

Amplitude determines the peak level of the waveform to be generated. When generating Noise, the amplitude is affected by the cutoff frequency selected.

WAVEFORM SHAPE

Waveform shape determines the shape of a single period of the waveform.

LENGTH

Length determines the output length of the generated waveform.

FREQUENCY

Frequency determines the frequency of the waveform.

When creating noise, the frequency specifies the low-pass cut-off frequency. For white noise, the frequency must be set to its maximum available setting.

Note that aliasing occurs with many of these waveforms when using high frequencies because they are not band-limited.

INSERT NEW WAVEFORM AT

Insert new waveform at determines the point where the generated waveform will be placed within the existing sound file.

THE DIRECTX MENU

The DirectX menu will appear whenever you have DirectX plug-ins installed on your computer. Additional DirectX menus will be created if the number of installed plug-ins becomes too great.

Consult the documentation of your plug-ins for information on their specific features

THE OPTIONS MENU

STATUS FORMAT

The current *Status Format* determines how Sound Forge displays position and length information.

SAMPLES

Use *Samples* to set the status format to samples.



TIME

Use *Time* to set the status format to show Hour:Minutes:Seconds:Milliseconds.



SECONDS

Use *Seconds* to set the status format to show seconds with six decimal place precision.

TIME & FRAMES

Use *Time & Frames* to set the status format to Hour:Minutes:Seconds:Frames, with the frame rate adjustable from the Edit Frame Rate dialog under the Special menu. In this mode, the hour, minutes, and seconds are in real time, independent of the frame rate.

ABSOLUTE FRAMES



Use *Absolute Frames* to set the status format to show frames and fractions of a frame.

MEASURES & BEATS



Use *Measures & Beats* to set the status format to show Measures:Beats.Quarter notes.

SMPTE NON-DROP



Use *SMPTE Non-Drop* to set the status format to SMPTE Non-Drop (30 or 29.97 fps). The standard rate for SMPTE Non-Drop when synchronizing to video is 29.97. However, in music applications, a rate of 30.00 is often used. You can enable using 30 instead of 29.97 under the Options menu on the MIDI/Sync page in the Preferences folder.

NOTE

When using a frame rate of 29.97, the time displayed will be different than real world time because thirty frames are still used to measure one second (see the Appendix section on SMPTE).

SMPTE DROP

Use *SMPTE Drop* to set the status format to *SMPTE Drop* (29.97 fps). Frames will be dropped at specific points in time to ensure that the hours and minutes displayed stay synchronized to real world time (see the Appendix section on SMPTE).

SMPTE EBU

Use *SMPTE EBU* to set the status format to *SMPTE EBU* (25 fps). This is the standard rate used by European television.

SMPTE FILM SYNC

Use *SMPTE Film Sync* to set the status format to *SMPTE Film Sync* (24 fps), which is a common slower film rate.

ANNOTATIONS

The names of markers, regions, and loops can be shown on the waveform display, along with dotted lines which run the entire height of the window...all with your permission, of course.

MARKER NAMES

Marker Names displays marker names on the waveform display.

MARKER LINES

Marker Lines displays dotted lines on the waveform display showing where markers exist.

REGION NAMES

Region Names displays region names on the waveform display.

REGION LINES

Region Lines displays dotted lines on the waveform display showing where regions exist.

LOOP NAMES

Loop Names displays loop names on the waveform display.

LOOP LINES

Loop Lines displays dotted lines on the waveform display showing where loops exist.

PLAY METERS

The following commands and options affect the *Play Meter* display. The *Play Meters* can be turned on from the *View* menu.

RESET CLIP

This *Reset Clip* function resets the clip indication on the *Play Meters*. After a clip occurs, the clip indication will remain lit until the *Reset Clip* function is invoked. You can also click on the clip light or levels to clear it.

RESOLUTION

Resolution settings display the minimum and maximum readings for the *Play Meters*. Changing the resolution of the meters is helpful for metering sound files of varying dynamism.

SHOW LABELS

The *Show Labels* option turns on or off the markings on the *Play Meters*.

HOLD PEAKS

The *Hold Peaks* option turns on or off the markings to indicate the highest peaks in the sound file.

HOLD VALLEYS

The *HoldValleys* option turns on or off the markings to indicate the lowest peaks in the sound file.

ALWAYS ON TOP

Turning the *Always on Top* option on forces the *Play Meters* to always remain on top even when another window is dragged over them.

TIME DISPLAY

These settings affect what the *Time Display* window will display.

POSITION

Checking the *Position* option will display the current cursor position in the time display window. The format of the display will depend on the *Status Format* setting above.

SYNC/TRIGGER STATUS

Checking the *Sync/Trigger Status* option will display the *Sync/Trigger Status* in the *Time Display* window.

PLAYLIST POSITION

Checking the *Playlist Position* option will display the playlist position in the *Time Display* window. The format of the display will depend on the *Status Format* setting above.

PASSIVE UPDATE

Checking the *Passive Update* option will reduce the overhead (processor time) needed to update the *Time Display* window. The smallest increments may not be exact but the major increments will stay perfectly locked. If you have a slower machine and are experiencing gaps in playback try checking this option.

VIDEO

COPY CURRENT FRAME

Enabling the *Copy Current Frame* option will copy the frame corresponding to the current cursor location to the clipboard.

WHITE, BLACK OR DEFAULT BACKGROUND

The *Background* option selects the color used surrounding frames in the *Video Preview* and video strips. The default color can be chosen from the in the *Options* menu on the *Preferences* page.

ANIMATE VIDEO STRIP

When *Animate Video Strip* is enabled, the video strip above the waveform display will come to life at the current cursor or play position. This feature is frame accurate, meaning that when you move the cursor, the correct frame corresponding to the cursor position will be displayed instead of the static overview frame.

When off, the video strip always shows the frame that corresponds to the left edge of each image in the video strip,

NUMBER FRAMES

Number Frames shows the frame number at the bottom left-hand corner if the video windows.

STRETCH TO WINDOW

Stretch to Window causes each *Video Preview* frame to fit exactly into the *Video Preview* window. This mode is generally the slowest because resizing a frame by non-integral fractions requires many calculations.

INTEGRAL STRETCH

When this *Integral Stretch* is enabled, the *Video Preview* frame will only be stretched by integral amounts. This is usually the fastest mode.

PRESERVE ASPECT

Check the *Preserve Aspect* re-size option when you want the original aspect ratio (horizontal vs. vertical size) of the video to be preserved.

PASSIVE UPDATE

Checking the *Passive Update* option will reduce the overhead needed to update the *Video Display* window. When on, the *Video Display* will only be updated when

the processor is idle, which in slower computers is much less frequent. This will help to keep the audio from skipping when a computer can't update video frames fast enough.

SCROLL PLAYBACK

Checking the *Scroll Playback* option makes the data window scroll automatically during playback. When the cursor moves off of the current window it will quickly scroll to show another full window of data.

Shortcut: F6

SCROLL SMOOTHLY

Checking the *Scroll Smoothly* option changes the behavior of the playback auto-scroll. Rather than jumping a window at a time, the cursor will slowly move back to the center and the wave data will scroll past it. This allows you to always see in front of the sound data that is being played.

Shortcut: Shift+F6

DRAG AND DROP SNAPPING

Drag and Drop Snapping toggles on and off the ability to snap to objects when dragging and dropping from window to window. With this on you can snap to major time ticks on the *Time Ruler*, markers and regions, .AVI video frames and cursor positions.

LOCK LOOP/REGION LENGTH

Lock Loop/Region Length forces the length of a region to remain constant when changing the start or end time of a region or loop.

Shortcut: Holding the *Control* key down while dragging *Ruler Tabs* also locks the length of a region or loop.

PASTE EVENTS

Enabling the *Paste Events* operation will transfer any events such as regions and markers from the source to the destination sound files when pasting, mixing, or crossfading. This occurs when using drag and drop or the clipboard.

If you are using another application to paste data into other than Sound Forge, these events may cause problems. Although a well-written application should ignore these events, if you are experiencing clicks at the beginning or end of the data you should disable this option.

SELECTION GRID LINES

When this option is enabled, all selections will be superimposed with grid lines that will split the selection into four equal parts. These lines make creating loops from existing material easier.

AUTO SNAP TO TIME



When *Auto Snap to Time* is enabled, any selection made will automatically be forced to the nearest whole time division.

AUTO SNAP TO ZERO



When *Auto Snap to Zero* is enabled, any selection made will automatically be forced to the nearest zero-crossing as defined by the *Editing* page of the *Preferences* folder.

Shortcut: *Control+B*

MIDI INPUT SYNC/TRIGGER



When *MIDI Input Sync/Trigger* is turned on, Sound Forge can be triggered from MIDI commands received through the MIDI Input port. The MIDI Input port is specified on the *MIDI/Sync* page in the *Preferences* folder (Options menu).

When this option is off, the MIDI triggers, Regions List triggers, and Playlist triggers specified will not be enabled. When this option is on, dialogs which specify MIDI triggers will also accept input from the MIDI Input port, allowing easy entry of MIDI Note and Controller values.

MIDI OUTPUT SYNC

Selecting the MIDI Output Sync item will enable Sound Forge to send MIDI time code out the MIDI Output port. The MIDI Output port is specified on the MIDI/Sync page of the Preferences folder (Options menu).



PRE-QUEUE FOR SMPTE/MTC

Selecting the Pre-Queue for SMPTE/MTC item will make Sound Forge open the wave device and pre-load data of the next region to be played in the Playlist. This is a good measure to take to insure that audio will begin playing the moment the designated SMPTE time is detected by Sound Forge when synchronizing to MTC. This keeps the time between triggering and playback to a minimum.



This option is automatically disabled when any other audio command is used such as Play, Stop, or Record.

NOTE

MIDI TRIGGERS

The MIDI Triggers option is used to control several Sound Forge functions using MIDI commands from external devices such as a MIDI Keyboard or sequencer. Triggers can be either MIDI Note-On/Off or Controller commands.

NAME

Name displays the name of the user-defined MIDI Trigger setup. You can save the current configuration using the Save as button.

EVENT

The Event window displays the available trigger operations in Sound Forge and, if enabled, triggering information. Selecting an event allows you to edit its trigger information.

EVENT

Event is the operation which will be performed by Sound Forge when a trigger is received.

TRIGGER

Trigger is the specific note or controller that will evoke the event it is assigned to.

CHANNEL

Channel is the MIDI channel the trigger is associated that will evoke the event.

TRIGGER

Trigger is where you enter the above information upon selecting an *Event* in the *Event* window.

NONE

Use *None* to select no MIDI trigger for the selected event.

NOTE

Note determines the selected event will be triggered by a MIDI note.

CONTROLLER

Controller determines the selected event will be a MIDI controller.

CHANNEL

Channel determines the channel to which the controller or trigger is assigned.

CONTROLLER/NOTE

Controller/Note determines the MIDI controller being assigned, if a controller is selected. The MIDI note if note is selected.

VALUE

Value determines the value the controller must reach in order to evoke the event.

ENABLE MIDI INPUT SYNC/TRIGGER (FOR TRAINING TRIGGERS)

If you have the *MIDI Input Sync/Trigger* mode enabled while in the *MIDI Triggers* dialog, you can auto fill the *Channel*, *Note*, and *Controller* values by pressing a key or controller on your *MIDI Keyboard*.

NOTE

MIDI Triggers are different from triggers in the Playlist and Regions List. When using triggers in the Playlist, Regions List, or MIDI Triggers, be aware that they can interact with each other to create unexpected results. Sound Forge first looks at the MIDI Triggers, then the Regions List, and then the Playlist when determining what to do when a MIDI command is detected.

PREFERENCES

Many user-configurable options are available in the Preferences dialog. To select an item, simply click on the folder tab or use the Tab key to place the focus on the folder name and use the right or left arrow key to navigate through the folders. Options available in the Preference folder include:

GENERAL

Use the *General page* to select general user options.

SHOW A TIP OF THE DAY ON STARTUP

When *Show a Tip of the Day on Startup* is enabled, a handy *Tip of the Day* will appear every time you run Sound Forge.

SHOW LOGO SPLASH-SCREEN ON STARTUP

When *Show logo splash-screen on startup* is enabled, a splash-screen will be displayed briefly upon startup.

SHOW BEAT MARKS IN DATA SELECTIONS

When *Show beat marks in data selections* is enabled, data selections will display beat marks.

ALWAYS ASK FOR REGION, MARKER, AND LOOP NAMES

When creating regions, markers, or loop points and this option is turned off, the *Edit Region* or *Loop* dialog will not automatically appear. A default name based on location will be used.

OPEN DEFAULT WORKSPACE ON STARTUP

If *Open Default Workspace on Startup* is enabled, Sound Forge will automatically open any files that were open when you last exited the program.

OPEN DRAG AND DROP SOUND FILES IN DIRECT MODE

Any files opened by dragging them onto the Sound Forge desktop (from outside of the program) will be opened in direct mode when this option is enabled.

OPEN COMMAND LINE SOUND FILES IN DIRECT MODE

When launching Sound Forge from the command line (DOS prompt), if this option is turned on any files in the command will be opened in direct mode. For example, if Sound Forge is being launched by another application (such as a multi-track editor), when this option is turned on Sound Forge will edit directly on the original file.

DEFAULT TO SLOW SCROLL WHEN DRAG SELECTING

In some very fast computers, automatic scrolling while selecting is too fast to use accurately. When this options is turned on, drag-selecting will cause a slow scroll.

N O T E

If you click the right mouse button while selecting, you can be toggle slow scroll on or off.

WARN WHEN MIX/CROSSFADE RATES MISMATCH

Enabling this option tells Sound Forge to warn you with a dialog prior to mixing or crossfading data which has different sample rates. Mixing or Crossfading data of different sample rates may produce unintended sounding results.

ALWAYS OPEN FULL-SCREEN

Enabling this option forces Sound Forge to always open in full-screen mode. If this option is disabled Sound Forge opens in the size and position it was in when it was last closed.

CONFIRM ON CLOSE

Enabling this option forces Sound Forge to prompt you with a confirmation message box prior to exiting Sound Forge.

If you are a system administrator you may want to enable this option for users who are new to Windows, and who may double-click on the close box by accident.

COMPATIBLE DRAW MODE (FOR BROKEN VIDEO DRIVERS)

Because Sound Forge's drawing routines are highly optimized, they increase the chance of causing little-known video card problems to arise. Some video cards have bugs in their drivers which can make your system lock up when Sound Forge tries to draw a waveform. If you experience a system lockup frequently in Sound Forge, it is likely that your video card is causing this problem.

Compatible draw mode uses a different method of drawing the waveform which, although not as smooth, puts less stress on the video card. With some video cards, this mode actually increases the draw speed. However, more flashing can occur.

This problem is known to happen with ATI's Mach64 card when using High Color (16-bit) in Windows 95.

COMPATIBLE SCROLL MODE

When this mode is enabled, Sound Forge will use a less optimized method of scrolling the Waveform Display. In some instances, this can reduce interference problems between your audio and video card during Smooth Scroll playback.

SHOW A TEXTURED BACKGROUND ON THE WORKSPACE

If this option is turned on, a stucco texture will be used for Sound Forge's application background.

IGNORE FACT CHUNK WHEN OPENING COMPRESSED WAV FILES

If this option is turned on, Sound Forge will ignore fact chunks in compressed WAV files. Compressed WAV files use fact chunks to specify how many actual samples are represented in the file. If a compressed file is improperly authored, this may cause some of the compressed data to not be loaded. If you suspect that not all sound data is being loaded from a compressed file, try checking this option and reopening the file.

SHOW ALL ACM CONVERT FORMATS IN SAVE AS

When enabled, Sound Forge will enumerate all compression formats available for each ACM compressor. However, this enumeration process can be slow on slow machines. For simple ACM codecs (Microsoft ADPCM, IMA ADPCM, etc.), this option is unnecessary (they have only one compression type per input) and could be disabled to speed up the selection

SHOW ACM DRIVER NAME IN SAVE AS

When enabled, Sound Forge will display the driver name next to each ACM codec. This is only necessary if you have more than one driver installed for a particular codec.

TURN OFF VOXWARE KEY FOR ASF

When enabled, you will need a software key supplied by Voxware to save ASF files using a Voxware codec. This is only necessary if you have a Voxware codec installed in your system with a higher version than the one included with NetShow 2.0.

SHOW FREE STORAGE SPACE ON STATUS BAR

When this options is turned on, the total amount of free disk space available on your specified temporary drive is displayed on the status bar.

SHOW SHUTTLE CONTROLS ON DATA WINDOW TRANSPORT

When enabled, Rewind and Forward buttons will appear on each Data Window's Transport.

AUTO-POWER MIDI KEYBOARD WINDOW

Enabling this option tells Sound Forge to open up the MIDI device assigned to the MIDI Keyboard (if it is not already open) when you press a key on the MIDI Keyboard. You may want to disable this option if you are using the same MIDI output device for MIDI synchronization or for your sequencer. If this option is disabled you need to press the On button on the Keyboard prior to using it to send notes.

IGNORE DEVICE CAPABILITIES IN MIXER WINDOW

Normally, when Sound Forge loads audio devices found in the system, it checks to see if the drivers support volume control. Those devices which report that they do not support volume control are grayed in the mixer devices list. However, some devices do not report their capabilities correctly. If this option is enabled Sound Forge will ignore the device capabilities and attempt to set the levels for any device you select in the Mixer window.

SHOW OLE DRAG SOURCE ON DATA WINDOW TRANSPORT

When checked, the *OLE Drag Source* control is displayed on the right side of the status bar in all data windows. This provides a convenient method of creating Sound Forge audio objects that can be embedded into other documents. If you do not use this feature and want to reclaim the space for the status bar, simply uncheck this option.

OFFSET VIDEO WHEN SAVING USING MPEG LAYER 3

Files saved with the MPEG Layer 3 codec have a slight amount of silence inserted at the beginning of the files. When this option is enabled, Sound Forge will insert video frames to compensate for this delay. This preserves audio and video sync.

OFFSET MARKERS/REGIONS WHEN SAVING USING MPEG LAYER 3

Files saved with the MPEG Layer 3 codec have a slight amount of silence inserted at the beginning of the files. When this option is enabled, Sound Forge will shift all markers and regions to compensate for this delay when saving.

ENABLE FORCED WRITE-THROUGH ON RECORD (WINDOWS 95 ONLY)

This option forces recording directly to the hard drive, eliminating large amounts of data to be cached. It is recommended that this option remain enabled for optimum recording performance.

DISPLAY

Use the *Display* page to set your display preferences.

NORMAL ZOOM RATIO

Normal zoom ratio specifies the default zoom magnification when loading a new sound file or using the *Zoom Normal* command. This allows you to control the maximum data that will be shown in a data window when a file is opened. Higher values show more data on initial file loading.

DEFAULT SOUND FILE WINDOW HEIGHT

Default sound file window height specifies the default data window height for a sound file on open.

DEFAULT VIDEO STRIP HEIGHT

Default video strip height specifies the default height of the video strip for an .AVI file on open.

COLOR PREFERENCE FOR

The color preferences section allows you set a custom color for a variety of graphics within Sound Forge. Each graphic has a default color or you can set a custom color by checking the Custom box and setting the component Red, Green, and Blue of each color.

For each mode of opening sound files, Normal, Direct, and Read Only, you can assign a different color for data window graphics.

Normal: Data Left

Specifies the color used to draw the left channel or mono wave data in each data window.

Normal: Data Left Selected

Specifies the color used to draw the left channel or mono selected wave data in each data window.

Normal: Data Right

Specifies the color used to draw the right channel wave data in each data window.

Normal: Data Right Selected

Specifies the color used to draw the right channel selected wave data in each data window.

Normal: Window Background

Specifies the color used to draw the background of each data window.

Normal: Selected Background

Specifies the color used to draw the background of selected data in each data window.

Normal: Center Line

Specifies the color used to draw the center line of the wave axis.

Normal: -6.0 dB Line

Specifies the color used to draw the line at half volume for the wave axis.

Normal: Play Cursor

Specifies the color of the cursor which is drawn while playing a file.

Direct: Data Left

Specifies the color used to draw the left channel or mono wave data in each data window.

Direct: Data Left Selected

Specifies the color used to draw the left channel or mono selected wave data in each data window.

Direct: Data Right

Specifies the color used to draw the right channel wave data in each data window.

Direct: Data Right Selected

Specifies the color used to draw the right channel selected wave data in each data window.

Direct: Window Background

Specifies the color used to draw the background of each data window.

Direct: Selected Background

Specifies the color used to draw the background of selected data in each data window.

Direct: Center Line

Specifies the color used to draw the center line of the wave axis.

Direct: -6.0 dB Line

Specifies the color used to draw the line at half volume for the wave axis.

Direct: Play Cursor

Specifies the color of the cursor which is drawn while playing a file.

Read Only: Data Left

Specifies the color used to draw the left-channel or mono wave data in each data window.

Read Only: Data Left Selected

Specifies the color used to draw the left channel or mono selected wave data in each data window.

Read Only: Data Right

Specifies the color used to draw the right channel wave data in each data window.

Read Only: Data Right Selected

Specifies the color used to draw the right channel selected wave data in each data window.

Read Only: Window Background

Specifies the color used to draw the background of each data window.

Read Only: Selected Background

Specifies the color used to draw the background of selected data in each data window.

Read Only: Center Line

Specifies the color used to draw the center line of the wave axis.

Read Only: -6.0 dB Line

Specifies the color used to draw the line at half volume for the wave axis.

Read Only: Play Cursor

Specifies the color of the cursor which is drawn while playing a file.

Regions: Marker

Specifies the color used to display marker tabs on the data window Time Ruler.

Regions: Section

Specifies the color used to display region tabs on the data window Time Ruler.

Sample Loop: Sustaining

Specifies the color used to display sustaining loop tabs on the data window Time Ruler.

Sample Loop: Release

Specifies the color used to display release loop tabs on the data window Time Ruler.

Record Prepared

Specifies the color used in the *Record* dialog to indicate Record Prepared status.

Record Pre-roll

Specifies the color used in the *Record* dialog to indicate record pre-roll status.

Record Recording

Specifies the color used in the *Record* dialog to indicate recording status.

Graph Axis

Specifies the color used when displaying an axis in a graph such as in *Paragraphic EQ*.

Graph Fill

Specifies the color used when displaying a filled graph such as in *Paragraphic EQ*.

Graph Fill Alternate

Specifies the color used when displaying filled graphs such as in *Multi-Band Dynamics*.

Envelope Axis

Specifies the color used when displaying the axis of envelope controls.

Envelope Highlight

Specifies the color used when displaying a highlight in envelope controls.

EDITING

Use the *Editing* page to set preferences related to editing and undo operations.

DISABLE TRIPLE-CLICKING TO SELECT ALL DATA

Turn this option on if you don't want Sound Forge to select all sound data when you triple-click. You might want to turn this off if triple-clicks are constantly being falsely detected when you make a selection and then try to do a drag operation. Otherwise, lowering Window's double-click threshold time might help.

SNAP TO ZERO-CROSSING SLOPE

Specifies how zero-crossings are detected.

Negative Slope

Zero-crossings are detected only on a negative slope.

Any Crossing

Zero-crossings are detected on both positive and negative slopes.

Positive Slope

Zero-crossings are detected only on a positive slope.

NOTE

Tip: It is usually best to use either the Positive or Negative Slope setting for zero-crossings, so that noticeable pops and clicks are not generated by cutting data when using Snap to Zero.

ZERO CROSS SCAN TIME (0 TO 1000 SAMPLES)

Determines the maximum time that Sound Forge will search for the next zero-crossing.

ZERO CROSS LEVEL THRESHOLD (0 TO 32,767)

Levels below this value will be considered a zero-crossing. This is useful if the file you are using has a DC offset, although a better idea would be to remove the DC offset.

DISABLE AUTO-SNAPPING BELOW 1:4 ZOOM RATIOS

When this option is turned on, auto-snapping will not occur when you are zoomed in at 1:1 or 1:2. This is useful if you commonly zoom in full to tweak selection points manually yet still want to use auto-snapping when zoomed out.

LIMIT NUMBER OF UNDO BUFFERS TO (0 TO 999)

If this is checked, it determines the maximum number of *Undo* operations that will be saved before they start to get deleted. Having many *Undo* operations available will use a lot of hard-disk space.

LIMIT NUMBER OF REDO BUFFERS TO (0 TO 999)

If this is checked, it determines the maximum number of *Redo* operations that will be saved before they start to get deleted. Having many *Redo* operations available will use a lot of hard-disk space.

DEFAULT ALL WINDOWS TO UNDO DISABLED

Enabling this option will keep Sound Forge from creating undo buffers for any operation which modifies a sound file. You are still able to turn the undo creation on by using the *Disable Undo* command (Edit menu) for any window, but by default no undo buffers will be created.

DRAG & DROP AUTO RISE DELAY (0.1 TO 2.0 SECONDS)

Determines the time necessary before a window underneath the cursor becomes active during drag and drop operations.

FILE

Use the *File* page to set preferences related to opening and saving files.

WARN IF SAVE CHANGES THE SOUND DATA

When you save a sound file to a different format, such as doing 16 to 8-bit conversions or compressing audio, you will not hear the conversion's effect on the sound until you re-open the file. If this option is enabled, you will get a message warning of this.

AUTOMATICALLY RE-OPEN IF SAVE CHANGES THE SOUND DATA

If this option is on, when you save a file to a different format Sound Forge will save the file and then re-open it. This will happen when you change bit size, channels, or compression format and it will allow you to listen to any changes in sound quality.

SECTOR ALIGN DATA FOR DIGIDESIGN'S SESSION 8 .WAV FILES

This option is only necessary if you are using Sound Forge to modify files for use with Digidesign's Session 8. Enabling this option forces Sound Forge to align the data portion of .WAV files on a sector boundary, which is needed when using .WAV files with the Session 8.

SAVE .VOC FILES USING NEWEST FORMAT (16-BIT CAPABLE)

Prior to Creative Labs' boards supporting the playback of 16-bit data, VOC files had no way of storing 16-bit audio data. With the addition of 16-bit playback sound boards to the Sound Blaster line of products, a newer version of the VOC file was created. With this option disabled Sound Forge will save all 8-bit files in the old format and all 16-bit files in the new format. With this option enabled it will save all files in the new format. Only enable this option if you have software which requires the new format for 8-bit data.

OPEN AND SAVE WITH SOUND FORGE 3.0 COMPATIBLE REGIONS

Sound Forge 3.0 had a minor bug when opening and saving regions in which the total length was off by one sample. Since we both opened and saved the files inconsistently with the published standard, it didn't affect anything. Sound Forge 4.0 fixes this problem, but in doing so this means that file created with Sound Forge 3.0 will have their region length display 1 sample smaller than before.

If this bothers you, turn this option on when opening a file saved in Sound Forge 3.0, turn it off, and then save the file. From there on after, the region lengths will not be off by one in Sound Forge 4.0 (but will be in 3.0!)

DEFAULT SDS CHANNEL WHEN SAVING

Specifies the channel used in SDS packets when saving audio as an SDS file.

DEFAULT SDS PATCH WHEN SAVING

Specifies the patch used in SDS packets when saving audio as an SDS file.

DEFAULT DIALOGIC .VOX SAMPLE RATE FOR OPEN

Specifies the default sample rate used when opening Dialogic .VOX sound files in Sound Forge. This option is necessary because Dialogic .VOX sound files do not contain this information. Typical settings for this option are 6,000 Hz and 8,000 Hz.

ASSOCIATE SOUND FILE EXTENSIONS

Pressing the *Associate Sound File Extensions* button from the *File* tab in the *Preferences* folder opens the *Sound File Associations* dialog. This dialog provides a method of associating audio files with Sound Forge. When a file is associated with Sound Forge, you can simply double-click on a sound file in the *File Manager* or the *Windows 95 Explorer* and it will open for editing.

File Types to Associate

The list displayed contains all sound file formats that Sound Forge supports. If a sound file extension is associated with an application, then the name of the application is displayed on the right. If a sound file is associated with Sound Forge, then a check mark is placed in the box to the left of the item. To associate an item with Sound Forge, simply check the item and select OK.

File Association Details

When a sound file extension is associated with Sound Forge, Windows will open Sound Forge to edit any files with the extension when selected. For example, if you double-click on a file with an extension that is associated with Sound Forge in the *File Manager* or the *Windows 95 Explorer*, Sound Forge will automatically run and the sound file will be opened for editing.

Select All

Pushing this button will select all sound files to be associated with Sound Forge.

Select None

Pushing this button will de-select all sound files so no files are associated with Sound Forge.

MIDI/SYNC

Use the *MIDI/Sync* page to select MIDI and synchronization user options.

OUTPUT

Output specifies the MIDI output device for synchronization.

This device is used to send SMPTE/MTC when *MIDI Output Sync* (Options menu) option is enabled.

INPUT

Input specifies the MIDI input device for synchronization and triggering when the *MIDI Input Sync/Trigger* (Options menu) is enabled.

This device is the device through which Sound Forge will receive all MIDI triggering and synchronization input. This includes *SMPTE/MTC*, *MIDI triggers*, and *Regions/Playlist triggers*.

USE 30 FRAMES PER SECOND FOR SMPTE NON-DROP

Specifies whether all *SMPTE Non-Drop* calculations performed by Sound Forge use a frame rate of 30 frames per second or 29.97 frames per second. You will only need to uncheck this option if you intend to work with video which expects true 29.97 fps *SMPTE Non-Drop*. Refer to the Appendix for a discussion of SMPTE issues.

BOUND RECORD TIME ON SMPTE RECORD SYNC

When this option is enabled, Sound Forge will not allow recording beyond the specified end time. This ensures that your record length is exact regardless of any inaccurate time code.

USE INTERNAL TIMER FOR SMPTE GENERATION

Specifies whether Sound Forge should use the internal timer for SMPTE generation rather than position values reported by the sound card driver. Since many sound cards do not report their position accurately it is usually better to use the internal timer for SMPTE generation.

INTERNAL TIMER RESOLUTION

Specifies the timer accuracy used for generating SMPTE when the *Use Internal Timer for SMPTE generation* option is enabled. The lower the resolution value, the more accurate the SMPTE generation will be. However lower values may also decrease system performance.

USE FREE-WHEEL FOR SMPTE LOSS

Enabling this option makes Sound Forge stop playback if the incoming MIDI time code signal stops being received. Otherwise, when the time code stops, Sound Forge will continue to play until the user stops playback manually.

Free-wheel Time

When the time code stops being received, Sound Forge will continue to play for this amount of time before stopping. During this time, if time code is found again, playback will continue.

Free-wheel Slack

This parameter controls how fast Sound Forge should expect time code updates before going into *Free-wheel* mode. If you have a fast machine, this value can be set to a lower value if you want Sound Forge to stop playback immediately when time code is interrupted.

ENABLE SMPTE PLAYBACK OFFSET

When outputting MIDI time code during playback, you can specify a SMPTE offset that will be added to the time displayed in Sound Forge's play counter. For example, if you want Sound Forge to generate MIDI time code starting at 01:00:00:00, instead of inserting 1 hour of silence to your sound file, you can specify that amount in this field.

When using Record Sync, you'll often want to set this value to the Input Synchronization start time. Sound Forge's ruler and play counter will not display this offset.

PERFORMANCE

Use the *Performance* page to set preferences related to opening and saving files.

TEMPORARY STORAGE FOLDER

The specified directory is used to store temporary files used by Sound Forge.

Sound Forge stores its temporary data in a temporary directory on your hard drive. Using temporary file space allows you to edit very large files and keeps Sound Forge from eating up large portions of RAM on your computer. Your temporary directory must have enough space to handle the total size of all files you plan to edit at one time along with space for any clipboard data and undo buffers.

If you change the *Temporary storage* folder, you will have to close Sound Forge for the change to take effect.

FREE STORAGE SPACE IN SELECTED FOLDER

Shows the amount of available disk space in the selected folder.

PEAK RATIO DEFAULT FOR NEW SOUND FILES

When drawing the waveform display, Sound Forge will use a peak file instead of scanning all of the data when the zoom ratio is above this level. Lower values allow for faster scrolling at those ratios, but increases the peak file's size. For example, setting this value to 1:128 instead of 1:256 will make the peak file two times larger but will make scrolling at 1:128 faster and scrolling at 1:65536 or higher ratios slower.

To calculate the size of the resulting peak files, you divide the size of the file by the *Peak ratio*. For example, a 100 MB sound file will need a 0.39 MB (100/256) peak file when using 1:256.

WAIT CURSOR ZOOM RATIO THRESHOLD

Specifies the amount of zoom above which an hourglass cursor will be displayed between data window repaints. If the mouse cursor seems to flicker on and off a lot when redrawing wave data, you may want to increase this value. This value is ignored during playback.

SHOW THE POSITION OF THE PLAYBACK CURSOR

When this option is enabled, the *Position* field in the status bar will show the cursor position during playback. Disable this option if you have a very slow computer or video card.

SHOW THE RECORD COUNTER WHILE RECORDING

When this option is enabled, the record time will show in the *Record* dialog while recording. Disable this option if you have a very slow computer or video card.

PLAYLIST

Use the *Playlist* page to select user options dealing with the *Regions List* and *Playlist*.

REGIONS LIST DISPLAY FORMAT

Specifies how the regions are displayed within the *Regions List*.

Start and End

Regions are displayed with the start and end time of each region.

Start and Length

Regions are displayed with the start time and length of each region.

PLAYLIST DISPLAY FORMAT

Specifies how the regions are displayed within the *Playlist*.

Start Only

Regions are displayed with only the start time of each region.

Start and End

Regions are displayed with the start and end time of each region.

Start and Length

Regions are displayed with the start time and length of each region.

SORT THE REGIONS LIST ALPHABETICALLY

Specifies whether regions are displayed in the *Regions List* in the order they were added or are sorted alphabetically.

PLAYLIST PRE-ROLL (0.0 TO 300.0 SECONDS)

When playing entries in the *Playlist*, this value specifies the amount of pre-roll which will be used. This allows you to easily hear the transition from one region to another without having to play all the way through the first region.

PREVIEWS

Preview preferences apply to all processing, effects, and tools which have the ability to play a preview. You can also modify preview options within any dialog which has preview by clicking on the system box of the dialog and selecting the *Preview Configuration* option.

LIMIT PREVIEWS TO

Specifies the length of audio which will be used when generating a preview. Keeping this value small lowers the amount of time needed to generate a preview when tuning effects or process values.

PRE-ROLL

When turned on, this determines how many seconds of unprocessed audio will be played before the processed selection. Use this to listen to the transition from unprocessed to processed data.

POST-ROLL

When turned on, this determines how many seconds of unprocessed audio will be played after the processed selection. Use this to listen to the transition from processed to unprocessed data.

FADE OUT LAST 10 MILLISECONDS (SNAP/POP ELIMINATOR)

Enabling this option tells *Sound Forge* to fade out the last 10 milliseconds of a preview effect so that large pops are not heard at the end of a preview buffer. With this option disabled, a preview could end with a large data value which can cause an audible and annoying pop.

LOOP PREVIEW CONTINUOUSLY

Enabling this option causes preview buffers to loop infinitely rather than playing just a single time. This can be useful when listening to the difference between the original sound and the sound after an effect or process has been applied.

REACTIVE PREVIEWING

Reactive Previewing is a preview mode in which any control changes in a dialog cause an automatic re-calculation and playback of the preview buffer. This allows for immediate aural feedback on the affects of a change.

This option is very useful when using a fast computer, limiting previews to a short time, and not using very slow processing effects. You can temporarily disable Reactive Previewing in a dialog by holding down the Shift key while making parameter changes.

STATUS

Use the Status page select display status options.

DEFAULT SMPTE FORMAT

Specifies the SMPTE type used when selecting a SMPTE format option from the Status Bar shortcut menu.

Non-Drop 30 FPS

Calculates status values using a 30 frames per second SMPTE code.

Drop 29.97 FPS

Calculates status values using a 29.97 frames per second SMPTE code. This is accomplished by dropping two frames every minute except for minutes 00, 10, 20, 30, 40, and 50.

EBU 25 FPS

Calculates status values using a 25 frames per second SMPTE code. This format is the European Broadcast Union time code.

Film Sync 24 FPS

Calculates status values using a 24 frames per second SMPTE code. This format is for standard 24 frames per second film media.

DEFAULT TIME FORMAT

Specifies the time format (time or seconds) used when selecting the Time option from the Status Bar shortcut menu.

DEFAULT FRAMES PER SECOND

Specifies the default frame rate used to calculate frame values.

Frame values are useful when trying to synchronize sound with animation. Most animation players specify a playback frame rate at which video frames are shown to the user. If you are using an animation that has a frame rate of 15.0 frames per second, you would set the frame rate to 15.0. When status values are shown by Sound Forge they will be shown in values of frames. This allows you to find the frame to which a given point in the sound file corresponds.

Frame rates can be any value between 1.0 and 1,000.0 frames per second.

DEFAULT BEATS PER MEASURE

Specifies the number of beats in each measure for displaying in measures and beats. For example 2/4 time would have two beats per measure.

DEFAULT BEATS PER MINUTE

Specifies the number of beats per minute, i.e. the tempo of a song for displaying lengths. For example, a tempo of 120 is 120 beats per minute.

RMS LEVEL SCAN TIME

Specifies the amount of sound data surrounding the cursor used to calculate the RMS level in the Levels toolbar.

PEAK LEVEL SCAN TIME

Specified the amount of sound data surrounding the cursor used when searching for a peak level to display in the Levels toolbar.

DISPLAY FORMAT FOR SMPTE CODE

Specifies how SMPTE values will be formatted for display by Sound Forge. Select one of the available preferences from the drop-down list.

TOOLBARS

Use to select which toolbars are displayed on the main screen. A check before the toolbar means that the toolbar is turned on. The following toolbars are available:

STANDARD

Contains standard file, and edit operations.

TRANSPORT

Contains record and play operations.

NAVIGATION

Contains operations for navigating within the data window. Some of the operations included are zooming, marking, moving the cursor, and moving to relevant points.

VIEWS

Contains operations for storing and retrieving data window views.

STATUS/SELECTION

Contains operations to set the display Status Format as well as the Snap To operations.

REGIONS/PLAYLIST

Contains the Regions/Playlist buttons as well as synchronization commands and synchronization status boxes.

PROCESS

Contains all processing operations.

EFFECTS

Contains all effect operations.

TOOLS

Contains all tool operations.

LEVELS

Displays the left and right channel levels at the cursor in a variety of different formats.

SHOW TOOLTIPS

The Show Tooltips check box turns on or off the help information box displayed next to the mouse when it is held over certain items.

WAVE

Use the Wave page to specify what wave devices Sound Forge should use for recording and playback.

NOTE

The capabilities of these devices may limit you from playing some files.

PLAYBACK

Specifies the device to be used when playing sound data.

Selecting a mapper device such as the *Wave Mapper* or *Microsoft Sound Mapper* allows Windows to select an appropriate device to use for the current sound data.

RECORD

Specifies the device to be used when recording sound data.

Selecting a mapper device such as the *Wave Mapper* or *Sound Mapper* allows Windows to select an appropriate device to use for the current sound data.

INTERPOLATE POSITION FOR INACCURATE DEVICES

Specifies whether Sound Forge should interpolate the position of the play position pointer during playback. If this option is not checked, Sound Forge relies on the device driver for the sound card to provide the correct position of the pointer during playback.

Many sound card drivers do not report their position accurately. If the play position pointer does not move smoothly you should enable this option. You may also want to contact your sound card manufacturer for updated drivers.

If your sound card driver does not have accurate position reporting, your *Drop Marker* and *Mark In/Out* positions will also be off when dropping markers while playing.

POSITION BIAS

Specifies an offset value which Sound Forge should add to the position values returned by the sound card wave driver during recording or playback.

Many sound card drivers do not report their position accurately. When you play sounds back, the position of the play pointer may seem to lag or precede the actual sound you are hearing. If this is the case, you may want to adjust this value. Most sound card drivers that report incorrect position give a value ahead of the actual sound.

To remove this bias, start with smaller values and work toward larger ones in increments of 4 (a value of 1 is equal to 512 bytes) until the play position pointer is accurate. Recommended settings are 4, 8, 16, and 32. These settings correspond to the typical bias values for these drivers of -2048, -4096, -8192 and -16384 (negative multiples of 2048) bytes.

If your sound card driver does not have accurate position reporting, your Drop Marker and Mark In/Out positions will also be off when dropping markers while playing.

N O T E

An inaccurate playback cursor is due to poorly written sound card drivers. If you experience this problem, call your sound card manufacturer and request that they get their act together.

TOTAL BUFFER SIZE (KILOBYTES)

Select the amount of RAM you would like to use for buffering sounds during play and record. Using a large value can optimize performance at high sample rates but may also cause a slight delay starting playback.

PRELOAD SIZE (KILOBYTES)

Determines the total amount of wave data that will be queued to the sound driver before playback starts.

VIDEO

Use the *Video* page to select options relating to Sound Forge's display of video.

VIDEO DECOMPRESSION MODE

Determines the bits per pixel (color resolution) which will be used by Sound Forge when it decompresses video frames before displaying. Note that when a high number of bits is used, the performance demands on your system are higher. The available modes are:

8 bits always

Sound Forge will always decompress the video to 8 bits per pixel (256 color).

Best for display

Sound Forge will decompress the video to the best format that matches your system's color capabilities.

Full decompress always

Sound Forge will always decompress the video to its original bits per pixel when displaying. This is useful if you want to ensure that the video frames you see contain all the color information stored in the file. Also, if you wanted to paste frames to the clipboard, this ensures maximum resolution.

ALLOW HARDWARE SCALING DURING DECOMPRESSION

When this option is selected the frames of the video can be scaled to a different size.

Only scale frames larger than

These controls specify the minimum size of frame that should be scaled.

Attempt to scale during decompression to

These controls specify the new, scaled size.

DISABLE SCAN FOR EXTERNAL MONITOR SUPPORT

When this option is enabled, Sound Forge will not look for an external monitor to display video on every time an AVI file is opened. If you do not have an external monitor, enabling this option will save some time.

OTHER

Use the Other page to select miscellaneous user options.

PLAY LOOPED ADJUST PRE-ROLL

Specifies the number of seconds that will be played before the loop end point when playing in *Loop Preview* mode. To play a loop in *Loop Preview* mode, control-click on the *Play Looped* button. This is useful when loop-tuning a long loop.

AUDIO EVENT LOCATOR

The following settings control how the *Audio Event Locator* plays audio. To use the *Audio Event Locator*, click on the overview (above the waveform display) and drag the mouse. Playback is looped around the position until the mouse button is released.

Pre-roll

Specifies the location where playback will begin prior to where the cursor is placed in the overview.

Loop Time

Determines the amount of time which will be looped during playback when using the *Audio Event Locator*.

CUT PREVIEW CONFIGURATION

When using the *Preview Cut* function in the *Edit* menu, the *Pre-roll* and *Post-roll* times specify the amount of audio data that is played prior to and after the data which is to be cut.

Pre-roll

Specifies the amount of data played prior to a cut region.

Post-roll

Specifies the amount of data played after a cut region.

THE WINDOW MENU

NEW WINDOW

Use this command to create a new data window which will be given a generic name. This command uses the sound parameters that were set when the *New* command in the *File* menu was last used.

Shortcut: Control+Shift+N

CASCADE

Use this command to arrange the data windows so they overlap with the title bar of each window remaining visible.

Shortcut: Shift+F5

TILE HORIZONTALLY

Use this command to arrange the data windows side by side without them overlapping.

TILE VERTICALLY

Use this command to arrange the data windows top to bottom without them overlapping.

Shortcut: Shift+F4

ARRANGE ICONS

Use this command to arrange the minimized data icons in the Forge workspace.

MINIMIZE ALL

Use this command to minimize all open data windows in the Forge workspace.

RESTORE ALL

Use this command to restore all minimized windows to their previous window size and position.

CLOSE ALL

Use this command to close all open data windows in the Forge workspace.

All open files are listed at the bottom of the Window menu. The active file is indicated by a check mark.

NOTE

THE HELP MENU

CONTENTS

Contents takes you to the online help table of contents.

Shortcut: F1

SEARCH FOR HELP ON

Use this to find help on a particular topic with the online help.

KEYBOARD SHORTCUTS

Keyboard Shortcuts displays a list of shortcuts – and Sound Forge is loaded with them.

TROUBLESHOOTING

Use this before calling the Tech Support line.

TIP OF THE DAY

This command will bring up the *Tip of the Day* pages in the help file. These are helpful and handy tips to increase your productivity with Sound Forge.

THE SONIC FOUNDRY HOME PAGE

This is a direct link to our home page on the World Wide Web. Invoking this command automatically launches your default web browser and takes you to the Sonic Foundry Home Page (<http://www.sonicfoundry.com>). Here you will find downloadable demos and updates, technical information, and news of what's new with the folks at Sonic Foundry.

You can also send us comments, suggestions, questions, sound files, gossip, whatever...via email from our web page or at feedback@sonicfoundry.com. We'd love to hear how Sound Forge is being used and your thoughts on how it can be improved.

ABOUT SOUND FORGE

Use this to display pleasantly puzzling information about Sound Forge, such as the software license owner, copyright and system information, program version and serial number, and the cool Sound Forge logo.

BATCH CONVERTER REFERENCE

FILE MENU

ADD

This command opens the *File Add* dialog for selecting the files to process. Files can also be added to the batch list by dragging files from the Windows Explorer and dropping them onto the Batch Converter window.

The File Add dialog allows one or more files to be added to the batch list. Select the files to add and click the Add button. The selected files will then be appended to the batch list as source files.

OPEN SOURCE

This command opens the selected source (un-converted) file in Sound Forge.

OPEN DESTINATION

This command opens the selected destination (converted) file in Sound Forge. The destination file can only be opened if a successful conversion has completed and the Status of the file is Done.

LIST ERRORS

This command displays a list of errors that were encountered during a conversion for the selected file. If no errors occurred, then this option is not available. Use this option to determine why the file was not converted.

PROPERTIES

This command displays various properties for the selected file.

BATCH SCRIPT

ADD BATCH SCRIPT

This command imports a batch script that was packaged by using the *Save Batch Script* command. It will install the selected file into the Batch Converter along with any new plug-in presets that may be required by the new script.

SAVE BATCH SCRIPT

This command allows a selected script to be packaged for transferring to different computers. The function creates a selected file with all the information necessary to install the selected script for use on another computer including user created presets for any plug-in. The file can then be passed to another user and installed using *Add Batch Script*.

CLOSE

Closes the Batch Converter window.

EDIT MENU

INVERT SELECTION

This command inverts the current selection. All files that are selected (highlighted) will be unselected; all files that are not selected will be selected.

SELECT DONE

This command selects all files that have a status of *Done*. Files are labeled *Done* after they have been successfully converted.

SELECT ALL

This command selects all files in the list.

REMOVE SELECTED

This command removes all files that are selected (highlighted) from the list.

REMOVE DONE

This command removes all files that have a status of *Done* from the list. Files are labeled *Done* after they have been successfully converted.

REMOVE ALL

This command removes all files from the list.

RESET SELECTED

This command resets the status of all selected files. After resetting a file's status, another conversion can be run on the file.

RESET ALL

This command resets the status of all files in the list.

OPTIONS MENU

CONVERSION SETTINGS

This command opens the *Conversion Settings* dialog. This dialog is used to specify the destination file format and audio processing functions that will be applied to the selected files.

FILE FORMAT

This field determines the output file type. If this option is turned off, then the source file format is preserved for the destination sound file. When the File Format is of Raw, RealAudio, RealMedia, Advanced Streaming Format, or Video for Windows, the *More* button provides control over the destination binary format.

DESTINATION RAW ATTRIBUTES

This dialog is reached when you select *Raw* from the File format list in the *Conversion Settings* dialog and press the *More* button. Use this dialog to save files into a file format that is not one of the standard sound file formats supported by Sound Forge.

Name

This list of file formats contains preset versions of raw file storage formats commonly used by other applications. You can create a new file format by using the *Save As* button to save the current storage parameters.

Sample Type

This option determines the format used to store each sample.

Format

This option determines the format of each sample value.

Channels

This option determines the number of channels stored in the file.

Byte Order

This option determines the order in which the high and low bytes of a 16-bit sample are stored. Little endian (low, high) is used by Intel microprocessors, while Big endian (high, low) is used by Motorola microprocessors.

Save As

Use this command to save the current parameters as a preset.

Delete

Use this command to delete a the current preset. The built-in presets can not be deleted.

REALAUDIO SETTINGS

This dialog is reached when you select *RealAudio (.ra)* from the *File Format* list in the *Conversion Settings* dialog and press the *More* button.

Desired Compression Type

This option determines the compression type used and the data rate required to play the file. If more than one type is used, multiple files are generated and stored in a RealAudio Bandwidth Negotiation-compatible folder.

Enable Selective Record Capabilities

When checked, RealAudio compatible players can allow users to save (record) RealAudio streams to their hard drives. Use this option only if it is acceptable for users to create copies of the RealAudio file. The RealAudio Player Plus from RealNetworks supports Selective Record capabilities.

Enable PerfectPlay Capabilities

When this option is enabled, users connected with standard modems can select to download a higher quality version or the audio that is normally only available over ISDN or LAN connections. The higher bandwidth RealAudio stream is then played back from the user's hard disk after the download process is complete. The RealAudio Player Plus from RealNetworks supports PerfectPlay capabilities.

Create event file (RAE) from region and marker labels

When this option is enabled, events can be embedded from the Regions List into the RealAudio stream. To embed an event, create a marker or region with a label starting with the following prefixes followed by a colon:

- RA URL: Specifies an URL (web address) to open
- RA Title: Sets the Title field in the RealAudio player
- RA Author: Sets the Author field in the RealAudio player
- RA Copyright: Sets the Copyright field in the RealAudio player

For example, if you create a marker with the label RA URL:<http://www.sonicfoundry.com>, the RealPlayer would open the Sonic Foundry home page during playback at the marker's location in the file. You can just as easily specify web pages. For example, RA URL:<http://www.sonicfoundry.com/default.html> would open the default.html page on Sonic Foundry's web site.

RealNetworks publishes extensive information on RealAudio Events on their web site (<http://www.real.com>). However, Sound Forge and the Batch Converter Plug-In are capable of generating all event files automatically rather than requiring several tools as described in RealNetworks' documentation.

Automatically Generate RealAudio Metafiles (RAM)

This option simplifies authoring files for the Internet by automatically creating RAM files that are compatible with the RealPlayer.

Create a single metafile for all files in the batch list

This option creates one metafile for all of the files in the batch list rather than creating a single metafile for each file.

RealAudio Server Path

This option defines the server path used as the base in RAM entries.

Use Alternate Destination Folder for RAM files

This option allows you to store your RAM entries apart from where the RealAudio files are accessed. The Browse button is a convenient way to select an existing folder.

Play Start Position

This option specifies the start of a segment that will be played from the RealAudio file. This information is stored in the RAM file. If the value is 0 (default) the entire file will be played.

Play Length

This option specifies the length of a segment that will be played from the RealAudio file. This information is stored in the RAM file. If the value is 0 (default) the entire file will be played.

REAL MEDIA SETTINGS

This dialog is reached when you select *RealMedia (.rm)* from the *File Format* list in the *Conversion Settings* dialog and press the *More* button.

Enable Selective Record Capabilities

When checked, RealMedia-compatible players can allow users to save (record) RealMedia streams to their hard drives. Use this option only if it is acceptable for users to create copies of the RealMedia file. The RealPlayer Plus from RealNetworks supports Selective Record capabilities.

Enable PerfectPlay Capabilities

When this option is enabled, users connected with standard modems can select to download a higher quality version or the audio and/or video that is normally only available over ISDN or LAN connections. The higher bandwidth RealMedia stream is then played back from the user's hard disk after the download process is complete. The RealPlayer Plus from RealNetworks supports PerfectPlay capabilities.

Create Event File (RAE) from Marker and Region Labels

When this option is enabled, events can be embedded from the Regions List (View menu) into the RealMedia stream.

To embed an event, create a marker or region with a label starting with one of the below prefixes followed by a colon.

- **RM URL:** Specifies a URL (web address) to open.
- **RM Title:** Sets the Title field in the RealPlayer.
- **RM Author:** Sets the Author field in the RealPlayer
- **RM Copyright:** Sets the Copyright field in the RealPlayer

For example, if you create a marker with the label **RM URL:http://www.sonicfoundry.com**, the RealPlayer would open the Sonic Foundry home page at the marker's location in the file during playback. You can just as easily open specific web pages. For example, **RM URL:http://www.sonicfoundry.com/default.html** would open the default.html page on Sonic Foundry's web site.

RealNetworks publishes extensive information on RealMedia events on their web site (www.real.com). However, Sound Forge is capable of generating all event files automatically rather than requiring several tools as described in RealNetworks' documentation.

Save Audio Stream

When checked, the file's audio stream will be converted and saved into the RealMedia file.

Audio Compression

This option selects the RealAudio compressor that is used when saving the file. If you need to save RealAudio files with Bandwidth Negotiation capabilities, you can automate the process using Sonic Foundry's Batch Converter Plug-In (available for Sound Forge 4.0 only).

When a compression type is selected, a brief description of it is displayed below the list box. The frequency response parameter of the description indicates how wide, in Hertz, the bandwidth of the audio in the RealMedia file will be. The bandwidth of a normal human's hearing is about 20,000 Hz or 20 kHz.

Save Video Stream

When checked, the file's video stream will be converted and saved into the RealMedia file.

Video Compression

Currently there are two types of RealVideo compression: RealVideo (Standard) and RealVideo (Fractal). Each type is optimized for different types of video and transfer rates. Detailed information on the two types is available on the RealMedia home page (www.real.com).

Video Compression Quality

This option allows you to compromise between quality of compression and lowest possible data transfer rate. Setting this level at 100 provides the best quality video with the greatest required bandwidth to stream. A setting of 1 gives you the lowest quality video but requires the least bandwidth to stream effectively.

Optimize frame rate

Enabling this option allows the RealMedia compressor to determine an optimized frame rate for the required data rate. It is highly suggested that this feature be enabled, since it allows the compressor to vary the frame rate over time depending on the content itself.

Video motion quality

This option determines if the encoder will focus on preserving the smoothness of motion or the quality of the image.

Frames per second

This setting allows you to change the frame rate of the video that is being compressed. A lower frame rate, such as three or five frames per second, will greatly lower the bandwidth requirements for streaming the video. Just type the desired frames per second into the edit box, or choose one of the presets from the drop-down list.

Re-size video frames to

Use this option to change the frame dimensions of the video. Several common sizes can be selected from the drop-down menu. If you'd like to specify dimensions not on the list, select (Custom) and type in the required proportions.

Desired maximum data rate (1.0 to 500.0 Kbps)

This setting is available only if you are saving a video stream. The data rate specifies the desired total kilobits per second (Kbps) for the RealMedia file (including audio and video). To increase the bandwidth available for streaming

the video, you can either lower the data rate for the audio, or increase the desired maximum data rate. Note that you must allow some room for inconsistent transfer rates. For example, a 28.8 modem cannot sustain 28.8 Kbps. Instead, you should target your maximum data rate for between 20.0 and 22.0 Kbps. The allowed data rates are displayed next to the Audio/Video compression options. After the file has been saved, statistics about the actual data rate will be displayed. Use this information to adjust parameters as necessary.

VIDEO FOR WINDOWS (.AVI) SETTINGS

These dialogs are reached when you select Video for Windows (.avi) from the File format list in the Conversion Settings dialog and press the More button.

Frames per second

This setting allows you to change the frame rate of the video that is being compressed. A lower frame rate, such as three or five frames per second, will greatly lower the bandwidth requirements for streaming the video. Just type the desired frames per second into the edit box, or choose one of the presets from the drop-down list.

Re-size video frames to

Use this option to change the frame dimensions of the video. Several common sizes can be selected from the drop-down menu. If you'd like to specify dimensions not on the list, select (Custom) and type in the required proportions.

Compression Options

This dialog is reached by pressing the Compression button.

Choose a Stream

Use this option to select the video stream for which you want to set the compression formats.

Selecting the audio stream in the Compression Options dialog will lead you to the Sound Selection dialog. This dialog is invalid in the Batch Converter. If you choose any of these options they are just ignored.

Interleave Frames

When enabled, this option determines how the audio and video streams will be interleaved. Interleaving the audio and video streams improves playback performance when using storage devices that are more efficient at providing data sequentially (such as CD-ROM disks). The *Frames* value indicates the interval between video and audio segments. The default is 1 for CD-ROMs, but can be increased for faster media.

N O T E

If you do not select any new parameters for video compression and interleaving, the Batch Converter will default to those selected in the original file.

Options

Press this button to view and/or edit video compression options for the selected stream.

Video Compression

This dialog is reached when you press the Options button in the *Compression Options* dialog.

Compressor

Select a compression format from any of the installed .AVI compression schemes. Compression formats are available from different vendors (not from Sonic Foundry), and also come included with Windows.

If the file is already compressed, you should select *No Recompression*. This will be faster and will also avoid the loss of quality inherent in most recompression.

Compression Quality

This control determines the final output quality. Usually, higher quality (close to 100) means less compression and fewer visual artifacts.

Key Frames Every

Some compression formats are based on key frames (sometimes called temporal compression). A key frame is a frame that is usually less compressed. Compressed frames that follow a key frame are smaller but slower to draw. This option is only relevant to video streams.

NOTE

A high number of key frames increases quality, but decreases compression. Also, if the number of key frames is very low, it will take a long time to perform decompression when drawing the video strip or moving randomly through the file.

Data Rate

Determines the data rate that will be required to play the compressed stream in real time. Lower data rates are less demanding on computer systems than higher rates. If the data rate is too high for a certain computer (usually because of a bottleneck such as a slow CD-ROM, network, or hard drive) the quality begins to suffer and glitching can occur.

Configure

Press this button to set more detailed configuration parameters specific to the compression scheme being used. Not all compressors provide this option.

Preview

Press this button to view a preview of the final output.

DATA FORMAT

Many file formats allow you to select between a variety of different data formats. For example, when the Audio Compression Manager (ACM) is installed, the Microsoft Wave file format allows for a variety of compressed data formats, as well as standard PCM data.

ATTRIBUTES

This option determines the attribute choices (if available) of the selected Data format.

ADD OR REMOVE SUMMARY AND EXTRA INFORMATION

When checked, the modifications selected in the *Summary and Extra Information* dialog are applied to the destination sound file. If this option is turned off, all Summary and Extra information in the source sound file is preserved, if possible, in the destination files. Some file formats do not let you store additional information in the file. In such cases this field is ignored.

The *Summary and Extra Information* dialog is opened by pressing the *More* button. All of the options apply only to Microsoft Wave and Video for Windows files with two exceptions. The *Remove Sampler Information* option is available for any file formats that support sampler information and *Summary Information* (Title, Copyright, and Engineer only) can be added to RealAudio, RealMedia and Advanced Streaming Format files, but not removed.

It is sometimes desirable to remove all extraneous information from sound files to save disk space or to ensure compatibility with software that mishandles extra data.

REMOVE SUMMARY INFORMATION

When checked, all *Summary Information* fields will be removed from Microsoft Wave and Video for Windows files.

REMOVE SAMPLER INFORMATION

When checked, all *Sampler Information* will be removed from sound files that support this information.

REMOVE REGION AND PLAY LISTS

When checked, all *Markers*, *Regions* and *Playlist* information will be removed from Microsoft Wave and Video for Windows files.

REMOVE UNKNOWN INFORMATION

When checked, unknown information chunks found in Microsoft Wave and Video for Windows files will be removed.

SUMMARY INFORMATION FIELDS TO ADD OR REPLACE

To add or replace a field, such as copyright information, make sure that it is checked (click in the box to the left of the field name) and enter the desired text into the *Contents of selected field* edit box. Any field that is checked but left blank (no text) will be removed. Unchecked fields are left unmodified.

N O T E

Summary Information is only relevant to Microsoft Wave and Video for Windows files. *Abridged Summary Information* (Title, Copyright, and Engineer) can be added to, but not removed, from RealAudio, RealMedia and Active Streaming Format files.

USE PLUG-INS

Enabling this option allows you to utilize processing features supplied by the built-in or DirectX plug-ins you may have installed.

CLASS

This option allows you to choose the Sound Forge menu where the function you wish to apply is located.

FUNCTION

This option allows you to select the particular function you wish to apply to the files.

PRESET

This option allows you to choose any existing preset for the function that you specified. If *(Untitled)* is selected, the last used parameters will be used. To ensure the use of a certain set of parameters, save the current parameters by using the Save As button in the plug-in's dialog.

SETTINGS

This button opens the corresponding function's dialog. This allows you to view the presets and change the last used parameters.

BATCH SCRIPT

This list contains the functions that will be applied to each file and their exact order. There are four buttons that you use to add, remove and move the desired functions in the list.

Double clicking on a function in the list will bring up the *Settings* dialog for the selected function. Also, you can use the Delete key to remove functions from the list.

ADD

This button adds the selected Plug-in to the *Batch script*.

REMOVE

This button removes the currently selected Plug-in in the *Batch script*.

MOVE UP

This button moves the currently selected Plug-in up one step in the Batch script.

MOVE DOWN

This button moves the currently selected Plug-in down one step in the Batch script.

SAVE AS

The *Save As* button saves a preset of the current settings in the *Conversion Settings* dialog. Presets will store all settings *except* Summary Information fields. This is done because Summary Information is usually common between different batches. If different Summary Information is required between batches, you will need to reenter the information for each batch.

DELETE

This button deletes the current preset from the Name list. Built-in presets can not be deleted.

PREVIEW CONVERSION

The *Preview Conversion* window, which is opened by pressing the *Preview* button in the *Conversion Settings* dialog, provides a quick method of testing different processing options on small regions of sound. You can keep the *Preview Conversion* window open while you change parameters in the *Conversion Settings* dialog making it easy to try out different settings. If no data compression options are chosen, then the preview created is a true representation of the final output.

ACTIVE FILE

Selects the file that will be used to create a preview clip. You can listen to this file by pressing the *Play* button. Only the first 50 files from the batch list are available in the Active file list.

>

This button plays the preview clip after it has been created. During playback, you can press this button again (it will turn into an S) to stop playback.

S

This button stops playback of the preview clip.

B

This button builds and plays a preview clip from the *Active* file using the current conversion settings.

COMMENT FIELDS

These fields allow you to enter comments about the conversion results and/or the parameters used.

PREVIEW START AND LENGTH

This control determines the region in the *Active* file that will be used to create a preview clip.

PLAY

This button plays the unaltered portion of the *Active* file for comparing a preview with the original material.

CLOSE

This button closes the *Preview Conversion* dialog.

CLEAR ALL

This button clears all of the preview buffers and releases any disk space used to store the preview information.

BACK

Switches the active window (keyboard focus) back to the *Conversion Settings* dialog. This is easier to remember than the Alt+F6 standard defined by Windows, though both methods work.

The previews that are created do not include data compression options. Therefore, if you choose a compressed format for the destination (i.e. Microsoft ADPCM or Dialogic VOX ADPCM), you will not hear the effects of the compression in the preview. The preview consists of only the processing performed before saving the destination file. If you need to preview the effects of data compression, simply select one or more files to convert, convert them, then use the *Open Destination* command in the *File* menu. Once the files are open in *Sound Forge*, you can listen to the true effects of the data compression.

NOTE

SOURCE RAW ATTRIBUTES

The Source Raw Attributes dialog is displayed by selecting the Source Raw Attributes option from the Options menu. This dialog is used to specify the format for all raw source files in the list. Using Raw file types may be necessary to open sound files that are not stored in one of the standard file formats supported by the Batch Converter.

All Raw files in a single batch list must have the same attributes. If multiple raw formats need to be converted, they must be converted as separate batches.

NAME

This list contains preset formats that are available. To create your own format from the selected parameters, press the *Save As* button.

SAMPLE RATE

This option determines the playback sample rate that will be used by Sound Forge when playing and processing the file.

SAMPLE TYPE

This option determines the format used to store each sample.

8 BIT PCM

This format is an 8 bit uncompressed linear format.

16 BIT PCM

This format is a 16 bit uncompressed linear format.

G.711 U-LAW

This option is a compressed format commonly used for telecommunications in the United States.

G.711 A-LAW

This option is a compressed format commonly used for telecommunications in Europe.

FORMAT

This option determines the binary format of each sample (for PCM sample types only).

CHANNELS

This option determines the number of channels stored in the file.

BYTE ORDER

This option determines the order in which the high and low bytes of a 16 bit sample are stored. Little endian is used by Intel microprocessors, while Big endian is used by Motorola microprocessors.

HEADER BYTES

This field specifies the number of bytes stored in the file before the sound data.

TRAILER BYTES

This field specifies the number of bytes stored in the file after the sound data.

DESTINATION FOLDER

This command displays a dialog for entering the destination folder that will receive the converted files. Entering a folder that does not exist will cause the Batch Converter to create the folder for you. The **Browse** button is a convenient way to select an existing folder.

If the destination folder is the same as the folder of the source files and the destination file name is the same as the source, then the destination files will be renamed to a new and unique name. This feature provides the greatest protection against accidental loss of sound files.

SHOW FULL PATH

This option switches between showing the full path to every file in the list and showing just the file name. This option can be useful if you are converting files from multiple directories.

SUPPRESS INITIALIZATION DIALOGS

When this option is enabled, any dialogs that would normally be displayed when certain processes are run during the batch will be suppressed. The default values for the dialogs will be used.

STOP ON ERRORS

This option alerts you of any errors which have occurred during processing by launching a dialog. Otherwise, errors are logged and processing continues without stopping. This option should be turned off if you want the batch to run unattended. When Stop on Errors is turned off, all destination file name conflicts are resolved by renaming the destination file to a new and unique name. The Batch Converter will always choose the safest resolution possible when running unattended.

LOG ERRORS TO FILE

When this option is enabled, all errors will be logged to a text file in the location specified by the *Error Log Destination* option.

RESET ERROR LOG

This command erases the contents of the current error log.

ERROR LOG DESTINATION

This option specifies the location of the error log.

REPLACE EXISTING FILES

When this option is selected, files in the destination directory that have the same name as the source file will be overwritten. When this option is disabled, a number will be appended to the file to make the name unique.

DELETE SOURCE AFTER CONVERSION

This option deletes source files immediately after a *successful* conversion. This option is used to save disk space during processing and is very safe if you are certain that you will be satisfied with the converted files. The source file is only deleted after the destination file has been successfully saved. If any errors are encountered during the conversion, the source file will not be deleted.

RENAME WITH LONG FILE NAMES

The Batch Converter the destination file name to a new and unique name if it conflicts with an existing file. With this option enabled, the new name chosen for the destination file will use the same name with a unique number appended to the end. The new file name can exceed the 8 character limit of non-32-bit operating systems.

When this option is turned off, then new file names will truncate the original file name to 6 characters and append a unique two-character number. This maintains compatibility with the 8.3 naming convention.

USE DIRECT MODE WHEN POSSIBLE

When this option is selected, all source files will be opened in Direct mode when possible. Direct mode saves time and disk space by not creating a copy of the source file to work on prior to conversion. This option can only be used in the following cases:

- The source and destination directories are the same and the format of the file is not being changed
- The source and destination directories are different and Use plug-ins is unchecked in the *Conversion Settings* dialog.

If either one of these conditions is not met, a copy of the file will be created to process.

REFRESH LIST

This command updates the batch list to make sure all of the files still exist.

CONVERT MENU

RUN CURRENT SCRIPT

Choosing this command starts the batch conversion process using the current script in the *Conversion Settings* window.

RUN MULTIPLE SCRIPTS

Choosing this command opens the *Run Multiple Scripts* dialog.

SELECT SCRIPTS

This window displays all of the Batch Scripts that are available. Click in the checkbox to choose to use that script.

RUN SCRIPTS

Pressing this button begins the batch conversion process using the selected scripts.

DESTINATION

Pressing this button allows you to choose a specific destination for each script individually.

AUTOMATIC BATCHING

Choosing this command opens the *Automatic Batching* dialog.

FOLDER TO MONITOR AND AUTOMATICALLY BATCH

Type the path or press the *Browse* button to choose the folder to monitor for new files. All new files added to this directory will be converted using the selected scripts. It may take up to 20 seconds before the batch process begins.

AUTOMATIC ARCHIVE FOLDER

Type the path or press the *Browse* button to choose the folder to send copies of the unprocessed files to. All new files added will be copied to this archive folder for safekeeping.

SCRIPTS

Pressing this button take you to the *Run Multiple Scripts* dialog. From here you will choose the scripts you wish to run on the files and the destination directories. Press the *Run Scripts* button to switch back to the *Automatic Batching* window.

START

Pressing this button tells the *Batch Converter* to begin monitoring the specified directory. This button will change to *Stop* once pressed. Pressing *Stop* will abort the monitoring process. Any files that exist in the folder will be batched as well.

STATUS

This text displays the current status of the *Automatic Batching* process.

STARTED

This text displays the date and time that the *Automatic Batching* was started.

FILES PROCESSED

This text displays the number of files that have been processed since the *Automatic Batching* was started.

CONVERT SELECTED ONLY

This command only converts the files that are selected in the batch list.

SPECTRUM ANALYSIS REFERENCE

OPTIONS MENU

SETTINGS

Changes the parameters used in calculating and displaying the graphs. The Spectrum Settings dialog can be reached from the Options menu or the Spectrum or Sonogram shortcut menus. You can store the current settings by selecting Save As. Stored settings are displayed in the Name drop-down list.

FFT SIZE

Size in samples of the analysis window and number of discrete frequencies analyzed. Higher numbers correspond to more frequency resolution at the expense of lower time resolution and longer computational time.

NOTE

A 128 point FFT performed on a 44,100 Hz file analyzes 0.003 seconds (128/44100) of sound while a 2048 point FFT analyzes 0.046 seconds (2048/44100). The 128 point FFT will result in 64 (128/2) discrete frequency bands with a bandwidth of 345 Hz (44100/128) per band. On the other hand, a 2048 point FFT results in 1024 discrete bands with a bandwidth of 22 Hz each.

FFT OVERLAP

This option controls the amount of overlap between FFT analysis windows. A lower setting decreases the number of distinct analysis functions performed, which decreases processing time. High settings allow for more analysis, which can result in extremely slow processing. Too small an overlap tends to make the Sonogram image look chunky, depending on the length of the selection.

SMOOTHING WINDOW

This option determines the window function applied to the input data before analysis. The window function has a strong effect on the sharpness of peaks in an FFT graph and the leakage into neighboring frequencies (the theory and correct nomenclature behind these functions is best left to advanced DSP textbooks). The effects of these windows can be easily appreciated by analyzing a single sine wave.

A rectangle window means no window is applied. Using a rectangular window results in a very sharp peak but high leakage.

A triangle window (also called a Bartlett or Parzen window) results in a bit less leakage than the rectangle window.

Hamming, Hanning, and Blackman windows are commonly used in audio applications.

Using the Blackman-Harris window results in the least side-band leakage of the six windowing of functions included at the expense of rounded graph peaks.

SLICES DISPLAYED

This option determines the number of FFT slices displayed. Each slice represents FFT size samples.

FORWARD/BACKWARD

When displaying multiple slices in the Spectrum Graph, the display is (from front to back) either Forward or Backward in time in the sound file depending on the selection of the radio buttons to the right of the slice selection.

SET SONOGRAM RESOLUTION

This option determines the number of FFT samplings used in the Sonogram, keeping the processing time and graph resolution constant. Increasing the samplings increases the horizontal graph resolution at the expense of longer processing time.

When disabled, the number of samplings is determined by the length of the selection and the Overlap. Using this mode can result in some very long processing times.

LOGARITHMIC GRAPHING

This option displays the x-axis in logarithmic mode, rather than linear mode. In logarithmic mode, more area of the graph is devoted to lower frequencies. This option is available in Normal display only.

FREQ. MIN.

This option determines the lowest frequency displayed in graphs when Zoom to Range is selected.

FREQ. MAX

This option determines the highest frequency displayed in graphs when Zoom to Range is selected.

CEILING

This option determines the highest amplitude level displayed in graphs when Zoom to Range is selected.

FLOOR

This option determines the lowest amplitude level displayed in graphs when Zoom to Range is selected.

HOLD PEAKS DURING MONITORING

When this option is selected, the highest value of each frequency bin will be held on the Spectrum Graph with a small line.

MAINTAIN LAST MONITORED VIEW

When this option is selected, the state of the Spectrum Graph when you stop playback will be maintained. If this option is not selected, the graph will reset to the cursor position.

AUTO UPDATE

When this option is selected, the Spectrum Graph will update for any change in the cursor position or selection of the current file.

MONITOR PLAYBACK

When this option is selected, the Spectrum Graph will update in real time as the current file plays. The response of the graph will depend on the current FFT size and the speed of your computer.

MONITOR RECORD

When this option is selected, the Spectrum Graph will update in real time to any input to your selected sound card. This feature makes it possible to do real time analysis of acoustics in rooms and the like.

SHOW POSITION

When this option is selected, the current position of the mouse will be displayed in a tool tip over the Spectrum Graph. This is handy for taking measurements of the graph.

SHOW NOTES

When this option is selected, the nearest musical note equivalent of the frequency that the mouse is over is shown in a tool tip over the Spectrum Graph.

PRINT DISPLAY

This command outputs the Spectrum Analysis window contents to a printer.

CLOSE

This command closes the Spectrum Analysis window.

DISPLAY MENU

ZOOM TO RANGE

This command sets the amplitude and frequency range to the *Display Range* maximum and minimum values specified in the *Spectrum Settings* dialog.

ZOOM OUT FULL

This command allows you to view the entire frequency and amplitude range (0 Hz to Nyquist frequency). This option is handy if you're lost somewhere in the graph.

NORMALIZE DB

This command sets the Spectrum Graph amplitude range equal to the maximum and minimum values in the spectrum graph.

LOGARITHMIC

This command displays the Spectrum Graph (x-axis) logarithmically.

N O T E

This mode is not available in the Sonogram display.

GRAB/PAN

This command enables you to slide up and down the Spectrum Graph by dragging the mouse vertically or horizontally. Cursor keys can also be used. This option is only available when you are zoomed in on a selection.

SYNC GRAPHS

Synchronizes both displays in a stereo file to allow you to view the same region of the FFT in both channels.

NORMAL DISPLAY

When selected, the Spectrum Graph (amplitude on y axis ,frequency on x axis) is displayed. If multiple slices are selected, a time axis is also displayed.

SONOGRAM (COLOR)

When selected, a color Sonogram Graph (frequency on y axis, time on x axis, amplitude represented by color) of the selected region is displayed.

SONOGRAM (B&W)

When selected, a black and white Sonogram Graph (frequency on y axis, time on x axis, amplitude represented by darkness) of the selected region is displayed.

LINE GRAPH

When this option is selected, a single, unfilled line is used to draw the Spectrum Graph.

BAR GRAPH

When this option is selected, a series of bars corresponding to the FFT bins is used to draw the Spectrum Graph.

FILLED GRAPH

When this option is selected, the Spectrum Graph is displayed as a filled graph.

N O T E

Some video drivers have problems displaying in this mode. If you encounter problems such as incorrect shading or very slow drawing, it is recommended that you use the Line Graph or Bar Graph options or change video drivers.

UPDATE!

Selecting the Update! menu causes the Spectrum Analysis Plug-In to re-analyze the current selection and update the frequency or Sonogram Graph.

If you want the graph to automatically update every time data is altered or the selection or cursor position changes, select the *Auto Update* option in the Options menu.

TROUBLESHOOTING

PLAYING AND RECORDING

When playing files from my SCSI drive, why does the sound have a warbling effect?

When playing files from a SCSI drive that is connected to a SCSI controller that utilizes a DMA channel (such as most older ISA controllers) you may experience a warbling effect in the sound playback. This can occur if your sound card uses DMA and its DMA channel is set to a higher number than your SCSI controller.

To fix this problem, set the DMA channel used by your sound card to a lower number than that used by your SCSI controller. For example, if your SCSI controller is set to DMA channel 5 and your sound card is set to DMA channel 7, try reconfiguring them so the sound card uses DMA channel 5 and your SCSI controller uses DMA channel 7.

When I start Sound Forge my *Play* and/or *Record* buttons/options are grayed out. Why can't I play or record?

When Sound Forge starts, it checks to see if you have a wave playback and record device installed in your system. If you don't, it disables the appropriate operations. If you are sure that you have a sound card installed in your system, yet Sound Forge does not recognize this, you need to check a couple of things. First check the *Wave* folder in the *Preferences* dialog. The two settings of importance are the *Playback* and *Record* drop-down lists. These options determine the wave playback and record devices Sound Forge will use. By default Sound Forge uses the *Wave Mapper* (it may also be called the *Sound Mapper*) device which should automatically pick an appropriate device for use by Sound Forge. If either the *Playback* or *Record* device is set to *(None)* then Sound Forge can not execute the corresponding operation. To learn more about configuring your wave devices refer to the *Wave Preference* folder in the Reference chapter.

I have sound files that don't want to play on my system. Sound Forge loads them fine, why won't it play them?

Sound Forge allows you to load many different types of sound files, including ones that your system may not be capable of playing. There are a variety of reasons that your sound card may not be capable of playing a file.

- Some sound cards only allow playback and recording at specific sample rates. Typically these rates are limited to 11,025, 22,050, and 44,100 Hz. If you try to play a file that is 32,000 Hz the sound card will not allow you to play it.
- If you have an older sound card it may only support mono playback. If this is the case simply convert any stereo files to mono and you will be able to play them.
- If you have an older sound card it may only support 8-bit playback. If this is the case simply change the sample size of the file from 16 to 8 bits.

Why is it when I play my sound files the play pointer that Sound Forge displays seems to be off from what I am hearing?

For Sound Forge to display the current position of the pointer relative to playback of audio data it must rely on your sound card's driver to provide the correct position. Many sound card drivers don't do a very good job of this. Typically the position these cards provide is both off by a fixed amount as well as reported only in large discrete steps. Sound Forge provides a way to compensate for this by setting various options in the *Wave Preferences* folder. Refer to the *Wave Preferences* chapter in the Reference chapter for more information.

Why can't I get Sound Forge to respond to levels from my microphone or my CD player?

Most sound cards come with multiple inputs from which you can record sound. These may include microphone, CD, auxiliary, MIDI, and line inputs. Although some cards record from all of these inputs at the same time, many cards make you choose which device is currently your record device (also known as the input source). These cards come with software which allows you to pick which device is the input source. You will need to refer to your sound card manual to determine how to set the input source.

Most sound cards also come with some kind of application which controls the input level of the recording device. You may need to increase the level of your input sources in order to record from them. Again, you will need to refer to your sound card manual to learn how to adjust these levels.

When I record or play files I hear what sounds like small clicks in the sound.

When Sound Forge is recording or playing audio data it moves large amounts of data between your hard drive and your sound card. This can be a tremendous amount of work for your computer depending on the kind of system you have. If Sound Forge can't move the data fast enough then you will hear small gaps between blocks of data that typically sound like clicks or puttering. There are a number of things you can do to reduce the overhead during the time Sound Forge is moving sound data, all of which are covered in the chapter called *Optimizing Sound Forge*.

When I record (or play) data at high sample rates, I can't get Sound Forge to stop by pressing the *Stop* button.

Sound Forge tries very hard to keep recording and playback of audio smooth (no gaps in your sound). So there may be delays between the time that you press the *Stop* button and the time that Sound Forge actually stops. Note that you can always press and hold the *Escape* key to stop record or playback no matter how bad the situation is. Just be a little patient and wait a few seconds. These delays are usually only observed when one or more components of your system can not keep up with the data rate. Typical faults are:

- Very slow (older) hard drives.
- Poorly written sound card drivers.
- Setting your *Storage* directory to a slow device like a network drive or a compressed hard drive.

If at all possible, use the fastest uncompressed storage device available with Sound Forge. For more information on configuring your system, see the chapter on *Optimizing Sound Forge*.

When playing sound files with rates above 44,100 Hz on my *Turtle Beach Multisound* or *Tahiti*, my system crashes.

This is caused by a severe bug in the *Turtle Beach Multisound* and *Tahiti* sound card drivers. You will experience the same crash no matter what software you use to play sample rates above 44,100 Hz — including *Media Player* and even *Turtle Beach's* own *Wave* application.

Contact *Turtle Beach* for updated drivers.

When playing sound files at 8,000 Hz on my *Turtle Beach Multisound*, the play pointer moves but I don't hear any output.

This is caused by a bug in the *Turtle Beach Multisound* driver.

Contact *Turtle Beach* for an updated driver.

When playing sound files at rates other than 11,025 Hz, 22,050 Hz, or 44,100 Hz the files sound wrong (the pitch is wrong).

Some sound drivers incorrectly play sounds with sample rates that are not supported by the sound card hardware. For example, a sound that is supposed to play at 13,000 Hz might get played at 11,025 Hz instead. This will cause the pitch to be lower and the playback time to be slower than it should be. The *Microsoft Windows Sound System Version 2.0* and many “compatible” sound cards contain this unfortunate bug.

Contact your sound card manufacturer for updated drivers.

When playing sound files sometimes when the sound starts I hear a blip, drop-out, or click.

Some sound drivers do not deal correctly with the pre-load buffer option of *Sound Forge*. If you have set your pre-load size in the *Wave Preferences* folder to a value greater than 0, try setting it back to 0. If this solves the problem then you have a driver that does not deal with this situation correctly.

Contact your sound card manufacturer for updated drivers.

FILE SUPPORT

When saving files as Macintosh AIFF and transferring them to my Macintosh, why doesn't the Macintosh recognize them as sound files?

The Macintosh uses special information within its operating system to keep track of the types of files it can access. Sound Forge can save the audio data in the correct file format, but it can't tell the Macintosh what type it is. To do this you must use a program on the Macintosh, like *ResEdit*, which allows you to change the file's type. You simply need to set the *Type* field to AIFF for the Macintosh to recognize the file as a sound file.

There is also a 3rd party utility called *AIF Typer* that is available for download from our web site. It is a fairly easy way to make your AIF files readable on a Macintosh.

Why won't Sound Forge open my compressed files like Microsoft ADPCM?

Sound Forge has built in support for the Microsoft Audio Compression Manager (ACM). This allows you to open files that are compressed with a variety of algorithms including Microsoft ADPCM, IMA ADPCM, and other third party compression schemes like *The DSP Group's TrueSpeech*.

If you can't open a Microsoft ADPCM file, you may not have the Microsoft ADPCM driver installed and enabled for the ACM. Please refer to the Appendices on Sound Forge and The Microsoft Audio Compression Manager for more information on how to use the ACM with Sound Forge.

After converting WAV files to Redbook audio using the software that came with my compact disc recorder, there are audible clicks at the end of the track.

Not all software for converting .WAV files to Redbook audio handles the .WAV file structure correctly. Some software simply skips the first 44 bytes of data in the .WAV file and then considers everything else in the file to be audio data. This means that if you have saved Summary information, Sampler information, or Regions and Playlist data in the .WAV file it will be interpreted as audio data.

If you are experiencing this problem simply use the *Save As* command in Sound Forge to save a new .WAV file with the *Save Summary Information*, *Save Sampler Information*, and *Save Regions/Playlists* check boxes unchecked. This tells Sound Forge to save only the minimum amount of information needed for a .WAV file.

My compact disc recorder uses a file type with an .RBK extension. What is an .RBK file and can I open it with Sound Forge?

RBK files are simply raw PCM sound data files. These files are used by some compact disc recorders to store Redbook audio onto a compact disc. RBK files are data only and contain no other information. They are stored as 16 bit, stereo data with a sample rate of 44,100 Hz. There are a number of things that make them different than a normal 16-bit raw file:

- The channels are stored in reverse order.
- The data is stored in Big endian format (Motorola).
- The data length of the file has to be a multiple of 2,352 bytes (1 CD sector).

The size of a CD sector is calculated as follows:

1 second of 44.1k, 16 bit, stereo audio = 176,400 bytes

CD audio has 75 sectors of data for each 1 second of audio

Sector size = $176,400 / 75 = 2,352$ bytes

To open an existing RBK file, simply use the raw file type in Sound Forge's Open dialog box and select the RBK preset. Remember that you will have to swap the channels if you want to maintain the correct left/right channel settings.

EFFECTS

After running an effect on a file, why is there a loss in volume level?

Most effects that can alter the volume level of a file have two fader controls, Dry Out and Effect Out (where "Effect" is the name of the function, like Reverb). These controls should be adjusted to get the proper balance between processed and non-processed sound and also determine the combined volume level. For example, if you run an effect with both controls set at 40% and the level after processing is too low, simply raise both the Dry and Effect Out faders an equal amount to raise the final level. However, take care to keep the output signal from clipping.

Why do some effects take so long to process a file?

Some processing functions used by Sound Forge are very math intensive. Examples include EQ, Pitch Shift without change of duration, and the Reverb effect. This means that millions of integer and floating-point arithmetic operations must be performed for each second of sampled data. If your

computer's CPU does not have a math coprocessor (in the Intel architecture that means if you have an i486 SX instead of an i486 DX or Pentium), floating-point operations will drastically slow you down and there's not much you can do except save up for a faster processor.

Of course, running effects on stereo, 44,100 Hz, hour long files will slow you down no matter what machine you are using. If you are going to be deleting or inserting space into these huge files, we recommend that you break the files into smaller ones and then mix the smaller files together once you are done processing them. This is especially painful when you insert or delete from the beginning of a large file, since Sound Forge must then shift the rest of the file over.

One way to get around disk-related editing delays is to use the *Regions List* and *Playlist* to make all arrangements non-destructively and then use the *Convert to New* function on your *Playlist*. Another useful feature in Sound Forge is the *Preview Cut/Cursor* function under the *Edit Menu* – use this function to verify your *Cut* and *Clear* operations before actually changing the file.

SAMPLING AND SEQUENCING

When using SMPTE/MTC to drive my sequencer, my sequencer sometimes loses sync.

Depending on the performance of your system (and your sequencer) you may experience SMPTE sync drop-out problems. This is typically due to too many things going on at once (i.e. more than your computer is capable of maintaining). For example, if you have wave playback, synchronization messages, sequencing, and perhaps more going on at the same time, you may experience sync drop-out.

To compensate for this, many sequencers allow you to set a freewheel time. This is a slack time between SMPTE messages during which the sequencer will continue to play even if SMPTE/MTC messages are being missed. By increasing the freewheel time you allow Sound Forge and your sequencer to be less tightly dependent on SMPTE messages, but this may also decrease your synchronization accuracy. The best recommendation is to start with a small freewheel time and, if you experience sync drop-out problems, increase the freewheel time until the drop-outs go away.

MIDI

After installing the Sonic Foundry Virtual MIDI Router in Windows 95, my system is having trouble with MIDI devices or the VMR.

If you have a large number of MIDI devices in your system or a multi-port MIDI board, such as an eight port MIDI interface, you may have problems installing additional VMR MIDI interface ports. Windows 95 has a limit on the number of MIDI devices that it can support. Depending on the installation, this number typically varies between 10 and 12 devices. If you have trouble with MIDI devices causing system errors after installing the VMR, you should either reduce the number of ports in use by the VMR, remove the VMR, or reduce the number of MIDI devices in your system.

.AVI FILES

When opening an .AVI file, if the audio stream format is 44.1 kHz stereo file Microsoft ADPCM, the audio track does not open correctly.

This is due to a known bug in *Video for Windows*. It only happens if the file is both 44.1 kHz, stereo, and compressed using Microsoft or IMA ADPCM. To get around this problem, resample the audio to 22 kHz before saving.

APPENDIX A: SHORTCUTS

This section contains a list of available shortcuts including:

- Keyboard Shortcuts
- Shortcut Menus
- Drag and Drop Shortcuts
- Other Mouse Shortcuts

KEYBOARD SHORTCUTS

The following is a comprehensive list of keyboard shortcuts available in Sound Forge.

GENERAL KEYBOARD SHORTCUTS

FILE COMMANDS

Press	To
Alt+Enter	Show the <i>Properties</i> dialog for the active data window
Control+N	Create a new data window
Control+O	Open a sound file into a new data window
Control+S	Save modified sound data back to the file
Control+W	Close the active data window
Alt+F4	Exit Sound Forge

DATA WINDOW MANAGEMENT COMMANDS

Press	To
Alt+0	Set input focus to the waveform display in the active window
Alt+1	Set input focus to the <i>Regions List</i>
Alt+2	Set input focus to the <i>Playlist</i>

Alt+3	Set input focus to the <i>MIDI Keyboard</i> window
Alt+4	Set input focus to the <i>Audio Device Mixer</i> window
Alt+5	Set focus to the <i>Time Display</i> window
Alt+6	Set focus to the <i>Video Preview</i> window
Alt+7	Set focus to the <i>Play Meter</i> window
Alt+8	Set input focus to the <i>Undo/Redo History</i> window
Alt+F5	Restore the <i>Sound Forge</i> application window
Alt+F10	Maximize the <i>Sound Forge</i> application window
Control+F5	Restore the active data window
Control+F6	Go to the next data window
Control+F10	Maximize the active data window
Control+Shift+F6	Go to the previous data window
Shift+F4	Tile the data windows vertically
Shift+F5	Cascade the data windows

EDIT COMMANDS

Press

Control+A

Control+C

Control+E

Control+F

Control+K

Control+Shift+K

Control+M

Control+T

Control+U

Control+V

To

Select all data in the active window

Copy the selected data onto the clipboard

Paste the clipboard contents into a new data window

Crossfade data from the clipboard with the active window

Preview Cut or *Clear* in the active window

Play to cursor with pre-roll

Mix data from the clipboard with the active window

Trim (crop) to the current selection

Toggle *Undo buffer* creation for the active window

Paste data from the clipboard into the active window

Control+X	Move (cut) the selected data onto the clipboard
Control+Y	Repeat last process, effect or tool
Control+Z	Undo the last action
Control+Shift+Z	Redo the last undone action
Delete	Clear (delete) the selected data; nothing placed on the clipboard

MISC. VIEW, OPTION AND SPECIAL COMMANDS

Press	To
S	Toggle current selection on and off
Control+B	Toggle <i>Auto Snap to Zero</i> on and off for the active data window
Control+Enter	Maximize the width of the active data window
Control+D	Show the <i>Set Selection</i> dialog
Control+G	Show the <i>Go To</i> dialog
Control+L	Show the <i>Loop Tuner</i> for the active data window
Control+R	Record new data into a data window
Escape	Stop or cancel the current action (including playback)
F1	Open the <i>Sound Forge Help</i> file to the <i>Contents</i> page
F10	Make the menu bar active
Shift+F10	Display a shortcut menu
F6	Toggle playback scrolling on and off
Shift+F6	Toggle smooth playback scrolling on and off
F8	Toggle <i>Drag and Drop Snapping</i> on and off
F7	Turn on <i>MIDI Input Sync/Trigger</i>
Shift+F7	Turn on <i>MIDI Output Sync</i>

WAVEFORM DISPLAY KEYBOARD SHORTCUTS

CURSOR MOVEMENT

Press	To move to
Right Arrow	The next screen pixel
Left Arrow	The previous screen pixel
End	The last sample visible in the waveform display
Home	The first sample visible in the waveform display
Control+End	The last sample in the data window
Control+Home	The first sample in the data window
Page Up	10% of the current view past the cursor position
Page Down	10% of the current view previous to the cursor position
Control+Right Arrow	2% of the current view past the cursor position
Control+Left Arrow	2% of the current view previous to the cursor position
Numpad +	The next sample
Numpad -	The previous sample
Control+Numpad +	10 samples past the current cursor
Control+Numpad -	10 samples previous to the current cursor

WHEN REGIONS, LOOPS OR MARKERS EXIST IN THE FILE:

Control+ Left Arrow	The previous region, loop or marker boundary
Control+ Right Arrow	The next region, loop or marker boundary

Selecting Data

Press	To select from the cursor to
Shift+Right Arrow	The next screen pixel
Shift+Left Arrow	The previous screen pixel
Shift+End	The last sample visible in the waveform display
Shift+Home	The first sample visible in the waveform display
Control+Shift+End	The last sample in the data window
Control+Shift+Home	The first sample in the data window
Shift+Page Up	10% of the current view past the cursor position

Shift+Page Down	10% of the current view previous to the cursor position
Control+Shift+Right Arrow	2% of the current view past the cursor position
Control+Shift+Left Arrow	2% of the current view previous to the cursor position
Shift+Numpad +	The next sample
Shift+Numpad -	The previous sample
Control+Shift+Numpad +	10 samples past the current cursor
Control+Shift+Numpad -	10 samples previous to the current cursor

WHEN REGIONS, LOOPS OR MARKERS EXIST IN THE FILE:

Control+Shift+Left Arrow	The previous region, loop or marker boundary
Control+Shift+Right Arrow	The next region, loop or marker boundary

NAVIGATION AND PLAYBACK

Press	To
Control+<Number>	Save a view in cell <Number> where <Number> ranges from 1 to 8
<Number>	Restore a view using cell <Number> where <Number> ranges from 1 to 8
Up Arrow	Increase time magnification (zoom in closer to data)
Down Arrow	Decrease time magnification (zoom out farther from data)
Shift+Up Arrow	Increase level magnification
Shift+Down Arrow	Decrease level magnification
Control+Up Arrow	Zoom Selection if a selection exists, Zoom In Full if no selection
Control+Down Arrow	Zoom Normal (zooms to default zoom ratio set in Preferences)
Control+Shift+Up Arrow	Pan data window up if zoomed in vertically
Control+Shift+Down Arrow	Pan data window down if zoomed in vertically
Numpad 5 or *	Switch cursor to opposite end of selection

Tab	Cycle stereo selection from left channel to right channel to both channels
Shift+Tab	Cycle stereo selection from both channels to right channel to left channel
[(open bracket)	Set <i>Mark In</i> at the current cursor position
] (open bracket)	Set <i>Mark Out</i> at the current cursor position
M or \ (back slash)	Drop a marker at the current cursor position
C or . (period)	Center the cursor in the waveform display
Spacebar	<i>Play</i> or <i>Stop</i> the current data window in default mode
Shift+Spacebar	<i>Play All</i> or <i>Pause</i> the current data window
Control+Spacebar	Switch play mode between <i>Normal</i> , <i>Looped</i> and <i>As Sample</i>
Enter	Pause playback; keep the cursor at the current position
Escape	Stop playback of the current data window
Control+K	<i>Preview Cut</i> (skip selection on playback with pre-roll)
Control+Shift+K	<i>Play to cursor</i> with pre-roll

REGIONS LIST AND PLAYLIST KEYBOARD SHORTCUTS

REGIONS LIST SHORTCUTS

Press	To
Spacebar	<i>Play</i> or <i>Stop</i> the active marker or region
Enter	Edit the active marker or region
Delete	Delete the active marker or region
` (grave accent)	Cycle through the <i>Regions List</i> display formats
R	Create region from the current selection
Shift+R	Create region without bringing up dialog
L	Create loop from the current selection
Shift+R	Create loop without bringing up dialog

PLAYLIST SHORTCUTS

Press	To
Spacebar	Play (Stop) from the active <i>Playlist</i> entry
Enter	Edit the active <i>Playlist</i> entry
Delete	Delete the active <i>Playlist</i> entry
+ (plus sign)	Add one to the active <i>Playlist</i> entry play count
- (minus sign)	Subtract one from the active <i>Playlist</i> entry play count
* (asterisk)	Add or remove a <i>Stop Point</i> on the active <i>Playlist</i> entry
/ (forward slash)	Toggle <i>Pre-roll</i> on and off for the <i>Playlist</i>
` (grave accent)	Cycle through the <i>Playlist</i> display formats

SHORTCUT MENUS

Shortcut menus are menus that appear when you click on certain areas of the screen. They allow you to quickly access important functions.

SOUND FORGE DESKTOP SHORTCUT MENU

This shortcut menu is reached by right-clicking on the desktop workspace of Sound Forge. This menu has options for sound file management on your workspace. You can open recently used files from the *Recent Files* option. The *Preferences* dialog can quickly be reached from this menu as well.

WAVEFORM DISPLAY SHORTCUT MENU

This shortcut menu is reached by right-clicking on the waveform display of a sound file. It is a way of quickly reaching the following commands and dialogs: *Zoom In Full/Normal/Out Full*, *Paste*, *Crossfade*, *Mix*, *Cursor Snap to Time/Zero*, *Cursor Go To*, *Scroll Playback*, *Properties*.

PLAYBAR SELECTION STATUS SHORTCUT MENU

This shortcut menu is reached by right-clicking on any of the playbar's selection status fields. It allows you to set the current status format.

REGIONS LIST SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Regions List*. It allows you to perform the following *Regions List* operations: *Add*, *Delete*, *Edit*, *Replicate*, *Split*, *Update*, *Markers to Regions*, *Open*, *Save As*, *Copy onto Clipboard*.

PLAYLIST SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Playlist*. It allows you to perform the following *Playlist* operations: *Add*, *Delete*, *Edit*, *Replicate*, *Stop Point*, *Pre-roll Playback*, *Open*, *Save As*, *Copy onto clipboard*, *Convert to New*.

TIME RULER SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Time Ruler*. It allows you to reach the *Add Marker/Region* or *Edit Sample* dialog for adding regions, markers, and loop points, as well as changing the status format.

DISPLAY PROPERTIES SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Edit Tool* selector at the upper left-hand corner of the data window. It allows you to set options described under *Properties | Display* in the *File Menu*.

LEVEL RULER SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Level Ruler*. It allows you to change the labeling used, as well as various level zooming operations.

ZOOM OUT FULL

This command zooms out vertically to view the maximum amplitude range.

ZOOM WINDOW

This command zooms vertically to optimally fit the entire waveform display.

ZOOM SELECTION

This command zooms vertically to optimally fit the current selection.

LABEL IN PERCENT

When this option is checked, the *Level Ruler* will be labeled in percent.

LABEL IN DB

When this option is checked, the *Level Ruler* will be labeled in decibels.

LEVEL TOOLBAR SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Levels* toolbar. It sets the status of the display for the toolbar.

INTEGER

This option displays the sample value at the cursor as an integer. The range is from -32768 to 32767 in 16-bit audio and -128 and 127 in 8-bit audio.

PERCENT

This option displays the sample value at the cursor as a percentage. The range is from -100 to 100%.

dB

This option displays the sample value at the cursor as a dB value. A value of 0 dB corresponds to maximum absolute amplitude and negative infinity (-Inf.) corresponds to complete silence. In 16-bit audio, -90.3 dB is the lowest possible dB value (sample value of 1).

PEAK

This option displays the peak sample value in the vicinity of the cursor as a dB value. The amount of time which is scanned by default is 50 milliseconds. This can be set on the *Status* page of the *Preferences* folder.

RMS

This option displays the Root Mean Square (RMS) of the sample values in the vicinity of the cursor as a dB value. RMS is a more accurate measurement of loudness than the peak value. The amount of time which is scanned by default is 50 milliseconds. This can be set on the *Status* page of the *Preferences* folder.

The maximum value of 0 dB is only achieved by a square wave (or DC offset) of maximum amplitude. A maximum amplitude sine wave reaches only -3 dB. Values below -96 dB can be achieved by combination of silence and very low-levels, and only complete silence will register negative infinity.

MARKER TAG SHORTCUT MENU

This shortcut menu is reached by right-clicking on a marker tag in the ruler.

GO TO

This command moves the cursor to a marker position.

DELETE

This command removes a marker from the *Regions List*.

EDIT

This command takes you to the *Edit Region/Marker* dialog.

UPDATE

This command changes a marker to the current cursor position.

REGION TAG SHORTCUT MENU

This shortcut menu is reached by right-clicking on a region tag in the ruler.

SELECT

This command makes the selection equal to the region end points.

DELETE

This command removes the region from the *Regions List*.

EDIT

This command takes you to the *Edit Region/Marker* dialog.

SPLIT

This command splits the region into two regions using the current cursor position as the split point.

UPDATE

This command makes region end points equal to the current selection.

LOOP TAG SHORTCUT MENU

This shortcut menu is reached by right-clicking on a loop tag in the ruler.

SELECT

This command makes the selection equal to the loop end points.

DELETE

This command removes the loop.

EDIT

This command takes you to the *Edit Sample* dialog.

UPDATE

This command makes the loop end points equal to the current selection.

STATUS BAR SHORTCUT MENUS

These shortcut menus allow you to change the playback rate, sample size, channels (*Stereo/Mono*), and status format by right-clicking on the status bar fields.

WAVE, MIDI, AND AUXILIARY MIXER SHORTCUT MENUS

This shortcut menu is reached by right or left-clicking on the *Wave*, *MIDI*, or *Aux* headings in the *Mixer* window. They allow you to select the *Wave*, *MIDI*, or *Auxiliary* device that the volume fader controls. Many sound cards do not allow for fader controls, in which case this feature will not work.

KEYBOARD MIDI OUT SHORTCUT MENU

This shortcut menu is reached by right or left-clicking on the *MIDI Out* button on the *Keyboard*. It allows you to set the device to which the *MIDI Keyboard* will transmit.

PLAY METERS SHORTCUT MENU

Right-click on the *Play Meters* to reach this menu will allows you to select the *Play Meter* resolution and other options including: *Reset Clip*, *Show Labels*, *Hold Peaks*, *Hold Valleys* and *Always on Top*.

RESET CLIP

This command resets the clip indicator. Clicking on the indicator does the same thing.

RANGE (-90, -78, -60, -42, -24 TO 0 DB)

This option determines the bottom-most level that will be displayed. A value of -90 dB is the absolute minimum peak value in 16-bit audio, corresponding to a sample value of +/- 1.

SHOW LABELS

This command displays labels and tick marks on the meter.

SHOW PEAKS

When this option is turned on, the maximum detected peaks are displayed for around 500 milliseconds.

SHOW VALLEYS

When this option is turned on, the minimum detected peaks are displayed for around 500 milliseconds.

ALWAYS ON TOP

Use this option to have the *Play Meters* window always be above any other windows.

PREVIEW CONFIGURATION SHORTCUT MENU

This shortcut menu is reached by right-clicking on any unused area in a dialog which has *Preview* capabilities.

PRE-ROLL

When enabled, sound data preceding the current selection is played before the processed data during previews.

POST-ROLL

When enabled, sound data following the current selection is played after the processed data during previews.

LEFT

Use this option to process left channel only when using a stereo sound file.

RIGHT

Use this option to process right channel only when using a stereo sound file.

BOTH

Use this option to process both channels when using a stereo sound file.

RESET ALL

This command resets all of the dialog's parameters to their default values.

SAVE CURRENT

This command saves the current dialog settings to a new preset.

CONFIGURATION

This command takes you to the *Preview Configuration* dialog.

RECORD DIALOG SHORTCUT MENU

This command this shortcut menu is reached by right-clicking anywhere in the *Record* window. The *Remote Record* shortcut menu is similar, with options not available in the dialog available from the menu.

RECORD METERS

This command allows you to set the record meter range and other options.

GAP DETECTION

This option determines what order of action to take if a gap is detected while recording.

BLINKING STATUS

This command enables the blinking record status underneath the record meters.

TIME DISPLAY SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Time Display* window. It allows you to set the display format for the *Time Display*.

POSITION

This option displays the position at the edit or play cursor.

SYNC/TRIGGER STATUS

This option displays incoming MIDI time code and trigger signals. The current MIDI Input port is set on the MIDI/Sync page of the *Preferences* folder.

PLAYLIST POSITION

This option displays the total playback time starting always at 00:00:00:000, not the play cursor time, when playing from the *Playlist*.

PASSIVE UPDATE

When this option is turned on, the *Time Display* status will only be displayed when the processor is idle. Use this when using a slow computer to prevent the audio from glitching.

VIDEO STRIP SHORTCUT MENU

This shortcut menu is reached by right-clicking on the video strip above the waveform display.

COPY

This command copies the selected frame as a bitmap to the clipboard. You could, for example, then paste it into the *Summary* information folder (*Properties* in the *File* menu) by right-clicking over the shot you want to copy.

GO TO

This command moves the cursor to the beginning of the frame.

ANIMATE

This command displays the frame closest to the cursor position.

NUMBER FRAMES

This command displays frame number on each frame.

VIDEO PREVIEW SHORTCUT MENU

This shortcut menu is reached by right-clicking on the *Video Preview* window. It allow you to set various options of the video display. A complete list and explanation of these options is available in the Reference section under the Options menu item.

DRAG AND DROP SHORTCUTS

Using drag and drop allows you to quickly perform operations crossing between open data windows, the *Playlist*, the *Regions List*, and the *Time Ruler*.

DRAG TO NEW

To create a new file from the current selection, drag the selection to an open area of the Sound Forge desktop. You can also drag regions from the *Regions List* to the desktop.

DRAG MIX

To mix a selection into another data window, drag the selection from the source to the destination window. Processing begins at the point where you drop the region in the destination file. When mixing a mono selection into a stereo file, you can mix the selection to both channels by dropping the selection on the destination centerline. Otherwise, the selection will be mixed into either the left or right channel only.

DRAG CROSSFADE

To crossfade a selection into another data window, drag the selection from the source to the destination window. Processing begins at the point where you drop the region in the destination file. Press the *Control* key while in the destination window. When crossfading a mono selection into a stereo file, you can crossfade the selection to both channels by dropping the selection on the destination centerline. Otherwise, the selection will be crossfaded into either the left or right channel only.

DRAG PASTE

To paste a selection into another data window, drag the selection from the source to the destination window. Processing begins at the point where you drop the region in the destination file. Press the *Alt* key while in the destination window. When pasting a mono selection into a stereo file, you can paste the selection to both channels by dropping the selection on the destination centerline. Otherwise, the sound data will be placed in either the left or right channel only, with silence in the other.

DRAG TO REGIONS LIST

To add the current selection to the *Regions List*, drag it to the *Regions List*. You can also drag regions from the *Regions List* to the *Playlist*.

DRAG TO PLAYLIST

To add a region in the *Regions List* to the *Playlist*, drag it from the waveform display or the *Regions List* to the *Playlist*. You can also drag regions within the *Playlist* to rearrange the playback order.

DRAG TO TIME RULER

To create a regions, drag the current selection to the *Time Ruler*.

OTHER MOUSE SHORTCUTS

SELECT ALL

Double-click on the waveform display to select the entire sound file. Triple-click when regions, loops or markers are present.

ZOOM BOTH TIME AND LEVEL

Double-click on the *Level Ruler* to zoom the current selection both vertically and horizontally. If no selection exists, the entire waveform display data is zoomed. Double-click again to *Zoom Out Full* vertically and *Zoom Normal* horizontally.

MAGNIFY MODE

To zoom in to a section, select an area while holding the *Control* key. This will switch Sound Forge to *Magnify* mode, which zooms to fit the selected area.

RETURN CONTROL VALUE TO DEFAULT

Double-click on a trackbar, fader, or spinner to set it to its default value.

FINE TUNE CONTROL VALUE

To fine tune a trackbar, fader, or spinner, hold down both right and left mouse buttons when dragging. This can also be done if you hold down the *Control* key while dragging.

PREVIEW

Shift-click on the *Preview* button to hear the original sound. This is equivalent to checking the *Bypass* check box.

Control-click on the *Preview* button to reach the *Preview Configuration* dialog.

MAIN STATUS BAR

Double-click on the *Sample Rate*, *Sample Size*, or *Channels (Stereo/Mono)* field to reach the *Properties* dialog.

SELECTION STATUS BAR

Double-click on the left-most field to reach the *Go To* dialog. Double-click on the other two fields to reach the *Set Selection* dialog.

GO TO MARKER

Double-click on a marker tag in the ruler to move the cursor to the marker's position.

SET SELECTION TO REGION/LOOP

Double-click on a region or loop tag in the ruler to change the current selection to the region or loop end points.

EDIT REGION OR MARKER

Double-click on a region or marker in the *Regions List* to reach the *Edit Region/Marker* dialog.

EDIT PLAYLIST

Double-click on a *Playlist* entry to reach the *Edit Playlist* dialog.

PLAY NORMAL BUTTON

Control-click to *Preview Cut* (skip selection) on playback. Control+shift-click to play to cursor with pre-roll.

SLOW AND FAST SELECTION SCROLL TOGGLE

When making a selecting past the end or beginning of the waveform display, while holding down the left mouse button, click on the right mouse button to toggle between fast scrolling and slow scrolling.

APPENDIX B: SOUND FORMATS

The file formats recognized by Sound Forge are listed below along with the extension used to identify them to the user. An arbitrary extension has been assigned for each of the formats used by computers other than the PC. Many of these formats when found on another type of computer will tend to use the extension .SND. Thus, even if a file has the .SND extension it may not be a Sounder/SoundTool file.

* Indicates this file format supports the saving of sample looping information.

Ad Lib Sample **(.SMP)**

This format is used by the Ad Lib Gold card, for its sampled instruments. It supports, 8-/16-bit, stereo/mono, and 4-bit Yamaha ADPCM compression.

Amiga SVX File **(.SVX/.IFF)**

This file type is found on the Commodore Amiga and is similar in substance to the Microsoft .RIFF Wave format.

Covox 8-Bit File **(.V8)**

This format is used with Covox software. It is an 8-bit mono uncompressed file format.

Creative Labs VOC File **(.VOC)**

This is one of the more commonly found sample sound formats found on PC-compatible computers. The .VOC format supports packed data that Sound Forge will unpack prior to importing the file. The .VOC format also supports information for silence, looping, and varying sample rates. When Sound Forge imports a .VOC file it uses the first sample rate found and ignores such information as looping and silence. Sound Forge will save .VOC files as unpacked data with none of the additional information such as looping or silence. Sound Forge gives you the ability to save .VOC files in either the old or new .VOC standard. The new .VOC standard supports 16-bit data. Refer to the Preferences section in the Reference chapter for information on configuring Sound Forge to always use the new .VOC standard.

Dialogic VOX File **(.VOX)**

This format is used with specialized voice data boards. It is a 4-bit mono ADPCM file format which expands to 16-bit data.

Gravis Patch File* **(.PAT)**

This format is used by the Gravis UltraSound Card. Sound Forge can only read and write Gravis Patch files which have single samples stored in them. Patch files with multiple samples are not recognized. Sound Forge recognizes Gravis Patch files that are of version 1.10. Older versions, such as 1.00, are not supported.

Intervoice (IVC)

This file format is for use with Interoice telephony systems. It is a mono format with a variety of compression schemes available.

Macintosh AIFF* (.AIF/ .SND)

This format is used on the Apple Macintosh to save sound data files. An .AIFF file is best when transferring files between the PC and the Mac using a network, since most network software will delete any resource information in a Macintosh file. Sound Forge can also read an .AIFF with a Mac binary header attached but will identify the file as a Macintosh Resource instead.

Macintosh Resource (.SND)

This format is used on the Apple Macintosh to save sound resources. When reading a file of this type it must have a Mac-Binary header. This header is attached to binary files transferred with common Macintosh file transfer programs. Sound Forge does not save to this file format.

MIDI SDS File* (.SDS)

This format is a file that contains the exact data that would be transferred to a sampler when using the MIDI Sample Dump Standard (SDS) protocol. It is a series of MIDI SysEx messages. MIDI SDS supports mono files only.

NeXT/Sun (Java) File (.AU/ .SND)

This is the standard format of sound files found on the NeXT and Sun Sparc station computers. This format has a large number of sub types of file data. Sound Forge reads and writes the most common data formats for these files, 8-bit linear, 16-bit linear, and μ -Law.

RealAudio/RealVideo (.RA)

This format is a compressed format designed for real time audio and video streaming over the Internet. It has a variety of compression options for different bandwidths.

SampleVision File* (.SMP)

This format is used with Turtle Beach's SampleVision software. It is a 16-bit mono file format. The .SMP format is meant to be used with sampling keyboards.

Sonic Foundry Sample Resource (.SFR)

This format is an old format defined by Sonic Foundry and is supplied for compatibility. It is not recommended for use in saving files.

Sound Designer 1* (.DIG/ .SD)

This format is used by Sound Designer 1 on the Macintosh.

Sounder/SoundTool File (.SND)

This format is used with the shareware applications Sounder and Sound Tool. It is very popular on bulletin boards and various online services.

Video for Windows (.AVI)

This is a Microsoft digital video file format. It supports multiple channel audio, but Sound Forge will only allow you to open one audio and one video stream at a time. When saving as an .AVI, many compression schemes are available for both the video and audio streams.

Wave (Microsoft) File* (.WAV)

This is the format for sampled sounds defined by Microsoft for use with Windows. It is an expandable format that supports multiple data formats and compression schemes.

Raw Files (*.*)

Raw files can be composed of any sound data. If the data is any form of 8- or 16-bit PCM (uncompressed) or G.711 A-Law or μ -Law (a form of companded compression), then Sound Forge can open and save the file.

APPENDIX C: SOUND FORGE AND THE MICROSOFT AUDIO COMPRESSION MANAGER

THE AUDIO COMPRESSION MANAGER (ACM)

The *Microsoft Audio Compression Manager* (more commonly referred to as the *ACM*) is a standard interface for audio compression in Windows. This interface allows applications like *Sound Forge* to use compression algorithms provided by other companies.

Sound Forge fully supports audio compression through the *ACM*. This enables you to use any *ACM*-compatible compression. The best part of this support is you don't have to learn anything new to use it! *Sound Forge* transparently opens compressed *.WAV* files and provides all available compression formats for *.WAV* files in the *Save As* dialog.

There are two major pieces to the *ACM* (though these pieces are tightly coupled). It is not necessary for you to know any details about these pieces to use the *ACM* from *Sound Forge*, so this information is supplied for those who want to know more. That being said, the three pieces of the *ACM* are:

- Audio data compression and decompression
- Transparent playback and recording of non-hardware supported audio files

AUDIO DATA COMPRESSION AND DECOMPRESSION

The first piece of the *ACM* allows you to compress and decompress audio data. Audio compression is used to decrease the amount of data required to represent a sound; this ultimately results in smaller sound files. However, there are drawbacks to using audio compression on your sound files:

1. Most audio compression algorithms will degrade the quality of the sound. This is referred to as *lossy* compression since information contained in the sound is lost when it is compressed. The amount of sound degradation is entirely dependent on the algorithm.

2. Audio compression requires more processing time than uncompressed data. The amount of processing time is dependent on the algorithm, as well as your hardware setup. The result is opening and saving compressed files will usually take longer than uncompressed files.
3. Compressed files are not as portable as uncompressed files. If you are distributing .WAV files in a compressed format, you must make sure the person receiving the files can use them. Also, not all audio software can use compressed .WAV files, which could make using other programs with Sound Forge inconvenient.

In Sound Forge, any compressed .WAV file can be opened as long as a compatible ACM driver is installed and enabled. If there is no compatible ACM driver available for a compressed .WAV file, Sound Forge will inform you of the problem when you try to open it.

Saving compressed .WAV files is as simple as choosing the compression algorithm in the *Format* drop-down list of the *Save As* dialog. Once a file has been saved as compressed, Sound Forge will always save changes to the file using the selected compression algorithm; you do not need to re-select the compression format each time you save. Of course you can change the compression format at any time with the *Save As* dialog. You can also revert to an uncompressed format.

Experimentation is probably the most important step in deciding whether or not to use audio compression for your projects.

TRANSPARENT PLAYBACK AND RECORDING OF NON-HARDWARE SUPPORTED AUDIO FILES

The third part of the ACM is probably the most significant (or at least the coolest). It is called the *Sound Mapper* and allows audio data formats that are not directly supported by your sound card to be played and recorded. Sound Forge lets you use the *Sound Mapper* by selecting it for *Playback* and *Record* in the *Wave* folder of the *Preferences* dialog. Additional configuration of the *Sound Mapper* for your system can be done using the Windows Control Panel.

Let's go through some examples of what the *Sound Mapper*, can do. If you have a sound file that is recorded at a bizarre sample rate like 22,257 Hz (perhaps a file that came from an external sampler), and the closest sample rate that your sound card supports is 22,050 Hz, then the sound file normally cannot be

played. You would have to change the sample rate of the file to 22,050 Hz before you could play it. However, changing the sample rate (without resampling) would cause the sound to play at a lower pitch.

The *Sound Mapper*, can play this sound file correctly without forcing you to resample the file first. The *Sound Mapper*, will map the sound to the best format possible and perform the resampling in real time. Yes Virginia, there is a Santa Claus!

Other examples include the ability to play 16-bit sounds on an 8-bit-only sound card. Or playing stereo sounds on a mono-only sound card. And believe it or not, you can record stereo files on a mono-only sound card. However, this last example is essentially useless because the *Sound Mapper*, will simply duplicate the mono input of your sound card in both channels. So keep in mind that the *Sound Mapper*, cannot create something that is better than what the sound card can supply.

The usefulness of the *Sound Mapper*, for non-compressed sound files can be summarized as convenience. You don't have to convert the sound to a supported format before you listen to it. Some of the tasks that the *Sound Mapper*, can perform (like format conversions while recording) are available not because they are useful in real world situations, but because the developers of the *ACM* knew it could be done. Much like climbing a mountain for no better reason than it's there.

So what about compressed sound files? Yes, the *Sound Mapper* allows you to play (and sometimes record) compressed sound files. Even on sound cards that do not support compression directly. So you can play a sound file that is compressed with Microsoft ADPCM or The DSP Group's TrueSpeech on any sound card without decompressing the sound file first. Many multimedia and entertainment software titles use compressed sound files to save disk space. The software then plays the files through the *Sound Mapper*.

We stated above that the *Sound Mapper* can sometimes record compressed sound files. The reason for not always being able to record compressed files (even when the compressed file can be played) is simple. Compressing sound data can be very computationally expensive (it takes a long time); the amount of time required is completely dependent on the compression algorithm and how it is implemented. Decompressing sound data is almost always faster than compressing the same sound data that makes it possible to play it back in real time.

It should be noted that Sound Forge does not play and record compressed sound files directly. Rather, all compression and decompression is performed while opening and saving the sound files. This limitation is actually insignificant; Sound Forge is the perfect tool for authoring compressed sound files for the ACM and many software authors are using Sound Forge for exactly this. Sound Forge saves compressed sound files using the best quality possible – something that cannot always be done in real time. So saving compressed sound files with Sound Forge will usually sound better than those recorded with audio compression.

Notes on Sound Forge's sound file compression support:

- The *Open* dialog allows you to preview compressed .WAV files if you have an appropriate ACM driver installed. However, you must have your *Playback* device set to the *Sound Mapper* for this to work (*Preferences* | *Wave* in the *Options* menu).
- When saving uncompressed audio data to a compressed format with the *Save As* dialog, it is a good idea to close the file and re-open it after saving. Since Sound Forge does the compression and decompression during saving and loading you will not be able to hear what the file sounds like with compression until after you've saved and reloaded the file.

APPENDIX D: SMPTE TIME CODE

The Society of Motion Picture and Television Engineers (SMPTE) time code may be one of the most misunderstood concepts among individuals within the music industry. After working with SMPTE time code for years, many people are still confused by the concept, so don't feel bad if you haven't got it all figured out. Hopefully this discussion will clear the mud.

The biggest problem with SMPTE time code is that, depending on whether you sit on the video or audio side of the fence, SMPTE time codes may mean different things to you. When dealing with SMPTE you will probably see five, perhaps six, different types of time code formats (six is for the people who are really confused).

Here is a description of each SMPTE time code format:

SMPTE 25 EBU

This SMPTE code runs at 25 frames per second and is also known as SMPTE EBU (European Broadcasting Union). The reason for having this rate is that European television systems run at exactly 25 frames per second.

SMPTE 24 FILM SYNC

This SMPTE code runs at 24 frames per second and is also known as SMPTE Film Sync. This rate matches a nominal film rate of 24 frames per second (the slowest speed possible for apparent continuous motion).

OK, those two are easy. Now things start to get a little crazy.

SMPTE 30 NON-DROP (AS USED IN THE AUDIO WORLD)

In the US, the 60 Hz power system makes it easy to generate a time code rate of 30 frames per second. This rate is commonly used in audio environments and is typically known as 30 Non-Drop. You will probably use this rate when synchronizing audio applications like a multi-track recorder or your MIDI sequencer.

NOTE

If all you care about is working with audio and not dealing with video, stop reading right here. We mean it! All you really need to know is that there are three different SMPTE rates you might want to use: SMPTE 24, SMPTE 25, and SMPTE 30 Non-Drop. However, be aware that SMPTE 30 Non-Drop in the video world runs at 29.97 frames per second.

TRUE SMPTE 30 DROP AND SMPTE 30 NON-DROP (AS USED IN THE VIDEO WORLD)

If you are planning to work with video, the frame rate of exactly 30 frames per second is never used. When NTSC color systems were developed, the frame rate was changed by a tiny amount to eliminate the possibility of crosstalk between the audio and color information. Even though it is still referred to as SMPTE 30 Drop or Non-Drop, the actual frame rate that is used is exactly 29.97 frames per second. This poses a problem since this small difference will cause SMPTE time and real time (what your clock reads) to be different over long periods. Because of this, two methods are used to generate SMPTE time code in the video world: Drop and Non-Drop.

In SMPTE Non-Drop, the time code frames are always incremented by one in exact synchronization to the frames of your video. However, since the video actually plays at only 29.97 frames per second (rather than 30 frames per second), SMPTE time will increment at a slower rate than real world time. This will lead to a SMPTE time versus real time discrepancy. Thus, after a while, we

could look at the clock on the wall and notice it is farther ahead than the SMPTE time displayed in our application.

SMPTE Drop time code (which also runs at 29.97 frames per second) attempts to compensate for the discrepancy between real world time and SMPTE time by dropping frames from the sequence of SMPTE frames in order to catch up with real world time. What this means is that occasionally in the SMPTE sequence of time, the SMPTE time will jump forward by more than one frame. The time is adjusted forward by two frames on every minute boundary except 00, 10, 20, 30, 40 , and 50. Thus when SMPTE Drop time increments from 00:00:59:29, the next value will be 00:01:00:02 in SMPTE Drop rather than 00:01:00:00 in SMPTE Non-Drop. In SMPTE Drop, it must be remembered that certain codes no longer exist. For instance, there is no such time as 00:01:00:00 in SMPTE Drop. The time code is actually 00:01:00:02.

When synchronizing audio to video, it is crucial that the SMPTE time code (30 Drop or Non-Drop) used in your sequencer or digital audio workstation is the same as the SMPTE time code striped onto the video. Only then will the SMPTE times on the video screen and computer monitor match exactly during playback.

In the audio world, people have started to call 30 Non-Drop (which runs at 29.97 frames per second) 29.97-Non-Drop to distinguish it from the 30 Non-Drop used between audio applications (which runs at a true 30 frames per second). SMPTE 30 Drop (as used in video) may also be referred to as SMPTE 29.97 Drop just to reiterate that the frame rate is actually 29.97 frames per second. It just depends on who you talk to.

However, you must remember that there is no difference between 30 Drop and 29.97 Drop time code.

HOW SOUND FORGE DEALS WITH THE ISSUE

For SMPTE 24, SMPTE 25, and SMPTE 30 Drop time codes, Sound Forge behaves as you would expect. The 24, 25 and 30 Drop frame rates are calculated and generated at exactly 24, 25, and 29.97 frames per second respectively. In 30 Drop rate, the dropped frames are inserted to make up for the time discrepancies with real world time.

For SMPTE 30 Non-Drop, the default method for Sound Forge is to calculate all SMPTE values at exactly 30 frames per second. If you are working with video and need SMPTE Non-Drop running at 29.97 frames per second, you can turn off the *Use 30 frames per second for SMPTE Non-Drop* option on the MIDI/Sync page of the *Preferences* folder. This will cause Sound Forge to use 29.97 frames per second instead of 30 frames per second for SMPTE 30 Non-Drop.

Some calculations for your benefit:

SMPTE 30 fr = 33.3333 ms per frame; SMPTE 29.97 fr = 33.3667 ms per frame

1 minute = 60 seconds * 30 = 1,800 frames per minute

1 minute = 60 seconds * 29.97 = 1,798.2 frames per minute

Difference per minute = 1.8 frames

1 hour = 3,600 seconds * 30 = 108,000 frames per hour

1 hour = 3,600 seconds * 29.97 = 107,892 frames per hour

Difference per hour = 108 frames

Thus in drop frame we need to catch up by 108 frames per hour.

For example, in Non-Drop mode the SMPTE time code of 01:00:00:00 will equate to a real world time of 01:00:03:18. Therefore, there is a lag behind real world time of three seconds and 18 frames after one hour.

Here is a nifty chart to show how SMPTE Drop works. As time progresses, two frames are dropped each minute except at the 00, 10, 20, 30, 40 , and 50 minute points. Notice that after one hour, 108 frames have been dropped. This number corresponds to the number of drop frames we calculated above needed to match SMPTE 30 (29.97 fps) time to real time.

NUMBER OF DROPPED FRAMES VERSUS MINUTES ELAPSED OVER ONE HOUR

Minutes	00	10	20	30	40	50
	xx:00:00:00	xx:10:00:00	xx:20:00:00	xx:30:00:00	xx:40:00:00	xx:50:00:00
0	0 frames	18 frames	36 frames	54 frames	72 frames	90 frames
1	2	20	38	56	74	92
2	4	22	40	58	76	94
3	6	24	42	60	78	96
4	8	26	44	62	80	98
5	10	28	46	64	82	100
6	12	30	48	66	84	102
7	14	32	50	68	86	104
8	16	34	52	70	88	106
9	18	36	54	72	90	108

APPENDIX E: GENERAL MIDI PATCH MAP

1-8 Piano	9-16 Chromatic Percussion	17-24 Organ
1 Acoustic Grand Piano	9 Celesta	17 Drawbar Organ
2 Bright Acoustic Piano	10 Glockenspiel	18 Percussive Organ
3 Electric Grand Piano	11 Music Box	19 Rock Organ
4 Honky-tonk Piano	12 Vibraphone	20 Church Organ
5 Electric Piano 1	13 Marimba	21 Reed Organ
6 Electric Piano 2	14 Xylophone	22 Accordion
7 Harpsichord	15 Tubular Bells	23 Harmonica
8 Clav	16 Dulcimer	24 Tango Accordion

25-32 Guitar	33-40 Bass	41-48 Strings
25 Acoustic Guitar (nylon)	33 Acoustic Bass	41 Violin
26 Acoustic Guitar (steel)	34 Electric Bass (finger)	42 Viola
27 Electric Guitar (jazz)	35 Electric Bass (pick)	43 Cello
28 Electric Guitar (clean)	36 Fretless Bass	44 Contrabass
29 Electric Guitar (muted)	37 Slap Bass 1	45 Tremolo Strings
30 Overdriven Guitar	38 Slap Bass 2	46 Pizzicato Strings
31 Distortion Guitar	39 Synth Bass 1	47 Orchestral Strings
32 Guitar Harmonics	4 Synth Bass 2	48 Timpani

49-56 Ensemble	57-64 Brass	65-72 Reed
49 String Ensemble 1	57 Trumpet	65 Soprano Sax
50 String Ensemble 2	58 Trombone	66 Alto Sax
51 Synth Strings 1	59 Tuba	67 Tenor Sax
52 Synth Strings 2	60 Muted Trumpet	68 Baritone Sax
53 Choir Aahs	61 French Horn	69 Oboe
54 Voice Oohs	62 Brass Section	70 English Horn
55 Synth Voice	63 Synth Brass 1	71 Bassoon
56 Orchestra Hit	64 Synth Brass 2	72 Clarinet

73-80 Pipe	81-88 Synth Lead	89-96 Synth Pad
73 Piccolo	81 Lead 1 (square)	89 Pad 1 (new age)
74 Flute	82 Lead 2 (sawtooth)	90 Pad 2 (warm)
75 Recorder	83 Lead 3 (calliope)	91 Pad 3 (polysynth)
76 Pan Flute	84 Lead 4 (chiff)	92 Pad 4 (choir)
77 Blown Bottle	85 Lead 5 (charang)	93 Pad 5 (bowed)
78 Shakuhachi	86 Lead 6 (voice)	94 Pad 6 (metallic)
79 Whistle	87 Lead 7 (fifths)	95 Pad 7 (halo)
80 Ocarina	88 Lead 8 (bass+lead)	96 Pad 8 (sweep)

97-104 Synth Effects	105-112 Ethnic	113-124 Percussive
97 FX 1 (rain)	105 Sitar	113 Tinkle Bell
98 FX 2 (soundtrack)	106 Banjo	114 Agogo
99 FX 3 (crystal)	107 Shamisen	115 Steel Drums
100 FX 4 (atmosphere)	108 Koto	116 Woodblock
101 FX 5 (brightness)	109 Kalimba	117 Taiko Drum
102 FX 6 (goblins)	110 Bagpipe	118 Melodic Tom
103 FX 7 (echoes)	111 Fiddle	119 Synth Drum
104 FX 8 (sci-fi)	112 Shanai	120 Reverse Cymbal
121-128 Sound Effects		
121 Guitar Fret Noise	124 Bird Tweet	127 Applause
122 Breath Noise	125 Telephone Ring	128 Gunshot
123 Seashore	126 Helicopter	

PERCUSSIVE CHANNEL 10 NOTE VALUES AND PATCHES

35 Acoustic Bass Drum	51 Ride Cymbal 1	67 High Agogo
36 Bass Drum 1	52 Chinese Cymbal	68 Low Agogo
37 Side Stick	53 Ride Bell	69 Cabasa
38 Acoustic Snare	54 Tambourine	70 Maracas
39 Hand Clap	55 Splash Cymbal	71 Short Whistle
40 Electric Snare	56 Cowbell	72 Long Whistle
41 Low Floor Tom	57 Crash Cymbal 2	73 Short Guiro
42 Closed Hi Hat	58 Vibraslap	74 Long Guiro
43 High Floor Tom	59 Ride Cymbal 2	75 Claves
44 Pedal Hi-Hat	60 Hi Bongo	76 Hi Wood Block
45 Low Tom	61 Low Bongo	77 Low Wood Block
46 Open Hi-Hat	62 Mute Hi Conga	78 Mute Cuica
47 Low-Mid Tom	63 Open Hi Conga	79 Open Cuica
48 Hi-Mid Tom	64 Low Conga	80 Mute Triangle
49 Crash Cymbal 1	65 High Timbale	81 Open Triangle
50 High Tom	66 Low Timbale	

APPENDIX F: MIDI NOTES AND FREQUENCIES

The following chart shows the musical pitch and frequency associated to each MIDI Note in Sound Forge. A4 is at 440 Hz and middle C is C4 at 261.63 Hz. These values can differ, often by one octave, from values used by other manufacturers. To avoid having negative values, some MIDI sequencers and sound modules use C0 as MIDI Key 0, which then makes middle C C3 instead of the more standard C4.

The frequencies below are based on equal temperament tuning.

Pitch	MIDI Key	Frequency
C -1	0	8.176
C# -1	1	8.662
D -1	2	9.177
D# -1	3	9.723
E -1	4	10.301
F -1	5	10.913
F# -1	6	11.562
G -1	7	12.250
G# -1	8	12.978
A -1	9	13.750
A# -1	10	14.568
B -1	11	15.434
C 0	12	16.352
C# 0	13	17.324

Pitch	MIDI Key	Frequency
D 0	14	18.354
D# 0	15	19.445
E 0	16	20.601
F 0	17	21.826
F# 0	18	23.124
G 0	19	24.499
G# 0	20	25.956
A 0	21	27.500
A# 0	22	29.135
B 0	23	30.867
C 1	24	32.703
C# 1	25	34.648
D 1	26	36.708
D# 1	27	38.890

Pitch	MIDI Key	Frequency
E 1	28	41.203
F 1	29	43.653
F# 1	30	46.249
G 1	31	48.999
G# 1	32	51.913
A 1	33	55.000
A# 1	34	58.270
B 1	35	61.735
C 2	36	65.406
C# 2	37	69.295
D 2	38	73.416
D# 2	39	77.781
E 2	40	82.406
F 2	41	87.307
F# 2	42	92.499
G 2	43	97.998
G# 2	44	103.82
A 2	45	110.00
A# 2	46	116.54
B 2	47	123.47
C 3	48	130.81
C# 3	49	138.59
D 3	50	146.83

Pitch	MIDI Key	Frequency
D# 3	51	155.56
E 3	52	164.81
F 3	53	174.61
F# 3	54	184.99
G 3	55	195.99
G# 3	56	207.65
A 3	57	220.00
A# 3	58	233.08
B 3	59	246.94
C 4	60	261.63
C# 4	61	277.18
D 4	62	293.66
D# 4	63	311.13
E 4	64	329.63
F 4	65	349.23
F# 4	66	369.99
G 4	67	391.99
G# 4	68	415.31
A 4	69	440.00
A# 4	70	466.16
B 4	71	439.88
C 5	72	523.25
C# 5	73	554.37

Pitch	MIDI Key	Frequency
D 5	74	587.33
D# 5	75	622.25
E 5	76	659.26
F 5	77	698.46
F# 5	78	739.99
G 5	79	783.99
G# 5	80	830.61
A 5	81	880.00
A# 5	82	932.32
B 5	83	987.77
C 6	84	1046.5
C# 6	85	1108.7
D 6	86	1174.7
D# 6	87	1244.5
E 6	88	1318.5
F 6	89	1396.9
F# 6	90	1480.0
G 6	91	1568.0
G# 6	92	1661.2
A 6	93	1760.0
A# 6	94	1864.7
B 6	95	1975.5
C 7	96	2093.0

Pitch	MIDI Key	Frequency
C# 7	97	2217.5
D 7	98	2349.3
D# 7	99	2489.0
E 7	100	2637.0
F 7	101	2793.8
F# 7	102	2960.0
G 7	103	3136.0
G# 7	104	3322.4
A 7	105	3520.0
A# 7	106	3729.3
B 7	107	3951.1
C 8	108	4186.0
C# 8	109	4434.9
D 8	110	4698.6
D# 8	111	4978.0
E 8	112	5274.0
F 8	113	5587.7
F# 8	114	5919.9
G 8	115	6271.9
G# 8	116	6644.9
A 8	117	7040.0
A# 8	118	7458.6
B 8	119	7902.1

Pitch	MIDI Key	Frequency
C 9	120	8372.0
C# 9	121	8869.8
D 9	122	9397.3
D# 9	123	9956.1
E 9	124	10548.1
F 9	125	11175.3
F# 9	126	11839.8
G 9	127	12543.9

APPENDIX G: RIFF WAVE CHUNKS

The following is an overview of the defined RIFF Wave File chunks which Sound Forge 4.0 stores and recognizes. Although Microsoft recommends that the audio data be stored after all other chunks, the only chunks stored by Sound Forge prior to the data are the `fmt`, `fact`, and `PAD` chunks.

The reasoning proposed for storing the data after all other chunks is if the data is modified, the other chunks do not have to be rewritten to disk. However, since the auxiliary chunks are usually much smaller than the data, the associated time to rewrite them is usually negligible. For example if a Playlist chunk is modified to add a new region we have to rewrite the complete contents of the data to insert the new Playlist chunk. However if it is stored at the end of the data only the Playlist chunk and any chunks following it need to be rewritten. If we were operating with 3 minutes of 44k stereo data then we would need to rewrite 30 megabytes of data to insert 1 region change. The auxiliary chunks typically are less than 100k so we only have to move 100k with changes to the audio data itself.

Another reason for storing additional chunks after the data is that some less robust software packages assume that all that will be found in a `.WAV` file prior to the data is the `fmt` chunk. This allows you to still read `.WAV` files to some degree with these packages.

Chunk	Use
PAD	This chunk is used by Sound Forge to align the DATA chunk on a sector boundary for use with products like Digidesign's Session 8.
fact	Used to specify the length of audio data in samples when data is stored in a format other than PCM.
cue	Used to specify the start of markers and regions.
plst	Used to specify a Playlist within a single wave file.
adtl	Used to specify region lengths and to provide text names for regions and markers.
smpl	Used to specify information for files used with samplers. Includes information on Release and Sustaining loops.
tlst	Used to specify external trigger events for Regions List and Playlist entries.
inst	Unsupported, but documented for consistency.

For more information on the above chunks, refer to the RIFF specification available from Microsoft.

The following chunk has been proposed by Sonic Foundry but has not been registered with Microsoft. This chunk is used to save trigger information in .WAV files created by Sound Forge.

TLST (TRIGGER LIST) CHUNK

Added: ??

Author: Sonic Foundry

Defined for : .WAV form

The <tlst-ck> trigger list chunk specifies a list of triggers which can be used to trigger playback of a series of cue points or Playlist entries. The <tlst-ck> is defined as follows:

```

|<tlst-ck>|  () tlst(  <dwTriggers:DWORD> // count of triggers
              <trigger(s)>...) // trigger table

<trigger>  struct
            {
            FOURCC  fccListReference;
            DWORD   dwName;
            DWORD   dwType;
            DWORD   dwTriggerOn;
            DWORD   dwFunction;
            DWORD   cbExtra;
            }

```

The <trigger> chunk:

fccList	Reference Specifies the list which this entry references. A trigger can reference either the <cue-ck> or the <plst-ck>.
DwName	When the trigger references the <cue-ck> this field specifies the cue point name. This value must match one of the names listed in the <cue-ck> cue-point table. When the trigger references the <plst-ck> this field references a Playlist entry by index. This value must be less than the total count of <plst-ck> entries.
DwType	Specifies the type of trigger. 0 - SMPTE Trigger 1 - MIDI Command Trigger. 2 - MIDI SysEx Trigger. (SysEx message)

dwTriggerOn

Specifies the value which will cause a trigger to be generated. This value will change value depending on the type of trigger.

For triggers of type SMPTE (0)

Specifies a SMPTE time offset for the cue. The format of this value is 0xhhmmssff. hh is an unsigned Hours value [0..23]. mm is an unsigned Minutes value [0..59]. ss is an unsigned Seconds value [0..59]. ff is an unsigned Frames value [0..numFrames].

For triggers of type MIDI Command (1)

Specifies a MIDI Command.

Specified as :

Channel Command Param1 Param2

Byte 0 Byte 1 Byte 2 Byte 3

Where Byte 0 is the MSB and Byte 3 is the LSB.

For channels a value of FF signals trigger on all channels.

For any other param other than the MIDI command byte FF signals a don't care value.

For triggers of type MIDI Sysex (2)

Unused, should be set to 0.

DwFunction	<p>Specifies the function of this trigger.</p> <p>0 - Play</p> <p>1 - Stop</p> <p>2 - Queue</p> <p>Trigger functions which are unsupported or unknown should be ignored.</p>
CbExtra	<p>Specifies the size in bytes of additional trigger information. For the case of a MIDI Sysex Trigger, this is the size of the matching MIDI sysex command immediately following the structure.</p> <p>Any value within a sysex that is set to FF is treated as a don't care. This allows response on multiple channels, values, etc.</p>

EXAMPLES:

Play when any MIDI note on is received:

FF 90 FF FF

Play when MIDI note 60 is received on Channel 1:

00 90 3C FF

Play when MIDI note 60 is received on Channel 1, stop when note off received for same note:

00 90 3C FF

00 80 3C FF

Queue when MIDI note 60 is received on Channel 2, Play when note off received for same note:

01 90 3C FF

01 80 3C FF

MARKERS, REGIONS, AND PLAYLIST ENTRIES AS USED BY SOUND FORGE

When creating or reading markers, regions and playlist entries in .WAV files, the RIFF documentation can be a little confusing. We have added this section to give software developers a better idea of how to deal with these chunks.

MARKERS

Markers are a position within a wave file that can also have an associated text name. Each marker in Sound Forge is stored as an entry within the <cue> chunk. The entry within the <cue> chunk stores the position, as well as a DWORD name field. The name field is used to reference the text name entry which can be found within the <adtl> chunk. For instance, if we had a marker at 0 which had a name of *Start* there would be an entry within the <cue> chunk with a sample offset of 0 and a name entry of 0x0001 (we have chosen 0x0001 for the example. The name entry can be any value from 0x00000000 to 0xFFFFFFFF) Within the <adtl> chunk we would have an entry of type <labl> (label chunk) with a name of 0x0001 and a text field of *Start*.

REGIONS

Regions are much like markers but they also include an associated length as well as a start position and text name. Each region is stored as an entry within the <cue> chunk which stores the position as well as a DWORD name field. The name field is again used to reference the text name entry which can be found within the <adtl> chunk. The difference between the marker and the region is that instead of using a <labl> chunk in the <adtl> chunk we use a <ltxt> chunk which allows us to store a length as well as a text name. For instance if we had a region at 0 with a length of 1000 samples and a name of *Beat 1* there would be an entry within the <cue> chunk with a sample offset of 0 and a name entry of 0x0001. (Again, we're using the name value of 0x0001 as an example.) Within the <adtl> chunk we would have an entry of type <ltxt> with a Name of 0x0001, a text field of *Beat 1*, and a length of 1,000. Within the <ltxt> chunk there is also a field of type FOURCC which specifies the purpose of this entry. For all regions in Sound Forge the FOURCC type is set to *rgn*.

PLAYLIST ENTRIES

Playlist regions are stored a little differently. They specify a reference to a <cue> for the start of each playlist region and a length and number of loops for each entry. Thus, our playlist entries can actually have lengths that are different than those stored within the Regions List. To handle this, for every playlist entry there will exist a corresponding region of the exact same length when a wave file is created by Sound Forge. This is to simplify the overall use and interface to the user. If we were to have a playlist which contained our previous example of a region of -Beat 1 playing 4 times, it would be stored in the playlist with a name field of 0x0001, a length of 1,000 samples, and a loop count of 4.

APPENDIX H: TIPS AND IDEAS

PSEUDO-QUANTIZING DIGITAL AUDIO WITH MIDI SEQUENCERS

There are many ways to use Sound Forge with MIDI sequencers. The following two tips deal with drum patterns, though any type of material could be adapted to these techniques. These examples are some of many possible uses for the MTC sync and MIDI triggering capabilities of Sound Forge and MIDI sequencers.

USING MIDI NOTE TRIGGERING TO PSEUDO-QUANTIZE REGIONS

1. Record some drum tracks into Sound Forge that are in a consistent tempo and measure scheme.
2. Split up the sound file into regions that are a measure in length. In the Regions List, assign a different MIDI Note On trigger to each region.

3. Select a region and open the *Edit Tempo* dialog in Sound Forge's Special menu. By selecting *Measure*, you can now see what the approximate tempo (BPM) is according to the number of beats you have entered and the length of the selection. Examine a number of different regions in this way and take an average BPM.
4. Now, set the sequencer's tempo to match the BPM found in Sound Forge's *Edit Tempo*.
5. Next, create a MIDI track with whole notes that correspond to the note numbers selected in the Sound Forge *Regions List* for each separate region. These notes should fall on downbeats of the measures so there is one note per measure. It is a good idea to make their lengths a few clock pulses shy of the next downbeat. Also, it may be necessary to set a negative time offset for the whole track to allow for any slight delay that MIDI note triggering is prone to have.
6. Now, playing back the track in the sequencer, while triggering Sound Forge, should give you a pseudo-quantized version of the sound file. Each new measure in the sequencer should trigger a measure of audio material in Sound Forge. Although this method has potential for less accurate start points than MTC SMPTE triggering, it will allow for easier tweaking and alignment within the sequencer.

USING MTC SYNCHRONIZATION TO PSEUDO-QUANTIZE THE PLAYLIST

1. Record some drum tracks into Sound Forge that are in a consistent tempo and measure scheme.
2. Split up the sound file into regions in the *Regions List* and *Playlist* that are a measure in length.
3. Select a region and open the *Edit Tempo* dialog in from the Special menu. By selecting *Measure* you can now see what the approximate tempo (BPM) is according to the number of beats you have entered and the length of the selection. Check a number of different regions in this way to find an average BPM.
4. Now, in the sequencer, set the tempo to match the tempo found in Sound Forge. Next, there should be a way (in the sequencer) to see what SMPTE time will correspond to whatever measure downbeat you choose. These

times can be entered as the Sound Forge playlist regions SMPTE start times. This procedure will allow the playlist regions to start at the downbeat of the sequencer measures when Sound Forge is set to follow MTC. You will undoubtedly need to tweak the start times and possibly the sequencer tempo, but should be able to get it all to sound correct.

5. This method of triggering is quite accurate, since Sound Forge knows what region is to be triggered next and can pre-load the data accordingly. It is a little precarious, though, to have to adjust the SMPTE start times every time the sequencer tempo is changed due to tweaking.

REALISTIC MIDI DRUM PARTS WITH MIDI TRIGGERING

If you do a lot of MIDI sequencing, you know how unconvincing and static MIDI drum parts can be, especially if you're not a drummer. Good results can be achieved through the use of triggering Sound Forge regions from a MIDI sequencer to go along with drum parts that have been created in the sequencer. If you have a sampler, here is a procedure that yields quite convincing drum tracks.

1. First, record into Sound Forge some drum tracks/sounds from an audio source that has an abundance of drums-alone material.
2. Run *Auto Region* on the sound file to automatically breakup the events into easily manipulated regions. To audition each region, open the *Regions List* and click on the *Play* button for each entry.
3. Create some MIDI drum tracks in your sequencer without worrying too much about fills, flams, etc. These things will be added by MIDI triggering Sound Forge.
4. Now that you have a basic drum track in the sequencer, go to Sound Forge and choose events that are distinctly drummer-esque from your *Regions List* (you might need to edit the lengths and positions of the regions). These are the types of events that in themselves do not make good single drum samples, or are too long to reasonably use with a sampler with limited amounts of RAM (for example: drum fills, flams, long crash cymbals and events that have a certain amount of human variation).
5. For each region you want to use, select a *MIDI Trigger* event in the *Edit Marker/Region* dialog. It is recommended that you use *MIDI: Note On - Queue*

/ Note Off - Play, as this will ensure that the timing will be accurate. The only drawback to this method of triggering is that the actual the start of the region playback will occur on the *release* of the MIDI note in the sequencer. That is, the place in time the triggered region should start playing will have to line up to the end of the MIDI event. In some sequencers, this is a little easier to deal with than others. This is a way for Sound Forge to pre-load the region data for accurate playback at the trigger point. See the Using Sound Forge section on MIDI triggering for a detailed explanation of setting up Sound Forge MIDI triggers.

6. Now, in your sequencer, assign a new track with the same MIDI channel as specified in the MIDI Triggers for the regions. At this point each of your drum regions in Sound Forge can be triggered from your sequencer with MIDI notes. You can decorate your drum parts with these events to add variation and drummer-isms that are not possible with single drum samples and MIDI alone.

RANDOMIZING PLAYLIST REGIONS

Some very interesting rhythmic patterns can be discovered by using the region scrambler (a.k.a. the Egg Beater) technique on a drum track outlined below:

1. Open a file with a drum pattern. Run the *Auto Region* tool using the *Drum beats* preset (minor tweaks may be necessary).
2. In the *Playlist* window, randomly rearrange the regions. The fastest (and most fun) way to do this is to randomly drag and drop each region up and down the *Playlist*. Do this for about 10 seconds or until the regions are completely scrambled.
3. Now, listen to the rearranged regions in the *Playlist*.
4. Another thing you can do is to raise or lower the play count in some regions. You can do so by using the *Numpad +* and *Numpad -* key.
5. Repeat steps 2-4 until you're positive you've created the most complex rhythmic pattern known to man.
6. Create a new file from the *Playlist* by using the *Convert to New* function. Now, play your new drum loop in *Looped* mode. Tap along to your new groove and try to figure out its time signature.
7. Cool, eh?

APPENDIX I: OBJECT LINKING AND EMBEDDING (OLE)

OLE support is only available in the 32-Bit version of Sound Forge.

NOTE

Sound Forge supports Microsoft's Object Linking and Embedding (OLE) technology for improved integration with other applications. This section discusses the implementation details and different uses for the OLE technology supported by Sound Forge. The following topics are covered:

- Introduction to Object Linking and Embedding (OLE)
- Using OLE with Sound Forge
- Tips, Limitations, etc.

INTRODUCTION TO OBJECT LINKING AND EMBEDDING (OLE)

Object Linking and Embedding (OLE) is a technology developed by Microsoft to allow independent applications to behave as though they are tightly integrated. Note that this is not an all-encompassing definition, but it will work for the implementation provided by Sound Forge. If you would like a more in depth look at OLE, there are many references available at most bookstores that discuss the technology.

By using the OLE support in Sound Forge, you can embed digital audio files used by Sound Forge into documents used by other applications (such as *Write*, *Microsoft Word* and *Excel*). Once a digital audio file is embedded into a document, you can then edit the audio with Sound Forge at any time by activating the embedded object (usually by double-clicking on it with the mouse).

USING OLE WITH SOUND FORGE

Sound Forge is an OLE server. What this means is Sound Forge can create OLE objects (serve them up if you will) that can be embedded (inserted) into applications that are OLE containers (the consumers of objects supplied by servers). Don't let the terms scare you. We'll go through some examples that will make everything clear as mud shortly.

The OLE objects that are created by Sound Forge are called Sound Forge Audio objects. These objects are really just links (oh no, another term!) to a sound file on your hard disk or network server. Therefore, when you embed a Sound Forge Audio object into *Write*, you are really only inserting a link (a directory and file name) to the actual sound file. We will explain why this is important to know in the *Tips, Limitations, etc.* section.

There are three methods that can be used to create Sound Forge Audio objects. When we mention an OLE container, we simply mean an application that can contain OLE objects (like *Microsoft Word* and *Excel*). The three methods are as follows:

1. Dragging and dropping a link from the OLE Drag Source control in Sound Forge to an OLE container.
2. Using the *Copy Object Link* command in the *Edit* menu of Sound Forge to place an object link onto the clipboard. This object link can then be pasted into an OLE container, usually through the *Paste Special* menu.
3. Invoking the *Insert Object* command from an OLE container and selecting the Sound Forge Audio object from the list.

DRAG AND DROP



The most convenient method of embedding a Sound Forge Audio object into a document is to drag a link from the OLE Drag Source control in Sound Forge and receive from the link on the document that you desire. Note that the application that you drop the OLE object on must be OLE 2.0 compatible.

Follow these steps to embed a Sound Forge Audio object into Microsoft WordPad using drag and drop:

1. Run WordPad.
2. Now run Sound Forge and open a sound file (TUTOR1.WAV in your Sound Forge directory will suffice).
3. Ensure that at least a portion of WordPad is visible while Sound Forge is open (if either application is maximized, you will need to restore them and resize them accordingly).
4. Locate the OLE Drag Source control in the lower right corner of the data window. It's on the right end of the selection length box and contains the letters OLE.
5. Click and hold the left mouse button on the OLE Drag Source control and then drag the mouse cursor while holding the mouse button down.
6. Now simply drag the mouse over the top of WordPad and drop it (release the mouse button). If you did everything correctly, a Sound Forge Audio object will now be embedded in your document.

A newly created sound file must first be saved before you can drag and drop the object into an OLE container. If the sound file in question needs to be saved, you will be prompted to do so when you click on the OLE Drag Source control.

There is an option in the *General* page of the *Preferences* folder to turn off the OLE Drag Source control (hide it). It may be desirable to turn this control off if you do not plan to use the OLE drag and drop support of Sound Forge.

COPY OBJECT LINK

Another method for creating a Sound Forge Audio object is to invoke the *Copy Object Link* command under the Edit menu in Sound Forge. This method is sometimes preferable to the drag and drop method because you don't need a mouse.

Follow these steps to embed a Sound Forge Audio object into *WordPad* using the *Copy Object Link* command:

1. Run *Sound Forge* and open a sound file (TUTOR1.WAV in your *Sound Forge* directory will suffice).
2. Select the *Copy Object Link* command under the Edit menu of *Sound Forge*. If the sound file needs to be saved, you will be prompted to do so.
3. At this point, a *Sound Forge Audio* object is on the clipboard and ready to be pasted into a document.
4. Run *WordPad* and select the *Paste* command from the Edit menu. If you did everything correctly, a *Sound Forge Audio* object will now be embedded in your document.

INSERT OBJECT

The last method for creating a *Sound Forge Audio* object is to use the *Insert Object* command from another application. This method provides a convenient way to get a new *Sound Forge Audio* object embedded into a document without having to first create it in *Sound Forge* and copy it into the document.

Follow these steps to embed a *Sound Forge Audio* object into *WordPad* using the *Insert Object* command:

1. Run *WordPad*.
2. Select the *Insert Object* command. In *WordPad*, use the *Object* option under the *Insert* menu.
3. From the *Object Type* list, select the *Sound Forge Audio* item and click *OK*. Note that *WordPad* has two options: *Create New* and *Create From File*. You will need to select the *Create New* option.
4. At this time, *Sound Forge* will run automatically to create the new object. A new window is automatically opened allowing you to copy or create any new sound data you wish. Note that you must create the sound data in the window supplied.
5. After creating any desired sound data, simply save the modifications and close *Sound Forge*. If you did everything correctly, a *Sound Forge Audio* object will now be embedded in your document.

After the Sound Forge Audio object has been created and embedded into the document, you can edit it at any time using Sound Forge by double-clicking on the object in your document.

TIPS, LIMITATIONS, ETC.

As stated earlier, Sound Forge Audio objects are implemented as links to existing sound files on your hard disk or network server. Because of this, there are some limitations you should be aware of before using the OLE support in Sound Forge.

LINKS CAN BE BROKEN, BUT NOT MENDED

The most important rule to bear in mind when using OLE with Sound Forge is that links can be broken, but not mended. When you create a Sound Forge Audio object, it stores only a link (which is simply a directory path) to the sound file. Sound Forge uses links because sound files tend to be very large and would usually not be acceptable to store along with your documents.

Because Sound Forge uses object links, moving, renaming or deleting the sound file associated with a Sound Forge Audio object will break the link. Attempting to edit the object in your document will fail because Sound Forge will not be able to locate the correct sound file.

To replace a broken link, simply delete the Sound Forge Audio object from the document and create a new one using the correct sound file.

If you need to move a sound file that is linked to a document, you can use Sound Forge to save the file to a different location and the object link will automatically be updated. However, you must open the sound file that you want to move as an object for this feature to work. To open a sound file as an object, simply double-click on the Sound Forge Audio object in your document; do not open the sound file directly. You can then use *Save As* to move the sound file. Sound Forge will prompt you about updating the object link after saving the new file.

SOUND FILES DO NOT TRAVEL WITH YOUR DOCUMENTS

When you copy a document that contains one or more Sound Forge Audio objects to another computer, the sound files that are linked to the objects are not copied with the document. If you need to copy a document to another

computer, you will need to copy all sound files associated with it as well. However, you must also ensure that the sound files have the exact same path on the new computer as they did on the old computer.

USING A NETWORK SERVER

To ease the problems associated with moving or sharing documents between computers, it is suggested that you use a network server to store all sound files that are embedded in the documents. If you do this, all computers that have access to the network server will be able to open and edit the Sound Forge Audio objects.

There are two methods of ensuring a common network server will work with shared documents. The first is to make sure you use the same drive letter when connecting to the network server on all machines. Remember that the full path must remain constant between computers.

The second and more preferable method is to use *Universal Naming Convention (UNC)* paths when embedding Sound Forge Audio objects into your documents. A UNC name looks like this:

```
\\server\share\<directory and file name>
```

Therefore, if your sound files are stored on `\\myserver\sounds`, open them in Sound Forge using UNC path names rather than a drive letter. You can create the Sound Forge Audio object using either the *Copy Object Link* command or the *OLE Drag Source* control and embed the sound files into your documents. By using UNC names, you eliminate the need to use a common drive letter on all computers that share the document.

If you are running Windows 95, UNC path names are extremely simple to use. Instead of using a drive letter when browsing in the *File Open* dialog, select *Network Neighborhood* in the *Look In* field and browse to the appropriate network server and share containing your sound files. This method of browsing uses UNC names exclusively.

SOUND FORGE IS NOT AN OLE CONTAINER

Although Sound Forge can create OLE objects to embed in other applications, the inverse is not true. Therefore, OLE objects created by applications such as *Microsoft Word* or *Excel* cannot be embedded into sound files open in Sound Forge.

APPENDIX J: USING CSOUND, MTU, IRCAM, BICSF AND EBICSF FILES

This section discusses how to open and save files of the type CSOUND, MTU, IRCAM, BICSF, EBICSF with Sound Forge.

INTRODUCTION

Although Sound Forge supports a large number of sound file formats directly, it does not support the CSOUND, MTU, IRCAM, BICSF or EBICSF file types. However, by using the Raw File Type capabilities of Sound Forge, sound data from these file types can be extracted.

The IRCAM or IRCAM-Gross format consists of a 1024 byte header prior to the audio data. This header contains standard information like the number of channels, sampling rate, and data format. It also can contain information such as the name of the sample and comments on what is included in the file. This is the format used by the MTU system and is often referred to as MTU files. IRCAM files support two types of data formats: 16-bit linear PCM and floating point data.

BICSF or EBICSF files (Berkeley/IRCAM/CARL Sound File or Extended BICSF) are an extension of the IRCAM format. The difference being that instead of using the standard IRCAM header they replace the first 28 bytes of the header with a standard NeXT/Sun header. This allows the IRCAM format to continue to hold additional information in its 1024 byte header and still allow the files to be read by software that supports the NeXT/Sun file format (such as Sound Forge).

OPENING AND SAVING FILES

BICSF AND EBICSF FILES

When reading BICSF or EBICSF Sound Forge will identify them as NeXT/Sun files. This is due to the fact that the header of the BICSF file has been modified to allow it to be read as a NeXT/Sun File. Sound Forge has no difficulty reading these files as long as they are in one of the supported NeXT/Sun data formats.

IRCAM, CSOUND AND MTU FILES

When reading any of these types you must import the files as a Raw data file. The best way to do this is to set up the correct parameters in the Sound Forge Raw File Type dialog box and then save presets labeled accordingly. Sound Forge's Raw File import will allow you to read these files as long as they are mono or stereo and are stored in 16-bit linear format. Sound Forge will not allow you to read the floating point format IRCAM files.

OPENING AN IRCAM FILE

To open an IRCAM sound file in Sound Forge, follow these steps:

1. Open the *File Open* dialog box.
2. Select the *Raw File* type.
3. Select an IRCAM file to load.
4. Press the *OK* button in the *File Open* dialog. This will bring up the *Raw File Type* dialog box.
5. Select the following parameters:

Sample rate: Set the most common sample rate you use (typically 44,100).

Sample type: 16-bit PCM

Format: Signed

Channels: Mono or Stereo (depends on the file)

Byte order: *Big endian* or *Little endian* (depends on where the file was generated).

Header: 1,024 bytes

Trailer: 0 bytes

6. Before selecting *OK* to load the file press the *Save As* button to create a preset. Type in an appropriate preset name and press *OK*.

Now whenever you wish to load an IRCAM file, you can load it as *Raw File Type* and then select the preset you saved to set all the necessary parameters to load the file. Note that if any of the files are not the same as your preset, you should change them in the raw dialog box. You may even want to create multiple presets.

It is important to note that the *Byte order* of files is not always the same for files generated by CSOUND. The CSOUND executables for the PC generate files which use Little Endian byte ordering. CSOUND for other platforms tends to generate files with Big endian ordering. MTU files also use Big endian ordering. The best policy is to try a file first in Big endian. If the sound data seems to be garbage, then try loading the file using Little endian. Again you may want to save presets for byte ordering, as well as mono/stereo, if you are receiving CSOUND files from a variety of other computers.

SAVING FILES WITH SOUND FORGE

When saving files of the above types you will not be able to save them as IRCAM, CSOUND or MTU. You will have to save any files of these types as one of the file formats supported by Sound Forge.

If you wish to save a file for use with software which supports the BICSF/EBICSF format you should save the file as a NeXT/Sun file which will not save any of the additional information found in a BICSF/EBICSF file but should allow the software to read the data as a NeXT/Sun file.

APPENDIX K: SAMPLE TRANSFER PROCEDURES

AKAI S1000

MIDI/SDS

The S1000 is capable of sending and receiving samples via MIDI/SDS in all methods implemented by Sound Forge. Provided that your MIDI connections are correct, the S1000 is basically ready to receive or send a sample. However, there are a few things to keep in mind.

The S1000 structures its sample memory in such a way that the *Logical send/receive sample number* in Sound Forge can be a little confusing. The problem is that the S1000 will name the newly received sample according to the *Logical send sample number* specified in Sound Forge. However, in the S1000, a new sample is placed into the next available memory location, not the logical send location.

For example, if you send the S1000 a sample with a *Logical send sample number* of 0, send another with the *Logical send sample number* of 2, and then send yet another with the *Logical send sample number* of 1, the S1000 will now have three samples next to each other in memory. One will be named MIDI 0, another named MIDI 2, and another named MIDI 1. If you want to use Sound Forge's *Get Sample* to retrieve these from the S1000, MIDI 0 would have to be referred to as *Logical receive sample number* 4, MIDI 2 would be *Logical receive sample number* 5, and MIDI 1 would be referred to as *Logical receive sample number* 6. What's going on here?

First of all, the S1000 has four waveforms permanently stored in memory location 0-3; sine, square, sawtooth, and pulse. The S1000 doesn't necessarily refer to these waveforms as being at these locations. We are referring to them as such because if you want to *Get Sample* from Sound Forge, these are the *Logical receive sample numbers* you would use to get these waveforms. The S1000 is simply taking the *Logical send sample number* and using it as a naming

convention for received samples. However, the *Logical send sample number* has no bearing on where it puts the sample.

When the S1000 receives a new sample from Sound Forge, it puts the sample into its next available memory location, deriving its name from the *Logical send sample number*. In our example, the S1000 put our first sample sent in location number 4 and named it MIDI 0, the next one in location number 5 as MIDI 2, and finally put MIDI 1 in location number 6. So bear in mind that when invoking *Get Sample* in Sound Forge, the *Logical receive sample number* does not necessarily correspond to the name of the sample in the S1000.

The S1000 will replace any sample in its memory with an incoming sample of the same name. That is, if you have sent a sample into the S1000 with *Logical send sample number 2* and the S1000 has named this MIDI 2, then if you send another sample and do not change the *Logical send sample number*, the original MIDI 2 will be replaced with a new MIDI 2.

In order to audition the sample from a MIDI source (such as the *Sound Forge MIDI Keyboard*), it will be necessary to go to the S1000 EDIT SAMPLE page and select the sample that you want to hear. You will need to be on this page in order to audition a sample that has not yet been incorporated into a program. While the S1000 is receiving the sample, the MIDI LED will flicker indicating that the transfer is in process.

The S1000 supports SDS handshaking. Therefore, any method of MIDI/SDS transfer will work. It is recommended that you turn the *Open loop* option off and *Send request* on. Be sure that the MIDI In and Out are connected between the S1000 and Sound Forge to use handshaking (closed loop).

On the S1000's MIDI EXCL page, set the type of transmission to SINGLE SAMPLE and sample protocol to STANDARD.

TO TRANSFER MIDI/SDS TO THE S1000 FROM SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within Sound Forge. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Send to* section, select Akai S1000 (MIDI/SDS).
4. In the *MIDI output* section choose the MIDI driver that is connected the S1000's MIDI input.

5. Now select the *MIDI channel*. (This should match the MIDI channel found in the S1000's MIDI/EXCL page).
6. Hit OK.
7. Hit *Send Sample*. The S1000 should now receive the sample.

TO TRANSFER MIDI/SDS TO SOUND FORGE FROM THE S1000

1. Having set the necessary *Start/Sustaining loop/Release loop* points within the S1000. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Receive from* section, select Akai S1000 (MIDI/SDS).
4. In the *MIDI output* section choose the MIDI driver that is connected to the S1000's MIDI input.
5. In the *MIDI input* section choose the MIDI driver that is connected to the S1000 MIDI output.
6. We recommend that you select *Send request (MIDI/SDS)*.
7. Now select the *MIDI channel*. (This should match the MIDI channel found in the S1000's MIDI/EXCL page).
8. Hit OK.
9. Hit *Get Sample*. Sound Forge should now receive the sample.

DIGIDESIGN SAMPLECELL II

TRANSFER TO SAMPLECELL

The SampleCell II is capable of receiving samples via internal interaction between Sound Forge and the SampleCell software. There are a few things to keep in mind.

You must have SampleCell software version 1.02 or above to transfer samples from Sound Forge to your SampleCell II card. If you have an earlier version contact Digidesign to acquire an update.

The SampleCell can only receive samples from Sound Forge. You can't retrieve samples from the internal memory of the SampleCell. This is usually not a

problem since the SampleCell does not allow you to record new samples. All samples in its memory are supplied from files you already have on your computer.

Sending a sample to the SampleCell from Sound Forge will replace any previous samples sent by Sound Forge.

In order to audition the sample from a MIDI source (such as the *Sound Forge MIDI Keyboard*), it will be necessary to set the MIDI output device to SampleCell 1. If you have more than one SampleCell the number will be adjusted accordingly for each card in your system.

You must be running the SampleCell application software to transfer samples from Sound Forge to the SampleCell.

All samples sent to the SampleCell from Sound Forge have the name FRGSCNEW.

The SampleCell ignores the *Sample number* and *Sample bias* fields used by Sound Forge.

TO TRANSFER SAMPLES FROM SOUND FORGE

1. Start the SampleCell software by double-clicking on the Sample Cell icon.
2. Having set the necessary *Start/Sustaining loop/Release loop* points within Sound Forge. Select *Sampler* from the Tools menu.
3. Select *Configure*.
4. In the *Send to* section, select *Digidesign SampleCell II*. You do not have to set any other options in the dialog. They are ignored by the SampleCell.
5. Hit *OK*.
6. Hit *Send Sample*. The SampleCell should now receive the sample.

When transferring samples to the SampleCell, a temporary .WAV patch file is created prior to sending the data. You should not save instruments or patches based on files transferred from Sound Forge to the SampleCell, as the SampleCell bases its instruments on the original .WAV file name. The downloading of samples from Sound Forge is meant to be used to create and tune single .WAV sounds within an instrument. If you save an instrument based on a temporary file created by Sound Forge, when you attempt to reload

the instrument at a later time, the SampleCell software will not be able to find the temporary file.

Do as much work as you can in Sound Forge when creating samples for the SampleCell. Download the sample to the SampleCell to listen to how it sounds. If you are happy with your work, save the file using the Save As item from the File menu, and then load the new .WAV file from SampleCell software to create new instruments.

KURZWEIL K2000

MIDI/SDS

The K2000 is capable of receiving a sample via MIDI/SDS. The K2000 will not send a sample via MIDI/SDS. In its manual there is mention of the K2000 bulk dump. This protocol is not supported by Sound Forge. Therefore, the Sound Forge Get Sample procedure does not work with the K2000 for MIDI/SDS transfers.

Sending samples to the K2000 via MIDI/SDS is really quite simple. Provided that your MIDI connections are correct, the K2000 is basically ready to receive a new sample. The K2000 can be sitting on any screen and will actually operate normally while the transfer is going on in the background. There are a few things to keep in mind.

The K2000 will automatically assign the incoming sample to the next available sample/RAM location if the *Logical send/receive sample number* is left at 0. This location is generally going to be in the 200-299 range (unless these are already filled). If you choose a *Logical send/receive sample number* of 0-199, the K2000 will automatically add 200 to the selected number (plus any necessary to reach the next available sample number). If you choose a number of 200 or more, the K2000 will add 1 (again, plus any necessary to reach the next available sample number).

In order to audition the sample from a MIDI source (such as the *Sound Forge MIDI Keyboard*), it will be necessary to create a Keymap within the K2000. This is done by hitting Preview in the Master/Sample page, or by calling up the sample within the Keymap section of a pre-existing program.

While the K2000 is receiving the sample, the MIDI LED will flicker, indicating that the transfer is in process.

TO TRANSFER MIDI/SDS TO THE K2000 FROM SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within Sound Forge. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Send* to section, select Kurzweil K2000 (MIDI/SDS).
4. In the *MIDI output* section choose the MIDI driver that is connected the K2000's MIDI input.
5. Now select the *MIDI channel*. (This should match the Basic MIDI channel found in the K2000's MIDI/RECV page).
6. Hit OK.
7. If you now hit *Send Sample*, the K2000 will start receiving the data.
8. To examine the new sample within the K2000, simply go to the *Master/Sample* page of the K2000. The sample that appears there as a default should be your newly received sample. Hit *EDIT*, and look to see that the waveform appears in the K2000's screen. Hitting *preview* from the *Master/Sample* page will create a new program based on this sample and will allow you to audition the sample from an external MIDI source.

SCSI/SMDI

Sound Forge is capable of sending and receiving samples via SCSI/SMDI to the Kurzweil K2000. This method is very fast and very simple once you have configured the SCSI hardware properly. Bear these few things in mind.

Unlike in MIDI/SDS sample transfers, the K2000 will not automatically adjust for sample number locations (adding 200, etc.). There are two ways of handling this fact;

1. You can select the specific *Logical send/receive sample number* as you would like it to appear in the K2000. Since the K2000 will add 1 to whatever number you have entered here, set the *Sample bias* in the *Sampler Configuration* dialog to -1.

2. If you set the *Sample bias* to 199, you can then refer to the *Logical send/receive sample number* as 0 through x. This will make the Sound Forge send of sample number 0 locate to the K2000 at 200, sample number 1 locates to 201 and so on.

When sending a sample to a K2000 RAM location that is already filled, the K2000 will not warn you that the location is already filled. Instead, it will replace the sample with the new one being sent. Be very careful when doing large numbers of sample transfers (SCSI/SMDI), so that you do not wipe out your hard work by neglecting to select a new *Logical send/receive sample number* for each separate sample.

When choosing SCSI/SMDI in the *Sampler* dialog, Sound Forge will verify that the SCSI bus is operational and may therefore take a second or two to respond to any user interaction.

When the SCSI/SMDI method is selected in the *Sampler Configuration* dialog, you only need to be concerned with setting the SCSI host and SCSI sampler. Any SCSI device that appears in these lists is being recognized by Sound Forge and should work correctly.

Make sure that your SCSI adapter is supported by Sound Forge for sample transfer procedures. If you are unsure as to whether your adapter is supported for SMDI transfers with Sound Forge, see *SCSI/SMDI Setup and Troubleshooting* at the beginning of this section.

The K2000's SCSI ID is set in its *MIDI/RECV* page.

TO TRANSFER SCSI/SMDI TO THE K2000 FROM SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within Sound Forge. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Send to* section, select *Kurzweil K2000 (SCSI/SMDI)*.
4. In the *SCSI host* section, select your SCSI adapter.
5. In the *SCSI sampler* section, choose the K2000 connected to your computer's SCSI bus.
6. Save your configuration if you like and/or hit OK.

7. Check to see that your *Configuration* settings are confirmed without error and now hit *Send Sample*. The K2000 will receive the sample very quickly.
8. To examine the new sample within the K2000, go to the *Master/Sample* page of the K2000. The sample that appears there as a default should be your newly received sample. Hit *EDIT*, and check to see that the waveform appears in the K2000's screen. Hitting *preview* from the *Master/Sample* page will create a new program based on this sample and will allow you to audition the sample from an external MIDI source.

TO TRANSFER SCSI/SMDI FROM THE K2000 TO SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release* loop points within the K2000. Select *Sampler* from the *Tools* menu.
2. Select *Configure*.
3. In the *Receive from* section, select *Kurzweil K2000 (SCSI/SMDI)*.
4. In the *SCSI host* section, select your *SCSI* adapter.
5. In the *SCSI sampler* section, choose the K2000 connected to your computer's *SCSI* bus.
6. Save your configuration if you like and/or hit *OK*.
7. Check to see that your *Configuration* settings are confirmed without error and hit *Get Sample* (remember to select the specific *Logical send/receive sample number* minus one, unless you have adjusted the *Sample bias*). *Sound Forge* will receive the sample very quickly.
8. When the transfer is complete, the sample will appear in the active sample window.

KNOWN PROBLEMS WITH THE K2000

The following is a known problem with transferring samples to K2000:

The K2000 reboots when in the *Master/Sample* page. This is a bug in the K2000's ROM. The K2000 will reboot if you transfer a sample to the same sample number twice in a row while the K2000 is on the *Master/Sample* page. The best way to get around this problem is to not transfer samples while you are using the *Master/Sample* page of the K2000.

PEAVEY SP

MIDI/SDS

The Peavey SP is capable of sending and receiving samples via MIDI/SDS in all configuration methods implemented by Sound Forge.

The SP is basically ready to receive a new sample provided that your MIDI connections are correct. The SP can be sitting on any screen and will begin to receive the sample upon initiating the transfer from Sound Forge. Keep the following things in mind:

- Make sure the SP's Get slave samp# is set to using: midi (sds) – this is located in the SP's Util menu.
- If you have set up the transfer configuration as a *Wait for request*, you will need to go to the SP's Util menu and instigate the request for sample dump. This is done by hitting Exec twice from the Util menu, on the SP, after telling Sound Forge to *Send Sample*.
- In order to audition the sample from a MIDI source (such as the *Sound Forge MIDI Keyboard*), it will be necessary to go to the wave menu on the SP and select the appropriate wave number that you are trying to audition (the last received sample should appear on this screen).
- If you are sending a sample using *Open loop*, it is not required that the sampler's MIDI out be connected to Sound Forge. However, it is recommended that you use *handshaking (no Open loop)* for faster and more reliable communications.

TO TRANSFER MIDI/SDS TO THE PEAVEY SP FROM SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within Sound Forge. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Send to* section, select *Peavey SP (MIDI/SDS)*.
4. In the *MIDI output* section choose the MIDI driver that is connected to the SP's MIDI input.

5. Select the *MIDI channel* (this should match the omni/poly channel found in the SP's Global menu, unless the SP is set to omni mode).
6. In Sound Forge hit the *Send Sample* button. The transfer will begin and both Sound Forge and the SP will display the amount of data left to go until the procedure is complete.
7. When the transfer is complete, you can hit the wave button on the SP and immediately audition the sample from the *Sound Forge MIDI Keyboard* or whatever MIDI source is connected to the SP.

TO TRANSFER MIDI/SDS FROM THE PEAVEY SP TO SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within the Peavey SP. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Receive from* section, select *Peavey SP (MIDI/SDS)*.
4. In the *MIDI input* section choose the MIDI driver that is connected to the SP's MIDI output.
5. In the *MIDI output* section choose the MIDI driver that is connected to the SP's MIDI input.
6. Select the *MIDI channel* (this should match the omni/poly channel found in the SP's Global menu, unless the SP is set to omni mode).
7. In order for Sound Forge to automatically get the sample without invoking the transfer from the SP's front panel, it is necessary to turn *Wait for request* on in the *Receive* section of the *Sampler Configuration* dialog. Otherwise, you will have to hit *exec* on the SP after hitting the *Get Sample* button in Sound Forge.
8. In Sound Forge, hit the *Get Sample* button. The transfer will begin and both Sound Forge and the SP will display the amount of data left to go until the procedure is complete.
9. When the transfer is complete, the sample will appear in the active data window.

SCSI/SMDI

Sound Forge is capable of sending and receiving samples via SCSI/SMDI to the Peavey SP. This method is very fast and very simple once you have configured the SCSI hardware properly. Keep the following things in mind.

- Make sure the SP's Get slave samp# is SCSI (SMDI). This is located in the SP's Util menu. Though the SP is capable of detecting what method is used with the incoming sample, it is highly recommended that the user select this manually, as crashes can occur when first sending MIDI/SDS followed immediately by SCSI/SMDI.
- The SCSI cable from your computer must be connected to the SP's SCSI B port.
- When choosing SCSI/SMDI in the *Sampler* dialog, Sound Forge will verify that the SCSI bus is operational and may therefore take a second or two to respond to any user interaction.
- When SCSI/SMDI method is selected, you only need to be concerned with setting the SCSI host and SCSI sampler in the *Sampler Configuration* dialog. Any SCSI device that appears in these lists is being recognized by Sound Forge and should work correctly.
- Make sure that your SCSI adapter is supported by Sound Forge for sample transfer procedures. If you are unsure as to whether your adapter is supported for SMDI transfers with Sound Forge, see SCSI/SMDI Setup and Troubleshooting in the *Sampling* chapter.
- The Peavey SP's SCSI ID is set by default to Device 0. This setting is made by means of jumpers within the unit itself. Every time the SP is powered on, it will reset itself to this setting. You can temporarily change the SCSI ID to another number from the SP's UTIL menu and will remain until the unit is powered off. If you find that you consistently want a SCSI ID other than 0, you should change the jumper settings inside the unit according to the directions in the Peavey SP manual. If you have a bootable SCSI hard drive its ID will also be 0 so you must change the ID of the SP for proper operation.

TO TRANSFER SCSI/SMDI TO THE PEAVEY SP FROM SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within Sound Forge. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Send to* section, select *Peavey SP (SCSI/SMDI)*.
4. In the *SCSI host* section, select your SCSI adapter.
5. In the *SCSI sampler* section, choose the Peavey SP connected to your computer's SCSI bus.
6. Save your configuration if you like and/or hit *OK*.
7. Check to see that your *Configuration* settings are confirmed without error and hit *Send Sample*. The Peavey SP will receive the sample very quickly.
8. When the transfer is complete, you can hit the wave button on the SP and immediately audition the sample from the *Sound Forge MIDI Keyboard* or whatever MIDI source is connected to the SP.

TO TRANSFER SCSI/SMDI FROM THE PEAVEY SP TO SOUND FORGE

1. Having set the necessary *Start/Sustaining loop/Release loop* points within the Peavey SP. Select *Sampler* from the Tools menu.
2. Select *Configure*.
3. In the *Receive from* section, select *Peavey SP (SCSI/SMDI)*.
4. In the *SCSI host* section, select your SCSI adapter.
5. In the *SCSI sampler* section, choose the Peavey SP connected to your computer's SCSI bus.
6. Save your configuration if you like and/or hit *OK*.
7. Check to see that your *Configuration* settings are confirmed without error and hit *Get Sample*. Sound Forge will receive the sample very quickly.
8. When the transfer is complete, the sample will appear in the active sample window.

This page was intentionally left blank.

GLOSSARY

A-Law

A-Law is a companded compression algorithm for voice signals defined by the Geneva Recommendations (G.711). The G.711 recommendation defines A-Law as a method of encoding 16-bit PCM signals into a non-linear 8-bit format. The algorithm is commonly used in United States' telecommunications. A-Law is very similar to μ -Law, however, each uses a slightly different coder and decoder.

Adaptive Delta Pulse Code Modulation (ADPCM)

ADPCM is a method of compressing audio data. Although the theory for compression using ADPCM is standard, there are many different algorithms employed. As an example, Microsoft's ADPCM algorithm is not compatible with the International Multimedia Association's (IMA) approved ADPCM.

Aliasing

Aliasing is a type of distortion that occurs when digitally recording high frequencies with a low sample rate. For example, in a motion picture, when a car's wheels appear to slowly spin backward while the car is quickly moving forward you are seeing the effects of aliasing. Similarly when you try to record a frequency greater than one half of the sampling rate (the Nyquist Frequency), instead of hearing a high pitch you may hear a low frequency rumble.

To prevent aliasing, an anti-aliasing filter is used to remove high-frequencies *before* recording. Once the sound has been recorded, aliasing distortion is impossible to remove without also removing other frequencies from the sound. This same anti-aliasing filter must be applied when resampling to a lower sample rate.

Amplitude Modulation

Amplitude Modulation (AM) is a process whereby the amplitude (loudness) of a sound is varied over time. When varied slowly, a tremolo effect occurs. If the frequency of modulation is high, many side frequencies are created which can strongly alter the timbre of a sound.

Analog

When discussing audio, this term refers to a method of reproducing a sound wave with voltage fluctuations that are analogous to the pressure fluctuations of the sound wave. This is different from digital recording in that these fluctuations are infinitely varying rather than discrete changes at sample time. *See Quantization.*

Attack

The attack of a sound is the initial portion of the sound. Percussive sounds (like drums, piano, guitar plucks) are said to have a fast attack. This means that the sound reaches its maximum amplitude in a very short time. Sounds that slowly swell up in volume (soft strings and wind sounds) are said to have a slow attack.

Audio Compression Manager (ACM)

The Audio Compression Manager (ACM) from Microsoft is a standard interface for audio compression and signal processing for Windows. The ACM can be used by Windows programs to compress and decompress .WAV files, as well as apply DSP algorithms to audio data.

Audio Event Locator

The Audio Event Locator is similar to a scrub function. However, rather than playing the sound file at a slow speed, it loops playback around the cursor position. This position can be selected by dragging the cursor around in the Sound Forge overview window.

Bit

A bit is the most elementary unit in digital systems. Its value can only be 1 or 0, corresponding to a voltage in an electronic circuit. Bits are used to represent values in the binary numbering system. As an example, the 8-bit binary number 10011010 represents the unsigned value of 154 in the decimal system. In digital sampling, a binary number is used to store individual sound levels, called samples.

Bit Depth

Bit Depth refers to the number of bits used to represent a sample. For example 8-bit and 16-bit are two common bit depths. Sound Forge uses the term Sample Size when referring to bit depth.

Byte

A byte refers to a set of 8 bits. An 8-bit sample requires one byte of memory to store, while a 16-bit sample takes two bytes of memory to store.

CODEC

Acronym for Coder/Decoder. Commonly used when working with data compression.

Channel Converter

The Channel Converter is a function in Sound Forge that converts files from mono to stereo and stereo to mono with independent level control of the new channels. This function can also create interesting effects by converting stereo files to stereo with various levels and inversion of channels.

Chorus

Chorus is an effect created by combining a signal with a modulating delayed copy of itself. This effect creates the illusion of multiple sources creating the same sound.

Clipboard

The clipboard is where sample data is saved when you cut or copy it from a data window. You can then paste, mix, or crossfade the sample data stored on the clipboard with another data window. This sample data can also be used by other Windows applications that support sound data on the clipboard, such as Sound Recorder.

Clipping

Clipping occurs when the amplitude of a sound is above the maximum allowed recording level. In digital systems, clipping is seen as a clamping of the data to a maximum value, such as 32,767 in 16-bit data. Clipping causes sound to distort.

Convert to 8-Bit

Converting a 16-bit sample to 8-bit is simply a matter of what to do with the extra 8 bits in the sample word. The three available options are truncation, rounding and dithering. These options are all explained in the Reference chapter under Convert to 8-Bit in the Process menu.

Copy Object Link (OLE)

A Copy Object Link command can be found in the Edit menu. This command allows you to create an Object Linking and Embedding (OLE) compatible link to the active sound file in Sound Forge for pasting into OLE compatible applications.

Crossfade Loop

Sometimes a sample loop cannot be easily created from the given source material. In these instances a crossfade can be applied to the beginning and end of the loop to aid in the smooth transition between the two. The Crossfade Loop function provides a method of creating sampling loops in material that is otherwise difficult to loop.

Cut-off Frequency

The cut-off frequency of a filter is the frequency at which the filter changes its response. For example, in a low-pass filter, frequencies greater than the cutoff frequency are attenuated while frequencies less than the cutoff frequency are not affected.

DC Adjust

DC Adjust is a function of the Record dialog in Sound Forge. It is the process to compensate for DC offset. *See DC Offset.*

DC Offset

DC offset is a phenomenon that occurs when hardware such as a sound card adds DC current to a recorded audio signal. This current causes the audio signal to alternate around a point above or below the normal -Infinity dB (center) line in the sound file. To visually see if you have a DC offset present you can zoom all the way into a sound file and see if it appears to be floating over the center line. Sound Forge can compensate for this DC offset by adding a constant value to the samples in the sound file.

Data Window

Each opened sound file in Sound Forge has its own data window. On top of each data window is a title bar displaying either the title of the sample or the name of the file. Also in each data window are the waveform display, time and level rulers, playbar and other tools that give you information and allow you to navigate throughout the entire sound file.

Decibel

The decibel (abbreviated as dB) is a unit used to represent a ratio between two numbers using a logarithmic scale. For example, when comparing the numbers 14 and 7, you could say 14 is two times greater than the number 7; or you could say 14 is 6 dB greater than the number 7. Where did we pull that 6 dB from? Engineers use the equation $\text{dB} = 20 \times \log (V1/V2)$ when comparing two instantaneous values. Decibels are commonly used when dealing with sound because the ear perceives loudness in a logarithmic scale.

In Sound Forge, most measurements are given in decibels. For example, if you want to double the amplitude of a sound, you apply a 6 dB gain. A sample value of 32,767 (maximum positive sample value for 16-bit sound) can be referred to as having a value of 0 dB. Likewise, a sample value of 16,384 can be referred to having a value of -6 dB.

Destructive Editing

Destructive editing is the type of editing whereby all cuts, deletes, mixes and other processes are actually processed to the sound file. Anytime you delete a section of a sound file in Sound Forge the sound file on disk is actually re-written without the deleted section. This is different than non-destructive editing. See *Non-destructive Editing*.

Device Driver

A device driver is a program that allows Windows or DOS to connect different hardware and software together. For example, a sound card device driver is used by Windows software to control sound card recording and playback.

Digital Signal Processing (DSP)

Digital Signal Processing (DSP) is a very general term describing anything that alters digital data. Signal processors have existed for a very long time (tone controls, distortion boxes, wah-wah pedals) in the analog (electrical) domain. Digital signal processors alter the data after it has been digitized by using a combination of programming and mathematical techniques. DSP techniques are used to perform many effects such as equalization and reverb simulation.

Since most DSP is performed with simple arithmetic operations (additions and multiplications), both your computer's processor and specialized DSP chips can be used to perform any DSP operation. The difference is that DSP chips are optimized specifically for mathematical functions while your

computer's microprocessor is not. This results in a difference in processing speed.

Direct Edit Mode

Opening a file from the Open dialog and selecting to Operate directly on the sound file box checked, opens the file in direct edit mode. This simply means that no temporary file is created for the file. The first time you open a file in direct edit mode, a peak data file will be created and stored in the directory that the file resides. From then on opening and saving the file will be almost instantaneous regardless of the size of the file. Please note, however, that any edits you make to the file are actually changing the file. With the unlimited undo capability of Sound Forge this shouldn't be a problem, but keep in mind that if the system crashes for any reason, the file will be saved in the state that it is in when it crashes and no undo will be available.

Dithering

Dithering is the practice of adding noise to a signal to mask quantization noise.

Drag and Drop

Drag and drop is a quick way to perform certain operations using the mouse in Sound Forge. To drag and drop, you click and hold a highlighted selection, drag it (hold the left mouse button down and move the mouse) and drop it (let go of the mouse button) at another position on the screen. For more information on where you can use drag and drop, look in Appendix A: Shortcuts.

Dynamic Range

Dynamic Range is the difference between the maximum and minimum signal levels. It can refer to a musical performance (high volume vs. low volume signals) or to electrical equipment (peak level before distortion vs. noise floor). For example, orchestral music has a wide dynamic range while thrash metal has a very small (always loud) range.

Endian (Little and Big)

Little and Big Endian describe the ordering of multi-byte data that is used by a computer's microprocessor. Little Endian specifies that data is stored in a low to high-byte format; this ordering is used by the Intel microprocessors. Big Endian specifies that data is stored in a high to low-byte format; this ordering is used by the Motorola microprocessors.

Equalization (EQ)

Equalizing a sound file is a process by which certain frequency bands are raised or lowered in level. EQ has various uses. The most common use for Sound Forge users is to simply adjust the subjective timbral qualities of a sound.

File Associations

This dialog allows you to associate sound file extensions (such as .WAV, .AU, .SND, etc.) with Sound Forge. This dialog is opened from the File page of the Preferences dialog.

File Format

A file format specifies the way in which data is stored on your floppy disks or hard drive. In Windows, the most common file format is the Microsoft .WAV format. However, Sound Forge can read and write to many other file formats so you can maintain compatibility with other software and hardware configurations.

Frequency Modulation (FM)

Frequency Modulation (FM) is a process whereby the frequency (pitch) of a sound is varied over time. Sub audio frequency modulation results in pitch-bending effects (vibrato). Frequency modulation within audio band frequencies (20 Hz - 20,000 Hz) creates many different side-band frequencies which drastically alter the timbre of the sound.

Frequency Modulation (FM) Synthesis

This type of synthesis relies on the principles of Frequency Modulation. You can use Sound Forge to create fresh sounds with the FM Synthesis tool.

Anything from an organ to whacked-out space trip sounds can be created with this powerful tool.

Frequency Spectrum

The Frequency Spectrum of a signal refers to its range of frequencies. In audio, the frequency range is basically 20 Hz to 20,000 Hz. The frequency spectrum sometimes refers to the distribution of these frequencies. For example, bass-heavy sounds have a large frequency content in the low end (20 Hz - 200 Hz) of the spectrum.

Hertz (Hz)

Hertz is the unit of measurement for frequency or cycles per second (cps).

Insertion Point

The Insertion Point (also referred to as the Cursor Position) is analogous to the cursor in a word processor. It is where pasted data will be placed or other data may be inserted depending on the operation. The insertion point appears as a vertical flashing black line and can be moved by clicking the left mouse button any where in the waveform display of a data window.

InterVoice Sound File Support

The InterVoice sound file format (.IVC), commonly used in telephony applications, includes G.711 μ -Law and A-Law, G.721 ADPCM (32 kb/s) and G.723 ADPCM (24 kb/s) data formats.

Loop Tuner

The Sound Forge Loop Tuner can be found in the View menu. The Loop Tuner shows the beginning and ending of Sample Loops butted together so that a smooth loop can more easily be made. While the loop is playing, you may move either end by zero-crossings or individual samples.

MIDI (Musical Instrument Device Interface)

Musical Instrument Device Interface (MIDI) is a standardized way for music equipment to communicate. Technically, it is an asynchronous serial communications protocol using a 31,250 Hz baud rate. MIDI is used to send control, timing, and data information between devices for triggering notes, time synchronization, and sample dumps.

MIDI Channels

MIDI allows for 16 discrete channels for sending data. When dealing with MIDI triggers, Sound Forge needs to know what MIDI channel to look at for receiving the trigger. The channel this information is sent to in Sound Forge depends on the device sending the MIDI messages.

MIDI Controllers

MIDI Controllers are a specific type of MIDI message. Sound Forge can use MIDI controllers to trigger events and playback of sound files. Consult your MIDI sending device to see what controller messages it sends.

MIDI Notes

MIDI Notes are a specific type of MIDI message. Sound Forge can use MIDI notes to trigger events and playback of sound files. Any MIDI sequencer or controller will send MIDI notes.

MIDI Port

A MIDI Port is the physical MIDI connection on a piece of MIDI gear. This port can be a MIDI in, out or thru. Your computer must have a MIDI port to utilize MIDI triggering in Sound Forge, or to output MIDI Time Code to an external device.

MIDI Time Code (MTC)

MTC is an addendum to the MIDI 1.0 Specification and provides a way to specify absolute time for synchronizing MIDI capable applications. Basically, it is a MIDI representation of SMPTE time code.

Markers

Markers are simply saved locations in the sound file. Markers are stored in the Regions List and can be used for quick navigation.

Media Control Interface (MCI)

The MCI is a standard way for Windows programs to communicate with multimedia devices like sound cards and CD players. If a device has a MCI device driver, it can easily be controlled by most multimedia Windows software.

Media Player

Media Player is a Microsoft Windows program that can play digital sounds or videos using MCI devices. Media Player is useful for testing your sound card setup. For example, if you can't hear sound when using Sound Forge, try using Media Player. If you can't play sound using Media Player, check the sound card's manual (do not call Sonic Foundry's technical support until you've called the sound card manufacturer).

Mix

Mix is a function in Sound Forge to allow multiple sound files to be blended into one file at user-defined relative levels.

Noise-shaping

Noise-shaping is a technique which can minimize the audibility of quantization noise by shifting its frequency spectrum. For example, in 44,100 Hz audio quantization noise is shifted towards the Nyquist Frequency of 22,050 Hz.

Non-Destructive Editing

This type of editing involves a pointer-based system of keeping track of edits. When you delete a section of audio in a non-destructive system the audio on disk is not actually deleted. Instead, a set of pointers is established to tell the program to skip the deleted section during playback. Sound Forge is a destructive editor except for when using the Playlist.

Normalize

To normalize a file means to raise its volume so that the highest level sample in the file reaches a user-defined level. Use this function to make sure you are fully utilizing the dynamic range available to you. Sound Forge also allows normalization to RMS power. This means that a scan will be done on the sound file and it will be raised in level so its RMS power will be equal to the normalization level. This is helpful for making multiple files perceptually as loud as each other.

Nyquist Frequency

The Nyquist Frequency (or Nyquist Rate) is one half of the sample rate and represents the highest frequency that can be recorded using the sample rate without aliasing. For example, the Nyquist Frequency of 44,100 Hz is 22,050 Hz. Any frequencies higher than 22,050 Hz will produce aliasing distortion in the sample if no anti-aliasing filter is used while recording.

Object Linking and Embedding (OLE)

OLE is a technology developed by Microsoft to allow independent applications to behave as though they are tightly integrated. This allows objects such as Sound Forge audio files to be integrated into other applications such as a Microsoft Word document.

Overview

The overview is the area on the data window directly under the title bar. The entire length of the overview represents the entire sound file. Cursor, selection, and position information is shown relative to the entire length of the sound file.

Peak Data File

A peak data file is what is created by Sound Forge when a file is opened in direct edit mode for the first time. This file stores the information regarding the graphic display of the waveform so that opening a file is almost instantaneous in direct edit mode. This file is stored in the directory that the file resides in and has a .sfk extension. If this file is not in the same directory as the file, or is deleted, it will be recalculated the next time you open the file in direct mode.

Playlist

The Playlist is a list of regions set to play in a specific order. The Playlist allows for non-destructive editing and rearranging of a sound file quickly and easily. Multiple versions of the Playlist can be saved in an external playlist file for easy comparison.

Play Meters

The Play Meters in Sound Forge are toolbar-like items that display the peak output levels of the sound file currently playing. These meters have selectable resolution and options to hold peaks and valleys. These meters will also be active during previews for easier tweaking and clip prevention.

Pulse Code Modulation (PCM)

PCM is the most common representation of uncompressed audio signals. This method of coding yields the highest fidelity possible when using digital storage.

Punch-In

Punching-in during recording means automatically starting and stopping recording at user-specified times.

Pre-roll/Post-roll

Pre-roll is the amount of time elapsed before an event occurs. Post-roll is the amount of time after the event. Pre- and Post-roll have various uses in Sound Forge. Pre-roll can be added to a crossfade preview to listen to the sound before the crossfade begins to give context to it. Pre-roll can also be used in the Playlist to hear previous regions when playback is initiated from the middle of the Playlist.

Preset

A preset calls up a bulk setting of a function in Sound Forge. If you like the way you tweaked that EQ but don't want to have to spend the time getting it back for later use, save it as a preset. All presets show up in the drop down list at the top of most function dialogs in Sound Forge.

Quantization

Quantization is the process by which measurements are rounded to discrete values. Specifically with respect to audio, quantization is a function of the analog to digital conversion process. The continuous variation of the voltages of an analog audio signal are quantized to discrete amplitude values represented by digital, binary numbers. The number of bits available to describe these values determines the resolution or accuracy of quantization. For example, if you have 8-bit analog to digital converters, the varying analog voltage must be quantized to 1 of 256 discrete values; a 16-bit converter has 65,536 values.

Quantization Noise

Quantization noise is a result of describing an analog signal in discrete digital terms (see *quantization*). This noise is most easily heard in low resolution digital sounds that have low bit depths and sounds like a shhhhh type sound while the audio is playing. It becomes more apparent when the signal is at low levels, such as when doing a fade out.

Reactive Preview

Reactive Previews allow for the tweaking of parameters in a function dialog while the preview is playing. When a parameter is changed the preview will automatically rebuild and continue playback.

Region

A region in Sound Forge is a subsection of a sound file. You can define any number of regions in a sound file which are stored in the Regions List.

Regions List

The Regions List is simply the list containing all of the regions and markers defined within the sound file. From this list you can preview and edit the regions as well as drag them to the Playlist or to the desktop to create new files from them.

Resample

To resample means to recalculate samples in a sound file at a different rate than the file was originally recorded. If a sample is resampled at a lower rate, sample points are removed from the sound file decreasing its size, but also decreasing its available frequency range. Resampling to a higher sample rate, Sound Forge will interpolate extra sample points in the sound file. This increases the size of the sound file but does not increase the quality. When down-sampling one must be aware of aliasing. See *Aliasing*.

Root Mean Square (RMS)

The Root Mean Square (RMS) of a sound is a measurement of the intensity of the sound over a period of time. The RMS power of a sound corresponds to the loudness perceived by a listener when measured over small chunks of time.

Ruler, Time

The Time Ruler is the area on a data window above the waveform display. It shows the horizontal axis units as well as marker, region, and loop tags.

Ruler, Level

The Level Ruler is the area on a data window to the left of the waveform display. It shows the vertical axis units as either percents or dB's.

Ruler Tags

Ruler tags are the small tab-shaped controls on the Time Ruler which represent the location of markers, regions, and loop points in the waveform display.

SCSI MIDI Device Interface (SMDI)

SMDI is a standardized protocol for music equipment communication. Instead of using the slower standard MIDI serial protocol, it uses a SCSI bus for transferring information. Because of its speed, SMDI is often used for sample dumps.

Sample

The word sample is used in many different (and often confusing) ways when talking about digital sound. Here are some of the different meanings:

- A discrete point in time which a sound signal is divided into when digitizing. For example, an audio CD-ROM contains 44,100 samples per second. Each sample is really only a number which contains the amplitude value of a waveform measured over time.
- A sound which has been recorded in a digital format; used by musicians who make short recordings of musical instruments to be used for composition and performance of music or sound effects. These recordings are called samples. In this manual, we try to use sound file instead of sample whenever referring to a digital recording.
- The act of recording sound digitally, i.e. to sample an instrument means to digitize and store it.

Sample Dump

A Sample Dump is the process of transferring sample data between music equipment. Because of the large amounts of data required to store digital sound, sample dumps may take a very long time when using the MIDI Sample Dump Standard (SDS). However, when using the faster SCSI MIDI Device Interface (SMDI) protocol, sample dumps can be performed many times faster.

Sample Dump Standard (SDS)

The MIDI Sample Dump Standard is a way to transfer samples between music equipment. Samples transferred with SDS are sent across MIDI cables at the MIDI data rate of 31,250 Hz baud. The slow data rate of MIDI makes SDS a painful experience when transferring large samples. SMDI is a much faster sample transfer method for musicians.

Sample Rate

The Sample Rate (also referred to as the Sampling Rate or Sampling Frequency) is the number of samples per second used to store a sound. High sample rates, such as 44,100 Hz provide higher fidelity than lower sample rates, such as 11,025 Hz. However, more storage space is required when using higher sample rates.

Sample Size

The Sample Size (also referred to as bit depth or word size), refers to the number of bits used to represent a single sample. Sound Forge uses either 8- or 16-bit samples. While 8-bit samples take up less memory (and hard-disk space), they are inherently noisier than 16-bit samples.

Sample Value

The Sample Value (also referred to as sample amplitude) is the number stored by a single sample. In 16-bit audio, these values range from -32,768 to 32,767. In 8-bit audio, they range from -128 to 127. The maximum allowed sample value is often referred to as 100% or 0 dB.

Sampler

A sampler is a device that records sounds digitally. Although, in theory, your sound card is a sampler, the term usually refers to a device used to trigger and play back samples while changing the sample pitch.

Send Mail

A Send option for electronic mail (e-mail) can be found in the File menu. This option makes it easy to send sound files directly from Sound Forge through electronic mail.

Shortcut Menu

A shortcut menu is a context-sensitive menu which appears when you click on certain areas of the screen. The functions available in the shortcut menu depend on the object being clicked on, as well as the state of the program. As with any menu, you can select an item from the shortcut menu to perform an operation. Shortcut menus are used frequently in Sound Forge for quick access to commands. An example of a shortcut menu can be found by right-clicking on any waveform display in a data window.

Signal-to-Noise Ratio

The signal-to-noise ratio (SNR) is a measurement of the difference between a recorded signal and noise levels. A high SNR is always the goal.

The maximum signal-to-noise ratio of digital audio is determined by the number of bits per sample. In 16-bit audio, the signal to noise ratio is 96 dB while in 8-bit audio it's 48 dB. However, in practice this SNR is never achieved, especially when using low-end electronics.

Sign-Bit

Data that has positive and negative values and uses zero to represent silence. Unlike the signed format, two's complement is not used. Instead, negative values are represented by setting the highest bit of the binary number to one without complementing all other bits. This is a format option when opening and saving RAW sound files.

Signed

Data that has positive and negative two's complement values and uses zero to represent silence. This is a format option when opening and saving RAW sound files.

Small Computer Systems Interface (SCSI)

SCSI is a standard interface protocol for connecting devices to your computer. The SCSI bus can accept up to seven devices at a time including CD-ROM drives, hard drives and samplers.

Society of Motion Picture and Television Engineers (SMPTE)

SMPTE time code is used to synchronize time between devices. The time code is calculated in Hours:Minutes:Second:Frames, where Frames are fractions of a second based on the frame rate. Frame rates for SMPTE time code are 24, 25, 29.97 and 30 frames per second.

Sound Card

The sound card is the audio interface between your computer and the outside world. It is responsible for converting analog signals to digital and vice versa. There are a million sound cards available on the market today covering the spectrum of quality and price. Sound Forge will work with any Windows-compatible sound card.

Status Format

The status format is the format by which Sound Forge displays the time ruler and selection times. These include: time, seconds, frames and all standard SMPTE frame rates. The status format is set for each sound file individually.

Tempo

Tempo is the rhythmic rate of a musical composition, usually specified in Beats Per Minute (BPM).

Trim/Crop

Trim/Crop is a function that will delete all data in a sound file outside of the current selection. This is a necessary function when dealing with samples to be played by a sampler to get rid of blank time at the beginning and ending of samples.

μ -Law

μ -Law (mu-Law) is a companded compression algorithm for voice signals defined by the Geneva Recommendations (G.711). The G.711 recommendation defines μ -Law as a method of encoding 16-bit PCM signals into a non-linear 8-bit format. The algorithm is commonly used in European and Asian telecommunications. μ -Law is very similar to A-Law, however, each uses a slightly different coder and decoder.

Undo Buffer

This is the temporary file created before you do any processing to a sound file. This undo buffer allows the ability to rewrite previous versions of the sound file if you decide you don't like changes you've made to the sound file. This undo buffer is erased when the file is closed or the command Clear Undo History is invoked.

Undo/Redo

These commands allow you to change a sound file back to a previous state, when you don't like the changes you have made, or reapply the changes after you have undone them. The ability to Undo/Redo is only limited by the size of your hard drive (*see Undo Buffer*).

Undo/Redo History

This is a list of all of the functions that have been done to a file that are available to be undone or redone. Undo/Redo History gives you the ability to undo or redo multiple functions as well as preview the functions for quick A/B-ing of the processed and unprocessed material. This list can be brought up from within the View menu.

Unsigned

Data that has only positive values and uses half the maximum value to represent silence. This is a format option when opening and saving RAW sound files.

Video for Windows (.AVI)

.AVI files are a format of digital video for Windows. Sound Forge allows you to open these files and edit the audio embedded within them.

Virtual MIDI Router (VMR)

A Virtual MIDI Router (VMR) is a software-only router for MIDI data between programs. No MIDI hardware or cables are required for a VMR, so routing can only be performed between programs running on the same PC. A VMR is normally used to synchronize two MIDI capable programs (for example, a VMR allows Sound Forge to drive a sequencer with SMPTE/MTC). Sonic Foundry supplies a VMR with Sound Forge called the Sonic Foundry Virtual MIDI Router.

Waveform

A waveform is the visual representation of wave-like phenomena, such as sound or light. For example, when the amplitude of sound pressure is graphed over time, pressure variations usually form a smooth waveform.

Waveform Display

The waveform display is the part of the data window which shows a graph of the sound data waveform. The vertical axis corresponds to the amplitude of the wave. For 16-bit sounds, the amplitude range is -32,768 to +32,767. For 8-bit sounds, the range is -128 to +127. The horizontal axis corresponds to time, with the left-most point being the start of the waveform. In memory, the horizontal axis corresponds to the number of samples from the start of the sound file.

Zero-crossing

A zero-crossing is the point where a fluctuating signal crosses the zero amplitude axis. By making edits at zero-crossings with the same slope, the chance of creating glitches is minimized.

Zipper Noise

Zipper noise occurs when you apply a changing gain to a signal, such as when fading out. If the gain does not change in small enough increments, zipper noise can become very noticeable. Sound Forge's fades are accomplished using 64-bit arithmetic, thereby creating no audible zipper noise.

This page was intentionally left blank.

A

A-Law 208, 384, 426, 475
ACM *see* Audio Compression Manager
ADPCM 212, 345, 383, 401, 404, 424, 428
AIFF 211, 401, 425
 transferring to Macintosh using 401
ACID Looping Tools 266-267
 Halve Selection 266
 Double Selection 266
 Shift Selection Left 266
 Shift Selection Right 267
 Rotate Audio 267
ACID Properties 267-268
 One-Shot 267
 Loop 267
 Disk-Based 268
Active Streaming Format 380
Add File 369
additive synthesis 118, 119, 331, 332
aliasing 281, 291, 312, 332, 475, 485
Amiga .SVX Files 424
Amplitude Modulation 295-296
 defined 475
Analog Recording/Playback 8-9
 defined 476
Annotations 335
Attributes 379
Audio Compression Manager 379,
 401, 427-430, 476
 description of 427-430
Audio Event Locator 476
audio & video sync 159
Auto Play 207
Auto Region 130, 315
Auto Snap
 Time 62, 340
 Zero 45, 340, 407
Auto Trim/Crop 274

Auto Update 393
Automatic archive folder 389
Automatic Batching 140, 388
.AVI Files 57-62, 69-72, 158,
 209, 230, 270, 404, 426, 492
.AVI Compression 60-61, 214
 compressors for 214
 configuring 215
 data rate 60, 215, 193
 interleave frames 214
 key frames 60, 159, 215
 settings 60
 quality of 215

B

background processing 158
bar graph 395
Batch Converter 137-141
 Add Files 137
 Add Processing Functions 138
 Convert 139
 Destination Format 138
 Destination Folder 139
 Processing Specific Files 139
 Save Settings 139
 Using Batch Scripts 139-140
 Window 137
Batch Converter Reference 369
 Convert Menu 388
 Edit Menu 370
 File Menu 369
 Options Menu 371
Batch Script 139, 370, 381
 Batch Script Add 140, 370, 381
 Batch Script Save As 140, 370
buffers
 Delay Buffer 222
 pre load buffers 156, 359, 400
 total buffer size 156, 322
 undo/redo 229-230, 352-353, 357

C

canceling operations 92, 158, 177, 407
ceiling 392
Center Cursor 271
changing sound formats 63
channels 54-56, 65-66, 92-93
 previewing 56
 single-channel editing 56, 84
 toggling 56
Channel Converter 66, 276-277, 477
check boxes 97
Chorus 107, 295-297, 477
Class 381
Clear 240, 261, 264, 406
Clear/Delete 32, 240, 407
clicking with the mouse 16-17
clicks 399
clipboard 32-34, 214-215, 246, 407
 copying data to 32, 232
 create new file from 34
 overwrite from 238
 pasting data from 33, 234-236
 playing contents of 246
 replicate from 238
clipping 477
Close 210, 370, 393
Close All 367
Compression 214-215, 217-219, 427-430
 data 215, 427
 Dynamic 66, 288
 options 214
Convert to 8-Bit 66, 277-278, 477
 Dither 278
 Round 278
 Truncate 277
controls 93-98
Conversion Settings 371
Convert Menu 388

Convert Selected Only 389
converting
 files 66-69
 formats 62-69
 sample rates *see sample rate*
Copy 31-33, 233, 406, 419
Copy Object Link 233-234, 454, 478
counters (play and record) 156
Crash Recovery 152-152
creating
 new file from Playlist 62, 134
 new window 30
 new window from the clipboard 34
Crop *see Trim/Crop*
Crossfade 32, 52, 235-236, 344
Crossfade Loop 181, 317-318, 478
 Graph 318
 Loop to Crossfade 318
 Preview mode for 318
cursor
 centering 270
 definitions 55, 287
 position 229, 327
Cut 34-35, 101, 232-233

D

DC Offset 85, 253, 278-279, 328, 352, 478
 finding 279, 328
 removing 278, 408
data compression 427-430, 477
data format 228-229, 379
Data window 18, 20-22, 478
 maximize width 22, 245
Delay/Echo 103-105, 233, 285, 297-301
 Multi-Tap 104-105, 298-301
 Simple 103, 297-298
Delete *see Clear*

- data 32, 240
 - file 193
 - region 131, 134, 260, 262, 410
 - silence between phrases 275
- Delete Source after conversion 387
- Destination folder 388
- direct edit mode 154-155, 207, 344, 480
- display
 - level ruler 20-22, 40, 412-413, 421, 488
 - Color Preferences 348-349
 - max zoom 39-41
 - time ruler 21, 22, 36, 39, 53, 412, 421, 488
- Display Menu 394
- Display Range 394
- Distortion 301
- Dither 67, 283-284, 476, 480
- drag and drop 339, 420-421, 480
 - creating new window 53, 344
 - operations 51-53
 - OLE 454-455
- drop-down lists 95-96
- Drop Marker 136, 255, 272
- ducking 54, 238
- Dynamics 109-112, 302-306
 - Graphics 302-304
 - Multi-Band 113-114, 304-306

E

- echo *See Delay/Echo*
- edit box spinner controls 95
- Edit Frame Rate 269-270
- Edit Menu (Batch Converter) 370
- Edit Menu (Sound Forge) 231-245
 - Clear/Delete 240
 - Copy 233
 - Copy Object Link 233-234
 - Cut 232
 - Disable Undo/Redo 245
 - Go To 242
 - Paste 234-235
 - Paste Special 235-239
 - Paste to New 239
 - Pre-roll to Cursor 241
 - Preview Cut/Cursor 240
 - Redo 232
 - Repeat 232
 - Select All 240
 - Selection 243-244
 - Tool 241
 - Trim/Crop 240
 - Undo 231
 - Undo All 245
- edit operations 32-37
- Edit Sample 268-268
- Edit Sample Rate 63
- Edit Tempo 270
 - destructive vs. non-destructive 127
- effects
 - applying to entire file 92
 - applying to section 91
 - applying to single channel 92
 - applying to stereo file 92-93
 - canceling 92
 - changing affected area 102
 - getting help on 93
 - length of time to process 157-158, 402-403
 - presets 98-99
 - previewing 100
 - repeating 101
 - shortcuts 101
- Effects menu 294-315
 - Amplitude Modulation 295
 - Chorus 295-296
 - Delay/Echo 297-301
 - Distortion 301
 - Dynamics 302

- Envelope 307
- Flange/Wah-Wah 308-309
- Gapper/Snipper 309-310
- Noise Gate 310
- Pitch 311-313
- Reverb 313-314
- Vibrato 314-315
- embedded information 69-72
- Envelope 307
- Envelope graphs 97-98
- Error Log Destination 386
- EQ 121-126, 279-283, 411
 - Graphic 124, 279-280
 - Paragraphic 121-123, 280-281
 - Parametric 125-126, 281-283
- Exit 231
 - confirming before exiting 344
- expansion 112, 289, 302
- F**
- FFT 390
 - overlap 390
- FM Synthesis 118-250, 191, 481
- Fade 283-285
 - Graphic 283-284
 - In 284
 - Out 285
- File Format 371
- File Menu (Batch Converter) 369
- File menu (Sound Forge) 205-231
 - Close 210
 - Exit 231
 - New 205-206
 - Open 206-208
 - Open Raw Data File 208
 - Open Video Stream 209
 - Properties 224-230
 - Save 211
 - Save All 224
 - Save As 211-224
 - Send 230-231
 - Video Save Options 213
 - Workspace 210
- file types 423-426
 - Ad lib 423
 - Amiga .SVX 423,
 - Covox 424
 - Creative Labs .VOC 424
 - Dialogic .VOX 69, 424
 - Gravis .PAT 424
 - Intervoice 425
 - Macintosh .AIFF 425
 - Macintosh Resource 187, 425
 - Microsoft .WAV 426
 - MIDI SDS 425
 - Next/Sun Sparc .AU 425
 - raw data files 206, 208, 212, 426
 - RealAudio 193, 425
 - RealVideo 425
 - SampleVision .SMP 425
 - Sonic Foundry .SFR 426
 - Sound Designer 1 426
 - Sounder/SoundTool .SND 426
 - Video for Windows 426
- files
 - can't load compressed 401
 - converting 64-69
 - formats 423-426
 - of type (*see file types*)
 - opening 206-209
 - play 25, 397-400
 - properties 224- 230
 - saving 211-224
 - saving for use with Session 8
 - 212, 353, 444
 - processed by Batch Converter 389
- Filled Graph 395
- filtering 121-126
 - Graphic EQ 24, 279-280
 - Paragraphic EQ 121-123, 280-281
 - Parametric EQ 125-126, 281-283

Find 320-322
Flange 108, 308-309
Floor 392
Folder to monitor and
 automatically batch 388
frame rate
 configuring 38
 setting for window 269-270
Frequency Maximum 392
Frequency Minimum 392
frequency modulation
 FM Synthesis 118-250, 191, 481
 Vibrato 314-315

G

Gapper/Snipper 309-310
General MIDI Patch Map 436
Go To 43,242
Grab/Pan 394
Graphic EQ 124, 279-280

H

hard drive use 153-154
Help
 on effects 93
 online 23-24, 367-368
Help menu 324-325
 About Sound Forge 368
 Contents 367
 Keyboard Shortcuts 368
 Search for Help on 367
 Sonic Foundry Home Page 368
 Tip of the Day 368
 Troubleshooting 368
Hold peaks during monitoring 392
horizontal trackbars 94

I

identifying unknown sound formats
 206-208
Insert Silence 285
insertion point 482
installation 3
Invert/Flip 285
Invert Selection 370

L

limiting 110-112, 302-306
Line Graph 395
List Errors 369
Lock Loop/Region Length 339
logarithmic 394
 graphing 392
Log Errors to file 386
loops
 creating 179
 creating/editing w/out loop tuner
 180
 editing 179
 fine-tuning loop end points 183
 Lock Loop/Region Length 184,339
 locking length of 339
Loop Tuner 180-184, 412
 looping techniques 180-181
 opening 181-182
 playing from 182
 real-time loop tuning 184
 saving loop points 184
 Sustaining and Release loops 182
 Zero-Crossing finder 183
loudness 7

M

- MIDI 193-204
 - configuring devices 184
 - enabling Output Sync 82-83, 341
 - enabling Input Sync/Trigger 83, 340, 337
 - notes and frequencies 370-373
- MIDI Keyboard 177-178
 - configuring output port/channel 178
 - opening 178
 - selecting instruments 179
 - setting MIDI device for 178
 - triggering file playback from 197-199
- MIDI /SMPTE 193-204
- MIDI Time Code (MTC)
 - generating from Sound Forge 204
 - synchronization with Sound Forge 201-204
- MIDI Triggers 128, 201-204, 449-450
- Macintosh
 - transferring files to 425
- magnification 39-42
- Magnify Mode 42, 241, 421
- main screen 15-24
- Maintain last monitored view 392
- Mark In 272, 410
- Mark Out 273, 410
- Markers 135-136
 - creating 135
 - dropping 136, 272
 - editing 136
 - starting playback from 136
 - to Regions 261
- Maximize Width 22, 245
- maximum amplitude
 - finding 326-327
- Mix 32, 35-37, 54, 237-238, 406, 484

Mixer

- setting devices to control 406, 416
- opening 249

Monitor Playback 393**Monitor Record 393****mono to stereo 64-65****mouse operations 16-18**

- clicking 16
- control-clicking 17
- double-clicking 17
- dragging 17
- dropping 18
- pointing 16
- shift-clicking 17
- slow-dragging 17
- toggle-clicking 17
- triple-clicking 17

Multi-Tap Delay/Echo 104-105, 298-300**Mute 286****N****Navigation 32-56**

- Go To 43
- using the Overview 22
- using Views 50-51

New window 30**Noise Gate 109-110, 310****Normal Display 395****Normalize dB 394****Normalize 286-289, 484****Nyquist Frequency 11, 485****O****Object Linking and Embedding (OLE)
22, 233-234, 453-458****OLE Drag Source 22, 347, 454-455, 458****online help 23-24, 367-368****Open 24, 206-208**

- raw data files 208

- read-only 207, 245
- Open Destination 369
- Open Source 369
- Optimizing Sound Forge 153-159
 - Audio and Video Synchronization 159
 - Background Processing 157-158
 - Editing in Direct Mode 154-155
 - Fast File Saving 155
 - Hard Drive Use 153-154
 - Meters 157
 - Passive Updating for Video and Time Displays 157
 - Playback Cursor and Record Counter 157
 - Setting Up Wave Devices 155-156
- Options 333-323
 - Annotations 335
 - Auto Snap to Time 340
 - Auto Snap to Zero 340
 - Drag and Drop Snapping 339
 - Lock Loop/Region Length 339
 - MIDI Input Sync/Trigger 340
 - MIDI Output Sync 341
 - MIDI Triggers 341-343
 - Paste Events 339
 - Play Meters 336
 - Preferences 343-366
 - Pre-Queue for SMPTE/MTC 341
 - Scroll Playback 339
 - Scroll Smoothly 339
 - Status Format 333-335
 - Time Display 337
 - Video 337-338
- Options Menu (Batch Converter) 371
- Options Menu (Spectrum Analysis) 390
- Overview 22, 48-50
- Overwrite 238

P

- PC speakers (synchronic wave devices) 10
- Pan/Expand 289-291
 - MS Mix 290
 - Output Gain 290-291
 - Pan left and right 289-290
 - Process Mode 289
 - Stereo expand 290
- Paragraphic EQ 121-123, 280-281
- Parametric EQ 125-126, 281-283
- Passive Update 157, 337, 338
- Paste 33-34, 53, 234-235, 406, 419
- Paste Events 339
- Paste to New 239
- Paste Special 235
 - Crossfade 235
 - Mix 237
 - Overwrite 238
 - Paste to New 239
 - Replicate 239
- Pencil Mode 242
- phaser 108, 308-309
- pitch 7
- Pitch Bend 311
- Pitch Shift 312-313
- play cursor/pointer 349, 350, 358
 - bias 272, 237, 363
 - interpolation 363
 - problems with 272, 273
- Play Meters 157, 250, 336, 416-417
 - Always on Top 337, 347, 417
 - Hold Peaks 253, 336, 417
 - Hold Valleys 253, 336, 417
 - Reset Clip 252, 336, 416
 - Show Labels 252, 336, 416
- Play modes 27-29
 - Play as Sample 28, 258
 - Play Looped Mode 28, 258, 365

- Play Normal Mode 28, 256-257, 423
- play counter 356-357
- playbar 21, 27-28, 411
- playing
 - crashes 400
 - files 25-26
 - files won't play 398
 - grayed buttons 397
 - Playlist 134
 - selection 26
 - sounds wrong 400
- Playlist 127-136, 217
 - adding regions to 132
 - arranging 133
 - changing region playcount 133
 - clearing 261
 - converting to new 265
 - copying onto clipboard 265
 - creating a new file from 134
 - deleting from 262
 - editing entries 263
 - opening 264
 - playing 134
 - pre-roll playback 264
 - removing region from 134
 - replicating 264
 - saving as 212, 265
 - setting display format 316-317
 - setting pre-roll time 263
 - setting stop point 264
 - setting triggers for 202-203
 - shortcut menu for 410-411, 412
 - with .AVI files 61
 - storage in .WAV files 448-449
- power (RMS) 286-289, 328, 361, 414, 487
- Preferences 343-366
 - Display 347-351
 - Editing 351-353
 - File 353-355
 - General 343-347
 - MIDI/Sync 355-357
 - Other 365-366
 - Performance 357-358
 - Playlist 358-359
 - Previews 359-360
 - Status 360-361
 - Toolbars 361-362
 - Video 364-365
 - Wave 362-364
- Pre-roll to Cursor 241
- presets 98-100, 381, 486
 - creating presets 99
 - storing custom presets 99-100
 - using presets 98
- preview 86, 100-101, 318, 359-360
 - configuring 100-101
 - cut 101
 - effects 100-101
 - video 58, 249, 338
- Preview Cut/Cursor 101, 240
- Print Display 393
- Process menu 274-
 - Auto Trim/Crop 274-276
 - Channel Converter 276-277
 - Convert to 8-Bit 277-278
 - DC Offset 278-279
 - EQ 279-283
 - Fade 283-285
 - Insert Silence 285
 - Invert/Flip 285
 - Mute 286
 - Normalize 286-289
 - Pan/Expand 289-291
 - Resample 291-92
 - Reverse 292
 - Smooth/Enhance 292-293
 - Time Compress/Expand 293-294
 - Volume 294
- Punch-In (mode) 88, 252, 486
 - using pre-roll/post-roll 88
- push buttons 96

R

- RIFF Wave Chunks 443-449
- radio buttons 96
- rate *see* sample rate
- raw files 426
 - opening 208-209
 - saving 212-213
- Read-only mode 207
- RealAudio 372, 373, 425
 - saving as 215
- RealMedia 374
- Rebuild Peak Data 274
- Record 80-90, 219-223, 348
 - available time 83
 - blinking status indicator 253
 - Calibrate 253
 - counter 157, 357
 - DC adjust 85-86, 253,
 - Drop Marker 136
 - gap detection 253
 - grayed buttons 397
 - levels 84
 - Meters 84-85, 252
 - modes 86-88, 251-252
 - Monitor 252
 - no response from levels 328
 - prepare 86, 254
 - Pre-roll/Post-roll 252
 - previewing 86-88
 - punch-in 88
 - remote 81-82
 - status 81
 - Sync 82-83, 253
 - to another window 82
 - to New window 81
 - won't stop with Stop button 399
- Redbook Audio 401
- Refresh List 387
- Regions List 127-136
 - add markers/regions 132, 258
 - arranging in playlists 133
 - Auto Region tool 130, 315
 - copy onto clipboard 262
 - clearing from playlist 261
 - creating regions 128-130
 - deleting regions 262
 - destructive vs. non-destructive editing 127
 - editing 131, 260-261, 263
 - editing with ruler tags 131
 - fast navigation through 131
 - locking length of a region 339
 - MIDI synchronization and triggering 128, 258-259, 260-261, 263
 - replicating 264
 - setting display format 350
 - shortcut menu for 410, 412, 314-315
 - sorting alphabetically 358
 - splitting 261
 - updating 261
- Registering Sound Forge 4-5
- Remote Record 89-90
- Remove All 371
- Remove Done 371
- Remove Selected 370
- Rename with long file names 387
- Repeat 101, 232
- Replace existing files 386
- Replicate 239, 261, 264
- Resample 291, 487
- Reset All 371
- Reset Error Log 386
- Reset Selected 371
- Restore All 367
- Reverb 106, 313-314
- Reverse 292
- Ruler
 - Level 20, 40, 412-413, 421, 488
 - shortcut menus for 421

- Time 2, 37, 39-40, 53, 412, 421, 488
 - ruler tags 21, 488
 - shortcut menus for 412
 - Run Batch Script 139-140, 388
 - Run Current Script 139-140, 388
 - Run Multiple Scripts 139-140, 388, 389
 - Run Scripts 139-140, 388
- S**
- SDS files 161-162
 - setting channel number 165, 354
 - setting patch number 354
 - SMPTE
 - configuring for use with video 355
 - discussion of 431-435
 - display format 360
 - enabling 341
 - losing sync 403
 - Pre-Queue Playlist regions 341
 - setting the default type 360
 - using Internal Timer for generation 356
 - sample looping *see* loops
 - sample parameters
 - MIDI unity note 263
 - setting loop count 269
 - setting loop type 208
 - sample rate 11, 384
 - changing play/record rate 291
 - resampling 291
 - sample type 384
 - Sampler Tool 162-174
 - Actual Send/Receive Sample Number 163
 - Configuration 166-172
 - Configure 164-166
 - Get Sample 163
 - Logical Send/Receive Sample Number 163
 - Send Sample 163
 - samplers 161-173
 - MIDI/SDS 161
 - SCSI/SMDI 162
 - configurations 166-172
 - external 161
 - internal 162
 - saving sampler configurations 174-175
 - selecting a sample bias 165, 173-174
 - tips and ideas 191-192
 - transfer procedures 462-473
 - Save 211
 - Save As 211-212
 - saving files 31, 155
 - video options for 60-61, 213
 - Select All 240, 319, 336, 370, 406, 421
 - Select Done 370
 - Selection 243--244
 - Set 43-44, 244
 - Snap Edge to Time 243
 - Snap Edge to Zero 244
 - Snap to Time 243
 - Snap to Zero 45, 243
 - Selection status fields 22
 - shortcut menu for 342
 - selections
 - changing 102
 - making 33
 - making with the keyboard 46-47
 - making with the mouse 45
 - playing 26
 - restoring 48
 - setting with dialog 43-44
 - single channel in stereo 55
 - togglng 56
 - Send 231-231, 490
 - serial number 4
 - Session 8 212, 353
 - Set sonogram resolution 391
 - Setting Dialog 390
 - shortcuts

- drag and drop 420-421
- keyboard 405-411
- mouse 421-423
- Shortcut Menus
 - Display Properties 412
 - Keyboard MIDI Out 416
 - Level Ruler 412
 - Level Toolbar 413
 - Loop Tag 415
 - Marker Tag 414
 - Mixer Window 416
 - Play Meters 416-417
 - Playbar Selection Status 410
 - Playlist 412
 - Preview Configuration 417-418
 - Record Dialog 418
 - Region Tag 414-415
 - Regions List 412
 - Sound Forge Desktop 410
 - Status Bar 415
 - Time Display 418-418
 - Time Ruler 412
 - Video Preview 419
 - Video Strip 419
 - Waveform Display 410
- Show Full Path 385
- Show Notes 393
- Show Position 393
- Simple Delay/Echo 103, 297-298
- Simple Synthesis 118
- Slices displayed 391
- Smooth/Enhance 292
- Smoothing window 391
- Snap to Time 243, 244
- Snap to Zero 243, 244
- Sonic Foundry Virtual MIDI Router
 - 193-197
 - configuring 195-197
 - description of 193
 - installing 194
- Sonogram 395
 - color 395
 - black and white 395
 - resolution 391
- Sound Designer 1 426
- sound formats
 - changing 63-66
 - editing 62-72
 - mono to stereo/stereo to mono 64
- Sound Mapper 155-156
- sound waves 6
- Source Raw Attributes 384
- Special Menu 250-
 - ACID Looping Tools 266-268
 - ASF Markers/Commands 265-266
 - Center Cursor 271
 - Drop Marker 272
 - Edit Frame Rate 269-270
 - Edit Sample 268-269
 - Edit Tempo 270-271
 - Mark In 272
 - Mark Out 273
 - Playlist 262-265
 - Rebuild Peak Data 274
 - Regions List 258-262
 - Toggle Selection 273
 - Transport Items 250-258
 - Undo/Redo History 274
- Spectrum Analysis 142-150
 - Analysing Spectral Data 142
 - Spectrum Graph 145
 - Sonogram 149
- Spectrum Analysis Reference 390
 - Options menu 390
 - Display menu 394
- Statistics 327-328
 - Average Value (DC Offset) 328
 - Cursor Position 327
 - Maximum/Minimum Position and Sample Value 327

- RMS Power 328
- Sample Value at Cursor 327
- Zero Crossings 328
- Status Bar 19, 346, 347, 358
 - shortcut menu for 422
- Status Formats 37-38, 333
- stereo files
 - editing 54
 - effects on single channel 56
 - panning between channels 289-291
 - playing channels 55-56
 - previewing channels 56
 - selecting data in 55
 - single channel editing 56
 - swapping channels 55
 - toggling selection 55
- stereo to mono 64-65
- Stop on errors 385
- storage folder 357
- Summary 70-71, 225-228
 - additional embedded information 69
 - fields described 225-228
- Suppress Initialization Dialogs 386
- Summary Information 379
- synchronizing 128, 201-204
- Sync Graphics 394
- Synthesis 118-121
 - DTMF/MF 120-212, 328
 - FM 118-120, 328, 481
 - Simple 118, 332
- T**
- technical support 5-6
- tempo
 - configuring for window 270-271
- temporary files 83, 151-152, 245, 274 357
- temporary file storage 151-152, 357
- text, embedding 69-71, 225-228
- time (duration)
 - altering without pitch change 116-117
 - changes 115-116
- Time Compress/Expand 116, 270, 293-294
- Time Display 249, 337
 - Passive Update 157, 337
 - Playlist Position 337
 - Position 337
 - Sync/Trigger Status 337
- time format *see* Status Format
- toggle selection 273
- Tool 210-211
 - Edit 33, 210
 - Magnify 42, 210
 - Pencil 210
- toolbars 19-20, 93, 214, 320-321
- Tools menu 315
 - Auto Region 315-316
 - Crossfade Loop 317-318
 - Extract Regions 319-322
 - Find 320
 - Preset Manager 322
 - Run Batch Script 323
 - Sampler 324-327
 - Statistics 327-328
 - Synthesis 328-333
- Tooltips 23, 362
- Transport
 - menu items 251-258
- Trigger List Chunk 444-447
- triggering 128, 197-201, 254, 258-259
- Trim/Crop 32, 35, 240, 274, 491
 - Auto Trim/Crop 274-275
- Troubleshooting 397-404
 - .AVI Files 404
 - Effects 402-403
 - File Support 401-402
 - MIDI 404

Playing and Recording 397-400
Sampling and Sequencing 403

U

Undo/Redo 231-232, 245, 422
 disabling by default 311
 disabling for a data window 232
Undo/Redo History 250, 274, 406,
 422
 Clear 274
 Clear All 274
Use Long File Names 387

V

.VOC files 353, 424
 saving using new format 353
vertical faders 94
Vibrato 314-315
Video 57-61, 230, 337-338
 Animate Video Strip 338
 backgrounds for 338
 Copy Current Frame 337
 Integral Stretch 338
 Number Frames 338
 Passive Update 157, 338
 Preserve Aspect 338
 Stretch to Window 338
Video Compression 60, 214
Video Preview 58, 249
Video Stream 57, 209, 199
 opening a 209
 saving a 213, 375
View menu 245-250
 Clipboard 246
 Focus to data window 248
 Full Screen 246
 Keyboard 249
 Loop Tuner 250
 Maximize Width 245
 Mixer 249
 Play Meters 250

Playlist 249
Regions List 248
Time Display 249
Toolbars 246
Undo/Redo History 250
Video Preview 249
Zoom Level 246
Zoom Time 247
Views 50-51
Virtual MIDI Router 193-197, 492
 configuring 195-197
 installing 194-195
Volume 294
 gain 294
 loss of when running effects 402
 setting 294

W

Wah-Wah 108, 308-309
warbling problems 397
Wave Devices 155-156, 362
 configuring 156-157
Wave Mapper *see* Sound Mapper
waveform display 21, 493
 shortcut menu for 411
Window menu 366-367
 Arrange Icons 367
 Cascade 366
 Close All 367
 Minimize All 367
 New Window 366
 Restore All 367
 Tile Horizontally 366
 Tile Vertically 367
windows
 active vs. inactive 31
 Maximize Width 245
workspace 210

Z

- zero-crossings 493
 - Auto Snap to 340
 - endpoints 181
 - finder 183
 - setting slope 351
 - Snap to 243
- Zoom Level 246
 - Out Full 246
 - Selection 247
 - Window 246
- Zoom Time 247-248
 - In Full 247
 - Normal 247
 - Out Full 248
 - Selection 248
- Zoom to Range 394
- Zoom out full 394