



Module 9

Software for Capturing, Editing, and Storing Digital Audio



ONLINE PROJECTS

Projects are designed to give you hands-on experience with the tasks in Table 9.1.



LINK

As noted in Module 8, the original popular MPEG audio format was MP3. The newer AAC and WMA formats offer improved quality and compression. Be sure to review the compression details for these audio files in Module 8 and the hardware possibilities in Module 10. When we refer to web audio or web music, we are referring to MP3, WMA, and AAC.

In the last module, you learned about acoustics and the distinction between digital and analog sound. We stressed the historic notions of sound synthesis and how these form the basis for modern sound design. You also learned the important sound file formats and became acquainted with compression techniques such as those found in MP3 and AAC files. In this module and in the many software modules to come, we describe how to apply many of these ideas to making and enjoying music.

Digital audio software is the fastest-growing category of music software today. Ten years ago, perhaps 15 software titles were widely used for digital audio work, mostly in professional studios and labs. Now, the number is well over 200, with additional titles hitting the shelves monthly. Software has now replaced much of the hardware used to capture and process digital audio and many musicians are working with this generation of digital audio software in professional studios, home setups, and Internet settings.


As with each of the software modules to come in this book, we stress the importance of matching software to your desired task. Table 9.1 displays a set of tasks that are commonly done when capturing, editing, and storing digital audio. The table suggests a likely setting for each task and some typical software titles that can be used. In the remainder of the module, we present information for each task.

Working with Audio on the Web


One of the more popular uses of digital audio with computers involves MP3 software and streaming media. In this section, we focus on the use and production of MP3 files and the newer formats, AAC and WMA. The popularity of these audio file formats, especially MP3, has grown enormously, fueled by Internet sites that enable sharing thousands of these files in all genres of music and the development of free and inexpensive software that encodes MP3 files from sound sources such as commercial CDs. An entire generation of small, very portable hardware devices known as MP3 players has developed to host these files, directly supported by a personal computer. Custom-created CDs that are filled with these files are easy to create with today's software tools and such discs can often be played in the same player that hosts the more traditional audio CDs.

TABLE 9.1 Tasks for Basic Digital Audio

Setting	Task	Typical Software
Home	<ul style="list-style-type: none"> • Download and listen to MP3 files • Create playlists and organize MP3 files • Create custom CDs for your MP3 files • Encode MP3 files from CDs and use ID3 tag data • Use and create streaming media 	Winamp (Win) Yahoo! Music Jukebox (Win) iTunes (Mac/Win) Media Jukebox (Win), Audion (Mac) Windows Media Player (Mac/Win), QuickTime Pro (Mac/Win), RealPlayer (Mac/Win)
Home and Studio	<ul style="list-style-type: none"> • Prepare your computer for importing sound from CDs; live recording with simple microphones • Record and edit a live and a CD segment of digital audio • Use basic editing, file management, effects-processing options • Experiment with advanced editing, file management, and effects-processing options 	Operating System Settings Audition (Win), Audacity (Mac/Win), Sound Studio (Mac) PeakLE (Mac), Sound Forge 9 (Win), Amadeus Pro (Mac) Peak(Mac), Sound Forge 9 (Win), WaveLab (Win)

**LINK**

See Viewport II, Module 5, for Internet software concepts, including information about legal downloads.

**LINK**

Peer-to-peer network computing, as well as client-server networking, is explained in Modules 5 and 7.

Obtaining Web Music Files

Large numbers of these audio files, especially MP3s, that are free and legally offered to the public can be found on the Internet and downloaded by using procedures defined by websites, FTP servers, and news groups. In using any of these sources, we stress that it is your responsibility as a user to verify that the content you download is, in fact, legally distributed. Locations such as these offer music CDs for purchase, but also provide MPEG audio files that can be auditioned or downloaded for use at a later time. Often, these sites have agreements with artists and recording companies that allow them to offer this music in the hopes that you might want to purchase a full CD.

Another popular source for free and legal web music files is the actual Internet site for the ensemble or artist you are interested in. For example, rock groups may actually encourage distribution of their music to increase concert attendance or CD sales.

Finally, another method of obtaining web music files is an approach that often is called “peer-to-peer” distribution. This was the technique used by the Napster site when it was launched in 1999. Music files are exchanged directly between personal computers with no central server needed; everyone shares. This approach was challenged in court by the recording industry as encouraging copyright infringement and the industry won cases that forced the Napster site to close in 2001. As we noted in Module 5, similar peer-to-peer software programs have emerged following Napster’s demise. We strongly advise against using this method of obtaining MP3 files unless you are very sure that the material is intended to be free.

Some sites on the Internet that once were portals for free MP3 distribution have now developed a commercial approach that encourage you to purchase MP3s



ASIDE

Napster was reborn as a month-by-month, subscription-based web music service after its assets were purchased by Roxio, Inc. Roxio then changed its name to Napster, Inc. in December 2004.

for your collection. This form of marketing music is becoming increasingly popular as an alternative to purchasing CDs at a local store and has contributed greatly to the decline of CD sales worldwide. Online purchasing of music in this way may become the dominant approach to music distribution in the coming years.

Organizing and Playing Web Music Files

Software for organizing and playing music files is readily available for all types of computers. Two such programs are displayed in Figures 9.1a and b. Winamp is a classic program that has had a long history of supporting MP3s and other file formats such as CD audio, WAV, and MIDI. The software is free and does not produce any pop-up ads on your screen. Windows Media Player comes free with the Windows operating system and is Microsoft’s all-purpose media player, supporting its proprietary WMA format as well as MP3, among others.

Take a moment to study the displays for both titles. For Winamp, we have captured this view as the software is playing back the second item in the playlist editor. A similar but more expanded playlist is displayed in Windows Media Player. A playlist is simply a collection of, in this case, MP3 files that make logical sense for you. For example, you may want to create a large playlist that has all the titles by the King Singers or you may want to create a playlist of various music titles that depict love or a favorite guitar player. Most MP3 software allows you to organize the music in any way you want. For Winamp, buttons on the bottom of the playlist window allow for this organization; Windows Media Player has similar controls.


One additional point about playlists relates to the information itself. The MP3 file format provides for what is known as “ID3” tag information to accompany data for the audio. This means that text for the artist, title, band name, genre, year, and other information can be included. The nature of the information allowed by the format is expanding all the time; additions may include song lyrics and even artwork, such as CD covers. Such information can be used by the software to enhance displays.



FIGURE 9.1A Winamp



FIGURE 9.1B Windows Media Player

**LINK**

For help understanding what 160 Kbps means, go back to Module 8.

**ASIDE**

The concept of “plug-ins” will return many times throughout the software modules in this text. Nearly all music software companies today make use of plug-ins as a way of extending the power and function of their titles.

The controls for playback in the upper left of Figure 9.1a are self-explanatory. They control the playback of not only MP3s but also inserted audio CDs and other digital audio files on the computer. The software shows the timing progress and also the attributes of the file playing (MP3 file encoded at 160 Kbps, stereo at 44.1 kHz). The equalizer (EQ) graphic at the bottom left is a common feature with MP3 players. Moving the slider buttons changes the filtering for frequencies in the various auditory ranges as described in Module 8. We will present more information about EQ later in this module.

The mini-browser window to the right, a special feature of MP3 players, is designed to seek information from the Internet to support the music playing locally, as well as provide commercial links. Notice the links to plug-ins. The company that distributes Winamp gives programmers enough code to customize extensions for the software. This allows expandability of components, such as special visualizations (changing graphic displays for visual interest), sound processing effects, or input/output controls.

Notice that in both the Windows Media Player and Winamp, reference is made to “radio” and “skins.” Many MP3 software titles offer expandability by hosting Internet radio streams that turn the software into a kind of special software radio station for different kinds of music. These may not be real radio stations per se, but rather streams of audio organized to be heard through the software. “Skins” is a special name for visual representations of the software interface. Applying such skins allows you to customize the look of the software to suit your mood or personality, but has no effect on the music heard.

Creating and Storing Your Own Web Music Audio

If all you could do with web music audio software was play and organize other people’s files, there would be far less interest in this digital audio category. Besides the clear advantages offered by file compression and high-quality sound, the ability to encode and store your own collections of digital music has contributed enormously to the MP3 craze.

Settings that show how to do the encoding are included within the software. Figure 9.2 gives a clear picture of how this is done in the iTunes software offered free for Macintosh and PC computers. In this graphic, we are importing music from an audio CD. In the middle screen, we can see the MP3 encoding. The bit rate for stereo files is set at 160 Kbps which, as you can tell from Module 8, Table 8.3, is slightly better than near-quality CD sound. The software provides for a range from 8 to 320 Kbps.


Variable bit rate encoding allows the software to change the 160 setting “on the fly” based on the complexity of the music it is sampling, thus offering more sophisticated matches between music and bit rate. This option may actually make the file bigger with no great difference in sound quality, so use your ears to test this option.

The software also filters all inaudible sounds below 10 Hz for added efficiency in size. The Joint Stereo setting reduces redundancy across the two stereo tracks and is most useful for bit rates under 80. The Auto Sampling Rate option matches the software rate to the original source (in this case 44.1 kHz) and is the recommended choice in nearly all situations.

The iTunes graphic also demonstrates important storage features for software such as this. Notice that the audio CD is mounted while the encoding process is

**TIP**

Portable digital audio players are called MP3 players, even though they now play other audio formats such as WAV, AIF, RM, and AAC.



LINK

Burning CDs is presented in more detail in Viewport IV, Module 12. Hardware for CD-R, DVD-R, and DVD-RWs is discussed in Modules 7 and 13.



ASIDE

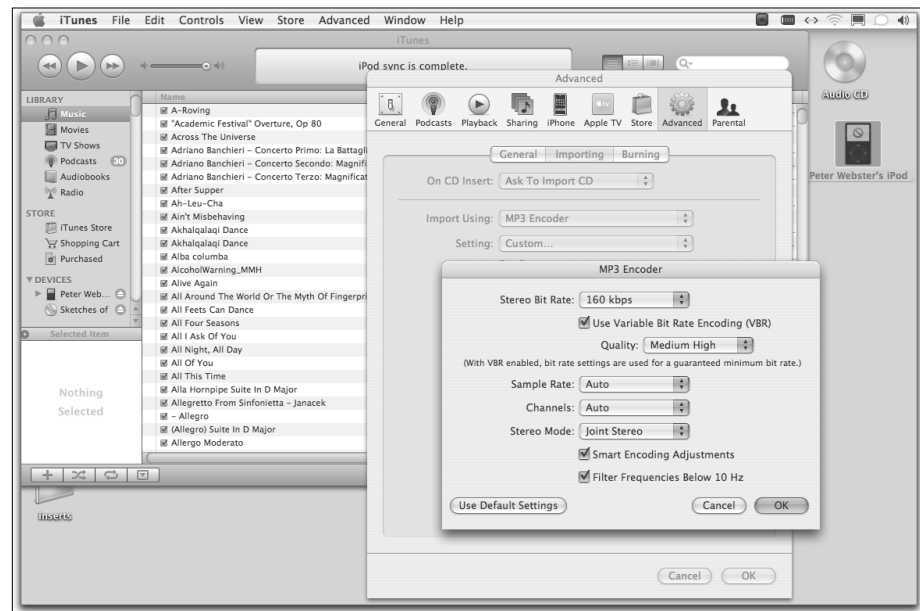
Cddb, one of the most complete databases used by the industry, was begun in 1995 by two software developers who simply wanted a method for maintaining their own CD collection. It is now run by Gracenote, Inc. and serves a worldwide audience.

occurring; however, the original CD would be ejected and replaced by a blank CD-R or CD-RW if you wanted to save an MP3, AAC, or WMA file back to a removable medium. Clicking on the “Burn” button would cause iTunes to seek out a writeable disc drive and begin the process of copying a single or multiple file set to disc. The process is very simple and is integrated into nearly all jukebox MPEG audio software titles.

The other object on the desktop in Figure 9.2 is an icon that represents a portable MP3 player, in this instance an Apple iPod. iTunes works closely with hardware players like the iPod by polling the device to see if the library of MP3 files on the computer matches the portable device. Since the assumption is that you want to synchronize the portable device with the library, it automatically does so, in this case by way of the Firewire serial connection (IEEE 1394, as described in Viewport II, Module 7). Various web music audio software programs support a wide array of CD-burning drives and portable MP3 players, but you do need to check these matches when making purchase decisions.

One last point before leaving this activity of personal MPEG audio creation: online database management. When you first insert a commercial CD into a computer running this type of software, you might have noticed that, if your machine is connected to the net, there is a small pause and, as if by magic, the titles appear for each track together with album information. Is this information on the CD? No, actually the information resides in a remote database that maintains all this information on commercial CDs released by the major recording companies. A unique ID number is sent to the database by your software and the proper information is sent back to be stored with your digital audio data. This is how the MP3 files generate the ID3 tag information without you doing a thing. For this reason, it’s always good to be connected to the net when ripping commercial CDs.

FIGURE 9.2 iTunes



Apple Computer, Inc. (Mac/Win)


Working with Streamed Media

In the previous section, we mentioned how many web music audio software titles can be used as Internet radio receivers. Digital audio (or video) that arrives at your desktop from remote sources in nearly “real time” is a common occurrence on the Internet. Some of these streams of information are, in fact, real radio stations using the Internet as an alternative to the local airwaves. But what is much more common are digital audio broadcasts from archive sources delivered to you on demand. Media content of all sorts is plentiful now on the Internet and will only grow as computers become more powerful and the delivery system more robust.

Streaming Audio in Action

Streaming audio works by sending a file to your desktop with enough data for you to begin to hear it before it all arrives to its destination. In fact, most streaming systems will not even leave the file on the listener’s computer; instead they only provide enough data for the music to be heard at any one time. Streaming can serve many people at once at many different times. All this is made possible by the serving software and advances in compression you learned about in the last module.

Figure 9.3 provides a view of streaming in action. You may have encountered this already in high school or college courses that use technology support. In this case, a music-listening course is taught with the support of a web-based course management system called Blackboard. The teacher has created a page with a link to music by the composer Messiaen. The user has clicked on this link and is now listening to music being delivered from a streaming server located somewhere on campus. With more than 80 students in the course, it is likely that the server is sending this same music to many other students at the same time.



LINK

Return to the last module to review the technical way that streaming works. Pay particular attention to file types such as MOV, WMA, RM, MP3, and AAC.

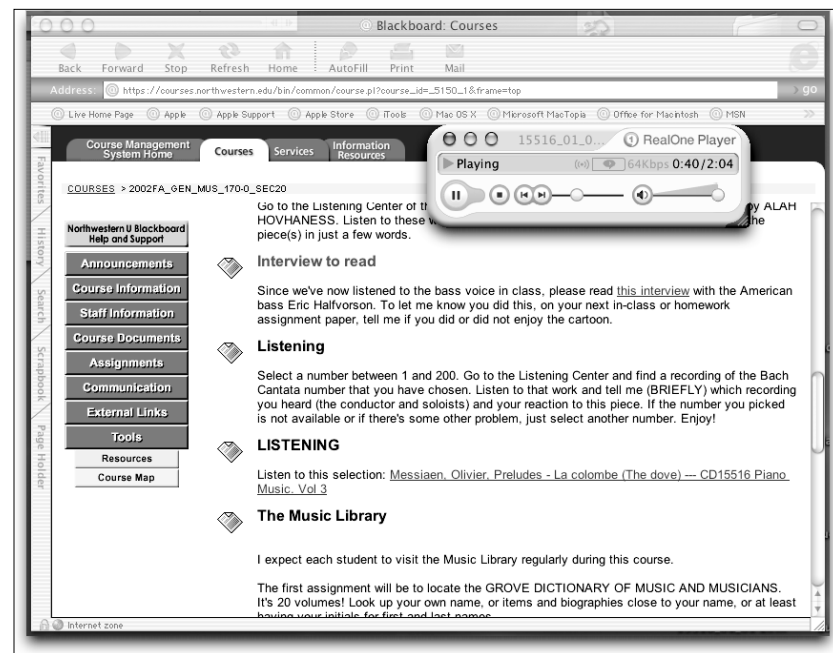


FIGURE 9.3 Course management software Blackboard serving as a trigger for the streaming software RealOne Player

Blackboard, Inc. (Mac, Win); RealNetworks, Inc. (Mac, Win)

Three Systems of Streaming

In addition to MP3, you are likely to encounter three dominant systems when working with streaming digital audio content:

- RealNetworks' RealOne
- Microsoft's Windows Media Series
- Apple's QuickTime

Each system has a player that works on both Macintosh and PC computers and each offers encoding options and server software. The RealNetworks and Microsoft systems generally require PCs as dedicated servers. Apple's solution runs from Macintosh computers. As a consumer, you need to be prepared for each of these systems by installing the free player application as part of the Internet software setup.

Listening to Streamed Media

Except for the MP3 players we reviewed in the previous section, you will encounter streaming audio most frequently as part of links in websites. The players that actually play the streamed media for each of three major systems of streaming may need to be installed separately after your Internet browser of choice is installed. These players can also run by themselves without the browser software.

Making Streamed Media

Each of the three systems provides encoding software to create media for its player. These encoding packages are designed to (1) compress the original file, (2) prepare the file for the desired delivery speed to the user, and (3) provide other options depending on the streaming system. When these steps are finished, the file is ready to be used on a server for hosting websites. The QuickTime Pro version of QuickTime software, for example, provides options for creating streamed media. The other systems do the same.

Preparing Your Computer for Digital Audio Recording

Take a look at Table 9.1 again. Up to this point, we have concerned ourselves with music listening or configuring audio and have done little with actual recording. In the next sections, we will describe basic recording and editing of single mono or stereo tracks and applying effects to the audio signal. We will also describe tasks that may likely be encountered in music studios. These are the tasks noted in the lower half of Table 9.1. Before you can do so, some simple preparations are necessary to make your computer ready to record.

PC Computers

To prepare your PC and its Windows operating system for sound input, you should check the settings for sound management for your sound card and CD. Settings in the Control Panel collection of resources for sounds will help you do this. You can control the volume levels for the CD, microphone, and Line In (if you were going to record into your computer from a tape recorder or another sound source). By

the same token, you can control sound out with sliders. You can “pan” or move the sound from one speaker to another for each setting. These settings are very important for PCs, so be sure to check these before doing basic audio capture, even though your editing software of choice might also be used to make these settings.

Macintosh Computers

Similar settings for Macintosh computers can be found in System Preferences. If the Macintosh has a special audio card, it's likely that the software for that sound card has been installed and it would show up as a resource. A volume-level control is included so that you can check the level.

Using Digital Audio Editing Software

This might be a good time to return to the last Viewport and study the EMT-1 Basic Computer Workstation chart (Figure 7.2) and then glance ahead to the next module to find the EMT-2 No-Frills Digital Audio chart, which adds speakers, headphones, and an MP3 player. There is nothing fancy or terribly expensive here for making digital audio files. These setups are perfect for producing single-track mono or stereo projects for:

- creating files for the Internet
- editing music for a digital movie
- converting, cleaning, and archiving your old tape or LP albums
- making a CD for teaching or entertainment
- producing a professional portfolio for CD or DVD

These specific objectives relate to the kinds of tasks anticipated by Table 9.1 and can be completed at home or in a studio with the help of the software described next.

What Is Digital Audio Editing Software?

The vast majority of software described so far has not permitted actually changing the content of the music. The software titles noted below do exactly this. Here are some common capabilities of such software:

- supporting single mono or stereo track
- recording a number of sound sources, including live, analog recording with microphones; “line-in” feeds from cassette decks, LP record players, or other analog sources; and digital material transferred from CDs
- creating as a minimum 16-bit, 44- or 48-kHz, stereo sound saved in traditional formats such as AIF, WAV, and MP3 or AAC
- allowing digital signal processing (DSP) functions that alter such aspects as dynamics, pitch, tempo, and timbre
- providing protection against altering the original sound file with a system of “nondestructive editing” or multiple levels of “undos”

Note that this is NOT the software to use when you want to record multiple channels of audio in a concert setting, create a numbers of loops for popular music,

combine MIDI data with digital audio, or create a digital sound from scratch. That software will be described later in this text.

Basic Capture and Display of Digital Audio Editing Software

Live Recording

Figure 9.4 provides an example of how this type of software manages to capture and display audio. Here, we are using Sound Studio. In capturing live audio, a microphone is used to record and the signal is displayed as a complex waveform in mono or stereo in the main window. Input levels are monitored in a separate window and can be adjusted if the signal is too loud or soft. Basic programs offer a wide variety of settings for bit depths and sampling rates up to 16-bit and 48-kHz stereo.

Microphone options for live recording are mentioned in Module 10. We cannot stress enough the importance of the best-quality original sound for all your work. It is tempting to settle for a cheap microphone, given the cost of computer hardware and software. This is a false economy in the long run. Knowledge of good recording techniques is a real plus and will make your work with software editors much more satisfying.



TIP

If you are importing sound from an older-style LP record player, be sure to invest in a small pre-amp designed for phono input to boost the signal strength as it comes into your computer.

Other Ways to Import Audio

In Figure 9.4, we have just imported (“extracted” or “ripped”) an audio track from an audio CD. The software is simply directed to the mounted CD and asked to import a track. Programs like QuickTime and iTunes can extract such audio, but digital audio editors will also provide you with a visual display and the tools to edit the audio to your liking.

Most digital audio software editors also allow you to capture audio directly from a “line in” source such as a cassette tape or older LP record. This type of software is an excellent choice for converting recorded audio from an analog source to digital form.

Of course, if digital audio files already exist, these can be opened directly by these programs. Typically, WAV, AIF, and MP3 formats are supported for import.

Typical Display

The central window display in Figure 9.4 can be changed to show greater resolution (smaller time units) and can be enlarged vertically to better show the amplitude. Selection of all or a portion of the waveform allows you to change the content (moving, cutting, or pasting), and by asking the program to “treat” the selected waveform in some way by applying effects, shown in coming sections. Most titles also display counters that show where you are in the file. The software also provides a top window that displays the waveform for the overall sound file.

There are a wide variety of digital audio editors to choose from. A few, like the Audacity software, are offered free or with a small shareware charge. Others, such as Sound Studio and Amadeus Pro and the entry-level versions of Peak and Sound Forge cost less than \$100. Still others, such as WaveLab, Audition, and the full versions of Peak and Sound Forge, cost more. The expensive programs offer wider ranges of options, including more built-in effects processing, extensive plug-in and sound management options, and support for very large audio files. As with all the software profiled in our text, you need to match the software to your task. The more expensive programs may not be necessary for what you wish to do.

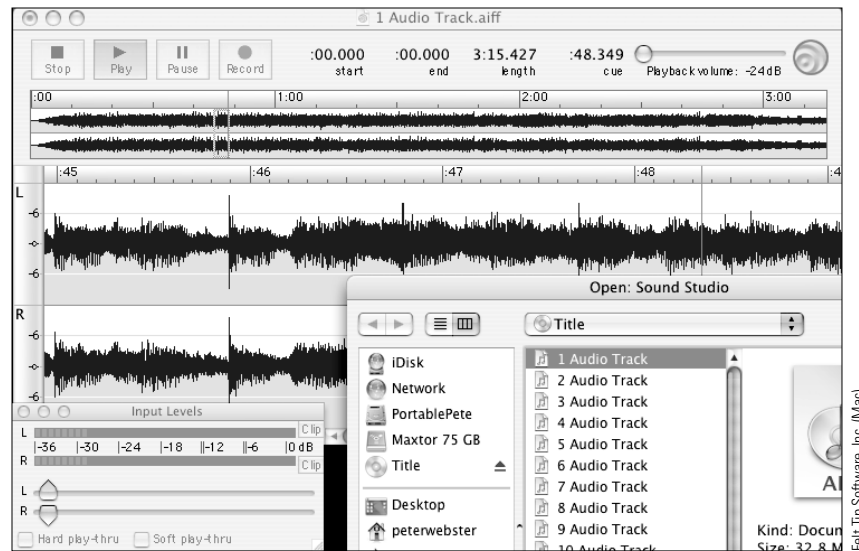


FIGURE 9.4
Sound Studio

To give a better sense of the many options beyond basic capturing and display, we have created a table to guide your study (see Table 9.2). Editing software accomplishes two types of tasks: (1) editing/file management and (2) effects processing. In the first instance, the work is about moving parts of sound files around and managing the results. For effects processing, the work involves changes that are often dramatic and most often associated with digital editors. In Table 9.2, we display these tasks in both basic and more advanced categories. The basic cells of the table contain features that you can expect in any software program of this type. The more advanced items can be expected in the more expensive titles, but there are plenty of variations that might surprise you as you review each program carefully.

Basic Editing and File Management

“Project” Folders and Nondestructive Editing

One additional point about importing and working with digital audio editors relates to the concept of a “project” file set. Some programs, such as Audacity, Peak, and Sound Forge, are optimized to leave the original sound file untouched and to create small editing files that actually record the changes you are making. Of course, in the end you need to save the finished file that “rewrites” the original data, but you can maintain this as a separate end product so that the original data is not altered. This “nondestructive” editing approach, also used in the multitrack recording programs described in Viewport IV, protects the original data and speeds the work process.

Cutting, Pasting, and Manipulating Sound

It is a joy to work with digital audio editing software when you are moving segments of audio around. Because you can “see” the actual waveform in some depth using “zoom in” features, you can find the exact points where an attack or release in the music exists. Selecting the portion to cut to the computer’s clipboard for moving elsewhere or just removing the portion from the total mix is a snap to do. Of course, much more is possible.

TABLE 9.2 Comparison of Basic and More Advanced Tasks in Digital Audio Editing Software

	Editing/File Management	Effects Processing
Basic Tasks	<ul style="list-style-type: none"> • Importing sound (live, “in-line,” CD) • Cutting, pasting, duplicating, slicing, splitting, realigning • Creating clips • Saving (standard formats) • Converting file formats • Other basic features such as mix, swap channels and invert, and resample 	<ul style="list-style-type: none"> • Volume changes (overall gain and normalize) • Fade in, out (linear and custom) • Echo/delay, chorus, and flange • Reverberation (room simulation)
Advanced Tasks	<ul style="list-style-type: none"> • Support for more advanced sound management (e.g., WDM, ASIO, DirectSound, CoreAudio) • Loop-creation tools and links to hardware devices • Waveform editing down to the single-sample level • Markers and playlist creation for CD rendering • Automatic triggers • Menu customization • Batch processing of files • Analyses of waveforms (spectral analysis) • Support for files with high bit sizes and resolution (e.g., 32-bit samples and 192 kHz sampling) and larger file sizes (2–4 GB single files) • Saving in larger number of formats (including RealMedia) 	<ul style="list-style-type: none"> • Noise reduction • Pitch and tempo shift • Cross-looping support • Pan • DC offset • Plug-in support for third-party (VST, DirectX, Audio Units) processing and effects (e.g., EQ, advanced reverb) • Unique effects by program (e.g., harmonization, acoustic mirror)

Figure 9.5 displays an approach that uses Audacity to duplicate and move digital audio in ways that create interesting results. We have selected a portion of a waveform and asked the computer to duplicate it in an accompanying track. Then, we highlighted this new clip of sound and “slid” it to the left so that it sounds a few moments earlier. This kind of fine adjustment in sound can be useful for special effects or for aligning music to exact points in a digital movie or animation. Many of the basic audio editors accept QuickTime or AVI movie file sound tracks for just this kind of task.

Saving and File Conversion

Audacity software includes saving options for standard file types. This program also offers options for the MP3 format. Note, too, the AIF format with track markers. This would be useful if you imagined using an audio track as a trigger for other events.

Digital audio editing programs are also excellent utilities to use for converting files to different formats. If you are developing sound files for the Web, for example, the same file can be saved in different formats for downloading or streaming options.

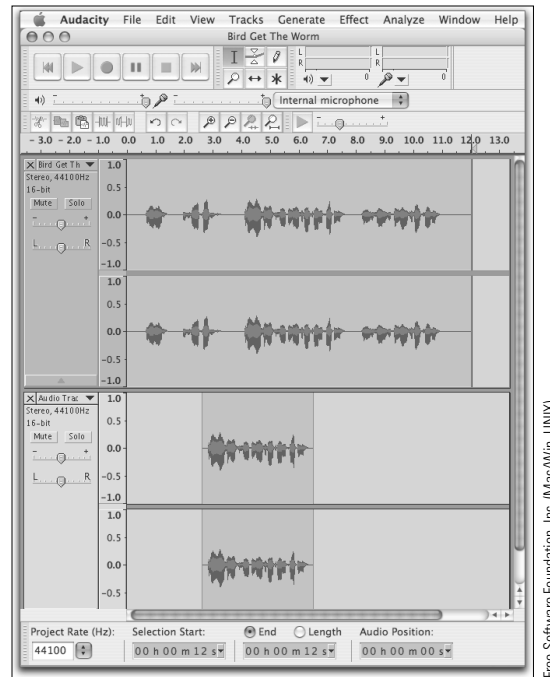



FIGURE 9.5 Copying and pasting in Audacity

Other Basic Management Options

Depending on the software, a number of other options might be expected in basic packages. Here is a listing with brief explanations:

- *Mix*. A mix function takes the contents of the computer's clipboard memory and mixes it directly into a selected region of audio. It does not create a new track, but simply merges the audio information. This might be useful when adding sound tracks to digital movies or simply making the audio more complex.
- *Swap Channels and Invert*. Swapping channels in a stereo signal might be useful in creating a special effect when the channels have significantly different content. Inverting an audio waveform has no effect on sound but does provide a display that places the negative peaks on the top and positive on the bottom, or vice versa.
- *Resample*. Resampling allows for changing the bit size and sample rate. If you are working with 24-bit audio at 48 kHz and you wish to create a standard CD with the format common for conventional CDs, resampling allows you to lower the setting to 16-bit, 44-kHz sound, which is required. Resampling can also make stereo files into mono ones if necessary.



LINK

Lower latency is better! Latency is a measure of how fast information flows between a computer and external devices like digital audio hardware. See Module 10.

Advanced Editing and File Management

More Advanced Sound Management

When analog audio is recorded by the computer and then stored for playback, the application software, the operating system, and the hardware have to establish a way of working with each other to manage the flow. This all falls under the heading of sound management and Macintosh and PC systems each have ways to handle this. Developments in hardware and software design by Apple, Microsoft,

and others have improved the speed of information flow, allowing everything to move faster with very low latency.

This has resulted in the development of standards for sound management such as ASIO, DirectSound, and CoreAudio. In Module 10, we will cover these software “drivers” for hardware in some depth. Here, we need to stress that the more advanced digital audio editors support the faster technology.

Figure 9.6 demonstrates this with the Peak software. The Record Settings dialog box under the Audio menu is where you choose the sound management system and other settings for how the computer is to handle the incoming numbers. Because Macintosh OS X natively employs the newer CoreAudio sound management system, the dialog in the upper right comes up when we click on the “Device and Sample Format . . .” button. Here we can set the way we wish the CoreAudio routines to handle data, including sample rate, channels, and bit depth.

The “Hardware Settings” button displays the choices for what devices to use. In this case, we are using the built-in equipment, but if we were using a different audio hardware interface such as an external card or special USB or Firewire device, we would see these listed here. This is the place, too, where we can set the “buffer size” for the incoming data. The bigger the buffer, the higher the latency (which is best described as the lag between capturing the sound and realizing it through the computer’s output back to us). The smaller the buffer, the more likely the system will miss important incoming information if the computer is processing other tasks. Similar settings are featured in a variety of PC products as well, using ASIO or DirectSound. The point of all of this is that the more advanced audio editors (and many of the more interesting digital audio programs described in coming modules) all work much better with audio when these sound management systems are used.

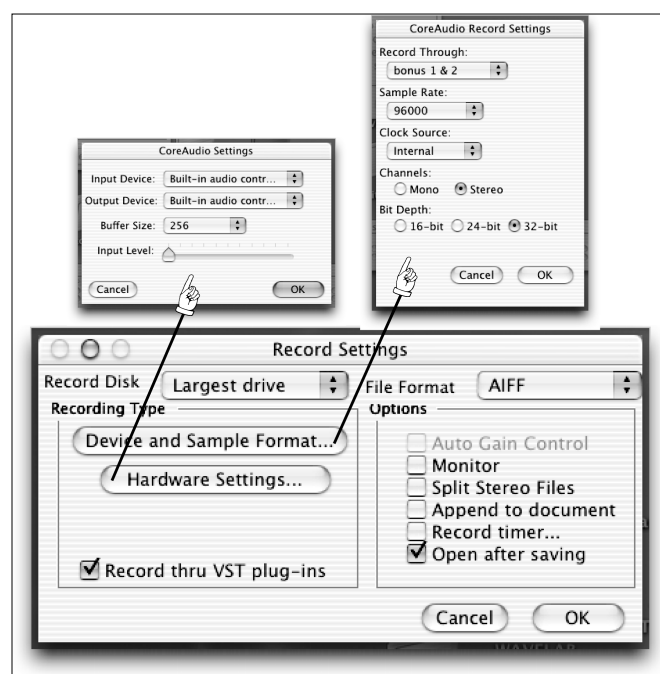




FIGURE 9.6 Advanced sound management in Peak


Berkley Integrated Audio Software, Inc. (Mac)

**LINK**

Wave shaping with sustained sound loops is first discussed in Module 8 in the context of digital wave synthesis techniques.

**LINK**

Check out Viewport V for more about MIDI.

**LINK**

We will present much more about digital audio loops in Viewport IV, Module 12, when we explain “sequential” loops in multitrack recording software.

“Sustained” and “Sequential” Loop Creation and Links to Hardware Samplers

Loops are simply recurring sections of digital audio, often with carefully crafted beginning and end points so that the loop sounds smooth when it is heard. There are basically two kinds of digital audio loops: sequential and sustaining.

Sequential loops are used extensively in music with repetitive structures and are often heard in popular music of many genres. Let’s say you have created a drum track that you want to serve as the basis of a composition. This might be a five- or 10-second sequence of two to four measures’ worth of recorded drum sounds that you want to repeat, over which you would layer a vocal or keyboard track. This kind of loop can be manipulated in many ways; the special software designed to handle this kind of multitrack recording is covered in Module 12. However, editing audio software of the kind described in this module can help fine-tune these tracks by adjusting tempo and beginning and end points for the loop.

Another type of looping is for a short section of digital audio, say a sample of an instrument timbre, that you would like to repeat for as long as a key on an instrument is depressed. Often, these loops last for only a second or less of time and contain envelope properties that sustain similar to that presented in Module 8. This sort of loop is sometimes called a “sustained loop” and can result in an “instrument sample” that can be used by hardware or software samplers. The creation of this sustained loop can be engineered using digital audio software.

Figure 9.7 shows this process using Sound Forge. We have selected a short section of a brass timbre and activated the edit sample window under the Special menu. Notice the ability to fine-tune the loop start and stop locations and the markers inserted into the main data window that mark the sustain and release portions. You can assign a MIDI value that establishes a pitch and fine-tune that value in the box below. The software lets you listen to the loop as you make these changes.

Once these loops are created, advanced software can send these digital files to hardware samplers. Both Peak and Sound Forge have links to allow transfer of data to hardware devices.

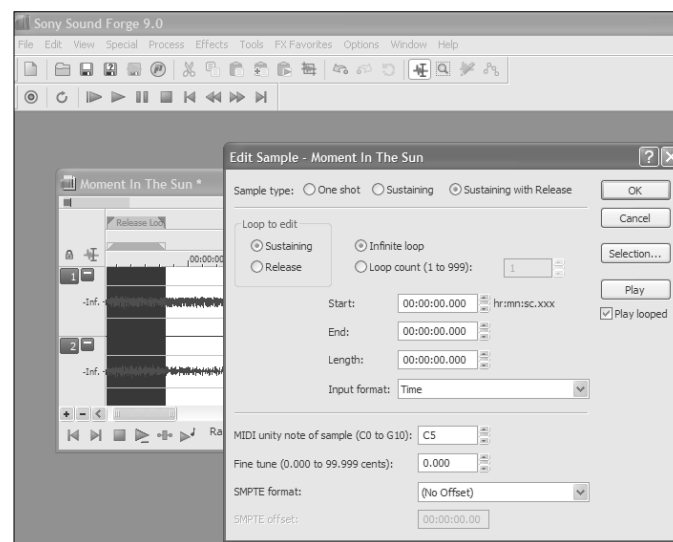


FIGURE 9.7 Sample editing with Sound Forge 9

Sony Pictures Digital Inc. (Wfn)

Waveform Editing

In addition to being able to control the start and stop points of a loop, other detailed editing is possible in any waveform. Figure 9.8 shows the ability of advanced software to actually control the waveform down to the most basic level with a pencil tool. Such editing capability allows you to modify the crossover points of a loop and edit a waveform in general if, in doing so, you are interested in altering sound for a particular reason. This might occur, for instance, if there is an unusual disturbance in a live concert recording and such editing might eliminate the phenomenon from the waveform.

Markers and Playlists

Another interesting way to manage digital audio is to create regions in an audio file for the purpose of creating separate sections. Imagine that you are preparing a digital audio file for presentation in class and you want to be able to isolate sections of audio so you can highlight your presentation by playing each individual section. Sections can be named and a playlist created for a custom CD that can help with your presentation.

Additional Editing and File Management Features

The following lists other advanced editing and file management features of digital audio editing software. Taken as a group, these features might well justify the cost of a more advanced program:

- *Automatic triggers.* Software can be set to start recording when a certain sound level is reached, or stop or pause at other designated times. Also, the software can be triggered by the computer or MIDI keyboard.
- *Customized menus.* Menus can be easily changed to meet the needs of projects that require repetitive tasks.
- *Batch processing.* A large number of files that need to be treated similarly with a series of operations (such as changes in file type, resampling, filtering, or volume change) can all be changed at one time.
- *Spectrum analysis.* Waveforms can be viewed as a spectrum in order to study the properties of the audio for noise problems (see Figure 8.2b in Module 8).

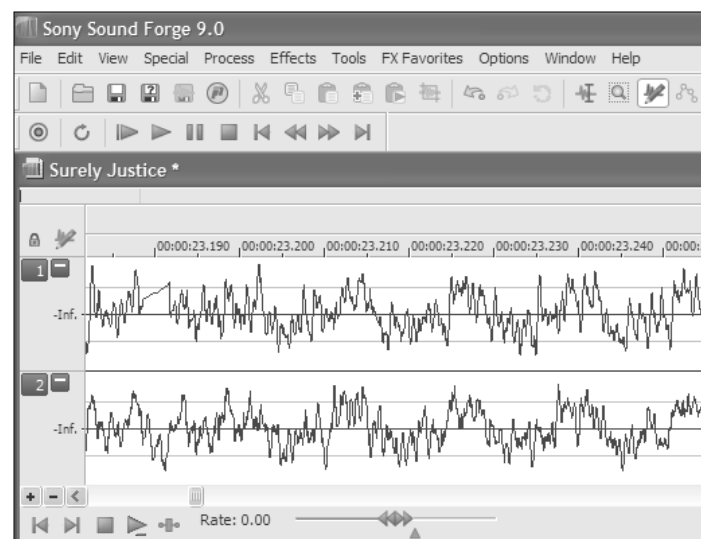


FIGURE 9.8 Editing a waveform with Sound Forge 9

- *File support.* Advanced programs can handle file resolutions equal to and in excess of 24 bits and handle sampling rates from 48 kHz to 96 kHz or higher. There are much higher limits to file size as well, often accommodating files as large as 4 gigabytes or close to six hours of CD-quality stereo sound.
- *File format.* Advanced editors can import and export far more types of file formats than basic software. Sound Forge, for example, supports some 15 file formats for import and 17 for export, including encoded compressed formats such as RealAudio, Windows Media, and MP3. This means that a program like Sound Forge can be used to master content for stand-alone or streamed formats without the need for additional encoding software.

Effects Processing

Up to now, we have focused on basic and advanced work that involves editing and file management. In these final two sections, we focus on changing the sound itself through effects processing. This was once only possible with specialized hardware devices or digital signal processing (DSP) chips on add-on cards. Such hardware solutions continue to be available for special requirements in high-end studios; however, most effects processing today can be handled completely within software such as the audio editors described in this module and in multiple track (Module 12) and sequencing (Module 17) software.

Study the common types of effects processing as described below. We have grouped these in five families with the first two (amplitude and time) being the most common and the others (frequency, timbre, and specialized) being more advanced. In the remaining pages of this module, we describe how some of these effects are used in basic and more advanced audio editing programs. In other modules to come, we will return to this topic in more detail.


- **Amplitude:** gain, tremolo, normalization, fading, cross-fading, compression, expansion, envelope or ADSR manipulation, panning
- **Time:** reverb, echo, delay, flanging, chorus, time shifting
- **Frequency:** various wah effects, phase shifting, pitch shifting
- **Timbre:** noise reduction/addition, distortion, clipping control, limiting, rectification, equalization, resonator
- **Specialized Effects:** harmonizing, vocalizing, ring modulation, Leslie speaker rotation

Amplitude

Perhaps the most used effect with digital audio files is amplitude change. Such activity takes the form of changes in the overall signal and changes based on fading. Amplitude changes, as most of the effects processing described here, are really accomplished by applying certain software-designed filters to the original sound. This is much like what was done in the early days of electronic music with hardware (see Module 8), but it is now accomplished with software.

Typically, a region of audio or the complete file is highlighted and an amplitude or “gain” change is executed for the selected region or file. The software displays a dialog box that contains the controls for raising or lowering the amount of gain. “Clipping” (exceeding the threshold of loudness, which creates distortion) is often controlled automatically by the software if you choose this.

A special kind of amplitude adjustment is called “normalization.” Normalized sound has been boosted in amplitude only after the selected waveform is analyzed



LINK
See Module 8 for a full explanation of decibels and the scales that govern it.

to determine the highest peaks of sound. The waveform is adjusted to be as loud as it can be without the highest peaks being distorted. The entire waveform is adjusted accordingly.

Two other effects are “compressor” and “expander.” These filters boost the lows and compress the highs in such a way as to alter the overall waveform. This may be useful when you want to have a very quiet section in a performance boosted or a loud section reduced to help create a better balance.

Figure 9.9 displays two applications of amplitude change. The highlighted section of the waveform to the right has been changed by the Amplify Volume dialog box below it. Notice that the volume is changed by about 5.1 dB, which is about an 80% increase (100% is the base figure of no gain in this software). It does not take much to make a significant gain in sound intensity. The Preview button allows you to hear what you are doing. The Blending option is really a kind of cross-fade function in this software that blends the volume change in and out with the material around it by a factor specified in milliseconds. The section of audio to the left has been normalized by just under 100%. Notice that the waveform fills the available bandwidth to the 83.2% level and that everything is adjusted in relation to the highest peaks. Normalization is treated separately for each channel.

Fading audio in and out is a typical adjustment in basic effects processing. There are countless uses for this in digital audio editing and it is accomplished in a similar way to changing volumes, described above. The section of audio on which the fading effect is to be applied is highlighted. You then tell the software if you want a linear adjustment (consistent rate of fade from start to finish) or if you want some kind of nonlinear shape such as a slow or fast rate. Figure 9.10 shows a custom,

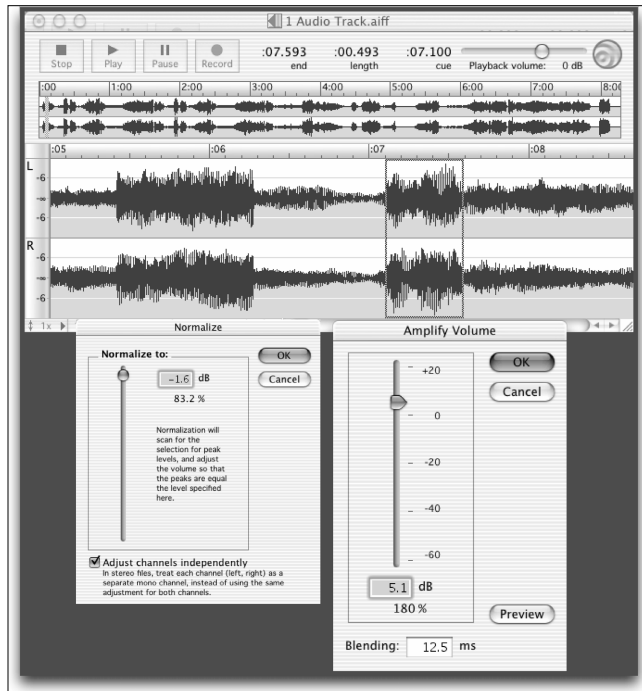


FIGURE 9.9 Amplitude adjustments in Sound Studio

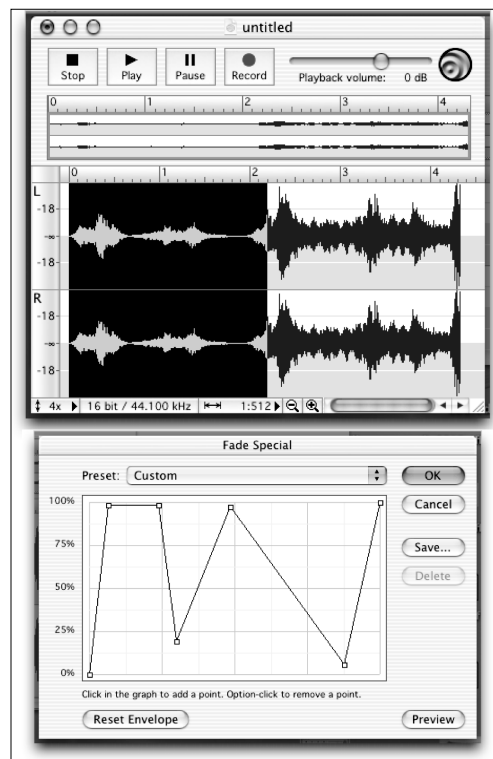



FIGURE 9.10 Custom fade in Sound Studio

Felt Tip Software, Inc. (Mac)



LINK
 These concepts have their roots in the early days of electronic music and the manipulation of analog sound with oscillators, filters, and patch cords. Be sure to link this with the material in Module 8.



TIP
 LFO = low-frequency oscillator

nonlinear shape. The data window displays the final fading envelope and the box below it has the envelope as we defined it. The more elaborate fades that are possible in the context of loop construction are noted in the following section.

Panning, or moving the audio signal's intensity from one stereo channel to another to give the illusion of moving the sound in space, is very common in digital audio editors. This effect can be found in many audio applications, including sequencers. It takes on a new life in surround-sound work, when more than two audio channels are involved; this is covered in more detail in Module 12.

Time

Many of the effects in this family work much the same way but each has a distinctive sound, depending on the settings chosen. To understand them, you need to understand the concept of “wet” and “dry” sound. “Dry” sound is the original audio and “wet” sound is the same audio fed back into the signal at a different time and intensity. Actually, the intensity of both wet and dry sound can be controlled in all three of these effects.

Study Figures 9.11a and b carefully. Echo/Delay and Chorus are displayed in Figure 9.11a. For the Echo effect, a delay time is specified. You can have the delayed audio take output as its input (feedback), creating an echo effect. Without the feedback checked, the signal would just have a delay effect in which you hear only one repeat.

Chorusing is more complex. Here, the idea is to offset the signal a bit to create a fuller and perhaps slightly out-of-tune signal. The wet signal, which is slightly delayed, is controlled by the LFO waveform shape, with the sine wave being

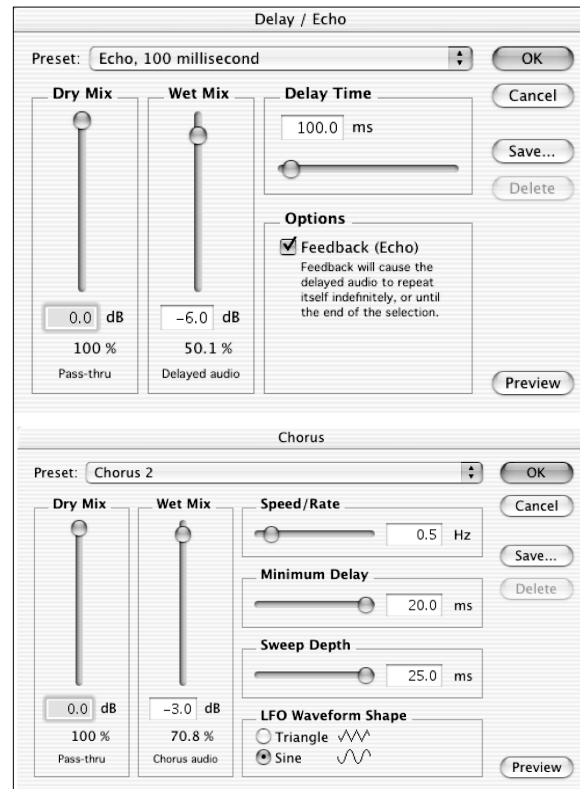


FIGURE 9.11A
 Echo/Delay and Chorus in Sound Studio



TIP

Good sound-reproduction hardware in the form of quality headphones or loud speakers is necessary for this work, even with very basic EMT-1 and 2 setups.

smoother than the triangle wave. The Speed/Rate scale controls the magnitude of the LFO's effect on the wet sound. Minimum Delay and Sweep Depth also control the timing and out-of-tune quality. Describing these settings in words is difficult; they must be experienced musically with various dry sounds as original source material.

Figure 9.11b demonstrates the Flange effect. This is very much like the Chorus effect but has shorter delay times. A sweeping whooshing sound is created by this effect and the rate and depth settings control the effect. Again, the effect must be used with real music to get a sense of the sound. Not all of these effects are included in basic software and each program treats them differently.

Reverb is often confused with echo/delay, but it really is different. Rather than combining wet and dry sound in a delay state, reverb recreates the acoustical properties you experience as a listener in a closed space. The effect tries to simulate the properties of live audio as it is heard bouncing around a room from walls, ceiling, floor, and the source itself. When you record live audio into the computer or import sound from another source, it may lack a sense of “depth” or “warmth” that you have come to expect in a live concert. Reverb adds subtle touches to recreate this sense of reality.

Each software program provides its own way of creating these artifacts in the sound. Figure 9.12 shows Sound Forge's own reverb effect. To experiment with the settings, an MP3 file that seems to warrant a “warm space” is loaded in. Adjustments in reverb amount, room size, and liveliness relate to how filters work to adjust the highs and lows of the signal while adding small amounts of delay for different parts of the spectrum. Using these settings is not an exact science and, as with all of the human judgments about effects processing, much depends on your ear and the experience you have had with sounds.

Advanced Effects Processing

The more advanced digital audio editors often contain more powerful built-in effects. As we have noted before, not all of these effects are found in expensive

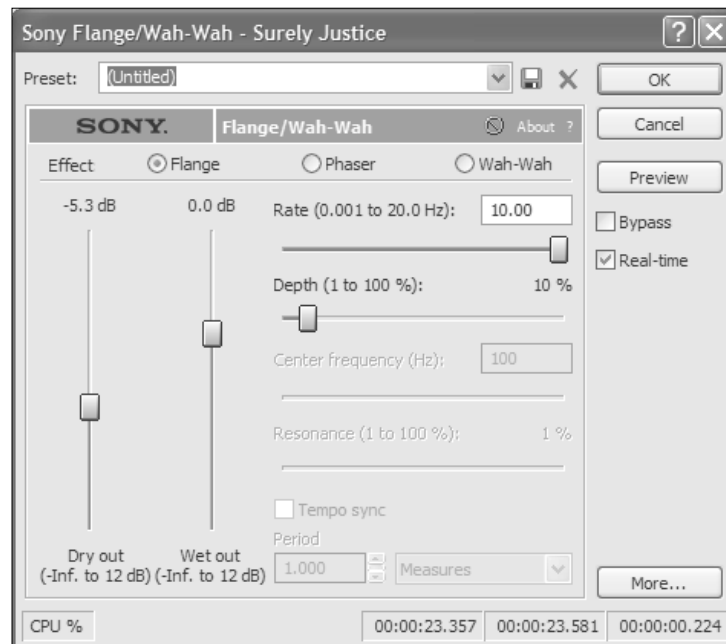


FIGURE 9.11B
Flange in Sound Forge 9

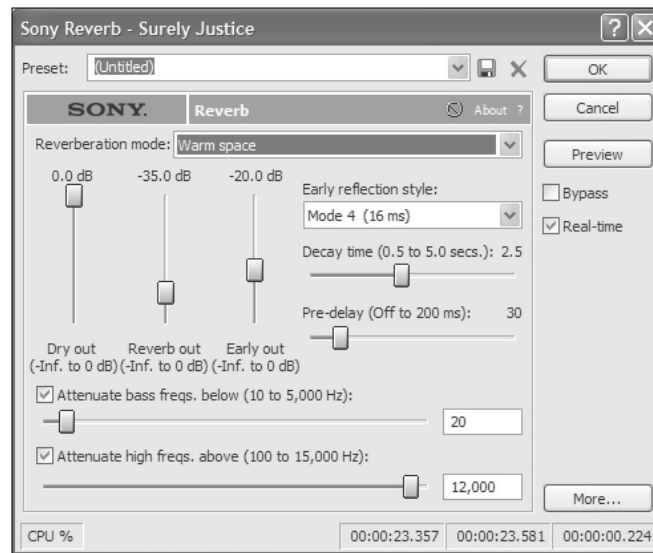


FIGURE 9.12 Reverb dialog box for Sound Forge 9

programs and not all effects are equally well engineered. Experimentation with your own ears is key.

Frequency and Time

A more advanced feature set involves shifting both pitch and tempo of selected regions in a waveform or entire files. There are many reasons why each of these procedures might be useful. For example, working with a sound track for a digital movie, you may be interested in stretching or constricting the time to fit the video image. If you are developing audio as an accompaniment for teaching, you may want to lower or raise the pitch to accommodate a voice range or an instrument's key.

Figure 9.13 provides an example of how this might be done. The example shows a dialog box for settings to stretch a region of audio. At the bottom of the dialog box, notice the options for stretching time but preserving the pitch and for shifting the pitch but preserving the time (tempo)! This notion of changing one dimension and keeping the other constant may seem trivial, but it is very powerful. This can be especially useful in working with digital audio. Pitch change is also possible.

Timbre: Noise and Silence

Both removing and adding noise and silence are important capabilities. One popular reason to invest in more advanced digital audio editing software is its ability to “clean up” a messy audio file. Many people own older recordings on cassette tape or even long-playing (LP) vinyl records. These sources often contain valuable, in some cases irreplaceable, data; however, the recording also contains background noise, tape hiss, and distracting crackles and pops.

Figure 9.14 shows a solution. The figure portrays noise removal for a selected region. This approach requires that you first “teach” the software what you mean by noise and then ask the program to filter a region based on that intelligence. Other advanced software programs use different techniques, including support for stand-alone programs such as plug-ins that use quite unusual and powerful algorithms to clean up sound. Many of these are quite expensive and are designed for highly skilled sound engineers.

FIGURE 9.13 Tempo shift in Audition

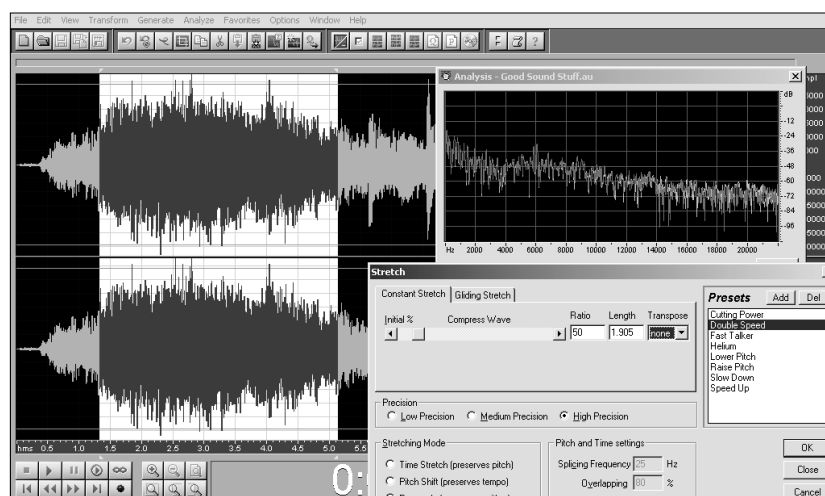
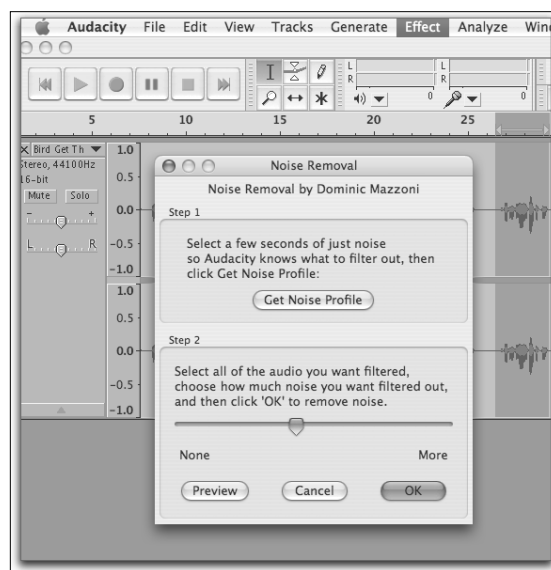


FIGURE 9.14 Noise removal in Audacity



Adding noise and portions of silence are two additional features possible with digital audio editors. Adding a touch of distortion may be useful in certain instances and inserting silence is often needed when dealing with sound tracks for film or television.

Timbre: Equalization

Another very important timbre effect is equalization (EQ). Equalization helps balance the frequencies within an audio spectrum by filtering out certain specified ranges of the spectrum. Figure 9.15 displays how this is done in Sound Studio with a 10-band EQ. More bands provide greater control over the frequency range; here we have chosen to cut the mid-range of a selected part of an audio clip. The sliders allow you to control the filtering level for the particular band. Module 17 has an example of a much more sophisticated EQ plug-in.

Advanced Amplitude: Crossfade Loop

A particular kind of fading effect useful in sustained loop construction is the crossfade loop. This is an advanced technique tied to the creation of loop start and end points. The shape of the crossfade can be controlled and the time can be specified

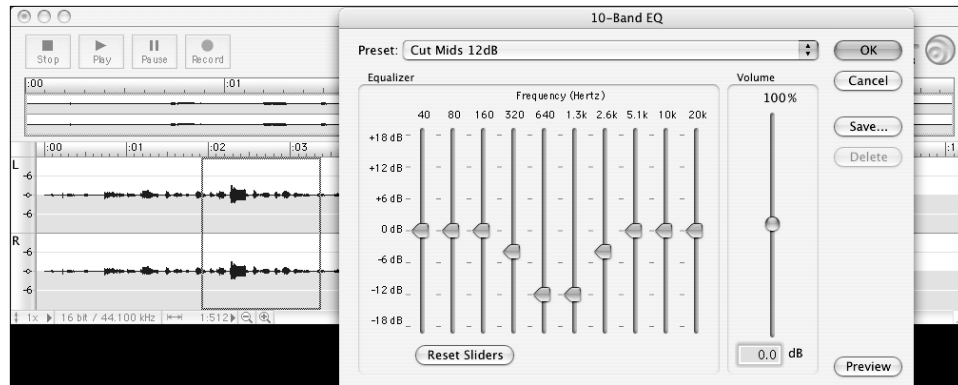


FIGURE 9.15 EQ adjustments in Sound Studio

as well. Most software lets you preview your choices so you can use your ear to test how the fading will sound in real time.

Specialized Effects

Certain advanced effects are peculiar to one program or another. For example, in WaveLab, a harmonization effect is included. This interesting effect creates a harmonic accompaniment when applied to a melodic line. You will see how music notation programs using MIDI can also create a harmonic context in Viewport VII, Module 20.

Another example of a rather unique effect is the Acoustic Mirror in Sound Forge. Related to the family of reverb effects, Acoustic Mirror uses real acoustical data from experiments within certain famous concert spaces. This information gets stored in a special formatted file with the .sfi extension. These files can be downloaded from the Sony Media Software website and found on the application

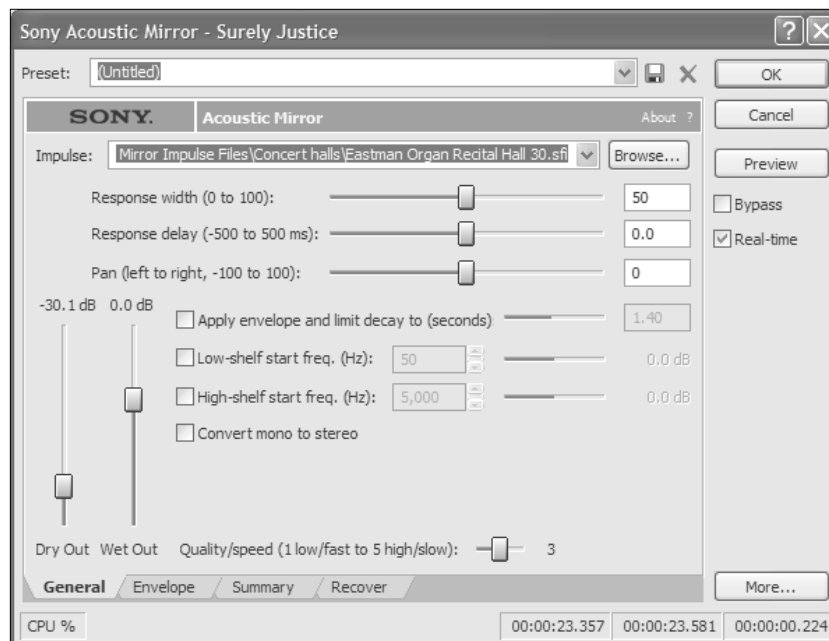


FIGURE 9.16 Acoustic Mirror in Sound Forge 9

SIDE BAR

Audio Utility Software

In addition to built-in and plug-in support for digital audio, a few audio utility programs are available as adjuncts to digital audio editing packages. These programs are distributed for minimal cost and offer special capabilities that might not be available elsewhere. Here are a few of these programs and a short description of what they do:

- Soundhack (Soundhack, Inc. [Mac]) Provides many processing and utility functions, such as time stretching, pitch shifting, and sonogram creation. Plays and converts many sound formats and records input from built-in Macintosh input. Changes values in the sound file header.
 - SoundConverter (Steve Dekorte, www.dekorte.com [Mac]) A batch sound format conversion program that supports many input and output formats, including ringtone formats.
 - Amazing Slow Downer (Roni Music, Inc. [Mac/Win]) Slows down music from 50% to 400% without changing the pitch. Works from a CD and supports MP3, AAC, AIF, and WAV formats.
 - Transcribe! (Seventh String, Inc. [Mac/Win]) Also slows down music without changing pitch. Supports AIF, MP3, and WAV files.
-

CD. When applied to a region of audio or an entire file, this effect can be used to simulate what the audio might sound like in that real environment. Figure 9.16 shows a dialog box that has a favorite concert hall featured.

Other effects specific to one program or another include:

- Ring modulator: multiplies two audio signals together, creating a somewhat “metallic” sound
- Phase Vocoder: allows you to alter the duration and/or pitch of an audio region
- Rappify: applies extreme dynamic filtering to a selection
- Reverse Boomerang: combines an original source with a version of itself played backwards

Plug-in Support for Digital Audio Editors

Plug-ins are small programs that work in tandem with larger “host” programs like Peak, Sound Forge, and WaveLab. They function in many ways, but most are designed to offer advanced versions of effects not available in the host program. We will provide much more detail on plug-ins and their formats in Modules 12 and 17.

Plug-in software for audio editors is typically placed in a folder in the same directory as the host program. Once there, the software becomes “available” to be used as part of the processing system of the audio editor.

Plug-ins can be used in combination. This allows the host software to use the processing power of two or more plug-ins (even from different vendors, if they conform to an accepted plug-in format such as DirectX or VST) together to make simultaneous changes in audio. For example, Figure 9.17 shows a reverb effect

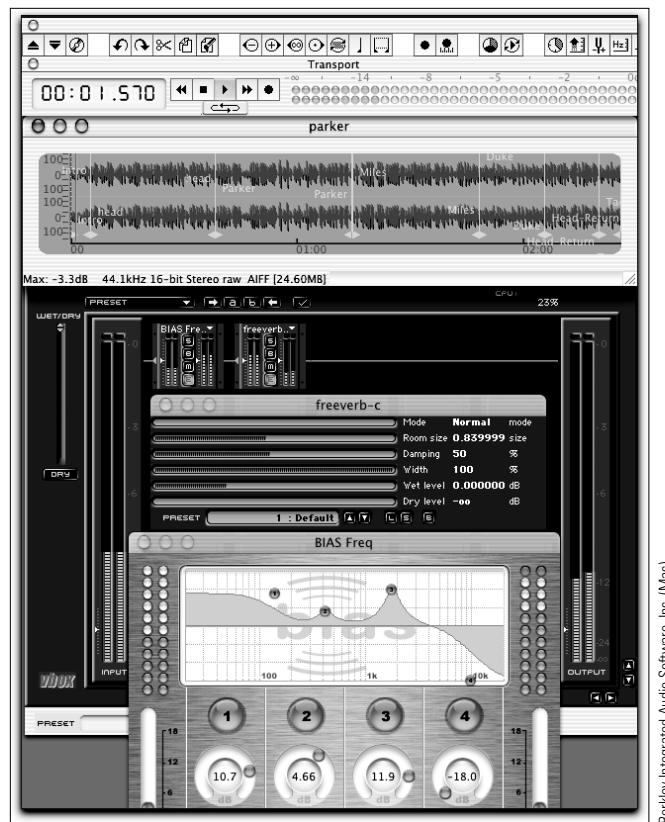


FIGURE 9.17 Multiple plug-ins in Peak

used together with an equalization (EQ) effect in the host program Peak. EQ is an effect that enables you to highlight frequencies across the spectrum. You probably use this option in stereo systems that use slider bars to adjust various levels of treble, mid-range, and bass. In the Peak example, notice the ability to change the amount of the effect “in play” in the small boxes in the center of the graphic. Many other features are offered by combining plug-in effects. Keep in mind that, as multiple plug-in effects are used, more powerful computer processing is necessary.

Berkley Integrated Audio Software, Inc. (Mac)