

Producing and Adding Sounds

Adding sounds to your Web pages is not particularly difficult. There are advantages to adding sound to your pages, but equally good reasons exist not to put sound in your pages. This chapter talks about how digital audio works, what sound formats work on the Web, how audio compression works, how streaming audio differs from regular audio, how to obtain and record sound files, how to use sound editing software, and, finally, how to add sound to your page with the `OBJECT` element and the `EMBED` element.

Understanding Digital Audio

Digital audio tries to represent sound, which is inherently analog (being a wave), with ones and zeros. When you digitize sound waves, you take samples of the wave at certain intervals. Just as you select the quality of images you scan by setting the resolution, you can select the precision of the sound captured by setting the bit-depth of the sound file. Eight-bit sound divides the spectrum into 256 levels. Sixteen-bit sound gives you 65,536 levels. There is also 32-bit sound, but it does not run over the Web and it is not highly accessible to client workstations because so few truly 32-bit sound cards exist.

The quality of sound your computer produces is dependent on five things:

1. The quality of equipment used to record the sound
2. The bit-depth of the recorded sound
3. The frequency of sampling (sample rate)
4. Whether the sound is recorded in mono or stereo (the number of channels)
5. The quality of equipment in the machine playing the sound back



In This Chapter

Understanding digital audio

Introducing sound file formats

Introducing streaming audio

Obtaining sound files

Recording sound files

Editing sound files

Compressing sound files

Adding sound files to your page



As with images and other types of multimedia files, the quality has to be balanced against the file size. You want to deliver the maximum quality sound for the minimum file size. Compression can help to some degree.

Recording equipment

Several factors figure into the quality of the sound you record. One is the recording environment. If you are going to record voice, you'll want to record in a room with good acoustics. You also need to be sure the bit-depth of the recording equipment is at least as high as the bit-depth at which you are recording. A smart system won't even let you *try* to record 16-bit stereo sound on an 8-bit mono system.

Recording bit-depth

The *bit-depth* of the sound recording indicates how precisely the sound that occurs at the moment of sampling is recorded. As with color palettes, the number of sounds you can choose from is limited to the palette you select. For 8-bit sounds, the sampled sound will be one of 256. For 16-bit sound, the sampled sound will be one of 65,536 sounds. Table 40-1 shows bit-depth and sample rates for common types of digitized sound.

Table 40-1
Sample Rates and Bit-Depth for
Common Types of Digitized Sound

<i>Quality</i>	<i>Sample Rate (kHz)</i>	<i>Bits per Sample</i>	<i>Mono/Stereo</i>	<i>Uncompressed Data Rate (Kbps)</i>
Telephone	8	8	Mono	8
AM Radio	11.025	8	Mono	11.0
FM Radio	22.050	16	Stereo	88.2
CD	44.1	16	Stereo	176.4
DAT	48	16	Stereo	192.0

Sample rate

Sample rates are measured in kilohertz (kHz). The more frequently you sample sound, the higher the sample rate, and the more your digitized sound will match the original sound. Of course, the more frequently you sample, the more information you have to store, resulting in a larger file size.

Mono versus stereo

The difference between mono and stereo is the number of sound channels in the file. For stereo, there is exactly twice as much data. Why does twice as much data exist? Because you are sending two channels of information: one for the left speaker and one for the right speaker.

Do you want to use stereo sound? This depends on two things: file size and likely equipment of the end user. Do you expect site visitors to have stereo playback equipment? If not, it doesn't matter. The real question is whether the sound you are delivering is audio for its own sake, in which the sound quality really matters, or whether the sound is part of something bigger, in which case, mono is probably fine.

Playback equipment

All this business about bit depth, sample rates, and mono versus stereo don't mean much if the people trying to play back your sound have old or low-quality sound cards. If people playing back your sound don't have a 16-bit sound card, they can't possibly play back all 65,536 sounds in the 16-bit spectrum. Their computers will automatically choose the next closest sound from the 256 sounds in the 8-bit spectrum. Bad speakers and a slow processor — that can't re-create the sound at the proper sample rate — can also result in inferior sound quality.

Introducing Sound File Formats

Four sound formats frequently used on the Web include the following:

- ♦ **.WAV** is the original Windows file format. Although it used to be proprietary to Windows, you can actually play it now under any of the newer browsers, even on a Mac or on UNIX.
- ♦ **.AIF** (or **.AIFF**) stands for Audio Interchange File Format. This common audio file format is also cross-platform, but doesn't take advantage of any compression.
- ♦ **.AU** is the original Sun standard; it enables you to make tiny sound files.
- ♦ **.MPEG** and MPEG Audio Layer 3 (**.MP3**) are part of the MPEG2 standard. With it, you can create small files while maintaining a pretty high sound quality. It is nonproprietary. You can achieve compression rates of 1:10 or 1:12 with CD-quality sound. On the downside, it does require a special player. A streamworks player will play streamed files. A lot of record companies are using the MPEG format because you can put CD-quality files on the Web. You also need special software to compress your audio files into **.MP3** format.

Other ways exist to put sound on the Web, including Shockwave, RealAudio, and QuickTime, but they are proprietary. They require plug-in software (as does .MP3, although it is nonproprietary), which can inconvenience your site visitors enough that they might not bother to get the plug-in or to try to hear the audio files you put so much time into creating.

Introducing Streaming Audio

Streaming audio has caught on so well that now everyone wants to call their product *streaming*. Consequently, two possible things are meant by the term *streaming audio*. Both kinds of streaming audio enable the visitor to your site to begin hearing the audio file before it finishes downloading, which is the usual way you listen to audio files on the Web.

The first involves a special server that delivers the audio over the Web using a proprietary protocol (not HTTP) to get it to the desktop. It also requires a special plug-in for the browser. The most popular streaming audio server on the market is the RealAudio server. You can learn more about it and about the technology in general at the RealAudio Web site at www.realaudio.com.

The other kind of streaming audio does take advantage of the HTTP protocol, so for lack of a better name, it is called *HTTP Streaming*.

Streaming audio is a big improvement over traditional audio, because the sound file can begin to play before it has completed downloading. Normally, a sound file won't begin playing until the entire file has downloaded. This was part of the reason for the initial impatience with audio on the Web.

How does streaming audio work? When you save a file for streaming, all the basic file format information is front-loaded in the file. This means the first things the browser computer receives about the audio file is enough information to begin playing immediately. In a nonstreaming file, the audio file information is spread out along the length of the file, so the receiving computer can't begin to play the file until the entire file is downloaded.

If you are going to use either kind of streaming audio, you might also have to make changes to the MIME type on your server. Check with your systems administrator to see if your server can handle the MIME type you will be using.

Obtaining Sound Files

You can legally obtain public domain digital music at some sites, but many of the sources you might consider using can get you into trouble with lawyers. You may not obtain sound files by capturing music off your favorite CDs, the radio, or cassettes. You may obtain sound files by recording your own musical works or your own voice, or by purchasing licensed production music or sound effects.

Will you get caught if you use music you obtain illegally or you even own legally, but you distribute illegally (and this is what you are doing when you publish music, even music from recordings you've purchased) on the Web? Probably. There are Web agents — software programs that crawl around the Web looking for sound files belonging to different recording companies — and these agents have been successful at locating illegally distributed music being published by even the most harmless organizations.

Being hassled by a big-name law firm in New York City with a cease-and-desist order is only half the trouble. If you are using music because you think it adds something spicy to your site — say, the theme music from *Mission Impossible* — it will be obvious to visitors to your site, unless your site appears affiliated with the movie in some way, that you have stolen the music and are using it illegally. Does this make your site look professional? No.

Recording Sound Files

If you choose to use a professional sound-editing tool, such as Sound Edit 16, you can record directly into that. Or, you can use whatever sound-recording software comes installed with your sound card. You do need a sound card and an input jack or CD-ROM. Because you know you shouldn't record music off your CD, you won't actually be using any music from your CD collection, but you might be using *clip sounds*, which are like clip art in that you can reuse them if you purchase the CD on which they are distributed.

Figure 40-1 shows the screen for Sound Edit 16 when you are recording sound. You can set the sample rate, the bit-depth, and mono versus stereo. The controls for record, stop, pause, and so forth should look familiar, if you have ever played a CD on your computer.

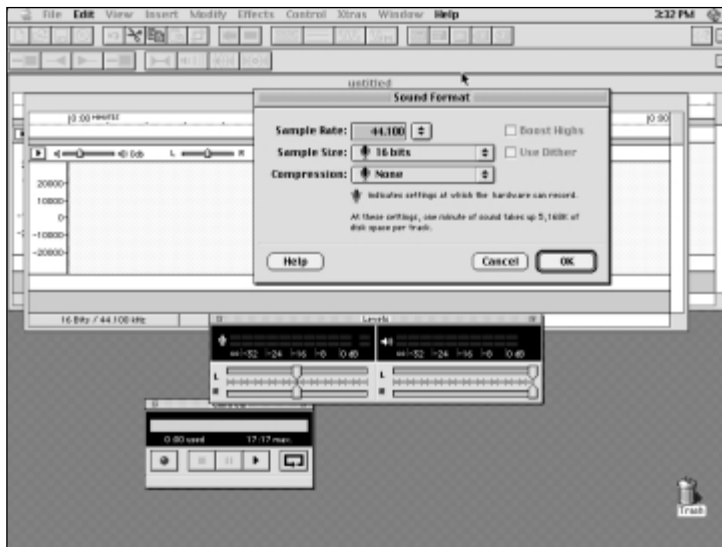


Figure 40-1: Recording sound with Sound Edit 16

Editing Sound Files and Adding Filters

Sound-editing software, unlike image-editing software, isn't usually cheap or free. One free package runs on the Mac, called SoundHack, which does a respectable job. One of the more reasonably-priced packages is Sound Edit 16, by Macromedia. What can you do with your sound files? You can add noise, remove noise, and add filters to create effects such as if you were singing the national anthem in a football stadium (or in space!).

Sound Edit 16 makes it easy to add special effects to your sound files. Just as easy is removing loud breath sounds from a voice recording and adding pauses. Figure 40-2 shows one of the filters you can apply in Sound Edit 16.

Compressing Sound Files

Once you have made all the trade-off decisions about your audio files, bit depth, sample rate, and mono versus stereo, you can still use compression to reduce the file size. Many compression schemes are available. Probably the two seen most often are MPEG Audio Layer 3 (also known as MP3) and RealAudio.

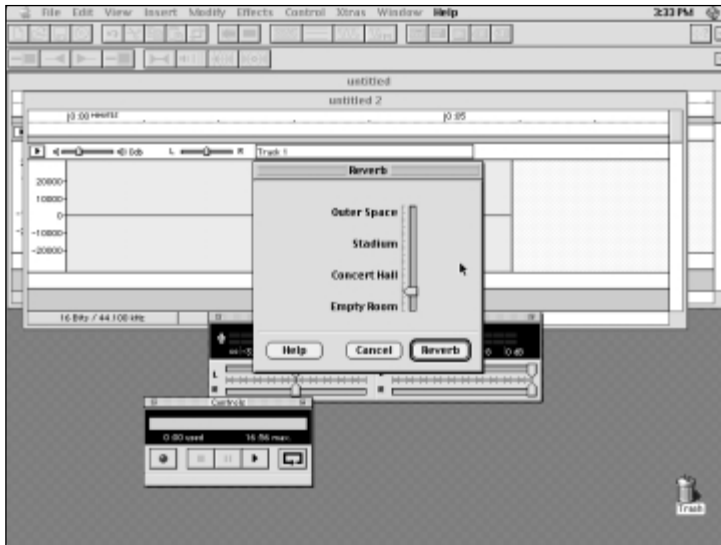


Figure 40-2: Adding effects with Sound Edit 16

RealAudio is proprietary, but is widely used and the compression tool is free, as is the plug-in. MP3 is not proprietary. The MPEG standard is probably the future of free audio on the Web. Their Web site (www.mpeg.org) is a treasure trove of information about audio compression and video compression, along with tools for delivering both on the Web.

Other compression standards exist, but they are not as compatible across platforms and your site visitors might have trouble playing the files back.

Adding Sound Files to Your Page

You have two choices for adding sound to your page with HTML: You can add inline sound, which starts to play as soon as your page loads, or you can add out-of-line sound that plays as a result of some action on the part of the visitor.

Inline sound

The official HTML 4 specification indicates that inline sound should be created with the `OBJECT` element, but as of publication of this book, the `OBJECT` element did not work to deliver sound in any of the major browsers. You can add inline sound using the `EMBED` element (which is deprecated, but works with all major browsers), with the value of the `src` attribute being the URL for the sound file. You'll also want to

use the `loop` attribute to indicate how many times to play the sound in the background. If you want the sound to play ad nauseam, you can set the value of `loop` to *infinite* (which should probably be *ad nauseum* to make it clear to Web authors that this is the effect of playing a loop of sound an infinite number of times).

```
<EMBED src="media/sound/annoying-beep.mpeg" loop="infinite"
autostart="true">
```

Or, the official way, with the `OBJECT` element:

```
<OBJECT data="media/sound/annoying-beep.mpeg"
type="application/mpeg">This is an annoying beep.</OBJECT>
```

Out-of-line sound

The alternative to inline sound is the preferred method of delivering sound, which is delivering sound only when the visitor to your page requests it. Many record stores on the Web let you play snippets from songs on albums they are selling. It wouldn't make sense to show a list of albums (based on your selection criteria) and then start playing snippets from all the songs from all the albums in a row. Out-of-line sound gives the visitor to your site the option to listen to your audio file or to browse in silence.

Adding out-of-line sound to your page is as simple as using the `A` element and pointing to the sound file with the `href` attribute.

```
<A href="media/sound/elvis-impersonation.wav" alt="me singing
Blue Christmas with the stadium filter">Listen to me singing
Blue Christmas at the Meadowlands</A>
```

From Here

Cross-
Reference

Jump to Chapter 42 and learn about plug-ins and ActiveX controls.

Proceed to Chapter 41 to learn how to create video and add it to your page.

Summary

In this chapter you learned how digital audio works, what the variables are that make for a better or a worse recording, and why even the best recording may play back poorly on a visitor's computer. You learned about sound file formats and why streaming audio is a good idea. You also endured a stern lecture about why you shouldn't use unauthorized sound files, even if you think you can get away with it.

This chapter discussed Sound Edit 16, which you can use both to record and edit the sound, along with adding effects and filters. As with most multimedia objects, you must understand quality decisions, sample rate, and bit depth to deliver the best quality sound for the smallest file size to your site visitors. Finally, you learned about your two options for adding sound to your Web page, and why out-of-line is probably the better approach.



