Chapter 3

MIDIs and Other Synthesized Music Files

Did you know that your computer's sound board includes a complete synthesizer, just like rock stars play? True, it doesn't have a keyboard, but all the internal circuitry is there. And you don't have to be a rock star to play it. All you need do is start a MIDI file, which generates music on your synthesizer. Many multimedia games use MIDI files for background music while WAVs provide action sounds in the foreground. You have probably visited Web sites that play MIDIs on your synthesizer while you browse. You can also build your own collection of MIDIs on your hard drive. Tens of thousands of MIDIs are on the Internet for you to download and enjoy — rock, pop, jazz, country, classical, ethnic, almost any category you can think of. In this chapter, I'll show you how to play and enjoy MIDIs and related files — karaoke, anyone?

What you'll learn:

- How your synthesizer works
- What the General MIDI standard is and how it affects you
- Important extensions to the MIDI standard: GS, XG, and DLS
- What a MIDI file is
- How to play MIDI files in Windows and Mac OS

- How MOD files are different from MIDI files (and why some people like them better)
- How adding lyrics to MIDI files creates karaoke (KAR) files
- Where to find tons of MIDI, MOD, and KAR files on the Internet

Synth I Met You, Baby

As I said at the outset, this book is not about hardware. But because synthesizers and MIDIs go together like computers and programs, I find it impossible to explain MIDIs without talking a bit about synthesizers. In fact, I'm going to take you on a brief historical tour of synthesizers and electronic music. Stop that moaning . . . it won't be too painful.

The first synthesizers, developed in the 1960s by Robert Moog and others, were meant to help composers try out their compositions without going to the expense and bother of hiring an orchestra by simulating, or *synthesizing*, the needed instruments in their studios. Moog was an electronics engineer, not a musician, and from an electronics standpoint, the synthesizer was a relatively simple device. An ordinary oscillator generated an electronic waveform, which was modulated for characteristics such as pitch (the note or tone), volume, reverberation, and *timbre*—the complex of overtones that give each instrument its unique voice. Add a keyboard and some buttons and wheels so the user can play notes and adjust the aforementioned characteristics, and voilá—instant orchestra.

But the synthesizer was not so simple from a musician's standpoint. It was child's play to reproduce the sound of an electronic organ, but not too many people wanted to compose tunes for solo electronic organ. No one knew, or could discover, how to adjust a synthesizer to recreate the rich timbre of a cello, guitar, kettle drum, or just about any other instrument. We still don't know. So, in a sense, that early synthesizer was a failure. But before it got dumped on the trash heap, musicians—primarily rock musicians—began experimenting with the "synth" as an instrument in its own right. The world has never been the same since.

Nowadays, a synthesizer doesn't necessarily come wrapped in a keyboard. It might have strings, valves, drum pads, or some other human interface. Or—the reason we're all here—it could be built into a computer sound board and be driven solely by software. And we've come full circle, because with the exploding popularity of MIDIs on the Internet, people once again want synthesizers to realistically synthesize orchestral instruments.

Today's sound board synthesizers employ two common technologies. At the low end, *FM synthesis* comes closest to the original technology, although today's device is digital rather than analog. Unfortunately, because we still don't know how to analyze or recreate timbre, most FM synthesis sounds pretty artificial. Only the cheapest sound boards these days employ FM synthesis. More expensive boards use *wavetable sampling,* which bypasses the timbre problem by using the real thing. A *wavetable* is a collection of recorded sound samples of real pianos, guitars, bassoons, kettle drums, violins, harmonicas, flutes, ukuleles, piccolos, accordions . . . well, you get the idea. The synthesizer pulls samples as needed from the wavetable and modulates them for pitch, reverberation, and other characteristics. The result is a lot closer to the real thing.

Don't expect your computer to sound exactly like the New York Philharmonic or the Glenn Miller Band just yet. It's close, but even a tin ear can tell the difference. Many synthesizer instruments are near perfect—listen to a piano piece played on your synthesizer and you'll think it's a recording. But string and wind instruments still have an artificial quality to them. Of course, a lot depends on the quality of the sound samples in the wavetable. We'll come back to that point a bit later in this chapter.

The General MIDI Standard

Let's go back for a minute to those hip musicians of the 1970s exploring their new electronic synths. It seems only natural that someone would eventually patch two of them together and control both from one console. And once you figure out how to connect two, why not four . . . eight . . . twenty? You might call this the pre-MIDI era, when synthesizers began sprouting ports and patch cords, but no standard dictated how they communicated. Each manufacturer developed its own interface, so you couldn't connect a Roland to a Yamaha, nor a Casio to an E-mu. To bring some order to the chaos, the *Musical Instrument Digital Interface* (MIDI) standard was developed to define the way digital musical instruments (now called MIDI instruments when they adhere to the standard) could communicate. The original standard, which was released in the early 1980s, was upgraded in 1991 to the *General MIDI* (GM) standard, which remains the current standard. The GM standard calls for the following requirements:

- A five-pin connector for devices that have ports. Your computer's sound board probably doesn't have a MIDI port, by the way. If you want to connect it to an external MIDI instrument, you would use a cable designed to connect the MIDI port on the instrument to a serial or game port on your computer.
- A bank of 128 instruments. The original standard did not define the instruments but merely allowed for 128 of them. Instruments are selected by number, not name. So a trio using instrument numbers 3, 17, and 25 sounded like piano, sax, and drums on one synthesizer but ukulele, bassoon, and piccolo on another. General MIDI solved this problem by standardizing the *patch map* (think of it as a wavetable) shown in Table B-1 in Appendix B on this books CD-ROM.

Note

On early synthesizers, you had to connect modules together via patch cables to create the sound you wanted. The word *patch* came to represent a particular timbre, or instrument, and is still used in that sense today by MIDI composers and musicians. Hence the term *patch map*.

- Sixteen output channels. Each channel represents one instrument, so you can have a maximum of 16 instruments playing at once. You can change a channel's instrument at will. Channel 1 may start as grand piano but change to flute, then acoustic guitar, then back to piano, and wind up playing alto sax.
- Channel 10 is reserved for percussion instruments. The note number, which identifies the pitch in any other channel, indicates the percussion instrument on channel 10. Table B-2 in Appendix B shows the GM percussion map.
- At least 24 voices. Multiple voices permit some channels to play more than one note simultaneously, as in guitar and piano chords. Many modern synthesizers exceed this standard, providing 32 or even 64 voices.
- *MIDI messages* to control the sequence of MIDI events. A typical message might turn on note number 100 in channel 7. The sequence of events—note 16 on in channel 3, note 65 off in channel 1, patch instrument 105 to channel 14—is so crucial that a recorded MIDI song is called a *sequence,* the software you use to record it is called a *sequencer,* and we say the song was *sequenced by,* rather than performed by, Jasper Doohingle. (Jasper might not be a traditional instrumental performer but someone who enters events into the sequencer via computer keyboard and mouse. Just think—you, too, could become a MIDI musician if you have the patience to enter thousands of events to create a single song.)

How GM became GS and XG (with a dash of DLS)

The GM standard does not go far enough for some people. They want more channels, voices, controls, effects, and instruments. To meet their demands, both Roland and Yamaha have devised their own extensions to GM — Roland GS and Yamaha XG. Both extensions provide for many additional instruments plus other extra features. Each company has recorded extensive banks of instrument samples for use with their standards. If you want to start an argument, ask a room full of MIDI enthusiasts which extension is better — GS or XG. In general, Roland's GS seems to be winning out. Microsoft has licensed the Roland instrument samples for inclusion with the GS-compliant software synthesizer that comes with Windows 98. Apple also has licensed the Roland samples for its GS-compliant software synthesizer included with QuickTime 3.0. I have the feeling that XG will eventually join the ranks of electronics also-rans, along with Betamax VCRs.

You'll often see the notation GM, XG, or GS next to a MIDI filename on an Internet site, indicating the standard for which the sequence was created. You don't necessarily need to avoid MIDIs that use a different standard than your synthesizer. A GS or XG synthesizer can play a GM sequence with no problems simply by using the GM instrument bank. On the other hand, a sequence that takes advantage of the GS or XG extras will still sound pretty good on your plain vanilla GM synthesizer. The additional GS and XG sound banks have been carefully designed to complement the GM bank so that a GS sequence will sound okay, but not enhanced, on a GM synthesizer. Other additional features, such as extra effects, are simply ignored. The sequence will sound a lot better on a GS or XG synthesizer, though.

In 1996, the MIDI Manufacturer's Association (MMA) extended the GM standard to include DLS — downloadable sounds. Rather than storing wavetables permanently in the sound board's ROM, with DLS a wavetable is stored in RAM and is easily replaceable. Thus you could change your wavetable whenever you wish. Or, as you'll see when we talk about MODs later in this chapter, a sequence could come with its own wavetable, giving the composer/sequencer total control over the way the music sounds. Games developers in particular take advantage of downloadable sounds to add their game sounds to your basic set. Downloadable sounds have long been used on Creative Labs boards, which incorporate E-mu synthesizers along with E-mu's SoundFonts downloadable fonts.

Are you now wondering whether your sound board's synthesizer is GM, GS, or XG compatible? What sound banks you're currently using and how to change them? Whether you can use SoundFonts or other types of downloadable fonts? And if so, how you download and install them? Here are some suggestions on how to find out.

- Read its manual.
- Read the box it came in.
- Locate its folder on your hard drive and read all the documentation files (.txt, .doc, .hlp, and so on).
- Check the manufacturer's Web site. (Use a search tool to find it, or just try **http://www.*companyname*.com**, as in `http://www.creative.com`.)
- Contact its manufacturer. (You can find contact information on the company's Web site.)
- Look for newsgroups about MIDI, synthesizers, electronic music, sound boards, and so on.
- If it was installed at the factory when you bought your computer,
  - Look for a folder on your hard drive containing information about your computer; read all the documentation files there.
  - Contact your computer's manufacturer.
  - Contact the dealer where you bought the computer.

Mighty Mini MIDIs

It took us a while to get here, but we have finally arrived at the main topic of this chapter, MIDIs—more correctly called MIDI files. They are even more correctly called Standard MIDI Files, or SMF, but you'll see that phrase just about as often as you'll see the president referred to as William Clinton or WJC.

Note

Most MIDI files today have the extension .mid. You will find some with the extension .rmi, from Microsoft's RMID format.

It's important to understand that a MIDI file is not a sampled recording, nor any other kind of recording. It can't capture your acceptance speech, Elmer Fudd singing Wagner's "I killed the wabbit," or the sound of an egg frying on the sidewalk. That's the job of the sampled file formats we talked about in Chapter 2. MIDI files store only MIDI messages, which a MIDI player can interpret on a MIDI instrument such as your sound board's synthesizer.

The big advantage of MIDI files over sampled files—I should say *humongous* advantage—is their tiny size. A 15K MIDI might produce more than three minutes of music. By contrast, a 15K WAV using low-quality settings (8-bit, 11 Kbps, one channel) lasts less than two seconds. Is it any wonder that people prefer MIDIs to sampled sounds on the Internet?

MIDI files come in three flavors:

- *Format 0* assembles all MIDI data on a single track, so even the simplest player can play it.
- *Format 1* provides multiple tracks for sequencers that can record and playback separate tracks, an important feature for people who create sequences. This is the most popular format.
- *Format 2* permits multiple tracks and multiple sequence patterns and is rarely used.

WAVS, MIDIs, & RealAudio • 7507-4/1 • ch03.2 • Chapter 3 • Proof 2 • MS • 9/30/98 • p.52

Most of the MIDI files you'll encounter on the Internet use format 1. A few use format 0. Your player should be able to handle both formats. You probably won't even be aware of the file format when you download and play a MIDI.

Interpreting a MIDI sequence requires a lot of processing power, and it has to be done quickly or the music doesn't come out right. A slower processor — even an early Pentium — would not have time to decompress the file as well. Hence MIDI files are not compressed. Fortunately, they're small files and compression is not important.

Playing back MIDI files

Once a MIDI file is stored on one of your local drives, you play it back just like any of the sound files discussed in Chapter 2. Windows Media Player, ActiveMovie, QuickTime, Jet-Audio, Mac's Sound Manager, SoundApp, and many other players all play MIDI files. Be sure to check whether a MIDI player is included with your sound board, too.

For Windows, I prefer the Jet-Audio MIDI player, shown in Figure 3-1, because it enables you to make albums out of your MIDIs. When you collect a group of MIDIs into an album, you can play them much like a CD, with repeat mode, random mode, and a programmed playlist. An album on the MIDI player can contain many types of sequenced files, not just MIDIs. For example, it can also contain MOD and karaoke files, which are explained later in this chapter. To make or play an album, see the "Digital Audio Player" section in Chapter 2.

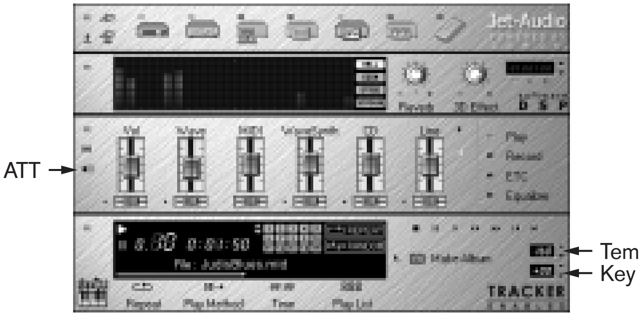

Figure 3-1 *Jet-Audio's MIDI player plays many formats of MIDI, MOD, and karaoke files.*

Want to sing along with the music? You may need the two controls at the right side of the player. The top control, called the Tempo control, changes the speed of the playback. Click the up arrow to make it faster or the down arrow to make it slower. The bottom control, called the Key control, transposes the music to other keys. Click the up arrow for a higher key or the down arrow for a lower one.

I included Mixer and Sound Effector in Figure 3-1 because they interact with MIDI Player. Mixer can be used to control the volume of the MIDI music in relationship to any other devices that are currently playing, such as WAVs from an application that you are using. You can use the attenuation (ATT) button, shown in the margin, to lower the volume while you talk to someone and then raise it again. Mixer's graphic equalizer has no effect on MIDIs. Sound Effector's spectrum analysis display responds to music played by MIDI Player. The sleep timer works with MIDI Player, but the DSP effects do not.

Other file types that contain MIDI sequences

The MIDI file format (or SMF) is not the only file format that contains MIDI sequences. MOD files, which have a wavetable built into the file, are gaining a near-cult following. And you can imagine how popular karaoke files have become.

MOD files

One way for a composer to make sure a MIDI sequence sounds exactly the same on everyone's system is to include the desired wavetable right in the file. The player uses the included sound samples instead of those provided by the sound board, and the music sounds right on everyone's computer. You can't include a wavetable in a MIDI file, but you can in a MOD file. In fact, that's the major difference between MIDI and MOD.

Note

MOD is short for *modules*.

While the included wavetable makes a MOD much larger than a MIDI, it is still significantly smaller than a WAV or other sampled recording. So why isn't MOD more popular? For one, because MODs are so much larger than MIDIs. For two, because MOD editors — which are called *trackers* not sequencers — tend to be hard to use, so many composers don't bother to learn them. And for three, because the several forms of MOD files aren't necessarily compatible with each other. A MOD player often can't play all the different types of MODs. You may run into the following MOD variations:

- .mod — Files tracked by ProTracker, FastTracker 1, StarTrekker, Noise Tracker, and others
- .669 — Files tracked by Composer 669
- .mtm — Files tracked by Multi Module Edit
- .s3m — Files tracked by ScreamTracker

- .stm — Files tracked by ScreamTracker 2.*xx*
- .it — Files tracked by Impulse Tracker
- .xm — Files tracked by various trackers

Tip

Notice the phrase *Tracker Enabled* on the Jet-Audio MIDI Player in Figure 3-1. That means it can play MODs as well as MIDIs.

Karaoke files

Add lyrics to a MIDI file and what do you get? Karaoke! KAR files are played on a karaoke player that shows the lyrics while playing the music. A good player uses some method to mark off the lyrics so you know where you are in the music. Some display a bouncing object, while others simply change the color of the words. As I write this, a few shareware karaoke players and editors are available, but I couldn't find one I liked well enough to include with the book. Internet karaoke is an up-and-coming field, so some excellent players/editors may be available by the time you read this. Check this book's Web site for recommendations:

`http://members.aol.com/jnfbooks`

In the meantime, you can play karaoke files with QuickTime, which is included on this book's CD-ROM. Both Windows and Macintosh versions of QuickTime 3 are on the CD-ROM. When you open a KAR file, QuickTime converts a KAR file to its own MOV format and then displays it as a movie with a sound track and a text track. This is not the ideal solution, but if you're a karaoke fan, it's better than no player at all. QuickTime marks off the lyrics by changing colors, as shown in Figure 3-2. Unfortunately, QuickTime doesn't enable you to change the speed or key of the music, so you have to sing along with the music as is.

Figure 3-2 *QuickTime times the karaoke lyrics to the music by changing the color of the text.*

Note

Jet-Audio's MIDI player plays KAR files and enables you to adjust the key and the speed of the music, but it does not display lyrics.

The fine print

(Please try to read this in a fast voice.) There is no charge for the basic version of QuickTime. The Windows version works with Windows 95, 98, and NT 4.0. It requires a 486-DX2/66 or higher, 16MB RAM, and a Sound Blaster-compatible card. Also, the following are recommended for better performance: DirectX 3.0, DirectDraw, and DirectSound. (These can all be downloaded from Microsoft's Web site.) The Macintosh version works with System 7.1 or higher. It requires 16MB RAM for PowerPC or 8MB RAM for 68K machines; 68K-based computers must also support Color QuickDraw.

How to play karaoke files with QuickTime:

1. With QuickTime active, choose File ➪ Open to open a browse box.
2. Locate and select the KAR file and then choose Convert. QuickTime displays a dialog box so you can save the converted file.
3. Change the name of the new movie file, if desired, select a folder for it, and choose Save. QuickTime opens the new movie file but does not play it.

Now you can control the song with these buttons:

 Play

 Pause

 Jump forward to next verse

 Jump backward to preceding verse

 Volume control

You can also move around in the song by dragging the slider. If a group of people are trying to read the lyrics, click Movie ➪ Double Size from the menu bar to double the size of the display. Choose Movie ➪ Normal Size to return to its regular size.

When you convert a KAR file to a movie file, QuickTime enables you to choose such things as font, size, and color for the text. If you don't want to use the default settings, choose the Options button in the dialog box. This opens a MIDI import options dialog box, where you choose the Text Options button to open the Text Import Settings dialog box. Figure 3-3 shows the Windows version of this dialog box, but the Macintosh version is nearly identical.

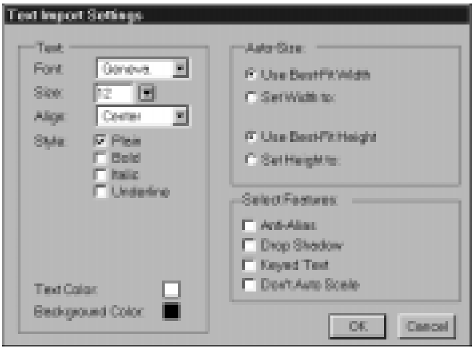

Figure 3-3 *When you convert a karaoke file with QuickTime, you can select text properties in this dialog box.*

In the Text group, you choose the font, point size, style, alignment, and the foreground and background colors. To change a color, click the colored box to open a color panel where you can select the desired color. The Auto-Size group affects the size of the text display window, not the size of the font. If you select the Auto-Size options, QuickTime selects the best size to display the text of the song. If you enter specific width and/or height, in pixels, text lines might have to be wrapped to fit the window.

By default, if the designated font size is so large that lines must be wrapped in the window, QuickTime adapts the font size to fit the window. Also, if you enlarge or shrink the window, QuickTime automatically adapts the font size to fit. But if you enable Don't Autoscale, then QuickTime is forced to use the designated font size, no matter what size the window is.

The Keyed Text option displays the lyrics on a white background. If the text is also white, it's invisible until its color changes. So instead of the lyrics appearing all at once and changing color, they pop onto the screen in time with the music.

You can download karaoke files from several sites on the Internet. To get you started, I have included a karaoke Webring in Appendix A on the CD-ROM. If you have a karaoke editor, you can create your own by typing in the lyrics for MIDI files that you already have. Plain-old vanilla QuickTime 3 plays karaoke movie files but cannot edit the text. Upgrading to QuickTime 3 Pro ($29.95) gives you the ability to create and edit text tracks in any MOV file, including karaoke movies.

What's Next?

We've touched a bit on the Windows sound features. Chapter 4 explores in greater detail what you can do with sound in Windows 95 and 98, including some hints on how to configure your MIDI synthesizer.