

4 } sfz and sfz+

First, let's deal with a bit of confusion: sfz and sfz. One is an instrument, and one is a file format (which the sfz player recognizes). To prevent confusion as much as possible, I will refer to the sfz player and the sfz format (from which we can create sfz definition files).

The sfz player was originally developed as a component for what was to be rgc:audio's next synthesizer. We now know the synthesizer to be Dimension (which is essentially four sfz players on steroids, strapped together). The sfz engine in its current incarnation as the Expression Engine has also been used as the basis for Cakewalk's most recent synthesizer: Rapture.

What started out as a component has found a life for itself as a high-quality SoundFont and wav/aiff/ogg file player. There are now two versions of the component:

- **sfz**. The sfz player—this is a freeware instrument.
- **sfz+**. This a commercial instrument with a more sophisticated interface.

These are both discussed in this chapter.

A Word or Two about sfz Format Files and SoundFont® Format Files

The sfz format and the SoundFont format are ways of arranging files so they can be replayed. Within the file format, a group of samples can be structured and arranged so they can be loaded into a synthesizer and deployed for a performance. The SoundFont format was developed by E-MU, and the sfz format was developed by Cakewalk (rgc:audio).

There are differences between the two formats, most notably:

- SoundFont format files are monolithic—the waves and the mapping information are stored together in one file.
- sfz format files are non-monolithic—each wave is separately stored, and the mapping file (which references the wave files) is a separate document.

CHAPTER 4 } sfz and sfz+

SoundFont format files can include a range of programs in the monolithic file, whereas sfz format files do not.

The sfz format is discussed in greater detail in Chapter 17, “The sfz Format.”

What Is the sfz Player?

The sfz player (see Figure 4.1) is a high-quality multi-timbral software synthesizer. The player is bundled with SONAR version 6 (as it was with version 5) and SONAR Home Studio version 4 (but not SONAR Home Studio version 4 XL, which has Dyad bundled with it). You also can download it for free from rgcaudio (www.rgcaudio.com).

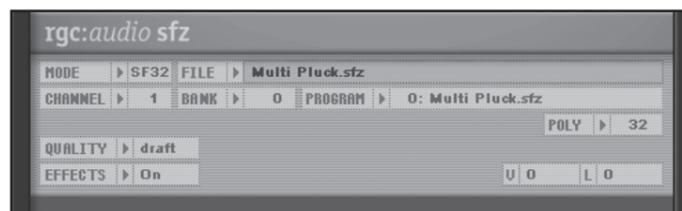


Figure 4.1 The sfz player: One of the highest-quality sample playback tools on the market, available for free from rgcaudio and bundled with SONAR version 5 and SONAR Home Studio version 4.

The sfz player is a high-quality sample playback engine featuring:

- Nine quality settings to allow control over CPU usage
- Two-hundred fifty-six voices of polyphony with unlimited layers per voice
- Multi-timbral operation
- Different loading modes to control memory usage and CPU overhead
- A multi-mode filter (6dB/octave low-pass and high-pass filters, 12dB/octave low-pass, high-pass, band-pass, and band-reject filters)
- Three LFOs
- Three envelopes
- Stereo chorus and stereo reverb effects
- ASIO microhost

The sfz player can load and play:

- sfz format (.sfz) files
- SoundFont (.sf2) format Banks
- Wave (.wav) files
- Ogg-vorbis (.ogg) compressed files

Single wave and ogg-vorbis files can be loaded directly and will then be playable over the whole keyboard. If you want to map a wave or an ogg-vorbis file to a specified region of the

■ Getting to Know the sfz Player

keyboard or set its pitch center, you would need to do this with an sfz file. Chapter 17 explains the mechanics.

Why Use the sfz Player?

I haven't tried all of the sample playback devices on the market; however, sfz is certainly one of the highest-quality tools that I have tried. It is even more impressive because it is free and uses the sfz format.

As a tool it performs only one task—it replays samples. It is designed to perform that single task, and it does it as well as you could ask.

■ Using the sfz Player as a Full-Featured Synthesizer

It isn't obvious from the interface, but the sfz player is actually a very flexible, highly specified synthesizer.

Here are some of the features you might not realize the sfz player can offer as a synthesizer:

- Unlimited choice of waveforms
- Unlimited number of oscillators
- Multi-mode oscillators (in which you can configure the number of oscillators as well as the tuning and the panning of each oscillator)
- Sub-oscillators
- A range of filters
- Envelope and LFO modulation sources
- Control over most parameters from MIDI hardware
- Two programmable FX units

All you need to go with your sfz player are a few waves and some sfz definition files.

If you're looking for a source of waves, then Audacity (<http://audacity.sourceforge.net>) might be a good choice for you. Audacity is a free wave editor with which you can create waveforms and then apply a range of transformations to those waves.

The combination of the sfz player, Audacity, and a bit of sfz definition file ingenuity on your part can create the ultimate freeware synthesizer.

Check out Chapter 17 for details about how to create, edit, and deploy sfz definition files within the sfz player.

Getting to Know the sfz Player

To use the sfz player, you obviously have to load a file first.

Loading Files

To load a file in the sfz player, you can perform one of the following actions:

CHAPTER 4 } sfz and sfz+

- Click in the File Name box and navigate to the appropriate file (which can be in .sfz, .sf2, .wav, or .ogg format).
- Click on the arrowhead to the right of the File label and select one of the recently loaded files.
- Drag and drop a file onto the interface.

If you have loaded a SoundFont Bank, you can then select the appropriate Bank and program by clicking on the arrowhead next to the Bank and program label or by clicking in the Bank/program name. As with other functions in the sfz player, when you click on the name, left clicks move one step up the list and right clicks move one step down the list.

You can only load one SoundFont Bank at a time (which may contain several Banks and many programs), and you can only load one sfz format file (in other words, one sfz instrument, one wave file, or one ogg file).

You can use the sfz multi-timbrally with SoundFont Banks by loading a different program for each MIDI channel. You can select the different MIDI channel with the arrowhead to the right of the Channel label (see Figure 4.2) or by clicking (left or right) on the selected channel number. As you change the channel, you will see the selected program name change.

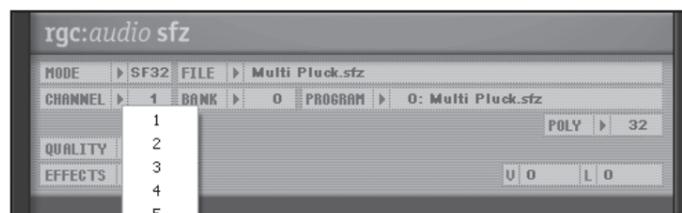


Figure 4.2 Selecting a different channel in sfz from the Channel drop-down menu.

Polyphony

The polyphony of the current program is controlled by the Poly selector (see Figure 4.3). This controls the number of voices the current channel will play. A voice can contain several layers (in essence, several samples mapped to the same key). The number of voices and layers are indicated by the V and L meters under the Poly selector.

Effects

You can switch the internal effects units (a Chorus and a Reverb unit) with the Effects selector. This is a global control that affects all channels (in other words, the FX units are either on or off for all parts); however, it is not necessary to send each channel to the effects unit.

When the effects are engaged, they can be controlled from your host.

■ Getting to Know the sfz Player

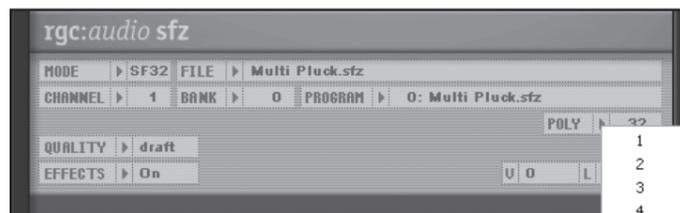


Figure 4.3 The Polyphony selector drop-down menu in sfz.

Preview Ribbon

As with all rgc:audio synthesizers, there is a preview ribbon across the bottom of the sfz player. As well as controlling pitch, this also generates velocity information—the lower you click on the ribbon, the higher the velocity.

Quality Settings

The Quality Settings selector is a feature that is found in both the sfz player and sfz+.

The Quality selector allows you to select the playback quality from standard (draft) to extremely high (72). To change the Quality settings you can either:

- Click on the current value—left-click to increase the quality and right-click to decrease the quality.
- Click on the arrowhead to the left of the current value. A drop-down menu will appear, from which you can select the appropriate Quality setting (see Figure 4.4).

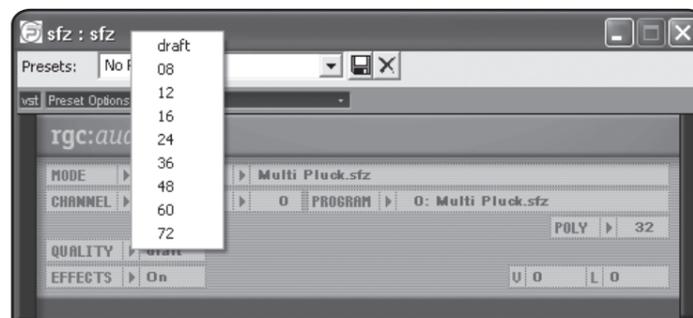


Figure 4.4 The Quality selector drop-down menu in sfz.

The sampling engine in the sfz player and sfz+ takes a sample of one pitch and changes its pitch based on the incoming MIDI data. This, in essence, is what all samplers do. The amount by which each sample has to be re-pitched depends on how many samples are included in

CHAPTER 4 } sfz and sfz+

the sample set. If every note in an instrument has been sampled, then no re-pitching is necessary. However if, as is likely, not every note has been sampled, then some samples will need to be re-pitched, perhaps radically.

When samples are re-pitched, digital artifacts are introduced. These are unwanted sounds that manifest themselves as distortion and as degradation to the audio signal. Different re-pitching algorithms lead to differing amounts of distortion. As a general rule, the more sophisticated the algorithm, the less the distortion. As another general rule, the more sophisticated the algorithm, the more CPU power it will use.

These two rules are born out in the sfz player and sfz+, so the Quality selector allows you to choose a balance between sonic fidelity and CPU hit.

Although draft quality is the lowest quality on offer, it still represents a good level of playback quality. To determine the right setting for you, you will need to use your ears and keep an eye on your CPU meter.

The Quality setting is a global setting that applies to all channels.

Mode Settings

The Mode setting (see Figure 4.5) can be changed by clicking on the current Mode setting or accessing it from the drop-down menu by clicking on the arrowhead next to the current Mode setting.

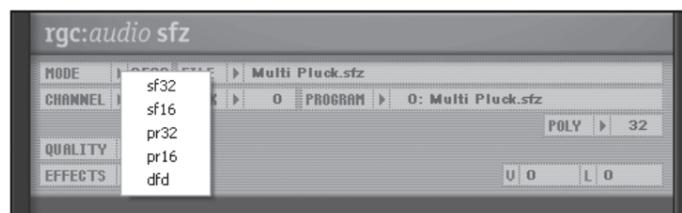


Figure 4.5 The Mode selector drop-down menu in sfz.

Conceptually, the Mode setting is not necessarily straightforward: It does not affect the quality of the output. However, it does affect the CPU usage and the amount of memory used when you load SoundFont Banks.

SF32 Mode

In SF32 mode, the sfz player will load the whole SoundFont Bank file into memory (that is, all the waveforms for all of the programs in the SoundFont Bank will be loaded). The waves will be converted to 32-bit during loading.

This option optimizes performance and flexibility so programs can be changed while voices are playing without interruption. However, this option uses a lot of memory—approximately two-and-a-half times the size of the SoundFont Bank.

■ What Is sfz+?

If there is insufficient memory, then the sfz player will switch to SF16 mode (see the following section), which will reduce the amount of memory required by about 50%. However, this also increases the CPU overhead.

SF16 Mode

In SF16 mode, the whole SoundFont Bank file is loaded, but all samples are loaded and played as 16-bit samples. This mode takes half the memory space required by SF32 mode, but increases the CPU overhead because the 16-bit files are converted to 32-bit as the samples are played.

PR32 Mode

In PR32 mode, only the samples required by the current program (preset) are loaded and converted to 32-bit. This mode only requires memory for the samples that will be used, hence the RAM requirement may be reduced considerably (depending on the size of the SoundFont Bank).

When running in PR32 mode, program changes will lead to all voices being stopped immediately before the next set of samples is loaded. This might lead to unwanted artifacts during playback.

PR16 Mode

PR16 mode is similar to PR32 mode except that the samples are loaded as 16-bit, so they take up less memory.

DFD Mode

DFD (or *Direct from Disk*) mode is the disk-streaming mode in the sfz player. Samples are streamed directly from the computer's hard disk without being loaded into memory. This mode takes the least amount of memory and allows sample sets that are larger than the available memory in the host computer. It also has the advantage of having a very short loading time.

Creating Programs in the sfz Player

To create programs that are playable in the sfz player, you need to manipulate sfz definition files. Check out Chapter 17 for details of how to create, edit, and deploy sfz definition files within the sfz player. You could alternatively create a SoundFont Bank for which you would need a specialized editor, such as Alive from SoundFaction (www.soundfaction.com).

What Is sfz+?

So sfz+ is just the sfz player with a pretty dress and a nice price tag, right? Wrong. There are several key differences between the sfz player and sfz+:

CHAPTER 4 } sfz and sfz+

- sfz+ will not load sfz format files, nor will it load ogg format files.
- sfz+ has a more detailed interface, allowing direct editing of many features. As you will see in Chapter 17, these editing options can be accessed in the sfz player by editing an sfz format file.
- sfz+ has multiple outputs.

Why Use sfz+?

sfz+ (see Figure 4.6) is perhaps the highest-quality SoundFont Bank player—the coupling of high-quality audio with a highly intuitive interface makes it an obvious choice for SoundFont playback. Although you can tweak SoundFont Banks in the sfz player, it is so much easier in sfz+.

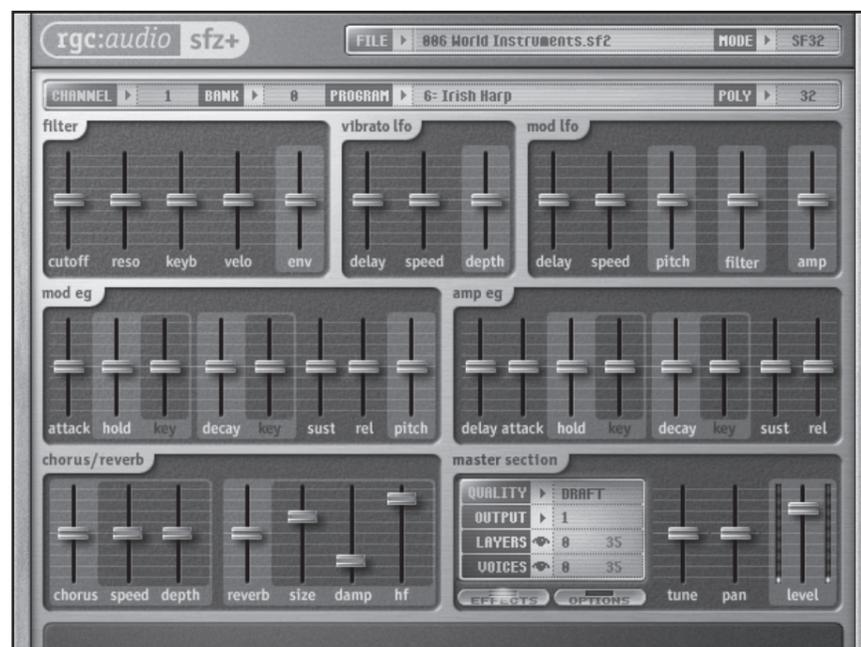


Figure 4.6 sfz+: more than an sfz player with a cool interface.

Getting to Know sfz+

Initially, sfz+ works in a very similar way to the sfz player. The following functions perform in exactly the same manner:

■ Getting to Know sfz+

- File loading.
- The Mode selector.
- The Channel selector (including the options to use sfz+ as a multi-timbral sound source).
- The Bank selector.
- The Poly selector.
- The Quality selector. As in the sfz player, the Quality selector is a global control.
- The Effects on/off function (although in sfz+ it is a button with a light that glows orange when the effects are activated). As in the sfz player, this control affects all channels.
- The Layers and Voices indicators (although these are labeled in full), which display the layers and voices used globally.
- The preview ribbon.

However, sfz+ also borrows many usability features from its big brother, Z3TA+. For instance, after you click and drag a slider to change its value, it will keep focus and can be adjusted by your mouse wheel or the PC keyboard. From your PC keyboard:

- The left and right arrows make very fine adjustments.
- The up and down arrows make normal adjustments.
- The Page Up and Page Down keys make large adjustments.

In addition, if you Ctrl-click on a slider (or double-click it), you will reset it to its default value.

Filter Section

sfz+ has a low-pass filter. The Filter section in sfz+ will adjust the filter parameters as instructed by any SoundFont Bank. You can further tweak the filter with these controls (see Figure 4.7):

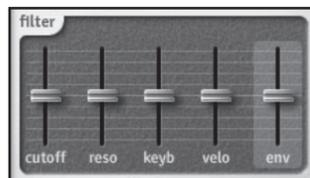


Figure 4.7 The Filter section in sfz+.

- **Cutoff slider.** This sets the filter's cut-off frequency.
- **Reso slider.** This sets the filter's resonance.
- **Keyb (key scaling) slider.** This determines the extent to which the filter's cut-off varies according to the pitch of the note played. This can have a positive or negative value, so if you set the slider to a negative value (that is, below the midpoint), higher notes will be less bright.

CHAPTER 4 } sfz and sfz+

- **Velo (velocity) slider.** This controls the extent to which the filter's cut-off frequency varies according to the incoming MIDI velocity. This can have a positive or negative effect, and with negative settings (that is, below the midpoint), an increase in velocity will lead to a cut in the filter's cut-off.
- **Env (envelope depth) slider.** This controls the extent to which the filter's cut-off frequency is controlled by the Mod EG (modulation envelope). Again, this can be both positive and negative, and with negative values the Mod EG can be set to close the filter.

Vibrato LFO

The Vibrato LFO (see Figure 4.8) is a dedicated LFO that modulates the pitch of the current channel. It has three controls:

- **Delay.** This is the time after a note is struck until the LFO will begin to affect the pitch.
- **Speed.** This is the speed of the Vibrato LFO.
- **Depth.** This is the maximum pitch variation introduced by the Vibrato LFO. This control is bipolar: When positive values are selected (in other words, the slider is above the midpoint), then the LFO will modulate the pitch normally. When negative values are selected (the slider is below the midpoint), the polarity of the LFO will be reversed.



Figure 4.8 The Vibrato LFO in sfz+.

Mod LFO

Working in addition to the Vibrato LFO, the Mod LFO (see Figure 4.9) is an LFO that can modulate three destinations: pitch, filter cut-off, and volume. The controls in the Mod LFO are:

- **Delay.** This is the time after a note is struck until the LFO will begin to affect the modulation destination.
- **Speed.** This is the speed of the LFO.
- **Pitch.** This controls the extent to which the Mod LFO modulates pitch.
- **Filter.** This controls the extent to which the Mod LFO modulates the filter cut-off.
- **Amp.** This controls the extent to which the Mod LFO modulates the volume of the program.

■ Getting to Know sfz+

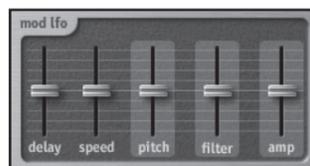


Figure 4.9 The Mod LFO—an assignable LFO—in sfz+.

The Mod LFO only works when SoundFont Banks are loaded. It has no effect if you have loaded a wave file.

sfz+ Envelopes

There are two envelopes in sfz+:

- **Mod EG.** This is the modulation envelope generator (see Figure 4.10), an assignable envelope that can control:
 - **Filter cut-off.** The extent to which the Mod EG controls the filter cut-off is set by the Envelope slider in the Filter section.
 - **Pitch.** The extent to which the Mod EG controls pitch is set by the Pitch slider in the Mod EG.



Figure 4.10 The Mod EG (modulation envelope generator)—an assignable envelope in sfz+.

- **Amp EG.** This is the amplifier envelope generator (see Figure 4.11), which controls the volume of the current channel.



Figure 4.11 The Amp EG (amplifier envelope) in sfz+.

CHAPTER 4 } sfz and sfz+

Each envelope has virtually the same controls (see Figure 4.12):

- **Delay (delay time).** This is the time after the key is struck before the envelope starts. (This is only available in the Amp EG.)
- **Attack (attack time).** This is the time after the key is struck (in the Mod EG) or the delay time ends (Amp EG) that it takes for the envelope level to go from zero to the maximum level.
- **Hold (hold time).** This is the time that the envelope remains at the maximum level at the end of the attack time.
- **Key.** This is the variation (positive or negative) in the hold time depending on the pitch of the incoming note.
- **Decay (decay time).** This is the time that it takes for the note to decay from its maximum level to the sustain level at the end of the hold time.
- **Key.** This is the variation (positive or negative) in the decay time, depending on the pitch of the note.
- **Sustain (sustain level).** This is the level at which the note will sustain while the key is held.
- **Release (release time).** This is the time it takes for the note to decay from the sustain level to zero after a key is released.

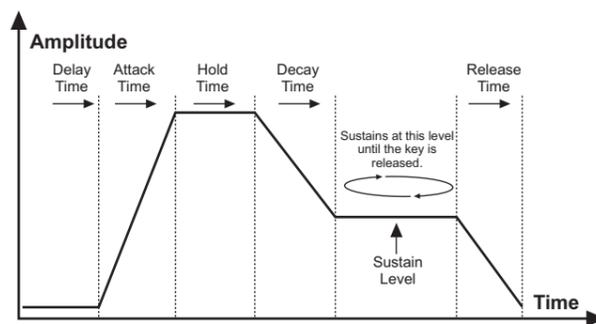


Figure 4.12 sfz+ features two envelopes: the Amp EG is a DAHDSR (delay time, attack time, hold time, decay time, sustain level, release time) envelope. For both envelopes, the hold time and the decay time can be modulated by the pitch of the played key (key tracking).

Effects Section

There are two effects units: a Chorus and a Reverb (see Figure 4.13). These are both global effects (in other words, there is only one of each unit). The main settings for each FX unit are global—you can control the amount of signal from each channel that is then sent to the FX unit. You can readily see when you are editing a global control; as you click on the slider, its knob glows orange rather than blue.

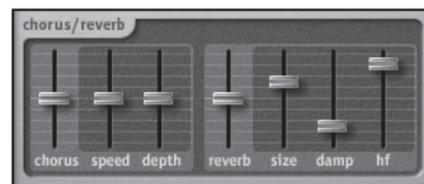


Figure 4.13 The Chorus/Reverb FX unit in sfz+.

Chorus

The Chorus unit has three controls:

- **Chorus slider.** This controls the amount of the channel output that is sent to the Chorus unit. It is set on a per-channel basis.
- **Speed slider.** This controls the speed of the LFO modulating the delayed signal. It is a global control.
- **Depth slider.** This controls the maximum amount that the Chorus LFO can modulate the pitch of the delayed signal. This is also a global control.

Reverb

The Reverb section has four controls:

- **Reverb slider.** This controls the amount of the channel output that is sent to the Reverb unit. It is set on a per-channel basis.
- **Size slider.** This sets the size of the simulated space, thereby controlling the decay time of the reverb effect. This is a global control.
- **Damp slider.** This sets the diffusion of the higher frequencies. It is a global control.
- **HF (high-frequency) slider.** This cuts the high frequencies in the reverb. It is also a global control.

These last two controls will need to be balanced to create a natural reverb effect. You will find that when the Damp slider is set to the maximum, the effect of the HF slider becomes less noticeable.

Master Section

The Master section (see Figure 4.14) contains several controls:

- **Tune slider.** This adjusts the tuning of the current channel program over the range of ± 100 cents (one semitone).
- **Pan slider.** This controls the placement of the program in the stereo panorama.
- **Level slider.** This sets the level for the current program.
- **Output control.** This allows you to choose the output for the current channel. You can choose any one of eight outputs; the ninth output is for the effects unit (if the effects signal is not sent to the main output). If you use multiple outputs, remember to ensure that there is a corresponding audio channel set up in your host for each audio signal.

CHAPTER 4 } sfz and sfz+



Figure 4.14 The Master section in sfz+.

The Options menu (see Figure 4.15) is selected by clicking on the Options button. This menu gives the following options:

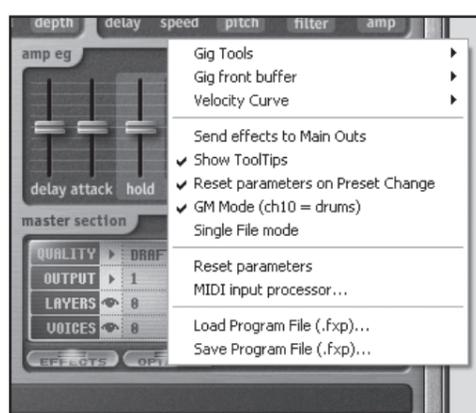


Figure 4.15 The Options menu in sfz+. You should note that the Gig Tools and Gig Front Buffer options are not available in the current version of sfz+.

- **Velocity Curve.** This gives you a range of velocity response curves so you can pick one that works best with your playing style. In addition, there are three fixed velocity responses.
- **Send Effects to Main Outs.** When this option is not enabled, the outputs of the effects units are routed to a dedicated output (output 9).
- **Show ToolTips.** This enables or disables the tool tips, which are displayed when you hold your mouse over a slider. Enabling this option takes a small amount of CPU resources.
- **Reset Parameters on Preset Change.** This will set all sliders to their default setting when a new program is selected from the SoundFont Bank.
- **GM Mode (ch10 = drums).** This will automatically select Bank 128 (if present) for channel 10 when a SoundFont Bank is loaded.

■ Getting to Know sfz+

- **Single File Mode.** This mode will load the same file into all 16 channels (when a new file is loaded).
- **Reset Parameters.** This resets all parameters to their default values for the current channel.
- **MIDI Input Processor.** This allows you to layer sounds and assign sounds to certain key and velocity ranges. In the MIDI Input Processor dialog, for each channel you can set:
 - **Ch.** This overrides the MIDI channel that will control that channel. So, for instance, you could set MIDI channel 2 to control channel 1 by clicking on the 1 in the Ch column, deleting it, and inserting 2.
 - **Trans.** This will transpose the channel in semitones. (Negative values can be inserted.)
 - **Lo Key and Hi Key.** These set the key range over which a part will sound. Notes outside of this range will not sound on this channel.
 - **Vel Off (velocity offset).** This applies an offset (whether positive or negative) of a fixed amount to the velocity of the incoming MIDI data.
 - **Lo Vel and Hi Vel.** These set the limit of the velocity range over which the channel will sound. Velocities outside of this range will not sound on this channel.
- **Save Program File (.fxp) and Load Program File (.fxp).** These will save and load your settings. These options will not save the underlying SoundFont Banks or wave files; instead, they will reference them, so if you transfer .fxp files between different machines, you will need to ensure you transfer your sound sets too. You can also load an .fxp program by dropping the file on the main sfz+ interface.

sfz+ allows you to save a file called sfz+default.fxp, which will automatically load a default set of sounds and settings when sfz+ is launched.

■ Finding SoundFont Banks for the sfz Player and sfz+

There are many free SoundFont Banks available on the Internet. Google the phrase “free soundfonts” and see what comes up.

As an alternative source for SoundFont Banks, take a look through your existing library of samples. If you’ve got anything interesting in there, it might be worth converting those samples into SoundFont format with a sample conversion utility.



1-59863-314-7_CH04_86_09/19/06

