

the **Sonata** user manual

by SoHa Sound Design

The Synthesizer

Welcome to Sonata.

Sonata is a VST soft-synth for strong single-shot leads and pads. Adept at stringy things and lots of simple subtractive sounds, but specializing in sparkly and glassy mallet sounds, deep basses, and warm and rich mid-range chromatic inventions, Sonata is a one-trick pony with distinctive character.

Sonata features two separate but similarly controlled generators, one for the sound's immediate attack and another for the fundamental tone. Both consist of rich harmonic spectra which are filtered to provide desired overtones. Each possesses an exponential decay envelope for amplitude and another for the filter, as well as a normal ADSR envelope for one or both. The output from these is then fed into a little magic black box (a complex setup of delay devices set relatively for each note) that shapes the sound, emphasizing desired frequencies and providing richer post-filter overtones, similar to the way an instrument's body will influence its sound. There are three global effects: a cross-delay, distortion, and a dense near-field reverb meant to accent near-field reflections.

Sonata is meant to produce sounds with consistent timbres in both high and low registers for better two-handed performance and ranged melodies. As a result, many of its interior units exist on a per-note basis and each note can get a little cpu-hungry. Since Sonata's polyphony is limited only by your computer's processor power, a slim 2.5GHz PC should handle about 16 note polyphony without much else going on. More meager 1.6GHz laptops running other VSTi's moderately should expect only about 6 to 8. Polyphony is limited in a manner meant to sound most natural, but is of course never perfect. We therefore highly recommend using a host with the freeze function, especially on less powerful machines.

~



Table of Contents

The Synthesizer.....	1
The Tone Synthesizer.....	3
The Attack Synthesizer.....	4
The Effects.....	6
The Controls.....	6
The Presets.....	6
Trouble Shooting.....	7
The Company.....	8

~

The Tone Synthesizer

Generally, this is the basic, underlying tone in the patch. While the Attack Synth is meant to decay away rather quickly, this one typically endures. When you play a note, the generators start up, mixing between square and sawtooth waveforms. There are three of these generators, and they may be set to slightly different frequencies to cause a fuller spectrum and phasing. The output is then fed into an exponential decay envelope on the amplitude, causing the sound to diminish in time. Lastly, the output is fed through a low-pass filter whose frequency diminishes in time just like the amplitude.

The Waveform:

The waveform is a combination of sawtooth and square waves. When the slider is at the top, it is fully saw; at the bottom it is full square; and in the middle, the signal is mixed. Why only two waveforms? Well, this synth is mostly subtractive, and the sawtooth and square waveforms contain a lot of a energy which can be filtered out for the desired effect. Other wave forms typically retain the harmonics of either a sawtooth or square, but the harmonics have different phases and amplitudes (it is recognized that some synths use radically different waveforms). The hope here is that Sonata will change the amplitude and phase of the harmonics enough with the filter envelopes and especially the delay chambers that further waveforms will not result in significantly different sounds. So, the provided two waveforms are what we found to be the best pair for Sonata.

The Unison Parameters:

There are three generators for the overall Tone generator, and they can be set to slightly different frequencies. The first generator will be exactly at the frequency of the note. The second will be higher as set by Offset. The third will be lower as set by Offset. Amp controls how loud the second and third generators are. Moderate values of Offset and high values for Amp work well for stringy sounds. Very small values of Offset and low Amp values work well for sounds like a piano (where there actually are three strings, each at a slightly different frequency, just like here). Use all high values for “trance leads” and other thick, phat, or otherwise widely-scoped sounds.

The Attack Shape:

This is an exponentially decaying envelope meant to emulate the way many sounds get quieter after the initial onset. Set the point at which you want the volume to start with Init. Set the point to which it should decay to “Final”. Set how quickly you want it to decay with the Decay parameter, where higher means “more decay”. Keep in mind that more decay means shorter sound. If you strike a marimba, there’s a loud initial sound. Then it decays with time down to zero. Use these parameters to capture that

decay.

The Filter Shape:

This is almost identical to the attack decay curve, except with a -24db low-pass filter. Sounds often lose high-spectrum frequencies before low, so set the beginning, high frequency with Init, the low with Final, and how much decay there is with Decay. However, it's not entirely obvious how high the frequencies are. A setting of 0 for "Init", for instance, refers to the fundamental of the note being played. This means you basically only get a sine wave at the fundamental. A setting of 1 means the cutoff will be 24 times the fundamental. That's very high. Everything in between is a linear interpolation. Things are a bit different for "Final". It's assumed you don't want your sound's emphasized cutoff frequency to be 24 harmonics higher than the fundamental all the time, so the Final frequency only goes from 1 to 8.

There's also a resonance parameter that controls how "sharply" the high frequencies are cut off, and how emphasized the cutoff frequency is. Low values are generally good for natural sounds and high are good for synthetic sounds, but it really just depends on how much you want the cutoff to be emphasized. A couple of tips: Use low resonance for cutoffs meant to gracefully move from the highest to the final, such as for a piano. Using high resonance with a slow decay and do some cool stuff, like emphasizing every harmonic encounter on the way down from high to low. Or, just set the final to the first or second harmonic (the octave of fifth above it), and high resonance will just emphasize these frequencies. Yummy on bass and good for warm sounds. Remember, however, to set this a little cooler than what you want, as the Attack will take the high-energy part.

The ADSR Envelope:

This is pretty standard, but I'm going to describe it anyway. Sounds always start with zero volume. Attack sets how long the sound takes to reach full volume, so 0 means "almost instantly", whereas 1 means "1 second". Decay sets how long it takes to come back down to the Sustain level from full volume. Again, 0 to 1 seconds. Sustain sets the volume from that the sound will hold after decay and before release. This stage lasts as long as you hold the note. Release sets how long it takes to diminish to zero when the note is released. If you release before Sustain is achieved, it will take just as long to diminish from where the amplitude left off as it would to diminish from Sustain.

The Attack Synthesizer

The Attack synthesizer is meant to provide the transient parts of the sound, so its parameters are tweaked a bit differently to provide more varied effect. Still, it's functionality is essentially the same as the Tone, so it can also be used as a second set of

generators (this is done for strings and things without sharp attack for a thicker sound). Below is what's different or how it can be useful.

The Waveform:

This is the same as the Tone synthesizer. Not too much difference that you can provide here. One thing that you can try is, when your tone is a square wave (bottom slider setting), setting this to more sawtooth. That way, the high-energy saw will disappear into the odd harmonics of the square, which can make for a nice effect.

The Unison Parameters:

The Offset here can be about 3 times larger than for Tone. This makes for less "precise" attack sounds that diminish. Perhaps you want your electric piano to start with a bit of a wallop and then achieve a steady state. Give the Attack synth more offset and amplify the offset generators more.

The Attack Shape:

This is the same as the Tone, except that it decays about twice as quickly in order to facilitate fast "striking" transients. This is generally where you make the Attack synth disappear from the signal, but if you're making a patch for which you want the Attack synth to remain throughout the life of the sound, then have it decay slowly (or set Decay to 0 if there is to be no decay). You can also put it under the ADSR envelope (discussed below).

The Filter Shape:

The decay here is also faster than for the Tone, but is otherwise the same. Just a couple of things to try: setting a very fast Attack, with high cutoffs and high resonance, but that diminishes into a cooler, more stable Tone can be nice. Or, use high resonance with a touch of noise and a fast (high) decay for a good, percussive "whack" effect.

The Attack Frequency:

This is a multiple for the attack frequency, useful for adding a little sparkle. But that's only when the multiple is an integer (1, 2, 3, or 4 in this case). Different intervals cause the Attack synth to be detuned, or just altogether different. Also, you can bind this parameter to the mod wheel and then control it as you play. This is used a great deal in the drones. Most of the time it should be an integer however.

The Noise Parameter:

This is pretty simple. It interjects some noise into the Attack synth. This is generally used with extremely fast decays to provide a percussive effect, but in some cases can be used all the time. Often best with a high resonance to target the noise around the cutoff frequency. Just a touch will typically do. Beware of extra low frequencies.

The FM Parameters:

These cause a small amount of frequency modulation in the Attack synth. The Freq parameter is pretty obviously the frequency by which the frequency of the Attack generator is modulated, and the Depth is the amount (as a multiple of the fundamental) that the frequency will vary. Furthermore, the Depth fades with the amplitude, so the louder your Attack envelope, the more variance you achieve. Therefore, setting high deviation at around 5Hz with an amplitude envelope that starts high and decays quickly to a low amplitude will cause something of a plucked-half-full-wine-glass effect. You can also set the FM frequency very low and make sure the amp envelope remains audible for an LFO effect on the frequency. In small doses, this might be good for stringy things. In larger doses, the effect becomes more unnatural.

The Under Global Envelope Button:

If you're making any kind of percussive, pluck, or otherwise one-shot sound, you would probably want the Attack to start strong and just diminish. But if you're using the attack as just another set of generators, then you might want it under the global envelope. This has been done for the strings and several other presets.

The Effects

The Delay Chambers:

These in some ways emulate the body of an instrument, and can really add life to a sound. Here's how they work, and then a few suggestions (you might want to skip to the suggestions).

Each note has its own set of Delay Chambers whose lengths are set relatively to the fundamental frequency of the note. If you set the delay time to 0 and the feedback to zero, you'll just get the original sound out of the other end. But, as you move a slider, you're causing the delay time inside that chamber to change from 0 to the length of the fundamental. Thus, setting each to 1 would be perfect resonance for a perfect, steady-state waveform if the chambers were perfect. However, Sonata, with its unison, filter, and FM parameters, doesn't produce perfectly steady waveforms, and the chambers themselves are not perfectly resonant, so this setting, though causing a lot of resonance, would not lead to infinite amplitude (it will just get very high). Setting the delay length somewhere between 0 and 1 causes other frequencies to be resonant inside the chamber.

For instance, setting a chamber to .5 will double the frequency that resonates therein – the octave in this case. It's possible to cause phase cancellation to a great degree here, which might make you synth sound an octave or maybe a 13th higher than you're playing. Or, using values of 0 with lots of feedback would cause the sound to “smear”, delaying the amplitude of each sample to the next sample and feeding back upon itself. This just makes the sound a little more full-bodied but doesn't change it much.

Here are some suggestions:

- More feedback means more effect. Just remember to turn down the volume before doing anything drastic.
- Set the length to very low values and turn up the feedback. You'll need to turn the Vol. Min. (see below) down so it won't clip.
- Try one knob full left and the other full right. Move each inward in symmetric intervals. When it sounds good, remember where you are, and go all the way to the center with it. Then go back and pick the best spot.
- For percussion, leave one knob at 0 or very low and adjust the other until you find what you're looking for.
- When featuring bass notes with a fast attack, keep delay values low or the sound will be noticeably delayed.
- Use the text box to type in exact values, like 0.333333333.
- Play. This is important.

The Cross-Delay:

A long-time favorite, the cross-delay, also known as a ping-pong delay, provides an echo that bounces from left to right according to two different delay amounts. Most patches that use this use a tempo-synced delay (the Tempo Sync button is on), so the echoes occur on some fraction of the beat, but some might just want a specific time delay (up to 1 second max). To turn the Cross-Delay on, click on the button at its base.

When Tempo Sync is on, the Delay length varies from a $\frac{1}{4}$ of a beat to 4 beats. It's important to note that a $\frac{1}{4}$ th of a beat is not a quarter note. Sonata doesn't care what your time signature is and doesn't ask. So, if you want sixteenth note echoes in 4/4, use $\frac{1}{4}$ of a beat as your delay time. When you want quarter note echoes in 4/4, use 1. When Tempo Sync is off, the Delay length simply varies from 0 to 1 seconds. This is true for both L (left) and R (right) delay lengths.

The Feedback controls how much of the outgoing signal goes back into the Delay with the input. A feedback of 1 will never end, and Sonata will happily let this happen. Values around 50% are typical. Lastly, the Amp parameter controls how much of the Delay signal is sent to the output. Too much might sound neat at first, but it's often useful to keep this low.

The Distortion:

Most are also pretty familiar with this kind of unit. It distorts and saturates the sound. You can turn it on with the bottom button. Then, Wet controls how much of the signal is distorted. Zero distortion is actually in the middle. Turning the knob clockwise distorts it one way, and turning the knob counter-clockwise distorts it with an opposite phase for possible phase distortion. Dry refers to the amount of (original) dry signal that comes through. It can be useful to have this set low and then, with Aftertouch on, suddenly distort your sound with some extra pressure at the keyboard.

The Reverb:

We're calling this a “near-field” reverb, meant to accent early reflections such as those an instrument's immediate environment, not a massive hall, are inclined to make. It's also pretty cpu-hungry, so if you want to use your big nice reverb instead of this one, turn this one off and go right ahead. It's here to provide a little thickness to the sound or to just be a tool of convenience before whipping out your other reverbs.

Turn it off with the lower button, set the amount of time before the reverb will become very quiet (theoretically, it never reaches zero) with Time, set the amount of reverb signal sent out with the Wet parameter, and, as before, set the amount of original signal passing through with Dry.

For the most useful effect, try keeping values low. Two second reverb time is plenty for this plugin, and a 5-10% wet-to-dry signal ratio (Dry is 1 and Wet is 0.1) will do a lot.

The Controls

The Volume Major Parameter:

This sets the volume in a very big way, from -inf. to +12db. Use it to adjust the volume coming out of the delay chambers to keep it from clipping. Clipping occurs immediately after this parameter. This parameter should not generally be controlled with a MIDI controller, as it is meant to be set and never change for a patch – it is not interpolated, so if you change it with a MIDI controller from 120 to 121, a small click will occur with an amplitude of $1/121^{\text{th}}$ of your base sound. Basically, don't control this parameter with MIDI. That's what the next one is for. This will often need to be very low (especially when the delay chambers have high feedback), so remember that holding shift allows more precise movement while dragging the slider in most hosts. Lastly, **any clipping (saturation) will occur immediately after this parameter and before Volume Minor.**

The Volume Minor Parameter:

This parameter sets the volume from -inf. to 0db and is applied after the Volume Major parameter. This means that **any clipping occurs before the sound gets to this function**, so if you have a really loud sound that you want to saturate, make it loud with Volume Major. Then use Volume Minor to control how loud it comes out. This way, even soft sounds can still be saturated. Also, it means you won't ever have to worry about your MIDI controller making the sound clip because it won't amplify your sound by more than 1. And lastly, this parameter is interpolated, so if you're controlling this and go from a MIDI value of 0 to 127, there will be no click. It will take a very small amount of time to linearly increase to the desired value.

The Panning Parameter:

Panning can be used and controlled within the synth to place the sound left or right. Setting the slider to full left, or 0, is full left sound. Setting 1 is full right. This parameter is also interpolated to safeguard against sudden clicks.

The Aftertouch:

After playing your note at whatever velocity you used, you sometimes want more volume. If your keyboard supports aftertouch, this is easy. Extra pressure equals extra volume – up to twice as much. Sonata supports aftertouch at a global level and on a per-note basis, so any aftertouch-capable keyboard should work. The aftertouch volume increase is interpolated in the same way as Volume Minor. Just remember to turn it on when you want it and off when you don't. You probably don't want your mallet tail to get louder when you accidentally apply a little extra effort.

The MIDI Learn Feature:

This is how Sonata binds MIDI controllers to the synth parameters. Just click the MIDI Learn button, wiggle the parameter you wish to control, and then wiggle the controller you want to use. That controller should now be bound to the parameter you wish to control. Further, the controller will be saved with the preset, so next time you open it up, the same controller still works. In this manner, the Mod Wheel has been bound to different parameters on many of the included presets. Note that only one parameter may be controlled by one MIDI controller, and only one MIDI controller can control a parameter. All applicable parameters can be controlled this way and you can control as many at one time as you want.

The Polyphony Limiter:

Your computer's processor power will limit how many notes can be played. As a result, you don't want to go overboard and suddenly hear your computer lose real-time-playability. The Polyphony parameters allow you to specify how many notes can be playing simultaneously, from 0 to 127, using the Plus and Minus buttons to increase or decrease the number of notes. Note that the maximum polyphony setting exists for each preset, and e.g. setting it to 4 in Preset A doesn't mean it will be 4 when you load Preset B.

Here's how it works: When the number of notes being played goes over what is set for the maximum polyphony, the oldest note is released, as if your finger had simply let up. The newest notes are still played immediately. Since release tails can be long, however, this is not enough. The polyphony manager in Sonata will only allow two notes to be in this normal release stage while the maximum number of notes are also being played. After that, Sonata will cut off the old released notes very quickly, but not to the point of clicking. Here's an example:

Let's say I select the synth string patch and set the polyphony to 3. I play C2, then D5, then E5. That's 3. Now I play F5 which goes 1 over max polyphony. What happens is C2 releases, just as if I had let up my finger. But the synth strings have a long release, so let's say I play G5 before C2 is done with its release. D5 now releases (the second note that was played). That leaves C2 and D5 in their release stages, and E5, F5, and G5 all on as normal. Let me put a logical split here:

If you play A5 before C2 is finished releasing, A5 will begin, C2 will release in about .0005s or so, and E5 (the third note played) will release as normal. So, if you set polyphony to 3, play C2 and then 5 high notes all at once, you should hear C2 disappear quickly.

But if C2 has finished releasing before you play A5, then playing A5 will release E5, leaving D5 and E5 in their release stages, and F5, G5, and A5 playing as normal.

This is meant to sound very natural without sudden disappearance of notes where possible. However, it's not perfect, so we recommend using the Freeze function if you have it.

The Text Input Box:

This is quite simply, really. When you move a parameter, the new value shows up in this box. You can click in this box and type something from 0 to 1, or one of the following ranges, to set the parameter specifically. Very useful to achieve extremely small or precise values.

The following few parameters should not have values entered as 0 to 1 (it would not be clear how to achieve the desired value using this range), but should use what is specified instead:

- Attack Frequency: use 1 to 4

- FM Frequency: use .001 to 20
- Reverb: use 0 to 5
- Distortion: use -1 to 1

~

The Presets

Crystal: presets 1 through 8

These are lovely chimey things, but they don't have to stay in the high end only. Some create wonderful textures dueling between a low octave strike and a playful arpeggio up top. Most parameters don't have any controls, but some control the wave mix to change the warmth of the sound.

Percussive: presets 9 through 16

These are somehow whacked, and the Tone usually decays away pretty quickly, but not always. More earthy textures than the Crystal sounds, but there still aren't really any control parameters here.

Strings and Horns: presets 17 through 24

These are lovely little synthy things. Not exactly what was heard in Carnegie Hall last weekend, but quite at home for simple synth stuff. Play the strings slowly. Play the brass in perfect fifths (e.g. C and G above it) and stay modal (only on the notes of the scale) for some simple, dated horn harmonies.

Electric Pianos: presets 25 through 31

Electrics only. Best when cool, but some bright stuff in there too. There aren't any control parameters by default, but try putting some on the release tail, with longer tails for the atmospheric parts and shorter for the more demanding ones.

Flutes: presets 32 through 35

These really play on the resonance of the Delay Chambers. In fact, it's crucial. In case you're wondering, the Dizi and Xiao are wonderful traditional Chinese instruments. The Dizi is high (around an octave above middle C) and happy. The Xiao is lower (around middle C) and more of the scholar's instrument. Play in a pentatonic scale (an

easy way to do this is to play only black notes).

Plucked: presets 36 through 40

Try adding just a touch of pitch bend to the beginning of a note strike for the Pipa, also a traditional Chinese instrument. Release the pitch wheel the moment you play the note. Also try playing in a pentatonic scale, easily accomplished by using only the black keys. The plucked sounds don't have quite the percussive attack of the Percussive section, but similar sharp beginnings and longer trailing ends.

Basses: presets 33 through 44

Deep – around C2. Add a pinch of EQ and distort to taste. No complicated classic rock bass riffs – just stick to the simple stuff. Sometimes these are good in treble too. Sometimes not.

Synths: presets 45 through 60

No music for airports. Sonata is just too digital and not typically mellow enough. Keep the atmospheres rich and moving. Take a ride to Saturn's moon, write the next Zelda-for-Gameboy soundtrack, add a touch of screaminess to your angst, or go spelunking. Control the volume on heavily distorted things. Don't scream all the time – just when you need to.

Play augmented seconds (enharmonically, a minor third typically followed by a half step) in your scale with the Idioglottal Synth. Try the mod wheel on Key String.

Drones: presets 61 through 64

The mod wheel will be your best friend here. Move slowly. Try adding further midi controllers to your attack pitch if it's not set already (or resonance if you're feeling particularly violent) and keep in mind that the pitch wheel can create very alien moments within these. Best kept mono, in octaves, or maybe fifths.

Trouble Shooting

1. The attack envelope doesn't seem to work when I re-strike notes quickly.
 - Perhaps you have a sound with a fast “AD” envelope. When you restrike, the ADSR envelope builds from its last amplitude. So, if you have a sound with a fast AD envelope, you'd miss it during a long attack envelope. However, when restriking, you're already at an audible amplitude, so that part you didn't know was there actually is. Try making this initial attack in the AD envelopes more mild so you're not producing sound you're not meant to hear.

2. Polyphony is letting too many notes go by.
 - When the maximum polyphony is reached, the synth releases the oldest notes, just like a regular note-up (except it releases regardless of the sustain pedal). However, if more than two notes are in this release stage when yet more notes are played, these notes are quickly cut out. This prevents some awkward cut-outs. So, when setting polyphony to a strict limit, keep these extra two notes in mind. See the Polyphony section for further details.
3. I'm holding down a note and changing the filter cutoff pitch (either Init or Final), but nothing is happening.
 - When a note is struck, the high and low points of it's envelope are set for it's entire life. You can't change them mid-way (the envelope would get lost). So, if you're changing the parameters, it will effect the next note you play.
4. It's called Sonata, but I can't play Moonlight. That needs 30-something note polyphony.
 - If you must have huge polyphony, you could play into your standard midi bank or some synth with negligible CPU, or something external, then point that midi to Sonata and freeze it. Not ideal, but better than compromised sound, right?

Credits

We'd like to thank several people, not the least of which are friends who put up with us during the development of this plugin. We'd like to thank the beta testers, who we will not list here. You all provided an excellent beta experience. And furthermore, we'd like to thank a few who provided the extra presets that appear in the factory bank (others have provided excellent banks which will be hosted and credited on the website):

Tim Conrardy (KVR: tconrardy)

John (KVR: JoMal)

Peter (KVR: pschelfh)

FrankieG (KVR: ToTCProductions)

The Company

This synth was programmed by Tucker McClure and Grant Hutchins – programmers, students, and avid electronic music fans. We can be contacted at the following addresses:

Tucker: tucker@sohasound.com

Grant: grant@sohasound.com

Or, visit us online at www.sohasound.com.