Sanestation - Manual

March 5, 2012

SaneStation is a subtractive synth, connected as follows: A voice begins as oscillators, is filtered with a filter before the amplitude is modulated by an envelope. Then, the signals from all voices belonging to a channel are collected, LPC is run, channel effects (insert effects) are run, and aux bues are filled with signal. The signals from all channels are collected, global effects (send effects) are run, and the entire mix is sent through a final compressor/limiter. The next chapters go into further details of each of these parts. (For full schematics, see appendix B.)

The synth features:

- Three oscillators per voice (sine, saw, square, triangle, noise, all properly bandwidth limited with correction for the Gibbs phenomenon, supporting coarse tuning, fine tuning and multi-osc/detune), with ring modulation and FM synthesis.
- Per-channel voice limits (monophony ala 303! :-)).
- Two LFOs with flexible routing, with restart on note-on.
- Per-voice 12dB/oct (selectable) filters, selectable from highpass, lowpass, bandpass, allpass and notch (including waveshaping versions of all of the above), with adjustable cutoff and resonance.
- Moog-style ladder filters, selectable from lowpass (6, 12, 18 or 24 dB/oct), bandpass, highpass or notch (all 24 dB/oct).
- Two per-channel effects: First is selectable from distortion/limiting, compression (with fully adjustable threshold, ratio, attack/decay times and post-gain), bitreduction, downsampling and extra filter; second selectable from stereo delay, chorus, flanger etc. via LFO.
- Portamento and pitch bend.
- Two aux send buses: one for reverb, one for delay effects
- Per-channel and global four-band parametric EQ, with freely chooseable type, cutoff, resonance/Q factor, and gain for each filter.
- A post-mixing multiband compressor/limiter.
- General, flexible influence/routing system.
- A full-featured VST interface for music creation.
- Retrigger box (with time stretching) for global glitching.
- Vocal compression vaguely resembling GSM.

1 Voice

1.1 Oscillators

SaneStation has three oscillators, which can be either off, sine, sawtooth, square, triangle or noise. Pitch can be controlled with **coarse tune**, setting tuning in half notes, and **fine tune**, which go from -100 to 100 cent. In addition, they are influenced by **portamento**, and can be affected by influences, which will be described later.

Multi-osc is also supported, splitting the oscillator into three or more oscillators that are slightly detuned from each other. Multi-osc is activated by setting the **multi count** controller larger than 1 (seven is a typical value), and the **multi spread** controller larger than zero. The spread value specifies detune in half notes from the center to the outmost oscillator; for instance, if it is set to 0.3 and the count is set to 7, the seven oscillators will be at -0.3, -0.2, -0.1, 0, +0.1, +0.2 and +0.3 half notes from the base pitch. Experiment with different values for different sounds, but remember that multi-osc uses a lot of CPU, since all sub-oscillators must be rendered individually. If you enable **random phase**, the sub-oscillators will start out of phase with each other. This negates the “sweeping effect” you typically get with multi-osc, which you may or may not want.

The oscillators are mixed as follows: If **ring modulator** is off, osc1 and osc2 are mixed together depending on **osc1-osc2** mix, where all left means osc1 only, and all right means osc2 only. If **ring modulator** is on, however, the oscillators are multiplied together (osc1-osc2 mix setting is ignored). Finally, osc3 is added, with amplitude determined by **osc3 level**. The audio is then sent to the filters.

1.2 Filters

Available filter types are: off, lowpass filter, highpass filter, bandpass filter (bandwidth is set using the resonance controller), notch (again, bandwidth is set using the resonance controller), and allpass filter. All filters are available in a **waveshaping** version, which usually introduces a few extra harmonies in the signal. For instance, the allpass filter is pretty useless without waveshaping, but can make a slightly unusual sound with waveshaping enabled. Cutoff follows the **cutoff** slider and any influences. If **key follow** is enabled, the height of the note is also added to cutoff. Resonance follows the **resonance** slider and any influences. The **filter depth** parameter sets the strength of the filter, from 0 to 48 dB/oct in steps of 12 dB/oct. The value specifies how strongly the filter dampens in the stop bands.

Below the filter, there is a frequency response display, showing the filter’s amplitude response as function of the frequency. This is accurate (ie., it is calculated from the real filter). The y axis has dB scale (the gray line marks 0 dB), and goes from -80 dB to 20 dB. The x axis shows frequency; it is also logarithmic, and goes from 20 Hz to 20000 Hz.

1.2.1 Ladder filters

The ladder filter is a digital version of the classic Moog ladder filter. It is nonlinear and has a different sounds from the other filters. Unlike the regular filters, where depth and type are independently setup, it comes in a limited set of configurations (all with the same CPU cost, again unlike the normal filters). These are: 6 dB/oct lowpass**,** 12 dB/oct lowpass**,** 18 dB/oct lowpass, 24 dB/oct lowpass, 24 dB/oct bandpass, 24 dB/oct highpass and notch. Note that the cutoff of these filters might not be 100% precise (they are more difficult to tune exactly then the other filters); they should, however, be pretty close.

The frequency response display is correct like the other filters; however, they cannot capture the inherent nonlinearity of the filter, so take it with a grain of salt. (It is also much more CPU-intensive to draw, so expect it to be somewhat less responsive currently.)

The ladder filter can, like the regular filters, be used as a channel effect; due to the nonlinearity this may or may not give interesting results.

The ladder filters are oversampled four times, and thus rather CPU intensive. Use them with care, and see section 4.1.

1.3 Envelopes

SaneStation has three available envelopes. Envelope 1 always influences the amplitude; in addition to this, all envelopes can be used freely in influences. The three envelopes have somewhat different parameter ranges: While envelope 1 is normal, envelope 2 is suited for fast envelopes, and envelope 3 for slowly changing envelopes.

1.4 LFOs

SaneStation has two LFOs per voice, both triggered on note start. The results are usually best if one chooses sine or triangle. LFO2 goes higher in frequency than LFO1. LFOs do not carry any meaning unless they are used in influences.

1.5 Influences

The influence system works as follow: For a given destination, one specifies a source and an influence factor. The source value multiplied by the influence factor is added to the destination. If one pulls the influence factor all the way to the right (you can see in the status box when this happens), the destination is instead modulated by the source.

Possible sources are: Nothing, envelope 1-3, LFO1-3, oscillator 1-3, and resonance. Possible destinations are: Oscillator 1-3 pitch, resonance, cutoff up to three times (for multiple funky effects), and amplitude. Note that envelope 1 is already automatically routed to the amplitude, so you won't have to set up that. A source can not be used as influence before it's rendered, which is why not all sources are available for all destinations.

A few examples:

1. 303-style filters: 303 lets the filter follow the amplitude envelope via the env-mod knob. SaneStation can emulate this by taking one of the cutoff destinations, set the source to "env1" and amount to a suitable value (the amount here corresponds to the 303 env-mod knob).

2. Vibrato: Tweak LFO1 to a suitable vibrato speed, set LFO1 as source for oscillator pitch (for whatever oscillators you're using), and set amount slightly away from zero.

3. FM between two oscillators: Set osc1 as source for osc2 pitch, set amount to a suitable value, and make sure osc1-osc2 mix is all the way to the right.

The influence controls are found in the mix tab of SaneVST.

2 Channel

Maximum number of voices per channel is determined by the **max voices** control. If you want a monophonic synth sound or convincing portamento, you would want to set this to one. Even for polyphone, it might be a good idea to put this as low as possible to use as little CPU time as possible.

After all voices corresponding to a channel are rendered, the result goes through LPC and the channel effects, if any. We have two types of channel effects: mono effects and stereo effects, and each channel can use up to one of each.

2.1 Linear Predictive Coding (LPC)

LPC is a variation of the speech coding system used in GSM, and works by splitting a source signal into another (more basic) source signal and a frequency spectrum. Classic LPC switches between a wideband pulse oscillator for voiced sounds (ie. vowels) and a noise oscillator for unvoiced sounds (ie. consonants), and runs that signal through the frequency filter to get the compressed sound back. However, in SaneStation we already have a full synth available for producing the source signal, so we only have to store the frequency spectrums. Thus, using LPC in SaneStation is a bit like using a vocoder: The carrier signal comes from a synth channel, which is then modulated with the sample currently chosen in the LPC module. Thus, LPC samples are compressed to 50 spectrums per second, where each spectrum is described by eleven coefficients; a pretty hefty amount of compression. As far as we know, no intro has used something similar before (excluding hooking onto OS-included codecs).

LPC is governed by two sets of controllers: The first are in the global tab, and control loading of samples into LPC slots. LPC slots are a global resource, and there are 127 of them, numbered 1-127. You choose a slot with the **lpc slot** control, and load an audio file with the **select wave...** button. The LPC modul accepts *mono* PCM wave files at 44, 22 or 11 kHz. If you give it something else, it will answer with a less than useful error message ala "File operation gave error 1", so don't do that. **Clear** empties the slot, and with **name**, you can give the sample a name. The name is by default set to the filename sans the .wav ending whenever you load a new audio file, and is always "(unused)" when no sample is loaded.

When a few samples is loaded, you most likely would want to use them as well. LPC is triggered by sending the **LPC select** controller (controller 94) with value equal to the number of the LPC sample you want to use. Zero is always "turn off LPC", and using an empty slot also gives the effect that LPC is disabled. However, if the select controller is triggered with some existing sample (or re-triggered with the same sample), that sample is started from the beginning, and used to modulate whatever is played on the channel. When playing past the end of a sample, or before beginning of a sample (yes, it's possible), one gets silence. In the mix tab, there's the **LPC speed** control, which goes from -2 to 2. 1 is normal speed, 2 is twice normal speed, -1 is normal speed in reverse, 0 is to stand completely still etc. When changing the speed, one changes how quickly the frequency filter is changing, but pitch is not changed — think time stretch, not changing the playback speed of a normal sample. In addition, the tab has an **LPC select** control. The controller is only sent to the channel when changing it or at patch change, so it's there mostly for quick preview of a sample.

A couple of tips:

- In order to save space, the sample is downconverted to 11 kHz before LPC encoding. (Ideally, resample the audio to 11 kHz yourself or just record it at 11 kHz in the first place; the downsampling is not superb, although more than adequate for most purposes.) This means that the audio spectrum is only correct between 0-5 kHz. 5-11 kHz is a reversed copy of the 0-5 kHz part, 11-16 kHz is an exact copy, and 16-22 kHz is a reversed copy again. This can sometimes be a desirable effect, and sometimes not. If you don't want it, set the mono effect to filter, the type to lowpass, the order to 12 or 24 dB/oct, the resonance to one (neutral), and the cutoff around 5500 Hz (apply your own best judgement here).
- If you want LPC encoded voices to be as clear as possible, make sure the carrier signal is rich in frequency, giving the LPC filter something to work with. Use square or saw, don't be afraid to use the ring modulator, and if you mix in some noise, "s"-es in speech sound better.
- The better (ie. clearer/less noisy) your input signals are, the better the output will be. Note that the samples are automatically (global) gain adjusted to try to minimize the compression artifacts; this can be somewhat unpredictable.

2.2 Mono Effects

The following mono effects are available: No effect, distortion, compression, lo-fi and filter. You have two slots which you can use as you'd like (e.g., distortion and compression after each other, distortion twice with no parameters, distortion only once, etc.).

2.2.1 Distortion

A simple soft-knee distortion module with four available parameters:

- **Threshold**: Sets the threshold the signal level can reach before the signal is affected by the distortion function.
- **Ratio**: Sets how strongly the signal above the threshold is to be clipped. Higher number = harder clipping.
- **Asymmetry**: Adds a DC offset to the signal before it is processed, adding even harmonics.
- **Post gain**: Level adjustment after the effect.

2.2.2 Compression

A hard-knee mono compressor with five parameters:

- **Threshold**: Sets the threshold the signal level can reach before the signal is affected by the compressors.
- **Ratio**: Sets how strongly the signal above the threshold is to be dampened. Higher number = harder dampening.
- **Attack time**: Sets how long the signal level must stay above the threshold before the compressor starts dampening the signal.
- **Release time**: Sets how long the signal level must stay below the threshold before the compressor stops dampening the signal.
- **Makeup gain**: Level adjustment after the effect.

2.2.3 Lo-Fi

A digital lo-fi effect. Hopefully this will come in handy at some point :-) It has the following parameters:

- **Bit reduction**: Reduces the amount of bits the signal uses. You can use this to simulate old trackers with 8-bit samples.
- **Downsampling**: Resamples the signal by a factor given by this controller. For instance, if the controller is set to 4, the signal will be resampled to 11 kHz (= 44 kHz/4), and then resampled back again. The resampling does of course happen without any form of filtering, for lovely grunge sound.

2.2.4 Filter

An extra filter, should the per-voice filter not be enough. The controllers are cutoff, resonance, filter type and filter order. Take a look in section 1.2 to see what they do.

2.2.5 Exciter

A very simple variant of what you could get out of old radio tubes — adds some extra oomph in the harmonics (and a bunch of DC offset as a side effect).

2.3 Stereo Effects

After the mono effect, the signal is mixed to stereo with panning controller by the **panning** control and volume controlled by the **channel volume** controller. Available stereo effects are: no effect, delay, chorus, and flanger.

2.3.1 Delay

- **Left delay:** Left channel delay, in milliseconds.
- **Right delay:** Right channel delay, in milliseconds.
- **Feedback:** Controls how much of the output signal is mixed back into the effect.
- **Wet level**: How much of the result is mixed into the main signal.

2.3.2 Chorus

- **Delay:** Base delay, in milliseconds.
- **LFO speed:** The speed of the LFO modulating the delay length.
- **LFO amount:** The strength of the LFO.
- **Feedback:** Controls how much of the output signal is sent back into the effect.
- **LFO type:** The type of the LFO oscillator. Continuous waveforms such as sine and triangle work best; the others will introduce clicks in the audio.
- **Wet level**: How much of the result is mixed into the main signal.

2.3.3 Flanger

- **Delay:** Base delay, in milliseconds.
- **LFO speed:** The speed of the LFO modulating the delay length.
- **LFO amount:** The strength of the LFO.
- **Feedback:** Feedback — higher value gives a more metallic sound.
- **LFO type:** The type of the LFO oscillator. Continuous waveforms such as sine and triangle work best; the others will introduce clicks in the audio.
- **Wet level**: How much of the result is mixed into the main signal.

2.4 Buses

SaneStation has conceptually three buses; the main bus, and two send buses for connected to the reverb and delay global effects. The controllers governing this are **reverb send** and **delay send**. All are post-fader.

2.5 EQ

Each channel has a separate four-band parametric EQ. Each band consists of a 12 dB/oct filter; these are freely choosable, but the most useful configuration is probably one **low shelf**, two **peaking EQ** and one **high shelf**. You can choose cutoff and resonance/Q factor freely; for the shelf and peakEQ filters, you also need a gain setting (from -20 to +20 dB; in the default of 0 dB, the filter has no effect save for possibly phase distortion).

There is also a global EQ section, which works just the same as the parametric EQ.

3 Global Mixing

Some effects are added at the very end. Note that if you try to change these while audio is being played, you can get somewhat unpredictable clicking — in other words, it's perfectly fine to change them while mastering, but you don't want to change them in the middle of a finished tune. (Generally, this should not be a problem; you shouldn't generally need to change reverb or compressor settings underway.)

3.1 Reverb

This is the reverb at the end of the reverb aux bus. The code originates in Freeverb. There are four parameters: **room size**, setting the size of the room the reverb simulates; **damp**, dampening of the echo as time passes; **width**, the stereo width; and **wet**, how much of the result is mixed into the final mix.

3.2 Delay

The contents of the delay aux bus go here. We have the following parameters: **left time**, left channel delay, in milliseconds; **right time**, right channel delay, in milliseconds; and **feedback**.

3.3 Retrigger Module

The retrigger module is almost last in the mix (after reverb, delay, and EQ contribution from the channels), and is meant for making glitchy sound. It is essentially a somewhat specially configured delay box with timestretching. As it is meter-synced and triggered by controller changes, it is generally hard to use without a sequencer (ie. directly from the GUI).

The retrigger module essentially works as follows: It records audio corresponding to a single beat, and then played it back in N time the speed (with time stretching, so there is no pitch difference), N times, ie. also in a single beat. While the audio is played, a new beat is recorded, which is then played the next beat, etc. Output from the box will thus always trail the input by a single beat.

The module has three parameters: **bpm**, **speedup factor** and **wet level**. The BPM controller governs how long each recording sequence is (for instance, bpm=120 means half a second is recorded at a time). **Speedup factor** is how much quicker the playback is (N=3 will for instance mean that the beat is played back at triple speed, three times). Negative values mean the buffer is played back in reverse (e.g.: N=-3 means the same as N=3, except that the audio is played backwards instead of forwards).

The box starts recording (and playing back) whenever the bpm or speedup controllers are changed. Typical usage would in other words be setting correct bpm and speedup=1 right at the start of the project (or at some earlier beat) to make the box start recording. Then, one beat before you want it to glitch, set e.g. speedup=2, and there will be glitched audio beginning one beat later. Accordingly, one must set speedup=1 one beat before one wants normal sound to resume. (As long as speedup=1 and the bpm controller is not changed, the box will function as a simple delay of exactly one beat.)

Finally, **wet level** will determine how much of the effect is mixed in. As previously mentioned, the output of the retrigger box is delayed, so most likely it's not so useful to set this at anything but 0.0 or 1.0, but it might be possible to create weird effects using other values.

Special case: If speedup=0 (the default setting), the retrigger box will eventually be turned completely off to save CPU. It will not record audio, nor play back any meaningful audio. In other words, this is only useful if you're not actually using it, ie. have wet level at 0.0 (speedup=0 with wet level above 0 will give undefined and unpredictable results). If you plan to use the retrigger box later, you'll need to set speedup=1 at least one beat before you actually need the output, to give it time to "start back up again".

NOTE: The retrigger box is **not code-complete**; its output is currently in **mono** (it only processes the left channel, and makes the right channel a copy of the right channel). Please let the coders know if you actually use it for something, and we'll lift this restriction.

3.4 Compressor

At the end of the mix, we have a four-band compressor/limiter. Its parameters are: **pre-gain**, pre-gain before the compressor; and then for each band: **makeup gain**, gain after the compressor; **threshold**; **ratio**; **attack time**; **release time**, all as documented in section 2.1.2.

Typical setting to use it as a limiter will be e.g. setting attack time all the way to the left, release time to somewhat above middle, ratio all the way to the right, threshold to -0.1 dB, and then adjust pre-gain and release time to taste. The band limits are adjustable; it is possible to make meaningless setups (e.g. band 1/2 crossover at 5000 Hz and band 2/3 crossover at 500 Hz), but if so, you just end up wasting a band (in this case, band 2 will not get any audio routed to it).

Finally, audio is hard-clipped before it is sent to the sound card.

4 Optimization

4.1 Speed Optimization

SaneStation is not a very slow synth — generally, it uses about 10% CPU time on a 1.2 GHz machine with SSE. (SSE routines take slightly more space, though, so they might not be enabled in a size-constrained intro.) If there is still lack of CPU for some reason, there are several things you can do:

- Filters take time to calculate, so if you can do without filter, or less sharp ones (e.g. 12 dB/oct vs. 48 dB/oct), it will save CPU. The same goes for EQ usage; each used band on a channel costs the same as a regular 12 dB/oct filter. Ladder filters are even more expensive (but all ladder filter configurations have the same CPU cost). If you do not need influences or key follow, consider using a channel filter instead (available as a mono effect), which will be cheaper if you have multiple voices on the channel.

- Also, you should not use resources you don't need — if you don't use an oscillator, set the type to None; if you don't use an influence, set the influence source to None. LFOs and envelopes will not be rendered unless there's something using them.

- Limit the use of channel effects as much as possible, especially flange and chorus, which are particularly slow.

- You should also turn down max voices as far as you can, to reduce the amount of voices calculated at once. This is especially useful if one or more of the oscillators for that voice is set to multi-osc, since seven oscillators must be rendered instead of one (multi-osc can be a CPU hog).

- If you can do with fewer bands in the multiband compressor, this will also help, but the synth needs to be recompiled to change the number of bands.

4.2 Size Optimization

The things that take up space in the final product (in addition to code) is instrument data and note data. Instrument data is reduced by using as few instruments as possible (including removing unused ones). To minimize note data, controllers take a lot of space, so limit the use of unneeded controllers. Time stamps and controller values are stored delta encoded, so quantized notes compress well, and controller slides also compress a lot better if they're as regular as possible.

LPC samples cost 225–250 bytes per second of sample material. (You can see an estimated size in the GUI when you load a sample.) Needless to say, reuse as much as you can.

4.2.1 Feature removal

Should you still be cramped for space, you can ask one of the coders to remove features from the synth to reduce the amount of code. (Of course, then the feature will not be available to you either.) Note that this is generally a last-ditch hack; it is not necessarily something that can be done cleanly in the code. Nevertheless, if you're over the limit and need to get under…

The synth and MIDI player takes roughly 12.5 kB after executable compression (with kkrunchy), in a size-optimized build (MSVC, no SSE support, no text display). Below is a list of roughly how much you can gain by removing certain features (again after executable compression). They are not an exact science, but should help to give you a hint of what is an expensive or a cheap feature size-wise.

Feature	Space savings (approx.)
LPC	1200 bytes
Retrigger	550 bytes
Multiband final compressor[1]	300 bytes
Multi-osc	110 bytes
Filter types, per type[2]	20 bytes
Filter waveshaping	110 bytes
Ladder filters[3]	650 bytes
Pitch bend	80 bytes
Portamento	60 bytes

[1] If disabling the multiband final compressor, it will be replaced by a single-band compressor that is otherwise identical.

[2] A filter type is e.g. notch or bandpass (so removing both would yield approx. 40 bytes). Note that lowpass and highpass cannot be removed, as they are used internally by the final compressor and other effects.

[3] Removing individual ladder filter configurations will save very little space (est. 5 bytes per configuration).

5 SaneStation and Renoise

SaneStation should work fine in most VST hosts, including Renoise, but as Renoise's VST support has a few somewhat unusual quirks, we will here give a few hints to get the most out of using the two together. It is assumed that you are already familiar with the Renoise interface, and that the SaneStation DLL is in your VST path.

The main challenge is to get multiple instruments working at the same time. SaneStation has 16 channels and one program (instrument) on each (there is no MIDI mapping), but by default Renoise only sends data to channel 1. There are two approaches, which we will visit in turn.

5.1 Multiple VST instances

This is probably the easiest to set up: Simply set up two Renoise instruments (e.g., 00 and 01), both with "SaneStation" as instrument. This will create two entirely separate instances of the synth, with separate interface etc.. However, this means that each instance will not see the other's sound data, which means that the mastering section is significantly less useful (each channel would be compressed separately). Also, programs are not shared between the instances, and this method uses more CPU than having a single instance. Thus, it is not recommended except for simple experimentation.

5.2 VST aliases

Renoise has a concept of "VST aliases", which are Renoise instruments that point to the same VST instance. The only difference between the two instruments (aliases) is that they can send to different MIDI channels. To set up an alias, first set up the VST as usual on instrument 00. Then, in instrument 01, you can now choose SaneStation under "VST aliases" as an instrument, and then set a different channel (e.g. 2). You now have two Renoise instruments sending data to different channels of the same synth.

By default, MIDI channels are set up so that channel 1 has instrument 0, channel 2 has instrument 1, etc.. (Changing MIDI programs in the Renoise "instrument settings" tab only controls which one is visible in the editor; it does not change the audio output.) If you want to change this, you can send a MIDI program change to one of the channels by using the C2 pan command; see the Renoise documentation for more information.

Using VST aliases is the preferred method, since it will allow you to master the song as a whole (using the final compressor), and will most closely match how the size-optimized player (as well as non-Renoise VST hosts) works.

6 Appendix A: MIDI Control of SaneStation

Most of the parameters in SaneStation are controllable using MIDI controllers.

6.1 MIDI controllers

This is the table of MIDI controller assignments and their ranges. The columns are as follows:

- Ctrl: Controller number. Will mostly match official MIDI mappings where it makes sense (see http://www.indiana.edu/~emusi
- Scope: Whether the controller affects only the current channel it is set in, or the entire synth globally. Global controllers can be set in any channel, although it is customary to set them in channel 1.
- Granularity: How often the controller is checked. There are three options:
  - Sample: Sample-accurate (or every Nth sample). Can be changed freely at any time.
  - Note-on: The controller value at the time of the note-on command is used. Changes to the controller will affect future notes, but not any notes still playing.
  - Fixed: Used for controllers that should not change during a song. You may change them as you wish in the GUI, but expect clicking or other artifacts when doing so. Attempts to change them as part of a song may lead to unpredictable effects in the final player, or even get ignored (some of the fixed parameters get hard-coded into the player at compile time).
- Mapping: Floating-point values the input controller values 0..127 are mapped to. Sometimes a unit is also given. “cubed” means that the mapping is not linear between the endpoints, but x^3-mapped, so there is more resolution closer to zero.

Ctrl	Scope	Granularity	Description	Range	Notes
7	Channel	Sample	Main volume	0.0..1.0	cubed
10	Channel	Sample	Panning	-1.0..1.0	
Bus 1 (reverb) settings					
12	Global	Fixed	Bus 1 reverb: Room size	0.0..1.0	
13	Global	Fixed	Bux 1 reverb: Damp level	0.0..1.0	
14	Global	Fixed	Bus 1 reverb: Width	0.0..1.0	
15	Global	Fixed	Bus 1 wet level	0.0..1.0	
Bus 2 (delay) settings					
21	Global	Fixed	Bus 2 delay: Left delay time	0.01..1.0	seconds
22	Global	Fixed	Bus 2 delay: Right delay time	0.01..1.0	seconds
23	Global	Fixed	Bus 2 delay: Feedback	0.0..0.99	
Bus 1/2 filter settings					
24	Global	Fixed	Bus highpass filter cutoff	20.0..500.0	Hz
Channel bus send					
25	Channel	Fixed	Channel→Bus 1 level	0.0..1.0	cubed
26	Channel	Fixed	Channel→Bus 2 level	0.0..1.0	cubed
27	Channel	Fixed	Bass level	-1.0..3.0	cubed
28	Channel	Fixed	Mid level	-1.0..3.0	cubed
29	Channel	Fixed	Treble level	-1.0..3.0	cubed
Final compressor					
30	Global	Fixed	Final compressor pre-gain	0.0..4.0	cubed
31	Global	Fixed	Final compressor band 1: Threshold	0.0..1.0	cubed
32	Global	Fixed	Final compressor band 1: Ratio	1.0..∞	cubed
33	Global	Fixed	Final compressor band 1: Attack time	0.000..0.020	seconds, cubed
34	Global	Fixed	Final compressor band 1: Release time	0.020..1.000	seconds, cubed
35	Global	Fixed	Final compressor band 1: Makeup gain	0.0..4.0	cubed
36	Global	Fixed	Final compressor band 1/2 frequency	50..15000	Hz, cubed
37	Global	Fixed	Final compressor band 2: Threshold	0.0..1.0	cubed
38	Global	Fixed	Final compressor band 2: Ratio	1.0..∞	cubed
39	Global	Fixed	Final compressor band 2: Attack time	0.000..0.020	seconds, cubed
40	Global	Fixed	Final compressor band 2: Release time	0.020..1.000	seconds, cubed
41	Global	Fixed	Final compressor band 2: Makeup gain	0.0..4.0	cubed
42	Global	Fixed	Final compressor band 2/3 frequency	50..15000	Hz, cubed
43	Global	Fixed	Final compressor band 3: Threshold	0.0..1.0	cubed
44	Global	Fixed	Final compressor band 3: Ratio	1.0..∞	cubed
45	Global	Fixed	Final compressor band 3: Attack time	0.000..0.020	seconds, cubed
46	Global	Fixed	Final compressor band 3: Release time	0.020..1.000	seconds, cubed
47	Global	Fixed	Final compressor band 3: Makeup gain	0.0..4.0	cubed
48	Global	Fixed	Final compressor band 3/4 frequency	50..15000	Hz, cubed
49	Global	Fixed	Final compressor band 4: Threshold	0.0..1.0	cubed
50	Global	Fixed	Final compressor band 4: Ratio	1.0..∞	cubed
51	Global	Fixed	Final compressor band 4: Attack time	0.000..0.020	seconds, cubed
52	Global	Fixed	Final compressor band 4: Release time	0.020..1.000	seconds, cubed
53	Global	Fixed	Final compressor band 4: Makeup gain	0.0..4.0	cubed
Retrigger					
54	Global	Sample	Retrigger: BPM	60.0..187.0	
55	Global	Sample	Retrigger: Speedup factor	-63..64	
56	Global	Sample	Retrigger: Wet level	0.0..1.0	
When channel effect 1 is set to ***distortion*** *(value 1)*					
57/78	Channel	Fixed	Distortion: Threshold	0.0..1.5	cubed
58/79	Channel	Fixed	Distortion: Ratio	1.0..20.0	cubed
59/80	Channel	Fixed	Distortion: Asymmetry	0.0..1.0	cubed
60/81	Channel	Fixed	Distortion: Makeup gain	0.0..8.0	cubed

Ctrl	Scope	Granularity	Description	Range	Notes
When channel effect 1 is set to ***compression*** *(value 2)*					
57/78	Channel	Fixed	Compression: Threshold	0.0..1.5	cubed
58/79	Channel	Fixed	Compression: Ratio	1.0..64.0	cubed
59/80	Channel	Fixed	Compression: Attack time	0.000..0.020	seconds, cubed
60/81	Channel	Fixed	Compression: Release time	0.020..1.000	seconds, cubed
61/82	Channel	Fixed	Compression: Makeup gain	0.0..20.0	cubed
When channel effect 1 is set to ***lo-fi*** *(value 3)*					
57/78	Channel	Fixed	Lo-fi: Bit reduction	0.0..24.0	bits
58/78	Channel	Fixed	Lo-fi: Downsampling factor	0..127	
When channel effect 1 is set to ***filter*** *(value 4)*					
57/78	Channel	Sample	Filter: Cutoff	20.0..20480.0	Hz, logarithmic
58/79	Channel	Sample	Filter: Resonance	0.0..7.5	
59/80	Channel	Fixed	Filter: Filter type	See 5.2	
60/81	Channel	Fixed	Filter: Order (in 12 dB/oct)	0..127	only 1..4 valid
When channel effect 1/1B is set to ***exciter*** *(value 5)*					
57/78	Channel	Fixed	Exciter: Amount	-1.0..0.98	cubed
Oscillators					
62	Channel	Sample	Oscillator 1/2 balance	0.0..1.0	
63	Channel	Sample	Oscillator 3 volume	0.0..1.5	cubed
Portamento					
65	Channel	Note-on	Portamento time	0.0..1.0	seconds, cubed
Filter					
74	Channel	Sample	Filter cutoff	20.0..20480.0	Hz, logarithmic
77	Channel	Sample	Filter resonance	0.0..7.5	
LFOs					
84	Channel	Sample	LFO 1 speed	0.000001..20	Hz, cubed
85	Channel	Sample	LFO 2 speed	0.001..125	Hz, cubed
When channel effect 2 is set to ***delay*** *(value 1)*					
86	Channel	Fixed	Delay: Left delay time	0.010..1.000	seconds
87	Channel	Fixed	Delay: Right delay time	0.010..1.000	seconds
88	Channel	Fixed	Delay: Feedback	0.0..0.9	
When channel effect 2 is set to ***chorus*** *(value 2)*					
86	Channel	Fixed	Chorus: Base delay	0.020..0.030	seconds
87	Channel	Fixed	Chorus: LFO speed	0.0..6.0	Hz
88	Channel	Fixed	Chorus: LFO amount	0.0..0.006	seconds
89	Channel	Fixed	Chorus: LFO oscillator type	See 5.3	
When channel effect 2 is set to ***flanger*** *(value 3)*					
86	Channel	Fixed	Flanger: Base delay	0.000..0.005	seconds
87	Channel	Fixed	Flanger: LFO speed	0.0..6.0	Hz
88	Channel	Fixed	Flanger: LFO amount	0.0..0.006	seconds
89	Channel	Fixed	Flanger: LFO oscillator type	See 5.3	
90	Channel	Fixed	Flanger: Feedback	0.0..0.9	
Channel effect 2					
91	Channel	Fixed	Channel effect 2: Wet level	0.0..1.0	
LPC					
94	Channel	Sample	LPC sample select	0..127	Sample number
95	Channel	Sample	LPC relative speed	-2.0..1.97	
Reserved (used for internal pitch bend mapping)					
96	Channel	Sample	-Reserved, don’t use- (pitch bend MSB)		
97	Channel	Sample	-Reserved, don’t use- (pitch bend LSB)		
Voice setup					
102	Channel	Note-on	Oscillator 1: Coarse tune	-64..63	seminotes
103	Channel	Note-on	Oscillator 2: Coarse tune	-64..63	seminotes

Ctrl	Scope	Granularity	Description	Range	Notes
104	Channel	Note-on	Oscillator 3: Coarse tune	-64..63	seminotes
105	Channel	Note-on	Oscillator 1: Fine tune	-1.0..0.99	seminotes
106	Channel	Note-on	Oscillator 2: Fine tune	-1.0..0.99	seminotes
107	Channel	Note-on	Oscillator 3: Fine tune	-1.0..0.99	seminotes
108	Channel	Note-on	Envelope 1: Attack time	0.0..2.0	seconds, cubed
109	Channel	Note-on	Envelope 2: Attack time	0.0..8.0	seconds, cubed
110	Channel	Note-on	Envelope 3: Attack time	0.0..0.125	seconds, cubed
111	Channel	Note-on	Envelope 1: Decay time	0.0..2.0	seconds, cubed
112	Channel	Note-on	Envelope 2: Decay time	0.0..3.38	seconds, cubed
113	Channel	Note-on	Envelope 3: Decay time	0.0..0.125	seconds, cubed
114	Channel	Note-on	Envelope 1: Sustain level	0.0..1.0	cubed
115	Channel	Note-on	Envelope 2: Sustain level	0.0..1.0	cubed
116	Channel	Note-on	Envelope 3: Sustain level	0.0..1.0	cubed
117	Channel	Note-on	Envelope 1: Release time	0.0..2.0	seconds, cubed
118	Channel	Note-on	Envelope 2: Release time	0.0..10.0	seconds, cubed
119	Channel	Note-on	Envelope 3: Release time	0.0..1.0	seconds, cubed
120	Channel	Note-on	Oscillator 1: Multi spread	0.0..1.0	seminotes, cubed
121	Channel	Note-on	Oscillator 2: Multi spread	0.0..1.0	seminotes, cubed
122	Channel	Note-on	Oscillator 3: Multi spread	0.0..1.0	seminotes, cubed
123	Channel	Note-on	Oscillator 1: Multi-osc count	1..255	always odd
124	Channel	Note-on	Oscillator 2: Multi-osc count	1..255	always odd
125	Channel	Note-on	Oscillator 3: Multi-osc count	1..255	always odd
98	Channel	Note-on	Oscillator 1: Randomize phase	0..127	0=false, >0=true
99	Channel	Note-on	Oscillator 2: Randomize phase	0..127	0=false, >0=true
100	Channel	Note-on	Oscillator 3: Randomize phase	0..127	0=false, >0=true

6.2 Filter type identifiers

As all filters come in normal and waveshaping types, they also have two numerical identifiers (FN and FW, for normal and waveshaping, respectively).

FN	FW	Description
0	—	None
1	10	Lowpass
2	11	Highpass
3	12	Bandpass
4	13	Notch
5	14	Allpass
6	—	Peaking EQ
7	—	EQ low shelf
8	—	EQ high shelf
18	—	Ladder 6 dB/oct lowpass
19	—	Ladder 12 dB/oct lowpass
20	—	Ladder 18 dB/oct lowpass
21	—	Ladder 24 dB/oct lowpass
22	—	Ladder 24 dB/oct highpass
23	—	Ladder 24 dB/oct bandpass
24	—	Ladder 24 dB/oct notch

6.3 Oscillator type identifiers

All oscillators are also available as LFOs, even when they do not make all that much sense (e.g., noise is not a very useful LFO).

No.	Type	Notes
0	None	ie., silence
1	Sine	
2	Saw	bandwidth limited
3	Square	bandwidth limited
4	Triangle	bandwidth limited
5	Noise	

7 Appendix B: Schematic Synth Overview

7.1 Conventions

```
+ = add
x = multiply (modulate)
? = scale-and-add, or modulate (user selectable)
A/B/C = either A, or B, or C, or nothing (user can select one)
```

7.2 Voice Diagram

```
      ENV/LFO                                ENV/LFO/OSC1                 ENV/LFO/OSC[12]
        |                                         |                            |
        |  pitch  ------           ------  pitch  |             ------  pitch  |
 note --?--------| OSC1 |         | OSC2 |--------?-- note     | OSC3 |--------?-- note
                  ------           ------                       ------
                    |                 |                           |
                    |  osc1/osc2 lvl  |                           |
                    |        |        |                           |
                     ---> mix/mul <---           osc3 vol         |
                             |                      |             |
                              ------> + <---------- x ------------
                                      |
                                      |--------------------------------------------------
                                      |                                                  |
                                      |                            flt cutoff            |
                                      |                                 |                |
                                      ↓         note (if key follow) -- +                ↓
                         enable  ------------                           |         ---------------
                      ----------| Waveshaper |   ENV/LFO/OSC/resonance -?        | 4x upsampling |
                     |           ------------                           |         ---------------
                     |                |          ENV/LFO/OSC/resonance -?                |
                     |                |                                 |                |
                     |                |          ENV/LFO/OSC/resonance -?                ↓
                     |                ↓                                 |             --------
                     |             --------  cutoff                    / \    cutoff |        |
                     |            |        |---------------------------   -----------| Ladder |
Waveshaping enable --|            | Filter |                                         | filter |
                     |            |        |---------------------------   -----------|        |
                     |             --------  resonance                 \ / resonance  --------
                     |                |                                 |                |
                     |                |                 ENV/FLO/OSC ----?                |
                     |                |                                 |                |
                     |                ↓                              flt reso            ↓
                     |  enable  ---------------                                  -----------------
                      ---------| De-waveshaper |                                | 4x downsampling |
                                ---------------                                  -----------------
                                      |                                                  |
                         ENV1         |                                                  |
                          |           |                                                  |
  ENV/LFO/OSC/resonance --?            ------------------> select <----------------------
                          |                                   |
                          |                                   |
                           --------------------------------- x
                                                             |
      note velocity ---------------------------------------- x
                                                             |
                                                             ↓

                                                         To channel
```

```
Available oscillator types: sine, saw, square, triangle, noise, or none
Available envelopes: ENV1/ENV2/ENV3, all ASDR but different parameter ranges
Available LFOs: LFO1/LFO2, all oscillator types, trigger on note start
```

7.3 Channel Diagram

```
                              Voices (added together)
                                        |
                                        ↓
                                      -----
                                     | LPC |
                                      -----
                                        |
                                        ↓
                                 ---------------
                                | Mono effect 1 |
                                 ---------------
                                        |
                                        ↓
                                 ---------------
                                | Mono effect 2 |
                                 ---------------
                                        |
                                        ↓
                                ------------------
                               |    EQ filters    |-- channel EQ settings
                               | (12 dB/oct each) |
                                ------------------
                                        |
                                        x -- channel volume
                                        |
                                        ↓
                                     -------
                                    |  Pan  |-- panning
                                     -------
                                        |
                                        | [All signals from this point are stereo]
                                        |
                                        |---------
                                        |         |
                                        |         ↓
                                        |   ---------------
                                        |  | Stereo effect |
                                        |   ---------------
                                        ↓         |
                         wet level --- mix <------
                                        |
            ----------------------------------------------------------
           |                            |                             |
           |                            x-- reverb send               x-- delay send
           |                            |                             |
           ↓                            ↓                             ↓
     To main bus                 To reverb bus                   To delay bus
```

Available mono effects: Distortion, compression, lo-fi, extra filter, exciter, or none
Available stereo effects: Delay, chorus, flanger, or none

7.4 Global Diagram

```
          Main bus               Reverb bus                         Delay bus
              |                       |                                 |
              |                       ↓                                 |
              |                   --------                              + <-----------------
              |                  | Reverb |-- room size, damp, width    |                   |
              |                   --------                              ↓                   |
              |                       |                              -------                |
              |                       x -- wet level                | Delay |-- L/R time    |
              |                       |                              -------                |
              |                       |                                 |     feedback      |
              |                       |                                 |         |         |
              |                       |                                 |-------> x --------
              |                       |                                 |
              |                       + -----------------------------------
              |                       |
              |                       ↓
              |                  -----------
              |                 |  Highpass |-- aux cutoff frequency
              |                 |   filter  |
              |                 | 12 dB/oct |
              |                  -----------
              |                       |
              + <---------------------
              |
              |---------
              |         |
              |         ↓
              |   ---------------
              |  | Retrigger box |
              |   ---------------
              ↓         |
wet level -- mix <------
              |
              ↓
      ------------------
     |    EQ filters    |-- global EQ settings
     | (12 dB/oct each) |
      ------------------
              |
              ↓
         -----------
        |   Filter  |
        |  (20- Hz) |
        | 24 dB/oct |
         -----------
              |
              x -- compressor pre-gain
              |
              ↓
    ----------------------
   | Splitting filterbank |-- crossover frequencies
    ----------------------
     |     |     |     |
     | B1  | B2  | B3  | B4   [only band 1 will be shown, band 2-4 are identically treated]
     |     |     |     |
     |    ...   ...   ...
     |
     ↓
    -------------------
   | Band 1 compressor |-- band 1 attack time, band 1 release time, band 1 threshold, band 1 ratio
    -------------------
             |
             x -- band 1 makeup gain
             |
             + <---- output from band 2 compressor (multiplied by band 2 makeup gain)
             |
             + <---- output from band 3 compressor (multiplied by band 3 makeup gain)
             |
             + <---- output from band 4 compressor (multiplied by band 4 makeup gain)
             |
           ------
```

```
 | Clip |
  ------
     |
     ↓

Sound card
```