

ZVORK

NOXIOUS



User Manual

version 1.5.0

Table of Contents

Introduction.....	4
A brief introduction to Additive Synthesis.....	5
Harmonic sounds.....	5
Inharmonic sounds.....	5
The Noxious engine.....	6
Oscillator.....	6
Partial Modifiers.....	7
Filters.....	7
Modulations.....	8
Front Panel.....	9
Knob default positions.....	9
Performance.....	9
Pitch Bend and Modulation wheel.....	9
Keyboard Mode and Glide.....	10
Sound.....	10
Architecture.....	10
Waveform.....	10
Variation.....	10
Stereo control.....	11
Width.....	11
Pan.....	11
Quality Settings.....	11
Partial Count.....	11
Precision.....	11
Display.....	11
Modifiers.....	12
Activation.....	12
Type.....	13
Influence.....	13
Specific Parameters.....	13
Keyboard Tracking.....	13
Tuning.....	13
Keyboard Tracking.....	14
Filters.....	14
Filter mode.....	14
Keyboard Tracking.....	14
Amplitude and Output.....	14
Filter mix.....	14
Modulation.....	14
LFO.....	14
Modulation Envelopes.....	15
Modulation Matrix.....	15
Sources.....	15
Destinations.....	16

Back Panel.....	18
CV Inputs.....	18
CV Outputs.....	18
Audio Outputs.....	19
Partial Modifiers.....	20
Level.....	20
Equalizer.....	20
Harmonic Mix.....	20
Noise.....	20
Randomize.....	20
Low Pass and High Pass.....	20
Feedback.....	21
Limit.....	21
Pitch.....	21
Tune.....	21
Scale.....	21
Offset.....	22
Harmonize.....	22
Noise.....	22
Randomize.....	22
Both.....	23
Repeat.....	23
Mirror.....	23
Randomize.....	23
Duplicate.....	24
Tips and Tricks.....	25
Adjust the quality according to your bandwidth.....	25
Stealing high frequency partials with the Offset modifier.....	25
Attack partials with inharmonic wave forms.....	25
A note on CPU / DSP performance.....	25
MIDI Implementation Chart.....	26



Introduction

Noxious is a polyphonic additive synthesizer. Additive synthesis consists of adding tens, hundreds or thousands of individual sinusoidal oscillators to create complex waveforms or sounds. In theory, additive synthesis is able to recreate any sound imaginable but is usually a challenge to manipulate.

Noxious takes an intermediate route to additive synthesis by offering a limited mean of manipulating these sinusoidal oscillators, called partials, through “modifiers” and classical filters.

A brief introduction to Additive Synthesis

A majority of synthesizers are based on a particular form of synthesis called “subtractive synthesis”. It consists in producing a spectrally rich wave form, such as sawtooth, square or noise and applying a filter on these wave forms to remove, or subtract, parts of their spectrum. By varying the types of filters and their parameters, they can produce a fairly wide variety of sounds.

Additive synthesis works the other way around. It consists in building this rich spectrum from scratch by adding up individual sinusoidal oscillators, called **partials** in Noxious. Sine waves have the particularity of producing sound in a single frequency band of the spectrum. By adding up multiple sine waves of different frequency, we can, in theory, reproduce any frequency spectrum. In fact, science demonstrates that any sound of any length can be created with an infinite number of sine waves, properly tuned and equalized. This of course isn't practical and in reality additive synthesis constrains you to a finite number of sinusoidal oscillators but with controllable pitch and level. Typical sounds consists of hundreds of sine waves so controlling each one of them is a bit of a challenge. We will see how Noxious handles the problem later on but in the mean time, let us see how typical sounds are built from these individual partials.

Harmonic sounds

Musical harmonic sounds, such as created by acoustic instruments or analog subtractive synthesizers, are produced by periodic wave forms. The frequency of these wave forms determine the pitch of the note. Typical periodic wave forms in subtractive synthesis include sawtooth, triangle and square waves. They are all made up of sine waves of frequencies multiple of the wave form's note frequency. This note frequency is called **fundamental frequency** in additive synthesis. For example, if you hit the A4 key on your keyboard, which corresponds to a frequency of 440 Hz, and produce a saw tooth wave form, this saw tooth is in fact made up of sine waves of frequency 440 Hz, 880 Hz, 1320 Hz or 1760 Hz (or $1 \times 440\text{Hz}$, $2 \times 440\text{Hz}$, $3 \times 440\text{Hz}$ and $4 \times 440\text{Hz}$). The multiples of the fundamental frequency are called **harmonics**. The sine waves also have a particular amplitude but these amplitudes only change the timbre of the wave, not its note frequency or harmonicity. This difference in amplitude is what makes a sawtooth sound different from a square wave.

Inharmonic sounds

The vast majority of sounds are not harmonic or tuned. Percussion or noise sounds for instance don't have a clear pitch and thus don't have partials tuned so simply as harmonic sounds. In fact most inharmonic sounds have partials tuned with no relation between one another. In reality, acoustic instruments are not totally harmonic in the pure sense of the term but are almost harmonic. Their partials are not strictly tuned to multiples of the fundamental frequency and this is what gives their richness as this introduces some slight chorus / beating effects. Typical inharmonic sounds are metallic sounds, such as mistuned bells or rigid strings, and these can be easily built from harmonic sounds by detuning some of their partials.

The Noxious engine

In this chapter we will describe in detail how Noxious lets you manipulate these partials to create a wide variety of sounds and how it relates to additive synthesis and even to subtractive synthesis.

Oscillator

Instead of creating partials from scratch, Noxious will allow you to start with a basic wave form from a predetermined list. These wave forms are not purely periodic, which is what you find in most synthesizers. Noxious can produce inharmonic wave forms. Being an additive synthesizer it can generate up to 200 partials (in highest quality) per wave form. Here is the list of available wave forms:

- **Saw tooth** : classical saw tooth with a rich full spectrum.
- **Saw tooth + Chorus** : a “fat” saw tooth consisting of triplets or duos of slightly detuned partials. The spectrum has a lower bandwidth than the normal saw tooth.
- **Square** : classical square wave with a rich albeit “empty” spectrum. Square waves have the particularity of lacking all harmonics with pitches of even multiples of the fundamental frequency. In fact, all wave forms lacking these even harmonics have the same “empty” quality to them.
- **Square + Chorus** : the same square wave with a narrower bandwidth but with slightly detuned partials for a “fatter” sound.
- **Pulse 1/10** : a classic pulse wave with a 1/10 pulse width.
- **Pulse 2/10** : same thing but with a 2/10 pulse width.
- **Pulse 3/10** : again but with a 3/10 pulse width.
- **Dirac** : this is an interesting wave form as it has all harmonics, all with equal level.
- **Triangle** : the third classical wave form with the same “empty” quality to it as the square wave but without its brightness.
- **Male Voice** : a slightly synthetic male choir.
- **Female Voice** : a slightly synthetic female “oh” sound.
- **Strings** : an orchestral string wave form.
- **Saxophone** : a tenor saxophone.
- **Nasal** : a very nasal sounding wave form with lots of medium frequencies.
- **Cymbal** : an inharmonic metallic cymbal sound.
- **Dark noise** : an inharmonic low bandwidth noise.
- **Razor** : imagine an electric razor and you get a good idea of what this sounds like.
- **Quasi Sawtooth** : like a sawtooth but with all the partials slightly detuned.
- **Fractal** : a synthetic wave with a lot of high frequencies

- **Shepard** : this wave has only octave separated partials, perfect for organ or clear sounds.
- **Alien** : a spectrum halfway between noise and metal with a slight pulsation.

Partial Modifiers

The partials produced by the oscillator can then be transformed by a stack of up to five modifiers A, B, C, D and E. The modifiers are applied in sequence and can modify the levels of the partials, their pitch or both. Each modifier can operate on a subset of partials. Here is the list of available modifiers:

- **Level – Equalizer** : this is the basic modifier as it lets you boost or cut partials as a classic parametric equalizer but with the added possibility of modifying only even or odd harmonics.
- **Level – Harmonic Mix** : remove harmonic or inharmonic partials.
- **Level – Noise** : add a time varying randomness to the existing partial levels.
- **Level – Low Pass 12dB** : apply a low pass filter with a 12 dB/octave slope to the selected partials.
- **Level – High Pass 12dB** : apply a high pass filter with a 12 dB/octave slope to the selected partials.
- **Level – Feedback** : apply a feedback comb filter to the selected partials.
- **Level – Limiter** : limit the level of the partials to a certain level by applying a soft clip to their level.
- **Pitch – Tune** : tune the partials to a different fundamental frequency.
- **Pitch – Scale** : scale the partial pitch away or towards a center frequency.
- **Pitch – Harmonize** : detune the partial pitch away or towards the nearest harmonic frequency. Very handy to transform an inharmonic wave form into an harmonic one.
- **Pitch – Noise** : add a time varying randomness to the existing partial pitch.
- **Combined – Repeat** : Repeat the low frequency partials across the spectrum.
- **Combined – Mirror** : Mirror Repeat the low frequency partials across the spectrum.
- **Combined – Randomize** : Shuffle both levels and pitches.
- **Combined – Duplicate** : Create detuned copies of the partials.

A more thorough description and documentation on the partial modifiers is available in a dedicated chapter.

Filters

After the partial modifiers, the wave form can be modified furthermore with two multi-state filters in series featuring resonance. This is classical territory if you know subtractive synthesis but the filters operate directly on the partials without touching the phase. Here is the list of filter modes:

- **Low Pass Type I** : cuts the high frequencies with available slopes of 6 dB, 12 dB, 24 dB and 48 dB per octave and with a high constant resonance value.
- **Low Pass Type II** : cuts the high frequencies with available slopes of 6 dB, 12 dB, 24 dB and 48 dB per octave but with a frequency variable constant resonance value.
- **High Pass**: cuts the low frequencies with the same available slopes as the low pass mode, also with a high constant resonance value.
- **Band Pass** : cuts frequencies that are not around a center frequency, with the same available slopes.
- **Band Reject** : cuts frequencies that are around a center frequency, with the same available slopes.
- **Negative and positive feedback comb filter** : reproduces the delay based effect.
- **Vowel** : emphasizes various spectrum frequencies (called formants in speech synthesis) so as to reproduce human vowels. The frequency parameter of the filter is used to determine the pitch or gender of the voice whereas the resonance lets you sweep through the various vowels.

Modulations

All the above filters and modifiers can be furthermore modulated with two LFOs, two modulation envelopes, amplitude envelope and performance controllers such as velocity, aftertouch or modulation wheel. A modulation matrix lets you build your own modulation routes with adjustable modulation strength.

Front Panel



Knob default positions

On Noxious, every knob corresponding to an important parameter has a small LED light underneath. When this light is on, it indicates that the knob is set to a non default value. This lets you quickly see what has changed from the default on your panel.

Performance

Pitch Bend and Modulation wheel

The range of the pitch bend can be adjusted with a maximum range of 2 octaves (24 semitones).

Keyboard Mode and Glide

There are three keyboard modes:

- *Polyphonic* : the default mode where each new pressed key creates a new voice. The maximum number of simultaneous notes can be adjusted with the **Polyphony** setting. Noxious has a maximum polyphony of 16 voices.
- *Mono-Retrigger* : when pressing a key, the synth retriggers the envelopes and changes the pitch if a previous key is pressed.
- *Mono-Legato* : this is the same as *mono-retrigger* but the envelopes are not retriggered.

By changing the *Glide* duration, you can add portamento effect that will make the pitch glide from the previous note to the next. This also works in *polyphonic* mode but in this case the portamento will take place only if you press a new key while previous notes are in release mode.

Sound

Architecture


In Noxious, you don't directly create partials. You start from a set of partials corresponding to a wave form which you then modify. These modifications are done through two "classic" filters and up to five "modifiers". All these are applied in sequence to the partials. This architecture is summarized by a schematic on the back panel, reproduced on the right of this page.

Waveform

The wave form is selected by navigating through the list of wave forms with the left or right buttons or by directly selecting it in a list by clicking on the current wave form label.

Variation

Each waveform has a number of variations (for example pulse width for the sawtooth waveform) that can be selected in a continuous manner. Sweeping smoothly between each variation is a first method of changing your sound over time.



Stereo control

Width

Each partial of a waveform can furthermore be spread of the stereo space. This is controlled through the waveform's "width" parameter. A zero width gives a very narrow sound (the default) whereas set to 100, partials are balanced across the entire stereo space.

Pan

Like the name suggests, this is the parameter used to offset the center of the waveform in the stereo space.

Quality Settings

Additive synthesis can be very CPU hungry on low end computers, depending on the number of partials manipulated and the rate at which they are modified. These settings let you strike a balance between quality and CPU performance or just push the quality to its maximum.

Partial Count

By setting the partial count to one of the three possible settings – Low, Medium or High – you can lessen the CPU load of your patches. Each quality setting corresponds to a certain number of partials per voice, which affects the spectrum of the waveforms.


Precision

On most occasions you won't notice a difference when changing the precision but when you are changing your spectrum at a very fast rate (with a high rate LFO for instance or a high rate noise modifier) you might want to try a higher setting. If, on the contrary, you are creating a very slowly evolving pad, you can switch to the low setting without noticing any difference. Changing the spectrum includes modifiers and filters so very fast modulation of filter cutoff can also be of higher quality if you push this setting to its highest value.

Display

When all partial modifiers are turned off (the default situation), the partial display lets you visualize the levels (in dB) and pitches of the selected wave form's partials. Due to a limited display size only the partials with pitches up to hundred times the fundamental frequency will fit in the display. Nonetheless, especially in high quality extra partials can exist after the far right of the visible spectrum.


If all the partials are harmonics (for instance with a saw tooth or square wave), the pitch display will be empty. If this is not the case, inharmonic partials will display a bar with an amplitude indicating its distance to the nearest harmonic in cents. If multiple partials overlap on the display (with very similar pitches), the partial with the greatest deviation from the harmonic will be used to display this pitch.



Sawtooth wave

The saw tooth as seen in the display. You can see the levels of each partials decrease as we go up in frequency.


The pitch display is a flat line which indicates that all partials are harmonics.



Square wave

The square wave has the same decrease in level but you can see that one out of two partials are missing.


Again, the pitch display is a flat line which indicates that all partials are harmonics.



Sawtooth with chorus

The levels of the partials are identical to the sawtooth wave so the level display has an identical aspect.

The pitch display, however displays some small deviations in pitch which corresponds to the slight detuning of each partials to produce the chorus effect.



Cymbal

Finally the cymbal waveform, which is absolutely inharmonic, has an even distribution of frequencies across the spectrum all with approximately equal level.

Their pitch also seems to be a bit random as we can see that there is a lot of deviation from the harmonics.

When modifiers are active, the display will reflect the state of the partials after all active modifiers are applied, just before being sent to the filters.

Modifiers

There are up to five possible active modifiers. Each modifier is applied in sequence, starting with modifier A and ending with modifier E. A modifier thus is applied to the output of the previous active modifier.

Activation


A modifier is active if it is enabled by pressing the button at the top left of the corresponding modifier section.

Type

The type of the modifier can be selected by navigating through the list with the left or right button, or by pressing on the label which displays the list of all available modifier types. For a complete list and description of these modifier types, please refer to the specific chapter at the end of the documentation.

Influence

A modifier is applied to a subset of partials. This selection (or partial influence) is controlled through the parameters *Center*, *Spread* and *Falloff*. The center parameter controls the center frequency of the modifier, the spread controls the frequency range of its influence and finally the falloff determines how this influence will decrease to zero between the center and the edge of the selection. The illustrations below present two modifiers with decreasing falloff.



Specific Parameters

A modifier has two specific parameter knobs which control the type specific parameter 1 & 2. A display at the right of the knob presents the name of the parameter for the current modifier type.

Keyboard Tracking

When keyboard tracking is set to 100, the modifier center and spread follow exactly the note frequency. This corresponds to the modifier being applied the same on the wave form partials regardless of the note pitch, which is usually what we want. When the tracking is set to 0, the modifier center and spread are constant across the note range and fixed in the spectrum. For instance, on an equalizer, this corresponds to a fixed peak or notch on the spectrum just like a normal equalizer or formant filter would behave. For values of the keyboard tracking under 0, the center and spread increase when the note frequency decreases.

Tuning

There are two knobs to define the basic pitch of the oscillator. The main one is used to adjust the pitch in semi tones, from -48 to +48. The second one lets you finely tune the oscillator in cents or hundredth of semi tones.

Keyboard Tracking

When set to 0 (the default), the pitch of the oscillator follows the note pressed on the keyboards. Set to 100, the pitch changes at twice the normal note rate : two keys separated by an octave will give a two octave oscillator pitch difference. Finally, when set to -100, the pitch is constant across all keys.

Filters

Filter mode

The filters can be applied to the modified waveform in serial (the default) or in parallel. In the latter case, each filtered version (or the original modified waveform, if a filter is not activated) is then remixed to form the final output.

Keyboard Tracking

This controls how the filter frequency is adjusted for each note. By default, the value is 50 which corresponds to a frequency adjusted exactly with the note frequency. When the value is 0, the frequency is identical along all the note range. Negative values indicate that the frequency increases with low notes.

Amplitude and Output

This section of the front panel groups all the controls associated with the volume. The amplitude envelope is a classic ADSR envelope with attack time, decay time, sustain level and release time controls.

Filter mix

The influence of each filter in the final output can be controlled here. To the left, only the influence of the first filter is heard, if it is on, and the right, only the influence of the second filter.

Modulation

LFO


Two low frequency oscillators (LFO) can be used to modulate one or many parameters with a bipolar curve. All shapes oscillate between -1 and +1. Both LFO's have the same settings:

- Frequency: controls the rate of the LFO in hertz or in note lengths, if it is synced to the tempo,
- Shape: there are six different available wave shapes:
 - Sine
 - Triangle
 - Saw

- Square
- Random
- Sampled random
- Tempo Sync : when pressed, the frequency of the LFO is adjustable in note lengths, with the start of a wave cycle synced to the tempo.
- Key Sync : when pressed (the default), the wave cycle starts when a key is pressed. If the button is unpressed, the LFO is free running and the same for all notes.
- Delay : delay the effect of the LFO up to five seconds.
- Destination : by clicking on this display you can select what the LFO will modulate from all the available modulation destinations (see the Matrix paragraph for the list). The knob to the left controls the strength of the modulation from -100 to +100.

Modulation Envelopes

With the amplitude envelope come two freely assignable modulation envelopes. They are a bit more configurable than classic ADSR envelopes as you can also adjust the start, attack and release levels. All levels (L1, L2, Sustain and L4) have values ranging between -100 and +100.



Just like the LFO's, each modulation envelope has a destination freely selectable from all the modulation destinations, with an adjustable modulation strength.

Modulation Matrix

The modulation matrix is where you can create extra modulations by routing modulation sources to modulation destinations. Each route has an adjustable strength parameter to define the influence of each source on the corresponding destination.

Sources

Here is the list of available modulation sources:

- Key pressed velocity

- Amplitude envelope
- Modulation envelopes 1 & 2
- LFO 1 & 2
- Aftertouch
- Modulation wheel
- Keyboard tracking
- Sustain pedal
- Expression controller
- Breath control
- External CV inputs 1 to 4 (see back panel)
- Random on note 1 & 2 : this delivers a completely random bipolar value on each note pressed.


Destinations

Here is the list of available modulation destinations:

- LFO 1 & 2 frequency
- LFO 1 & 2 level : this will modify the strength of the modulation set in the corresponding LFO section. It will not modify the strength of other modulations as set in the modulation matrix and having the LFO as the source.
- LFO 1 & 2 delay
- Modulation envelope 1 & 2 level : this will modify the strength of the modulation set in the corresponding Modulation Envelope section. It will not modify the strength of other modulations as set in the modulation matrix and having the Modulation Envelope as the source.
- Modulation envelope 1 & 2 attack, decay and release times
- Filter 1 & 2 cutoff frequency
- Filter 1 & 2 resonance
- Filter mix
- Waveform variation
- Oscillator pitch
- Oscillator output level
- Oscillator pan
- Oscillator width
- Amplitude envelope attack, decay and release times
- Modifiers A-E center

- Modifiers A-E spread
- Modifiers A-E falloff
- Modifiers A-E parameter 1
- Modifiers A-E parameter 2

Back Panel



The back panel exposes the input / output CV and audio output sockets.

CV Inputs

Some of Noxious' parameters can have their values modulated by control voltage signals coming from other devices. Trim knobs let you control the amplitude of the modulation. Four special CV inputs (labeled CV 1, 2, 3, 4) are by default not routed to a specific parameter. It is up to you, through the modulation matrix, to route them to a particular modulation destination.

CV Outputs

Noxious exposes for internal parameter to CV outputs:

- LFO 1 & 2 : this the value of the LFO if it were not key synced.

- Modulation envelope 1 & 2 : the value of the modulation envelope of the first pressed voice is sent to this output.

Audio Outputs

The panned stereo output is sent to the left and right audio socket. If only the left socket is connected to another device, Noxious sends the mixed mono signal to the left socket.

Partial Modifiers

Level

Equalizer

This is the most basic of all modifiers and resembles the classic peak / notch equalizer in that it enables you to increase or decrease the levels of the influenced partials. Parameter 1 (Gain) controls the gain. Parameter 2 (Parity) enables you to modify all influenced partials, only even harmonics if set to +100 or only odd harmonics if set to -100.

Harmonic Mix

This modifier lets you control the mix between harmonic and inharmonic partials. It has two parameters: *Mix* and *Narrow*. *Mix* controls which of the two sets of partials to keep. On zero all partials are kept. On +100, inharmonic partials are attenuated while on -100, the harmonic ones are attenuated. The discrimination of partials between harmonic and inharmonic is controlled with the narrow parameter. At +100, the harmonic and inharmonic selection (depending on the value of the mix parameter) is extremely narrow and only partials which have frequencies almost exactly harmonic are considered harmonic. When the narrow parameter is a -100, the difference between harmonic and inharmonic is less constrained.

Noise

This modifier randomly changes the level of the influenced partials over time. The range of the random gain is controlled with the *Amount* parameter while the rate at which this evolves is set with the *Rate* parameter. The randomness is voice dependent so if you strike a chord, each note will have a different spectrum. Be careful as this modifier can be CPU intensive.

Randomize

The effect does not vary over time. The *Amount* parameter controls the amount of shuffling of the partials levels while the second specific parameter *Seed*, determines the noise distribution. Changing this will completely change the random values in unpredictable ways.

Low Pass and High Pass

These modifiers act like the respective filter types but their effect are weighted by spread, center and falloff. You can for instance apply low pass effect on just a part of the spectrum by correctly setting the center on the desired zone. The first parameter controls the filter cutoff frequency and the second resonance. But be aware that, like all modifier parameters, their value isn't influenced by the Keyboard Tracking value of the modifier. To do so, you explicitly create a link in the modulation matrix.

Feedback

This modifier applies a positive or negative comb feedback effect to any section of your spectrum.


Limit

To prevent certain partial levels from becoming too high (especially when using the Equalizer modifier with a high gain value), you can apply this modifier. The first parameter determines the threshold maximum value of the partial levels. The effect is softer than a hard clip.

Pitch


Tune

The Tune modifier lets you adjust the pitch of the influenced partials by a particular relative amount. The two parameters *Semi* and *Fine* changes the pitch by semitones or cents.




Scale

This modifier changes the pitch of the partials relatively to the modifier center. The center defines a central frequency and the *Amount* and *Fine* control how close the partial pitch is moved towards or away from this central frequency. If the amount is set to +100, the partial pitch is exactly the central frequency. If it is +50, it will be halfway between its original pitch and the central frequency. Finally, at -100, the partial pitch will be at double the original frequency distance from the center frequency.



Offset

Whereas the Tune modifier changes the pitch in a relative manner, the Offset modifier adds or subtracts the same frequency value to all the influenced partials. This can have the effect of changing the harmonic rank of the partials, for instance. The *Amount* and *Fine* parameters control the frequency value added or removed from the partial pitch.



Harmonize

The Harmonize modifier is extremely useful to transform a non harmonic wave form (such as Cymbal or Dark Noise) into an harmonic one (or vice versa with negative parameter values). With the help of the *Amount* parameter it moves the partial pitch towards the closest harmonic frequency when this parameter is positive or away from the closest harmonic when this parameter is negative.

Noise

This modifier randomly changes the pitch of the influenced partials over time. The range of the random tuning is controlled with the *Level* parameter while the rate at which this evolves is set with the *Rate* parameter. The randomness is voice dependent so if you strike a chord, each note will have a different spectrum. Be careful as this modifier can be CPU intensive.


Randomize

The effect does not vary over time. The *Amount* parameter controls the amount of shuffling of the partials pitches while the second specific parameter *Seed*, determines the noise distribution. Changing this will completely change the random values in unpredictable ways.

Both


Repeat

For this modifier the selection parameters behave somewhat differently. First of all, only the Center and Falloff parameters are used. The Spread parameter has no effect. The center frequency sets the upper limit of the partial spectrum that will be repeated. All the partials above this limit will be replaced by copies of the low spectrum partials and offset in frequency by multiples of the center frequency. The *Falloff* controls the progressive attenuation of these copies. The first specific parameter *Harmonic* forces the center frequency to the closest harmonic frequency when set to +100. This has the effect of ensuring that the partial copies remain harmonic if the originals were. At value -100 the center frequency is unconstrained. The second specific parameter *Mix* controls the dry / wet mix of the effect.



Mirror

The Mirror modifier has a very similar effect to the Repeat modifier but instead of copying the partials with offset, it also flips the partial frequencies every odd numbered copy. Odd numbered copies thus have their lowest frequency partials equal to the highest frequency partials of the original partials. However complicated this may seem, this mimics the effect of sampling the wave form with a sampling rate equal to the center frequency.



Randomize

The Randomize modifier is a fast way to create brand new wave forms by adding noise to both level and pitch of the partials. The effect does not vary over time. The *Amount* parameter controls the amount of shuffling both in level and pitch while the second specific parameter *Seed*, determines the noise distribution. Changing this will completely change the random values in unpredictable ways.

Duplicate

This modifier steals the partials of lowest amplitudes and reuses them to add one new partial for each partials in the zone of influence. This “new” partial can be detuned with the *Interval* and *Depth* parameters. The first parameter sets the interval in semi-tones between the new partial and the original one. The second parameter does the same but in cents. These new partials are in fact existing partials that are retuned and set to the correct level based on their level and frequency. The modifier will try to reuse partials with low levels and high frequency. Because of this, rapid changes of spectrum in drastic ways using modifiers may result in audible disappearance and reappearance of partials if a Duplicate modifier is applied in the final stages of the modifier stack.

Tips and Tricks

If you want to start making your presets for Noxious, there are a few tips, tricks and starting points you might want to know.

Adjust the quality according to your bandwidth

The higher the quality the heavier will playing Noxious tax your CPU power. But increasing partial count adds high frequencies that might not be audible, especially if you are using a low pass filter further down the line. So as a final touch, try adjusting to the lowest partial count.

The same thing applies to the precision. Try changing to the lowest quality to hear if there are any differences. It will most noticeable if there are very fast variations of you spectrum, because of high LFO or envelope rates, for instance.

Stealing high frequency partials with the Offset modifier

When your sound doesn't really need those high frequencies (because you're a huge fan of dark resonant muted sounds) you can still use them for other purposes. With the *Offset* modifier you can move your partials down to lower frequencies and use them for good effect, like fattening your sounds if you have previously slightly detuned them with the *Tune* modifier or add a touch of inharmonicity during the attack phase.

Attack partials with inharmonic wave forms

Starting with an inharmonic waveform is sometimes a great idea if you want to create nice percussive sounds with a musical tone. If you apply a *Harmonize* modifier on an inharmonic wave form and pull the intensity parameter to 100, you will hear the sound transforming into a completely harmonic wave. Likewise, if you apply a *Harmonic Mix* modifier instead and put the mix parameter to 100, playing with the narrow parameter will enable you to progressively remove the inharmonic partials.

With that in mind, it is easy to link one of the envelopes to the intensity or mix parameter, depending on which modifier you are using, to make a sound that starts with a rich inharmonic spectrum and gradually changes to a harmonic one.

A note on CPU / DSP performance

The CPU load of Noxious is directly proportional to the number of audible partials. Very high frequency partials are not computed to save CPU time. Some wave forms, the offset and tune modifiers, can lead to a great number of partials in the lower part of the spectrum (The Strings, Dark Noise and Razor wave forms are the perfect examples). This implies that most of the partials will always be audible and thus need more CPU power to render. There is always a balance between spectrum richness in the lower end and CPU performance. The most CPU demanding patch would be all the partials tuned or offset in the lower part of the spectrum, in high partial count and high precision quality settings.

MIDI Implementation Chart

MIDI Controller #	Parameter
4	Pan
5	Pitch - Semi
7	Pitch - Fine
8	Amplitude envelope Attack time
10	Amplitude envelope Decay time
12	Amplitude envelope Sustain level
13	Amplitude envelope Release time
14	Oscillator Waveform
15-19	Modifier A-E On / Off
20-24	Modifier A-E Center
25-29	Modifier A-E Spread
30-31, 33-35	Modifier A-E Falloff
36-37, 39-41	Modifier A-E Type
42-46	Modifier A-E Parameter 1
47-51	Modifier A-E Parameter 2
52, 53	Mod. Envelope 1 / 2 Start level
54, 55	Mod. Envelope 1 / 2 Attack time
56, 57	Mod. Envelope 1 / 2 Attack level
58, 59	Mod. Envelope 1 / 2 Decay time
60, 61	Mod. Envelope 1 / 2 Sustain level
62, 63	Mod. Envelope 1 / 2 Release time
65, 66	Mod. Envelope 1 / 2 Release level
67, 68	Filter 1 / 2 Type
69, 70	Filter 1 / 2 Frequency
71, 72	Filter 1 / 2 Resonance
73, 74	Filter 1 / 2 On / Off
75, 76	LFO 1 / 2 Key sync.
77, 78	LFO 1 / 2 Tempo sync.
79, 80	LFO 1 / 2 Frequency
81, 82	LFO 1 / 2 Frequency tempo synced
83, 84	LFO 1 / 2 Shape

85, 86	LFO 1 / 2 Delay
87	Volume
88	Keyboard mode
89	Glide
90	Oscillator stereo width
91	Waveform variation
92	Filter mode
93	Filter mix

More information at <http://www.zvork.fr/audio>

© 2013 Olivier Prat