

SpectrumWorx Manual

Surge Synth Team

Contents

Getting Started	4
Welcome!	4
What is SpectrumWorx?	5
What is “Spectral” Processing?	5
Installing SpectrumWorx	6
Windows	6
macOS	6
Linux	6
Building from source	6
How to use SpectrumWorx	6
The Interface	7
1 — The Main window	7
2 — The Module Bank	8
3 and 4 — Settings and Presets	8
About that CPU	8
Caution: watch those levels!	8
The Main Window	9
In, out, mix	9
Module name, parameter name and value	10
Module gain, wetness and frequency range	10
External audio	10
LFOs	11
Period	11
Waveform	12
Phase	12
Range	13
The Module Bank	14
Loading modules	14
Muting, deleting and rearranging	14
Settings: Engine	16
A brief diversion into how SpectrumWorx does its thing	17

Frame size	17
Overlap factor	18
Window type	18
Ripple factor	19
What is the best engine setup?	19
Input mode	20
Settings: GUI and About	21
Mouse over reaction	21
LFO update behavior	21
Load last session on startup	21
Hide cursor on knob drag	21
Set window and menu opacity	22
Settings: About	22
Presets	23
Location	23
Save, Save as, Delete	23
Ignore external audio	24
List	24
Description	24
The Modules	25
Modules: Pitch	26
Pitch Shifter	26
Pitch Follower	26
TuneWorx	26
Pitch Magnet	26
Sumo Pitch	26
Pitch Spring	26
Octaver	27
Modules: Timbre	28
Bandpass	28
Bandstop	28
Ah-ah	28
Smoother	28
Sharper	28
Centroid	28
Tonal	29
Atonal	29
Modules: Time	30
Freeze	30
Slicer	30
Wobbler	31
Reverser	31
Imploder	31

Exploder	31
Modules: Space	32
Frecho	32
Frevcho	32
Freqverb	32
Modules: Phase	33
Robotizer	33
Whisperer	33
Phasevolution	33
Phlip	33
Modules: Loudness	34
Gain	34
Exaggerator	34
Denoiser	34
Quiet Boost	34
Freqnamics	34
Modules: Combine	35
Talking Wind	35
Convolver	35
Ethereal	35
Vaxateer	35
Shapeless	36
Colorifer	36
Merger	36
Blender	36
Insertor	36
Burrito	37
Modules: Phase Vocoder	38
To PV and From PV	38
The variants	38
Modules: Miscellaneous	39
Armonizer	39
Slew Limiter	39
Shifter	39
Swappah	39
Quantizer	39
Credits	41
The original team	41
Special thanks	42
Post scriptum	42

Getting Started

SpectrumWorx is a modular spectral processor. It decomposes your audio into the frequency domain, lets you stack and chain effects on that spectrum, then resynthesizes the result.

It was originally developed and released by Little Endian, where development ended in 2016. The plugin was open sourced in 2024, and the Surge Synth Team picked it up in 2026 to modernize it. It is now available as CLAP, VST3, AU, and a standalone application on Windows, macOS and Linux.

To ask questions, report bugs or issues, or help develop SpectrumWorx, visit the Surge Synth Team Discord server.

A PDF version of this manual is also available for download.



Figure 1: The SpectrumWorx main window and module bank

Welcome!

Thank you for installing SpectrumWorx, one of the most advanced audio processors currently available! We are sure you will find SpectrumWorx to be an invaluable addition to your sonic toolkit. Though it shares a name and a philosophy with the original SpectrumWorx 1 plug-in from which it was derived, this version 2 of SpectrumWorx has been rebuilt from the ground up to better suit (and take advantage of) today's advanced technologies.

What is SpectrumWorx?

SpectrumWorx is a modular spectral processor. You may be familiar with modular synthesizers, wherein the user is allowed to select the types of components to be installed and how they should be routed. SpectrumWorx works in a similar fashion, except that it acts as an effects device. SpectrumWorx is not an instrument. Using an attractive, intuitive interface, it allows you to decide which processing modules should be installed at a given time and how they should be arranged. Each module is a sophisticated audio processor in and of itself, with its own, unique selection of adjustable parameters. Most of these can be modulated automatically using independent SpectrumWorx LFOs. In addition, certain modules can be side-channel driven, using either the built-in SpectrumWorx external sound file support or by utilizing side-channel routing in hosts that support it.

There are over fifty individual modules to choose from and as many as five may be loaded into SpectrumWorx at any time. You can arrange and rearrange them as you like without interrupting the signal being processed, making it ideal for processing on the fly.

What is “Spectral” Processing?

We’ll keep this bit basic, no need to break out your calculators and slide rules. After all, you probably want to get on with playing with SpectrumWorx — you didn’t come here for a science lesson!

You are almost certainly familiar with the standard sorts of effects processing used by professional recording engineers on nearly every commercial recording. Your DAW (Digital Audio Workstation) probably came with a bundle of the usual suspects: reverb, delay, phasers, compressors, equalizers and so on. These “bread and butter” effects originally derived from hardware devices and are therefore somewhat limited in scope. Essential, yes, but hardly the stuff of revolution. The effects included with SpectrumWorx, however, manipulate signals in ways that were simply not possible before advances in modern computing technology provided developers with the processing power required to dig deep into a digitized audio signal.

The effects in SpectrumWorx draw upon complex algorithms designed to analyze a signal and break it down into its many frequency components. Once that has been done, we can manipulate and rearrange those components at will. In this way, we can get at the very heart of the sound, manipulating it bit-by-bit. We can adjust the pitch in unusual ways, or affect the timing in interesting ways... or both. The entire frequency spectrum is at our disposal. If you know something about synthesis techniques, then you are probably nodding your head and saying to yourself “that sounds a bit like additive (or granular, or frame, or FFT) synthesis” and you’d not be far off the mark.

In days past, such analytical audio acrobatics were performed strictly offline. The engineer was forced to sit back and stare at an hourglass on the screen while the computer chewed on the numbers. Today, these things happen virtually instantly and such effects may be used in real-time, just as a guitarist might use a distortion box, tweaking and adjusting even as the notes ring out. SpectrumWorx does its magic in real-time and, unlike some processors, won’t interrupt the signal (or the inspiration!) with any annoying clicks or dropouts.

Installing SpectrumWorx

Installing SpectrumWorx is as easy as downloading the build from the product page and running the installer, which will guide you through the process.

Windows

The installer offers the CLAP and VST3 formats and the standalone application.

macOS

The installer offers the CLAP, VST3 and AU (v2) formats and the standalone application.

Linux

A Debian package offers the CLAP and VST3 formats and the standalone application. The individual components are also available packaged in a ZIP file.

Building from source

Steps for building SpectrumWorx yourself are in the README of the GitHub repository.

How to use SpectrumWorx

SpectrumWorx is an effect. It doesn't make any sound on its own (unless you are using a module that loads and plays a sample), and it should be opened as an effect in your signal path. Different hosts provide different methods for using effects, and you'll have to consult your host's manual to figure out how this is done. A standalone application is also provided if you want to run SpectrumWorx without a host.

SpectrumWorx can be applied to virtual instruments, recorded audio or even a live audio input. Some latency is inevitable, but many hosts will compensate for any delay introduced into the signal (again, you'll need to consult your manual to find out whether this is automatic or if you have to do it manually). The amount of delay depends primarily on the SpectrumWorx engine setup, while certain modules that are inherently delay-based (like the Reverser) can add to the amount of perceived audio delay.

The Interface

A lot of musicians tremble with trepidation when they hear the term “modular”. It brings to mind towering monolithic synthesizers festooned with hundreds of knobs, patch points and obscured by dozens of dangling cables. Fortunately, the SpectrumWorx GUI is a study in user-friendliness. Everything you need at any given moment is right in front of you, and there are only a few (and rarely needed) bits tucked away where they won’t interfere with your work until you need them.

Let’s take a look at the various components that make up the SpectrumWorx user interface, using the numbered screenshot guide as a reference.



Figure 2: The four SpectrumWorx windows, numbered

1 — The Main window

The Main plug-in window is labeled as 1. Some of the items in the Main plug-in window are always present, including the In, Out, and Mix knobs, Presets and Settings launch buttons and the area used to load external audio (Load and Unload buttons). Other, variable fields will change, depending on which module is selected, and even which parameter on the selected module is selected. The name of the currently selected module appears at the top, while the currently selected parameter and its value will be displayed just below the module name.

Below the currently selected module and its parameter value you will find the three controls related to the module which is currently selected, and those are: Gain knob, Wet knob, and Frequency range slider. This means that each module has separate gain, wetness and range.

Below the shared module parameters you will find the LFO section. The LFO settings are unique to each selected parameter, and the various fields in the LFO section will reflect the values specific to that selected parameter. We’ll talk more about LFOs, what they are, and how you can use them a little later.

2 — The Module Bank

You will find the Module Bank to the right of the Main plug-in window labeled as **2**. This is the area into which the many different SpectrumWorx effects modules may be inserted. In our numbered screen grab we have four modules loaded into the Module Bank (a Slew Limiter, Frevcho, TuneWorx and a Slicer). As you can see, our Slew Limiter for example has a total of two controls, one combo box and one knob, each controlling a different parameter. Each of the controls can, as we suggested above, be modulated by its own LFO.

3 and 4 — Settings and Presets

We've labeled the Settings window as number **3** and Presets as number **4**. These windows are left hidden until you need them.

You will access the Settings window when you want to adjust the engine's various analysis parameters, alter the GUI, or read the "About" page to find out the software version and who is behind SpectrumWorx.

The Presets child window you can open when you need to load or save presets.

About that CPU

SpectrumWorx uses complex algorithms designed to take advantage of today's powerful desktop environments. As such, it can consume a lot of resources. However, the SpectrumWorx engine, with its emphasis on quality as well as stability, is still an efficient performer. The processing required will vary dramatically, depending on which modules are installed in the Module Bank (most modules consume only up to 1%) and upon your Engine settings. For example, adjusting the Overlap, Frame length and/or the number of channels can significantly affect the hit on your computer. We'll discuss those settings momentarily.

Caution: watch those levels!

As is the case with any modular system, routing and processing signals through SpectrumWorx can sometimes provide unintended results. The signal can get loud, with the potential to damage your speakers and your hearing. Please exercise the same discretion when using SpectrumWorx as you would with any equipment that can produce excessive amplitudes. SpectrumWorx is all about experimentation and we encourage you to dive right in. However, you should be ready to hit that Bypass button. Better yet, keep your overall monitoring levels turned down until you have achieved the sound you want to make.

The Main Window

We've already briefly discussed the main plug-in window, but a little more detail might be in order. As we've said, many of the components that occupy the main plug-in window change to reflect the selected module and module parameter's settings. For instance, a module's knobs will give you an idea of how they are set, but the details of those settings are revealed in the main plug-in window.




Figure 3: The SpectrumWorx main window

In, out, mix

These knobs are always in view, and they tell us pretty much what we'd expect, given their names. The **input** knob is used to control the signal level coming into SpectrumWorx, while the **output** knob adjusts the gain of the outgoing signal. The **mix** knob controls the ratio of the “dry” and “wet” (unprocessed and processed) signal output. A value of 0% produces an entirely unprocessed signal, while 100% will allow us to hear only the processed audio. A 100% wet signal is desired when SpectrumWorx is used on a mixer's auxiliary “send” — though you can certainly crank it up to full-blast whenever you like!

Module name, parameter name and value

As we said earlier, the top-most field in the main plug-in view reveals the name of the currently selected module. Here, it is the Slicer module. Below that is a parameter field which tells you the name of the currently selected parameter on that module. In this case, “Slice time”. Just below that is the value of that parameter.

Module gain, wetness and frequency range

These three controls are available for every module. They are shown only if a module is selected.

- **Gain** controls the gain applied to the signal at module output.
- **Wet** specifies how module input and output signals are to be mixed.
- **Frequency range** specifies exactly what the name says: the frequency range inside which the module operates. Outside the range, the signal stays unchanged.

All four parameters — gain, wetness, start frequency, and stop frequency — are controllable by the LFO too.

External audio

Many SpectrumWorx modules work by processing and combining main audio input with side channel audio. These side channel files are loaded (and unloaded) via left (and right) mouse click in the “External audio” area of the main plug-in window.

There are a few known limitations:

- Audio files are not streamed, but are fully loaded into memory, hence their length is limited by available memory.
- 8 bit, 16 bit and packed 24 bit sample resolutions are supported.
- Mono and stereo channel configurations are supported.
- All plugin programs share the same sample.
- SpectrumWorx automatically tries to resample the external sample to match the main channel’s sample rate. However, the resampling will only work if the target sample rate is 48 kHz or lower.
- Changing the active sample rate (in your host) does not automatically resample a loaded sample; it has to be manually reloaded.

One final important note to remember is that a loaded external audio file takes precedence over the signal received from a side channel input.

LFOs

LFOs are a common modulation source. They are found on synthesizers and effects processors, and are used to produce rhythmic, periodic effects such as tremolo, vibrato or auto-wah. “LFO” is short for “Low Frequency Oscillator” and that is precisely what it is: an oscillator that runs (in most cases) at frequencies below the audible range.

Each knob-based, combo-box based or LED-based parameter in a given module can be modulated by its own LFO. This allows for some very complex sonic manipulation. When you select a parameter, the LFO settings in the main plug-in window will change to reflect the values of the LFO attached to that specific parameter. The LFO is activated with the blue LED button just to the right of the LFO label — this button is only visible when a parameter is selected.

Period

This parameter determines the speed (or frequency) of the LFO. There are two ways to adjust an LFO’s period, the “free” mode and the “synchronized” mode. You will most often use the latter.

In “free” mode simply drag the slider to choose your desired value. In practice however, you’ll usually want your LFOs to follow the rhythm of your track and this is where the “synchronized” mode kicks in. In synchronized mode SpectrumWorx uses time, tempo and time signature information from your host DAW to turn the period slider into a snapable control that will allow you to walk through and select only the periods that are meaningful for the current tempo and time signature (i.e. whole number ratios of the beat and bar durations). The range of the period slider is also automatically adjusted according to the active time signature. The minimum duration is set to 1/8th of the current time signature’s beat/base note duration and the maximum duration is always 16 bars. For the typical 4/4 measure this gives you values from 1/32 to 16/1.

Instead of always using power-of-two ratios as is the usual (and wrong!) practice, the values that fall between the durations of one beat and one bar are also carefully calculated based on the current time signature. Example: for the 4/4 time signature you get to set the value of 1/2 bar — which is meaningless for let’s say a 9/8 time signature, where you get to set 1/3 of a bar, or in a third case 1/2, 1/3, 1/4 and 1/6 for 12/8 time signature. Time signatures with prime number numerators have no whole number bar duration dividers so, for example if you have 7/8 time signature, you’ll jump straight from 1/1 to 1/7, then to 1/14, 1/28 and 1/56 (which is the duration of the 8th part of a quaver expressed as a part of the duration of a bar in the 7/8 time signature). Might look complicated, but it really isn’t!

As you have probably noticed, SpectrumWorx synchronizes its LFOs to the bar duration, as opposed to the beat duration. This might require a bit of getting used to but we hold this to be more musically correct and it plays better with creating presets that are supposed to work across different time signatures. So if you, for example, create a preset with an LFO synced to the duration of a bar (1/1) under a 4/4 time signature, when you load it under a project with a 3/4 time signature it will still be synchronized to a bar (as opposed to 4/3 of a bar).

The story, of course, does not stop here. The jump from “free” to “synchronized” is not a simple on/off switch. To turn on host tempo synchronization SpectrumWorx provides the usual three modes, independently selectable using the **N**, **T** and **D** text buttons. The **N** represents a normal

mode that enables straightforward host tempo synchronization described above. For more complex rhythms, the **T** and **D** buttons also enable snapping to triplet and dotted notes, respectively. When all three sync modes are deactivated, the LFO switches to “free” mode.

As usual, for this magic to work correctly, SpectrumWorx relies on the host to correctly implement and provide the timing information functionality. There are some hosts that do not provide timing information at all; for such hosts SpectrumWorx will assume a 4/4 measure with a 120 BPM tempo and set its LFOs to free mode and disable the sync buttons. If you load a SpectrumWorx preset with a synced LFO in such a host, SpectrumWorx will load the preset, adjust any synced LFOs for the mentioned “default tempo and measure” but leave the sync buttons disabled. This is done so (as opposed to automatically switching all LFOs to free mode) to prevent the destruction of the preset’s LFO synchronization setup by subsequent saving (accidental or automatic) of the preset while in such a host.

Finally, it is important to note that SpectrumWorx, as a frequency domain effect, inherently processes data in blocks (or frames) and, as such, cannot update its LFOs (i.e. their phase) for each sample but for every block. What this means in practice is that, unlike for time domain effects, the SpectrumWorx LFOs do not change perfectly smoothly, and that not all period values make sense given a particular setup — values comparable to or smaller than the current block size will cause the LFO to exhibit aliasing effects. Whether or not the effects of “quantized” LFO updating will be audibly noticeable depends on the period you set and on your host and SpectrumWorx engine setup. Currently, LFOs are updated every frame or every (host) data buffer block, whichever is larger. You only need to make sure that your LFO periods do not go near or below this value — unless of course you want the “undefined behavior”!

Waveform

Each LFO can output one of several different waveforms. This waveform has a lot to do with how the effect will sound. For instance, the regular and smooth **Sine** wave will result in a steady, even rise and fall of the destination parameter, while the **Square** wave will switch abruptly between the high and low parameter values as determined by the Range setting.

In addition to Sine and Square you get **Triangle**, **Sawtooth**, reverse Sawtooth (**htootwaS**), **Exponent**, three random waveforms (**Hrandom** — random and hold, **Grandom** — random and glide, and **Whacho** — sort of a totally-random waveform), and two impulse waveforms (**Dirac up** and **Dirac down**).

You can make a pretty good guess at how each will sound by looking at the little graphic that goes along with each of them.

Phase

Phase control enables the waveform to start from some shifted, non-zero value. It can go from -50% to +50%.

Range

This allows you to tell SpectrumWorx exactly how much of the LFO you would like to apply to a given parameter's value. It takes the form of a slider with dual handles. The left-most handle controls the lowest range of the LFO, while the right determines the maximum effect to be applied.

For instance, say you want to apply an LFO to the Start frequency of the Slew Limiter module from the first screen grab. You want the LFO to drive it all the way to the fullest range, but never dip below the halfway point. You'd set the left slider handle about halfway across, and the right one all the way to the right. The minimum and maximum range values will of course differ depending on the selected parameter.

The Module Bank

The Module Bank is designed to hold modules. You can have five modules at any given time. The signal flows from the left-most module to the right.




Figure 4: The module bank with the module menu open

Loading modules

Modules are loaded by clicking the arrow at the left of the Module Bank and selecting a module from the categories shown in the pop-up menu.

Muting, deleting and rearranging

In addition to the individual parameters, each module has a **mute** switch, which takes the form of a blue LED at the bottom. On the very top of each module is an **x** used to delete that module from the Module Bank. Deleting a module to the left of any other modules will cause the remaining modules to slide over to occupy the deleted module's former position in the Bank.

Any module can be dragged and dropped, making rearranging a sound a snap. Dragging a module into a new position most often results in an entirely new sound!

Settings: Engine

Hold on to your beanies, boys and girls, this is where we start to get a little technical! The Engine settings are where you tell SpectrumWorx exactly how much effort to put into its algorithms, as well as how its inputs and outputs should be arranged. We'll start at the top and work our way down.



Figure 5: The Engine tab of the Settings window

There may be times when you find that you are putting a lot of strain on your CPU and you might feel the need to fiddle with the Engine a bit to relieve some of the stress on your computer. There may also be times when you intentionally dial in settings that will compromise the incoming signal in interesting ways. You should note that the engine settings are saved in presets, so don't be surprised if you fire up a new preset only to find that your engine settings are suddenly changed!

All engine parameter changes take effect immediately and it is not required to stop the playback beforehand; SpectrumWorx was designed and thoroughly tested to take almost anything you throw at it. However, unlike for plugin stability, we cannot guarantee sonic stability. If your playback is active, depending on the changes you make and on the loaded modules, you may hear (very) loud transient noises — which may also sound cool sometimes — until the signal stabilizes with the new settings.

Note that changes to the engine setup parameters affect the plugin's latency and/or its I/O configuration and therefore require cooperation with the host software, which means that the stability of the whole process relies not just on SpectrumWorx but also on the host. It may happen that you run into a rare host that does not like these things to change while playback is active (especially in a very rapid manner, like fast scrolling through the presets with different engine setups) in which case you can try stopping the playback/transport before tweaking the engine parameters.

A brief diversion into how SpectrumWorx does its thing

To fully understand how some of the Engine settings work, you need to know a little about how SpectrumWorx itself operates. SpectrumWorx uses what is referred to as the "Short Time Fourier Transform", which we'll abbreviate as "STFT" (to avoid Carpal Tunnel Syndrome!). The learned among you may be familiar with DFT, or "Discrete Fourier Transform", or FFT (Fast DFT) and will know immediately what we're talking about.

Put as simply as possible, a Fourier analysis breaks down a signal into individual frequency components. In audio, you might use such a process to identify and separate individual harmonics or groups of harmonics. If you understand that any given sound may be broken down to nothing more complex than a pile of simple sine waves, each with its own frequency and phase, then you can begin to see how powerful such an analysis can be! This is, as we stated earlier, the basis for additive synthesis.

In the case of STFT, short segments of sound are analyzed, rather than the whole signal. This has a few benefits. First, and most obviously, it is faster and therefore more amenable to real-time manipulation. And it has less delay. But what is most important, it is also similar to the way our own hearing works. The ear similarly analyzes only a short segment of audio at a time (on the order of 10-20 ms worth). Therefore, to perform a spectrum analysis having time and frequency resolution comparable to human hearing, we must limit that window of time accordingly. The proper way to extract a short time segment from a longer signal is to multiply it by a window function. We'll talk about window functions a bit later.

Frame size

Now, if you've read the above paragraph, you may be able to understand how setting the Frame Size affects our signal. A shorter frame means that we will be able to capture fast changes of the input signal, while a larger frame size will provide a more accurate frequency analysis (since STFT gives us just an approximation of the input spectrum). This comes from the fact that STFT presupposes a fixed spectrum for the entire duration of the frame.

Your signal's quality as well as the overall sound is going to greatly depend on the Frame Size setting. Some modules are even targeted to specific frame sizes. Phasevolution, for example, will give best results with short frames (for 44100 and 48000 Hz sampling frequencies).

SpectrumWorx offers seven different frame sizes, ranging from 128 samples to 8192 samples. This figure directly coincides with the number of frequency components that will be produced.

If you choose a frame size of 128 samples in the time domain, then this chunk will be approximated with 128 frequencies (as we said, STFT is just an approximation). A value of 128 is in fact rather low,

but the trade-off in this case would be that your time resolution would be high, resulting in an ability to capture fast changes without smearing them in time. If the input sampling rate is 44100 Hz, then 128 samples would represent less than 3 ms of signal — which is a very good time resolution. At the same time, an input frequency spectrum range of 22050 Hz would be approximated with only 128 uniformly spaced frequencies, i.e. one frequency component every 172 Hz, which is rather poor.

Conversely, a setting of 8192 would provide excellent frequency resolution, but wouldn't be so great in the time domain — chunks analyzed would be 185 ms long in this case. On the other hand, an input frequency range of 22050 Hz would be approximated with 8192 frequency components, which is a great frequency resolution of less than 3 Hz!

A good compromise between time and frequency resolution would be a frame size of 1024 or 2048. This also is the best match to the human auditory system's time resolution.

It is important to note that shorter frames consume more CPU than longer ones, because of the overhead of running the whole process more often.

Overlap factor

We just keep hitting you with these technical terms, don't we? Suffice it to say that the longer the overlap, the smoother the results across time. SpectrumWorx allows you to choose four values: 0% (no overlap), 50% overlap, 75% overlap and 87.5% overlap. A setting of 75% or higher is required for proper pitch shifting. The "best" or "optimal" overlap also depends on the window type you choose, but 75% is typical and recommended.

If the window and STFT are applied to non-overlapping partitions of the signal, a significant part of the signal is lost due to most windows exhibiting small values near the boundaries (not true for the Rectangular window!). To avoid this loss of data, the transforms are usually applied to overlapping signal sequences. The overlap factor tells us how much we can advance the analysis time origin from frame to frame. In SpectrumWorx, an overlap of 75% means that it takes 4 steps to move one whole frame ahead in time.

Longer overlaps also mean more processing, so your CPU meter is going to go up accordingly. An overlap of 87.5% (factor 8) combined with a short frame length might just not be too healthy for your CPU!

Window type

We've already discussed the importance of the window's frame size, but the window's "shape" is also a factor in achieving good results. The main problem is that windowing causes the signal's Fourier transform to have non-zero values (commonly called spectral leakage) at frequencies other than those really present in the signal. Windowing is in fact used to shape and reduce the spectral leakage! We could say that windowing describes how the signal is "seen" by the analysis algorithm.

Think of two sinusoids with different frequencies: spectral leakage can interfere with the ability to distinguish them spectrally. If their frequencies are dissimilar, then the leakage interferes when one sinusoid is much smaller in amplitude than the other. That is, its spectral component can be hidden by the leakage from the larger component. But when the frequencies are near each other,

the leakage can be sufficient to interfere even when the sinusoids are equal strength; that is, they become unresolvable.

The easiest windowing operation is simply truncating the signal on the left and right at frame boundaries. This can be modeled mathematically as a multiplication of the signal by the so-called Rectangular window. The Rectangular window has excellent resolution characteristics for signals of comparable strength, but it is a poor choice for signals of disparate amplitudes. This characteristic is described as low-dynamic-range.

At the other extreme of dynamic range are the windows with the poorest resolution. These high-dynamic-range low-resolution windows are also poorest in terms of sensitivity; that is, if the input waveform contains random noise close to the signal frequency, the response to noise, compared to the sinusoid, will be higher than with a higher-resolution window. In other words, the ability to find weak sinusoids amidst the noise is diminished by a high-dynamic-range window.

In between the extremes are moderate windows, such as Hamming and, the default window in SpectrumWorx, Hann. The Hann window provides a good compromise and works well for general audio analysis, so it can be a good place to start!

In summary, choosing a window involves a tradeoff between resolving comparable strength signals with similar frequencies, and resolving disparate strength signals with dissimilar frequencies. SpectrumWorx provides nine different window types to choose from: Hann, Hamming, Blackman, Blackman Harris, Gaussian, Flat top, Welch, Triangle and Rectangle. Go experiment!

Ripple factor

You'll notice the Ripple factor displayed at the bottom of the page. This tells you if the current engine settings (Frame Size, Window Type, Overlap) will result in a so-called "perfect reconstruction" of the sound, or if some ripple will be introduced by the engine itself regardless of the modules being used. Successive windowed frames should overlap in time in such a way that all data are weighted equally — this produces no ripple.

Different window types have different frame sizes and overlap factors where they produce no ripple. For example, a Rectangular window can advance (overlap size) by $(\text{Frame Size})/k$ samples, where k is any positive integer, and a Hann or Hamming window can use any step size of the form $(\text{Frame Size}/2)/k$.

What is the best engine setup?

In SpectrumWorx there is no "best engine setup". Modules differ vastly in the engine parameters they work best with. It is up to the user to find that happy compromise between frequency precision and time resolution, or between sonic integrity and bringing your computer huffing and puffing to its knees.

Input mode

This bit is mostly self-explanatory. SpectrumWorx offers a variety of input and output modes tailored to various needs. By default, 2 in / 2 out (Stereo) is assigned, but you can switch that over to mono operation, or choose one of the modes designed for use with side-chaining effects.

When supported by the host, side-chain input modes are a powerful alternative to the built-in external sound file functionality. In fact, you will probably want to use the host's side chain facilities as the side chain source for your setups and presets with two-input modules in your final production phases, as this will give you far greater control and power over what and how exactly is fed into the SpectrumWorx side chain input. You will have to consult your host's manual to find out if it supports side chain routing and how to set it up.

In addition to side chain modes, the mono modes can be useful in situations where you are working with individual mono tracks and wish to achieve the maximum sound fidelity (primarily in terms of phase coherence and stereo imaging). You can load up multiple SpectrumWorx instances with each working on a different track; you will, for example, notice that Phase Vocoder based modules sometimes give better results that way.

The SpectrumWorx engine is not actually limited to the basic I/O modes which are made easily accessible through the SpectrumWorx GUI and the Input mode parameter. If your host provides the means, you can set SpectrumWorx up to work with more complex channel arrangements, such as 8-in / 4-out modes.

Settings: GUI and About

As you've probably guessed, this is where you can tell SpectrumWorx how you'd like its interface to behave, and affect aspects of its appearance.

Mouse over reaction

This determines how SpectrumWorx displays the various parameter settings in the main plug-in window when using your mouse in the rack. You have three options here:

- **Never** only displays the information for the currently selected parameter.
- **Module selected** will cause the information for all knobs of a selected module to appear when you mouse over them.
- **Always / nothing selected** will display information for any knob and/or module any time the cursor passes over it.

LFO update behavior

This setting determines how and under what circumstances SpectrumWorx updates its various parameter displays if those specific parameters are being modulated by an LFO. Basically, it can look as if any LFO-affected knobs are being tweaked all by themselves.

We probably needn't tell you what setting this function to **Never** will do. However, setting it to **Always** will result in all LFO-driven knobs being updated all of the time. It can look pretty cool, too, though it will use slightly more of your CPU. **Control selected** will display the LFO action only for a selected knob, while **Control active** will display the action of any control you might have moused over, given the appropriate settings of the Mouse over reaction function.

Load last session on startup

A simple on/off switch that does just what it says on the label. Erm, that is, your last session will be loaded up when you instantiate the plug-in. Keep in mind that a session includes all settings, including Engine settings. If, for instance, you change the frame length, that change will only be remembered if you have this option selected.

As already mentioned in the Input mode section, this need not work as expected in all hosts, as certain hosts can and will override your last session settings with the settings they have saved in their project files (which incidentally might often be the same settings as the ones saved in your last session preset).

Hide cursor on knob drag

Switch this on if you want the cursor to disappear while you turn a knob. Otherwise, leave it off.

Set window and menu opacity

This slider will determine the opacity of your Settings and Presets windows as well as pop-up menus. All the way left makes them fairly transparent, while all the way right makes them entirely opaque.

Settings: About

Find out what version you are running and who is behind the superlative SpectrumWorx by clicking the **About** tab.

Presets

Like the Settings window, the Presets browser is a collapsible “child” window that appears along the left side of the SpectrumWorx interface. As you’d expect, it is used to search for, load, save and manage the many factory presets, along with any you might want to create for yourself.




Figure 6: The Presets browser

Location

The directory where you keep your presets is displayed at the top of the window.

Save, Save as, Delete

Save, Save as and Delete buttons are available just below the preset location. When you elect to **Save as** your own preset, you’ll be provided with an opportunity to give it a name. Use the **Save** button to overwrite an existing preset, which you need to select first. Use the **Delete** button to remove any existing or selected preset that you no longer want to have available to you.

Ignore external audio

The **Ignore external audio** option allows you to skip the loading of external audio files from other peoples' presets, or the saving of the currently loaded external sample into your presets. This is useful if you want to test existing presets with your own audio sample.

List

The preset list shows the preset files and the subfolders available in the current folder for easy navigation. Single clicking on a preset name automatically loads the preset. A double click or a press of the enter key gives you the opportunity to rename a preset.

Description

There is a description field just below the preset list. This is a place where you can write (just single click inside the field) or read descriptions of, or comments about, selected presets. This goes a long way towards helping to keep all of the presets manageable.

The Modules

Modules are the heart and soul of SpectrumWorx! Fifty-seven of them provide you with countless ways of creating effects. As we already said, each module is a sophisticated audio processor in and of itself, with its own, unique selection of adjustable parameters. The following pages describe each module in some detail.

Modules are grouped into nine categories in the load menu:

Category	Count	What's in it
Pitch	7	Shifting, following, tuning and oscillating pitch
Timbre	8	Filtering and reshaping the spectrum
Time	6	Freezing, slicing, reversing and decaying
Space	3	Frequency domain echo and reverb
Phase	4	Flattening, randomizing and flipping phase
Loudness	5	Gain, dynamics and noise
Combine	10	Everything that mixes main with side channel
Phase Vocoder	9	True-frequency versions of the pitch and eximploder modules
Miscellaneous	5	The ones that fit nowhere else

Remember that every module also has its own Gain, Wet and Frequency range controls in the main window, and that almost every module parameter can be driven by its own LFO.

Modules: Pitch

Pitch Shifter

This module is as straightforward as it gets. It shifts the pitch of the incoming signal in semitones and cents. You can go anywhere between two octaves higher and two octaves lower than the original signal. It can also act upon a selected frequency range; any frequencies that fall outside of this range will not be affected. The selected band of frequencies is determined by the **Frequency range** slider located in the Main window.

Pitch Follower

The main channel is pitch-shifted to the dominant frequency detected in the side channel (via **External audio** or via **Settings / Input mode**). The **Speed** parameter determines how quickly the pitch-shifting to the target frequency should be done, in order to provide a smoother result.

TuneWorx

You know the sound: it's the classic Autotune effect. The main channel's pitch is detected and shifted to the nearest selected semitone. If all semitones are selected, you'll get a chromatic scale. Check out the Autotune subfolder in the Presets browser for some musical scale presets!

Pitch Magnet

This one forces the input channel pitch to the pitch selected with the **Target** knob. The **Strength** knob determines the speed with which to reach the target pitch.

Sumo Pitch

What happens when your channels are pitted against one another like Sumo wrestlers? The input signal and side channel are analyzed with the internal pitch detector. Both are then pitch-shifted towards each other and towards a center pitch. The speed of the shift is selectable with the **Speed** knob. After both signals have been shifted, they are blended together. The amount of blending is controlled with the **Blend** knob: when set to 0, only the main channel is output; a setting of 100% sends only the side channel to the output.

Pitch Spring

This is a pitch oscillator. You can choose the **Depth** and **Period** (rate) of the oscillation. There is also a direction combo box where you can select how the "spring" oscillates:

- **Up** will cause the spring to oscillate between the incoming pitch and a positive (higher) pitch determined by the Depth knob's setting.

- **Down** will cause it to swing from the incoming pitch to a negative (lower) pitch as determined by the Depth knob.
- **Symmetric** will result in the pitch oscillating back and forth from a negative pitch, through the original pitch of the incoming signal, and towards a positive pitch.

Octaver

This is an octaver effect: it adds two (or just one) octaves onto the original signal. It can add an octave that is two octaves down or one octave down from the original signal (to turn your guitar into a bass, for example), or it can add two or one octave up from the original sound to achieve a kind of harmonizer effect (a 12 string guitar, for example). It can also mix down and up octaves. Each octave has a gain control, and both octaves share a common cut-off frequency to limit the operation to only the low-passed signal. Fun!

Modules: Timbre

Bandpass

This is a simple bandpass filter with a variable range. An **Attenuation** knob controls how much of the signal is filtered. **Start frequency** and **Stop frequency** controls (located in the Main window) control the lower and upper range of the bandpass.

Bandstop

A simple bandstop or “notch” filter. The user controls attenuation and the start and stop frequencies that will be suppressed. It’s the exact opposite of the Bandpass module.

Ah-ah

This is sort of like a wah-wah. It filters a region with a variable center frequency. You must use an LFO to modulate the **Center frequency** knob for the intended effect. We recommend an LFO with a period of 0.5 seconds and a range from 300 to 1300 Hz. **Width** determines the width of the moving region, and **Gain** controls how much gain is applied to it.

Smoother

Smooths the frequency spectrum of the incoming signal, causing it to lose sharpness and detail for a blurred image.

Sharper

The opposite of the above module: this one sharpens the spectrum. The smoothed signal is subtracted from the original to arrive at the sharpened signal. The amount of sharpening is controlled by the **Sharpness** knob, which represents the length of the averaging filter in Hz. The sharpness is increased with the **Intensity** knob and then hard-limited with the **Limit** control to make it more stable.

Centroid

Centroid is an adaptive, self-tuning bandpass filter which follows the center frequency with a controllable symmetric band curve (**Border slope**). The module determines the time-variable center frequency with different processing modes:

- **Centroid** finds the weighted center frequency.
- **Peak** finds the highest peak in the spectrum.
- **Dominant** estimates the dominant pitch.

Frequencies outside the **Bandwidth** region around the centroid frequency are removed; the rest inside the band is smoothed with the Border slope from center to borders.

Tonal

This one lets through only highly tonal frequencies. First the module finds peaks in the spectrum, then it estimates the peak strength (how strong the peak is when compared to its neighbors), and if the strength is above the value determined by the **Peak strength** knob, and if the peak is stronger than that determined by the **Global threshold** or **Local threshold** settings, the peak will be passed through. The rest of the signal is attenuated by the **Attenuation** control.

Atonal

Similar to the Tonal module described above, but instead it attenuates the peaks and allows the non-peak components to pass through.

Modules: Time

Freeze

A module with an attitude! Freeze is a time freezer: press the **Freeze** button and the incoming sound will freeze in its current state, and will stay that way until another **Freeze** or until it is sent back to normal state via the **Melt** button. The transition from one frozen state to the next, and from frozen to normal state, is controlled with the **Transition time** knob — this enables the transition to be smooth and gradual.

Slicer




Figure 7: The Slicer module

This is one of the coolest modules in the collection. It slices the signal into chunks that are alternated with one of three different options, defined by the combo box on the module's face:

- **Sample & Hold** fills the gaps with a “frozen” sample of the last frame from the main input before the slice began.

- **Silence.**
- **Side channel** or external audio.

The **On time** knob determines the length of the main input slices (normal mode), and the **Slice time** knob controls the length of the silence, Sample & Hold or side channel slices. Some very cool effects can be obtained by using very small values.

Wobbler

As you can probably guess, this module creates a sinusoidal-based “wobbling” sound by changing the amplitudes of the input frequencies. The **Amplitude** knob defines the intensity of the wobbling (or how much the sound seems to go up and down). An **Offset** parameter is provided so that you can shift the whole input down before the wobbling stage — it might therefore be advisable to start with a negative Offset value that reflects the positive setting of the Amplitude setting. The **Period** function defines the rate of wobbling.

Reverser

This one breaks the input down into chunks and reverses them. For instance, the name “Spectrum-Worx” might be output as “cepSmurtxroW”. There is only one parameter: **Chunk length**. You can achieve some supremely alien effects with this one. One warning note: Reverser introduces delay. However, that delay can actually be used to create some interest as well.

Imploder

Imploder freezes incoming high amplitudes and introduces a slow decay towards the selected **Limit** level. Additionally, the frequency is gradually altered, controlled by the **Glissando** knob; glissandos are given in cents per second. **Decay** governs the time needed to decay from the current amplitude value down to the Limit value. Amplitudes lower than the **Gate** parameter are left untouched. Amplitudes are decayed until a higher signal is received at the input or until the Limit is reached.

Exploder

Pretty much the same idea found in the Imploder: this one also freezes amplitudes, but instead of introducing decay it ramps them up towards the **Limit**. Glissando is applied in the same way as it is in the Imploder. Amplitudes are ramped up until a louder signal is received at the input or until the Limit is reached. Amplitudes lower than the **Gate** parameter are left untouched.

Modules: Space

Frecho

This is a frequency domain echo that functions with pitch. The **Distance** control determines the perceived distance of the reflection; the minimum is set to 17 meters, resulting in 34 meters for the sound to return to the signal's perceived source, simply because any less than that and we wouldn't have echo but reverberation. **Absorption** controls how much of the sound is absorbed and how much is reflected back.

Using the **Pitch** function you can pitch-shift each echoed sound, and that is precisely what Frecho is all about! The values are given as semitones per echo. The **Mix** knob controls how much of the original sound is heard as compared to the echoes: a setting of 0 provides no echo, while a value of 100 provides only the echo.

Frevcho

Similar to Frecho, but the echoes are reversed! The controls are the same as in Frecho, but the sound is quite unique and different.

Freqverb

Similar to Frecho, but this time it is a frequency domain reverb, with an ability to pitch-shift every single reverberation. In other words, there are so many echoes in such a short time period that our ears can't differentiate between them.

The **Life time** function governs how long the sound reverberates. The longer the setting, the more sound builds up — it can even be driven into saturation! The **Room size** parameter is similar to the Absorption function in Frecho, and controls how much attenuation is applied to each reflection. **Pitch** functions the same as it does in Frecho: each reverberation can be pitch shifted, and once again the values are given as semitones per echo. The **Mix** knob controls the ratio of affected and unaffected signal.

Modules: Phase

Robotizer

Produces a classic robotic effect in the usual manner: by flattening the phase of the signal. Don't forget that you can adjust the desired start and stop frequencies, wetness and gain via the three controls in the Main window! Effects run the gamut from pseudo-vocoder effects to unusual phase-cancellation filtering.

Whisperer

This one provides a sort of whispering effect by randomizing the phases. Note: you'll need a short frame size for proper operation.

Phasevolution

This one creates an accelerated change in phase over time, which results in a sort of pulsating pitch effect. As with the Whisperer, this one demands a short frame size for proper operation. It also works better on signals with lots of clear harmonics. The **Period** knob controls the rate at which the phase comes full circle.

Philip

As you might guess, this module flips the phase, but don't confuse it with the "flip" button on your mixing console — this one creates some pretty unusual effects! The effect is quite variable, depending very much on the harmonic content of the incoming signal. Odd harmonics tend to result in a more dissonant sound, while even harmonics are more consonant and pleasant to the ear. Also, the shorter the frame size, the more audible the effect will be.

Modules: Loudness

Gain

Really, this is just a simple gain function. Try it with a selected frequency range, as determined by the Start and Stop frequency slider — and randomize this fellow with an LFO!

Exaggerator

With its single knob, the Exaggerator is a deceptively simple-looking module. However, the results are anything but standard! Exaggerator applies an exponential function to magnitudes, emphasizing or suppressing peaks. The **Intensity** knob is bipolar, with no effect when set to the 12 o'clock position. If a positive setting is dialed in, the higher the amplitude of the incoming signal the more it will be amplified. If a negative setting is selected, the higher the amplitude of the incoming signal the more it will be suppressed.

Denoiser

This effect applies the standard de-noising formula. SpectrumWorx's algorithm uses the side channel input as a noise footprint, or a sum of main and side channel. The main channel will be attenuated when the side channel contains noise (in the latter case). The **Amount** knob controls the intensity of the de-noising algorithm.

Quiet Boost

This module boosts low gain components. It amplifies signal below the desired **Threshold**, working as a sort of expander in the time domain. The **Ratio** of expansion can be selected, and there's a **Noise threshold** control — the module won't boost anything below this threshold, because we don't want to boost all the noise!

Freqnamics

Freqnamics provides dynamics processing in the frequency domain, just like the Quiet Boost module described above, except this one compresses rather than expands. Freqnamics acts as both a limiter and a noise gate. Use the **Limiter** knob to set the threshold of the hard limiter, and the **Noise gate** knob to set the threshold of the noise gate.

Modules: Combine

Every module on this page mixes the main input with a side channel, which you supply either by loading a file into the External audio area or by routing a side chain from your host.

Talking Wind

This is a classic vocoding effect, but with a SpectrumWorx twist! As you probably already know, vocoding uses the frequency spectrum and amplitude of one signal (the modulator) to modulate another (the carrier). Often a human voice will be used as the modulator and then applied to a sound such as white noise to produce the familiar “talking wind” effect.

The Talking Wind module offers a handful of controls. The **Envelope border** knob controls the perceived smoothness of the modulator. An **Envelope gain** knob can be used to add some gain to the envelope to amplify the effect. The **Gain** knob is provided to tame the final output, since it can often become quite loud. The carrier is loaded as an external audio file or fed into a side channel input; the modulator is passed through the main input.

Convolver

A multi-effect: a convolution reverb or a vocoder, depending on the **Phase** parameter! If Phase is set to **Sum**, this is in fact a convolution reverb, where the impulse response can only be as long as the frame size. The IR is fed via external audio.

Convolver has two basic ways of operation, triggered and continuous. The triggered option grabs the IR exactly at the time frames when the **Grab IR** button is pressed — this should be LFO controlled! The IR is then constantly convolved with the input signal, until the button is pressed again. In continuous mode the IRs are constantly taken from external audio, so every frame ends up with a different IR, which produces pretty chaotic results.

The **Phase** parameter controls what happens with the main and side channel phases. If **Sum** is selected, phases are added, representing a “normal” convolution. If **Side** is selected, phases are taken from the side channel, enabling a vocoding effect. The third option takes phases from the main channel, which is also a vocoding effect, but this time the main channel becomes the carrier.

Ethereal

This is a strange one! It compares the side channel signal with that of the main channel, and it replaces the main signal with the side signal if certain conditions are met (main minus side greater than or less than a **Threshold**). It can replace magnitudes or phases, or both.

Vaxateer

Like Ethereal, this is another module that provides conditional replacement of the main signal with that of the side channel. It works by calculating the RMS value (Root Mean Square) of the target

signal — which can be the main or side channel — and then, if the **Swap condition** is satisfied, replaces the main signal with the side channel content. There are eight available swap conditions, selectable via a combo box.

Shapeless

An interesting effect that transfers the frequency “shape” from the side channel to the signal passing through the main input. A **Shape width** parameter regulates the coarseness of the shape estimators.

Colorifer

This one also transfers the frequency shape from the side channel to the signal arriving through the main input. It calculates main divided by side over the selected band width, and then applies this “color” to the main channel. There is an option to also replace the phase of the main channel with that from the side. This produces some extremely bizarre and inspiring effects — for the experimentally minded among you!

Merger

This effect provides conditional combinations of the main signal with that of the side channel. It copies the side channel to the main channel if the selected conditions are met, and depending on the **Threshold** setting. The conditions are available in the combo box.

Blender

Provides linear blending of two channels. In the modules above, some frequencies in the main channel were replaced with side channel frequencies; in the Blender, all frequencies are instead combined. The **Amount** knob controls how much of each channel is sent to the output. A setting of 0 results in only the main channel passing through, 50 means both channels are equally considered, and 100 results in only the side channel material being heard.

Inserter

We’re getting pretty esoteric here! This module inserts a band of frequency bins from the side channel into the main channel. The **Source** value is the position (given in percentages) from where the spectrum is copied. The **Destination** is the target position in the main channel’s spectrum. The **Size** knob controls how much of the spectrum is copied.

“bw%” refers to bandwidth percentage. For example, if a signal is sampled at 44100 Hz then 50% would be half the spectrum, which is half of the highest frequency (22050 Hz) — so 50% would be 11025 Hz. You can select whether magnitudes, phases or both will be copied.

Burrito

The Burrito module copies the side channel to the main channel from random locations. The **Range period** parameter controls how often it will randomly change positions. The **Target range** value defines how many locations are to be replaced, or added if **Sum** is selected in the combo box.

Modules: Phase Vocoder

All of these are special versions of modules that also appear in other categories. However, Phase Vocoder modules operate differently under the hood.

“PV” stands for “Phase Vocoder”, but you should not confuse this with the familiar sort of “robot voice” vocoder — though it should be obvious that SpectrumWorx can provide all manner of robotic voices. Normally, SpectrumWorx draws upon Fourier transform algorithms to work its magic in the frequency domain. The Fourier transform is used to transform time samples to the frequency domain, so we end up with various amplitudes and phases of the resulting sine waves — but the sine waves produced by the Fourier transform are only approximations; they don’t represent the “true” frequencies present in the input signal.

Phase vocoding, on the other hand, can be used to find the true frequencies. After the true frequencies are found, we can manipulate amplitudes and frequencies rather than simply amplitudes and phases. This opens up new territory for effects processing. For example, accurate pitch shifting is possible since we know the true frequencies.

To PV and From PV

To PV and **From PV** are the pair of modules that enter and leave the phase vocoder domain. Modules placed between them in the bank operate on true frequencies rather than on magnitude and phase.

The variants

The remaining seven modules are the phase vocoder versions of modules documented elsewhere in this manual. Their controls are the same, so those descriptions are not repeated here:

- **Pitch Shifter (PV)** — see Pitch Shifter
- **Pitch Follower (PV)** — see Pitch Follower
- **TuneWorx (PV)** — see TuneWorx
- **Pitch Magnet (PV)** — see Pitch Magnet
- **Pitch Spring (PV)** — see Pitch Spring
- **Imploder (PV)** — see Imploder
- **Exploder (PV)** — see Exploder

Many modules that are not designed to work in the phase vocoder domain can provide some surprising results when used in it, so experiment — but proceed with caution!

Modules: Miscellaneous

Armonizer

A familiar pitch harmonizer like that found on the Eventide Clockworks devices, this module adds a pitch-shifted copy of an incoming signal to the original. Use the **Interval** knob to select how many semitones to shift the signal. The **Blend amount** knob controls the blending of original and shifted signals.

Slew Limiter

Limits the speed of the change of magnitudes. It's kind of like the familiar slew limiter in the time domain (used to create portamento), but here it's in the frequency domain of course! As with any slew limiter, you can select the direction of the slew using the combo box, and the **Slew rate** itself.

Note: if the signal is silent at the beginning and then abruptly rises, it might take some time for the slew to produce a sound — if the slew rate is short, the signal will climb slowly from silence to the target frequency.

Shifter

This module shifts frequency bins along the frequency axis to produce unusual clangorous and metallic effects. You can control the **Target** (whether you want to shift only magnitudes, only phases, or both), the amount of shift with the **Offset** knob (in %bw, percentage of the bandwidth), and use the **Tail** combo box to determine what shall be done with the tail data.

An example from the factory floor: if your input is, say, “ABCDEFGHJKLMNOP” and is shifted by three places, then **Leave** will produce ABCBCDEFGHI, **Clear** will produce 000ABCDEFGHI and **Circular** will produce JKLABCDEFGHI.

Note: in some cases Shifter will produce silent output due to sinusoids cancelling out at some shift ranges. This happens only when both magnitudes and phases are shifted, and depends very much on the engine settings (Overlap and Window type).

Swappah

Swaps the three frequency bands determined by the band borders. Swap order is determined by the **Swap order** combo box selection. For example, “High Low Mid” copies the high to low, mid to high, and low to mid.

Quantizer

This quantizes the spectrum. The range is cut into bands of a given width, and all frequencies in any region will have the same amplitude, determined by the first frequency component in that band.

The **Origami** knob allows frequency components to differ from the first one, but to linearly ascend towards the end of the band.

Credits

We've reached the conclusion of the user's guide, but we hope that this is just the beginning for SpectrumWorx. We have no doubt that you will be able to create great sounds and great music with the tools provided by SpectrumWorx, and we can't wait to hear what you do with it!



Figure 8: Little Endian

The original team

SpectrumWorx was developed and released by Little Endian Ltd.

Original idea	Alexey Menshikov
Programming	Domagoj Saric
DSP expertise	Ivan Dokmanic
Graphic design	Matija Bosnjakovic
Documentation	Scot Solida
Directing	Danijel Domazet

Special thanks

To the beta-testers: Martin Walker, Cinning Bao, Sander Heijndijk, Bill Davies (tonecarver), Robert DeNeefe, Mads Ljungdahl, and many others.

As with all folks who have decided to open source great commercial products at the end of their development, the Surge Synth Team is very grateful to Little Endian for making that decision. The original source dump is at github.com/LittleEndianLtd/SpectrumWorx.

Post scriptum

A very special thanks to Jean Baptiste Joseph Fourier (21 March 1768 – 16 May 1830) and to all who have supported the development of SpectrumWorx, the ultimate sound mangler!