HARMONIC SHIFT OSCILLATOR

For the Generation of Harmonic and Inharmonic Spectra

Manual Revision 1.0.1

NEW SYSTEMS INSTRUMENTS

SPECIFICATIONS

Size	12HP
Depth	25mm
Power Consumption	+12V 80mA −12V 75mA
Tuning Range	7Hz–50kHz
Tuning Accuracy	~7 octaves
Output Range	~+9dBu–~+12dBu, −8V–+8V peak
Input Impedance	20kΩ (FM and HS), 100kΩ (others)
Output Impedance	150Ω
Output Drive	2kΩ (min), 20kΩ+ (ideal)

INSTALLATION

Before installing the module, make sure the power is off. Attach the power cable to the module and to the bus. Double check the alignment of the red stripe (or the brown stripe for a multicolor cable) with the markings on the module and the bus. The red stripe should correspond with −12V, as is standard in Eurorack. Check the documentation of your bus and power solution if you are unsure. Screw the module to the rails of the case using the provided screws. (M2.5 and M3 size screws are provided.)

New Systems Instruments modules all have keyed headers and properly wired cables. But please remember to double check the other side of the cable for proper installation with the bus. Additionally, if using a different power cable, note that not every company wires modular power cables such that the red stripe will align properly with a keyed header. While our modules are reverse polarity protected as much as is practical, it is still very possible that you could damage the module, your power supply, or another module by installing the power cable improperly.

Lastly, please fully screw down the module before powering on your case. The electronics are potentially sensitive to shorts, and if the module is not properly attached to a case, there is a risk of contact with conductive or flammable matter.

BASIS

The Harmonic Shift Oscillator creates a sound from abstract parameters describing the characteristics of that sound. Specifically, given angular frequency ω, harmonic level L, and harmonic stride S, this produces two waveforms according to the following equations:

$$\sum_{n=0}^{\infty} L^n \sin((nS+1)\omega t)$$

$$\sum_{n=0}^{\infty} L^n \cos((nS+1)\omega t)$$

These are the real and imaginary components of the complex waveform:

$$\sum_{n=0}^{\infty} L^n e^{i(nS+1)\omega t}$$

EXPLANATION

All waveforms can be constructed by adding together a set of simpler waveforms at various amplitudes and frequencies. The relationship between frequency and amplitude for a given complex waveform is known as that waveform's *spectrum*. Generally this spectrum is given in terms of *sine waves*, which our ear hears as pure tones.

When the spectrum only has nonzero amplitudes at frequencies that are integer multiples of a fundamental frequency, that spectrum is said to be *harmonic*. Otherwise, the spectrum is *inharmonic*.

The formulae above each have two parts: the part expressing the frequencies of the components, and the part expressing the amplitude of each of those components. The Harmonic Shift Oscillator will produce sounds where the nth harmonic is $nS + 1$ times the fundamental frequency, ω. When S is an integer, $nS + 1$ will also be an integer, and you'll get a harmonic spectra. Otherwise, you'll get an inharmonic spectra. The other part of these equations is the amplitude of each of these waves: L^n. Spectra sound "brighter" or "darker" depending on how much high frequency content is present. Maximum brightness is achieved when L is 1, and so every harmonic has the same volume. The analog limitations mean this is not quite possible in practice, but very bright sounds are still achievable. Lowering L to 0.5 would produce a spectrum with harmonics at relative amplitudes 1 for the first, 0.5 for the second, 0.25 and 0.125 for the third and fourth, etc.

By controlling L, you directly control how bright or dark the sound is. By controlling S, you control the character of the spectrum. And lastly, by controlling ω, you control the fundamental frequency.

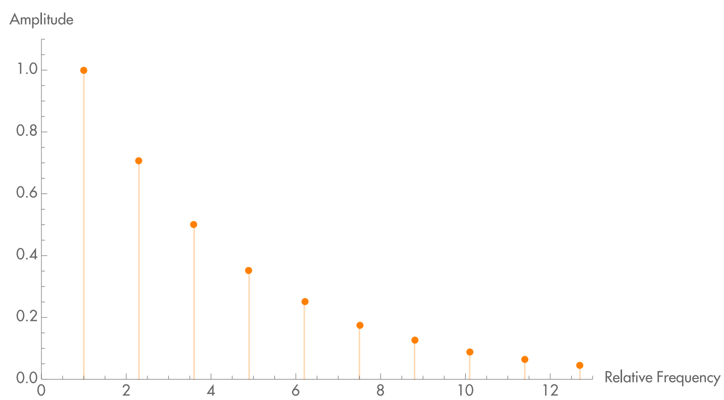

HSO Spectrum at $S = 1.3$ *and* $L = 0.707$

INTERFACE

1. Coarse Frequency – Control the frequency, ω, ranging from about 7Hz to 50kHz.

2. Fine Frequency – Control the frequency, ω, ranging about a fifth above or below the frequency determined by the Coarse knob.

3. Harmonic Level – Control L, the brightness of the spectrum, ranging from 0 to 1.

4. Harmonic Stride – Control S, the spacing between each harmonic, ranging from 0 to about 4.

5. Frequency Modulation Attenuator – Attenuator for #8, the FM input.

6. Harmonic Stride Modulation Attenuator – Attenuator for #9, the HS input.

7. Harmonic Level Modulation Attenuator – Attenuator for #10, the HL input.

8. Frequency Modulation – Modulate ω, the fundamental frequency. Typical range of ± 8 octaves, or 1.6 V/octave.

9. Harmonic Stride Modulation – Modulate S, the spacing between each harmonic. Typical range of about ± 4.

10. Harmonic Level Modulation – Modulate L, the brightness of the spectrum. Typical range of ± 1.

11. V/Octave Input – Control ω, the fundamental frequency, using a V/Octave scale.

12. ⊢ Output – Outputs one phase of the waveform. This is the cosine, real, or +90° component.

13. ⊥ Output – Outputs the other phase of the waveform. This is the sine, imaginary, or 0° component.

USING THE HARMONIC SHIFT OSCILLATOR

The Harmonic Shift Oscillator offers a huge field of possible waveforms, and the best way to get to know it is to use your ears and walk through the wavescape by turning the knobs. Unlike many other methods of inharmonic synthesis, the Harmonic Shift Oscillator is very intuitive, and won't require much theoretical knowledge to get good outcomes.

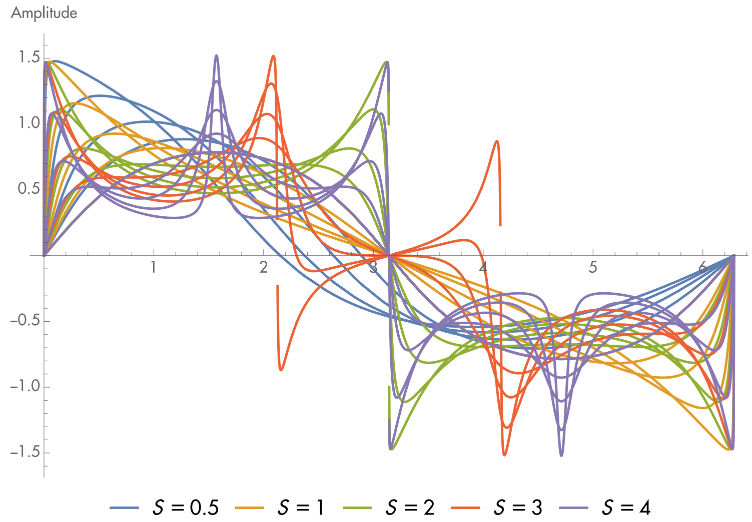

HSO waveforms for various values of S and L.

That being said, if all those possibilities are overwhelming, here are some tips to get started.

First, turn the LEVEL knob somewhere between halfway and 4 o'clock; and plug one of the outputs to your mixer, or however else you get audio out of your modular. To get harmonic sounds, you need to set S to an integer. S is set to 1 when the STRIDE knob is about halfway, or vertical. To set S, slowly turn the STRIDE knob, listening for the beat frequencies. When they slow down (they won't ever *completely* stop), you'll know you've reached your target. After that, they'll start speeding up again. The STRIDE knob is very sensitive, so take your time! When you've reached 1, the output will sound like a rich, "normal" waveform. Then slowly increase S to 2, which is found when the STRIDE knob is just before 3 o'clock. It will sound hollow, or flimsy. Continue to set S to 4, which is found just before the end. It should sound really sparse, and almost digital.

In addition to these, you can get a harmonic sound with S at 0.5 (at about 9 o'clock), which will sound like it's an octave lower, and at 3, which sounds a little strange and is more difficult to find than 2 or 4.

Next, try experimenting with inharmonic frequencies, setting the STRIDE knob somewhere in between the harmonic values. When you turn the STRIDE knob a bunch, the harmonics will move but the fundamental frequency will not. Because of this difference in *motion*, humans tend to pull these two sounds apart and perceive them separately. To perceptually stick them back together, move the fundamental frequency, especially with a V/Octave controller.

All the parameters of the Harmonic Shift Oscillator can be modulated. Modulation of harmonic stride can produce nice percussive tones, while modulation of harmonic level adds dynamism. Further, the outputs of the Harmonic Shift Oscillator function well as inputs to its various modulation capabilites. In this way, incredibly rich and varying soundscapes can be created from simple controls.

HARMONIC AND INHARMONIC SPECTRA, TUNING, AND BEAT FREQUENCIES

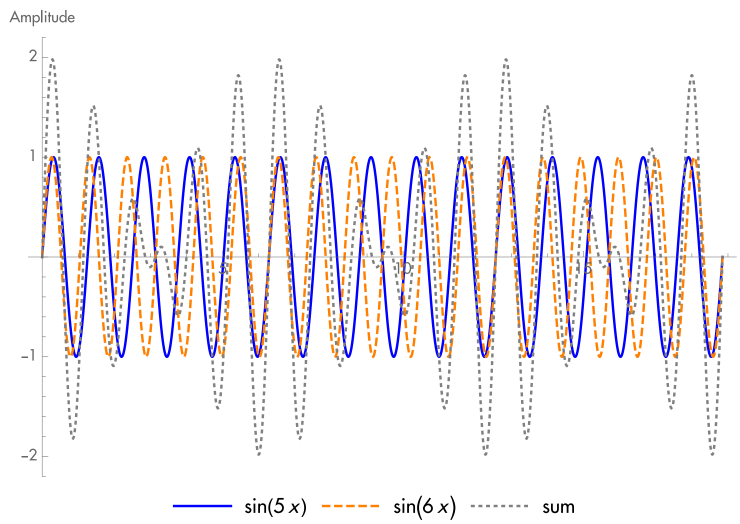

Constructive and Destructive Interference

Any two waves interact with each other to create constructive and destructive interference. That is, when the crests of one wave line up with the crests of another wave, the overall waveform gets louder. Alternately, when the crests of one wave line up with the troughs of another wave, the overall waveform gets quieter. The distance between the crests of two different waves is the *relative phase* of those waves. When two waves of two different frequencies interact with each other, that relative phase changes at a constant rate, as the slower wave keeps falling further behind the faster one. This continual change in phase produces a cyclic change in amplitude, known as a *beat frequency*. This frequency will always be equal to half the difference between the frequencies of the two waves being considered. If we ignore the sign of this waveform, focusing just on the change in amplitude itself, this is just the difference in frequencies. We can express it like this:

$$\sin(a) + \sin(b) = 2\cos((a - b)/2)\sin((a + b)/2)$$

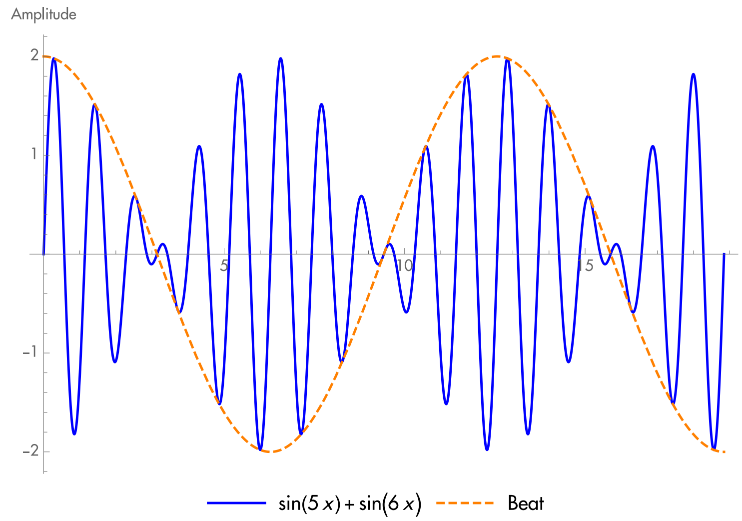

Beat Frequencies

Note that these are two different ways of perceiving the same phenomenon, not two different phenomena. Sometimes we perceive things one way or the other, but more often it's a mixture of both. Thus two very close waveforms might be perceived as a single wave at the average frequency, $(a + b)/2$, with a slow beat frequency at $a - b$ (twice the frequency of of the cosine modulator, since we're just hearing the absolute value of the amplitude). On the other hand, once these waves have been sufficiently separated, we'll notice two dissonant frequencies as well as the beat frequency.

When the two waves being considered are waves in a *harmonic* series, we can rewrite the equation like this:

$$\sin(n\omega t) + \sin((n + J)\omega t) = 2\cos(J(\omega/2)t)\sin((2n + J)(\omega/2)t)$$

Where n and J are two integers.

As we can see, both the beat frequency and average frequency of a harmonic series are themselves frequencies in a harmonic series of half the frequency. Thus, both ways of perceiving these two waveforms give rise to waveforms within the same basic harmonic series.

When the series is inharmonic, that is no longer the case. Beat frequencies other than those present within the spectrum can be perceived. But also, the overall perceived frequency can be difficult for the human ear to determine. This difficulty is what enables some inharmonic sounds to be used outside of a harmonic context (for example, purely rhythmically). Other inharmonic sounds—most of those produced by the Harmonic Shift Oscillator—will still be perceived as tonal, but the relationship of that tone to the fundamental frequency will be complicated. These sounds should be carefully tuned by ear.

PHASE AND THE TWO OUTPUTS

The Harmonic Shift Oscillator includes two different outputs in orthogonal phase with each other. The phase of the ⊢ output trails the ⊥ output by one quarter turn, or 90°. Thus, while these two outputs have similar spectra, they have peaks and troughs which occur at different times.

The perception of phase in humans is complex. While we do not seem to be able to directly detect absolute phase at all, phase differences between our two ears play an active role in our perception of the spatial characteristics of a sound. Further, these phase differences have other effects on the sound that *are* directly audible. First, the different placement of peaks and troughs cause these two sounds to saturate differently at different times, resulting in frequency content that varies with time, but which is different for each output. Second, the interaction of each of these outputs with another signal will produce beat frequencies which are out of phase with each other by one eighth turn. Lastly, because the phase of each individual component is orthogonal, combining the two waveforms does not produce a comb filter effect, but merely a new waveform at an intermediate phase. However, the differentiated application of the two waveforms to a multidimensional delay environment, such as natural or artificial reverberation, results in a more complex comb filter pattern than would the application of either signal alone.

While these two outputs are useful for spatialization, generally that is not best achieved by simply assigning one to each channel of a stereo output. Further, there are other textural uses of this phase difference beyond the needs of spatialization.

THE SPECTRA OF NATURE

It is not the intention of the Harmonic Shift Oscillator to mimic any natural sound. Nevertheless, it has certain characteristics that align with natural acoustic phenomena.

Were there no other factors, a single dimensional resonator would produce a perfectly harmonic waveform proportional to its length, mass, and elasticity. But in nature, there are always other factors. Columns of air have endcaps and bodies that have different elasticity, mass, and volume than their contents. Strings have anchor points which themselves are pulled and prodded while the string vibrates in space. As a result of this, many acoustic instruments produce slightly inharmonic sounds, where the space between harmonics is stretched or shrunk to just less than or just more than 1, or just less than or just more than 2.

Since it's not really possible to dial in an *exactly* integral harmonic stride, the Harmonic Shift Oscillator generally ends up in this barely inharmonic state. Further, just how inharmonic is readily tunable. This makes the Harmonic Shift Oscillator sound "alive" in a way that most oscillators don't (and to mimic this liveliness, other oscillators must be doubled or tripled).

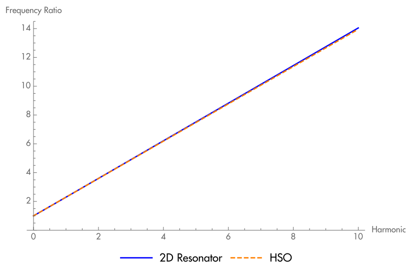

The Frequencies in a Two Dimensional Resonator and Those Producible by the HSO

A two dimensional resonator—a drum head, for example—has vibratory modes proportional to the zeros of Bessel functions of the first kind. Although these resonators can't be exactly modeled as a sum of frequencies proportional to $Sn+1$, we can get remarkably close. Adjusting to place the fundamental frequency at 1, the first 5 modes of a two dimensional resonator would be at: 1, 2.295, 3.598, 4.903, 6.209. Tuning S to exactly match the first harmonic, we get the first five modes for the Harmonic Shift Oscillator as 1, 2.295, 3.591, 4.886, 6.182, which is a difference of +0, +0, +4, +6, and +8 cents, respectively.

As far as the brightness of natural spectra, generally this brightness changes as a function of changing resonance, as in speech, but also as a function of overall energy of oscillation, as in a plucked or hammmered string. While the first can be modeled with a filter, the second is better modeled by modulation of L, where an exponential envelope results in greater rates of decay of the higher than the lower harmonics, while keeping the overall exponential character of that decay for each harmonic.

THE HARMONIC SHIFT OSCILLATOR AND CONVENTIONAL ANALOG SYNTHESIS

Conventional analog synthesis—"subtractive synthesis"—begins with four waveforms: *sawtooth*, *square*, *triangle*, and *sine*. Leaving aside the sine wave, we can express these waveforms according to their spectra:

$$\sum_{n=0}^{\infty} \sin((n+1)\omega t)/(n+1) \quad \textit{(sawtooth wave)}$$

$$\sum_{n=0}^{\infty} \sin((2n+1)\omega t)/(2n+1) \quad \textit{(square wave)}$$

$$\sum_{n=0}^{\infty} \cos((2n+1)\omega t)/(2n+1)^2 \quad \textit{(triangle wave)}$$

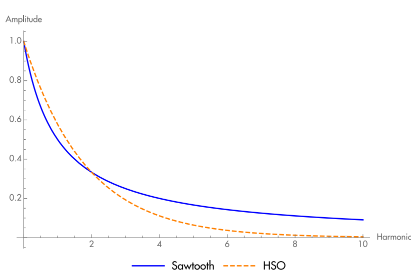

Harmonic Levels of a Sawtooth wave and an Approximation by the Harmonic Shift Oscillator

It is fairly straightforward to reproduce these spectra with the Harmonic Shift Oscillator. First, we can see that the saw contains all harmonic frequencies, while the triangle and square have only odd frequencies. We can produce this with the Harmonic Shift Oscillator by setting S to 1 or 2, respectively. The levels of these harmonics can only be approximated, however. The first five harmonics of a Sawtooth wave $(1/(n+1))$ are 1, 0.5, 0.333, 0.25, 0.2. Matching the first and third harmonics ($L = 1/\sqrt{3} = 0.577$), the harmonic shift oscillator gives 1, 0.577, 0.333, 0.192, 0.111. Whichever harmonic you choose to match, the Harmonic Shift Oscillator has a slightly greater amplitude of prior harmonics, and slightly lower amplitude of subsequent harmonics. This is similar for the spectra of the triangle wave, which has the first five harmonics ($1/(n+1)^2$): 1, 0.111, 0.04, 0.020, 0.012. Again, matching the third harmonic ($L = 0.2$), we get 1, 0.2, 0.04, 0.008, 0.002.

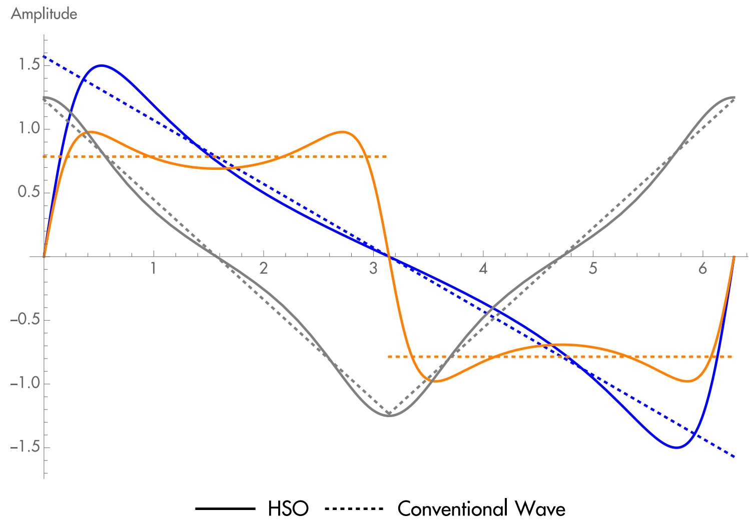

Typical waveforms vs. phase-corrected HSO variants

While these spectra are different enough to give the Harmonic Shift Oscillator a consistently distinct sound, in the practice of subtractive synthesis these traditional waves are rarely used entirely on their own. Instead, the sound source is shaped by mixing together different waveforms, and then passing the result through a chain of filters. In some cases, by using the Harmonic Shift Oscillator one can forego this whole process and arrive at the desired spectrum directly. But, filters offer a substantially different and interesting method of shaping sounds from that provided by the Harmonic Shift Oscillator. Fortunately we don't have to choose. The Harmonic Shift Oscillator is intended to complement rather than replace subtractive methods. It works well with filters and the other tools of subtractive synthesis.