TRIPHASE OSCILLATOR

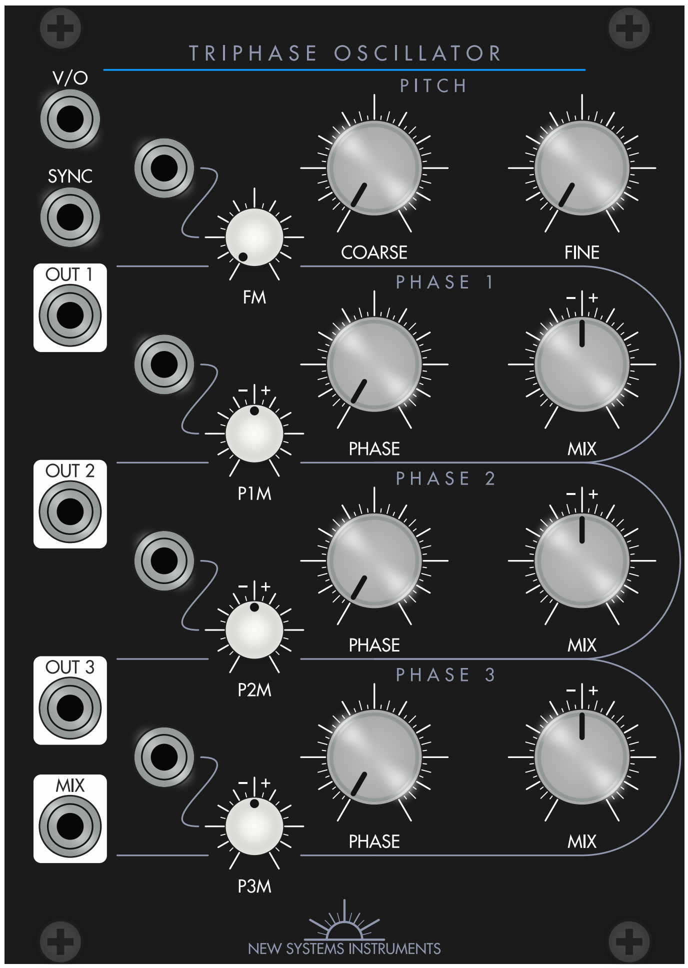

Complex Oscillator

Manual Revision 1.0

NEW SYSTEMS INSTRUMENTS

SPECIFICATIONS

Size	18HP
Depth	25mm
Power Consumption	+12V 50mA −12V 50mA
Tuning Range	1.2Hz–61kHz
Tuning Accuracy	~8 octaves
Output Range (each)	−5V–+5V peak
Output Range (mix)	−8V–+8V peak
Input Impedance	30kΩ (fm), 50kΩ (phase), 100kΩ (v/o), 22kΩ (sync)
Output Impedance	330Ω (out), 3kΩ (mix)
Output Drive	20kΩ+

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 wire 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 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.

OVERVIEW

The Triphase oscillator produces three different sawtooth waves at the same pitch, with independent control over the phase of each wave. Each output can be taken individually, or through a bipolar mixer. When mixed, the harmonics of the three waves reinforce or dampen each other, giving a spectrum with evenly spaced peaks and troughs throughout the frequency range.

Phase cancellation in this pattern—known as a comb filter—is a feature of acoustic spaces, the resonators of acoustic instruments, multiple voices playing in unison, etc. Phase cancellation gives sound a sense of depth, spaciousness, or thickness. While there are many ways to achieve this, each with their own sets of advantages and disadvantages, directly synthesizing phases (the approach of the Triphase Oscillator) has the advantages of precise control, minimal detuning, and maximal depth of the effect.

Phase cancellation can also be used to bring spectra in and out of focus, removing certain harmonic multiples (such as all even harmonics) and reinforcing others. It can also alter the relative level of the harmonics generally, moving from the strong fundamental and $1/n$ slope of a saw to the thin sound where all harmonics are the same volume. Phase cancellation thus forms an important part of the general synthesis toolkit.

HOW TO READ THIS MANUAL

This manual is intended to be an in depth resource for continuing exploration as you continue your journey through sound and synthesis. While the controls of this module are simple, this manual provides a deep analysis of how those controls perform in a wide variety of contexts and applications. It is not at all required for you to read the whole manual before using the module. Read according to your own learning style. I recommend reading over the Overview and Interface section, then going through Quick Start and trying out some patches. From there, browse through the rest of the sections and read what interests you. If you already understand a lot of the underlying theory, you can look through the Model and Parameters sections at the end and find the equations that govern the module's behavior, but these sections assume a lot of prior knowledge.

A note on the mathematics: many sections of this manual contain mathematical equations. These equations aid understanding, but where possible the text was written in such a way that you should be able to skip past them and still understand the basics. So read through with confidence, even if you don't understand the math yet. There are a few sections specifically focused on mathematics where that is not the case, but you can safely skip over these sections, too.

I do recommend you make an effort to learn and understand the mathematics of synthesis. Just like music theory is the language of music, mathematics is the language of synthesis. You can be an excellent musician without knowing any music theory, and an excellent synthesist without knowing any mathematics. But it's very difficult to speak precisely about music without music theory, and it's very difficult to speak precisely about synthesis without mathematics. Think of it as a tool to help you learn from others and teach them in turn.

INTERFACE

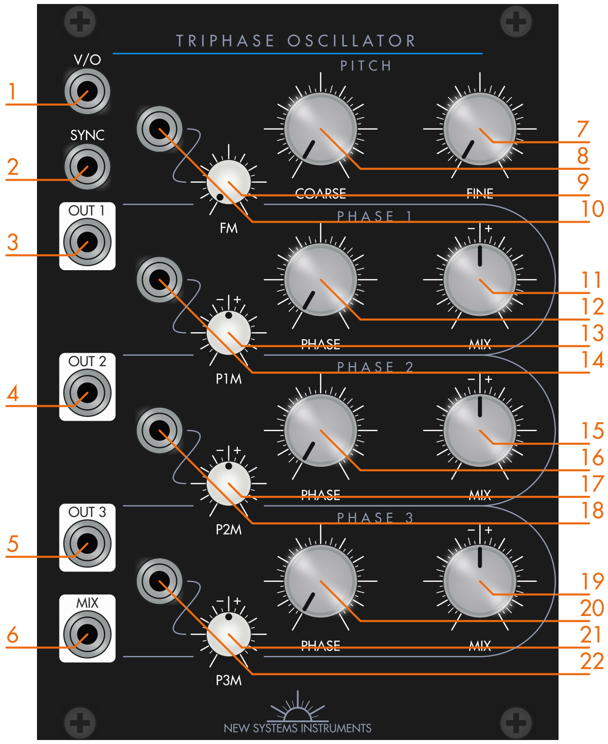

1. Volt per octave input – Controls the oscillator frequency with an exponential scale.

2. Sync input – The oscillator is reset when the sync input rises above ~ 1.2V. Note that if one of the phases would occur after the oscillator resets, that phase will be silent.

3. First wave output – An independent output for the PHASE 1 wave.

4. Second wave output – An independent output for the PHASE 2 wave.

5. Third wave output – An independent output for the PHASE 3 wave.

6. Mix output – A mix of all three waves.

7. Fine tuning knob.

8. Coarse tuning knob.

9. Frequency modulation attenuator – An attenuator for the frequency modulation signal. Note that to maximize the range of this control, this is an attenuator rather than an attenuverter, as with the other CV inputs.

10. Frequency modulation input – Linear CV control over the pitch.

11. PHASE 1 mix – Controls the amount and polarity of the PHASE 1 wave in the mix output.

12. PHASE 1 phase – Controls the relative phase of the PHASE 1 output, with a range of just under -180°–+180°

13. PHASE 1 modulation attenuverter – An attenuverter for the phase modulation signal for the PHASE 1 wave.

14. PHASE 1 modulation – This signal adjusts the phase by 36° per volt (180° for 5 volts, 360° for 10 volts), with the attenuverter all the way up. If the knob and CV would push the phase beyond ±180°, the phase gently folds back towards the center.

15–18. Controls for the PHASE 2 wave.

19–22. Controls for the PHASE 3 wave.

QUICK START

Pulse Width Modulation: Set PHASE 1 MIX all the way right, PHASE 2 MIX all the way left, PHASE 3 MIX centered and take the MIX output. The relative value of the PHASE for each wave determines the width of the pulse. To modulate, apply CV to one of the PM inputs and adjust the attenuverter. For a sound with less pitch change in the fundamental, apply the same CV to both PM inputs, and adjust the attenuverters in equal and opposite directions.

Supersaw: Set all MIX knobs right, center the PHASE controls, apply LFOs liberally to the PM inputs, and take the MIX output. Since phase is relative, you can leave one of the phases unmodulated, although modulation of all three will produce a slightly more complex sound.

Stereo Chorus: In a stereo mixer, put OUT 1 into both channels, OUT 2 into the left, and OUT 3 into the right. Apply subtle modulation to all three phases. To increase the width, increase the phase separation between PHASE 2 and PHASE 3.

Octave: You can get an octave signal in two ways: first, set the PHASE 1 and PHASE 2 MIX all the way right, with PHASE 3 MIX centered. Adjust the two phase knobs to be about 9:30 and 2:30 (just above horizontal). Nudge the phase until you can only hear the octave. To reach the octave in another way, from the above settings adjust the PHASE 1 MIX to 3 o'clock, and the PHASE 2 MIX to 9 o'clock. This should give you a square wave. Then set the PHASE 3 MIX all the way right, with the PHASE matching the PHASE 2 wave.

Third Harmonic: To get only the third harmonic, an octave and a fifth above the oscillator rate, set all MIX knobs right, then set the three PHASE knobs to noon, 9 o'clock, and 3 o'clock.

PHASE AND COMBINING SINE WAVES

It is recommended that you read Principles of Sound and Music on nsinstruments.com, or review it if you need to, particularly the sections on "Anatomy of a Wave" and "Phase, Delay, and Mixing." This section will just be a quick reminder on the principles involved.

When two waveforms meet, they constructively or destructively interfere according to whether peaks of the same polarity meet, or peaks of opposite polarities meet. For most waves, this interference changes the shape of the wave and the spectrum of the sound. But two sine waves of the same frequency always create another sine wave with some phase and amplitude determined by the phase and amplitude of the other waves. That is, with sine waves, the shape is not affected, only the amplitude and phase.

Because of this, it is often easier to talk about how the combination of spectra at different phases produces a new spectrum with different amplitudes for different frequencies, than to try and figure out how alterations in the shape of some wave will sound. Additionally, for sine waves a 180° phase difference and a reversal of polarity are the same thing. We can, then, state the effect of combining two different spectra pretty simply.

When sine waves at the same amplitude and same phase meet, the waves reinforce each other and we get a sine wave of twice the amplitude at the same phase and frequency. When sine waves at the same amplitude but 180° apart meet, the waves cancel each other perfectly to zero. In between these two phases, the amplitude of the resulting wave grows gradually from zero to double. Note, however, that the effect of relative phase on amplitude is not linear. Combining sine waves 90 degrees apart will produce an amplitude of 1.4 (i.e. $\sqrt{2}$), rather than 1, and 1 is reached when the phase difference is 120°.

When sine waves of unequal amplitude meet, the effect of phase difference is lessened. The most intuitive way to think about this is to imagine that the equal parts of the sine waves are mixed together, and then the remainder is added to this mix. While in reality that remainder will again be different in phase from the mixed sine waves, and so the process would repeat itself, in practice not much is lost by thinking about the remainder as just a little extra amplitude that is mixed in.

PHASE AND SPECTRUM

A sawtooth wave can be analyzed as a spectrum of sine waves where every sine wave is the same phase, and each harmonic is $1/n$ the amplitude of the fundamental frequency. That is, the first harmonic, the fundamental, has an amplitude of $1/1$, the second has an amplitude of $1/2$, the third $1/3$, etc. We can express this same thing with the following equation:

$$\mathrm{saw}(\omega t) = \sum_{n=1}^{\infty} \sin(n\omega t)/n$$

When we combine two sawtooth waves, this is the same as combining their spectra. Since in the last section we already saw how to think about combining sine waves of difference amplitudes and phases, the only remaining piece of this puzzle is to see what a shift in the phase of the whole sawtooth wave does to the phase of each individual component of the spectrum. This is easily derived.

We can represent a sawtooth at a different phase by shifting the whole sawtooth in time by some portion of a cycle, that is, by replacing ωt with $\omega t + \phi$. This gives us the modified equation:

$$\mathrm{saw}(\omega t + \phi) = \sum_{n=1}^{\infty} \sin(n(\omega t + \phi))/n$$

We can then rewrite the latter term to pull the ϕ out of the parenthesis, giving:

$$\sum_{n=1}^{\infty} \sin(n\omega t + n\phi)/n$$

So in our spectrum, each harmonic has a phase difference consisting of the phase difference of the saw as a whole, multiplied by the harmonic number. This makes intuitive sense when we think about it. Since the second harmonic is twice the frequency of the first, its phase takes up half the time, and so a small phase difference is twice as significant. Since these waves are periodic, a phase difference of 360° is the same as 0°, and the differences loop around.

As an example, consider a phase difference of 45°. This would give us harmonics with the following phase differences:

First Harmonic	45°
Second Harmonic	90°
Third Harmonic	135°
Fourth Harmonic	180°
Fifth Harmonic	225°
Sixth Harmonic	270°

Eighth Harmonic	360°
Ninth Harmonic	45°
...	

In this particular sequence, we can see that by combining waves 45° apart, we would get a somewhat reinforced first and second harmonic, a somewhat reduced third harmonic, a completely canceled fourth harmonic, a somewhat reduced fifth harmonic, reinforced sixth, seventh, and eigth harmonics, and then the same repeating pattern for the remaining harmonics.

When we change the polarity of the wave, this adds 180° to each harmonic. This would transform the above table as follows:

First Harmonic	225°
Second Harmonic	270°
Third Harmonic	315°
Fourth Harmonic	360°
Fifth Harmonic	45°
Sixth Harmonic	90°
Seventh Harmonic	135°
Eighth Harmonic	180°
Ninth Harmonic	225°
...	

While the increase in phase difference between any two harmonics is independent of the polarity, the polarity shifts the sequence as a whole. Consider two waves that are 180° apart, but opposite in polarity. That would give us the following phase differences of the harmonics:

First Harmonic	0°
Second Harmonic	180°
Third Harmonic	0°
Fourth Harmonic	180°
Fifth Harmonic	0°
...	

As you can see, this cancels all the even harmonics, leaving the odd harmonics untouched. With sawtooth waves, this will produce a square wave. These completely canceled harmonics are the reason this wave shape sounds "hollow," like it is missing something. It is!

Any phase difference at all produces a series of phase difference in the spectra of the waves such that frequencies cancel, reinforce, cancel again, reinforce again, etc. Because the shape of this graph resembles a comb, the effect on the spectrum of combining waves of different phases is known as a comb filter. Here are the reductions and enhancements on the first 100 harmonics with various phase differences:

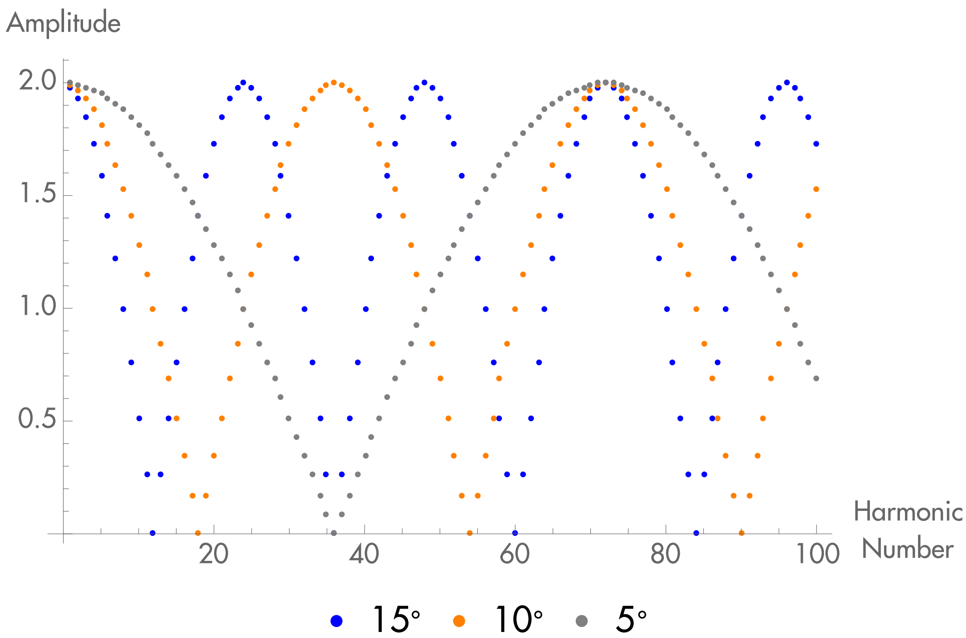

The relative levels of the first 100 harmonics when the phase difference is 15°, 10°, and 5°.

While the shape of the comb filter doesn't change, as the phase difference becomes smaller, the tines of the comb grow wider, with the limit at zero, where the tines are infinitely wide and all harmonics reinforce, and 180 degrees, where the comb has notches for every even harmonic. Viewed in a more conventional frequency response, the comb looks like this:

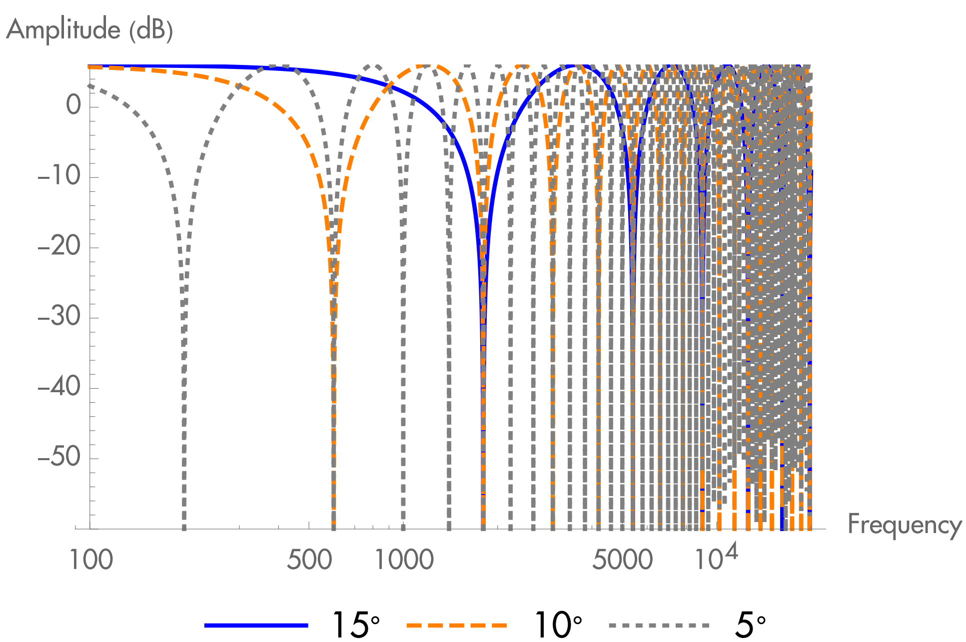

The filtering effect of combining sawtooth waveforms at 100Hz.

FREQUENCY AND DYNAMIC PHASE

Phase and frequency are related, such that a steadily changing phase is exactly the same thing as an increased frequency. Intuitively, imagine a sawtooth wave that is completely stopped—its frequency is zero. If we now advance the phase of sawtooth at a constant rate, we are effectively "playing back" the sawtooth, and the rate of playback determines the frequency. Similarly, when the sawtooth already has a frequency, "playing back" the already moving sawtooth will increase the frequency, while decreasing the phase steadily ("playing back" in reverse) will decrease the frequency.

The effect of phase on frequency is often a big part of the sound of phase modulation. For example, it provides the detuning sound of a chorus effect. Nevertheless, sometimes it is desired to have phase modulation without affecting the frequency. This will be possible when we are able to change the *relative* phase of two waves, while keeping the *absolute* phase of the combined wave constant.

Note that absolute phase does not exist, strictly speaking. That is, there is no aspect of the universe that allows us to differentiate between two times except through the *relative* timing of two different events. Asking when something occurs is a meaningless question unless something else has or will occur. Nevertheless, perhaps counterintuitively, while absolute phase does not exist, absolute change in phase does exist. In fact, the name of absolute change in phase is just *frequency*. Thus, on the one hand, we can only measure the relative phase of two signals. For example, we can't say that one is at phase zero and one is at phase 90°; we can only say that they are 90 degrees apart. But, given that they begin at the same frequency, if one of them is changing in phase, we know which one is changing—it's the one with an increased or decreased pitch. If you think there is something either fishy or profound about a universe where the absolute change of a value is detectable while the absolute value itself is not detectable, you are correct.

To understand the effect on frequency of changing phase, we can add a rate of phase change to our equation. Although in practice the rate of phase change will generally be determined by a moving LFO that changes speed and direction periodically, we will simplify by imagining a steadily changing phase, at some number of degrees per second. If we can can minimize the effect of this steadily changing phase, this will also minimize the effect of an LFO, envelope, or any other signal.

The following is the equation for the harmonics of a sawtooth oscillator, now with an element that represents a steadily changing phase:

$$\mathrm{saw}(\omega t + \phi + \varphi t) = \sum_{n=1}^{\infty} \sin(n(\omega t + \phi + \varphi t))/n$$

We can put this in two different useful forms:

$$\sum_{n=1}^{\infty} \sin(n\omega t + n(\phi + \varphi t))/n = \sum_{n=1}^{\infty} \sin(n(\omega + \varphi)t + n\phi)/n$$

The first equation shows us the effect the changing phase has on the phase of each harmonic. That effect is what we would expect to see. The constant and the changing phase difference add together, and then behave like they normally would for that harmonic. The second equation, however, shows us how the changing phase affects the frequency. For example, a phase changing at the rate of 360° per second will detune the wave by 1Hz.

What matters, however, is the phase of the combined waveforms. When two sine waves of equal amplitude combine, the phase of the resulting wave is the phase in between the two phases. In this case, then, you can remove the absolute motion of the resulting phase by having each phase move at an equal rate in opposite directions. When the amplitudes are not equal, the overall phase will stay closer to the louder wave, and consequently if the phase of the louder wave moves less than the phase of the quieter wave, the effect on frequency will be minimized. Note, however, that the motion can be perfectly canceled only when the amplitudes are exactly equal, or of equal magnitude and opposite polarity.

Given the folowing two harmonics at any amplitudes, moving at arbitrary rates:

$$a\sin(n\omega t + n\phi_a + n\varphi_a t) + b\sin(n\omega t + n\phi_b + n\varphi_b t)$$

The frequency of the resulting harmonic will be increased by:

$$n\frac{a^2\varphi_a + b^2\varphi_b + ab(\varphi_a + \varphi_b)\cos(n(\varphi_a - \varphi_b)t)}{a^2 + b^2 + 2ab\cos(n(\varphi_a - \varphi_b)t)}$$

While this exact equation is unlikely to be useful while using the Triphase Oscillator, we can still learn a few things from it. The $\cos(n(\varphi_a - \varphi_b)t)$ terms indicate that even with a constant motion of phase in each signal, the changing phase relationship between the two signals will cause the overall frequency shift to change over time. This extra term prevents us from eliminating frequency shift, but we can still minimize its effect, according to a some concept of what minimizing its effect looks like.

First, note that if we change the phase of both waves at the same rate and in the same direction, there is no frequency shift due to their interaction, and all the frequency shift is due to the changing phase. While this does not modulate the phase *difference*, and so it is not usually what we want, it can nevertheless be useful to limit interactions when mixing with a third wave.

Second, we might want to ensure that the DC component of this frequency offset is zero. Counterintuitively, this occurs when the entirety of the changing phase shift between the two waves takes place in the wave with the least amplitude, and the wave with more amplitude has no changing absolute phase shift.

However, minimizing the DC component does not minimize the excursion from center pitch. This is because the change in the pitch is not symmetric. It will stay fairly close to zero on one side, and then briefly shoot far to the other side. In order to minimize these excursions, we would set each phase modulation control as follows:

$$\varphi_a = \frac{b^2}{b^2 - a^2}\varphi_{a-b}, \varphi_b = \frac{a^2}{b^2 - a^2}\varphi_{a-b}$$

In practice, this just means that we modulate each wave inversely proportionally to its magnitude in the mix, such that the majority of the phase difference is placed on the smaller waves.

With three moving waveforms, or even two moving and one fixed waveform, the effects of a changing phase on the pitch can be quite complex. It can be characterized according to the following equation:

$$n\frac{a^2\varphi_a + b^2\varphi_b + c^2\varphi_c + ab(\varphi_a + \varphi_b)\cos(n(\varphi_a - \varphi_b)t) + ac(\varphi_a + \varphi_c)\cos(n(\varphi_a - \varphi_c)t) + bc(\varphi_b + \varphi_c)\cos(n(\varphi_b - \varphi_c)t)}{a^2 + b^2 + c^2 + 2ab\cos(n(\varphi_a - \varphi_b)t) + 2ac\cos(n(\varphi_a - \varphi_c)t) + 2bc\cos(n(\varphi_b - \varphi_c)t)}$$

While there's only a couple more constant terms, the cosine terms have exploded to include all three ways these waves can interact: a with b, a with c, and b with c. Adding one more wave triples the interactions. Note that while the effect of n on the constant part of the equation is still just a multiple (just sufficient to keep all detuning harmonic), n will also affect how the three cosine waves combine, such that at certain times some harmonics might be more detuned than others.

With three interactions, we have three phase differences to distribute amongst the phases of each wave. As such, it is arduous to even think through how to adjust the absolute phase changes while preserving the relative phase changes. It's best to just note a few inexact principles. As these principles also hold for the case of two waves, we provide it as a summary of the entire section.

MINIMIZING FREQUENCY SHIFTS: SUMMARY

–When two waves are mixed with equal magnitude, regardless of polarity, the effect of a relative phase change on frequency can be eliminated by having each wave change its phase in equal and opposite directions.

–When two waves are mixed with unequal magnitude, the constant offset of the frequency change is eliminated by assigning all of the phase change to the wave mixed at a lower volume.

–When two or more waves have the same phase change in the same direction, the change in pitch is unaffected by the interaction between the waves, and is only dependent on the rate of phase change.

–When two or three waves are mixed with unequal magnitude, the effect of a relative phase change on frequency is minimized by distributing the phase change such that waves mixed at a lower volume get more of the phase change and waves mixed at a higher volume get less.

MODEL AND PARAMETERS

The Triphase Oscillator provides three outputs with phase control, according to the following equations:

$$
\begin{aligned}
y_1 &= \mathrm{saw}(\omega t + \phi_1 + \varphi_1) \\
y_2 &= \mathrm{saw}(\omega t + \phi_2 + \varphi_2) \\
y_3 &= \mathrm{saw}(\omega t + \phi_3 + \varphi_3)
\end{aligned}
$$

These are mixed together for the mix output, according to the following equation:

$$y = m_1 y_1 + m_2 y_2 + m_3 y_3$$

ω is proportional to FM, COARSE, FINE, and V/O via the following relation:

$$\omega \propto (FM + BASE) 2^{COARSE + FINE + V/O}$$

Where the FM input gives linear modulation up to a maximum of between 0 and double the base frequency. COARSE ranges over approximately 13 octaves, and FINE ranges over 3 octaves. With both centered, the pitch is approximately C4, or 260Hz.

ϕ is a constant phase set by each PHASE knob. These range from approximately -180° to 180°. φ is an additional phase modulation signal, coming from the P1M, P2M, and P3M CV inputs and attenuverters. At maximum range, each PxM input affects the phase at a rate of 36° per volt, or ±180° for a standard 5 volt peak wave.

A separate output is provided for each y_x wave, as OUT 1, OUT 2, and OUT 3. These waveforms are at a 5V peak amplitude.

The mix output is a sum of all three waves, affected by the m_x parameter via each MIX knob. These are bipolar knobs ranging from -1 to 1.