Sound Basics

Koblo

Understanding Sound and Synthesizers

Version 2.0
Printed in Denmark.

Contents

Basics

Basics

Welcome

This book explains basic concepts about sound and synthesizers.

Read it if you would like a better understanding of what is going on inside synthesizers.

This book is provided as a public service to planet Earth.

Keep vibrating!

What is sound?

All the sound you hear is actually movement of air particles as they swing back and forth, following the motion of the object producing the sound: the strings of your guitar, your voice or a loudspeaker.

The movements of the air particles determine the type and timbre of the sound.

In analogue electronic form, sound is represented by oscillating electric voltages corresponding to an acoustic sound.

In digital form, sound is represented by numbers.

Samples

The process of recording a real sound to your computer is called sampling. The analogue sound waves are first converted into digital numbers by an ADC (Analog-to-Digital Converter) which produces a stream of numbers increasing and decreasing in the same way as the original sound wave. The rate at which the numbers are generated is called the sample rate. The amount of information used in each number determines the resolution. A higher sample rate and resolution gives better sound quality.

For example, the sampling rate of an audio CD is 44.1 kHz, and the resolution is 16 bits.

The sequence of numbers in a sample are usually stored either in the computer memory or on a hard disk. You can play back the sample using a sound editor, an audio sequencer or - even better - one of Koblo’s digital synthesizers.

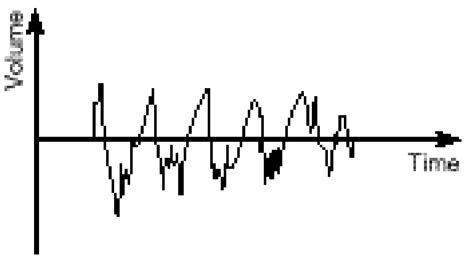

An audio sample

Wavetables and waveforms

If you cut out a very short piece of a sample containing only one period of the original sound, you get a wavetable. Such a wavetable can be played back looping it over and over at different speeds giving you different pitches. The shape of the wave contained in a wavetable is called a waveform.

A wavetable

Frequency

By definition, oscillation means something that repeats. While a guitar string repeats its movement many times while playing just one note, we say that the frequency is high.
The time is takes a string to do one complete vbration back and forth (to complete one cycle) is called the period of the tone.

The related frequency is the number of periods the string produces per second. It is measured in cycles per second (cps) or Hertz (Hz). In other words, higher frequency means faster repetition.

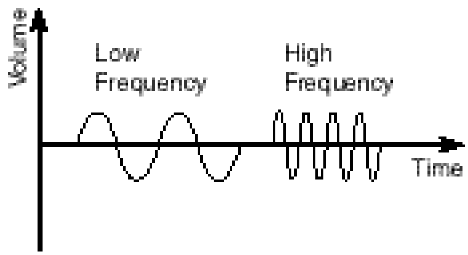

Low and high frequencies

The pitch of the sound is a musical term for frequency often expressed as a note name: g sharp, c flat, etc. Sometimes MIDI note numbers are used, e.g. 32, 33, etc.

The human hearing ranges from about 20 Hz to about 20000 Hz. Frequencies below 20 Hz are considered ‘low frequency’.

Octave, semitone, cent

A **semitone** is the smallest difference between notes on guitars, flutes, pianos, MIDI keyboard and most other western instruments.

A **cent** is just 1/100 of a semitone, so it is an extremely small interval that gives you high precision of tuning notes and instruments. Some synthesizers require you to tune more precisely, in so called cents. Do not expect to be able to hear a pitch change of only one cent - a quarter tone is 50 cents. Try this first.

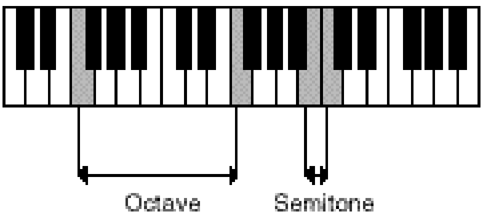

Octave and Semitone

One **octave** is 12 semitones. Raising or lowering the frequency by one octave always corresponds to doubling or halving the frequency. That is why two notes playing one octave apart sounds especially harmonic.

Unipolar and Bipolar signals

A sound signal can either be unipolar or bipolar, depending on the range of the signal. The signals otherwise look the same.

A signal is considered **unipolar** if it is always positive or zero. In synthesizers an unipolar signal often means that the signal has a range between 0 and 1.

A signal is considered **bipolar** if it can be both negative, zero and positive. In common synthesizer language a bipolar signal often means a signal with a range of -1 to 1.

Note that sound signals are most often bipolar. Unipolar signals are most useful as control signals.

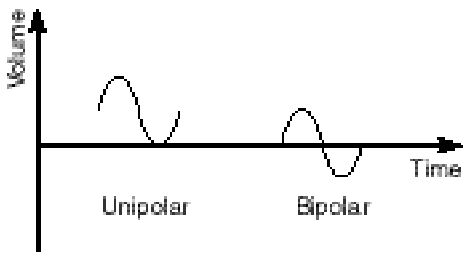

Unipolar and bipolar signals

Oscillators

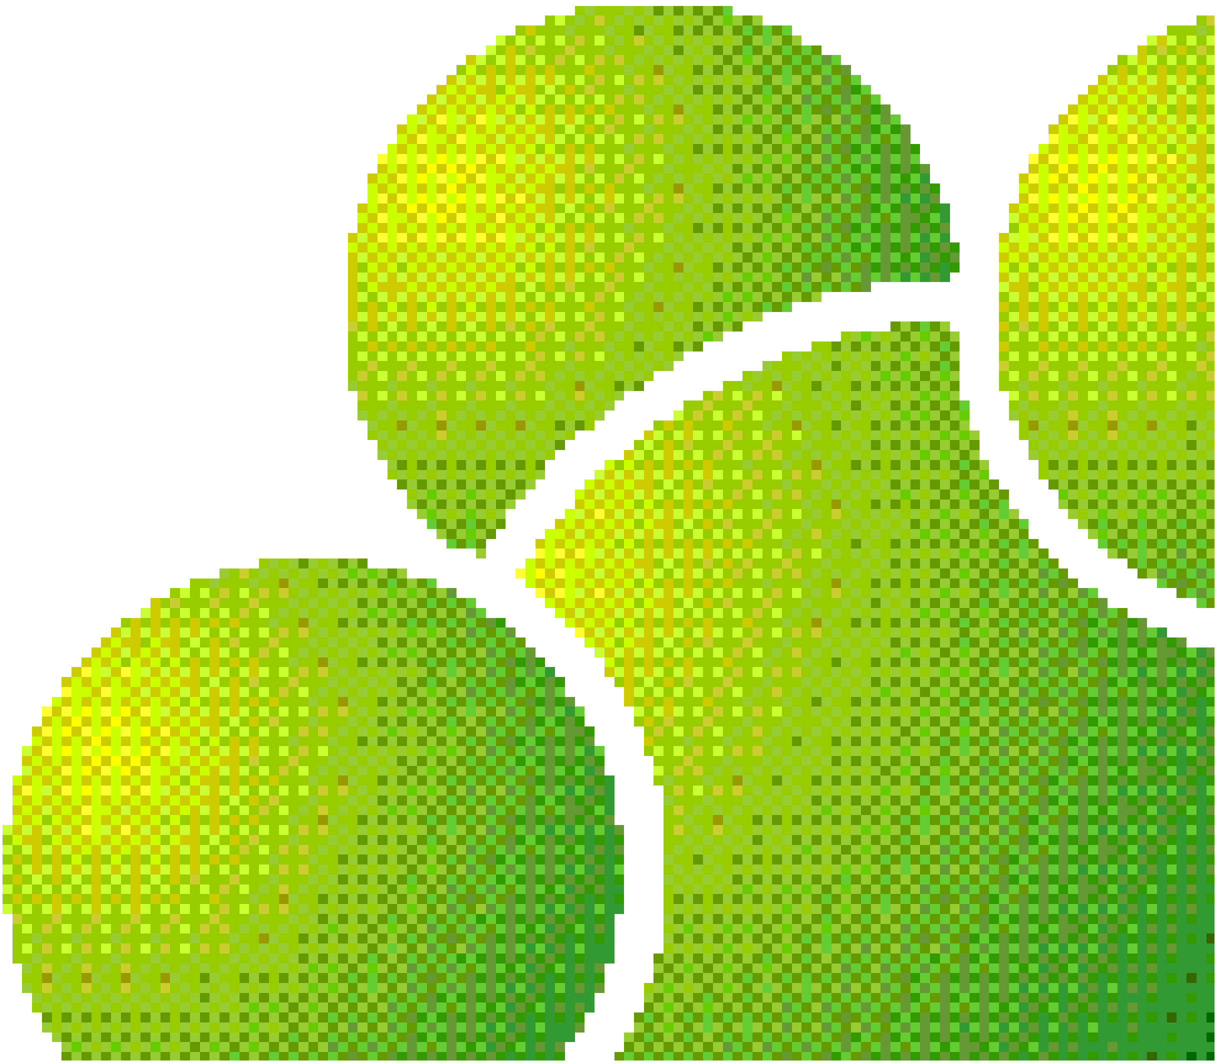

Oscillators

Introduction

The basic building block in synthesizers is the **oscillator -** sometimes called a tone generator. In digital synthesizers it produces a stream of oscillating digital numbers. In digital synthesizers it is often implemented using a wavetable that you can control in various ways: amplitude (volume), frequency (pitch), and phase (micro timing).

Sine waves

The most simple form of oscillation is the sine wave, also called harmonic oscillation The sound of a sine wave is very ‘pure’ and ‘simple’.

The sine wave is produced naturally in many places, for example in a swinging pendulum.

All sounds may be considered to be composed of sine waves at different frequency and amplitude.

Partials

If we take a harmonic, musical sound and analyze the frequency contents, we find that it most often contains a lot of different frequencies called partials. The lowest frequency is called the **fundamental frequency** or the base frequency.

All the other frequencies are multiples of the fundamental frequency, which means that, if the fundamental frequency is 100 Hz, the others will be 200, 300, 400, 500 Hz, ... etc.; only their amplitudes will vary.

Different timbres result from different patterns of amplitudes of the individual partials.

Spectrum

A graph showing the amplitudes of the various frequencies is called the spectrum of the sound.
Note that harmonic sounds contains only frequencies that are integer multiplies (2,3,4,5,6,7....) of the base frequency (1).

Common waveforms

Sine waves contain only the base frequency. That is why it sounds ‘pure’ or ‘simple’.

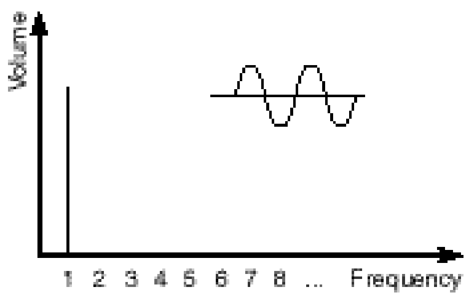

Sine spectrum

Sawtooth - sometimes just called Saw **-** contains all partials, odd and even. The amplitudes of the partials decrease with frequency. Because it contains all partials, it sounds ‘rich’ and ‘buzzy’.

Triangle - contains only odd partials, like the square wave. However, the amplitudes of the partials decrease with increasing frequency much faster than in the square. This is why it sounds much softer that the square and more like a sine wave.

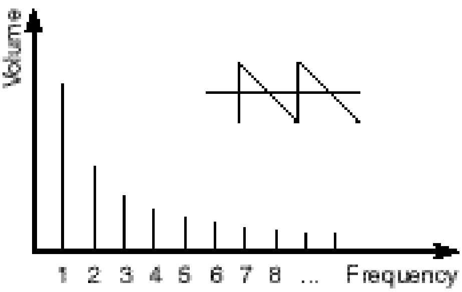

Saw spectrum

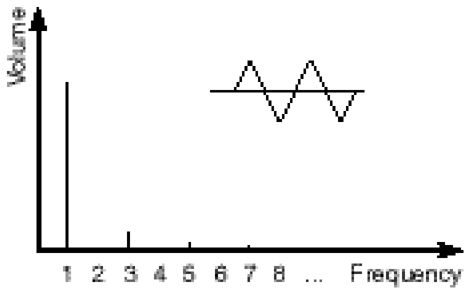

Triangle spectrum

Square - sometimes called Rectangle **-** contains only odd (and no even) partials. They decrease in amplitude with frequency, just as with the sawtooth.
Because it contains no even partials, it has a ‘hollow’ sound.

Noise - or more precisely White Noise **-** contains all frequencies changing randomly, all the time. It has no partials since the frequencies are not even multiples of each other.
It also has no base frequency.
Pink noise has less high frequencies than white noise; brown noise has even less.

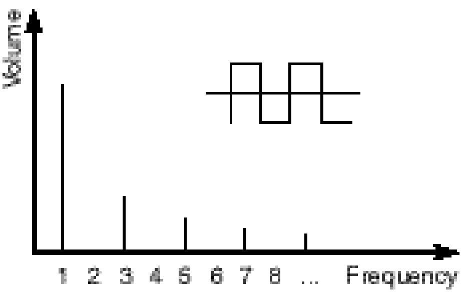

Square spectrum

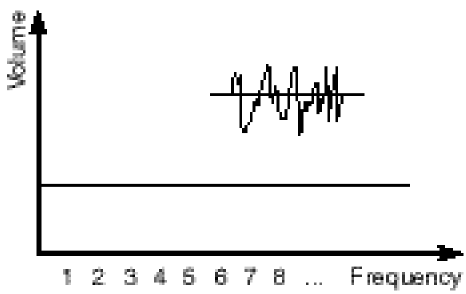

White noise spectrum

Filters

Filters

Introduction

A sound filter works by changing the frequency contents of the input signal by filtering away or amplifying certain frequency ranges. It is the most important module on many synthesizers.

There are an almost unlimited number of different filter types, but they can be broken down to four main types:

Low Pass filter

Low frequencies will pass through the filter accordingly. Higher frequencies will be attenuated. The frequency at which the filter starts to attenuate is called the cut-off frequency. The rate of attenuation is called the slope of the filter and is expressed in dB/octave (or sometimes in number of poles in the filter - see below).

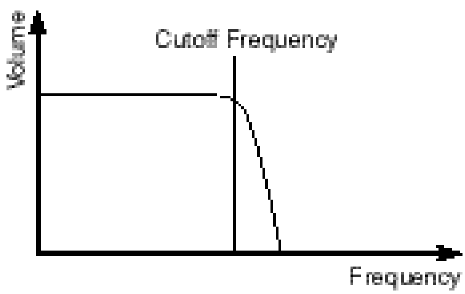

Low pass filter

High Pass filter

This is the opposite of the low pass. High frequencies pass through, and low frequencies are attenuated.

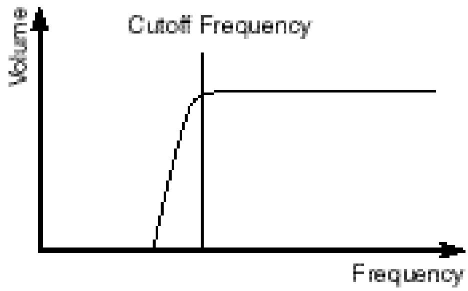

High pass filter

Band Pass filter

In this filter type only a range of frequencies passes thru the filter. The filter is characterized by the band width and center frequency. Frequencies above and below the center frequency (adjusted with the band width also called window) are attenuated.

The smaller the band width, the fewer frequencies go through the filter without being attenuated.

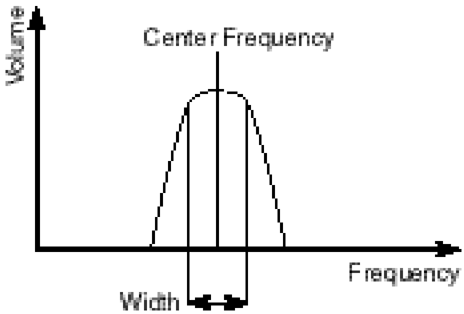

Band pass filter

Band Stop (Notch) filter

This is the opposite of the band pass filter. Frequencies in the band near the center frequency will be attenuated.

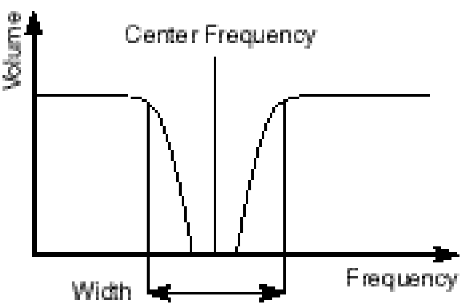

Band stop filter (Notch filter)

Resonance

Some filters have the ability to amplify certain frequencies. This is called resonance.

In resonant low / band / high pass filters the resonance will amplify frequencies around the cut-off or center frequency.

Some filters allows you to adjust the amount of resonance.

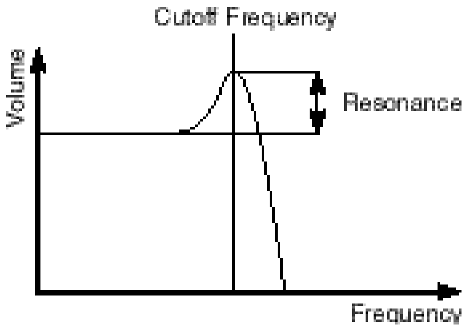

Filter Resonance

Filter Poles

The number of poles in a filter indicates the number of resonance points. If you put both poles at the same center frequency, you get an even stronger resonance.

The number of poles also determines the filter slope or filter roll off , measured in db/octave. A higher number of poles results in a steeper frequency response curve around the cutoff point.

For example, considering a low pass filter, the filter slope expresses how quickly the frequency response falls to zero above the cutoff frequency.

Each filter pole adds a maximum of 6 dB/octave to the filter roll off value. This is why a 4 pole filter is also called a 24dB/octave filter.

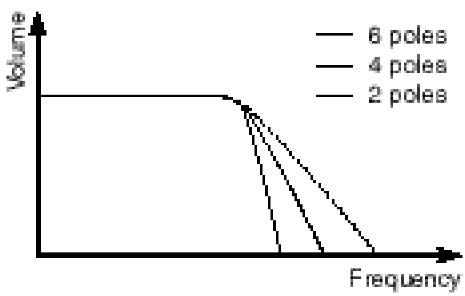

Filter poles

Comb filter

You can think of a comb filter as a whole bank of narrow band pass filters having resonances at multiples of the base frequency of the filter. If you have a resonance at 75 Hz, you will get other resonances at 150, 225, 300, 375, 450, 525, 600, etc. Hz. This will be the equivalent to a harmonic note, and you can actually use comb filters to make noise sounds into harmonic sounds.

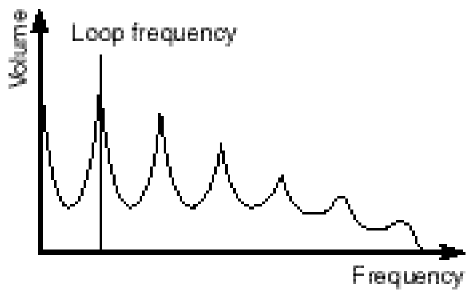

Comb filter

Filter combinations

When you have more than one filter, you can combine them in two ways:

• Parallel (side by side)
The sound goes through all the filters simultaneously, and the outputs of the filters are mixed.

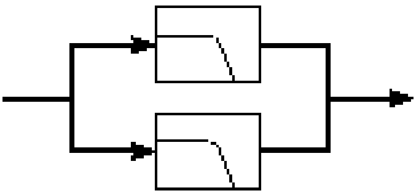

Parallel filters

• Serial (after each other)
The sound goes through one filter first, then the next, etc. Each step can filter a part of the signal away; be careful not to filter everything away.

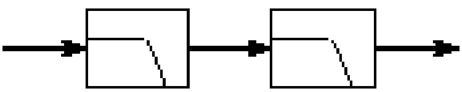

Serial filters

Envelopes

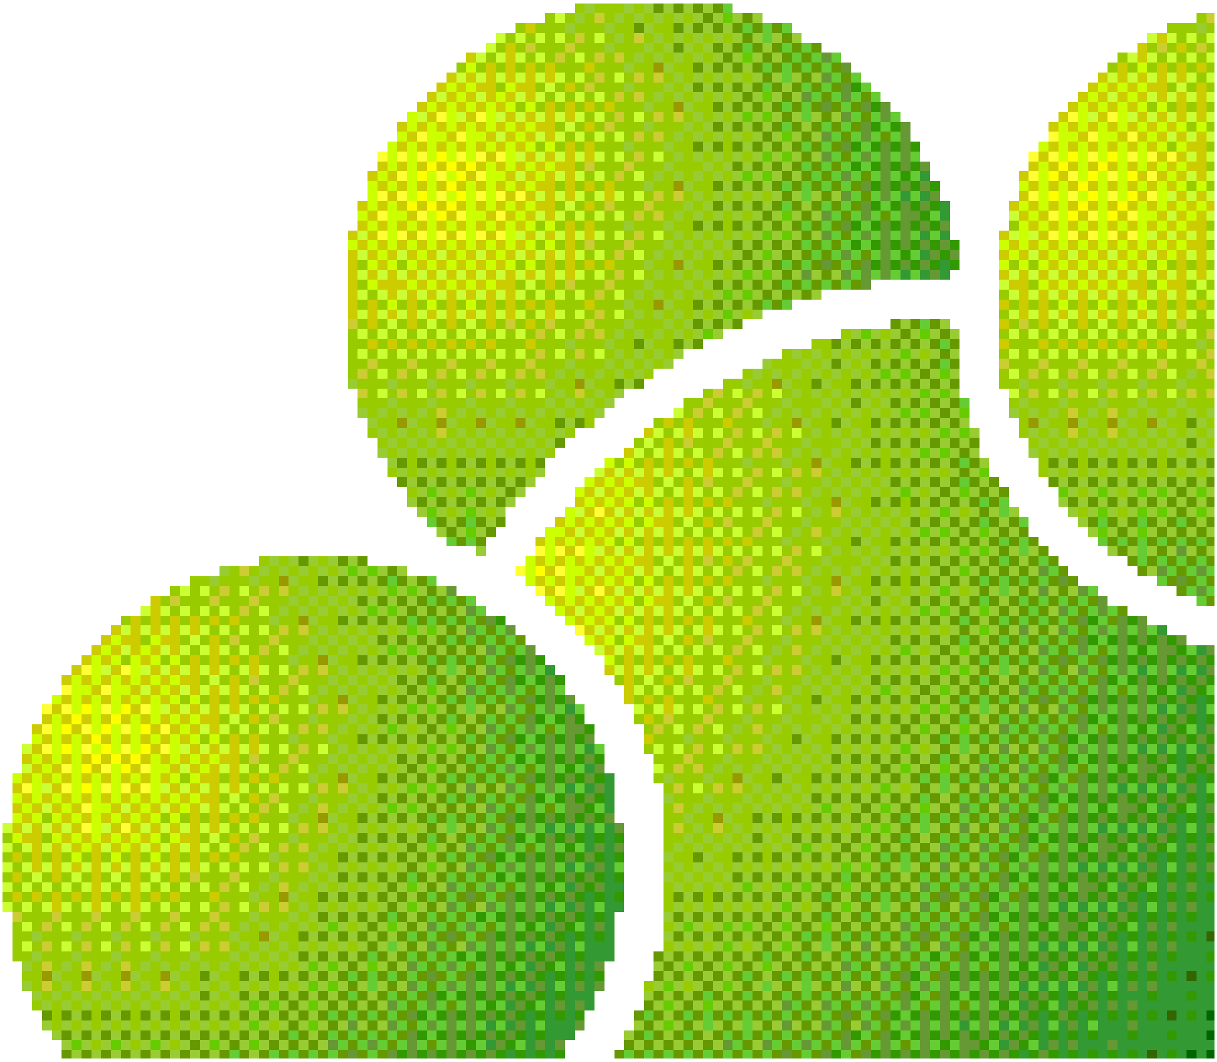

Envelopes

Introduction

An envelope is a unipolar, low frequency signal most often used to control the amplitude of an oscillator or the cutoff of a filter.

ADSR envelopes

Think of the amplitude development of a single piano tone:

• The hammer hits the string - the tone rises to full amplitude.
• The sound stabalizes into a steady note, at a certain level.
• You release the piano key - the tone dies out.

These 4 elements are called Attack Time, Decay Time, Sustain Level and the Release Time, often abbreviated as A, D, S and R. This is why a simple envelope generator is often called an ADSR.

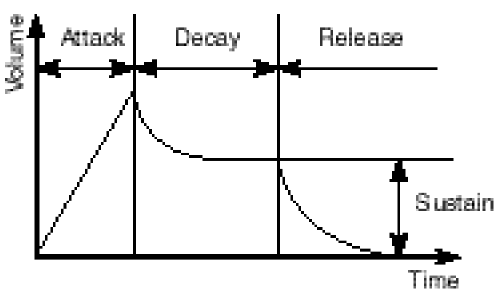

ADSR Envelope

Variations

Envelopes sometimes contain additional stages and may be inverted.

In an inverted ADSR envelope, the output decreases through the attack stage and increases to sustain level during the decay stage. In the release stage, the output from the envelope increases to maximum level.

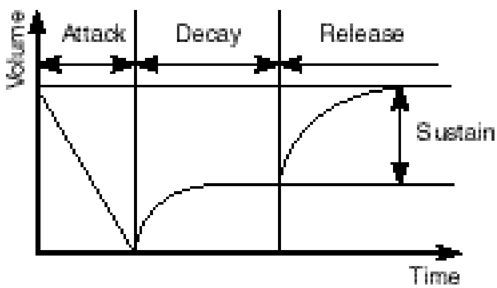

Inverted ADSR Envelope

LFO’s

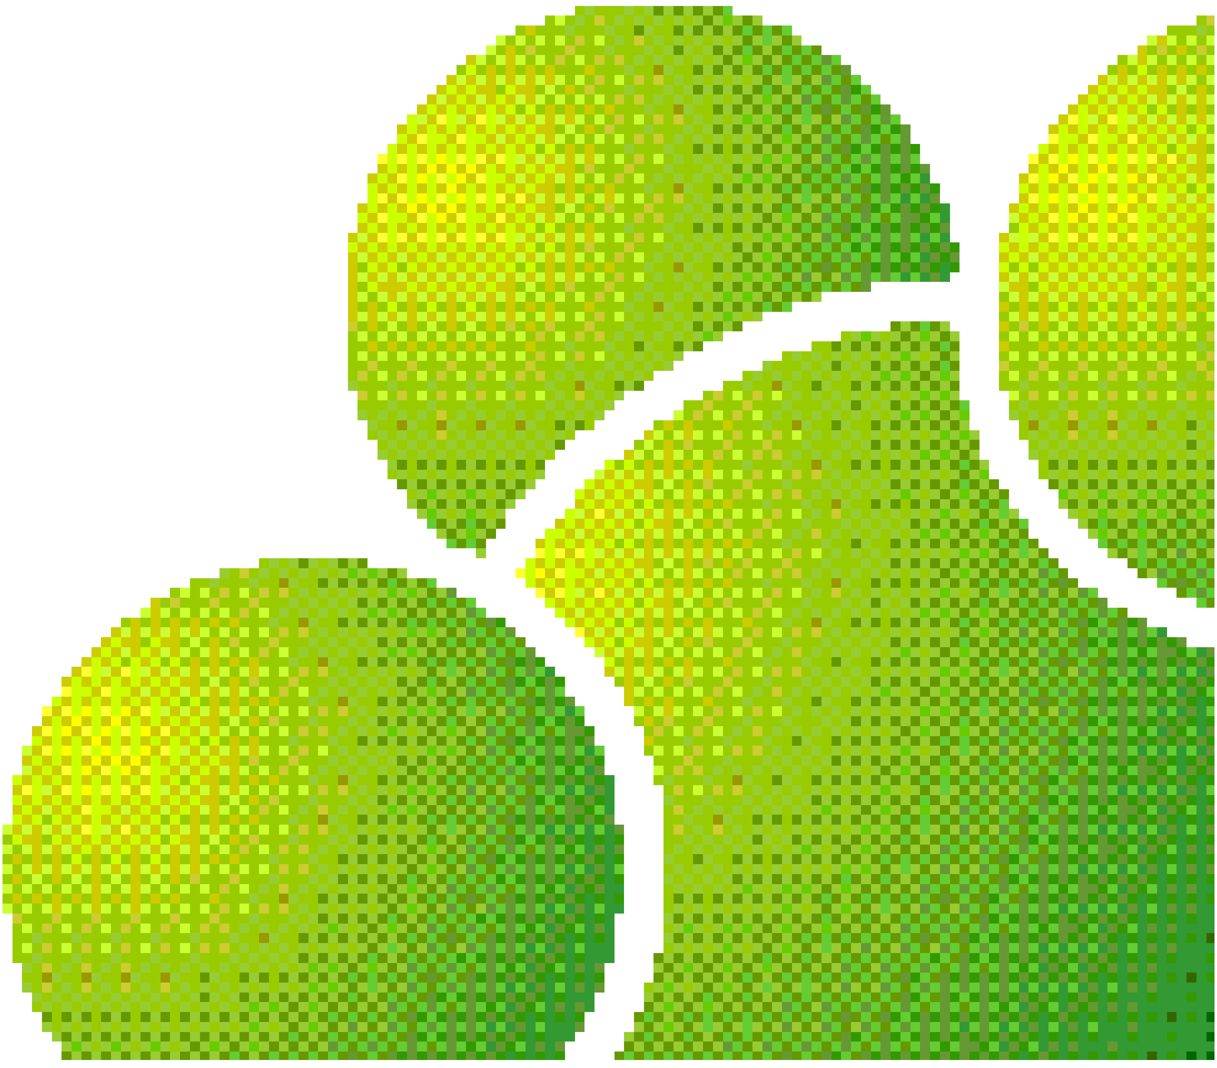

LFO’s (Low Frequency Oscillators)

Introduction

In principle there is no difference between a LFO and a standard oscillator except the way they are used. As the name states, an LFO has a low frequency, making it suitable to control things like vibrato, tremolo, filter sweeps, pan, etc.

Shapes

Typical shapes (or waveforms) of a LFO include Ramp Up, Ramp Down, Triangle, Sine, Random.

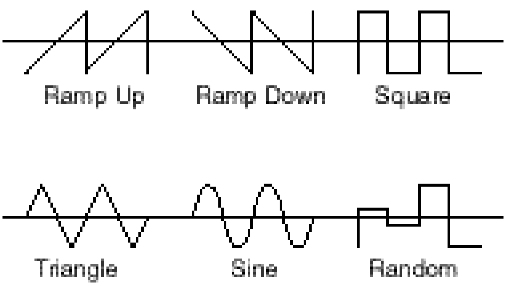

Common LFO Shapes

LFO's can sometimes be set to produce either unipolar or bipolar output. the polarity is useful for modulation of different sounds elements.

For example the unipolar output of an LFO can be used to control the amplitude of the volume, and the bipolar output of the LFO can be used to set the pan position so the sound moves between left and right channel.

Modulation

Modulation

Introduction

In theory, to modulate means to use a source to change something in the destination.

Most modulation elements requires two inputs and generates a new third output depending on the two inputs.

Modulation is in the concept of electronic sounds the way to connect different sound manipulating elements in a synthesizer to get an new often more interesting sound.

Using an envelope to modulate the amplitude or a LFO to modulate the filter cutoff, are examples of using low frequency sources. Often the value of MIDI controllers, pitch bend, wheel or aftertouch, can be used as low frequency modulation sources.

FM and AM are examples of using a high frequency modulation source.

AM (Amplitude Modulation)

Amplitude modulation is the process of changing the amplitude of sound. When you control the amplitude of an oscillator with the output from another oscillator, you actually perform amplitude modulation. For example, if you use a sine wave with a frequency of 200 Hz to control a sine oscillator at 800 Hz, a magical thing happens: you get sound that is a mix of three sine waves at 600 Hz, 800 Hz, and 1000 Hz, instead of two.

The two new frequencies are the difference and the sum of the two input frequencies.

If you are using more complex input waveforms, the output wave forms will also be more complex.

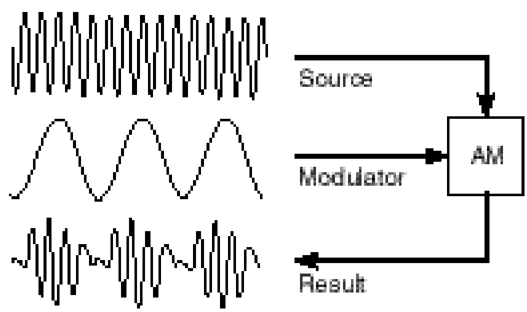

AM

FM (Frequency modulation)

Frequency modulation does even more magical things; two tones can actually produce a lot of different frequencies. If you control the frequency of an sine oscillator at 800 Hz with the output from another sine oscillator at 200 Hz you will create a sound that is a mix of sines at: 200, 400, 600, 800, 1000, 1200, and 1400 Hz!

If you are using more complex input waveforms, the output will be even more complex.

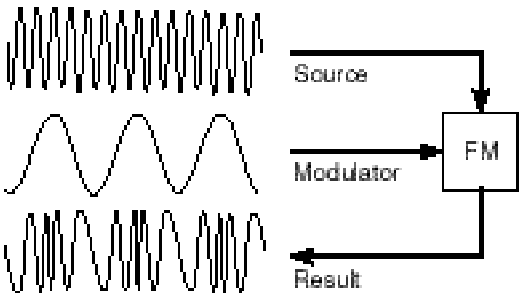

FM

RM (Ring Modulation)

Ring modulation is very similar to amplitude modulation, but the 200 and 800 Hz tones produce only two frequencies: 600 and 1000 Hz. The “center” frequency of 800 Hz disappears. This effect is interesting and a bit weird. As you put in two frequencies, two other frequencies are produced instead!

If you are using more complex input waveforms, the output will also be more complex.

PM (Phase modulation)

Phase modulation is the process of modulating the phase. The result is very similar to FM.