Basic Synthesis

Basic Synthesis: Part 1 – Oscillators

Without **oscillators**, your synthesizer will make no sound (well, actually it can - but that's in a future article). Oscillators produce the original sound which you can hear in your patches, which is then fed through the rest of the signal path in the synthesizer. Oscillation itself is the production of a certain type of waveform, which produces a different sound depending on the shape of the waveform. The waveform is constantly 'run' depending on the speed/pitch of the note - so if an oscillator is set to a low enough pitch you will eventually hear gaps due to the slow speed of oscillation (see LFO below). Common oscillator waveforms are:

- **Saw Wave** - shaped like the teeth on a saw blade, this produces a very common sharp, biting tone.
- **Square Wave** - looks like a (near) perfect square, produces a reedy, hollow sound.
- **Pulse Wave** - a variation on the above, the pulse wave is half as wide as a square wave, and has the unique ability to have its width modulated (called 'Pulse Width Modulation').
- **Triangle Wave** - unsurprisingly shaped like a triangle, this sounds somewhere in between a saw wave and a sine wave.
- **Sine Wave** - a smooth rising and falling shape (like a horizontal 'S'), this produces a mild, soft tone.
- **Noise** - not exactly a waveform, but a source of sound produced by a certain color of noise.

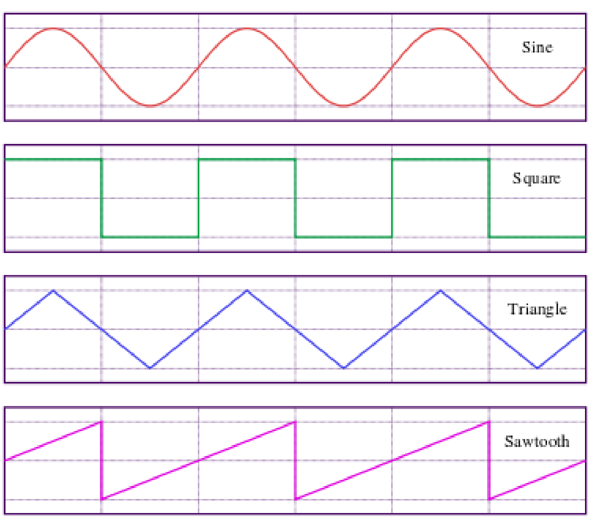

To start off with, that's all you need to know. To spice things up a bit though, depending on the number of oscillators a synthesizer has (usually 2 or 3) you can mix waveforms together! And not only that, but you can tune them differently from each other. This tuning can occur in octaves, semitones, and also in cents - which is a 100^{th} of a semitone, and adds a swirling 'detuned' sound created by multiple oscillators which are cents apart from each other.

LFO, or 'Low Frequency Oscillator' is a special kind of oscillator which oscillates at a frequencies so low you cannot hear it - unless you deliberately tune it into the standard hearing range. It is used to modulate other parts of the synthesizer, such as the pitch of an oscillator, or the frequency of the filter. LFOs still use standard waveforms just like oscillators, but because they operate so slowly the variation in time between the start and finish of the waveform is clearly noticeable - for example, with a sine wave you are able to hear the smooth ascending and descending nature of the waveshape. Used properly, this adds animation and a moving texture to the sound in your synthesizer.

Basic Synthesis: Part 2 – Filters

Filters are one of the most important parts of sound creation, and they are a foundation to the whole concept of 'subtractive' synthesis. The filters do exactly that – they filter out part of the sound, leaving you with a reduced portion of it, which sounds very different to the whole portion. The main control on any filter is the filter frequency, or 'cutoff', which is the key point at which all frequencies are cut off – be it all frequencies which are above, below, in between, or outside of the cutoff point. Common filter types include:

- **Low Pass** – the most common type of filter, the low pass allows all frequencies below the cutoff point to pass through.
- **High Pass** – the opposite of the low pass filter, the high pass filter allows all frequencies above the cutoff point to pass through.
- **Band Pass** – allows a band on frequencies to pass through in the centre, but stops all frequencies outside of this band.
- **Band Notch/Reject** – so called because it looks like a notch, this filter stops a band of frequencies in the centre from passing through.

Each of these filter types can have a different number of attenuation slope, usually 6dB, 12dB, 18dB, and 24dB per-octave – with the higher number being more effective. This is because of the increase in the steepness of each slope per-octave, so a 24dB filter is twice as effective as a 12dB filter, in relation to higher frequencies being higher pitched – e.g. 2000Hz is an octave higher than 1000Hz. Likewise, having more poles in the filter attenuates the signal more, so a 4-pole filter will create a duller, more muted sound than a 2-pole filter which is less effective in reducing frequencies.

Filters also have a feature called '**resonance**', which basically boosts the frequency the cutoff point is currently set off. The more you boost it, the most the set frequency increases, in most cases to the point where the filter will 'self-oscillate', meaning that it creates its own sine wave – the pitch of which can be controlled by changing the filter frequency. On its own, resonance is useful for giving a sound a little more high-end, but it can really create spectacular sounds when used in conjunction with an envelope or an LFO, often creating what is known as a 'filter sweep'.

Basic Synthesis: Part 3 – Envelopes

Envelopes are the key to the articulation of your sound. Without them your patch will immediately start off at full blast, and stay there, and then disappear all of a sudden when you let go of the key. Envelopes, although difficult to understand at first allow you to change that, so you can create expressive and dynamic sounds with your synthesizer.

The standard envelope is in 4 main stages, described below:

- **Attack** – the sound rising up to its maximum level. If it's set to nothing, the sound plays at full blast straight away, whereas if you set it quite high then the sound gradually fades (good for string sounds).
- **Decay** – this is how long the sound stays at the level the attack brings it up to. If it's set as high as it will go, it will stay at the maximum level forever (rendering the sustain stage useless).
- **Sustain** – this is the level that the sound stays at after the decay stage has passed. Some synthesizers also have a dedicated 'sustain time' setting, which decays the sustain stage after an adjustable amount of time too.
- **Release** – a bit like reverb at the end of your sound – it is how long the sustain level takes to die down to silence. Set the release to nothing and you won't get that effect – it will be instant.

The key to programming envelopes on your synthesizer is to practice and visual the envelope. Think of it like a graph with 4 stages, and you are plotting points higher or lower on the domain as the sound changes through the stages of the envelope. The sound rises through the attack, it dies down through the decay, it stays at the selected level in the sustain, and when you let go it disappears through the release.

Filter envelopes are just that – they articulate the filter using an envelope. To do this, there is usually a knob dedicated to the filter frequency (cut-off point) just for the filter envelope. Turn the normal filter cut-off down, and turn the filter envelope cut-off higher, and then program the filter envelope like a normal envelope. Takes a while to figure out, but its worth it – for example, turn the sustain on the filter envelope off, forget about the release and attack (set them to zero), and make the decay short – then turn up the release and wow! The thing to remember here is that its not the volume (amplitude) that your changing, it's the cut-off frequency with the envelope, so you can create great effects with clever programming – such as emulating brass sounds by increasing the attack on the filter envelope.

Basic Synthesis: Part 4 – Effects Part 1

Effects, while not technically a stage of any form of subtractive music synthesis, are commonly found on the end of most signal paths from synthesizers. Effects can come in an internal form (built into your synthesizer), and in an outboard form (where it is a separate effects box, such as a rack mount unit or pedal). While internal effects are useful parts of synthesizers as they save on space and power, they are usually not as high quality as a dedicated external effects unit, but will suffice for most users. Bear in mind when purchasing a vintage synthesizer, that many older models (from the 1970s and 80s) will most likely not have any form of built in effects.

Unison: In similar fashion to effects not technically being part of sound synthesis, unison is not really an effect, but more of a fixed option within your synthesizer. The effect of unison is simply to multiply the signal which is being produced by the synthesizer, and in most cases to detune each signal against the other (just like detuning oscillators). This creates a much bigger sound, and is especially useful for making up a lack of oscillators. Common numbers of signal multiplication in a unison effect can be 2, 4, 8 or even higher. Some synthesizers have become even more creative in their use of a unison effect, such as the Clavia Nord Lead 3 which pans 2 signals to the left of the stereo field, 2 to the right, and keeps 1 in the centre, while detuning all against each other, creating a massive stereo-wide effect.

Chorus: A similar result to the unison effect, but with a much more complex sound modification process, the chorus effect is a popular technique to embellish and enlarge a synthesizer's sound. Some would argue that the technique of the chorus effect is identical to that of unison or detuning oscillators, and this is true in part, as it can be physically demonstrated in acoustic instruments such as the piano or guitar, where multiple strings are played at the same pitch, but are ever-so-slightly out of tune to create a subtle warming effect to the sound.

The standard interpretation of chorusing within the synthesizer signal chain however is of an artificial effect where the signal is copied and mixed with several copies of it itself which have their pitch (at a very small level) constantly swept by an LFO. The effect can also be used in stereo by panning the delay effect within the stereo field, providing that the delay-based pitch sweeps are offset from each other. Because an LFO is used to control the effect, the rate can be adjusted, as well as the feedback and depth of the effect. This 'artificial' chorus effect is also one of the methods used in the flanging effect.

Basic Sound Synthesis: Part 5 – Effects Part 2

Distortion: A popular effect when used on electric guitars, distortion is the process of boosting the synthesizers signal over the limit, to the point where it clips and can remove parts of the audio range. It is also possible to distort signals so that additional harmonics in the frequency range are created. Distortion, or ‘overdrive’, are not common effects used on synthesizers due to the unpleasant clipping of the signal, although used creatively it can produce interesting effects
.

Phaser: In phasing, the signal is fed through an all-pass filter which creates peaks and notches (highs and lows within the frequency spectrum). When done statically, this creates modifications in the signal from the synthesizer, but to create a moving effect an LFO is used to sweep the comb filter, creating the standard phaser sound. Phasers can be used in various different stages, such as 2, 4, 6, 8, 16, or even more, with the more stages being used the more effective the signal.

Flanger: A variant of the artificial chorus and phaser technique, flanging was originally produced by slowing down identical reels of tape, which would produce a sweeping effect of the signals becoming out of sync with each other – but only by a very small amount. Likewise, in artificial flanging a delay technique is used to multiply the signal and constantly change the delay between the signals within milliseconds. This delay effect creates a jet plane-style sound which unlike phasing creates a comb filter effect with the harmonic bumps and dips of the signal being in series with each other. Like the chorus and phaser effects, flanging is modifiable by changing the rate of the sweeping delay, as well as the feedback and depth of the signal. When the feedback of the flanging signal is turned up too high however, it resonates to an unpleasant degree.

Delay: One of the most commonly used effects in sound synthesis; delay (or ‘echo’) is the relatively simple process of copying the original signal from the synthesizer and playing it several times between different time intervals. More exotic techniques can be used, such as reversing the delay playback, and creating feedback loops which run an infinite period of time.

Reverb: Originally created manually by plates and springs, reverb is one of the oldest effects in music in general, and is frequently used in sound synthesis. The effect is similar to creating sound in a large room, where the sound will reverberate due to the size of the room. The effect is created digitally by using many delays, and using accurate simulations of room sizes and their effects, which creates the fake effect of the sound hitting the listener’s ears first and loudest, followed by the reverberation effect afterwards.

Basic Sound Synthesis: Part 6 – Arpeggiators and Sequencers

Arpeggiators and **sequencers** have been part of synthesizers since the 1970s, when basic patterns could be played without having to manually press the keys down, allowing complex melodies to be played with 100% accuracy at any speed. Nowadays, sequencers and arpeggiators have become infinitely more complex, mostly due to the fact that they are programmable from a software interface, where a whole song can easily be programmed with a MIDI keyboard and a mouse. The arpeggiator or sequencer built into your trusty synthesizer is not to be ignored however, and the following information will hopefully give you some idea of the power such functions posses.

Arpeggiators are designed to create small melodic patterns from the notes the player gives the synthesizer. For example, the user may hold down the C, E, and G keys on the synthesizer, and with the arpeggiator active it would repeatedly play the selected keys in various different ways depending on how it was programmed. Arpeggiators can play patterns by ascending or descending through the notes like a scale, or by randomly choosing selected notes to play. Modern arpeggiators can do even more, such as perform sub-arpeggiations within the main arpeggio, or use a mask to mute certain notes within the pattern.

Sequencers (sometimes referred to as ‘step sequencers’) are the bigger brother of arpeggiators, allowing full melodic sequences to be programmed. These can be triggered with just a single key press, or in some cases with only a special trigger button. Sequencers for synthesizers come in two hardware forms. The first, being built into the synthesizer itself is not as powerful as its dedicated hardware alternative. The sequencer in a synthesizer will usually only be capable of holding 16 or 32 steps, as well as tempo and shuffle settings. Hardware sequencers are a dedicated hardware unit which can hold many more steps per pattern, as well as chaining patterns together to form a complete song. They are frequently also able to perform modulations to not just the pitch, but also the articulation of the sound – or indeed any aspect of the synthesizer’s architecture. They are an extremely powerful tool when used to their full capabilities, as well as fun to play with. It is not uncommon to see entire songs played with multiple sequencers connected to synthesizers, performing complex actions such as sending MIDI or CV/Gate triggers signals to all devices

Basic Sound Synthesis: Part 7 – Miscellaneous Stuff

Monophonic/Polyphonic: Early synthesizers were only capable of producing one note at a time, but through evolution of technology they became more and more powerful (and cheaper to produce), to the point that modern synthesizers can play a near-infinite amount of notes simultaneously (like a piano). However, most synthesizers still have the option of playing in a monophonic or polyphonic mode. Quite simply, a monophonic (think 'mono') synthesizer can only play one note at a time, and a polyphonic one can play two or more. The monophonic setting can be useful for some lead and bass sounds, as it prevents two keys from being held down accidentally and overlapping each other.

Portamento/Glide: Gives the ability to slide between notes. The effect of portamento is best used on monophonic sounds to create a bending between notes, but can also be used polyphonically when played in the style of block chords. The time (and occasionally scale) of the glide effect can be changed, allowing different severities of 'bending', from a simple glide, to a very slow sweep between adjacent notes (useful for SFX).

(Simple) FM: Although 'frequency modulation' is an entire form of synthesis itself, the basic concepts of its sound creation methods can also be used in subtractive synthesis. FM effects are produced in a similar fashion to a sped-up LFO, whereby the rate is set so fast that the oscillator produces an audible pitch. The resulting oscillator is then used to modulate the pitch of another oscillator, giving a generally non-harmonic sound, which is very sharp and biting. Due to the lack of pitch stability in analogue components, FM effects cannot be reliably used on analogue synthesizer keyboards.

Oscillator Sync: The standard effect of oscillator sync (technically called 'hard sync' in this case), is the use of 2 oscillators – one which is the master, and the other the slave. The master oscillator runs as standard with its waveform, but the slave oscillator will be running faster or slower than the master. When the master oscillator is triggered, the slave oscillator will also start again through its waveform regardless of whether it has finished its cycle or not. Because the oscillators will be at different parts of the waveform cycle, strange harmonic effects are created when the two oscillators are triggered together.

Ring Modulation: One of the more well-known musical features from the past, ring modulation in music synthesis is the simple process of taking two oscillator inputs and multiplying them against each other depending on their frequencies. Because of the non-harmonic result, this is a good technique for creating dissonant, percussive sounds.