Synthesizers

Synthesizers come in many shapes and sizes, but a great many of them share similar concepts. Understanding these will allow you to program them more intuitively and get improved results. Here we will look at Subtractive synthesizers, a method of synthesis which is used in most analogue synthesizers and as a result has become synonymous with “synth”. Many of these concepts are used in Samplers as well (as you will see in the Sampling chapter on page 177).

History

While there were some experiments made in the late 1800s and first half of the 1900s, the first programmable synthesizer was the RCA Mark II Sound Synthesizer, made in 1958 at the Columbia-Princeton Electronic Music Center in the USA, gaining some credibility with a number of composers creating music for it specifically. The first commercially available synthesizer was sold by Moog in 1964, although the real breakthrough was the 900 Series, a range of modules which

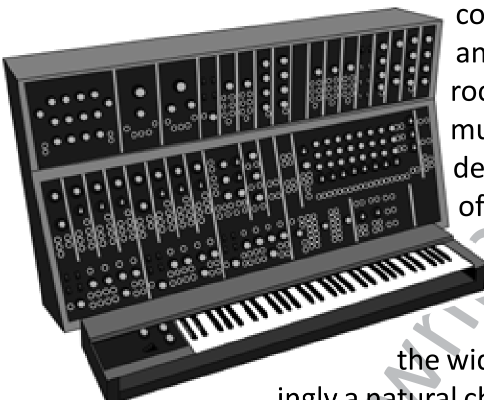

could be connected together to create diverse and complex sounds which became the bedrock on which the future sound of electronic music would be built. While there were many detractors, some musicians saw the potential of these new instruments - not to imitate other instruments, but as a new instrument in its own right. Wendy Carlos’ seminal “Switched On Bach” demonstrated this to the wider public, with Bach’s complex music seemingly a natural choice for such innovative sounds. Inevitably, other companies entered the growing market, and Moog’s high-end expensive and complex devices were undercut by simpler synthesizers leading Moog to produce the classic Minimoog, a synth which was not modular but gained a huge reputation for its sound which still carries on today, despite it being monophonic, meaning it could only play one note at once - a synth which can play more than one note is polyphonic, although the number of notes will depend on the synth in question. Japanese companies such as Yamaha, Korg and Roland released models which through the 1970s and 1980s redefined the sound of many forms of popular music, and the sounds of the synthesizer were transformed from being the unusual and experimental to becoming mainstream and popular. During the 1980s alternative forms of synthesis came to prominence, augmenting the subtractive synthesis model with additive, Frequency Modulation and other forms, and eventually in many cases incorporating an element of sampling by using waveforms from real instruments. By the late 1980s the original wave of analogue synthesizers were unfashionable and unwanted in many corners, but they became popular once more with the advent of House

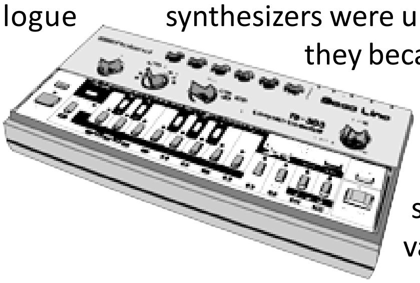

music, leading some previously unloved devices such as Roland’s TB-303 (which was originally a tool to provide bass accompaniment to solo musicians) to become hugely popular and extremely valuable. The increasing power of computer technology led to the possibility of simulating the behaviour of a synthesizer in real-time during the 1990s, and this led to the decrease in value of many physical synths as many felt the sound of the simulated versions of them was adequate and convenient enough to warrant their use over the originals - certainly there are far fewer who could afford an original Moog Modular compared to the £200 or so for the software emulation which is smaller and more reliable aside from any other factors. While there is clearly still a place for hardware-based keyboards it seems that software is where the main developments will be; after all many hardware synthesizers now are running software internally, and there are very few 'real' hardware synthesizers left.

Hardware

The RCA Mark II filled an entire room with rack-mounted components, all of which made up the synthesizer. They could be connected in many different ways and connecting ('patching') everything together was a complex and long-winded job, but necessary before any sound could be generated. The Moog Modular that followed nearly 10 years later was still a formidable piece of equipment; while it was considerably smaller than the RCA unit (due to its use of solid-state circuitry instead of valves), it still required patching between the various modules to allow a sound to be heard (and indeed programmed), and changes between sounds normally meant a considerable amount of re-patching. As many owners often made use of the same combinations of oscillators, amplifiers, filters and envelope generators, synthesizers were released which had them connected in a fixed manner, with the sounds being altered solely by changes to the parameters of each component using knobs (as the original modular synths had featured). This made them much more immediate to use and led to a rise in popularity as a non-technical user could quickly experiment and get a sound from them. This led to the rise of many now-popular models from manufacturers such as Moog (the miniMoog, seen right), Roland (the Jupiter and Juno series), Korg (Polysix, below left) and Yamaha (CS-80).

Up to this point all of the hardware had been analogue - in general the devices were all simple ones, initially being made from individual transistors, resistors, capacitors and inductors, but later on the popularity of the synth market led to some manufacturers such as Curtis making some elements of synthesizers (such as filters) available as a chip instead of a larger number of individual components. While this was expedient at the time, they are now discontinued and extremely rare, so repairs of some synths that use them can be expensive if not impossible! During the 1980s, digital technology began to be used; at first it was solely for controlling the analogue parts of the synthesizer (that still made the sounds), but increasingly the technology became available to create the sounds digitally - instead of using components to make the audio, it was calculated using mathematics and computer programming, and then converted into analogue audio for the listener. This led to synthesizers which were more stable in terms of tuning and behaviour, but generally made them far less user-friendly; to keep costs down the number of physical controls was reduced in many cases, and in some there were no knobs or sliders, just an LCD readout and a number of buttons to press to change settings, making programming a far cry from the original tactile experience that it was, and leading to many giving up on programming sounds and using presets instead; this combined with the complexity of some forms of synthesis (such as FM - the Yamaha DX7 seen right being the most famous of them all) led to many 1980s recordings making use of the same sounds. One of the advantages of such digital synthesis and control was that many synthesizers became multi-timbral. This is when a synth is capable of producing more than one different sound (timbre) at once. We take this for granted now with many synthesizers and the advent of General MIDI, but it was a groundbreaking feature when first introduced in the 1980s.

The advent of House music led to a resurgence in popularity of many analogue synths, as the ability to alter sounds readily in real-time meant they were ideal for making music where the music was largely static, but the sound changed; the altering of filter settings is something which is still popular in many electronic forms of music 20 years later. This time was the peak of the 'physical' synthesizer, however; advances in computer technology meant that a general purpose computer (such as a PC or Mac) had enough processing power to be able to simulate the sound generated by an entire synthesizer, calculating what would happen at any given moment with enough accuracy to replicate the sound to all but the most discerning of ears, and the fashion for simulating 'real' instruments inside a computer progressed hugely during the 1990s, with previously desirable analogue equipment becoming unwanted once more. This trend has continued throughout the start of the 2000s, with the ever-increasing power of computers meaning it's not only possible to simulate one of a given synthesizer, but many at once; one often overlooked advantage of this approach is that you are not just buying one synth when you purchase a software synth, you are buying as many as you can get your computer to run at once.

Structure

There are several different types of synthesis in use today, but the most common one is called subtractive synthesis – this is where an initial waveform is generated that is harmonically complex (see page 256 for details of this concept), and then this is filtered to reduce the amount of harmonics present, and the volume of the output is also controlled. Many of the classic analogue synthesizers revered today used exactly this method to produce their sounds, and in addition a great many software-based synthesizers (such as VST Instruments) use such a method.

Their structure looks is shown on the next page. The signal flows from left to right – the MIDI input goes to the oscillator, and this generates the appropriate sound which is then passed through an amplifier and filter. These use ADSR envelopes (a way to alter a parameter over time) to control their outputs. In addition there is also an extra oscillator (the LFO) which generates low-frequency waveforms which can be applied to any of the three main sections of the synthesizer to control their output. When the LFO is applied to the oscillator, it produces vibrato because the output frequency is changing. When applied to the Amplifier, it will produce a tremolo effect (a rapid variation in volume), while applying it to the filter will produce a cyclic change of the filtering of the sound, altering the brightness.

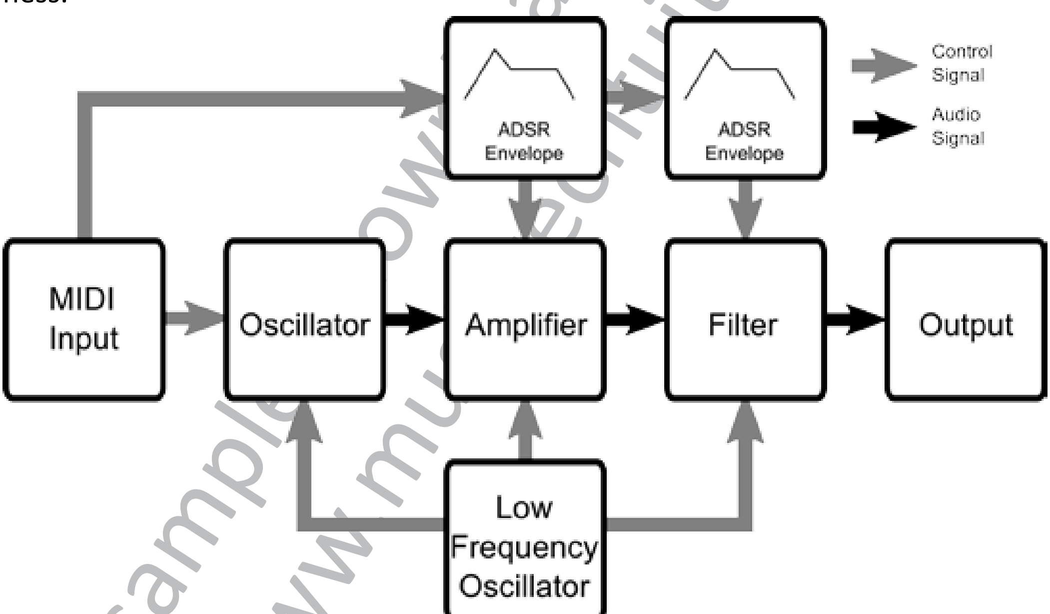

Some synthesizers feature two (or more) oscillators to allow each note played to produce more than one tone – it's possible to offset these in pitch (producing a doubling effect), detune them (producing a chorus or out-of-tune effect), and often to change the waveform of each as well, producing different effects.

Editing a synthesizer

To demonstrate the principles of editing such a synth, we will use the freely-available CM101 VST Instrument – create one now (`Devices > VST Instruments`) and it should look like this.

Here we see the CM101 with its default settings, which give a rather non-descript sawtooth wave sound (see page 255). However, looking around it you should see that the sections present in the schematic diagram we saw earlier are present – from left to right, there are Oscillators (two, but they both work in the same way), then master section (which is not covered in the diagram, but allows you to control the master output level, pitchbend range and the number of notes the CM101 can play at the same time), the amplifier section (with ADSR), an LFO below it, and then the filter section with its own ADSR and some related controls.

The simplicity of this synthesizer and its sections belies the amount of sonic editing potential that is present. To illustrate that, we will first alter the Oscillator to produce a waveform that is rich in harmonics, and then use the amplifier and then the filter to illustrate some of the possibilities that this setup allows.

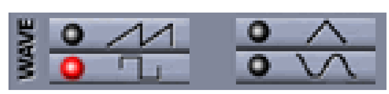

Firstly change the waveform to the square/pulse using the button shown to the right – this is a square wave, and has a sharp sound which will make hearing the filter's effect (later on) much easier. The 'width' of the two parts of the pulse wave can be varied with the "Pulse Width" control, but this should stay at the default position at the moment (50:50).

Finally we only want to hear Oscillator 1 at the moment, so turn the level of Osc 2 down to 0 from the default of 100. Playing this via MIDI will give a bright, clear tone which is reminiscent of early TV video games. Because of the shape of this waveform, it is very rich in harmonics – as well as the fundamental frequency there are a wide range of harmonics present as multiples of the original frequency. This gives us lots of potential to use the filter to alter the sound, but firstly we will look at the amplifier section.

The Amplifier Section

The amplifier section takes the basic oscillator waveform and amplifies it to create a more complex sound which evolves over time. This is achieved by the amplifier being controlled by an envelope (the ADSR). This is a method of controlling something over time, and usually it has four parameters – Attack, Decay, Sustain and Release – which sometimes lead to it being known as an ADSR Envelope, or even just an ADSR. Three of these parameters (A, D and R) are time-based, while Sustain is a level, as is shown below, with the parts of the envelope described next to it.

The Envelope Generator

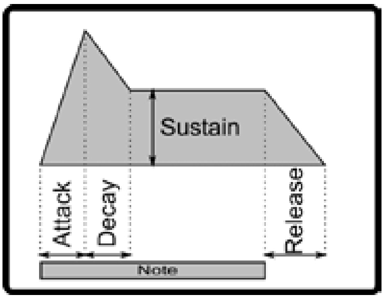

Attack	The time between hitting the key and the level reaching maximum
Decay	The time taken from the maximum to reach the Sustain Level
Sustain	The level which is maintained while the key is held down
Release	The amount of time for the envelope to reach zero once the key is released

To hear this in action, we will need some MIDI data to be sequenced and then played back by the CM101, to allow concentration on the changes the controls make. For this example, sequence a series of 16th notes at around 120bpm – this will mean you will hear the effect of the Amplifier. The notes themselves are not important, just that they are present, something like the sequence shown to the right.

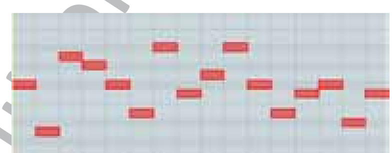

Next, we will examine the settings on the CM-101 Amplifier section - set the controls of the AMP section as shown in the left-hand column. Next there is a diagram of the envelope's output, and finally a description of the settings and their effect.

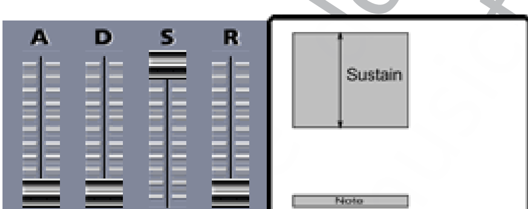

With the controls set as shown, the Amplifier is effectively either ON or OFF – as soon as the key is pressed the A and D times expire, leading directly to the sustain level, which is set to maximum. Releasing the key leads to an instant stop to the sound.

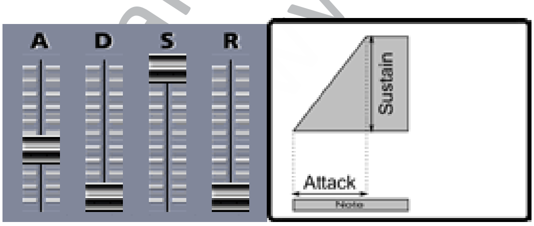

Moving the A control upwards will make each note 'fade in'. As long as the A time is less than the note length, it will reach maximum and then stay there, as shown.

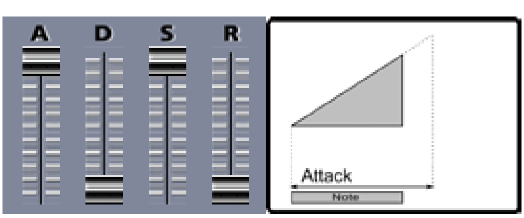

However, if the A control is made too long for the notes that are being played, the maximum level will never be reached (shown with the dotted line in the right-hand diagram).

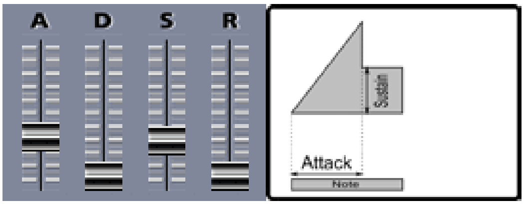

Next, set the S control to around 25%. This will mean as soon as the attack is over, the amplifier will instantly drop to the sustain level, as seen to the left.

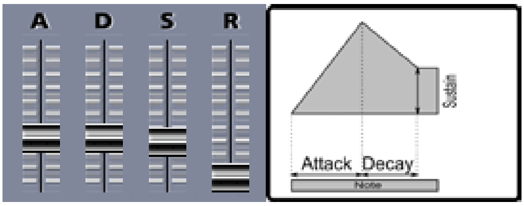

Now set the D control to around a quarter of the way up. It will now take some time for the amplifier to reach the sustain level after the maximum has been reached - this is useful for simulating the decay of real instruments.

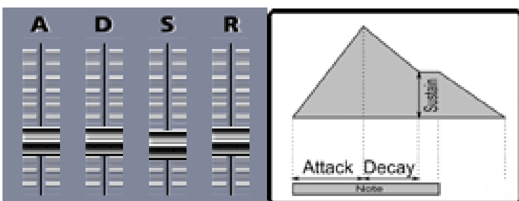

Now add in the R control - this will lead the sound to fade out slowly after the key has been released - this is often used to add more realism to sounds (very few 'real' sounds stop instantly), or to create more pad-like sounds which keep sounding after the note has been released.

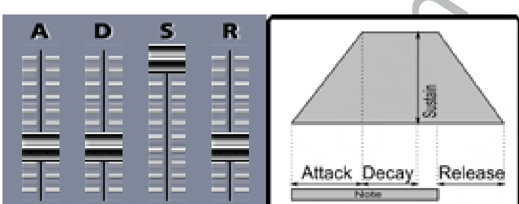

If the S control is set to maximum, it makes the D setting redundant - it doesn't matter how long it takes to change between the end of the attack (100%) and the sustain level (100%), and the end result looks like the diagram on the left.

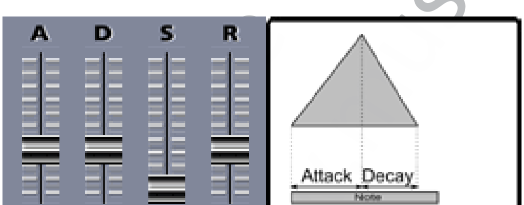

Conversely, if the sustain level is set to 0, then (providing the note is long enough) the release control will do nothing as the sound will already have stopped; only the A and D controls will make any difference.

The last settings will make a sound which is of a fixed length. Using this technique it's possible to create sounds which are percussive - with a short attack and a medium decay they will fade like a 'real' percussion instrument.

With these simple controls it is possible to create a wide range of envelopes for controlling volume and to simulate the performances of many 'real' instruments.

The Filter Section

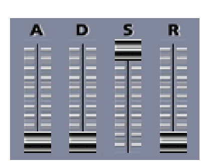

Before starting this section, make sure that the amplifier section is set to its default settings (seen to the right) – this is to make sure that you are only hearing the effect of the filter at the moment and nothing else. On the next page we see the filter in its default state – it is not filtering anything out, because the cutoff

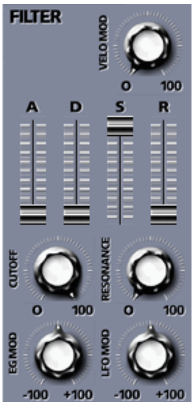

frequency is set to maximum (100).

To make the filter work, it is best to sequence and repeat a short (2 or 4 bars) section of short notes (such as the one shown for the Amplifier examples), and then alter the `Cutoff` control while this plays back. You should hear the sound alter quite dramatically as you reduce the cutoff frequency.

Tip: Be careful with your monitoring level when using the filter, particularly with resonance in the next section as it is possible to get extremely loud sounds, which could damage your hearing if you are not careful!

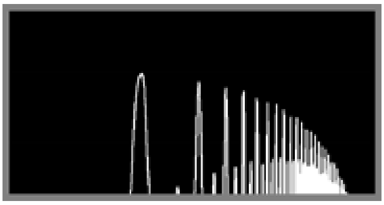

Here is a frequency analysis of the original sound. We can see the fundamental tone on the left, and then all the harmonics, which extend into the high frequency range, giving this a sharp sound.

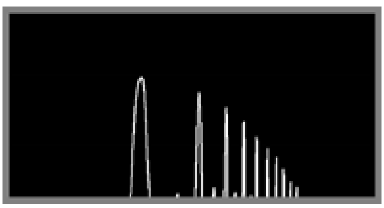

The cutoff frequency has been set to around the half way mark, and the result is this - there are far fewer harmonics present, and the sound is much duller, approximating the tone of a triangle wave.

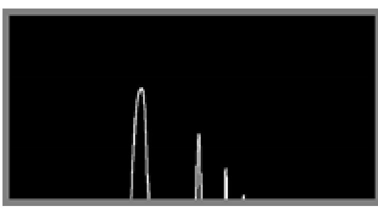

The cutoff frequency has been reduced even further, and leads to a very dull, near sine-wave sound which has none of the rich, bright tone of the original sound.

The effect of the filter will vary depending on the original waveform you are filtering; as we saw in the audio chapter on page 256 a sine wave has no harmonics, so it will not change tone. Square and Sawtooth waves have lots of harmonics, and so will be considerably altered by the low pass filter removing them progressively.

Resonance

The CM101's filter is like many others, in that it has control of the resonance of the filter - this means it can boost the frequencies present at the cutoff point. This can make the sound of the filter a great deal more dramatic, and as seen later on in this section it is possible to create some interesting and varied sounds, but first we will take a look at the effect of resonance.

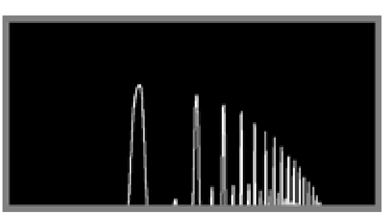

Here we see the original frequency analysis of the sound. The cutoff frequency is below that of the original example above, and some harmonics are being filtered out.

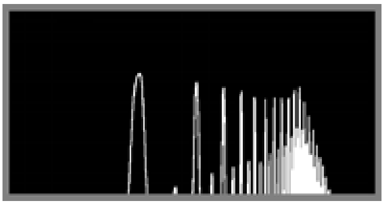

Now the resonance control has been turned up slightly, and the effects can be seen and heard - the sound is brighter, and we can see that the highest harmonics are being boosted.

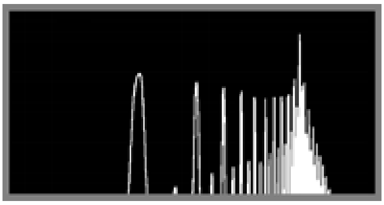

Here the control has been set to maximum, and there is a clear peak where one harmonic is being boosted more than all the others, but in addition the entire high frequency area is being boosted (compare this with the first diagram).

To demonstrate the interaction between cutoff and resonance of this effect, set the resonance control to around the 50% mark and move the cutoff frequency control, starting at 100% and moving all the way down to 0. This resonant low-pass filter sound is one that is extremely common and has been used widely both on synthesized sounds and also on samples (see the Samplers chapter – page 177 – for details of this) Note that with the correct values of cutoff and resonance it's possible to produce sounds which are quite different from the original waveform.

However, to get the traditional analogue "squelchy" sound, an extra element is needed – the filter needs to alter over time, so the next step is to use the ADSR envelope to control the filter so it opens/closes over time. This involves an extra step than with the amplifier, as we need to set how much effect the envelope will have on the filter, and this is done with the EG MOD control - for the time being set this to a positive value (+50% or so) – this means "Envelope Generator Modulation", or rather "how much effect the Envelope has on the filter" – with this set to zero, it would have no effect, something which has led to much confusion by many in the past!

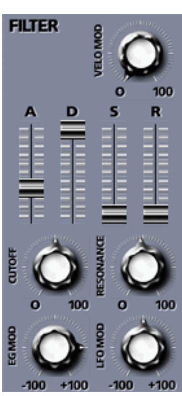

Next, it's time to make the envelope actually do something – its default settings don't do anything interesting. Move the A control to around a quarter of the way up (somewhere around 300mS) while playing your sequence – the Filter section should look something like the one to the left, with cutoff and resonance around 50% if the settings you had left it at didn't work straight away.

You should hear the familiar "squelchy" analogue synthesizer sound. Note that if you move the control too far, you may hear less as the filter never fully opens (this will depend on what you have sequenced – if you have played long notes then there is time for this to happen, while short notes will not have that).

You can also experiment with the filter's S, D and R controls, and see what effect they have – again remember that their effect will be dependent on what you have sequenced, in much the same way as they were with the Amplifier section.

Next, set the `EG Mod` control to -100%. This will 'invert' the ADSR envelope, so that instead of the filter opening with time, it will start out open and close instead. This will lead to some different sounds one more, again reminiscent of the heyday of analogue synth usage. Even with the cutoff set to 100% you may well get an interesting effect.

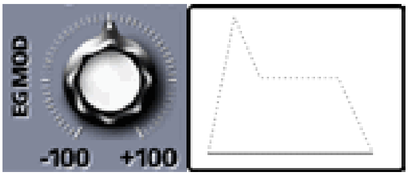

EG Mod at zero is the default – the filter settings will not change over time. The envelope is shown is a dotted line, but the filter is unchanged (solid line) - the Envelope has no effect on the filter.

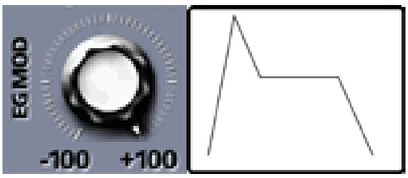

EG Mod at maximum (100%) – the filter will follow the ADSR envelope that is set up for it, allowing timed control of the filter's cutoff frequency, giving a sound which will brighten over time and then hold at the sustain level.

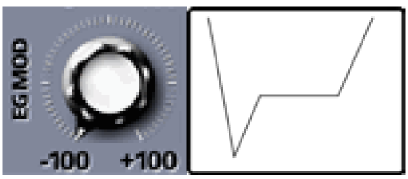

EG Mod at the inverse maximum (-100%) – the filter will follow the envelope, but 'upside-down', starting off open and then closing with time, making a sound which will become more muffled.

It is also possible to control the filter via two other sources – firstly, Velocity. Each MIDI note you play or sequence has velocity, and this can be used to open the filter further, using the `VELO MOD` control. When set to a high value this will mimic the tendency of real instruments to become brighter when played harder.

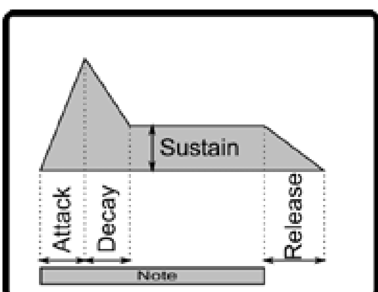

Here the note that is played is a low-velocity one, and as a result although the filter follows its usual ADSR envelope, all levels are reduced as result of the velocity. This leads to a note that is duller in tone as the filter doesn't open as far as it does when....

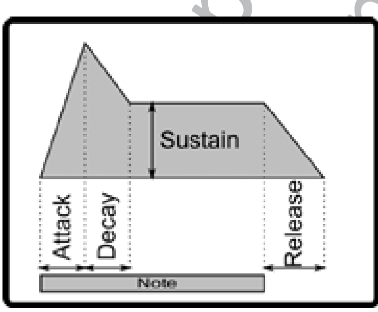

... the note is high velocity. This causes the envelope to act more strongly on the filter, and open it up further, giving a brighter, clearer sound. Routing note velocity to the filter envelope is an important technique when trying to emulate real instruments as well as making synthesized tones more expressive.

Finally, it's also possible to control the filter with the Low Frequency Oscillator.

Low Frequency Oscillator (LFO)

The LFO is an oscillator which runs constantly, and as the name says, at low frequency – typically it will be from around 0.01Hz (one cycle every 100 seconds) to 20Hz (20 cycles per second), although each synthesizer varies.

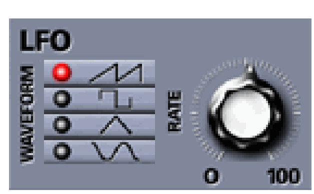

Typically a number of different waveforms can be selected – in the case of the CM101 there are four, Sawtooth, Square, Triangle and Sine.

The LFO doesn't generate audio itself (it is at too low a frequency), but is usually used to control other areas of the synth. If it is applied to the oscillator section's frequency, then the result is vibrato as the LFO will change the output of the oscillator over time. If it is applied to the amplifier (something which the CM101 doesn't allow, sadly), then the result is tremolo. If it is applied to the filter, the result can vary greatly from a subtle change in brightness to bizarre, special-FX type sounds.

The routing possibilities and functions of an LFO vary greatly between synthesizers – some (like the CM-101) offer basic facilities, while others have a wide array of options, and sometimes several LFOs, which can be synched to MIDI or the sequencer's tempo to give more musical applications.

Changing Waveforms

The next area to consider is the initial sound that is generated – i.e. the Oscillator section. The CM-101 oscillator section is fairly typical, and is shown to the right. Some synths only have one oscillator, while some have several, and their functions can vary, but the CM-101 covers all the basic, common functions. From the top we have:

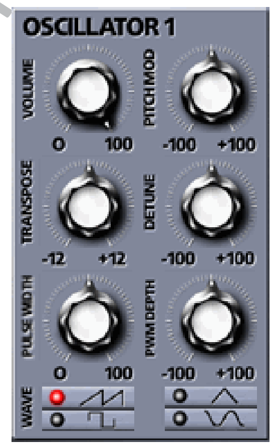

Volume – this controls the volume of the signal that is generated. By doing this we can control the overall level of the sound, as the Amplifier Envelope can only work with this volume level, not increase it.

Pitch Mod – this controls how much effect the LFO has over the pitch of the notes generated, or in short, vibrato. In the case of the CM-101, 100% means a change of an octave, giving a very wide range of vibrato which can be useful for extreme effects sounds.

Transpose - the pitch of the MIDI note can be transposed over a range of

±12 semitones. While this in itself may not seem that useful (after all, you could just transpose the MIDI data), when using two oscillators together, it allows for doubling and harmony to be created.

Detune – this is an extension of the Transpose control – it allows any tuning to be achieved via its ±100 cent range (i.e. a semitone each way). Subtle detuning of one oscillator gives a rich, warm sound mimicking the poor tuning of original analogue synthesizers.

Pulse Width – this control alters the width of the first part of the square wave, from 0% to 100% - the effect of the control is shown below. Note that if the PW was either 0 or 100% there would be no sound at all.

PW = 10 %　　PW = 50%　　PW = 90%

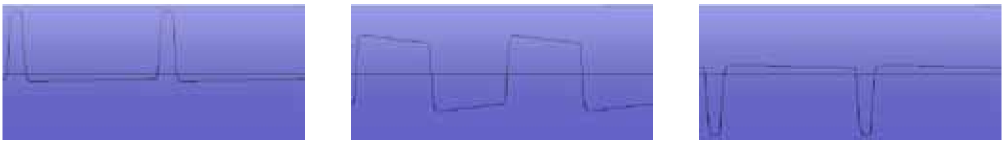

The pulse width alters the tone of the sound greatly, from very thin (at either extreme) to thick, but all of these settings create lots of harmonics which makes them particularly suitable for filtering.

PWM Depth (Pulse Width Modulation) – this control allows the LFO to control the pulse width, which can produce some sounds which alter greatly in tone as the pulsewidth changes. Using a slow LFO waveform can produce some evolving textures, while quicker settings can make for a detuned sound.

Two Oscillators at once

As mentioned above, using more than one oscillator has creative possibilities – it is possible to create a number of effects, using the settings outlined below, once OSC2's Volume control has been set to make it audible too (set both to 100 to start out with)

Slight detune on OSC 2 (around 6 cents) gives a thicker sound, mimicking a chorus effect, multiple players or an out-of-tune analogue synth.

The 'classic' fifths harmony sound, which plays the fundamental with OSC1 while playing a harmony seven semitones higher with OSC2.

By setting OSC 2 to an octave below OSC 1, an octave-doubled sound is achieved, thickening the texture greatly. This can be combined with the detuning shown above for a larger sound again.

With this example, the two oscillators are set similarly, but OSC2 has some PWM Depth modulation applied (as well as being set to be an octave below OSC 1), so the square-wave sound alters in thickness while still retaining some depth as OSC1 is always playing a straight square wave sound.

There are clearly a myriad of possibilities when using two oscillators, and their combinations are almost endless, especially when waveform types are taken into account. Experiment with the controls to see what you can create and discover.

Exercises

Listen to waveforms

Being able to identify different basic waveforms by ear is an important skill - hopefully at first the sine wave will sound quite different to the others, but hearing the difference between the others outlined on page 255 can take some practice. While there are a near-infinite number of waveforms, being able to recognise the basic ones will help you in being able to analyse sounds and program synthesizers to mimic them.

Use the amplifier's ADSR

The ADSR is the cornerstone of making synthesized sounds dynamic; without one the sound would be static with respect to time. Spend some time practicing using the ADSR to create sounds and mimic the dynamic changes of real instruments, and mastering the controls. Many people find the A and R components are easy to understand, but remembering that the signal level always reaches maximum at the end of the A phase, and that S is a level, not a time is something to keep in mind, as well as the D control becoming irrelevant if the S control is at 100% (as there is no change in level, the time taken is meaningless).

Use the filter

Filtering is the next important concept in subtractive synthesis (and indeed many forms of modern music production). Being able to recognise different kinds of filter is important; while a great many filtered sounds that are heard will be low pass filters (often with resonance), being able to recognise high-pass filtering and band-pass filtering is important, and with practice should be second nature. Remember that not all synths give a choice of filter type; while the CM101 (see page 169) offers only low-pass, others such as the A1 offer a wider selection, so experimentation is important.

Use the filter's ADSR

While a static filtered sound is more interesting that an unfiltered one, the change of filter cutoff over time is important to create truly interesting and dynamic sounds. The ADSR is the same as featured in the amplifier section, but controlling the tonal balance gives a completely different result. While experimenting it's important to ensure that the amplifier ADSR settings are not stopping you from hearing the effects of the filter ADSR; if the amplifier ADSR has already gone to 0% you won't hear what's happening with the filter as there is no signal.

Use multiple oscillators

A single-oscillator synth gives a wide range of sounds, but adding extra oscillators gives many more options - it's possible to detune the two oscillators relative to each other, both fine-tune and interval-based tuning will create interesting effects. There is also the option of using different waveforms for the two oscillators, giving a much wider range of tonal options. Listening to the effects of this as you experiment is important, as it will allow you to recognise where others have used similar techniques.

Summary

Subtractive Synthesis (as outlined above) gives a huge range of timbres, and there are many, many virtual synthesizers available today, some of which are new designs, some of which mimic hardware devices of yesteryear. In addition, 'real' synthesizers still exist, and knowing the basics of how to program any one of these devices will give you the ability to create unique sounds, mimic sounds that you have heard and also provides a basis for many activities that are done with conventional samplers. Familiarity with the controls and facilities of each synthesizer is a pre-requisite for being able to program them successfully, and this takes time and experimentation to fully master. However, even casual familiarity with features such as filters will allow quick programming of many effects that have become trademarks of modern music production. It is important to remember that even the most complex subtractive synthesizers which offer a huge range of modulation possibilities (such as Native Instruments' Massive) and modular synthesizers such as the Moog Modular (whether the real thing or Arturia's software equivalent) are based on the principles shown in this chapter, and only when mastering programming of a synth with simple architecture (such as the CM-101 and its real-world equivalent the Roland SH-101) can the programming of a more complex instrument be possible.