Synth Programming (pt V)

Our final look at effects, arpeggiators and more...

OVER THE LAST four months, if you've been following this series and trying out some of the ideas on your own kit, you should have some idea how to make more interesting sounds than average. You should by now be able to make sounds which are completely average on their own, but still fit perfectly into a track.

This month we're going to end by looking at effects, arpeggiators, and some of the more extreme synthesis options you can use if you want to try some more advanced ideas.

So let's start with those effects. Effects can be as much a part of synth programming as the oscillators and filters can. Especially now that many effects options are just as sophisticated as those you'll find in a separate effects box. The usual suspects include:

Delay: usually just echo-o-o-o. You can set the time, the mix (full effect 'wet' or no effect 'dry') and the amount of feedback, which sets the number of repeats. Often useful on basslines to add some extra in-between-notes interest - use a time that's in sync with the tempo - and on pads to thicken them out.

Flange: jet-plane style whooshy sound. Have this sweeping away on a sequence line, or on a pad. The more treble there is the in the original sound, the more obvious the effect is. Settings are modulation amount and rate (controlled by an internal LFO), and depth. And resonance - also known as feedback. Turn this up for a metallic effect.

Chorus: like flange, only thicker and runnier. Mostly used on pads. Uses similar controls as flange.

Phasing: like flanging, only more analogue sounding. Almost always great on basslines.

Distortion: there are basically three sorts. Amp simulators pretend to sound like guitar amps. Overdrive and 'pure' distortion just grunge the sound up. Digital distortion eliminates the detail in the sound by throwing most of the (digital) bits away and sounds really nasty; it's great for crunching up samples.

EQ: tone controls or, in other words, bass and treble. Sometimes you get a mid as well.

Enhancer: makes the sound more brash and bright.

Reverb: creates an artificial space around the sound - from subway to concert hall to aircraft hangar.

Compression: fattens up the sound.

What a combination

Some synths have combined effects as well, which give you separate effects on the left and right output channels, or two effects in series. Some also have individual effects for each voice in a multitimbral setup, and then one or two 'global' effects that add that extra something to the whole mix.

For example, the Access Virus has an effects chain that includes delay, distortion, a phaser, and chorus on the main mix. On each individual voice you can have a phaser, distortion, bass boost, chorus/flanger and a ring modulator.

So how do you put all these together? And are they icing or are they cake? That depends on the kind of patches you're using. It's tempting to throw everything into a patch to try and make it memorable, but that's usually a bad idea. The key thing to remember here is - as mentioned last month - don't listen to single sounds, but to how to they sit together in a mix.

As a rule of thumb, chorus and reverb tend to push things back a little, so they're great for adding atmosphere.

Distortion tends to push things forwards so they stand out more.

A second rule of thumb is to think of effects as being an essential part of the patch, rather than something stuck on afterwards. Try dialling in the effects while you're programming, rather than as an afterthought. Otherwise, a complete guide to everything that's possible would take an entire article, but for some starters take a look at the Effects specs box on p102. And then try them.

Into the rhythm
No synthesis series would be complete without mentioning arpeggiators. They're excellent tools, very easy to use and program, and a lot of synths have them. Some kinds of dance music - Goa and other flavours of trance, ambient and some older rave styles - would be impossible without them. So here's the beginner's guide.

The basic, bog standard plain vanilla arpeggiator works by playing a sequence based on notes you're holding on the keyboard. (Some synth modules have arpeggiators too - they work instead on MIDI information you send to them.) Slam down any chord pattern, and the arpeggiator will play it - up, down, up and down, at random, over a single octave, or over two or three or four.

Typically there's a 'hold' option, which means the arpeggiator keeps playing even when you let go. And MIDI clock sync can sometimes be divided down, so the arpeggiator plays half, quarter, eight, 16 or 32 to the bar. More sophisticated arpeggiators let you use 3/2, 3/8, 3/16 subdivisons for triplets for more interesting rhythms.

Most people tweak and twiddle with arpeggiators for a while, and then get bored with them. They're not that easy to play by hand. If you play lots of chords with three notes, and then play one with four notes, the rhythm changes awkwardly. The same thing happens if you hit an extra note by accident.

The way round this is to sequence the notes, instead of playing them by hand. 90% of this year's (and every year's) Ibiza classics use this technique. You can re-record the MIDI output of the arpeggiator if you want to edit it. In fact, this is an excellent thing to do anyway, because arpeggiators can be a fantastic source of ideas for bass and sequence lines.

Some synths have more interesting arpeggiation options. The Quasimidi Sirius has a range of preprogrammed patterns (which Quasimidi calls 'motifs') that either repeat a chord, or split it up in more interesting ways that go beyond playing the component notes. You can also program your own motifs. Another variation is to create a step-time rhythm pattern that only plays on certain beats. If you combine this with the riffing abilities of the typical arpeggiator, you can create hypnotic sequences that never repeat.

But things become really colourful when you combine all of these to create a multitimbral sequence. Synths like the Waldorf Microwave and the various larger Novation machines let you do this. On these you can have a different type of arpeggiation happening on lots of different patches at once. You can combine drum patches with synth and bass sounds to create a mega-patch that plays an impressive mini-sequence. All you have to do is hit a single note on the keyboard to start it off.

Confused? Here's an example. Let's start with a bass drum patch on channel 1. Set this up so the arpeggiator plays on quarter notes. Then hi-hats on channel 2 playing 16ths. This gives an excellent foundation for the rest of the multi-patch. What you do next is up to you. A good option is to create a rhythm motif, and use it with a bassline. Another excellent option is to create a patch with a very slow LFO sweep on the filter, and some resonance to give it a bit of bite.

Better still, do both. Using trial and error you can soon build up a complex rhythmic pattern that you can play with one finger. If you have other synths with arpeggiators you can use those as well, to build up an even more complex effect. It doesn't take long for a complete track to more or less write itself.

The hard stuff
Can you journey further into the world of weird sound possibilities and create limpid landscapes of shimmeringly seductive gorgeousness, drums with more aggression than an SAS unit facing down an alien invasion force, and basslines so chunky you could quarry out some rooms and live in them? In a literal sense, probably not. But if you want to go further than the average sound designer, here are some things to do, techniques to try, and tools to buy.

Vocoders: vocoders work by splitting an input ('analysis') sound into frequency bands, working out the volume in each band, and then applying those volumes to an equivalent set of frequency bands in a different ('synthesis') sound. When you do this, magic things happen.

If the analysis sound is a voice, and the synthesis sound is a string pad, you get that talking pad effect that's been

done to death. Try talking over basslines and drum parts too. Or don't talk at all; you can more or less vocode anything with anything (even itself) and still get an interesting result. You can also get interesting noises by retuning the different filters so they don't work on the same range of frequencies.

Resynthesis: this is where you pull a sound apart, a bit like sending light through a prism, hack around with the 'colours', and then put it back together again. With resynthesis you can morph sounds, exaggerate the way they change, tune unpitched sounds so they ring like a bell, blend sounds (so you can make talking drum loops - only with much higher quality than you'll get from a vocoder), create totally bizarre noises by randomising the 'colours', or vintage retro effects by throwing away a big chunk of the sound, like a kind of mega EQ.

Additive synthesis: is a bit like resynthesis. When you do resynthesis it turns out that the colours you get can be reduced down to sets of sine waves, each with their own frequency and level. In other words, if you had enough sine wave oscillators, each with an independent pitch and level control, you could make any sound you wanted - real, or imaginary.

In practice, you need hundreds or thousands of oscillators to do this well. That's for each note. And if you're trying to synthesize real instruments, you also need to remember that the sounds changed according to pitch and velocity. So high quality synthesis of real instruments isn't possible yet. But you can still get some interesting and unusual effects like this.

Granular synthesis. If you've ever timestretched a sample you've already used granular synthesis. Granular synthesis works by chopping a sound up into tiny segments. You can then control the envelope, volume and the pitch of each segment. You can 'granulise' samples or start from scratch with little oscillator-produced blips. It sounds a simple technique, but it has a lot of applications.

By scrunching up or expanding the gaps between granules you get timestretch. By changing their pitch you get pitchshift. By randomising the pitches you get a kind of 'sound cloud' effect. By randomising the times you can get monster reverb-like sounds. By using sine waves and precisely controlling the envelopes you can create interesting vocal-like effects. (This is sometimes known as formant synthesis.)

Physical modelling: Instead of using waveforms, physical modelling actually simulates in software how all the mechanical or electronic bits of a synth work together. So to make a guitar sound you create a complicated mathematical model of how the strings vibrate (this bit is called the 'exciter'), and how the soundboard amplifies the sound (this is known as the 'resonator'), and so on.

It's very heady stuff, which is why most
physical-modelling synths just give you a set of presets to play with, that you can tweak in a minimal kind of way.

And in spite of the complications, it's not that hard to get your head round; 'exciters' work like oscillators, 'resonators' work like filters. The rest is more or less trial and error.

How to get hands-on
Apart from vocoding, which is available in many mid- to high-end analogue-style synths, these and other effects are mostly stuck in the realm of the experimental. This is bad news for anyone who wants instant gratification. But good news for people who want to make sounds that are literally like nothing ever heard before.

As for getting your hands dirty with all this synthesis business, you really need to get your hands on some serious synths and for this the Internet is a blessing. Some of these toys are free, some are cheap, a couple are seriously expensive but here's a rough list:

Csound. Big, bad, totally obscure, totally free, multi-platform, and really only for über-nerds. This is (literally) the granddaddy of all other soft synths. It does more than everything, but it's slow, clumsy, almost perversely difficult to use, and maintained by a rather cliquey group of academics who will laugh at you if you don't have a PhD.

Most soft synths have oscillators, but Csound has simple oscillators, sample-playback oscillators, high-quality oscillators, click oscillators, additive synthesis and resynthesis oscillators, wavetable oscillators and an option which creates oscillator-like sounds by calculating the orbit of a planet in a binary star system.

To get the best from it you need to wrap it up in some smarter high-level software. Linux weenies or Mac users will want to use something called Cecilia, which sorts out some of Csound's shortcomings. (Unfortunately there's no Windows version of Cecilia.) More information can be found at www.notam.uio.no/ internt/csound/Csound/TITLE.html

Composer's Desktop Project (CDP).
UK-based resynthesis package. It's command line based (yes, that's right, you have to type in filenames and commands by hand) but it will do most resynthesis-style things. And there are attempts to create a more up to

date Windows-y interface. While it's a bit clunky by modern standards, it's also not too expensive at just over £100. And it can be fun to dabble with.

MetaSynth. Mac-only resynthesis toy, with the unique ability to turn pictures into sounds (and vice versa). It's a fantastic product, but there's some confusion about whether it's still available. Try www.uisoftware.com for details.

Kyma system. Forget expensive analogue modulars, this is currently the ultimate commercial synth system. At any price. It's a soft synth that uses its own hardware in the form of a rack full of plug-in cards, a bit like the Creamware Pulsar, only bigger and better. You can get started for around £2,000, while a fully expanded model costs around £7,000. Kyma is the thinking person's Csound. It does most things that Csound does, only it's about a million times easier, quicker and more fun to use. For details take a look at: www.symbolic sound.com/kyma.html.

Going hard
Meanwhile, if you don't fancy donning the lab coat and specs, and going down the intricate software route, why not go hard and get yourself a real synth! Kawai K5000 series. The only hardware synth to do proper additive synthesis the K5000 rack synth is now available second-hand for around £400 or less (check our Marketplace ads starting on p128). It's capable of some unique far-side noises and well worth watching out for as they do come up from time to time.

Yamaha VL-series synths. This is the acceptable face of physical-modelling synthesis, starting from the hugely overpriced and commercially unsuccessful VL1, via the VL1M and VL7, and eventually ending up with the VL70 and even some VL-ready options in the Yamaha budget soft synth range. You might find a lot of the sounds suffer from kind of diminishing returns effect, so the closer you get to the sound of a real instrument, the more obvious the shortcomings and differences. But if you can put that niggle aside, there's plenty of physical 'phun' (sorry!) to be had here.

Korg Z1. Another physical-modelling synth which achieved 90% and a Platinum award in FM61. Surprisingly it's not as popular as the Trinity before it (or the Triton since) but it's still worth looking out for and trying out.

AAS Tassman. OK, so it's not a hardware synth, but it's still physical-modelling even if it is software. It also has some analogue-style features. We'll actually be reviewing this relatively new synth next month, but in the mean time, find out more from www.applied-acoustics.com.

Richard Wentk Future Music 09/00