

▶ SYNTH SECRETS series

▶ SOUND SYNTHESIS Books

▶ WANT MORE TIPS &amp; TECHNIQUE INFO? Visit the SOS FORUM

▶ SOS SOUND ADVICE Tips

▶ GLOSSARY: Tech Terms explained

▶ Learn Music Production Study Online with Berklee College of Music

## SYNTH SECRETS

### Part 15: An Introduction To ESPS And Vocoders

Technique : Synthesis

Published in SOS July 2000

 [Printer friendly version](#)

**Gordon Reid** turns his attention to the effects that can be achieved when subtractive synthesis components are applied not to the output from oscillators, but to real-world sounds -- such as human speech.

Over the past few months I've described oscillators, filters, amplifiers and envelope generators -- the devices that make up the architecture of a typical analogue 'subtractive' synthesizer. I've even demonstrated that you can use these to recreate methods of synthesis that are normally associated only with digital synthesizers. However, no matter how basic or how radical the concepts we've discussed, they have all shared one trait: the initial sound has been generated within the synthesizer. Usually, the synthesizer's internal oscillators have been the culprits, although self-oscillating filters have also reared their heads. Today, however, we're going to step beyond this limitation, and see how you can use external signals in subtractive synthesis.

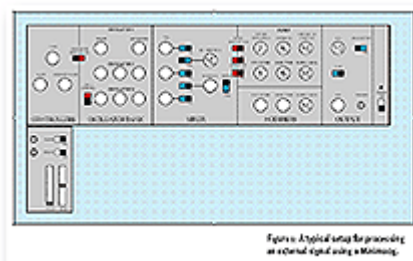


Figure 1: A typical setup for processing an external signal using a Moog.

### Getting The Signal In

Many of you will own contemporary analogue (and 'virtual analogue') synthesizers that offer signal inputs, but if you look at the back panels of vintage instruments you'll find that the idea is hardly a new one. If we return to the first non-modular synthesizers, we find that even the Minimoog offered a signal input as one of the five sound sources in its Mixer. Consequently, if you switched off the Moog's three internal oscillators and its noise generator, you could pass just the external signal through its filter and amplifier (see Figure 1).




Figure 2: A simple form of pitch/CV converter.

It may look very simple in principle, but even this arrangement has numerous subtleties and pitfalls. For example, the Minimoog has no Initial Level control for its amplifier. This means no sound can pass through the synth until you press a key. However, when you press a key, the amplifier's gain increases according to the contour generator's settings, allowing you to chop the external audio



Figure 3: A more sophisticated pitch/CV converter.

Monday 7th March 2005

Login here

Sub PIN or Email

Password

Remember me Stay logged in 

Login

**Forgotten your password?**  
[Request a reminder](#)

**Not registered?**  
[Register Now for FREE](#)

**No https access?**  
[Login here](#)

▶ [Show old-style menus](#)

Current Print Magazine: [click for Contents](#)

**Other recent issues:**

- ▶ February 2005
- ▶ January 2005
- ▶ December 2004
- ▶ November 2004
- ▶ October 2004



▶ [Recommended Reading](#)  
[PAUL WHITE Books](#)

**Screenshots too small?** Click on photos screenshots and diagrams in articles (aft August 2003 issue) to


into 'enveloped' bits and pieces.

In contrast, if the filter's cutoff frequency control is greater than zero, some low frequencies will pass through it regardless of whether you press a key or not -- though you still won't hear anything until you press a key, because the amplifier's gain is zero. When you press a key, the filter's dedicated contour generator will alter the cutoff frequency (or not, according to the Amount control) thus allowing a greater or lesser amount of the external signal's spectrum through the Moog. Furthermore, you can use the filter's resonance to emphasise parts of the external signal's spectrum, just as you would if the internal oscillators generated the sound. The advantage of this is obvious: whatever the nature of the external signal, you can use the Minimoog's filter to create new timbres.

Despite these abilities, the Minimoog is still very limited in the way that it can process an external signal. After all, there's no way to determine the pitch of the sound passing through it, nor can you determine when you hear the sound, other than by pressing a key. What we want is the ability to control the synthesizer by applying an external signal -- maybe using a guitar to 'play' it, or using your voice to make it 'speak'. So let's introduce the two 'modules' that make this possible: the envelope follower and, first, the pitch-to-voltage converter.

### External Signal Processing (1): The Pitch-to-Voltage Converter

Let's look at the idea of controlling and playing your synthesizer using a guitar. Oh yes... and don't forget, this is pure analogue technology we are considering, so we can't cheat by introducing a MIDI-to-CV converter! What we need instead is the aforementioned pitch-to-voltage converter. More commonly called a pitch/CV converter, this is a device that accepts a monophonic signal (ie. a signal with only one pitch present) at its input, determines its pitch, and then produces an appropriate control voltage at its output (see Figure 2).



Neat, isn't it? Indeed, if everything were as elementary as this, the worlds of electronics and synthesis would be very simple indeed. Unfortunately, pitch/CV converters can be fooled by stray signals and background noise, causing glitching. To overcome this, we add two sub-modules. The first of these is a simple audio amplifier called an input level or input gain control. This ensures that the external signal enters the converter itself at an optimal level. The second is a Slew Generator, or 'Lag Processor' (see Figure 3).

We haven't encountered slew generators before in Synth Secrets, but their effect might be obvious from the figure. In essence, they slow down transitions from one voltage to another, thus 'slewing' any abrupt changes over a period. (In a sophisticated system, the amount of slewing will be governed by a Slew Rate control.) As you can see from Figure 4, slowly changing signals (such as the low-frequency sine wave) pass unaffected, whereas signals with sharp transitions become




Figure 5: A sophisticated analogue pitch-to-CV converter.

open a Larger View window for detailed viewing/printing.

rounded. Interestingly, the single, sharp spike disappears completely. OK, this is an idealised description, but it gives you the general idea.

You should now be able to see that the slew generator is simply a low-pass filter, albeit one with a handful of specialised uses. (On most analogue synthesizers, it's a 6dB/oct low-pass filter with cutoff frequency variable in the range 0Hz to approximately 1kHz.)

You will most commonly encounter a slew generator in its role as a 'portamento' circuit, smoothing the transition from one keyboard-generated pitch CV to another. However, when it's hooked up to a pitch/CV converter, the slew generator's *raison d'être* is to remove the inevitable glitches that occur when the pitch detector loses lock on the desired signal. (Without the slew generator, the output CV would jump around wildly until lock was re-established.)

Moving on, we can enhance Figure 3 and improve the converter's performance still further. To be specific, we add a band-pass filter to create a narrow 'pass band' of accepted frequencies (see Figure 5). This reduces the risk of extraneous signals or high-amplitude harmonics confusing the pitch detector.

You might think that we now have everything needed to control our synthesizer using a guitar, but even the sophisticated converter in Figure 5 is not sufficient for this. Yes, it provides a monophonic CV that you can use to determine the oscillators' pitches, but it tells the synth nothing about the changing loudness of the notes (or even when they occur). For this we need something completely different...




Figure 7: The output from an envelope follower.




Figure 8: A more sophisticated envelope follower.

## External Signal Processing (2):The Envelope Follower

We have already determined that the pitch/CV circuit will provide a CV for the oscillators, so we now need something that can provide a CV to control the synthesizer's VCF and/or VCA. This something is an Envelope Follower (strictly speaking, a 'peak amplitude follower') -- a circuit that measures the amplitude of the positive peaks of the waveform.

If you're interested in the electronics (and because it's such an elegant solution) take a peek at the simple envelope follower in Figure 6. This works in a deliciously simple fashion. If successive peaks of the input signal are of increasing amplitude, the capacitor is charged up, and if the peaks are decreasing in amplitude the capacitor discharges at a rate determined by the value of the variable resistor. Of course, there is a small discharge between successive peaks even if the amplitude of the overall envelope is increasing but, if you choose your component values carefully, you can create the output shown in blue in Figure 7. As you can




see, this is remarkably similar to the true signal 'envelope'.

As with our pitch/CV converter, there's nothing stopping us from making the envelope follower more sophisticated, again adding an Input Gain Control and a slew generator to improve its performance. The latter of these will smooth out the 'bumps', making the output CV even more like the envelope of the signal (Figure 8).


### Putting It All Together

Now let's put it all together. Figure 9 shows an external signal -- which can be the output from a guitar, a microphone, a CD player, or whatever -- split and directed down two signal paths. The four blocks in the upper path are those shown in Figure 5, whereas the lower path is provided by Figure 8. Note that in the following figures I have used blue arrows for audio signals and black arrows for control voltages. I hope that this makes things clearer.





Now look at the CVs' destinations. Clearly, the pitch of the input is controlling the pitch of the oscillator, whereas the loudness of the input is controlling the cutoff frequency of a VCF and the gain of a VCA. The envelope follower is, therefore, replacing the contour generators that you would find in a conventional configuration. So, while the oscillator is providing the basic timbre of the output, the incoming signal is determining the pitch as well as articulating the new sound.

So there we have it... the perfect way (in theory) to control your synthesizer using an external signal such as a guitar. But while Figure 9 offers some interesting musical possibilities it is limited in one important way: the external signal has a very limited ability to modify the timbre of the output. In fact, the only way to affect the timbre is by patching the pitch CV or loudness CV to the cutoff frequency input of the filter. So, is there a way to make your external signal determine both the loudness and the timbre of the synthesized sound? Of course there is, and you probably have an example tucked away in one of your synthesizers or effects units. It's a vocoder.



### Introducing Vocoder

Imagine playing a percussion track -- from tape, or from a rhythm machine -- into an envelope follower. As you might guess, the follower will generate a succession of decaying pulses that you can use as envelopes or as triggers. If you then play another external signal such as a sustained organ chord through a VCA that is controlled by a contour generator that is itself triggered by the envelope follower... well, I'm sure that you get the idea. The organ now 'plays' the chords rhythmically and in perfect synchronisation with the percussion track, whether you lift your hand from the keys or not! (See Figure 10.)



Extending this idea further, you could attach a microphone and use your voice to articulate your keyboard playing -- allowing sound through the VCA when you say (or sing) words, but creating silence between words (Figure 11). Unfortunately, the envelope follower is simply following the amplitude peaks of what you say or sing... the words themselves are completely irrelevant. To put it another way, it doesn't matter what frequencies are present in the signal, it is only the total amplitude that determines the output.

To understand how great a limitation this can be, imagine that you have a signal with no frequencies present below 1kHz but so much signal above 1kHz that it registers as the maximum that the circuit can handle. In this scenario, the output from the envelope follower will also be a maximum. Now let's turn this example on its head: there is no signal above 1kHz, and maximum below this frequency. The envelope follower will again produce a maximum CV output. Ouch! Despite the fact that the signal in the first example contains only high frequencies while the second contains only low frequencies, the result is the same.

So, here's an idea: let's split the signal into two paths, and place two filters before a pair of envelope followers (see

Figure 12). In this scenario, low-frequency signals cause Envelope Follower 1 to generate a CV, while high-frequency signals cause Envelope Follower 2 to generate a CV. If we now send these CVs to a pair of VCAs, we can configure our synthesizer so that its response to incoming signals is 'frequency-sensitive'. For example, we can use the frequency content of the input signal to determine the relative amplitudes of the signals generated by two oscillators (Figure 13). Intriguing, yes?

However, we don't necessarily need to use multiple VCOs to generate interesting effects. Instead, we could take a complex signal and pass this through a second bank of band-pass filters to split groups of harmonics into a number of separate signals, each occupying a defined band of frequencies.

Now look at Figure 14. As before, the envelope followers raise and lower the gains of the appropriate VCAs in sympathy with the frequency content of the input (this, by the way, is a 'Spectrum Analyser'). However, instead of controlling the amplitudes of signals generated by independent oscillators, the VCAs now control the amplitudes of the harmonics in each of the frequency bands.

This is a hugely important result. Imagine that the signal presented to the envelope followers is your voice (we will call this the modulator) and that an oscillator generates the signal presented to the second bank of band-pass filters (we will call this the carrier). In this case, the carrier provides the basic tonality of the output, but the modulator determines its frequency content and amplitude. In other words, the modulator articulates the carrier. Voila! We have designed a vocoder.

### More Advanced Vocoding




Figure 14: A simple 4-band vocoder.




Figure 15: Using external signals for both carrier and modulator.

Many vocoders generate an internal carrier wave -- usually a sawtooth because it's the most harmonically rich of the simple waveforms, although a rounded pulse wave would be more suitable for vocal sounds, because this is closest to the raw waveform produced by the human vocal chords. An internal white noise generator is another good choice for speech resynthesis, because this contains all the frequencies in the sound spectrum. However, better vocoders offer two external signal inputs -- the first for the modulator, and the second for the carrier (see Figure 15).

Of course, there's nothing to say that the modulator must be a vocal signal. You can use guitars, other keyboards, any acoustic instrument (provided, of course, that you convert its sound into an electrical signal using a microphone or transducer) or even the outputs from CD players and the radio. Likewise, the carrier can be any signal. Indeed, using the same signal as both carrier and modulator provides one of the most interesting vocoder effects. If you want an even more radical configuration, you could even place pitch/CV converters and oscillators at one or both of the inputs.

Before finishing, I would like to describe two final enhancements to our vocoder. If you look at the front panels of units such as the Roland SVC350, you will see a bunch of faders. These scale the CVs produced by the envelope followers, allowing you to tailor the vocoder's response, accentuating or attenuating the outputs of specific bands. Less visibly, but no less importantly, some vocoders (the Roland VP330 springs to mind) replace one of the band-passed carrier signals with a noise generator. This is very important for correct articulation of sibilants and consonants -- the short, noisy sounds (for example, the letters 'd',



lips and tongue rather than the vocal chords. We can add the front panel controls and noise generator as shown in Figure 16.

OK, even with these enhancements, there are limitations. In particular, the low resolution of the band-pass filters -- typical roll-offs are 6dB/oct or 12dB/oct -- ensures that the output articulation retains only a remote

### The Genesis Of The Vocoder

A research physicist named Homer Dudley invented the Vocoder (the VOice EnCODER or, according to some commentators, the Voice Operated reCORDER) when he was working at Bell Laboratories, New Jersey, in 1939. He developed the machine as a research device for experimenting with audio compression, primarily to improve the voice-carrying capabilities of his employer's telephone lines.

Following WWII, Dudley visited Bonn University and met a chap named Werner Meyer-Eppler who, at the time, was the Director of Phonetics at the University. Meyer-Eppler recognised the relevance of the vocoder to electronic music, and subsequently used it within a number of compositions that would eventually become the 's' inspiration for the German 'Elektronische Musik' movement. Understandably, the sound fidelity of Dudley's creation was limited, primarily by concept has remained unchanged to this day.

semblance of the original. But with sufficient bands (10 plus noise is enough) a vocoder is easily good enough for 'Mr Blue Sky'!

So that's it. Easy, huh? Well... yes it is. When Walter Carlos recorded the score for *A Clockwork Orange*, he used off-the-shelf Moog filters, oscillators, envelope followers and VCAs to resynthesize (ie. vocode) the choral sounds. Given the cash, space and patience, you could do the same. So here's this month's Synth Secret...




Figure 16c: Modifying the vocoder's response.

*If your synthesizer has an external signal input plus an envelope follower and a pitch/CV converter, it is a much more powerful and flexible instrument than it would otherwise be. You should experiment with them!*

Alternatively, you could simply stick with conventional emulations of woodwind, strings and brass, or continue to create silly bleeping and squelchy noises. But where's the fun in that? **SOS**

Published in SOS July 2000

top ▲

[Home](#) | [Search](#) | [News](#) | [Current Issue](#) | [Articles](#) | [Forum](#) | [Subscribe](#) | [Shop](#) | [Readers Ads](#)  
[Advertise](#) | [Information](#) | [Links](#) | [Privacy Policy](#)



SHOP HERE

Buy Great Products  
At Great Prices!

In association with Amazon.co.uk

**Sound On Sound, Media House, Trafalgar Way, Bar Hill, Cambridge CB3 8SQ, UK.**  
 Email: [sos@soundonsound.com](mailto:sos@soundonsound.com) | Telephone: +44 (0)1954 789888 | Fax: +44 (0)1954 789895

All contents copyright © SOS Publications Group and/or its licensors, 1985-2005. All rights reserved. The contents of this article are subject to worldwide copyright protection and reproduction in whole or part, whether mechanical or electronic, is expressly forbidden without the prior written consent of the Publishers. Great care has been taken to ensure accuracy in the preparation of this article but neither Sound On Sound Limited nor the publishers can be held responsible for its contents. The views expressed are those of the contributors and not necessarily those of the publishers.

Web site designed & maintained by PB Associates | SOS | Relative Media

Problems with this site? Email [webmaster@soundonsound.com](mailto:webmaster@soundonsound.com)