The Technology of Auto-Wahs / Envelope-Controlled Filters

The goals:
This paper is intended to explain how a particular category of effect works, including the variations in design that are currently in production, or have previously existed. The desired outcomes are that the reader will:

a) know how to make any effects they own behave more in the way they want (including injecting a note of surprise or uncertainty),
b) have some idea about whether a given commercial effect (including vintage units) will do what they want,
c) have some idea whether a DIY project can do what they want, and have some intuition about how it will sound before they invest the time and effort making it,
d) have some idea about potential mods to existing DIY or commercial units
e) have some new concepts in their arsenal to use when thinking about other types of effects.

I'd like to thank ***RG Keen*** for his diligence and graciousness in drawing and providing some of the schematic fragments included. Readers are also recommended to take a look at the excellent 1980 article on ECF's from **MIX** magazine, co-authored by legendary studio guitarist ***Elliot Randall*** and the Les Paul of ECF's, ***Mike Beigel***, and reprinted at Elliot's web-site: `www.elliot-randall.com/ecf.htm`. The article overlaps a bit with this one, but gives a nice view of what a deluxe unit would do. They also have some nice graphics which flesh out some of the concepts discussed here (though only touched on there), so it is a recommended complementary read, if only for the pictures. Thanks to ***Aron Nelson*** for the link.

What they do:
Auto-wahs, more properly called *Envelope-Controlled Filters* (ECF), shape the tone of your instrument (the filtering part) in response to how loud the sound is going into them (the envelope part). They are often called auto-wahs or auto-filters to distinguish them from conventional foot-operated pedals, since they can be used in a set-and-forget manner. Although they can certainly mimic some of what a wah-wah does, they are substantially more flexible in many ways, and have a broader palette of tone shaping capabilities than a wah-wah does. Like anything automatic, however, they don't always respond exactly the way you want them to at any given moment.

Some early history:
They were once the darlings of funk and disco music, and seemed to fade away from popularity as disco's star descended (though they never quite disappeared). However, everything old is eventually new again, and the return of disco-like tonalities and timbres with contemporary dance music has brought back filtering with a vengeance, and with it the envelope-controlled filter. The connection with funk and dance music is *not* purely coincidental but is a natural alliance. Both wah-wah's and ECFs allow the guitarist to accentuate the beat of the song, and in effect, act like an additional percussionist. (*If you think of it, some types of percussion are very similar in their action to an ECF, in terms of the relationship between sweep and beat. For example, the first analog drums - the Syndrum being a prime example - and traditional Indian tablas both have the capacity to sweep across a range of emphasized frequencies to accentuate or articulate/complement the beat.*)

The origins of the ECF lie in the early days of modular analog synthesis. Long before FM and sampling, the most common form of electronic music synthesis was subtractive synthesis. That is, one started out with a standard oscillator-produced waveform (square, ramp, triangle, etc.) that was relatively rich in harmonics, and used filtering to remove harmonics and provide timbres in between the original waveform and a pure sinusoidal waveform. The filters and amplifiers were controlled by envelope-generators that produced a rising and falling control voltage which could make the filters have different cutoff frequencies and the amplifiers modulate the volume. The envelope voltage, in turn, was initiated by each keyboard keypress, such that the cutoff frequency of the filter varied in a predictable manner. By treating each note as a reproducible musical event, one could attempt to simulate things like plucked or blown notes.

Of course, what impressed the popular mind most was the most blatant and exaggerated filter use, and there can be no better example than the signature 4-pole lowpass filter of the Minimoog; the "*bwow*" that set the standard. What distinguished this particular tone from what we could already do with wah-wah pedals was the fact that it could be fast and note-specific, and ***very*** resonant, and guitarists wanted it. Some synthesizers (e.g. the Korg MS-20) started to come with external signal processing capability, which would allow them to apply their filter(s) to something other than the keyboard voltage and on-board VCO's. A module referred to as an "envelope follower" was part of this. This module followed the amplitude of whatever you put into it, and attempted (with varying degrees of success) to put out a voltage that corresponded to the natural amplitude envelope of the original signal source.

This didn't automatically put guitarists in the same arena as keyboard players. Keyboards had the advantage that they were hardwired to detect note onset, and could generate a trigger pulse that initiated the envelope generators controlling the filters and amplifiers reliably, consistently, and quickly. How would guitars perform the same function? Unlike electronic keyboards, guitars are used to produce notes that are more likely to go electronically undetected, by virtue of their softness or slurring. Making any circuitry sensitive enough to detect the onset of such notes would also make them sensitive enough to produce false triggering. We didn't want notes to be triggered by events like string or fingerboard noise. Enter envelope-controlled filters. Although it might not be possible to reliably initiate a trigger and envelope voltage with each note strummed or plucked on a guitar, envelope control of the filter was enough to make a satisfactory synth-like sound.

A brief history of the early days of ECF's can be found in the history of Mu-Tron and Musitronics at Mike Beigel's web-site. The article is an excerpt from ***Art Thompson***'s excellent book "**Stompbox**", and can be found at: `www.beitec.com/stomp.htm`. Mike's patent documents for the Mu-Tron III can be found there.

The Basic Characteristics:

A block diagram of an ECF is shown in **Figure 1**. This is not every conceivable ECF ever produced, but does capture most of them. What we see is that there is usually some kind of combination buffer/booster stage at the input. This will bring the signal level up to a point where it is likely to be easily detected by the envelope circuitry, and well above any extraneous signal sources, such as hum. From the buffer/booster stage, it is split and sent to the filter, and to the envelope detection/follower circuitry.

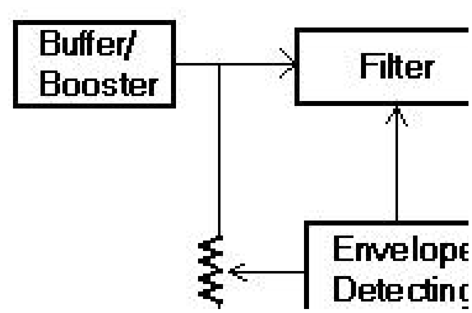

Since you never know how big the incoming signal will be, there is usually a sensitivity control between the input stage and the envelope detection circuitry. This will allow you to jack up the sensitivity if you have weak pickups or a less imposing picking style, or turn down the sensitivity if your signal is very hot or if you just want subtle changes despite intense strumming. In some cases (e.g., the Mutron III), the adjustment comes in the form of a gain control for the first stage, rather than a fixed gain followed by an attenuator, as shown here.

The envelope-following circuitry provides a unipolar signal (i.e., is either all negative or all positive, relative to ground) that is proportional to the amplitude of the incoming signal. Typically, the envelope follower will have some components that shape the envelope signal in some way, either slowing down the attack or onset, or adding some lag so that the envelope does not descend too quickly. Think of the envelope follower as being a bit like the record level meter on your tape deck, or the tachometer on your car dashboard. It responds to the overall average over a moderately brief period, rather than the absolute value at each fraction of a second.

The output of the envelope follower is connected up to some sort of control element. The control element, in turn, is part of a filter, whose bandwidth parameters or centre frequency changes as that element changes. In effect, the filter operates like someone twiddling a knob on a note by note basis. In some cases, the filter itself may be reconfigured to different types, or it may have control over its resonance or selectiveness. In a few odd cases, the envelope can be combined with other sources of modulation, such as input from a footpedal providing a second control voltage, an LFO, or some master control voltage that several devices are synced to. It is certainly possible to have a more open-ended design, but here we veer off into modular analog synthesis. Most commercial stompbox ECF's are closed systems, generally confined to the on-board envelope follower, and in a few cases (e.g., the BOSS Dynamic Filter), an external expression pedal.

Producing The Envelope Voltage

You might think that the filter is the guts of an ECF, but the true heart and soul of any ECF is really the envelope-follower. What gives different ECF's their unique feel is fundamentally determined by the characteristics of the envelope detection circuitry. If you're planful, and intent on using FX in a tactical way, you may try to get the filter to sweep and emphasize in synchrony with the beat. The time it takes to get from where you are to where you want to be depends on how effectively the envelope detector can respond to rapid changes. If the envelope follower "hangs up" for a while at the extreme end of its output range, then it may not be ready to respond to your upcoming pick strokes either. It's not just a question of beat synchrony either. Many players use "opening up" of the high end as a means to inject more emotion and expression, whether it's done by lifting the butt of your hand for certain notes, picking harder to get more distortion, or picking harder to get the filter to sweep up a little farther. If the filter is not ready to respond to your intentions, then it will be exasperating and impose too much on your style. Depending on the envelope following characteristics, a given ECF may feel more "urgent" or "lazier" than another. So, while you can squeeze a filter sweep out of any old ECF (assuming levels are right), there is little guarantee that the way in which it responds to notes and silences is ideally suited to the kind of material you play *all* the time. An ECF with lousy envelope following characteristics will likely be used very sparingly, like a ring modulator. An ECF with envelope response that is just what you want can easily integrate into your style and stay on for as long as many would use other mainstays like phase shifters, wahs, or flangers.

There are more and less complex ways of extracting an envelope signal. In the attempt to get things to run off a 9v battery for as long as possible and keep production costs down as much as possible, designers and manufacturers tend to use compromise designs, and keep parts count down. As a result, any given ECF will likely use an envelope follower that has certain design compromises. The bottom line is that unless you have a very consistent playing style, it is unlikely you can be happy with one single ECF, unless. Much like with fuzzes, polygamy is the order of the day when it comes to ECF's, partly because of the variety of filters out there (and hardcore analog modular folks will tell you that different filter designs *do* sound different), but also because of the variety of envelope follower traits. Even when the same basic envelope follower design is used (and 2 or 3 basic designs account for a large portion of what's out there), change of a few basic component values can yield very different response patterns.

Extracting an envelope is not as easy as you'd think. What kinds of problems are faced in accomplishing this task? First, the output signal of the envelope follower has to be of sufficient amplitude (current OR voltage) and appropriate polarity to drive the control element. Not much problem there. In most cases, this simply means providing appropriate gain within the envelope follower itself, or feeding it with an already boosted (or attenuated) signal. Gain setting can usually be adjusted by a single, easily identified resistor or pot. Polarity is set by either inverting the envelope signal somewhere along the way, or by using diodes to select only signals of a specific polarity, or both.

The envelope has to have appropriate rise and fall characteristics. In a sense, envelope followers almost always lag behind the signal they are detecting because their job is essentially to *describe* the average signal, rather than simply follow it like an amplifier or buffer. Unless the envelope follower has controls built in for adjusting attack and decay, you'll find that the amount of lag or instantaneous quality will vary from model to model. Although there is considerable variation among commercial products with respect to attack and decay characteristics, most will provide an envelope signal that responds with maximum swing over a period of 50msec or less, and drifts back to baseline over a period of 500msec or less.

In general, having a fast attack time is desirable... up to a point. For a wide variety of reasons, faster attack times help to define the relationship of the effect to the beat of a song. Faster attack allows for synchrony with the beat, since it is easier for the player to predict where and when maximum sweep will occur (and end) in relation to the beat. Besides, slow attack is easy to do with a wah; fast attack isn't. On the other hand, having only a whip-fast attack available is not the ideal state of affairs either, since sometimes you want a gentler attack that mimics a wah (freeing up your foot for other duties), or a fast attack may sound too insistent with a slower song. Though there may be some room to stretch with respect to decay, excessive decay tends to ruin the degree of definition of the effect, unless one is into nonrhythmic music, or uses an ECF strictly for "rolling-wash" type rhythm backgrounds (which a slow decay can do *very* nicely). I may be wrong, but my gut sense is that while attack differences on the order of 3-5msec can make a difference in one's playing, decay time differences of up to 50msec may be negligible, with respect to their impact on playing. I mention this because it may be sufficient to have a slow/fast decay time switch, but more necessary to have a tunable control for attack.

Apparent attack: It's important to distinguish between *apparent* attack, and actual attack time. Apparent attack? Think of this as the amount of perceived sweep per unit of time; khz/msec. Farther sweep over the same amount of time will seem like a faster attack, with more emphasis. Conversely, having farther to sweep back down will seem like a slower decay. Perfect case in point is the Dr. Q, which has two filter types, selectable by switch. Although it's the same envelope follower and envelope signal in each case, one of the filter types sweeps much higher than the other, and appears to have a much longer decay time than the other. In fact, it doesn't. It falls back at the same rate, but just has farther to go. Apparent attack is partly a function of the control element used (some respond more readily than others), and the filter design. Sometimes even small changes to a basic filter type can change its sweep ratio by a factor of 10 or more. And, of course, the sweep distance travelled is also a function of the signal fed in, and the drive or sensitivity settings of the device. A subtler and simple "bAh-oo", coming from decreased pick attack, will sound like a less pronounced attack than a "b-WOW" within the same time frame. I couldn't vouch for it, but I imagine that an ECF with greater resonance or emphasis around the cutoff frequency may also appear to have a faster velocity of attack because the frequency sweep is a little more "in your face". I may also be wrong about this one too, but ECF's that can sweep downward as well as upward, seem to "sound" faster when going downward. (At least to me. Of course, I also see dead people...eating fudgsicles and doing the Lindy hop...so make of it what you will.).

Some ECF's have a way of adjusting the apparent attack by adjusting the "start" and "stop" points of the filter sweep, such that the filter can be made to sweep over a narrower range within the same period. Some units, such as the Electro-Harmonix Microsynth for guitar or bass, have such controls labelled in exactly that manner. In other cases, you may see controls for "initial frequency" and "sensitivity" which can be tinkered with to determine where the filter starts sweeping from, and how far it sweeps from that point. Many of the more common commercial units, however, lack such controls. For example, the Mutron has a "range" switch to shift the overall range it sweeps across, but within each range, the filter always sweeps from the lowest possible centre-frequency in that range to the highest point possible (or highest to lowest if set to downward drive), given the drive setting and picking strength. Modifying a Mutron to adjust the sweep depth is discussed in the **Useful Mods** section towards the end of this paper.

Ripple and its sources: All envelope followers do their job best when they output a signal that provides smooth voltage or current changes that are free of what is called ***ripple***. When an AC-to-DC adaptor does not provide proper rectification and filtering to turn 60hz AC into steady DC, we hear this as hum superimposed on the audio signal of whatever the adaptor is plugged into. Filter the AC better, and the hum is reduced. Envelope followers are essentially operating like AC adaptors, converting the AC of the audio signal into something close to a DC voltage. When audio frequency content comes through as part of the envelope signal, we tend to hear this as a distortion-like sound, sometimes almost a hint of a ring-modulator type of sound. We hear it like that because it literally IS modulating the filter. Instead of going "bwow", the filter ends up going "bwaggeta-wiggeta" or "brrwowrrrrr" (insert Scottish brogue here) as the ripple causes momentary fluctuations in the filter center frequency at audio frequency rates. Ripple is most likely to be heard during the decay phase of the filter sweep. Partly because the attack is too fast for ripple to be heard, and partly because the decay portion won't occur unless you stop playing, so your attention is easily focussed on the modulation. It's a bit like the way a need for a wheel alignment becomes much more apparent as your take your foot off the gas and grind to a wiggly halt!

What is the relationship between envelope follower characteristics and ripple? If an envelope follower is set for an attack time of 10msec, then it will be unable to rise fast enough to respond to track any envelope fluctuations faster than 100hz (1000/10 = 100). So any ripple in the envelope signal above 100hz simply won't impact much on the filter. Conversely, if the decay time is 330msec, then the envelope follower won't generate much ripple above 3hz during the decay phase. (Nice examples of this are illustrated in the Randall-Beigel article bookmarked at the top of this paper). The general principle is that slower attack time and longer decay time equals less ripple. The problem they create is that excessively long times impede responsiveness and playability. The trick is to identify attack and decay times that balance playability with ripple rejection. The choices that designers make in this regard is part of what gives every commercial unit its own flavour. For example, some control elements are more sluggish by nature, and so can tolerate a bit more ripple in the envelope signal without audible effect. In effect, they provide post-follower envelope-signal filtering. This gives the designer the option to go with faster attack and decay times in the follower itself without penalty.

The attack and decay-times can often be easily measured from the existing components, as long as you know how to identify them. The principles described here can often be applied to other units. How do you spot them? Look for diodes, since they will be used in every type of envelope follower. Then, look for the first small-resistor and polarized-capacitor combination after the diodes, as in **Figure 2**, on the right. The resistor will typically be less than 1k and the cap will generally be something between 3u3 and 100uf. For example, the Dr.Q/Quack uses a simple half-wave rectifier, wherein the negative half-cycle of the audio input is essentially lopped off by two diodes between the op-amp and control element (transistor). In between the two diodes, you'll see the resistor-capacitor combination. The 100ohm resistor sets the attack time, and the 10uf capacitor sets the decay time. If you look at the Mutron, you'll see almost exactly the same thing.

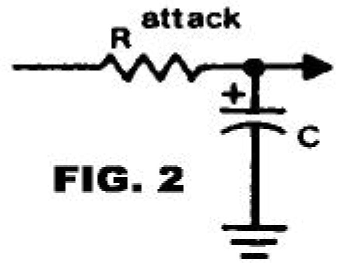

FIG. 2

Calculating the attack time is fairly easy. Simply multiply the resistance (in megohms) by the capacitance (in microfarads). So, using this formula, the attack time for a Dr. Q/Quack using a 100 ohm resistor and 10uf cap is .0001 * 10 = .001 seconds = 1 msec (although, quite frankly, it feels slower than this). For the Mutron, the component values yield an attack time of .00033 * 4.7 = .0016 seconds = 1.6 msec. Calculating the decay time is somewhat more complicated. Ideally, decay time, or the time taken to discharge the capacitor in Fig 2, will depend on the resistance in parallel with the capacitor, and follows the same formula (e.g., 10uf and a 10k resistor gets you a 100ms decay/discharge time). Most commercial units I've seen lack such a resistor, so the discharge/decay time is more a function of the internal resistance of the cap, and the input impedance or resistance to ground of whatever follows it. Maybe you can do it, but I sure as hell can't figure it out. In all instances using this configuration, though, doubling the values of each component will double their respective times, and halving their values will halve their times, so you can season to taste based on that formula. On my own Dr. Quack, I subbed a 50ohm resistor in series with a 1k pot to get a range of attack times from very fast to much slower than stock. Although the 10uf cap provides a nicely responsive decay time, I found it resulted in too much ripple for my tastes, so I added another 22uf cap with a toggle switch to put it in parallel with the existing 10uf cap. This provides a wide range of feels. Note that in envelope followers using this simple RC arrangement for setting the time constants, increasing the cap for decay will also increase attack time, since it interacts with the attack resistor.

Readers with even a bare minimum of electronics knowledge will realize that this RC pair is your basic single-pole low-pass filter, so OF COURSE it will result in there being less ripple to the envelope signal. What should also make sense is that the larger the decay-setting capacitor, the longer the decay time, and the less audible the ripple. The effect is entirely analogous to cranking up the filter capacitor in a power supply to get rid of AC hum. As a side-note, this envelope-follower configuration is often used in many contexts beyond ECF's that use side chains to control things, such as compressors, limiters, noise expanders, duckers, etc. The same rules for identifying relevent components and changing their values can often apply there as well.

It's worth noting that there are several different sources of ripple. One fairly straightforward source comes from inconsistencies in the overall amplitude of the guitar signal fed to the envelope follower. A "well, duh!" example would be feeding an ECF with the output of a tremolo set for bubbly fast speed. How would the follower be able to provide a smooth envelope corresponding to strength of picking? A less obvious example would be chords where one or more strings may accidentally be muted or activated, or both. The resulting signal would have variations in level superimposed onto it. A third unsuspecting source would be "beats" created by old or mis-shapen strings (I prefer high-profile frets, and this frequently results in "dented" high E's and B's that don't vibrate smoothly), or beats created by two or more dissonant notes. You may hear them as "sour", but the envelope follower hears them as a "wiggly" envelope. These sources of ripple come from the original audio signal, which means you can improve tracking and reduce ripple without having to tamper with the insides of the ECF itself by simply attending to the input signal quality. For instance, while compression removes some of the dynamics of the signal (generally frowned upon; see section on **Dynamics** below), eliminating some, but not all, of the dynamics can work wonders for smoothing out audio ripple and making an ECF behave more predictably, the same way it can improve the tone of a distortion device. As always, everything in moderation.

A second class of ripple comes from the envelope detection process itself. In **Figure 3**, you'll see the hypothetical output of a quick and dirty envelope follower in response to a brief audio signal (assume you are looking at a brief instant of an audio frequency triangle wave, not an LFO). Half-wave rectification results in only positive peaks being conducted. This produces a signal that gets bigger and smaller, in reference to a stable ground. Passing only half of each wave, however, is not the same as producing a DC voltage. The RC combination, discussed earlier, acts as a low pass filter to smooth out the envelope signal, and provides greater smoothing the higher the frequency content fed into the envelope follower. Fortunately, as guitar players, we tend not to play many notes with fundamentals below 150hz or so, so ripple produced by half-wave rectification is not a problem across much of the guitar's range. Obviously, for bass (or 7-string axeslingers) it's a much bigger problem. Once again, you CAN "fix" the problem by having a sluggish envelope follower, but that's not why you bought an ECF, is it? A better solution is to use full-wave rectification (FWR), which results in the negative (or positive, if you feel like it) half-wave being "folded over"and combined with the other half-cycle, as shown in **Figure 3**.

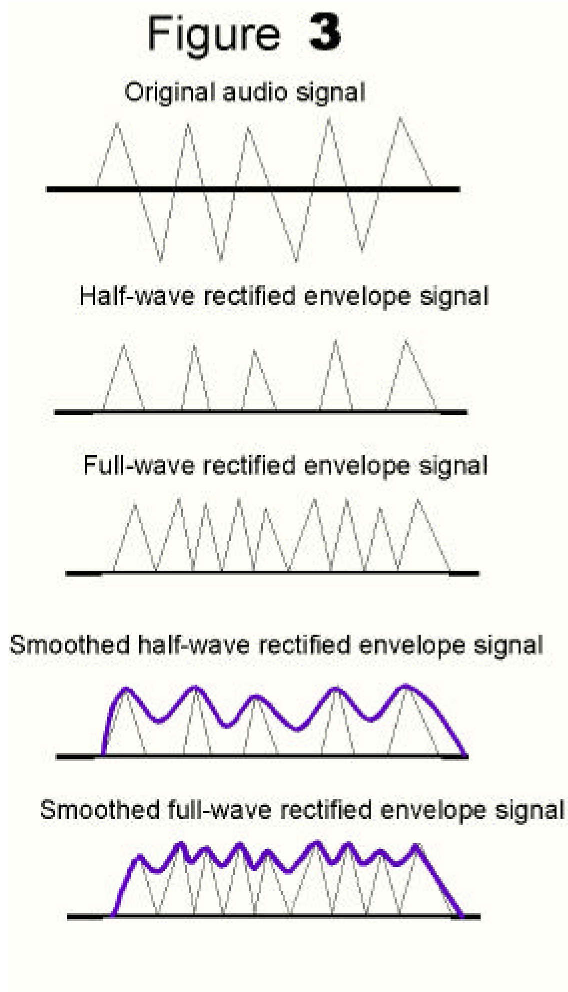

In effect, this produces an envelope output which has ripple at twice the frequency of a half-wave rectifier. This requires less lowpass filtering to turn the rectified signal into something audibly closer to DC, which means you get to keep your fast attack and decay, without inheriting the ripple that normal accompanies them. Lest you think that the depiction in Fig 3 is accurate, bear in mind that the smoothness of the transitions will still depend on the attack and decay settings, and the frequency of the notes.

Sadly, in an effort to keep costs down and package size small, few companies use FWR, but there are other things you can do. All other things being equal, audible effects of ripple are also most likely to occur if the control element is of the type that can respond very fast. Photocell-based ECF's are relatively immune from audible ripple, simply because many photocells tend to return to their MAX-resistance/OFF value more slowly than they swing to their MIN-resistance/ON value, imposing a longer decay period without additional circuitry. As noted above, upping the decay cap to smooth things also results in increased attack, so having a device which permits smoothing out the decay, independent of the attack, is kind of nice. The same envelope follower circuit used with a much faster-responding transistor as the control element may well result in considerable audible ripple unless the decay period is deliberately smoothed out with additional circuitry or larger value capacitors. The pros and cons of various types of control elements are described in more detail in another section.

Side-chains and other envelopes:
"Side chains" are circuits that provide an output which, in a sense, *describe* an input signal. If the input signal is like *this*, then the side chain signal will be like *that*. An envelope extractor or detector is a side chain that describes (i.e., is proportional to) the amplitude of the input signal. Normally, the ECF derives its envelope signal from the instrument you are playing, but there is often no reason why that has to be the case. A perfect analogy is the use of an external input with a noise gate. Normally, in noise gates, there is an electronic switch or fader, which is governed by the envelope of the instrument being fed into it; when the signal exceeds a predetermined level, the gate opens up. However, the gate (and many provide for this option) can be "keyed" by other signals. For example, in the classic David Bowie tune "Let's Dance", producer Nile Rodgers mixed down the entire horn section and ran it through a gate, using his guitar to control the gate. The result had an interesting feel; not quite horns, not quite guitar. In many (though not all) instances, modifying the ECF to accept external signals is quite feasible. This would permit the filter to be swept by a wide variety of sources other than your instrument. You could control the filter with your voice (via a mic), a drum machine, a radio, the rest of the band, or what the hell....a wireless mic hidden in the washroom of the bar you're playing in. Just in case you get some wild ideas, this is NOT the same thing as a vocorder. Vocorders use the *frequency* content of the side chain signal to determine the frequency content of the input signal. An externally keyed ECF would vary the frequency content based on the *amplitude* of the input. In fact we can construct a nice little 2 x 2 matrix here that shows the relationship between several different kinds of devices:

	Modify input amplitude	**Modify input frequencies**
Use side chain amplitude	noise gate, ducker, compressor, limiter	keyed ECF
Use side chain frequencies	de-esser	vocorder

Figure 4 shows a practical example with Jack Orman's *Dr. Quack* adaptation of the Electro-Harmonix Dr. Q. The DQ consists of a bandpass filter, an envelope follower, and a transistor linking the two. The Sensitivity control takes its signal directly from the input shared with the filter, but there is no reason why it has to. Connect the input of the pot to a jack and feed it an external signal as in **Figure 4,** and the external signal will trigger the sweep of the filter. You may want to put a cap ahead of the pot to filter out any possible DC coming from the external source (although most of the time there won't be any, but it's still good practice). Depending on the amplitude of what you feed in, you may also find that the gain in the envelope follower is insufficient. If that's the case, you can up the value of the feedback resistor from 2.2meg to 2.7 meg or even 3.3meg, using the trimpot to make fine adjustments. You can also stick a cap in parallel with this resistor to provide a little bit of additional lowpass filtering to reduce ripple. A value of 220pf in parallel with 2M2 will get you a rolloff starting around 330hz.

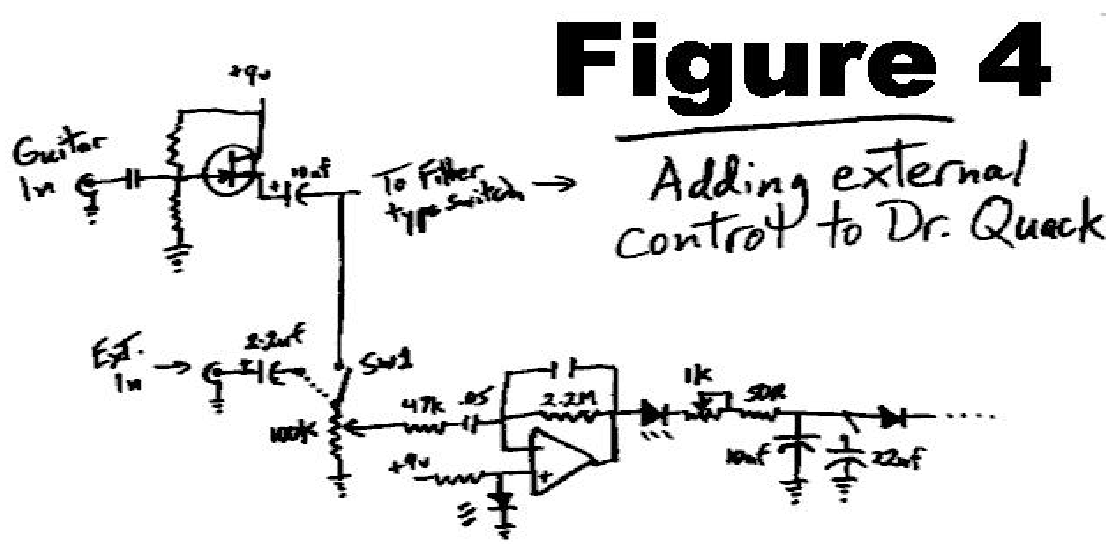

Jack Orman's improved version of the DQ, called the Dr. Quack, leaves the filter essentially untouched, but adds a FET buffer ahead of the filter and envelope follower (shown in **Fig. 4**), in addition to fixing up a few other things. The 100k Sensitivity control connects up to the output of the buffer, but is still able to be disconnected from there and driven directly from an external input, without screwing up the filter portion. The same general logic can be applied to virtually any ECF with a Sensitivity pot between the input and envelope follower. The only qualification would be that if the normal input stage provides some gain, that gain will have to be made up somewhere else when disconnecting the Sensitivity pot, perhaps requiring a small op-amp or single transistor stage. Normally, there would also be some concern about DC entering the envelope follower, but the .05uf cap at the input to the op-amp filters that out, so we're okay here. You'll note that I've also added the suggested envelope control mods mentioned earlier.

The Mutron III, which uses an adjustable gain stage at the input for the filter and envelope follower, rather than a fixed gain stage and Sensitivity pot, as in the Dr. Quack, is a different situation. In this case, adapting it for side chain use would necessitate a few other changes. First, since we have no idea of the level of the external source, good practice would suggest a gain stage immediately following the external input. The Mutron normally does all the signal boosting at its input stage, so if we want something for the envelope follower to follow, we'll have to do likewise with the gain stage for external inputs, to bring it up to a reasonable level. You can do this by simply creating a duplicate of the input gain stage (see **Circuit Analysis** below), connected to a separate input jack, followed by a similar type of switching arrangement, as per what's shown for the DQ. You'll end up with a unit that has two paths (filter and envelope follower), each preceded by its own pre-amp to optimize functioning.

One of the bonuses of this extra stage for external input is that we can do some lowpass filtering to try and reduce envelope ripple without having to be concerned about signal quality for audio purposes. The 10pf feedback capacitor you see in the input stage (see the schematic), in parallel with the 120k feedback resistor, normally forms a lowpass filter with a very high rolloff (starting well above the hearing range) to keep stray RF signals at bay. Given that this particular input stage doesn't have to do double duty as a booster for the audio and envelope paths, though, we have the liberty to take a bite out of the ripple at the booster stage, before it even gets to the envelope follower. If you increase the 10pf cap to 4700pf, the rolloff becomes roughly 280hz (though at 6db/octave, which isn't particularly steep). This still leaves an awful lot of room for audible ripple, but on the other hand, it still permits a reasonably fast attack time. In tandem with the RC combo in the envelope follower, a fair amount of envelope ripple can be removed, yielding a potentially smoother sound.

If you want a deluxe external input, consider installing a toggle switch to select between different values of the feedback cap, to introduce different degrees of slurring/filtering of the external input signal. Expect a tradeoff between rolloff frequency and attack capabilities. Readers should note that a common tactic in modular synthesis is to use a noise source, severely lowpass filtered (e.g., rolled off above 20hz), as a source of random modulation. There is no reason why the same logic can't be applied here, especially if the external signal is not a note-based type like the guitar itself (e.g., radio or CD). If you have the option of selecting notch filtering instead of bandpass or lowpass, a sluggish response and low resonance can deliver a nice ethereal tone. The insertion point of the external signal is at the positive end of the 2u2 capacitor at the input of the envelope follower stage (see diagram below), which normally takes its feed from the input/gain stage of the Mutron. A simple SPDT toggle switch can select between the output of the normal gain/input stage of the Mutron, and the added external input.

Control Elements

A big part of what goes into the design of an ECF is the choice of a control element; the components that are made to vary with envelope amplitude. As the table below shows, there are a variety of types of control elements to choose from. These include:

- *field effect transistors* used as variable resistors either directly or optically coupled to LED's,
- *photoresistors* driven by incandescent bulbs or LED's or by individual LED's,
- *operational transconductance amplifiers* such as the CA3080, CA3094, and LM13700,
- *switched filters* driven by clocked analog switches such as the CMOS 4066
- *bipolar transistors* such as the 2N3094 or 2N5088.

Table 1 - Comparison of Control Elements

Element	Cost	Speed	Noise	Current	Matching	Avail	Dist'n	Flexibility	Space
FET	++	++	-	++	+	++	- -	+	++
LED-FET	- -	++	-	++	++	-	- -	+	+
Lamp-Photocell	-	- -	++	- -	- -	+	++	++	- -
LED-Photocell	- -	-	++	-	++	-	++	++	+
OTA	-	++	+	++	++	-	+	+	- -
Switched Filter	++	++	+	++	++	++	++	+	+
++-Very good +- Not bad -- So-so - -- Problematic									

I've provided a very rough rating scheme for each control element along a number of dimensions. These aren't everything, but they include most of the things that a design and product would be sensitive to:

- *cost*: typical street price (may vary, depending on where you live)
- *speed* at which they can respond to rapid signal changes, both in terms of initial attack time, and recovery or decay time (NOTE: fast speed also means susceptibility to ripple distortion; see below)
- *noise* and hiss
- *current* consumption
- *matching*: the ability to easily find two or more units that are matched for characteristics
- *availability*: how easy it is to find them (i.e., will most walk-in retailers carry them?)
- *distortion*: how easily they handle large signals
- *flexibility*: how easy it is to incorporate them into different filter configurations or place them

anywhere in a filter design (e.g., if they can only be used as control elements tied to ground)
- *space*: the overall space occupied by the element and awkwardness of placement of components (e.g., multiple photocells coupled to a single lamp)

I'm sure there will be objection from somebody about almost every single little box in the table, but the basic point of the table is that by selecting a given control element, one inherits a number of constraints that are somewhat different from those associated with another control element. For example, choosing to use a photoFET as a control element, over a lamp/photocell combo, means that one may have to pay a little bit more, be a bit more attentive to signal level throughout the audio portion of the device, and do something to offset the distortion resulting from envelope ripple. The good side is that the board can pretty much be laid out the way one wants, and current consumption will be low, which may allow the designer to do other things with whatever reserve juice is in a 9-volt battery, or simply have the freedom to use it with batteries and not a wall-wart or on-board transformer (and all the cost and hum elimination hassles that incurs). When you consider that some manufacturers are more or less committed to a particular chassis size and style (e.g., BOSS, DOD), control elements may also be selected in order to permit the complete design to fit in the requisite space, the same way that controls and switches are.

The other thing to consider about the control element is that these represent the major distinction from, and advantage over, wah-wah's. In a typical wah-wah, the user sweeps a single control element (usually a pot) with their foot. Having only one controllable element limits one considerably with respect to filter design, which is why so many wah-wahs tend to sound similar in their functioning (although, to be fair, Morley has been using optical control of wahs for years). Many types of filters require simultaneous control of two elements, which can be a bit problematic in some instances, though not exceedingly so. For example, the state-variable filter in the Mutron III requires two elements to be changed simultaneously. In the Mutron, this is done with photocells, however it can also be done with a dual-ganged pot. In fact, if you've ever seen Craig Anderton's *Electronic Projects for Musicians* book, his "Super Tone Control" project is exactly that. Mount it in a footpedal with a dual-ganged pot controllable by the foot mechanism and you have a helluva wah. That digression aside, in many instances, it is possible to construct more complex filters using envelope control of control elements. The *Moogerfooger Lowpass Filter* is a prime example, requiring at least 4 control elements to sweep the frequency of a 4-pole filter (those elements may well be built into a custom chip, but that's another matter). If the designer is willing to use AC power, rather than relying on the current and voltage swing limitations of batteries alone, the possibilities become wide open.

Filters

Envelope follower characteristics amd control elements are obviously important, but filters and filter modes are what you bought/made the thing for, so let's talk about them. There are many types of filter designs, and ways of achieving the same type of filtering, but most filters can be classified into 4 basic types or modes:

- lowpass (LP) - only frequencies <u>below</u> a certain value pass unimpeded
- highpass (HP) - only frequencies <u>above</u> a certain value pass unimpeded
- bandpass (BP) - only frequencies <u>between</u> a certain minimum and maximum value pass unimpeded
- notch (AKA band-reject) - combination response of HP and LP, like the complete opposite of BP; frequencies between a certain minimum and maximum value are filtered out.

There are always exceptions to any rule, but the overwhelming majority of wah-wah pedals usually are bandpass filters, meaning that there are limits to the kind of sound textures you can create with a wah. All the more reason to have both a wah pedal and an ECF. Each has different capabilities.

The selectiveness of a filter plays a big role in the way that it directs the listener's attention. A more selective filter puts the region around the cutoff frequency front and centre, by boosting that part of the spectrum many db above everything else.. Selectiveness will depend on the number of poles/stages the filter has, in addition to the resonance or gain built into the filter. So a single-pole lowpass filter will provide a slope of 6db/octave starting at the designated rolloff frequency. Higher frequencies WILL pass, but at a reduced amplitude. You will clearly here a filter sweep, but you will also hear frequencies outside the cutoff frequency range. A more complex 12db/oct (2-pole), or 4-pole (24db/oct) lowpass filter will also let high frequency content through, but at a drastically-reduced amplitude, such that there will seem to be no frequency content above the rolloff.

With few exceptions, it is rare to see ECF's with more than a 12db/oct slope. The typical simplicity of filters stems from the number of control elements required to construct different filter topologies. It is rare to see more than 2 control elements being driven by the envelope follower, and fairly typical to see only one such element. With one control element, a 6db/oct bandpass filter is the most likely filter type used. Having two control elements is likely to be associated with a state-variable filter (as in the Mutron) or a 2-pole lowpass filter (as in the MXR Envelope Filter).

What is the relationship between selectiveness and tone? In some respects, one would think that a steeper rolloff would be preferable, but that isn't necessarily the case. For instance, many ECF's use a bandpass design. Since BP filters cut out low end as well as high end, having a shallow slope can preserve some of the body or meat of the notes as the filter sweeps. A more selective filter would tend to wipe out any bass once the filter has swept up even midway over its range. Conversely, though, a notch filter could probably stand a bit more selectivity, since a broad notch can "gut" the fundamental as the filter sweeps. The problem that accompanies this move, though, is that steeper notches are often created by cranking up the resonance, which will often result in increased just above and below the notch frequency, which may not be what you want (although it may well BE what you want).

Lowpass filters have the advantage of retaining the fundamental of all notes played as the filter sweeps, so that the sound is never thin. They also have the advantage that if the resonance is increased enough, they shift attention to the cutoff frequency in a way that mimics bandpass filtering. So, you get the pleasing emphasis of a BP filter with the meat of a LP. This is one reason why the Moogerfooger Lowpass can appear to be so flexible, despite having only one filter type. Also one reason why bass players prefer ECF's that can provide a LP mode.

HP filters can be very dramatic sounding, particularly when in downward drive mode. As they sweep, they go from being little more than upper harmonics and signal grit to consisting of more and more of the lower harmonics, and finally fundamental. With a rapid attack, it sounds like the notes are being sucked into existence. Think of LP filters as adding (or subtracting, depending on direction) brightness as they sweep, and HP filters as adding (or subtracting) body as they sweep. The most typical use of HP filter mode is likely with a downward sweep direction, whereas LP and BP filter modes can sound pretty good sweeping either direction. One of the reasons is that for this is the relative duration of the attack (short) vs decay (long) portions of the envelope. When sweeping down, a HP filter zips to the range of the fundamental, hangs around there, then slowly sweeps back upward (leaving only middle and upper harmonics) as the note fades out. You get to hang onto the meat of the note during the best parts of the envelope. When sweeping upward, you miss out on the fundamental if the unit is set for even moderate sweep. Meanwhile, the filter hangs around for a while at the high end, with no bottom at all, while all the good stuff is happening lower in the spectrum. Then, it slowly sweeps back down just in time to miss all the fun. You can see why it's possible to get more pleasure out of a downward-sweeping HP filter than an upward-sweeping one. Note that this is completely separate for what they may do for you as a *sound effect*, used to produce individual sounds, as opposed to a stompbox integrated into your rhythm or lead playing.

Whatever the type of filter, resonance or "peak" controls will adjust the emphasis or apparent selectivity of the filter. Not all ECF's will include resonance controls, although many have the capability of adding them. Depending on the filter design, increasing the resonance may increase distortion, and may also have an impact on overall output volume, since resonance usually involves changing the gain within the filter. On higher end ECF's, such as modular VCF's linked to envelope followers, higher resonance can verge on self-oscillation at the cutoff frequency. If you like Fatboy Slim or the Chemical Brothers, you've heard this sound. The type of filters that can do this however, are unlikely to run off a 9V battery, so they are unlikely to appear in stompbox form, though some higher end stompbox ECF's (Lovetone Meatball, Frostwave Funk-A-Duck) offer this feature.

One intriguing type of filter is that used by the old MXR Envelope Filter; one of my favourite ECF's of all time. This uses a 2-pole lowpass filter whose cutoff frequency is set by varying the duty cycle (on-vs-off time) of a pair of CMOS switches. A high frequency clock (outside the hearing range, I gather) switches a pair of resistors in and out of circuit for two simple lowpass sections. The envelope follower provides a signal which determines the duty-cycle of the clock signal, and the switches. Since the cutoff frequency is set by the proportion of time that the resistors are "on" (i.e., connected by the CMOS switches), varying the duty cycle varies the average resistance over time, akin to varying the actual resistor value. For whatever reasons, the MXR unit produces a delightfully responsive envelope and full-bodied tone, with very little noticeable distortion or envelope ripple. (Steve Giles has been applying his best mad scientist skills to this one, and has been experimenting with a number of modifications that can make this lowly 2-knobber a real competitor with the mighty Mutron.) I guess the point of this digression is that how a filter works can sometimes obligate a whole other set of design features that alter its playability.

<u>The Role of Signal Level and Dynamics</u>:

External Loops and Filter Dynamics The degree to which the filter sweeps certainly depends on how you set it, and how you play, but it is also a partial function of the dynamics of the input signal. The more dynamics there are to the signal, the broader the sweep, for any given setting of the sensitivity or drive control. Consequently, many players prefer to stick ECF's close to the start of their signal chain, since FX like compressors and fuzzes deliberately remove dynamics from the signal. On the other hand, many players will also tell you that ECF's simply sound richer when there is more harmonic content to filter, so there are sonic advantages to sticking an ECF *after* a fuzz Electro-Harmonix (and especially Mike Biegel) were smart enough to recognize this, and updated their Q-Tron pedal with an effects loop, as have other manufacturers like Lovetone. This allows the user to plug the guitar directly into the pedal and tap the initial guitar signal as the source of the envelope (with full dynamics), but stick fuzzes, flangers, or what not, before the filter section so that more interesting timbres can be played with without sacrificing dynamics.

Wait...I hear wheels rolling and gears turning....yup I hear them louder now. Not only does adding an external envelope input to an ECF allow you to modulate it with other sources. But, if you were to plug your guitar into an active splitter, and run one of the splitter outputs to the external input of your ECF, then you could stick any number of doohickeys between the guitar and input to the filter without sacrificing dynamics. Congratulations. You just figured out how to make an FX loop.

Level: A wah-wah will wah, and an EQ will equalize, regardless of signal level, as long as you twiddle the controls. The signal to noise ratio and the distortion characteristics may not be the greatest if the level is too low or too high, but it will work as intended. Not so with an ECF. The envelope follower is designed to assume a particular signal level and detect variations within that range. If the signal you feed it is hotter than it expects, you may find that only the lowest 10% of the sensitivity pot's rotation is useful, and anything over that will send the filter right to the limits of its sweep and keep it there a long time. If the signal is much weaker than it expects, you may find that the filter fails to react unless the sensitivity is up full tilt. Both types of situations are likely to be interpreted by users as a busted or nonfunctional or simply undesirable effect. Meanwhile the effect is just waiting for its preferred type of signal to show what it can do.

There are several fixes to the problem of poorly matched signal level:

- Pick harder (or softer), although that may conflict with your playing style or what the song needs.
- Make sure your guitar volume is up full (or down), although this may conflict with your plans for how you want to use your guitar volume to drive other things, or how you planned out volume levels over the course of a song/piece.
- Use hotter pickups or an on-board pre-amp, or ditch the pre-amp and hot pickups for something softer. Not without its drawbacks.
- Use some kind of outboard/stompbox pre-amp with an adjustable output level. If you can't find a suitable one, most stompbox compressors with adjustable output level are able to sub for this (see below), as can EQ boxes.
- Stick a distortion/overdrive device ahead of the ECF (although this has some other effects; see below).
- Find out what components set the gain of the envelope follower and tweak them for more gain.

Keeping the optimal signal level in mind, it is unwise to stick any device ahead of an ECF that is likely to alter the volume in a way that alters the envelope of each note or strum. For the ECF to sweep in sync with your picking, the initial attack of the note must be the strongest part of the signal. Any alteration of that, and the ECF will start to behave in unusual ways. This may be what you want, but then again it might not. Examples of this would include any effect that sweeps back and forth automatically such as a stompbox tremolo, a phase shifter, a flanger, and a univibe. Although the tremolo is the only one that explicitly alters the volume, the notches introduced into the signal by these other devices may well result in a volume drop for some notes at some points, and cause the ECF to behave erratically. Again, your aesthetic choice.

Dynamics: ECF's are dynamic devices. They depend not only on signal **level** but on **variations** in signal to work. If the signal doesn't vary, then the effect won't sweep. So, while you don't want the overall signal level to be too high or too low, you generally want it to vary a lot.

Effects that can reduce dynamics include limiters and compressors set to clamp the volume at a very steady level, and fuzzes that clip very heavily, or simply result in a more compressed sound. If the signal level fed to the ECF is not that strong to begin with, and the dynamics are compressed in some manner, then you can crank up the sensitivity all you want, and pick as hard as you want and may not hear any effect whatsoever. Conversely, if the sensitivity is turned up, and the signal lacks dynamics, you won't hear much sweep because it can hang-up at the extreme point of the sweep.

On the other hand, you can use this to your advantage if you want the sweep to be subtle, or if you want to be able pick very hard without getting pronounced sweeps. One of my favourite sounds involves use of a compressor ahead of an ECF. Especially if it's a bad compressor. A bad one? Yes. Compressors and limiters have their own envelope followers, that are used to adjust the gain or level of the signal instantaneously. Bad ones will often exhibit a phenomenon called "breathing", that involves a clumsy recovery from the initial attack, sometimes resulting in a slight gain over the life of a note. Although it makes vocals sound unnatural, it can add an otherwordly sound to things sometimes (I especially like it on the rear pickup of a Tele, but that's another thing). In our case, it can add a gentle openning-up of a note or chord after you've picked/strummed it, by getting the ECF to sweep very slowly and very subtly. In some respects, what would be nice would be an integrated dynamics processor and ECF, such that you could have constrained dynamics for the lower 2/3 of input signal levels, and expanded dynamics for the upper 1/3. This way you could play fluidly with little filter sweep, and easily "push the envelope" (sorry, couldn't resist) by picking just a bit harder to take you over the threshold when you need it. The bottom line here is that getting an ECF to do interesting things, and especially do them the way you want, depends on keeping one eye on the overall dynamics of your signal chain.

Dynamics is generally an issue insomuch as it affects sweep width. Sweep width, however, depends on the translation of envelope follower signal into changes in the control element. It's almost always possible to take the puny contrast between peak and valley of a highly compressed signal and electronically translate it into a more substantial contrast. You can do that by either increasing the gain of the envelope follower, or by increasing the gain of the input signal. This will have no impact on the attack and decay characteristics of the compressed input signal itself.

Common Controls:
These are the most commonly observed controls on commercial ECF's. Those controls labelled with an asterisk are part of the envelope detection circuitry.

- *Sensitivity, Gain, or Depth**: Adjusts input signal going to the envelope detection circuitry, and determines how much the filter responds to your playing. In some cases (e.g., Mutron), it also determines the signal level going to the filter itself. Less sensitivity means less sweep for the same strum. Important for matching responsiveness of effect to differing input levels (e.g., post fuzz).
- *Q, Resonance, or Peak*: Adjusts filter emphasis or resonance at the turnover frequency.
- *Initial or Start Frequency**: Point at which the filter starts its sweep.
- *Drive, or Direction**: Selects whether the filter starts its sweep going down from a starting point or going up.
- *Range*: Usually selects between two or more sets of capacitors that determine the overall range of sweep
- *Filter Type/Mode*: Selects between basic types of filters. Usual configurations are bandpass/lowpass, with highpass and notch thrown in sometimes. In some cases (e.g., Dr. Q) the switch may select between two versions of the same basic filter category.

Less common controls might include:

- *Attack**: The time taken for the filter to sweep from its resting frequency to the point of maximum sweep (wherever that happens to be).
- *Decay**: The time taken for the filter to settle back to its resting level.
- *Sweep**: Pans between inverted and non-inverted versions of the envelope signal. Since the centre position usually mixes inverted and non-inverted envelope signal and results in cancellation, this type of control effectively becomes a combined sweep width and direction control.
- *External CV/triggering*: Allows something other than what the user is playing to drive and control the filter circuitry. This might be processed as if it were an instrument (i.e., it goes through the envelope follower circuitry) or or treated as if it were the output of the envelope follower circuitry (i.e., it directly drives the filter)

An Example: Circuit analysis of the Mutron III

There were many others in the 70's, but the Musitronics Mutron III established many of the more common features to be found on many commercial, and probably homebrew, ECF's of the time, and since then. Let's take it apart.

Input Stage The first diagram shows the input stage. The op-amp labelled A1 is the inverting input buffer and gain stage. The 120k resistor in the feedback loop sets the gain of this stage, in conjunction with the 1meg pot labelled "Gain". The gain can be set from an attenuation of about 90% (i.e., 10% of the input level, for very hot signals) to a gain of about 40. Pretty versatile, with respect to the range of input signals it can handle, although bear in mind that the gain is applied to both the signal going to the filter and the signal going to the envelope follower.

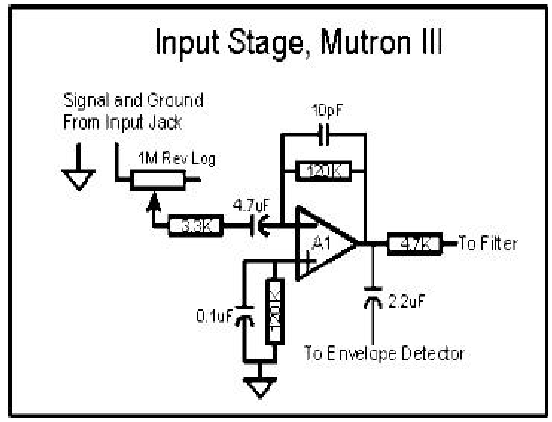

Any attempt to increase the input signal to improve the S/N ratio automatically results in a stronger envelope signal, and any attempt to keep signal clipping at bay in the filter will reduce the envelope signal. There is no way to simply reduce the responsiveness of the envelope follower section without changing the signal level. Fortunately, since the Mutron uses photocells, which are relatively immune to level-related distortion, this isn't a big problem in terms of maintaining a clean signal.

Filter Stage: The filter section has what is referred to as a state-variable filter, made up of op-amps A2-A4. This is a very standard filter that shows up in many applications beyond the Mutron. Craig Anderton used the state-variable configuration for his Super Tone Control project in the two *Electronic Projects for Musicians* books. The state variable filter can produce any of 4 different types of functions that can be tapped at different outputs:

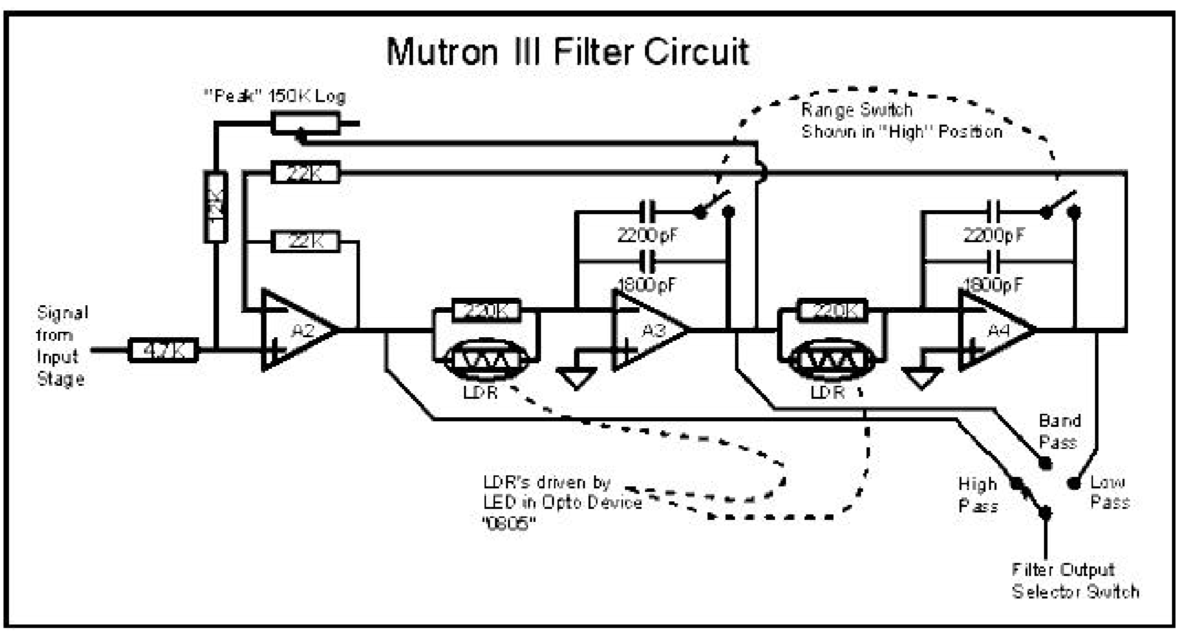

You can select which of the first three types of outputs you want via the "mode switch". If HP and LP are selected at the same time, the output is a notch-type filter. This is done on the upgraded version of the Mutron that Electro-Harmonix markets as the Q-Tron and Q-Tron+. If you own, or have made, an ECF that has a similar kind of mode switch with LP, BP, and HP settings, chances are that it is a state variable filter, and you can add the notch function by providing some sort of switching arrangement that lets you select both HP and LP at once. In contrast to the wah-like sound produced by the other filter modes, notch mode produces an effect much like an envelope-controlled phaser; largely because what a phaser does is introduce notches.

The frequency at which all of this happens is set by the value of the capacitors in the feedback loop of A3 and A4, the resistor between A2 and A3, and the resistor between A3 and A4. The capacitors need to be matched, and so do the resistors. You can tune the corner frequency of the filter by varying either the resistors OR the capacitors simultaneously. The Mutron sets a maximum resistance via the two 220k resistors linking the op-amps. In parallel with each of these two resistors are LDR's (photocells), whose resistance varies depending on the light that falls on them. They are contained in the mystery component labelled 0805 in the drawing. As the light shining on the photocells varies, so will the combined resistance of the fixed resistors and photocells, and the corner frequency of the filter's output will change for all of the various types of outputs it provides.

The overall *range* of filter frequency sweep can be changed by simply changing the value of the capacitors in the feedback loop of two of the op-amps. The Mutron does this by adding in a second pair of caps in parallel to lower the range set by 1800pf caps. Since the Mutron has no real way of setting the "start" or initial frequency of filter sweep, the range switch is pretty handy. What some folks do is use a switch with more than 2 positions, such as a 2P6T rotary switch. In a rather cumbersome way, this simulates being able to manually tune where the filter starts sweeping from. Bear in mind that changing the range simply revoices the filter higher or lower, but does not change the *amount* of sweep in any way. You can refer back to an earlier section of this paper for some ideas on how to alter the range of sweep. Some ECF's, such as the PAiA *Motion Filter*, which use a single control element to vary frequency, allow one to tune start frequency via a pot. The standard approach used in analog synth filter modules is to sum a DC voltage with the envelope voltage. A similar kind of arrangement can be done with many stompboxes. Most companies have neglected this feature, however.

The envelope follower: Below, you can see the envelope follower and envelope invertor section. The Mutron uses what is referred to as a precision half-wave rectifier. That is, it provides an AC voltage that is unipolar - varying between ground (0 volts) and some positive value. The circuit fragment around A6 serves as the driver for the LED that shines on the photocells, and also serves as a mixer stage for the two sources of control voltage. Two? Yes. To enable downward drive, the Mutron taps the positive end of the power, and *subtracts* the envelope voltage from that positive DC voltage. With the drive switch in the Up position, the positive DC voltage is out of circuit, and just the envelope follower output is used. The 330 ohm resistor and 4.7uf cap set the attack and decay time constants. Off to the far right, you'll see a resistor of undeclared value. Why the "mystery component"? Simple. Photocells, by their nature, are not predictable animals. You may be able to match a pair....maybe...but knowing in advance how much current will be required to drive an LED to shine on them to produce a given resistance is not an exact science. The "mystery" resistor is selected to provide the right amount of optical drive from the LED. It IS an LED, remember, so you should expect its value to be somewhere on the order of a few kilohms at most.

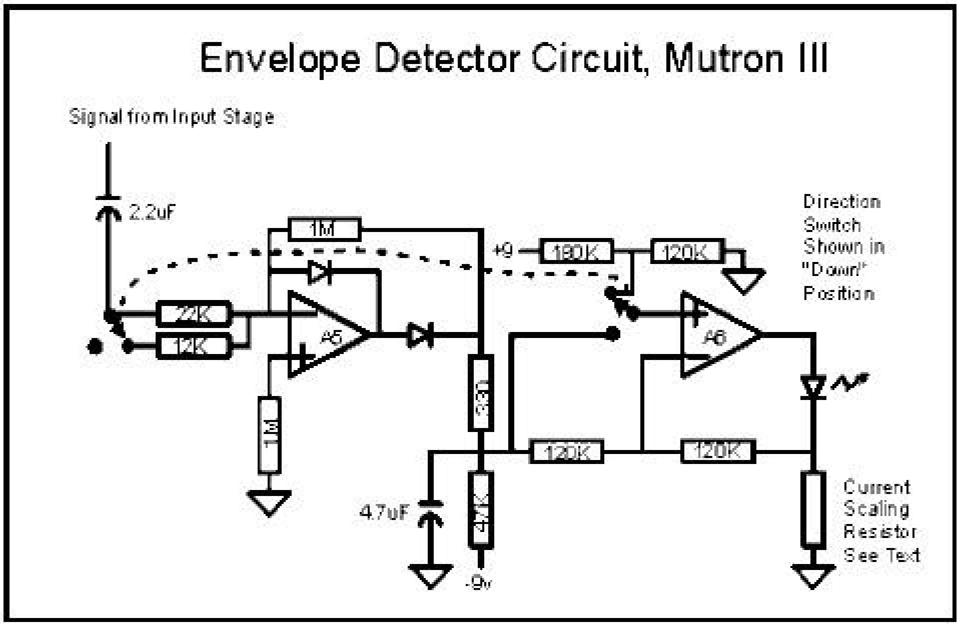

Useful Mods

Is there such a thing as the perfect version of any effect? Nah! There are always mods you can make to any effect to squeeze more tones and options out of it. What kinds of things can be done to the garden variety ECF to provide other sounds? We've touched on some already, but here is an extended list of tweaks to consider, all of which are feasible on many (though not all) ECF's:

- Change attack and decay time constants, improve envelope ripple; benefits already discussed.
- Change sensitivity/gain of envelope follower: benefits already discussed.
- Change envelope signal source; provide option for FX loop: benefits already discussed.
- Provide envelope output for syncing multiple devices: allows for player control of true stereo effects using different envelope-controlled devices. Some higher end ECF's offer this capability. Sometimes you do want a couple of things synced to the same envelope signal. Since you can't do that when you run them in series (remember, the filter will change the envelope of what comes out the output jack), it's nice to have the option of sending the envelope follower output to a jack where it can be sent to a couple of pedals.
- Provide more filter modes. This doesn't work for everything, but state-variable filters can offer 4 filter modes. If the stock unit only offers 2 or 3 (usually LP+BP, or LP/BP/HP), you can easily tap the remaining possibilities.
- Change frequency range of filter. The range of any given ECF doesn't always match the instrument you intend to use it with. Fortunately, most filter structures used in commercial effects can have their range of sweep altered by changing the values of the two range-setting caps (you might even improve tone by being more accurate in selecting matched-value caps than the manufacturer).
- Provide initial frequency tuning. Changing range is not the same as where you start out in that range. Handy, and often easy, to add a control which adjusts where you start your sweep. You can fake it to some extent with a multi-position range switch, but a high/low range switch and initial frequency control provide more accuracy.
- Provide sweep width/depth or maximum sweep control. Generally pretty easy to do, and an important complement to having adjustable envelope characteristics.
- Mix filter and straight signal. Many players, *especially* bass players, and *especially* folks using a bandpass-design ECF, find the sound is too thin once the sweep has gone up a ways. By adding a mixer stage after the filter, and being able to adjust the balance of straight (unfiltered) sound and

filter, you can keep the body in your tone, and still get a noticeable sweep sound. Craig Anderton's "Retro-Stereo" module (available on several schematic archives) is a clever little doo-hickey that provides sum and difference signals using the straight and filtered sound, to give two qualitatively different sounding outputs (hence the "stereo"), but synced with respect to what frequency they are at in their sweep. Simple. Clever. Love it.

- Mixer for normal and inverted envelope signals. Some ECF's have a drive direction control. Others use a normal and inverted envelope signal, and pan between them. The EH Y-Triggered Filter works like this. Set to the mid-point, the unit doesn't sweep, because the normal and inverted version, in equal measure, cancel each out. Turned in either direction, the control increases the amount of sweep of that type. You get direction and sensitivity in one knob. Clever.

If your ECF has no provision for setting start and stop frequencies, there are ways of accomplishing it via simple mods. For example, in the Mutron, (*see circuit analysis below*) the filter frequency is set by two range-setting capacitors, in tandem with two fixed resistors that are paralleled with two photocells. The combined resistance of the fixed resistors and photocells determines the specific frequency. If you wanted to decrease the effect of the photocells (i.e., so that less sweep would result), you could either reduce the value of the fixed resistors (say from 220k to 150k), or you could stick fixed resistors in series with the photocells so that their paralleled resistance never drops below a given value. Optically-controlled ECF's, or those that may use FET's as simple voltage-controlled resistors, are certainly the easiest to modify in this manner. Those using other categories of control elements (e.g., units thatuse OTA's) may be somewhat more complicated to work out the particulars for. Bear in mind that such resistance changes will also alter the overall frequency range, so you may want to offset this by changing the values of the two caps in the appropriate direction (upward or downward). Experiment with additional small values in parallel to find what you like.

In many instances, it may be difficult or at least awkward to figure out how to constrain the sweep of a given unit. The simplest fix is usually to identify two matched capacitors that are used to set the range of the filter, and sub other values to shift the range over a bit. For example, the Dr.Q/Quack uses two 4700pf (.005) caps to set the filter range. If you want to shift the sweep range downward, tack on another 1000-2200pf or so The principle is pretty generic. I've done this with the MXR Envelope Filter, the DQ, Anderton's Bi-Filter Follower (which automatically implies implementation on the EH Baseballs, and Seamoon units), and it can be done (and IS done) on every Mutron clone. An interesting side-note. On units which use a pair of BP filters (Baseballs, Bi-Filter Follower) changing the range of one filter section will spread out the two centre frequencies on the two filter sections. Spread them out far enough, and the unit starts to sound like an envelope controlled phaser. Nice.

<u>Interesting Things You Can Do With an ECF</u>:

There is a whole lot more to the universe of ECF's than playing "Disco Duck" or "What I Am is What I Am". For starters, much like a foot-operated wah-wah, ECF's have a different impact when placed before a distortion device, than when placed after one. I like to think of fuzzes as a kind of additive synthesis, where what you feed in determines the type and proportion of harmonics added at the output. Placed ahead of a fuzz, an ECF further alters the kind of harmonics that come out by emphasizing certain frequencies, and clipping them more than others. Since the harmonics generated extend beyond the passband of the ECF, you'll hear this as a kind of animation to an otherwise broadband sound, rather than any sort of obvious wah. Placed after a fuzz, the same device now eliminates harmonics outside of the passband, providing a more obvious "wah" or "ow" (depending on how you set drive/sweep direction).

One of my favourite uses of an ECF is to stick it in front of a fuzz set for moderate clipping, but stick a compressor ahead of the ECF. Normally, this is the kind of thing you don't want to do, since removing dynamic responsiveness makes the ECF sound static. What we'll do, though, is set the compressor for only moderate compression, and adjust the envelope follower sensitivity for moderate responsiveness. Given that the compressor is likely to have its own envelope "issues", it will likely take a bit longer for the ECF envelope detection to settle down. The interaction between the two envelope detection circuits will result in the ECF waivering around the sweep start point, without really sweeping in a broad obvious way. Moreover, it won't sound terribly systematic or note specific, but it won't be quite random either. If you have the initial frequency set for a kind of vocal range, the effect will be what Frank Zappa used to call "pillowy". The only example I can find on record is some of the early guitar work of Todd Rundgren with Utopia. That may not be how he did it, but that's what it sounds like. I've never tried it, but I'd be curious to know what an ECF sounds like placed in front of an octave fuzz. Conceivably, the octave effect would seem to come and go.

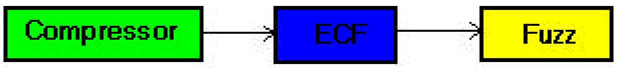

Another treat I stumbled onto made for a passable attempt at simulating tape reverse effects. The central feature of an effective tape reverse effect is to make all those things that happen at the start of a sound occur at the end and all those things that happen at the end occur at the start. Although you can do part of that with your pinky and guitar volume, or with a volume pedal, what you can't do with your hand is mimic the timbral change of backwards sounds.

Here's what I did. I sent the guitar to a compressor. Again, used to constrain dynamics, and provide another source of envelope follower recovery time. From the compressor, it went to a 6-band EQ, set for bass cut and mid-boost, and from there to an ECF; in this case an MXR Envelope Filter set for a slow attack time, and moderate sensitivity. The MXR-EF sweeps up, and has a fairly meaty, rather than thin, sound. With attack set for slow, it would take a few tenths of a second for the filter to reach the midrange and treble.

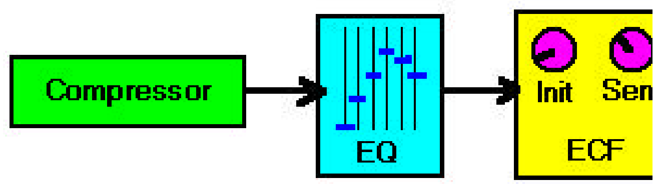

With the EQ set the way it is, there wouldn't be enough signal to really push the envelope follower hard, so you shouldn't expect a wide sweep. The compressor works by boosting the signal, and using an envelope follower to drive a gain-reduction control element. If the control element in the compressor has a slow recovery time, then what will happen is that the signal will appear to retreat and gain output slightly shortly after you pluck or strum hard.

ECF hall of fame:

While there are lots of one and two-knob wonders, there are also some absolutely stunning and functionally complete units out there. Some historical brightlights would include:

- The ***Musitronics Mutron III***, and its descendants, the Electro-Harmonix Q-Tron and Q-Tron+. There have been a number of workalikes of the Mutron over the years, which I suppose attests to its attractive set of features. The Ibanez Auto-Wah and the Univox Funky Filter both have the exact same set of features as the original Mutron. I used to own a Funky Filter, and it was a Mutron every inch of the way. In recent years, Mike Biegel updated the Mutron design for Electro-Harmonix, eventually incorporating a send/return loop so that dynamics could be preserved by extracting the envelope signal prior to any additional effects processing. Very handy. Those familiar with Craig Anderton's book *Electronic Projects for Musicians*, will know that Craig includes instructions for providing envelope-control of his Super Tone Control. In fact, the resulting unit is essentially a Mutron-type state-variable filter.

- The ***Beigel Sound Labs Envelope Controlled Filter***. This is kind of the holy grail of ECF's, with very few actually built and sold. This was a rack mount unit that had everything the Mutron had (no small coincidence, given the designer), plus much better fine control over the envelope, and better interfacing capabilities.. If you find one in a yard sale, call me...collect.
- The ***Electro-Harmonix Microsynth*** for Guitar or Bass. The EH-MS provides for an automatic swept filter sound of the straight guitar and several other derived tones (fuzz, octave). The begin and end points of the filter sweep can be used to set both the direction of sweep and the width of sweep. Although the guitar's natural envelope is not used to drive the filter, it is used to drive the VCA and provide some degree of dynamics. Not a dedicated ECF, strictly speaking, but the absence of a number of other features made the swept filter the most prominent feature of the MS, and one of the major reasons for using it. The Korg PME series of modular effects had an ECF unit with an almost identical set of controls. The more recent BOSS Bass Synth carries on in the same vein. Clearly EH struck a chord.

The resurgence of interest in ECF's has also produced some very classy and thoughtful products:

- ***Moogerfooger Lowpass Filter***: This is a stompbox version of the very filter that got people hungry for a guitar version of what a synth could do in the first place. Bob Moog has put a 4-pole lowpass filter in a classy footbox with top notch external control capabilities and Moog tone.
- ***Lovetone Meatball***: A twiddler's wet-dream. This British baby has more knobs than you can shake a stick at, including excellent control over envelope characteristics, and an external loop (like the Q-Tron+) for using the pre-loop dynamics to control an additionally processed signal.
- ***Frostwave Funk-A-Duck*** and ***Chunk Systems Agent 00Funk***: A pair of Australian gems. Not quite as tweakable as the Lovetone but both have solid thick classic tones, with good envelope control, and a high resonance filter for getting those Korg MS-20 sounds that have populated dance music as of late. The Duck has pretty much all the controls of the Mutron, plus a little more. The 00Funk has control voltage input and output for driving, and being driven by, other devices. Both worth looking into.
- ***Z-Vex Seek-Wah***: An intriguing combination of sequencer and triggered wah, this baby doesn't neatly fit into any category. In a loose sense, it corresponds to, and is sort of synched to one's playing, but doesn't correspond directly to the signal envelope. Rather, the envelope signal that drives the filter section can be shaped by 8 trimpots. More synth than stompbox in its pedigree, but that suits me fine.