ELECTRONOTES 224

Newsletter of the Musical Engineering Group

1016 Hanshaw Road, Ithaca, New York 14850

Volume 23, Number 224 February 2015

VARIABLE-SLOPE FILTERING REVISITED

-by Bernie Hutchins

INTRODUCTION

One useful function of ***Electronotes***, our main publication, recently, (since 2010 and most recently in 2013 and 2014) has been in a "review" mode. Some of this reviewing is mostly a summary or tutorial effort. Yet at times (six of them) I have used the term "Revisited" in titles to suggest that the topic under discussion was to be one for which we decided to pick up loose ends; or to some degree expand. In thinking of topics which I considered true revisits, I thought in particular of Through-Zero FM in EN#206 [1], Pulse-Width Modulation (as an ongoing periodic process – not just envelope driven) in EN#216 [2] and Self-FM in EN#217 [3]. These were topics I considered to have been left hanging for various reasons. All three were "analog" topics that we were lately able to explore further with digital methods of analysis (programming). This is not to say we necessarily could propose a practical digital means of synthesis.

To this revisit list (officially denoted as revisited, or just de facto) we here add the notions surrounding Variable-Slope (V-S) Filtering. A few weeks back John East asked about this – wanting to know if any definite conclusions about its implementation, and its usefulness had ever been made. I could not come up with an immediate answer to either question. So this did appear to be a loose-end topic worthy of a revisit.

Thinking that there were perhaps three ***Electronotes*** articles on the V-S topic (I had in mind articles by Lester Ludwig and Normand Provencher, I did a search through the back issues and was surprised to find many more articles, including the fact that we actually had a V-S option for a VCF in the ENS-76 series of designs in EN#72 [4], and this was a good lead going backward to what seemed to be the first mention in EN#59 in Dec. of 1974 [5].

In the EN#59 article I mention the first suggestion of a V-S filter as being from Lester Ludwig in what was probably summer of 1974. I do remember his asking. He was never embarrassed to suggest unconventional approaches at least in a "why can't we" cloak. The notion of somehow averaging integer multiples of 6db/octave for desired slopes was something we considered at that time. Also mentioned in EN#59 was a second impetus of the idea based on an Aud. Eng. Soc. paper by David Luce (then at Moog) which was presented at the spring 1975 AES convention and later published [6]. I do not have this paper at hand, but my description reminds me that David was remarking on how the spectra of "real" (acoustic) instruments had roughly a fixed frequency start to a low-pass and a cutoff slope that became shallower as amplitude got larger. That is, a louder tone had more high frequencies. This was not drastically different from the notion that as a tone got louder a cutoff frequency increased. With the "standard patch" of a synthesizer it was common to have similar control envelopes (very often, the same envelope) for amplitude (to the VCA) and for the cutoff of the VCF.

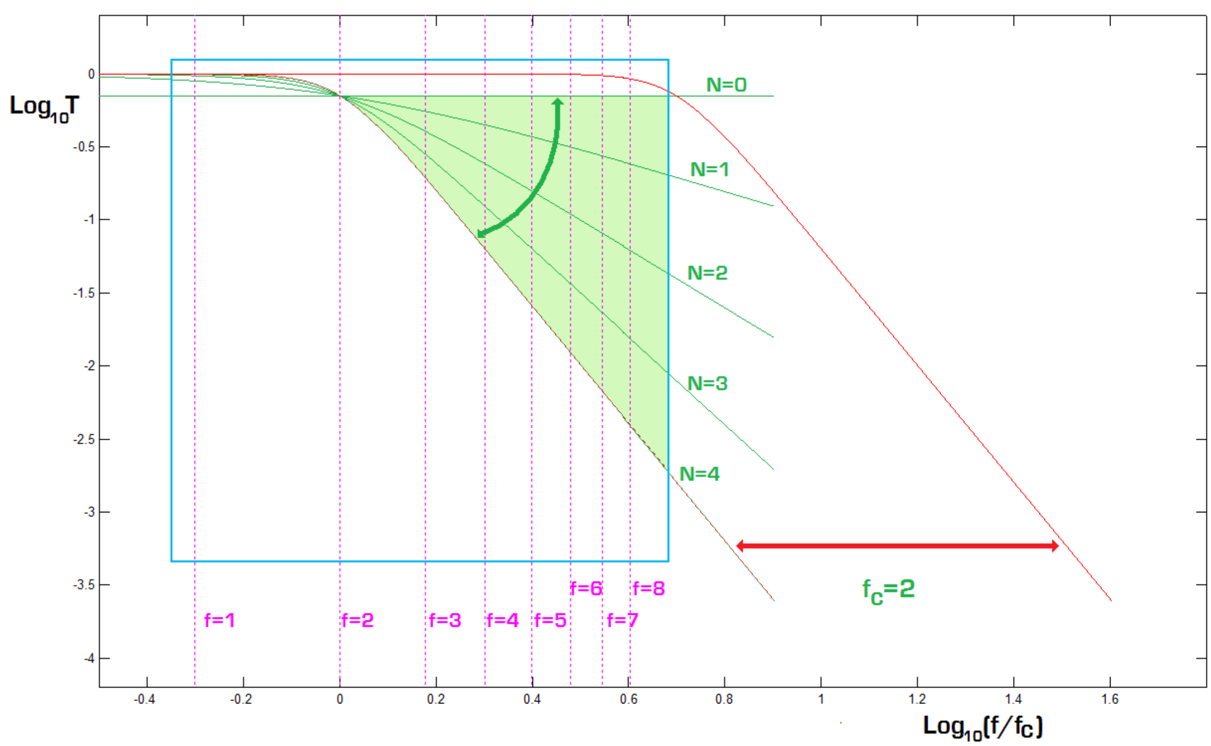

At this point it will be most useful to show a figure of what a VS filter does and how it differs from our more usual filters, and this is shown in Fig. 1 which is a fairly complicated diagram. Note that the frequency response curves here are done "log-log" so that the asymptotic roll-offs are straight lines. The responses here are all Butterworth (BW) for convenience (great convenience) as we shall see. Thus we have a typical lovely collection of BW responses – the green curves for orders N=0, N=1, N=2, N=3, and N=4. The N=4 curve is actually overplotted with dashed red. All five responses have a BW cutoff of f_c=2, where they all go through a magnitude of $1/\sqrt{2}$ = 0.7071 (or -0.1505) on the log scale. We have chosen f_c=2 so that when we consider a harmonic spectrum as the input to these filters, we can have two harmonics (f=1 and f=2) in the passband, so that there is always at least one harmonic to reinforce the pitch of f=1, while harmonics 3-8 are in the roll-off region.

In addition to the five green BW responses with cutoff 2, we also plot the red curve to the right that is a 4th-Order BW with cutoff 10. Note that this was chosen to pass all 8 harmonics. Our eight-harmonic waveform is thought of as being obtained just by summation – it is not a commonly found waveshape.

With all this setup in mind, we can suppose that we want to initiate a musical tone with the filter in the red/green N=4 position which has significant rejection above the 2nd harmonic. We then suppose that as the tone progresses, we want to let more harmonics through the filter. [At the same time, we likely have a larger amplitude in mind. Perhaps the amplitude even jumps up suddenly and then immediately begins to decay as would be the case with so many percussive instruments.] That is, we want to let the spectral energy in the green-shaded triangular region through. There are two ways to do this as shown.

Looking mainly inside the blue box, if we leave the filter characteristic alone, we can move the cutoff from f_c=2 to f_c=10 as suggested by the heavy red arrow. This works, and it very much the way most VCF's work in a "standard patch". [Well, it would not be unusual if the filter characteristic also had an enhanced corner typical of a Moog four pole low-pass.] Note that the slope does not change.

In contrast, the second way would be to have the slope "flap" upward (heavy green arrow), "hinged" at the cutoff, in which case the four BW cases above show steps crossed in the transition. Fully up (N=0), the response becomes flat. We will consider this flat like the cutoff 10 BW – the difference in gain can be adjusted if this is an issue. The point is that the triangle is let through.

As the tone ends, typically the low-pass would go back down, or the flap would go back down. This return to the low amplitudes with <u>higher harmonics being reduced faster</u> is typical of acoustic sounds. In addition to all the simplifications here, we probably should note that since we are changing the filter with time, it is certainly NOT time-invariant, and a quasi-stationary notion of frequency response is necessary (as it always has been in our VCF work).

So that's the basic idea.

As we review, we find some major theoretical offerings by Normand Provencher [7], Lester Ludwig [8], and Lee Powell [9] with additional brief notes by Normand [10,11], and these are valuable, although not easy reading. What we don't see, even in connection with the working circuit [4] is commentary on how well it works (sounds!). Not unlikely there was an element of expectation that as more and more people tried this, we would have useful evaluations. But, perhaps this was too extensive a construction to expect many people to have tried it. In fact, we don't know if anyone else tried it! As for just reporting on what we heard initially, that is not here either. For certain a part of this absence is the difficulty of describing, in words, what something sounds like. This was a common problem with many things we tried.

In thinking about this today, I have wondered a couple of things. (1) If we compare V-S to the standard VCF tone generation, is there a distinct difference? (2) If so, can we make the case that the V-S result is at least of comparable virtue, if not better? (3) How does the V-S compare, over all, to all the methods in our bag-of-tricks?

What has happened in the nearly 40 years, since the V-S was first suggested by Lester, is that we have some vastly more efficient tools (digital generation) than we had originally. It turns out, as we see below, that we can calculate results that we might have needed to examine with analog hardware, and we can even calculate results that are not available with analog hardware (like the continuous BW we use below). This means that we don't even need to simulate a V-S filter digitally – we just calculate the signal.

DIRECTLY CALCULATING THE SIGNAL

We did not say exactly how we calculated the frequency response curves in Fig. 1, but all we did was use the standard formula for the BW low-pass response:

$$|T(f)| = \frac{1}{\sqrt{1+(\frac{f}{f_c})^{2N}}} \qquad (1)$$

where f_c is the cutoff frequency and N is the order. Equation (1) will prove to be the key to the calculations that follow. In Fig. 1 we plotted a log-log representation to make the point about the straight line slopes that result in such a case. Fig. 2 and Fig. 3 will also be calculated using equation (1) to make two points. First, Fig. 2 is a corresponding linear-linear plot of the same BW responses, which is at least as traditionally published as Fig. 1. Fig. 3 makes the (perhaps astounding) point that, for the equation, N need not be an integer. Thus we can achieve a closed-form expression for a variable slope by plugging in non-integer values of N. The asymptotic slope would be 6N db/octave. (This does not mean that we could design an actual circuit and breadboard it on our benches.) So we will provisionally set aside the issue of whether or not this is a "real" BW filter and take equation (1) as a convenient formula for computing the magnitudes of harmonics to be used in our additive synthesis approach to studying the V-S filtering approach.

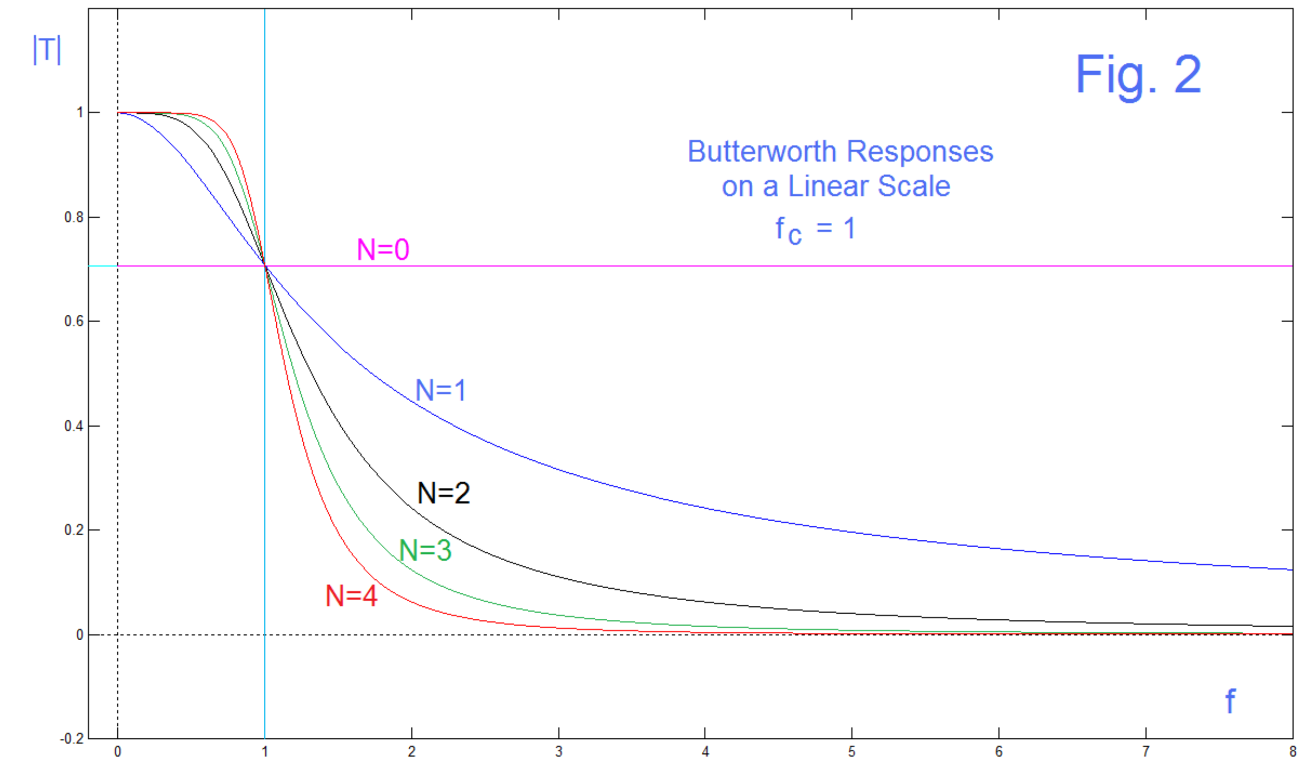
|T|
Fig. 2
Butterworth Responses
on a Linear Scale
f_c = 1
N=0
N=1
N=2
N=3
N=4
f
1
0.8
0.6
0.4
0.2
0
-0.2
0
1
2
3
4
5
6
7
8

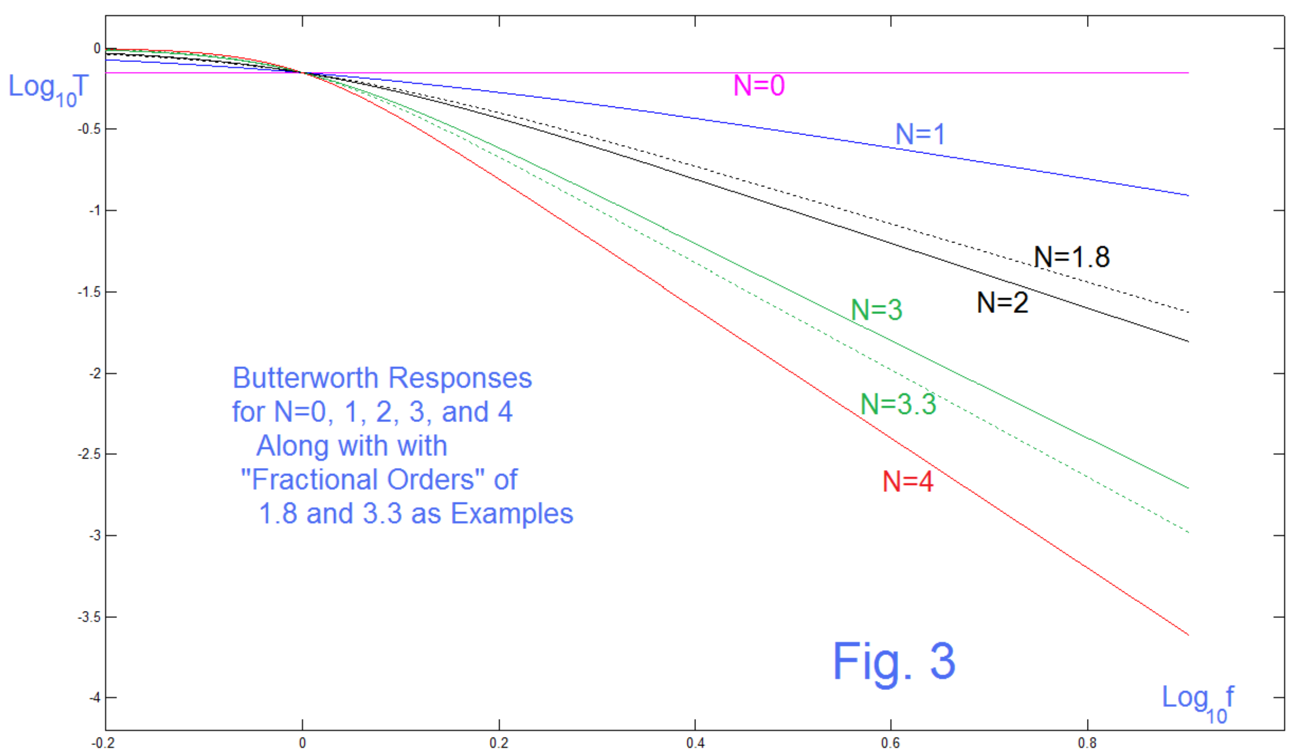
$Log_{10}T$
N=0
N=1
N=1.8
N=2
N=3
N=3.3
N=4
Butterworth Responses
for N=0, 1, 2, 3, and 4
Along with with
"Fractional Orders" of
1.8 and 3.3 as Examples
Fig. 3
$Log_{10}f$
0
-0.5
-1
-1.5
-2
-2.5
-3
-3.5
-4
-0.2
0
0.2
0.4
0.6
0.8

Keeping in mind that equation (1) is strictly applicable only to a Linear-Time-Invariant (LTI) system, in which case we are not correct to interpret it further than at best "quasi stationary", we will nonetheless think of a time varying filter with magnitude (we are ignoring phase) of |T(f,t)| where the cutoff f_c or the order N varies (N no longer restricted to integers).

There are thus two time-varying cases. In the conventional case of the usual VCF we have in mind that the cutoff frequency is a function of time. Typically in our music work this cutoff varies slowly, like once for a tone of duration lasting a good fraction of a second to several seconds. Taking the cutoff to be some function of time $f_c(t)$, equation (1) becomes equation (2) as:

$$|T_c(f,t)| = \frac{1}{\sqrt{1+(\frac{f}{f_c(t)})^{2N}}} \qquad (2)$$

To our embarrassment, it would seem that over the years we have given little thought to $f_c(t)$! Fig. 4 shows our usual "standard patch".

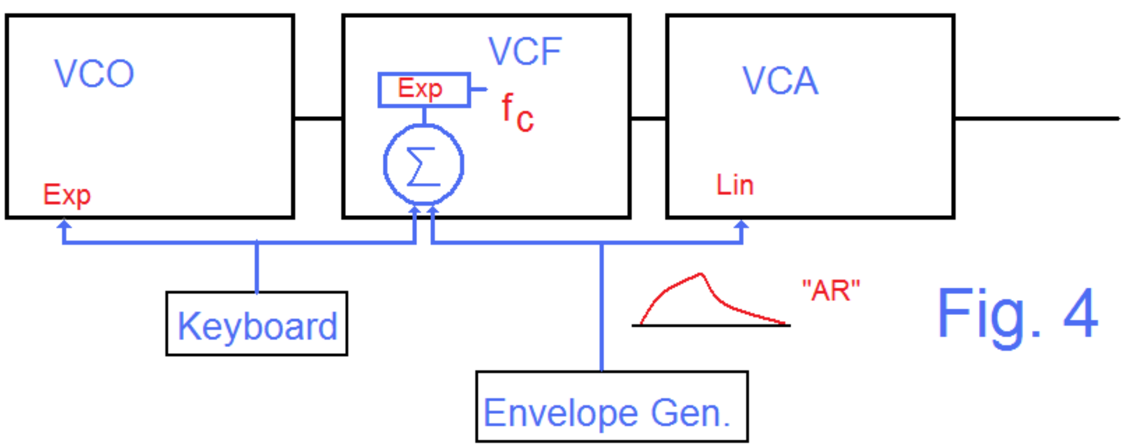

Fig. 4

In Fig. 4 we show the usual cascade of a VCO (Voltage-Controlled Oscillator) into a VCF and then a VCA (Voltage-Controlled Amplifier). The VCO produces a frequency controlled (typically) by a keyboard by an exponential relationship to a linear control voltage, and we also feed this frequency control to the VCF so that the cutoff of the filter tracks the VCO. But the VCF is "dynamic" in that it changes its cutoff during any one tone, and hence has a second control voltage supplied by an envelope generator. As suggested, this second input is typically summed with the linear control prior to exponential conversion that produces the cutoff f_c. Also typically the envelope is the same as the one that controls the amplitude via the VCA. It could be a separate envelope, but typically all envelopes have exponential sections of seen in the "AR" (Attack-Release) as in Fig. 4. The important thing about the time-varying exponentials of the envelopes is that they give natural-sounding amplitude decays (Release here) as opposed to linear decays that end too abruptly. [Physical system (acoustic musical instruments) decay at least approximately exponentially.]

Thus in Fig. 4, we have two thought-out connections. (1) It makes sense that the keyboard controls both the VCO and the VCF in an exponential manner. They "track". Note well that this is an exponential function of voltage (not of time). (2) Further it makes sense that the VCA controls amplitude as a linear function of an exponentially time-varying envelope. But we have, more or less ad hock, connected the amplitude determining envelope (or similar) to the exponentially-controlled VCF by summation. We need not quibble with this choice, especially at this time in history, and likely this is a situation where the general nature of what is done, not the exact details, is what matters. The point is that we really don't know what $f_c(t)$ should be!

A somewhat parallel situation is found with the V-S filter, except there would seem to be no pre-established reason to have an exponential relationship already in place. The slope would not seem to have to track the keyboard, although the cutoff f_c would continue to do so. The envelope to the VCF would instead go directly to a slope control. There is no apparent reason not to start with a linear relationship to slope control, probably in response to something like the AR (time exponential) envelope. This is not a situation dictated by tradition (there is no tradition for V-S). And since we are just programming anyway, we can experiment quite easily.

Making the slope variable is a simple matter of making N a function of time N(t) so that equation (1) now becomes equation (3):

$$|T_s(f,t)| = \frac{1}{\sqrt{1+(\frac{f}{f_c})^{2N(t)}}} \tag{3}$$

We are now in a position to use an additive synthesis approach of calculating amplitude functions |T(f,t)| for each frequency f of interest, multiplying them by corresponding sinusoidal waveforms, and summing to simulate the filtering. We can then look at and listen to results.

THE EXPERIMENTS

We now attempt an experiment. This will be a Matlab program that creates a file to study, and rather than just presenting the program we want to show a diagram that better illustrates the structure of the program, as we are doing a number of subtle things, and we need to be aware of where choices were made (and where alternative choices are certainly possible). Fig. 5 shows the scheme.

As suggested, the test setup here is basically additive synthesis simulating a filter, and as such, we have in mind a set of input components f_k each multiplied by corresponding frequency response magnitudes T_k, and then summed before being amplitude shaped. For example, the input frequencies might be the harmonics of Fig. 1. The point is that the frequency response is time varying and derived from an envelope by a complex

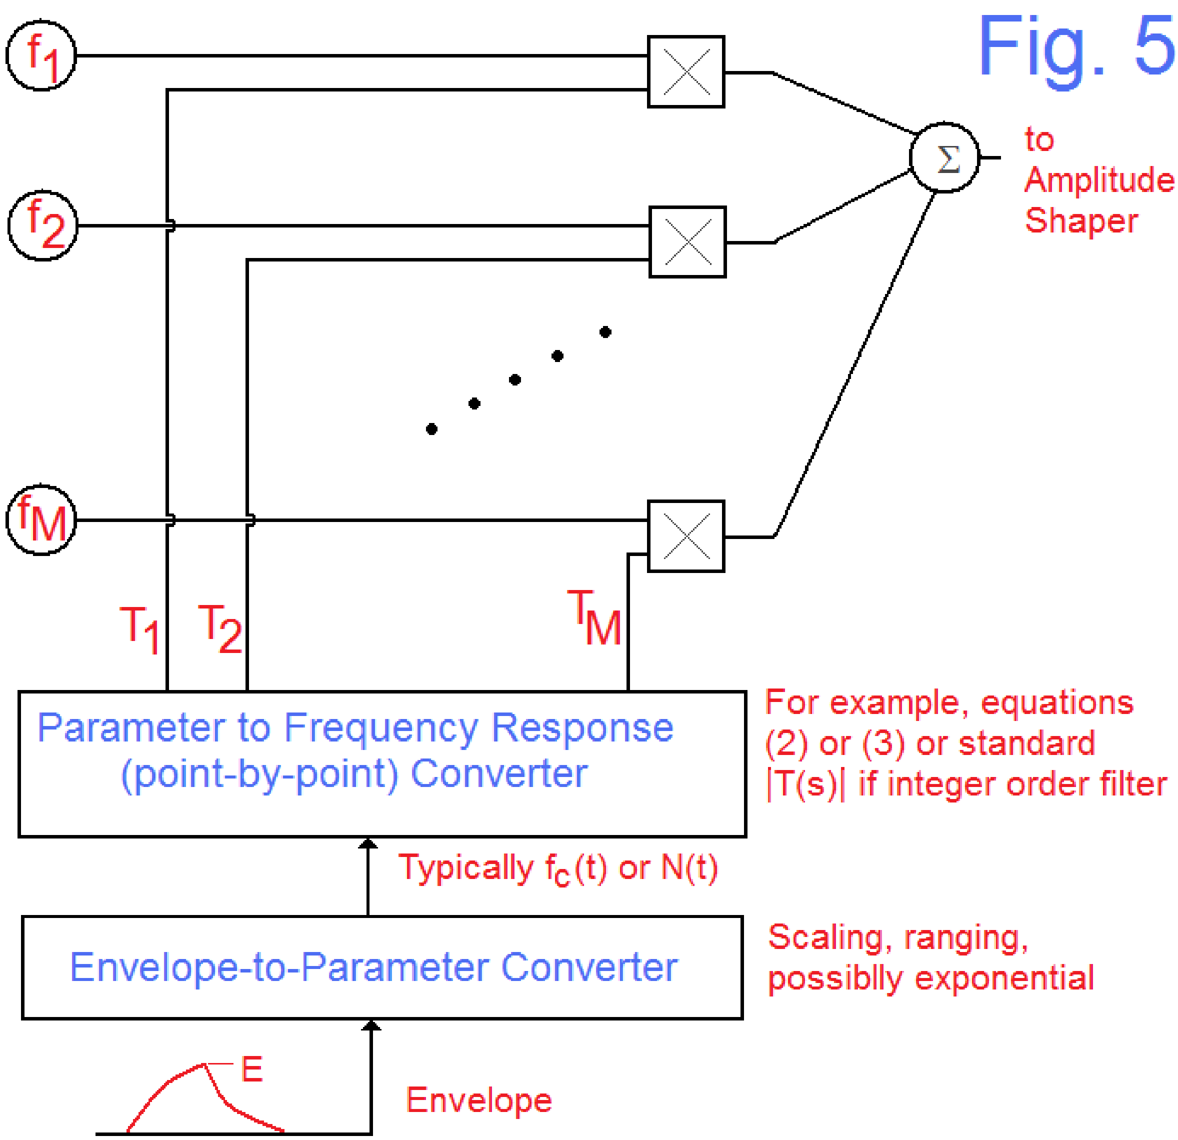

relationship. We show this processing from the envelope to the T_k, k=1,2,…M in two steps. First there is the need to take an envelope and scale/range it to some standard range (the "Envelope-to-Parameter Converter" box) appropriate to the box above it. For example, perhaps we need to adjust to a typical standard VCF range to fit a variable cutoff frequency $f_c(t)$, equation (2). So perhaps $f_c=2$ is at the bottom of the envelope (at 0), (the second harmonic), and the top of the envelope (E), is where we want a cutoff $f_c=10$, with a linear control relationship (as in Fig. 1). If on the other hand, we have a V-S filter the bottom of the envelope might be the N(t)=4 case while the top of the envelope might be the N(t)=0 case. That is, the V-S filter "flap" opens as the envelope rises. Once this is set, the actual $T_k(t)$ can be calculated from equations (2) or (3). Or we might be implementing the equivalent of a Moog 4-Pole VCF in which case we would use a more general equation for the frequency response, which would work fine since the order would be a fixed integer. Thus each of the T_k; T_1, T_2, …T_M of Fig. 5 is a time-varying control curve, calculated for a particular study.

It turns out that the "tools" available to do this experiment are not complicated. The problem is deciding exactly what to do. The following are considerations to be understood here:

(1) We desire here to hear the results, which means that the tones under study must be in the audio range and on the order of a second long. For example, perhaps 300 Hz for 1.5 seconds, which would be 450 total "cycles". Thus we can't expect to see the details if we plot the whole thing.

(2) Also in terms of listening to results, we don't want to have just a gated tone (Off/On/Off) but something with a proper envelope. The simplest meaningful envelope is simply an instantaneous attack (On) followed by an exponential decay, that typical of percussive and/or very short tones. A more complicated envelope tends to confuse our observations – at least for a start.

(3) Here we are concerned with time-varying filtering. We will not be trying to understand an associated frequency response. Instead we will have curves of amplitudes vs TIME (not frequency) which tell us how particular components evolve with time (the T_k of Fig. 5). The curves may resemble frequency response curves, but are different.

(4) While we will be imposing an envelope before listening to synthesized tones, when we plot the tones we will do this plotting in small sections (just hundreds of samples), without the envelope, so that changes in amplitude do not confuse us.

(5) Particularly as we will be playing our results, it will be useful to specify our harmonic components to be filtered in terms of typical frequencies, so instead of f, 2f, 3f,…8f we prefer things like 300 Hz, 600 Hz, 900 Hz, . . . 2400 Hz.

EXPERIMENT 1: JUST AMPLITUDE ENVELOPED

Can we form a tone from 8 harmonics and impose an exponentially decaying envelope on it? Of course we can. We choose (arbitrarily) 16,000 samples at a 8000 Hz sample rate, with frequencies of 300, 600, 900, 1200, 1500, 1800, 2100, and 2400 Hz. This is strongly pitched at 300 Hz, and about 2 seconds long. To this we apply an exponentially decaying envelope e^{-3t}. The result is shown in Fig. E1-a, and as mentioned, we basically only see the envelope since the details are lost in the 16,000 samples. Fig. E1-b shows the first 12.5 milliseconds of the waveform prior to multiplying by the envelope. [Because it is so short, the enveloped version is an identical plot here.] We see in Fig. E1-b the sum of the eight sinewaves (blue) and the 300 Hz sinewave (red). Note that the higher frequency components (to 2400) are approaching half the sampling frequency (half being 4000). So while adequately sampled, there are only about three samples/cycle and the plot LOOKS rough.

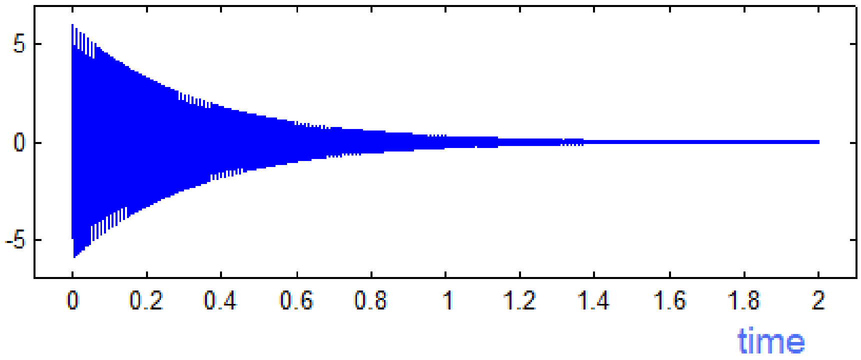

Fig. E1-a

Eight Harmonics Added and Enveloped

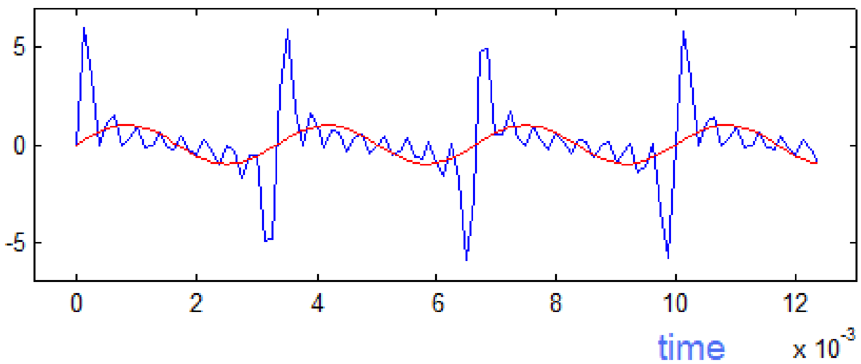

Fig. E1-b

First 12.5 ms of E1-a

EXPERIMENT 2: FIXED FILTER

In Experiment 1 we just established a baseline case, and have done no filtering. All eight component frequencies started out at amplitude 1 and decayed exponentially. In Experiment 2, we just do a fixed low-pass. Here we choose a fixed 2nd-Order BW with the cutoff at the second harmonic of 600 Hz. This is only a bit more exciting than Experiment 1. Fig. E2-a shows 12.5 ms of the waveform that corresponds to Fig. E1-b. Note that a good deal of the higher harmonic content is removed (blue curve) and we also again show the 300 Hz fundamental in red. Fig. E2-b shows the eight harmonic amplitudes as a function of time, and this is not too interesting as they are all constant in time (the filter is, after all, fixed) but we do see the roll-off of the higher harmonics. If we made a corresponding figure for Experiment 1, it would have been one line at amplitude 1.

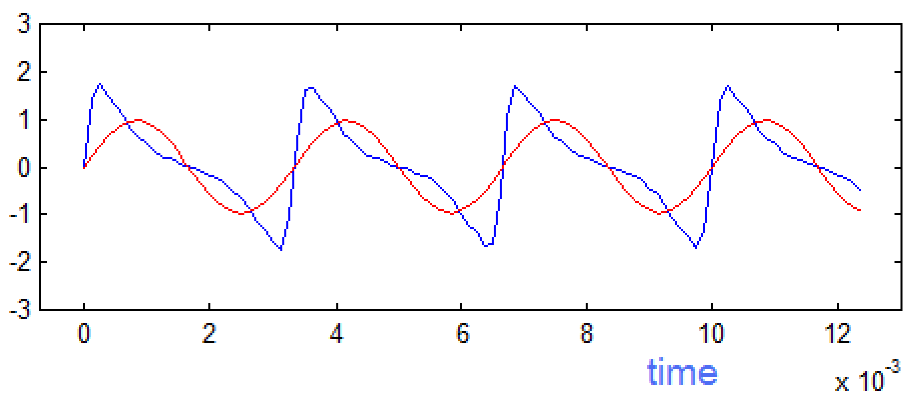

Fig. E2-a

Low-Pass of Eight Harmonics

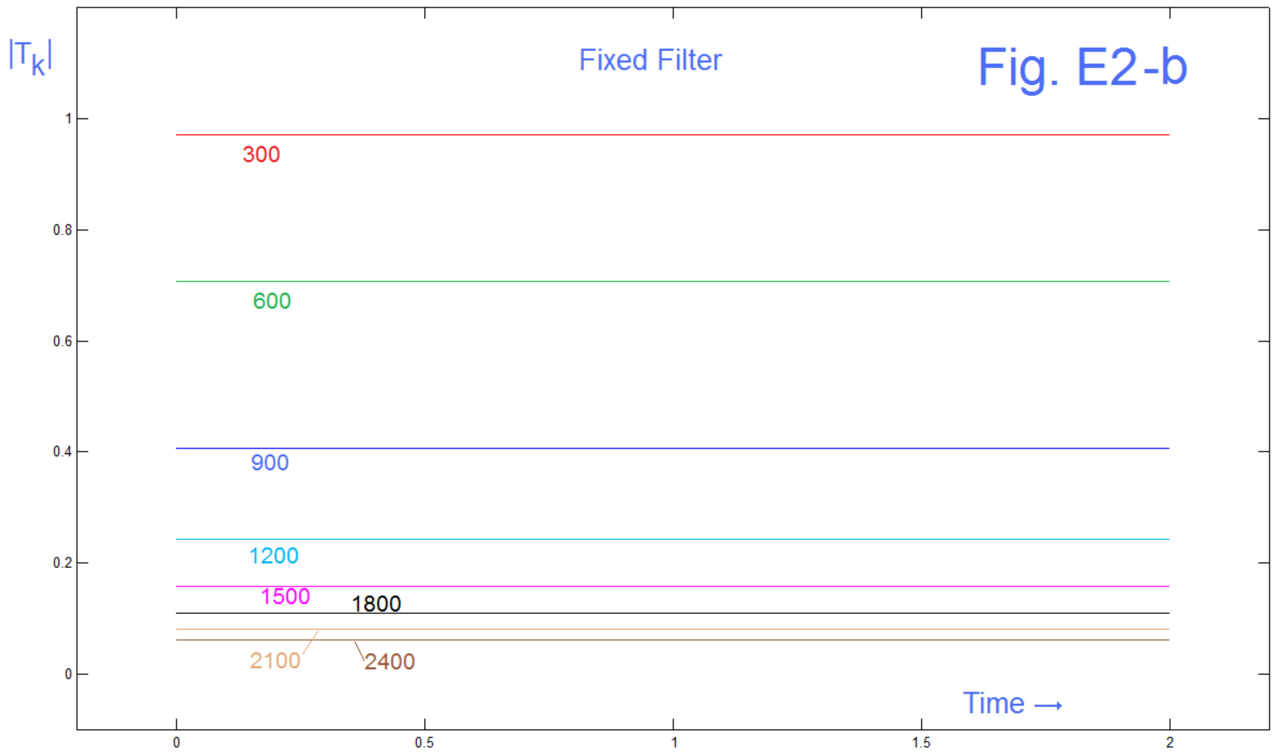

EXPERIMENT 3: MOVING FILTER (VCF)

At this point, we have still just done two control experiments, leading to just an ordinary fixed filter. Next we want to make the filter cutoff move. Thus we need an envelope which we will convert to a control parameter, f_c in this case, and then to a corresponding frequency response using equation (2) with N fixed at 4, and this gives us the quasi-stationary frequency response values for each of the eight components.

To be explicit, we set time t, the envelope e (Fig. E3-a), and the cutoff fc as:

t = 0 : 0.000125 : 2 - 0.000125

e = exp(-3*t)

fc = 600 (1 + 4*e)

so as the envelope drops exponentially from an initial value of 1 down to 0, the cutoff drops from 3000 down to 600 (Fig E3-b). Note that as time progresses (left to right in Fig. E3-c) the upper harmonics fall off first (as a function of time), except for the 300 Hz red component which is well below the cutoff at all times and remains at 1. Note that when the cutoff reaches 600 Hz (at 2 seconds), the corresponding green curve is at about 0.707.

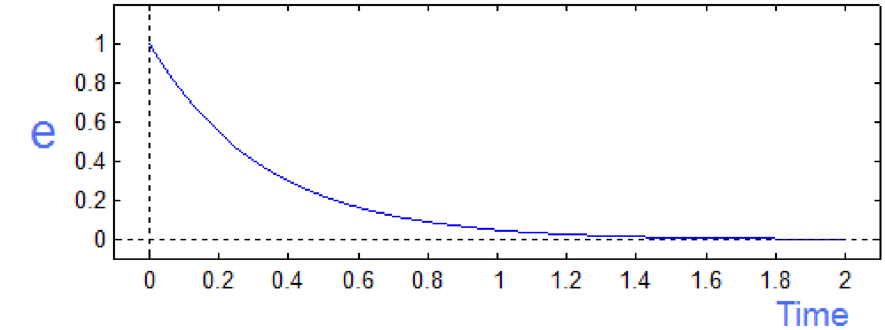

Fig. E3-a

Envelope

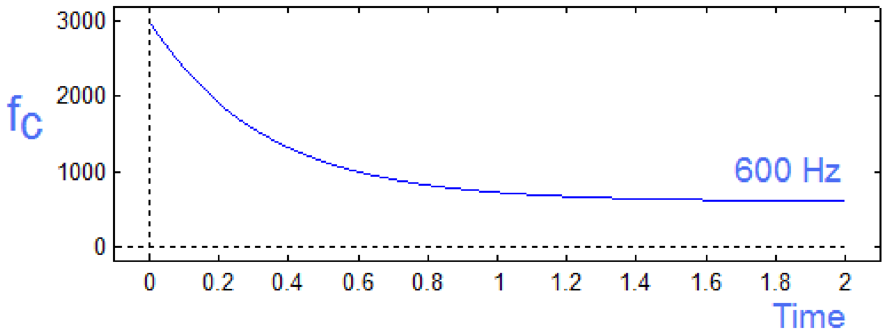

Fig. E3-b

VCF Cutoff

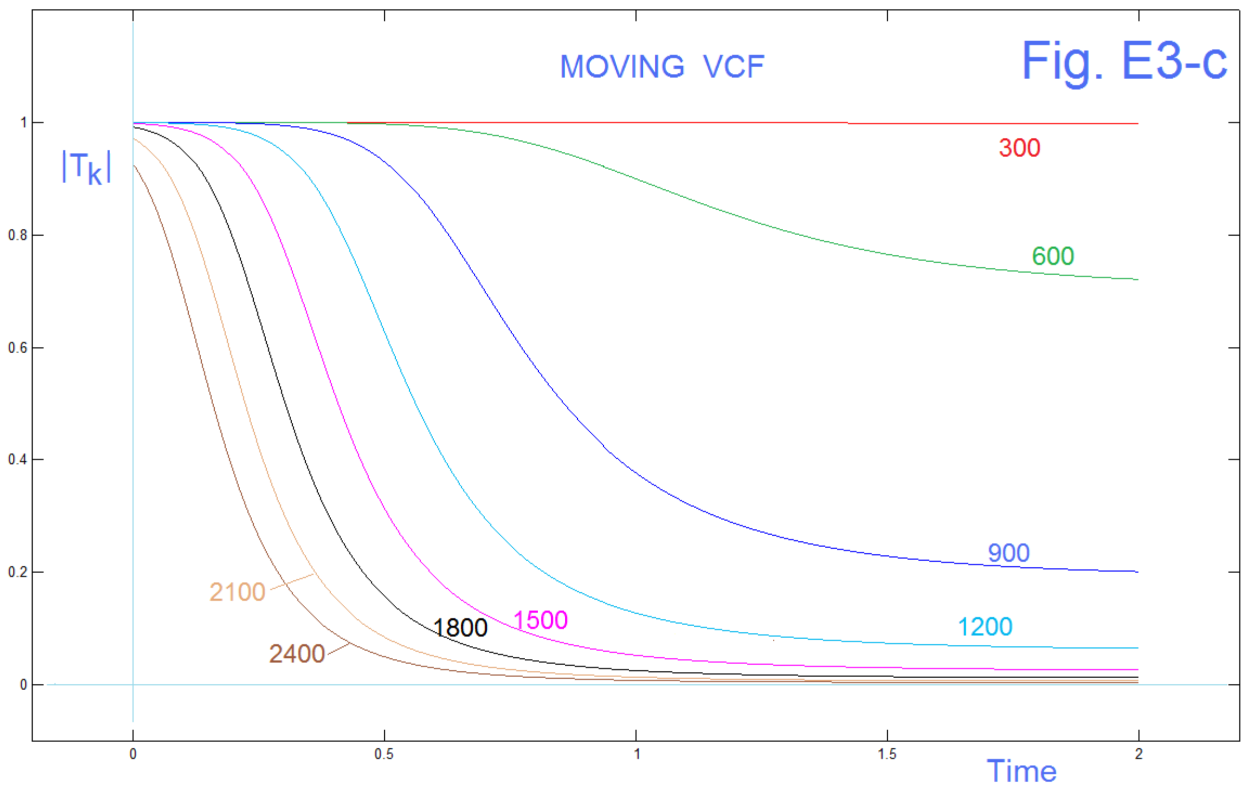

Fig. E3-c

Here Experiment 3 is a basic example of what occurs with a VCF in a "standard patch" with two exceptions which we need to keep in mind. First, for simplicity we are using only an exponential decay (Release) for the envelope, not a full ADSR (Attach-Decay-Sustain-Release). This could be implemented easily, and would cause the result to be reflected, to pause, and to change rates possibly. Secondly VCF's used in practice would likely have a peaking (resonance) near the cutoff (see bandpass experiment below).

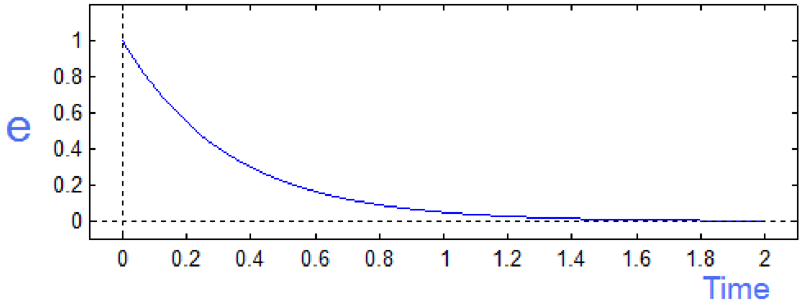

Fig. E4-a

Envelope

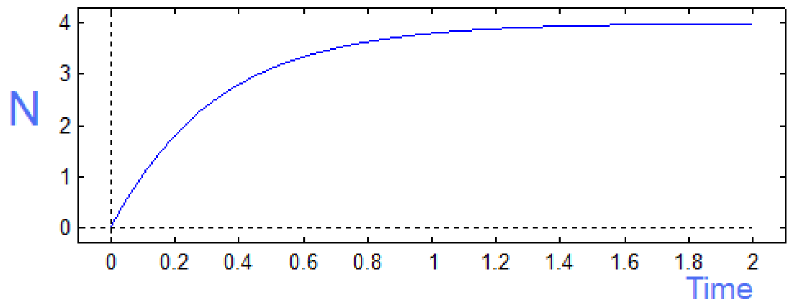

Fig. E4-b

Order

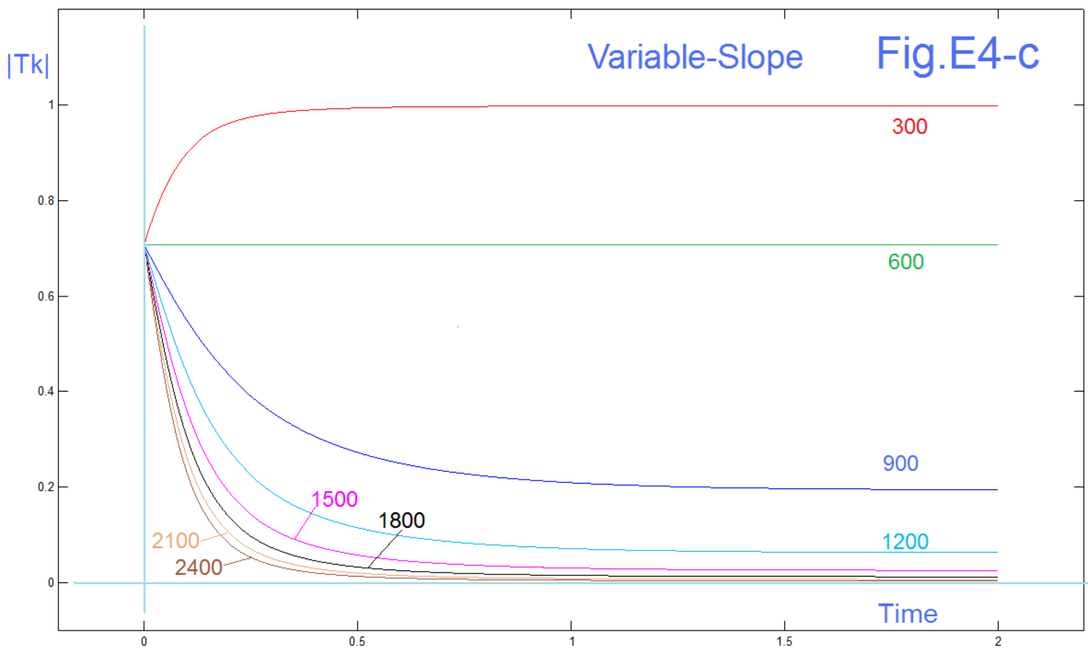

Fig.E4-c

EXPERIMENT 4: VARIABLE-SLOPE

Here in Experiment 4 we have a V-S example. We begin with the same (Fig. E4-a) envelope as was used in Fig. 3a except here we then derive from it not the cutoff f_c but rather the variable-order N(t) (Fig. E4-b) giving the variable-slope according to:

N = 4 - 4*e

which means that we start with 0 slope (e=1) and the slope increases to N=4 (24 db/octave) as the envelope e decays to 0. Thus, as shown in Fig. E4-c, we start with everything getting through and close down the "flap" to just a 4th-Order BW – the same ending point as Fig. E3-c. We see a very similar distribution as the time approaches 2 seconds in both figures. The difference is in the rejection processes. In Fig. E3-c, we see the higher components disappearing somewhat more successively than in Fig. E4-c where they go out in more like a parallel attenuation. It is also true that the fact that the first harmonic (300 – red curve) went up first was a surprise. This is because the initial 0th-order BW was compelled to go through 0.7071.

EXPERIMENT 5: BAND-PASS

In a bit we will comment on what these V-S tones sound like. But for the moment, we note that all the results for Experiments 1 through 4 have been listened too and that the results are just noticeably different, but not too conclusive. In particular, the V-S result is far less "striking" than the very familiar resonant VCF sound so characteristic of the synthesizer. Thus we might want to compare this to something like the Moog 4-Pole LP VCF [12]. A somewhat easier experiment that turns out more revealing is to use a Band-Pass. This we can do easily by employing the parameters of Experiment 3, Fig. E3-b, as the center frequency of a band-pass. This means we can choose a Q (of 2) and compute the transfer curves for the band-pass by equation (4), reference [13], similar to what we have used above for LP. The results are shown in Fig. E5.

$$|T_B(f)| = \frac{Q}{\sqrt{1 + Q^2(\frac{f}{f_C} - \frac{f_C}{f})^2}} \qquad (4)$$

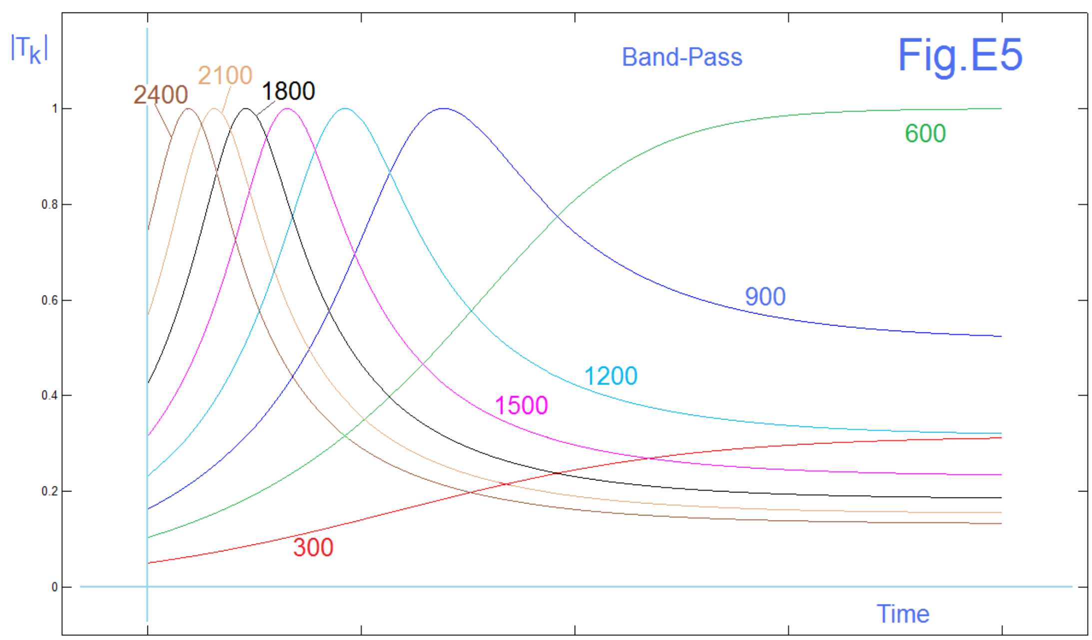

What we have in Fig. E5 turns out to be somewhat different looking than the low-pass cases. (We comment immediately that the sound does have the "wwwwoooooowww" twang we are familiar with from the resonant low-pass analog filter.) Note that here we have a band-pass of moderate Q sweeping from a center frequency of 3000 Hz at time 0 down to about 600 Hz as time approaches 2. This is especially evident from the green curve for 600 Hz, where the filter ends up. Further we see the filter band-pass sweep over the harmonics at 2400 Hz, 2100 Hz, ... 900 Hz. It never gets down to 300 Hz, although we see the red curve becoming sufficiently strong by time 2 that we understand a strongly pitched 300 Hz tone. The point of the experiment was to identify the familiar "interesting" synthesizer sound, which commands our attention, mainly <u>with the resonance</u> and not so much with the moving (but flat, Butterworth) low-pass cutoff.

THE SOUNDS OF THE EXPERIMENTS

We are at the point here where we have the pictures and understand them pretty well. What do the results sound like. Here we are in the usual difficult situation of describing in words what something sounds like. If we could do it well, we would record the sounds as an audio file and post that file. But the differences in sounds are subtle and tend to escape the recording process. Perhaps someone can do this well enough, and we would welcome such an effort. In the mean time, we offer the Matlab file that produces the figures and the sound files (below). The code is straightforward enough that it could probably be adapted to other program languages. Indeed, actually using the program would offer a much wider variety of examples for the interested researcher.

There are five sound files generated here, corresponding to the Figures E1 through E5. There is really an "Experiment 0" where we just listen to two full seconds, full amplitude, or all eight harmonics – sharp, jarring, and artificially unpleasant, suitable only as an alarm tone. By imposing, the exponentially decaying envelope (Experiment 1), it becomes just tolerable, something suitable for a musical composition although not the most pleasant of tones. We know that this amplitude shaping is necessary – you can't just suddenly end a tone without creating a "click". It remains our observation that you <u>can</u> suddenly <u>start</u> a tone (exactly as we have done here) and that sounds fine. This is almost certainly because we are accustomed to hearing percussive musical instruments (not to mention impulse induced sounds all day long from a multitude of sources). You can have gradual start-up envelopes (a commented-out line in the program). If the make the start-up (attack) too long, it sounds like a pipe-organ starting up. Nothing new in Experiment 1.

Jump now to Experiment 5 where we used a band-pass. The point here was to understand that it is the sweeping resonant peak that is responsible for the most characteristic "woosh" sound of the analog synthesizer. This sound was what was new and exciting – not like any acoustic instrument in particular. It is this exploitation of these circuitry possibilities, along with some perhaps unintentional side-effects (such as clipping) that made the synthesizer so exciting. All this despite what we knew or pretended to be investigating based on traditional acoustical instruments.

This leaves Experiments 2, 3, and 4: all low-pass cases. Because we did not employ extra resonance, or other "tricks" (except through always-interesting errors!) we were within the realm of imitating traditional acoustical tones. [Indeed, in citing David Luce's work, this was explicit.] In consequence, none of the three experiments is going to jump out as a new, never-before-heard synthetic sound. To some degree, this means boring! In another sense, the subtle results are very instructive. In particular, we wish first to know if a V-S tone is heard to be distinctive from one for which the cutoff moves, and secondly, if the V-S tone is in any sense "better". The answer to both questions can be "Yes" although some listeners would find it marginal. All the more reason why it is perhaps difficult to record.

Here we have not done enough investigation. Perhaps what is most important here is the laying out of possible tools. So we will begin by describing the three tones exactly as produced by the program. Recall that Experiment 2 was a fixed filter, Experiment 3 a usual VCF type, and Experiment 4, the V-S filter. All three are similar when considered relative to the edgy, harmonically-rich Experiment 1 and the animated "woosh" of the sweeping resonant Experiment 5. As produced by the program, I can hear no real difference between the fixed filter (Experiment 2) and the VCF case (Experiment 3), except one can just detect a change of timbre at the very end of Experiment 3. Both are much more listenable (relative to Experiment 1) due to the reduction of the harmonics. The V-S case (Experiment 4) sounds different from the other two in being a very natural sounding "ping" with a just enough animation in the upper harmonic evolution. Based on this, we might think we have a winner.

If we compare Fig. E3-c to Fig. E4-c we notice that the harmonic evolution is much the same at the ends. In the middle, however, the V-S case decays much faster. It matter not that the imposed amplitude envelope (e) is exactly the same, as the variations of the individual harmonic amplitudes out of the filter itself are different and dominate the apparent amplitude. The tone in Experiment 4 seems shorter. Since we do see this rather abrupt decay in Fig. E4-c, we might try to adjust this. The decay rates out of the filter are more or less accidentally related to the control envelope by our "converters" and so, one way to make up apparent differences is to change the decay constant to the envelope. The quick test here is to make the envelope of the V-S case longer, using e^{-2t} instead of e^{-3t}. Doing this, all three of Experiments 2 through 4 sound more similar.

Again, we have not so much found any definitive results as indentified some tools and some parameters. Further, all this was done with no investment in building any hardware. More time and a more careful study seems to be indicated.

I also made a preliminary effort at an internet search. This did not yield anything solid – at least not at a first try. I used the search term "Variable Slope Filter" and did find some links. Most of these were from an electronic music viewpoint, including what seem to be available products. Not much in the way of discussing technology. Some mentioned the Electronotes articles. The other search term I tried was "fractional order butterworth" and this looked interesting, including some very recent (2013) IEEE stuff (paywalled!). Two articles [14, 15] of apparent academic origins do come up, and are included here to encourage someone else to follow up. While there are suggestions (like mention of Sallen-Key) that suggest, perhaps, that someone is thinking about hardware, I did not find any circuits. But I didn't try all that hard.

REFERENCES

[1] B. Hutchins , "Negative Frequencies – and Through-Zero Frequency Modulation", ***Electronotes***, Vol. 22, No. 206, Dec. 2010
http://electronotes.netfirms.com/EN206.pdf

[2] B. Hutchins, "Revisiting Pulse-Width Modulation," ***Electronotes***, Vol 23, No. 216, July 2013
http://electronotes.netfirms.com/EN216.pdf

[3] B. Hutchins, "What About: Self-Frequency-Modulation," ***Electronotes***, Vol. 23, No 217, Sept. 2013
http://electronotes.netfirms.com/EN217.pdf

[4] B. Hutchins, "Some Theory on the Design and Application of Variable-Slope Filters," and "ENS-76 VCF Option 3", ***Electronotes***, Vol. 8, No. 72, Dec. 1976, pages 3-11 and pages 12-14 respectively

[5] B. Hutchins, "Phase Shift Networks and Applications, ***Electronotes***, Vol. 7, No. 59, Nov. 1975

[6] David Luce, "Dynamic Spectral Changes of Orchestral Instruments," ***J. Audio Eng. Soc.***, Vol. 23, No 7, Sept 1975

[7] Normand Provencher, "Notes on Noise, Fractional Calculus, and Variable-Slope Filters," ***Electronotes***, Vol. 11, No. 108, Dec. 1979.

[8] Lester Ludwig, "Uniform Variable-Slope Processes," ***Electronotes***, Vol. 12, No. 109, Jan. 1980

[9] Lee Powell, "A Solution to the Mystery of the Variable-Slope Filter," ***Electronotes***, Vol. 12, No. 112, April 1980.

[10] Normand Provencher, "A Brief Note on Variable-Slope Spectra and Self-Modulation," ***Electronotes***, Vol. 13, No. 121, Jan. 1981

[11] Normand Provencher, "A Few More Variable-Slope Ideas," ***Electronotes***, Vol. 13, No. 123, March 1981

[12 B. Hutchins, Revisiting Some VCF Ideas – And a Few New Ideas," ***Electronotes***, vol. 23, No. 215, May 2013
http://electronotes.netfirms.com/EN215.pdf

[13] B. Hutchins, "Frequency Response of a Bandpass Filter," Electronotes Application Note No. 271, Feb. 15, 1983

[14] Achary, Anish et al, "Extending the Concept of Analog Butterworth Filter for Fractional Order Systems" (unspecified source and date)
http://arxiv.org/ftp/arxiv/papers/1210/1210.8194.pdf

[15] Li, Ming, "Approximating Ideal Filters by Systems of Fractional Order" ***Computational and Mathematical Methods in Medicine*** Volume 2012 (2012), Article ID 365054,
http://dx.doi.org/10.1155/2012/365054
http://www.hindawi.com/journals/cmmm/2012/365054/

PROGRAM – For Documentation Purposes

```
% vs.m  Variable Slope - Experiments for EN#224
%
%
%
%  ORIGINAL TONE OF EIGHT EQUAL HARMONICS - ENVELOPED
t=0:.000125:2-0.000125;
S1=sin(2*pi*300*t);
S2=sin(2*pi*600*t);
S3=sin(2*pi*900*t);
S4=sin(2*pi*1200*t);
S5=sin(2*pi*1500*t);
S6=sin(2*pi*1800*t);
S7=sin(2*pi*2100*t);
S8=sin(2*pi*2400*t);
Sorig=S1+S2+S3+S4+S5+S6+S7+S8;
Sorig1=Sorig;
e=exp(-3*t);
%e=[ [0:.001:.999],e(1:14801)];
Sorig=Sorig.*e;
```

```
figure(6) % Fig. E1-a and E1-b of Report
subplot(211)
plot(t,Sorig)
axis([-.1 2.1,-7,7])
subplot(212)
plot(t(1:100),Sorig(1:100))
hold on
plot(t(1:100),S1(1:100),'r')
hold off
axis([-.0007 .013 -7 7])
%
%
%
% ADD FIXED LOW-PASS
N=2
fc=600;
T1=1 ./ sqrt( 1+(300/600).^(2*N)   );
T2=1 ./ sqrt( 1+(600/600).^(2*N)   );
T3=1 ./ sqrt( 1+(900/600).^(2*N)   );
T4=1 ./ sqrt( 1+(1200/600).^(2*N)   );
T5=1 ./ sqrt( 1+(1500/600).^(2*N)   );
T6=1 ./ sqrt( 1+(1800/600).^(2*N)   );
T7=1 ./ sqrt( 1+(2100/600).^(2*N)   );
T8=1 ./ sqrt( 1+(2400/600).^(2*N)   );
figure(1)  % Fig. E2 of Report
plot(t,T1*ones(1,16000),'r')
hold on
plot(t,T2*ones(1,16000),'g')
plot(t,T3*ones(1,16000),'b')
plot(t,T4*ones(1,16000),'c')
plot(t,T5*ones(1,16000),'m')
plot(t,T6*ones(1,16000),'k')
plot(t,T7*ones(1,16000),'r')
plot(t,T8*ones(1,16000),'g')
hold off
axis([-.2 2.2 -.1 1.2])


S1=S1*T1;
S2=S2*T2;
S3=S3*T3;
S4=S4*T4;
S5=S5*T5;
S6=S6*T6;
S7=S7*T7;
S8=S8*T8;
SLP=S1+S2+S3+S4+S5+S6+S7+S8;
SLP=SLP.*e;
figure(10) % Fig. E2-a of Report
subplot(211)
plot(t(1:100),SLP(1:100))
hold on
plot(t(1:100),S1(1:100),'r')
hold off
axis([-.0007 .013 -3 3])
%
%
%
```

```
% MOVING LOW-PASS - ORDINARY VCF
fc= 600*(1+ 4*e);
figure(7)   % Fig. E3-a and E3-b of Report
subplot(211)
plot(t,e)
hold on
plot([0 0],[-1 2],'k:')
plot([-1 3],[0 0],'k:')
axis([-.1 2.1 -.1 1.2])
hold off
subplot(212)
plot(t,fc)
hold on
plot([-1 2],[0 0],'k:')
plot([0 0],[-1000 4000],'k:')
hold off
axis([-.1 2.1 -200 3200])
N=4
T1=1 ./ sqrt( 1+(300./fc).^(2*N)   );
T2=1 ./ sqrt( 1+(600./fc).^(2*N)   );
T3=1 ./ sqrt( 1+(900./fc).^(2*N)   );
T4=1 ./ sqrt( 1+(1200./fc).^(2*N)   );
T5=1 ./ sqrt( 1+(1500./fc).^(2*N)   );
T6=1 ./ sqrt( 1+(1800./fc).^(2*N)   );
T7=1 ./ sqrt( 1+(2100./fc).^(2*N)   );
T8=1 ./ sqrt( 1+(2400./fc).^(2*N)   );
figure(2)  % Fig. E3-c of Report
plot(t,T1,'r')
hold on
plot(t,T2,'g')
plot(t,T3,'b')
plot(t,T4,'c')
plot(t,T5,'m')
plot(t,T6,'k')
plot(t,T7,'r')
plot(t,T8,'g')
hold off
axis([-.2 2.2 -.1 1.2])


S1=S1.*T1;
S2=S2.*T2;
S3=S3.*T3;
S4=S4.*T4;
S5=S5.*T5;
S6=S6.*T6;
S7=S7.*T7;
S8=S8.*T8;
SMLP=S1+S2+S3+S4+S5+S6+S7+S8;
SMLP=SMLP.*e;
%
%
%
```

```
% VARIABLE-SLOPE LOW-PASS
%
% or change decal constant
e=exp(-3*t);
%
%
N=4-4*e;
figure(8)  % Fig. E4-a and E4-b of Report
subplot(211)
plot(t,e)
hold on
plot([0 0],[-1 2],'k:')
plot([-1 3],[0 0],'k:')
axis([-.1 2.1 -.1 1.2])
hold off
subplot(212)
plot(t,N)
hold on
plot([-1 2],[0 0],'k:')
plot([0 0],[-0.3 4.3],'k:')
hold off
axis([-.1 2.1 -0.3 4.3])
fc=600;
T1=1 ./ sqrt( 1+(300/600).^(2*N)   );
T2=1 ./ sqrt( 1+(600/600).^(2*N)   );
T3=1 ./ sqrt( 1+(900/600).^(2*N)   );
T4=1 ./ sqrt( 1+(1200/600).^(2*N)   );
T5=1 ./ sqrt( 1+(1500/600).^(2*N)   );
T6=1 ./ sqrt( 1+(1800/600).^(2*N)   );
T7=1 ./ sqrt( 1+(2100/600).^(2*N)   );
T8=1 ./ sqrt( 1+(2400/600).^(2*N)   );
figure(3) % Fig. E4-c of Report
plot(t,T1,'r')
hold on
plot(t,T2,'g')
plot(t,T3,'b')
plot(t,T4,'c')
plot(t,T5,'m')
plot(t,T6,'k')
plot(t,T7,'r')
plot(t,T8,'g')
hold off
axis([-.2 2.2 -.1 1.2])


S1=S1.*T1;
S2=S2.*T2;
S3=S3.*T3;
S4=S4.*T4;
S5=S5.*T5;
S6=S6.*T6;
S7=S7.*T7;
S8=S8.*T8;
SVS=S1+S2+S3+S4+S5+S6+S7+S8;
e=exp(-3*t);
SVS=SVS.*e;
e=exp(-3*t);
%
%
%
```

```
% MOVING BAND-PASS
fc= 600*(1+ 4*e);
Q=2
T1=1./ sqrt( 1+(Q^2)*(300./fc - fc./300).^2   );
T2=1./ sqrt( 1+(Q^2)*(600./fc - fc./600).^2   );
T3=1./ sqrt( 1+(Q^2)*(900./fc - fc./900).^2   );
T4=1./ sqrt( 1+(Q^2)*(1200./fc - fc./1200).^2   );
T5=1./ sqrt( 1+(Q^2)*(1500./fc - fc./1500).^2   );
T6=1./ sqrt( 1+(Q^2)*(1800./fc - fc./1800).^2   );
T7=1./ sqrt( 1+(Q^2)*(2100./fc - fc./2100).^2   );
T8=1./ sqrt( 1+(Q^2)*(2400./fc - fc./2400).^2   );
figure(4)  % Fig. E5 of Report
plot(t,T1,'r')
hold on
plot(t,T2,'g')
plot(t,T3,'b')
plot(t,T4,'c')
plot(t,T5,'m')
plot(t,T6,'k')
plot(t,T7,'r')
plot(t,T8,'g')
hold off
axis([-.2 2.2 -.1 1.2])
S1=S1.*T1;
S2=S2.*T2;
S3=S3.*T3;
S4=S4.*T4;
S5=S5.*T5;
S6=S6.*T6;
S7=S7.*T7;
S8=S8.*T8;
SBP=S1+S2+S3+S4+S5+S6+S7+S8;
SBP=10*SBP.*e;
%
%
%
% PLAY SOUNDS
%
pause
sound(Sorig,8000)
pause
sound(SLP,8000)
pause
sound(SMLP,8000)
pause
sound(SVS,8000)
pause
sound(SBP,8000)
```

*   *   *   *   *   *