

# ELECTRONOTES 215

Newsletter of the Musical Engineering Group

1016 Hanshaw Road, Ithaca, New York 14850

Volume 23, Number 215

May 2013

## REVISITING SOME VCF IDEAS – AND A FEW NEW IDEAS

### STATE-VARIABLE STABILITY

#### INTRODUCTION:

When it comes to voltage-controlled filters VCF's as used in analog music synthesizers, only two basic design schemes have been popular: (1) The classic "Moog Four-Pole Low-Pass" and (2) the "State Variable" multi-mode (Fig. 1). First we can ask why it has boiled down to these two.

While there are many many styles of active filters, and many of these have found a niche in which they are most useful, the main thing that needs to be considered about a VCF is of course the voltage-control. That is, the frequency characteristics of the filter need to be variable, and to vary in response to a voltage (usually to a current, actually) which can be varied faster and automatically (and usually with more accuracy) than with manual control (like turning a knob). Achieving this control is the fundamental consideration. This is the equivalent of achieving one or more (more usually two or four) voltage-controlled resistors. For practical reasons, this is often difficult. The most useful approach to achieving voltage-control of an overall filter has been to drop back to building from simple voltage-controlled building blocks (like integrators and first-order low-pass sections) and accept any limitations that result. Musically, these limitations are not severe at all.

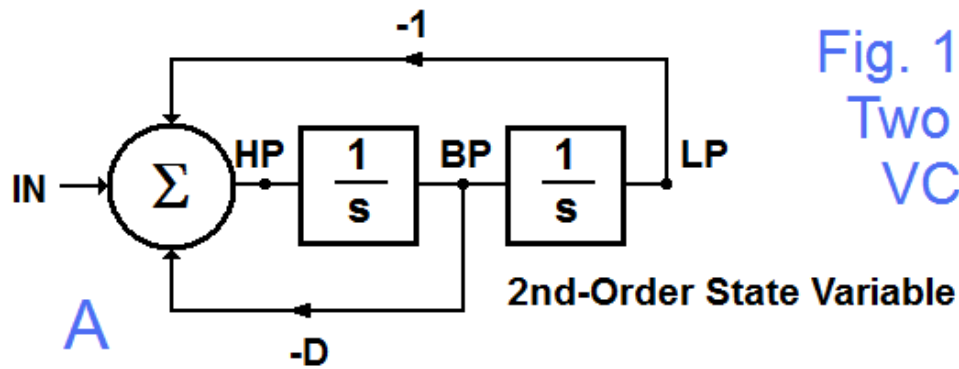
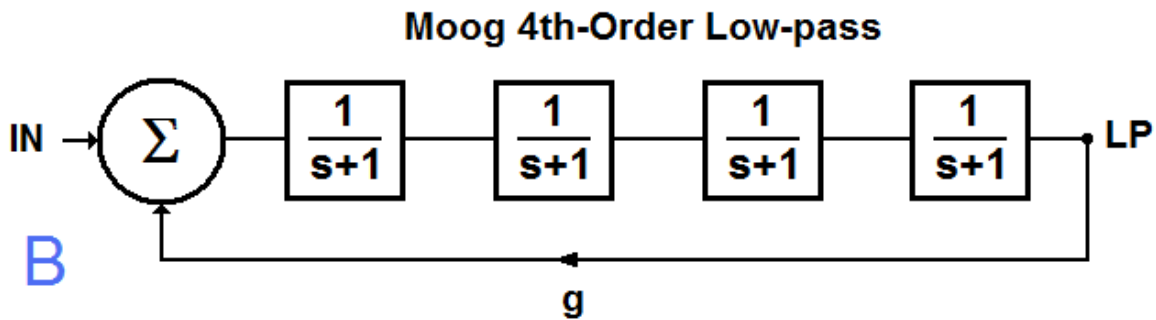


Fig. 1  
Two Common  
VCF Structures



We often hear about a person being an “inventor” when what this person really does is more properly called engineering – of a product – often more as an entrepreneur. In the case of the four-pole low-pass (Fig. 1B) however, Bob Moog is quite properly called the inventor. In support of this statement, for one thing, there was a famous patent that resulted from the particular implementation. But perhaps more importantly, Bob did more or less tinker until he got things to work. He used the classic approach of cascading a filter section for higher order, and then by inspiration from radio work (regeneration, or positive feedback, possibly inspired from his work with Theremins) tried positive feedback to sharpen the response. This was essential, and successful.

In this case of the four-pole low-pass, there was really NOT a corresponding tradition of a non-voltage-controlled version (no active filter using this approach). This was probably because there were better ways of achieving a four-pole response (typically, a Butterworth response). Everyone knew about the Butterworth, but it was not obvious how to efficiently make it voltage-controlled. Further, as it turned out, there may be good musical reasons why the four-pole with feedback was at least “good enough” and perhaps superior for the synthesizer. Much literature with regard to the four-pole is available [1-5]. Here we will

remark, with regard to the stability properties, that the avoidance of instability (or the active decision to achieve it – i.e., oscillation) became well understood. A pair of poles (two of the four) could be placed on the  $j\omega$ -axis with a feedback gain of -4. The Moog VCF was frequently used with this oscillation, or at least in a highly resonant mode.

So the Moog four-pole low-pass was achieved using voltage-controllable building blocks (the first-order low-pass) with an unconventional feedback scheme not found in the fixed filter bag of tricks. On the other hand, an even simpler voltage-controlled building block, the integrator, was available to be exploited. This time, there WAS a fixed prototype active filter – the so-called “State Variable Filter” (SVF) already in use (Fig. 1A). It was basically also known from the days of “analog computers” being two integrators and a summer in a loop. It was an analog “computer” in the sense that it simulated a second-order response. And it “simulated” it by producing voltages corresponding to the magnitudes of the variables (“state-variables”) of the system. In simulating the variables, it WAS itself a second-order system! [Today we think of simulation as something that is done on a digital-computer, usually written in a higher-order program language. What the SVF was to analog synthesis is pretty much what the digital technique of “physical modeling” is to digital synthesis. Sometimes the best way (sometimes only way) to examine something, particularly a chaotic system, is to “realize” it on a computer, sample-by-sample.]

While well described, initially the SVF was not widely used as an active filter. Simply this was because it required three op-amps, and achieved only second-order, while configurations using only one op-amp would often work as well or better (Sallen-Key, MFIG, etc.). And op-amps were initially expensive. Less obvious initially was the fact that the SVF was not a good performer from a sensitivity point of view (component tolerances mattered). The shining feature of the SVF had to be its “biquad” (bi-quadratic) potential. Using it, one could achieve (via an additional summer usually) a transfer function that had a numerator that was second – order – having powers of  $s$  that were  $s^0$ ,  $s^1$ , and  $s^2$ . Our readers will recognize this as suggesting that the filter could produce not just a low-pass response ( $s^0$ ), but a bandpass ( $s^1$ ) and a highpass ( $s^2$ ), simultaneously.

Thus, a synthesizer designer looking to construct a VCF naturally considered the SVF (likely in part to avoid patent issues). On the downside, it was only 2<sup>nd</sup> order. On the upside, it offered two additional responses. Both these issues were somewhat superficial. A full study of the 4-pole showed that two of the four poles dominated the response anyway. It did offer 24db/octave (rather than just 12 db/octave) roll-off, true enough. [Some will argue that there is something else “mystical” about the 4-pole that does not show up in the engineering specifications.] As for the three responses of the SVF, this also was a dubious advantage. Many (most?) of the useful musical sounds achieved tended to be low-pass (like their

acoustical precursors). Bandpass was sometimes used, but virtually the same thing could be had with a highly resonant (corner-peaked) low-pass. And it seemed, no one wanted high-pass.

## **SENSITIVITY WITH VARIABLE PARAMETERS – AN EXPERIMENT**

It is not too difficult to draw block diagrams and even to make a real design assuming ideal components. When it comes to real filters with real components, a wealth of information relating to “Sensitivity” is available [6]. For the most part, this pertains to what we might call “tweaking” or “fine tuning”, or it can be intentional “overdesign” knowing that actual performance will drop away from the nominal. This might relate to random occurrences of actual values of resistors and capacitors (up or down about a nominal marked value), or the general unidirectional “failure” of some active component to perform ideally (in reality, they all have speed limitations).

Even these are not too difficult for fixed specifications of an active filter. When it comes to a variable filter however, things can get very slippery. That is, we can make a fixed filter stand still where we want it, but when we try to make it offer a similar or corresponding performance over a range of one of its design parameters (typically the cutoff frequency) than additional complications are the rule. These can be so complicated that we often resort to tinkering right at the bench (like a few small phase-lead capacitors inserted) rather than engineering analysis [2]. Our electronic music VCF’s are not only variable, but variable on the time scales of individual musical notes. You can’t take these out of the case and adjust them! They have to be at least stable over a predictable range.

In a recent review [7] we looked again at this problem. A theory of compensation was examined. Does this work, and if so, how does it compare to a intuitive approach? Below we give brief results of using a test circuit for a non-voltage-controlled state variable filter (Fig. 2). We have taken some experimental data, but this is not meant to be, nor is it, conclusive. The circuit was breadboarded with LF13741 op-amps – BiFET input and 741 output. We chose these slow op-amps because we were looking to demonstrate limitations. It was also tried in part with LF351 op-amps. The circuit with the three switches shown does not really exist. The changes were made by moving wires. Some reader may want to take more data.

Fig 2 shows the test setup for the SVF. The full setup has six op-amps, the three blue ones being the original SVF structure, and the three red ones are the “active compensation” additions. In the setup, as mentioned, we show “switches” although these are just the

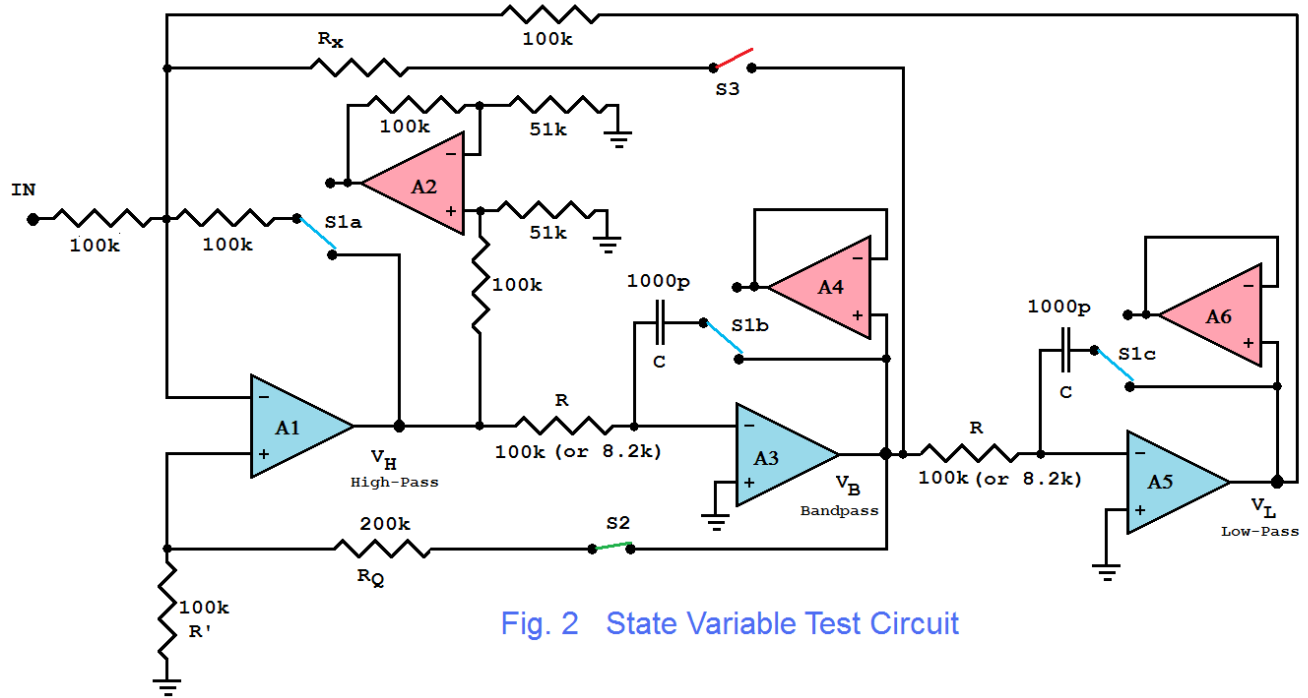
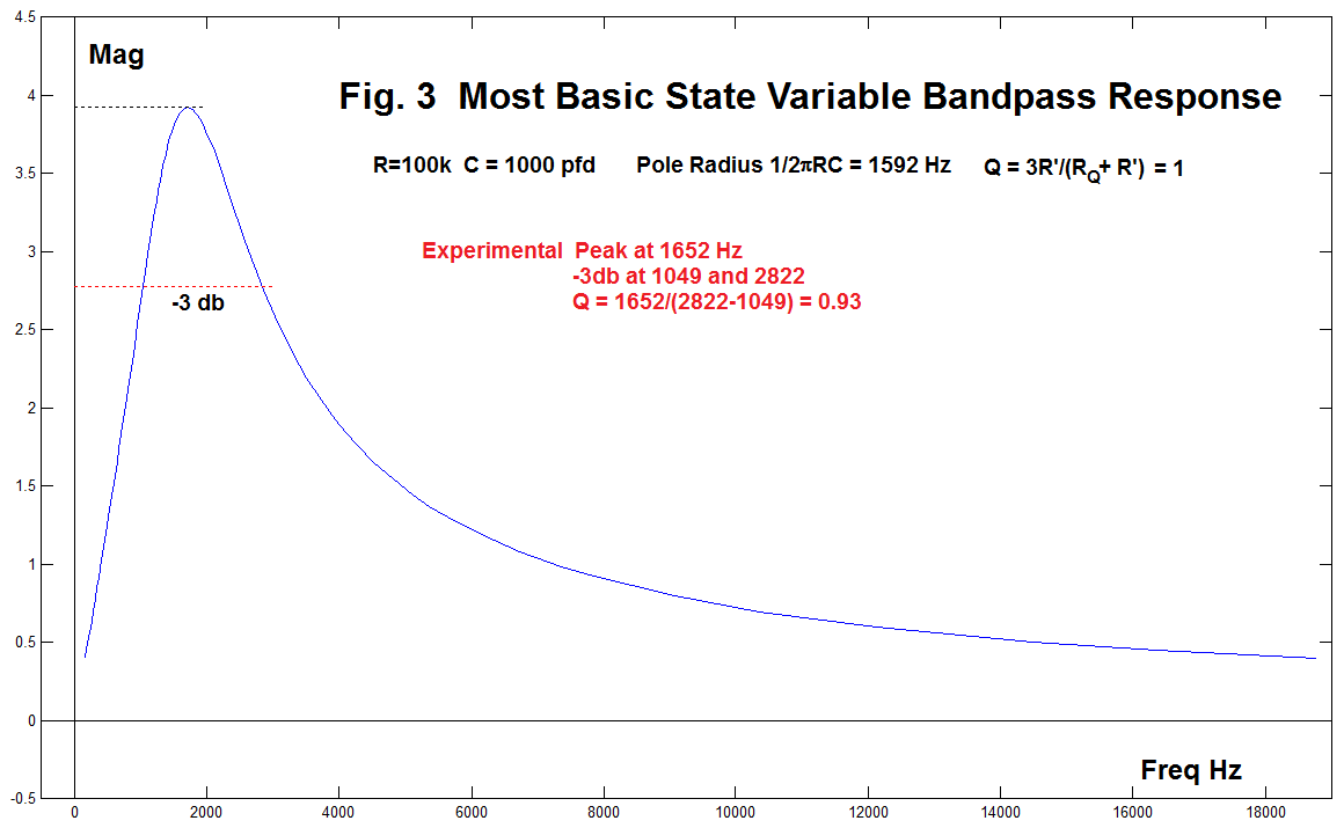


Fig. 2 State Variable Test Circuit

changes of jumpers on the breadboard, not actual switches. With the switches as shown, the red op-amps in the feedback loops are bypassed, so their outputs don't go anywhere. The switches S1a, S1b, and S1c (blue) are all to be thrown as a unit. The switch S2 can be used to disable the damping (by opening the green switch, which is closed here). With  $R_Q = 200k$  as shown, the  $Q$  is set for  $Q=1$ , a very modest response. [Note that  $Q = 1/D$ , and here the gain  $D$  is determined by the voltage divider,  $R' / (R' + R_Q)$  times a non-inverting gain of 3 from the (+) input of A1 to its output]. If we open switch S2 (green), all feedback (all damping,  $D$ ) from the bandpass output is blocked. In this case, the  $Q = 1/D$  is theoretically infinite, and the structure theoretically is a sine-wave oscillator. Thus, we want to observe if we open switch S2 if the filter oscillates: it doesn't although the  $Q$  shoots up. The red switch S3 shown as open here is not a part of a regular SVF. Instead it is a "cheater" path that allows us to switch in some positive feedback. Note that this is a path from the bandpass (already inverted because of the inverting integrator) to the (-) input of A1, hence positive feedback relative to that to the (+) input of A1 from  $R_Q$ . Note that the resistors  $R$  can be either 100k (a test at low frequencies) or 8.2k (a test for high frequencies). By placing the three S1 switches in the upper positions, we can test the same basic SVF with the compensated integrators and summer. Lot's of options to test.

The first test (shown in Fig. 3) is the baseline, original structure (three op-amps), low frequency ( $R=100k$ ), and  $Q$  of 1, green switch closed. In this case, the frequency at the center of the bandpass should be  $1/2\pi RC = 1592\text{Hz}$ . Experimentally we got the center of the bandpass at 1652 and a  $Q$  of 0.93.



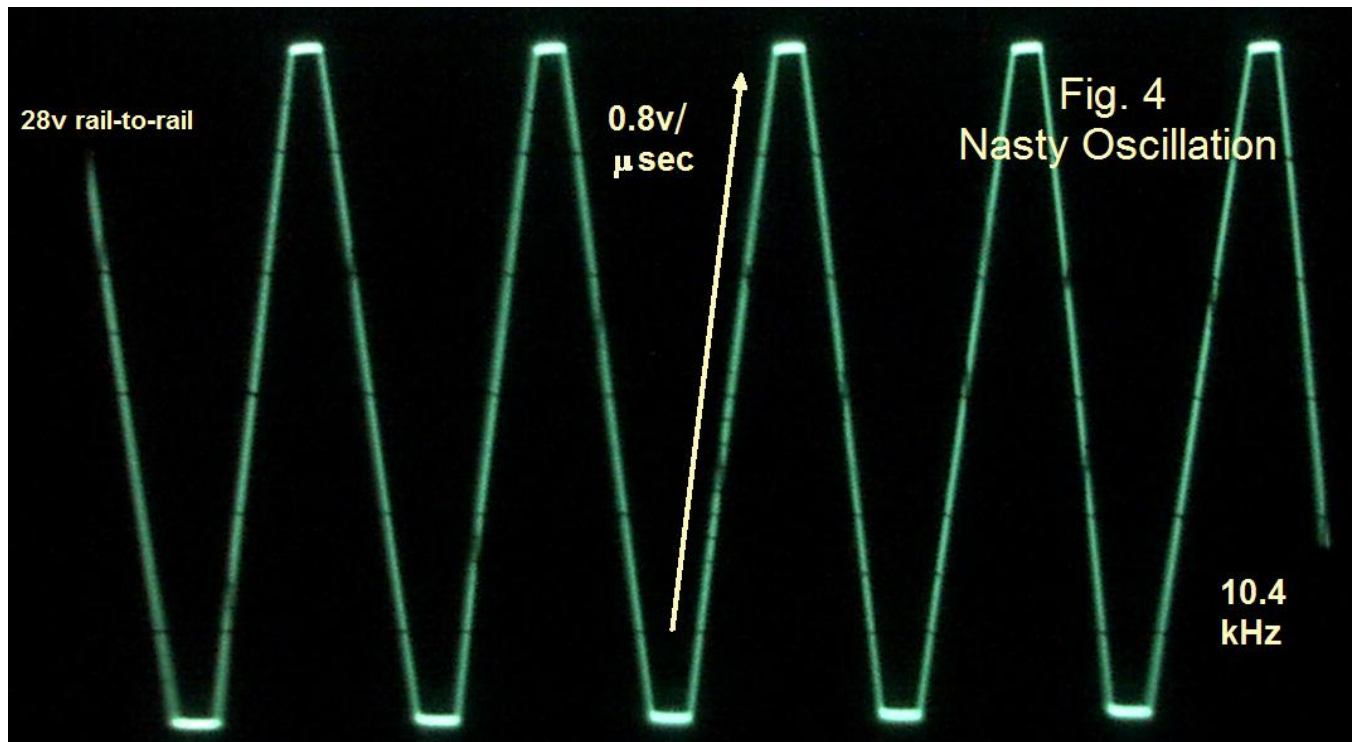
Our second test is to see if the circuit actually does work with the six op-amps. It does, and looks very similar to the three op-amp version. Experimentally, the center frequency was 1715 Hz and the Q was 0.95. Since this is only one measurement, no significant difference should be inferred between the three and six op-amp versions based on these results.

In the third test, we removed the damping by opening the switch S2. As mentioned, in theory we get exactly a sinusoidal oscillator. In practice, the circuit did not oscillate. In these two cases, we just measured the Q. The three op-amp version had a center frequency of 1685 Hz and a Q of 73. The six op-amp version had a center frequency of 1699 Hz and a Q of 36. Since these did not oscillate with all damping removed (a result consistent with years of similar findings, but often discussed) we added the “cheater positive feedback using  $R_x$  through switch S3. In the case of the three op-amp version, a value of  $R_x$  of 10M caused a nice sinewave oscillation at 1692 Hz. An oscillation of slightly wavering amplitude was achieved with  $R_x=12M$ , and no oscillation at  $R_x = 15M$ . If  $R_x$  was set down to 5.1M, the sinewave clipped (too much positive feedback). With the six op-amp version,  $R_x=5.1M$  gave a clean sinewave at 1689 Hz, and oscillation did not occur at  $R_x = 10M$ . The conclusion is that at low frequencies, the state-variable does not self-oscillate, but requires a tiny bit of added positive feedback.

For the fourth test, we went to high frequency by changing R to 8.2k. Originally I tried 1k, but the results were far too unstable. Put another way, the expected Q-enhancement and corresponding self-oscillation was clearly observed when we tried for a 159 kHz center frequency by making R=1k. Note right here that it was NOT the case that the instability went away with the six op-amp version (perhaps it was worse – see below).

So we tried 8.2k for R, which would be a theoretical center frequency of 19,419 Hz (top of audio range). With the three op-amp version, this gave a center frequency of 19,851 Hz and a Q of 0.86 (still nominally 1). With the six op-amp version, the center frequency was 19,181 Hz and the Q was 1.32. Not what we expected?

Keeping in mind that the results here are very sketchy (we emphasize we are mainly suggesting doing these test and the test circuit), this is disappointing. But, we learn from observations. It is true that we were looking for a Q enhancement with the three op-amp version at high frequencies, and hoped the six op-amp version would correct this. Instead, we saw a slightly higher Q with the six op-amps. Moreover, two other things were seen which may provide a clue. The waveforms with six op-amps also showed what was apparently a superimposed high-frequency oscillation. That is, it was fuzzy, and in fact looked like what one often interprets as a weak harmonic coming through an extra strong high-frequency bandpass. I could not find this higher bandpass however, and the effect did not seem to depend on input frequency (as would be true if this were harmonics). Secondly, with six op-amps, we also



found a nasty oscillation that could occur (be triggered) by high frequencies well above the cutoff (although the oscillation was lower, about 10.4 kHz) and by high input amplitude. Please see Fig. 4. This oscillation was clipped, and the sides looked slew-limited, so something broadly characterized as “non-linear” was going on (see note at end [8]). Recall the “jump resonant” phenomenon [9], a known result of slew-limiting, such as we have with these LF13741’s. This oscillation might or might not go away if we reduced the input signal to zero, or we might need to turn off the power supply. It can be also cleared out by randomly messing with the input conditions. (Keep in mind that it is always possible to completely cancel an oscillation by inputting a single cycle of the oscillation inverted.) This certainly reminds us of overflow oscillations in digital filters! The experiment is trying to tell us something.

At this point the “box the toy came in is more interesting than the toy”. But this would require a lot more study, which I may or may not get to eventually. Here are a few more observations.

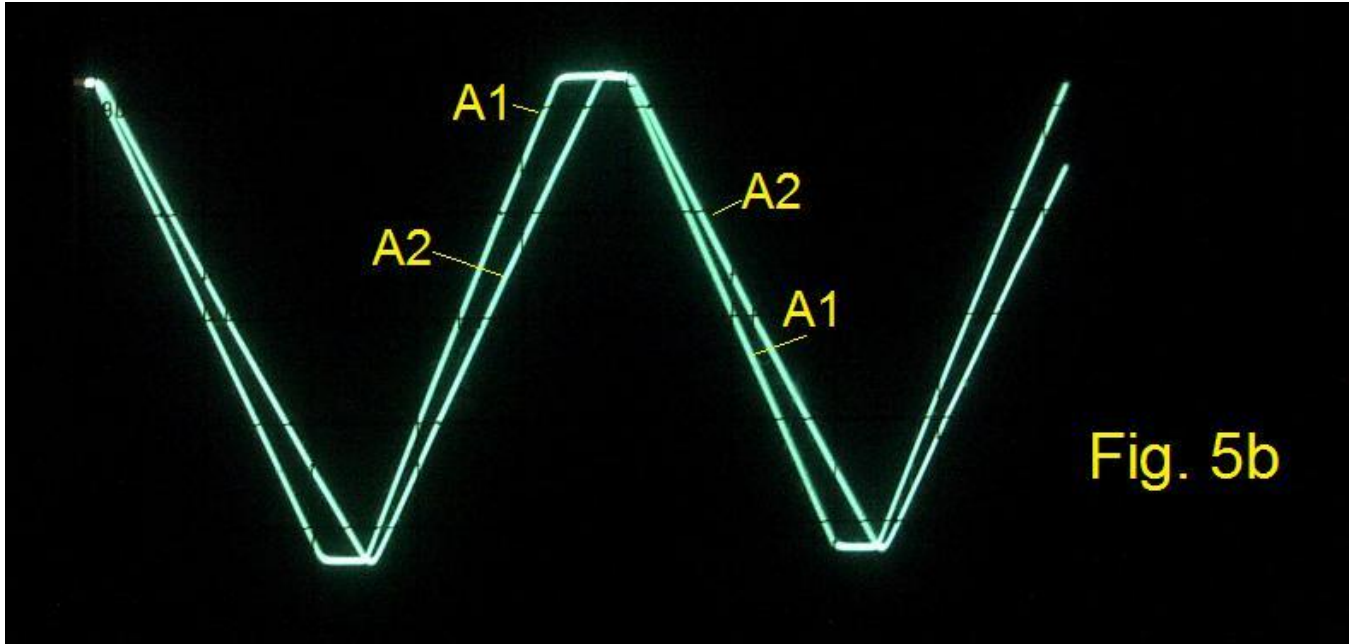

(1) The slew limiting at 0.8 volts/ $\mu$ sec seemed a bit high. Indeed, the data sheet gives 0.5 volts/ $\mu$ sec – the same as an ordinary 741 (denoted “typical”). Yet the slope on the scope and the frequency/amplitude observations agree. An independent test did give 0.8 v/ $\mu$ sec. Indeed, if we drive an open-loop SL13741 with a square wave of frequency 10,400 Hz, it is virtually identical to Fig. 4.

(2) A curious finding that seemed initially to be highly significant was that the waveform at the (+) input of A2 is not the same as 1/3 the output of A2 (Fig. 5a). A2 is supposed to be a gain of 3. Fig. 5b shows the similar comparison of the outputs of A1 and A2, which should be the same in the ideal case (A1 is the more symmetrical one). The differential inputs to A2 look like Fig. 5b, but are only 1/3 of the voltages. Thus we see that the difference between the inputs drives the A2 op-amp into slew limiting. But wait! If we had, for example, a series of op-amp followers in series, and the last one appears to be slew limited, which one is it that it actually slew limited? Well, the one that is (by chance) actually the slowest of the lot.

Accordingly, it at first seemed that since the voltage at the (+) input of A2 was not of sufficient slope to be slew limited (indeed it is only about 1/3 of what is necessary) that slew limiting occurs with A2, and that this may be the cause of the nasty oscillation. But it is as likely (in general) that any other op-amp in the loop is the slowest one. So this is a symptom (or a clue) but not an answer.

(3) Likewise, the other two followers (A4 and A6, which do actually have a gain of 1), like A2, have a difference between input and output waveshapes similar to Fig. 5b. This is a mess,





and probably switching op-amps around would change things. As a first test, a much faster LF351 can be substituted and all the waveform change. It is also possible to make all the

op-amps except one be the faster LF351s, so we know that the single LF13741 causes the slewing. This test was not done with enough care to report here.

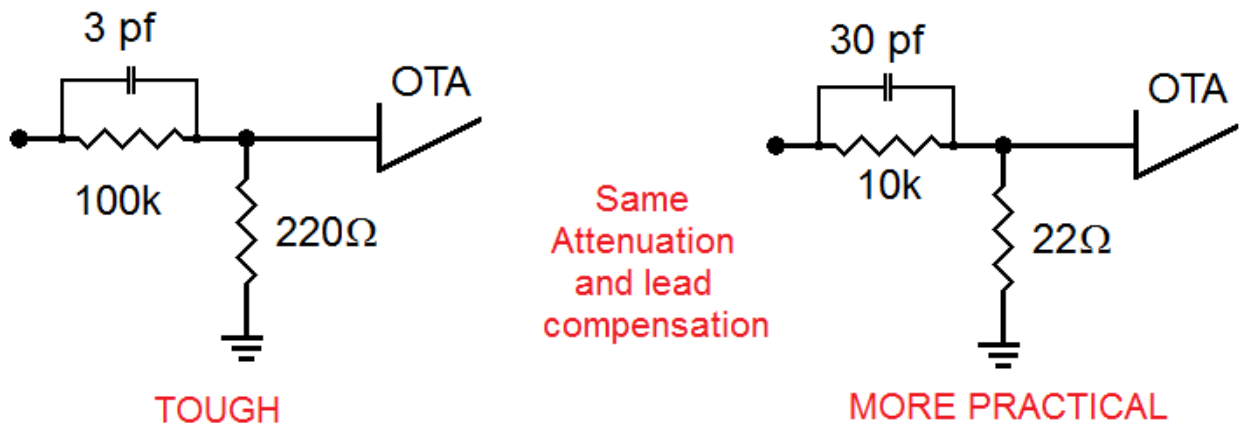
At this point, we have introduced an interesting toy, and will stop at that for now. What have we found that we conclude and want to note:

(1) At low frequencies, no automatic self-oscillation occurs by removing D. You have to add a tiny bit of positive feedback to get an oscillation.

(2) Q enhancement, from even a small Q, occurs at high frequency. The active compensation method did not fix this in an obvious or simple manner. Thus.....

(3) Lead compensation with shunt capacitors [2] is still the best practical solution. But we need to be prepared to reduce the resistance of the input leg of the input dividers to keep capacitors larger than stray. I believe all our current designs do this. Let's make this clear. We have the need for input attenuation to the OTA devices, and in our original work, this was a voltage divider of 100k to 220Ω. This required shunt capacitors on the order of 3 pfd. Capacitors this small are too small! Stray capacitances can already be this large. Accordingly, the same attenuation ratio could be achieved with 10k to 22Ω dividers, and the shunt capacitor then becomes something more manageable, like 30 pfd. This is shown in Fig. 6. Making the input leg even smaller would make the capacitor even larger and more convenient, but loading the previous stage might be a problem if we went to, say, 1k (not to mention the problem of finding 2.2Ω resistors).

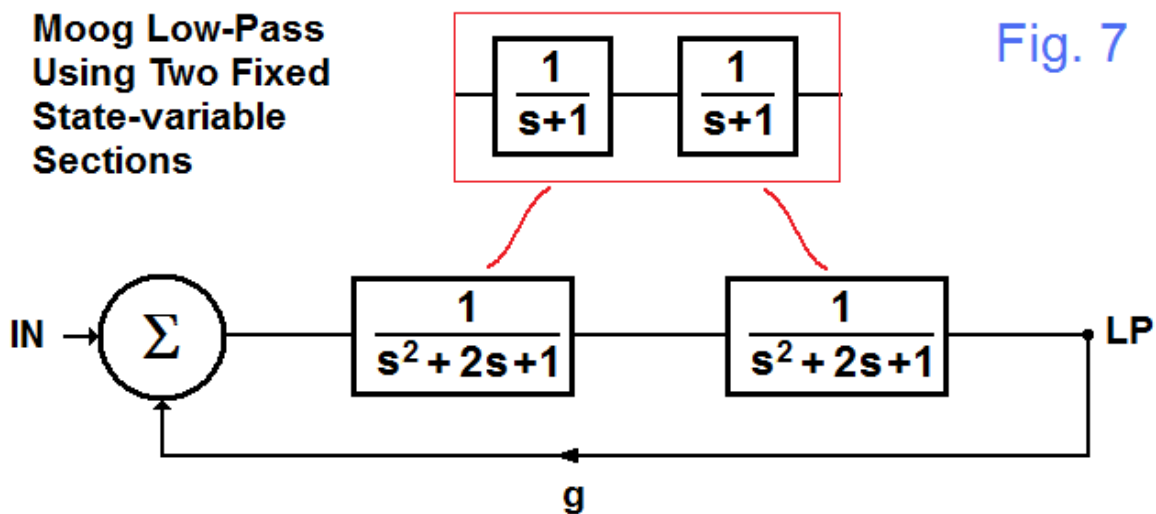
**Fig. 6 Shunt Compensation**



# REVISITING: STATE-VARIABLE TO 4-POLE

## CHOOSING BETWEEN STATE-VARIABLE AND FOUR POLE

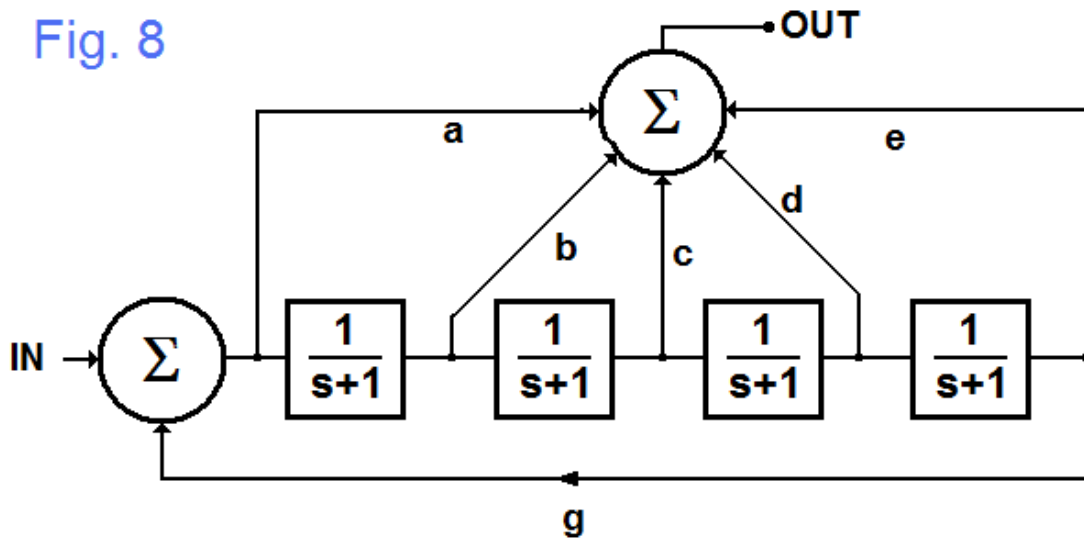
As we have mentioned, most commercial (and home-brew) analog synthesizers use either the state-variable or the 4-pole (Fig. 1) [1-4]. Many home builders have constructed both, or indeed many options. Another consideration is (was?) the availability of certain dedicated electronic music IC's for VCF's which have four transconductor sections. (At the same time, IC's with single OTA's are harder to find, while duals are more common.) In consequence, the notion of transitioning between S-V and 4-Pole approaches remains attractive. Fig. 7 shows an obvious configuration [5]



Here the two boxes in the lower portion of Fig. 7 are state-variable sections. We have set these to a fixed damping ( $D=2$ , or  $Q = 1/D = 1/2$ ), which correspond to two real poles at  $s = -1$ . The two second-order state-variable sections thus are (each) just the same as two wimpy first-order sections in series (red box). We assume that the cutoff frequencies are the same, equal to 1, as well. Thus Fig. 7 performs the same as Fig. 1B. In a bit, we will consider the possibilities of using different dampings in two different sections. For the moment, we have just done the Moog 4-Pole is a different way. The purpose was to better use available chips [5], and to avoid the dilemma of having to choose either the S-V or the 4-Pole.

First however, there are two other ways of manipulating designs related to a Moog 4-Pole structure [3]: one is to use feed-forward paths to implement zeros for bandpass, high-pass, etc. (essentially as we usually do for notch with state-variable), as indicated in Fig. 8, and the second is to use a series (cascade) of state-variable sections, the approach in Fig. 9.

Fig. 8



Feed-Forward Realization of a General 4th-Order Denominator

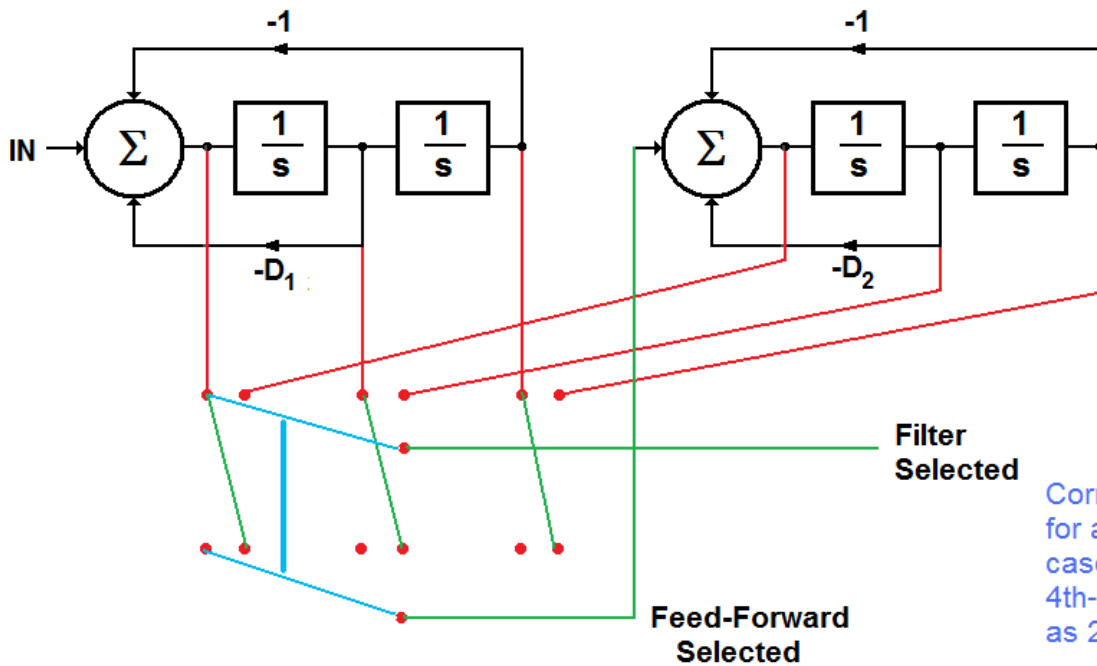


Fig. 9


Correct switching for a state variable cascade giving 4th-order as well as 2nd-order

The feed-forward approach (Fig. 8) is likely well-enough described in the previous reference [3] so will be left here just for review. The cascade approach is “trivial” but ripe with options. I appreciate a recent question from Nick Keller, who pointed out that this option was implemented (or at least attempted) in a factory-mod of a commercial synthesizer. This mod apparently did not use the second (lower) six pole ganged switch. In [3] we suggested only a four-pole switch (the four-pole switch not to be confused with a four-pole filter!), Fig. 21 of that paper, where only LP was done to 4<sup>th</sup> Order. I hope that our presentation, although clear, was not instigation for the confusion.

But to the present, we should mention that the cascade, even as correctly done in Fig. 9, is NOT and can not be identical to the Moog low-pass (except with  $g=0$ ). It may offer perfectly good filtering, but it is not Moog. Recall (and see Fig. 11 and Fig. 14b below) that for the original Moog, as one increases  $g$ , the four poles move (famously) outward from  $s = -1$  (that is, from a normalized cutoff frequency of 1, moving on the corners of a square. So first of all, they move to different pole radii, while as shown (Fig. 9), both S-V sections have the same pole radii. Additionally, with the Moog, two poles move toward the  $j\omega$ -axis, while the other two move away from the  $j\omega$ -axis. With considerable effort, you could make Fig. 9 be a Moog 4-pole, but of course, you would just simply use Fig. 1. Again, you might like Fig. 9 better than the Moog. [For example, making  $D_1 = 2\cos(22.5^\circ)$  and  $D_2=2\cos(67.5^\circ)$  you could get a 4<sup>th</sup>-Order Butterworth. Musically however, a corner-peaked response like the Moog is more familiar.]

## SEPARATE DAMPINGS WITH MOOG-LIKE FEEDBACK

Here we look at an obvious idea which I don't recall our doing, or seeing before. Here we simply modify Fig 7 by allowing the dampings to be different, as seen in Fig. 10.



**Four-Pole Lowpass Series with Feedback**

**Fig. 10**

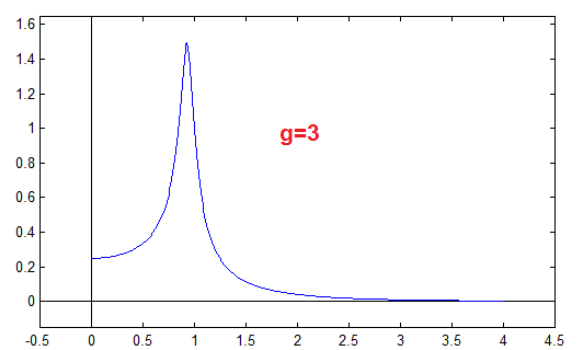



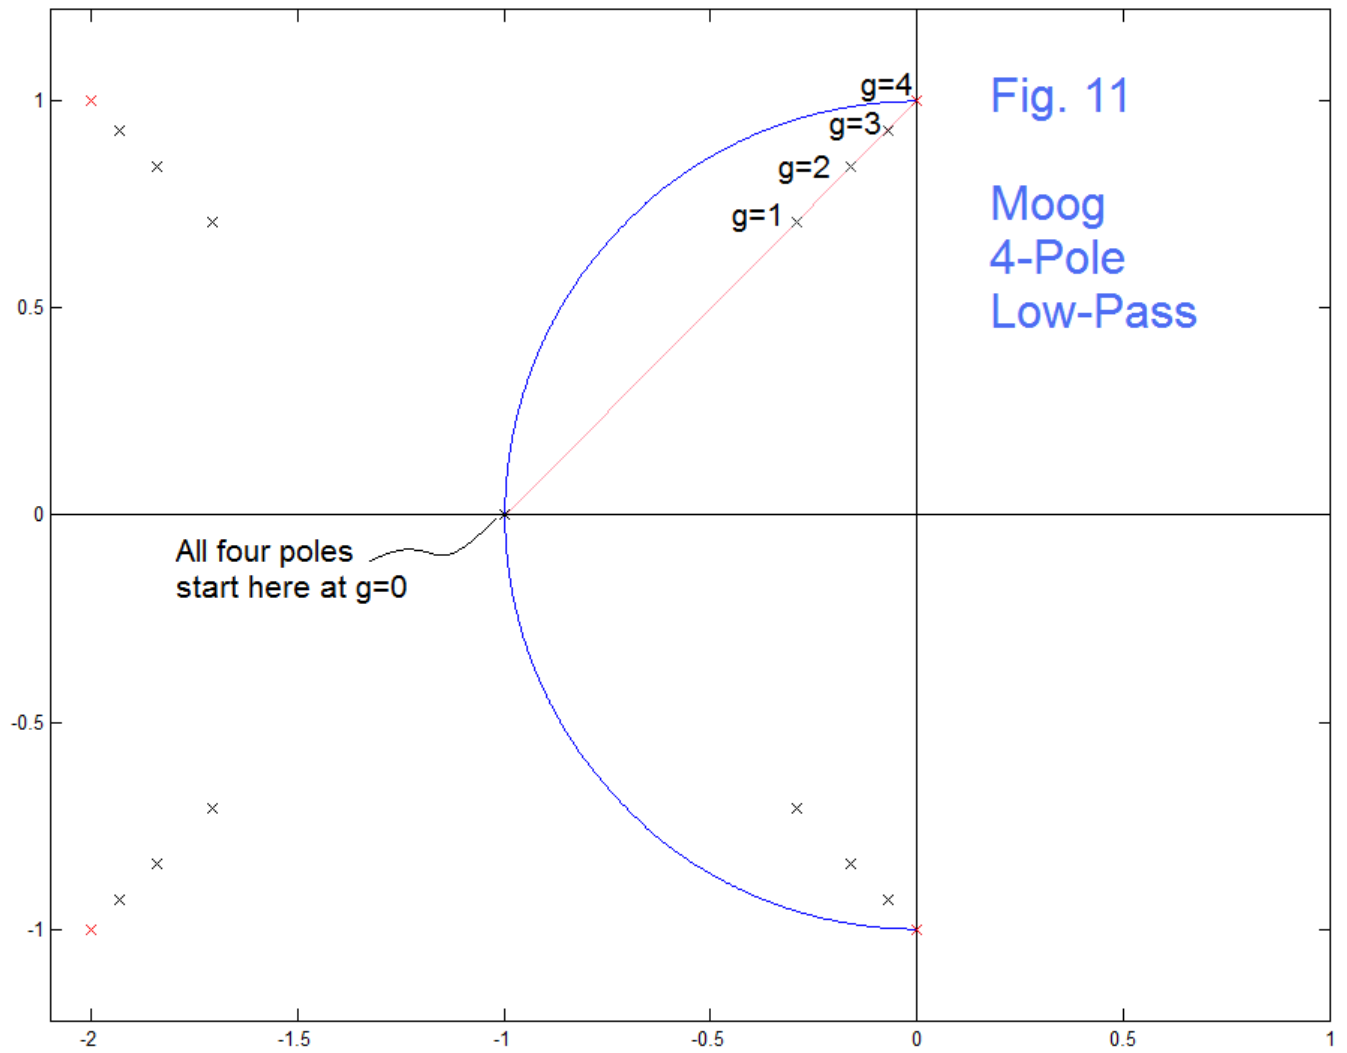
What happens? To cut to the chase, we get a result very similar to the Moog 4-Pole, with a pair of poles moving to the  $j\omega$ -axis, but for lower values of  $g$  (for dampings less than 2). This is perhaps testimony to the power of positive feedback and the general validity of Moog's notion of achieving corner-peaking. Neat. Reassuring.

Here the tool used is Matlab (simple code at end) and finding root loci is trivial, relative to 35 years ago! Here is what the test program does: For a choice of dampings  $D_1$  and  $D_2$ , it searches value of  $g$  starting at  $g=0$  until a pair of poles reaches the  $j\omega$ -axis. This value,  $g_{\max}$ , incidentally, is the product of  $D_1$  and  $D_2$ . (Need to prove this.) Then, for five values of  $g$ , of 0,  $g_{\max}/4$ ,  $g_{\max}/2$ ,  $3g_{\max}/4$ , and  $g_{\max}$ , chosen as representative, it plots the poles and the corresponding frequency responses. Fig. 11 shows the result for the original Moog 4-Pole: So the four poles start at  $s = -1$  (all poles start at a radius of 1) for  $g=0$  and two of the four end up at  $\pm j$  for  $g = 4$  (all poles for maximum  $g$  shown in red for stop). In addition, at the bottom of the plot are four frequency response plots, for four of the five choices of  $g$  ( $g_{\max}$  itself is not plotted since that would blow up). This is of course familiar to us from our previous studies.

Three additional sets of starting positions are considered here. Fig. 12 shows what happens if we start the poles not at  $s = -1$ , but at the 2<sup>nd</sup>-Order Butterworth positions  $s = -0.7071 \pm 0.7071j$ , with the  $\pm$  sign thus representing two poles, and each pole used twice, for a total of four. We see the poles initially on the unit circle, each double pole splitting, with two moving toward the  $j\omega$ -axis and two away – very similar to the Moog square array. Quite lovely to look at! In fact, the similarity of the frequency responses is also highly evident.

The array of Fig. 12 does not start out 4<sup>th</sup>-Order Butterworth, but a 4<sup>th</sup>-Order cascade of two 2<sup>nd</sup>-Order Butterworth sections. This point is better made by Fig. 13 where the original poles ( $g=0$ ) are 4<sup>th</sup>-Order Butterworth, on the unit circle, spaced 45 degrees apart. Here when we increase  $g$ , as before, two poles (the ones that were closest) move toward the  $j\omega$ -axis and two move away. Note the flatter response at  $g=0$  for the true 4<sup>th</sup>-Order Butterworth – a result we have often given in the past. However, the main finding from the frequency response plots at the bottom of Fig. 13 is the similarity to the previous results. That is, corner-peaking is powerful and a dominant performance.

The final example (which uses a slightly modified version of the Matlab program listed below to better plot a detailed locus) is shown in Fig. 14a (with the original Moog 4-Pole by the same program in Fig. 14b). Here we have put the original poles at  $s = -0.5$ ,  $s = -1$ ,  $s = -1.5$ , and  $s = -2.5$ , separated, but all real. Note how the poles merge and then split into conjugate pairs, and in the now familiar manner, two head for the  $j\omega$ -axis while the other two head away. This is a familiar behavior for a pole locus.



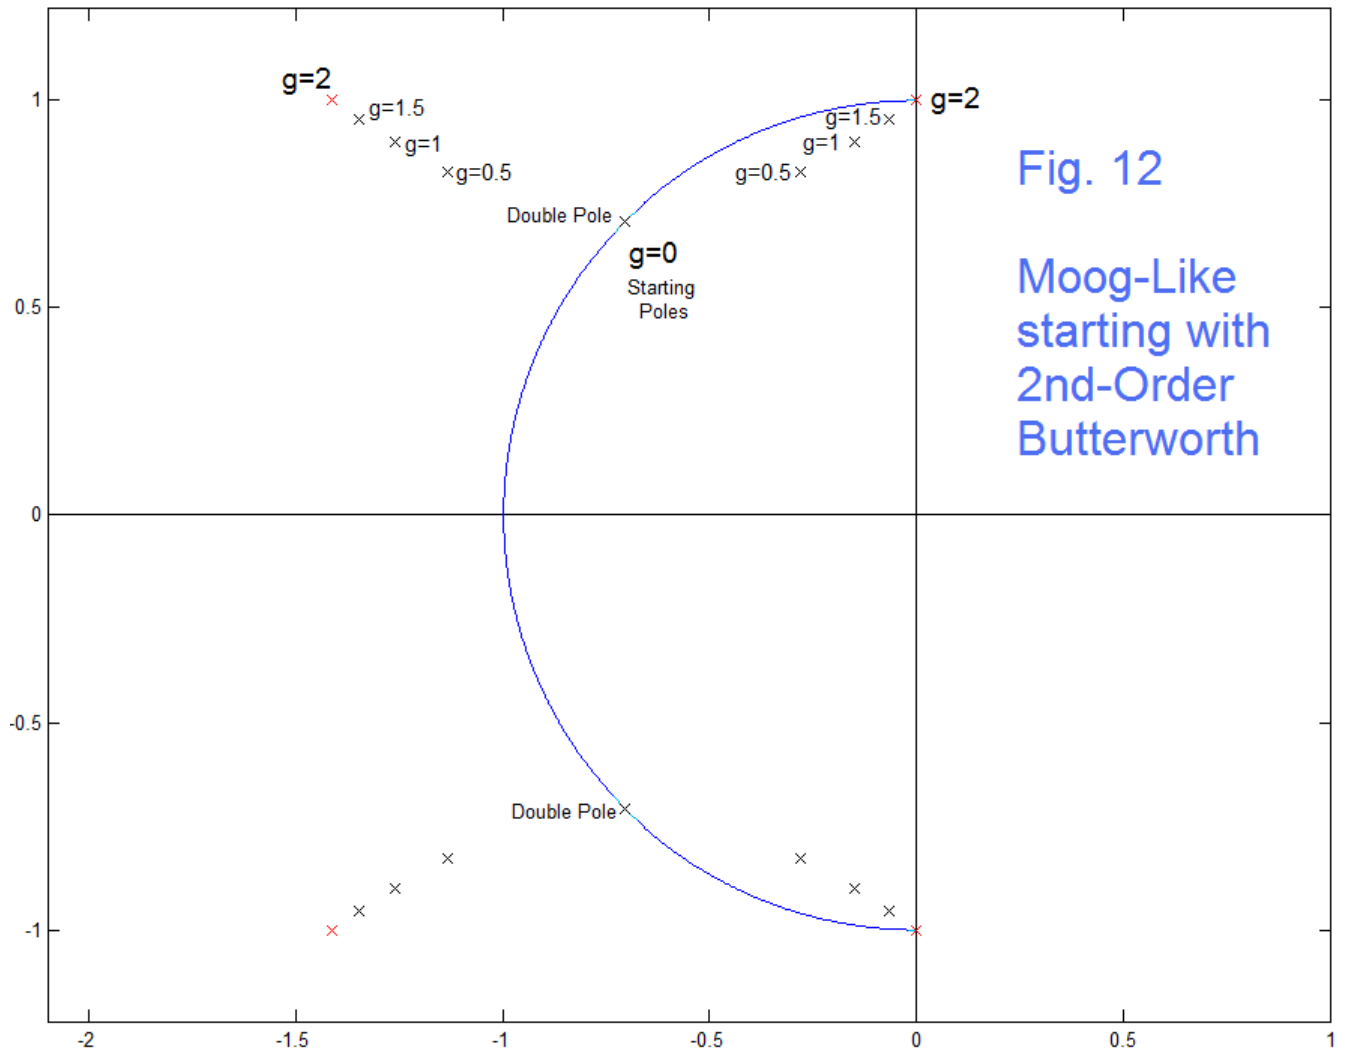






Fig. 12

Moog-Like  
starting with  
2nd-Order  
Butterworth





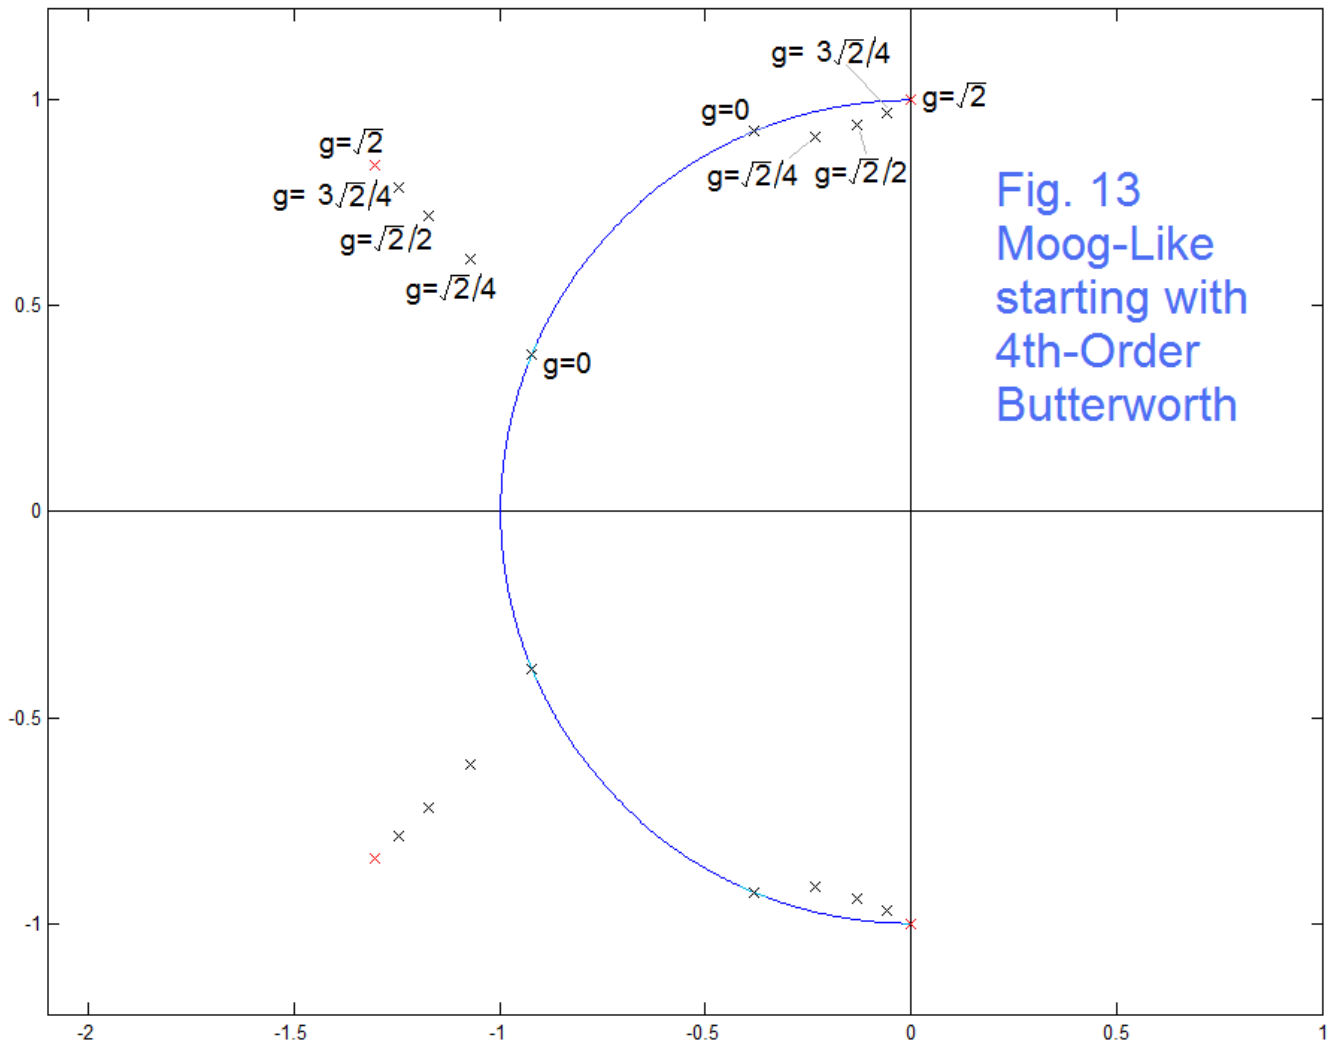



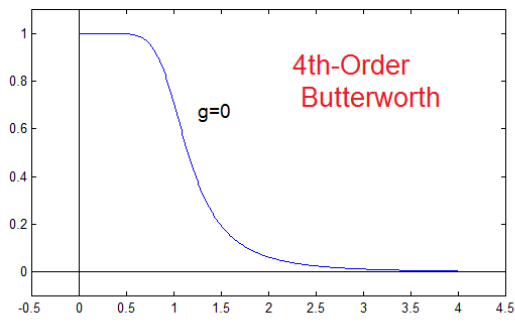

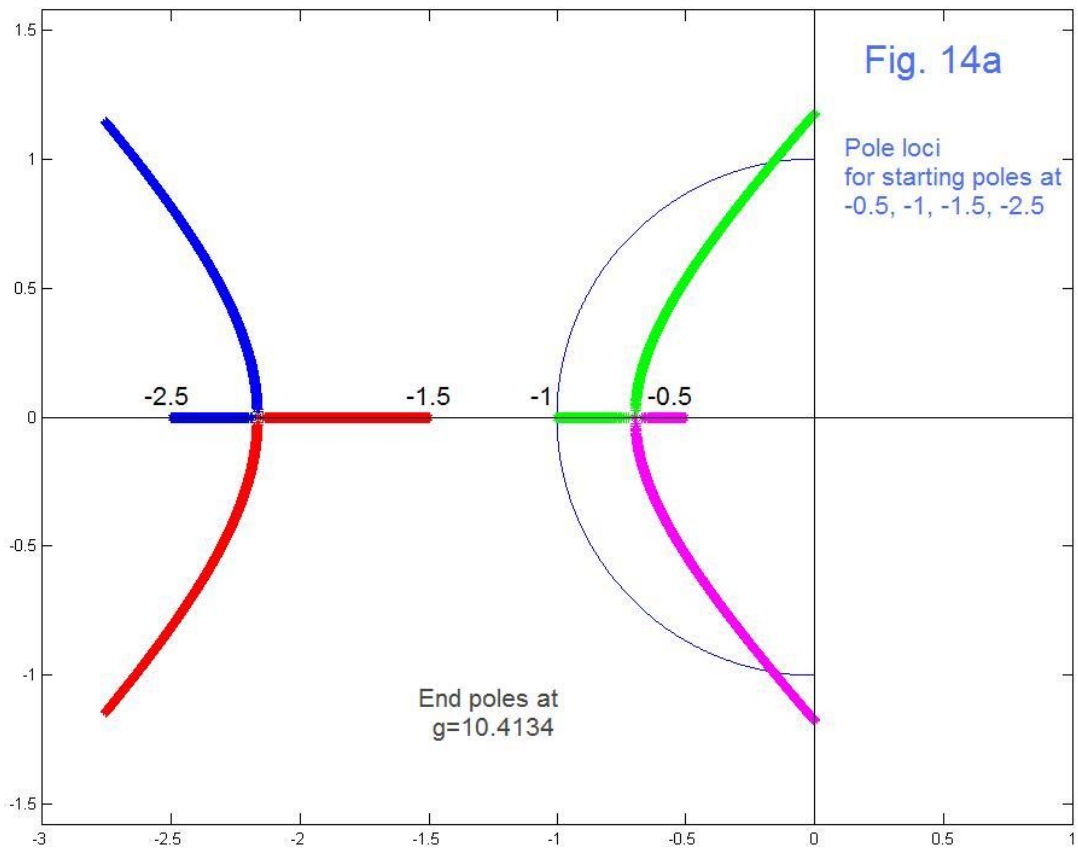


Fig. 13  
Moog-Like  
starting with  
4th-Order  
Butterworth





## CONCLUSIONS

Here in the spirit of a “revisit” we have looked at some familiar things, but have also looked at several interesting new things (and left some issues unanswered). One overall finding is: WE HAVE BEEN DOING THINGS PRETTY WELL ALL ALONG. With the state-variable, we have seen that the shunt lead compensation (the added capacitors on the attenuators) really are simple and effective. With the Moog 4-Pole, it was the feedback and corner peaking (Moog’s intuition) that is still the main story. Today, we have additional tools, but the existing practice is apparently something which we can still embrace.

## REFERENCES:

- [1] B. Hutchins, Musical Engineer’s Handbook, Electronotes (1975) Chapter 5d on VCF designs.
- [2] B. Hutchins, “The ENS-76 Home-Built Synthesizer System – Part 5”, VCF options., Electronotes, Vol. 8, No. 71, Nov. 1976 (Includes *ad hoc* lead compensation). See also the VCF options in the Electronotes Builder’s Guide and Preferred Circuits Collection.
- [3] B. Hutchins, “Additional Design Ideas for Voltage-Controlled Filters,” **Electronotes**, Vol. 10, No. 85, January 1978 (Location of 4 poles, feedforward, 4-function cascade).  
<http://electronotes.netfirms.com/EN85VCF.PDF>
- [4] B. Hutchins, “The Migration of Poles as a Function of Feedback in a Class of Voltage-Controlled Filters,” **Electronotes**, Vol. 10, No. 95, Nov. 1978. See also R. Bjorkman, “A Brief Note on Polygon Filters,” **Electronotes**, Vol. 11, No. 97, January 1979 available at:  
<http://electronotes.netfirms.com/EN97VCF.PDF>
- [5] B. Hutchins, “Integrated Musical Electronics, Part 3, Better Use of VCF Chips”, Electronotes, Vol. 14, No. 143, Nov. 1982 (two SV to Moog 4-pole)
- [6] B. Hutchins, “Analog Signal Processing, Chapter 7: Passive and Active Sensitivity”, **Electronotes**, Vol. 20, No. 195, July 2000  
<http://electronotes.netfirms.com/EN195.pdf>
- [7] B. Hutchins, “Revisiting: Compensation of Op-Amp Circuits for the Effects of the Op-Amp Poles”. **Electronotes**, Vol. 22, No. 214, Feb 2013  
<http://electronotes.netfirms.com/EN214.pdf>

[8] B. Hutchins, "Additional Notes on Band-Pass Filters, *Electronotes*, Vol. 10, No. 91. July 1978, see pp 18-20

[9] Note on "non-linear": We have mentioned over the years, most recently at:

<http://electronotes.netfirms.com/ENWN4.pdf>

that linearity has several meanings.

## Appendix – Matlab Program

```
% nm1
clear
clf
D1=2*cos(22.5*pi/180)
D2=2*cos(67.5*pi/180)

%D1=sqrt(2)
%D2=sqrt(2)

%D1=2
%D2=2

% find gmax
g=0;
rp=-100;
while rp<0
    poles=roots([1 D1+D2 (D1*D2+2) D1+D2 (1+g)]);
    p1=poles(1);
    p2=poles(2);
    p3=poles(3);
    p4=poles(4);
    rp=max([real(p1) real(p2) real(p3) real(p4)]);
    g=g+0.00001;
end
gmax=g
```

```

% now find poles at multiples of gmax/4
k=0;
for g = 0:gmax/4:gmax
    k=k+1
    g
    poles=roots([1 D1+D2 (D1*D2+2) D1+D2 (1+g)]);
    p1k(k)=poles(1);
    p2k(k)=poles(2);
    p3k(k)=poles(3);
    p4k(k)=poles(4);
end

% plot the poles
figure(1)
a=0:pi/1000:pi;
y=cos(a);
x=-sin(a);

plot([0 0],[-10 10],'k')
hold on
plot([-10 10],[0 0],'k')
plot(x,y,'b')
%
plot(p1k,'xk')
plot(p2k,'xk')
plot(p3k,'xk')
plot(p4k,'xk')
plot(p1k(5),'xr')
plot(p2k(5),'xr')
plot(p3k(5),'xr')
plot(p4k(5),'xr')
%
axis([-2.1 1 -2 +2])
axis('equal')
hold off
figure(1)

% now plot the frequency response (but not for gmax which becomes
infinite)
f=0:.01:4;

```

```

figure(2)
%
subplot(221)
g=0;
M=freqs(1,[1 D1+D2 (D1*D2+2) D1+D2 (1+g)],f);
plot([-1 5],[0 0],'k')
hold on
plot([0 0],[-10 100],'k')
plot(f,abs(M))
Mmax=max(abs(M));
axis([-0.5 4.5 -0.1*Mmax 1.1*Mmax])
hold off
%
subplot(222)
g=gmax/4;
M=freqs(1,[1 D1+D2 (D1*D2+2) D1+D2 (1+g)],f);
plot([-1 5],[0 0],'k')
hold on
plot([0 0],[-10 100],'k')
plot(f,abs(M))
Mmax=max(abs(M));
axis([-0.5 4.5 -0.1*Mmax 1.1*Mmax])
hold off
%
subplot(223)
g=gmax/2;
M=freqs(1,[1 D1+D2 (D1*D2+2) D1+D2 (1+g)],f);
plot([-1 5],[0 0],'k')
hold on
plot([0 0],[-10 100],'k')
plot(f,abs(M))
Mmax=max(abs(M));
axis([-0.5 4.5 -0.1*Mmax 1.1*Mmax])
hold off
%

```

```
subplot(224)
g=3*gmax/4;
M=freqs(1,[1 D1+D2 (D1*D2+2) D1+D2 (1+g)],f);
plot([-1 5],[0 0],'k')
hold on
plot([0 0],[-10 100],'k')
plot(f,abs(M))
Mmax=max(abs(M));
axis([-0.5 4.5 -0.1*Mmax 1.1*Mmax])
hold off
%
figure(2)
```