The Pulse-width Modulation System

a very detailed study with design recommendations

- James L. Tonne W4ENE

In the amateur radio amplitude-modulation arena, the use of pulse-width modulation (PWM) seems to be an increasingly popular modulation scheme. The efficiency of the system can be very high and the quality of the transmitted signal can be very high as well. This document was written in an effort to explain several aspects of that modulation system along with some suggested design methods. Various kinds of screen captures from *Elsie, ACTLPF* and *LTspice* are used to enhance the paper. The schematic captures allow the reader to pursue the various circuit tests and examinations for themselves, perhaps modifying the constants as needed to accurately match his own set of circumstances.

Your comments are welcome. You may drop a line to the author at
Sales@TonneSoftware.com

Last text update: 30 August 2021
Link to active lowpass designer updated 16 September 2021

The Basics

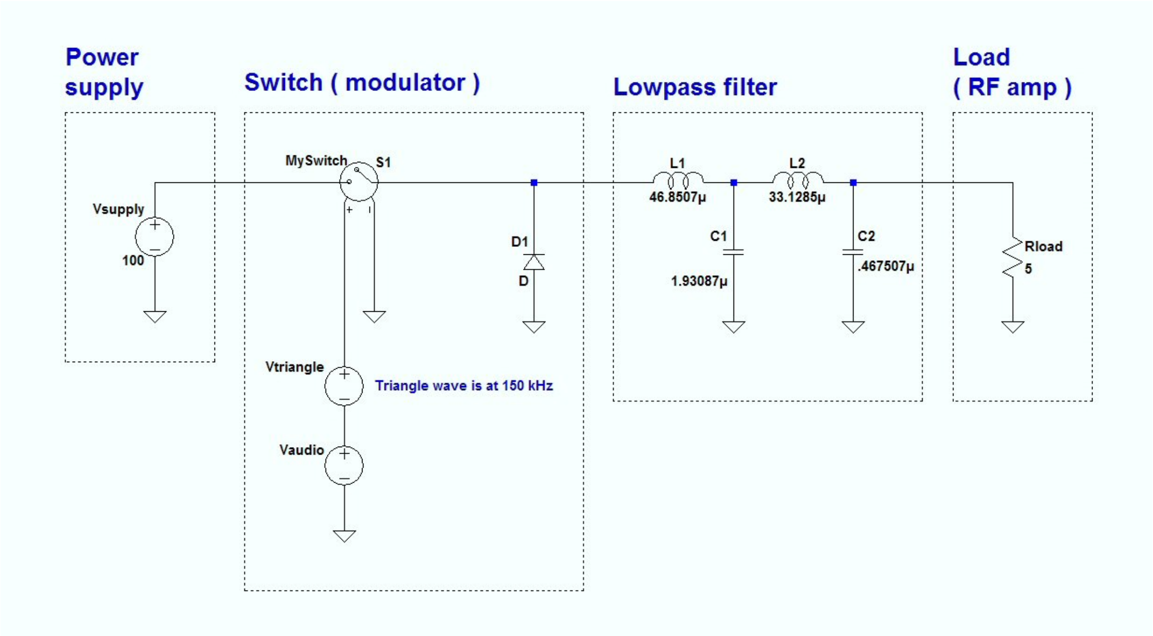

Fig. 1 - The essentials of a pulse-width modulator

This is the system schematic as rendered in *LTspice*. At the left we have the power supply. Its DC output is shown here as being set to 100 volts, a voltage commonly seen in modern solid-state systems. It must be at least twice the voltage that is applied to the RF amplifier without modulation. That supply is connected to a series switch, which block is labeled here as 'Switch.' The diode labeled 'D' is part of the switching system. (A second version of this switching scheme uses a second switch in lieu of the diode.) The switch output is applied to the next block, labeled here as 'Lowpass filter.' That filter passes DC (the unmodulated voltage) and audio (the modulation). But this filter's primary purpose is to reject the serious amount of unwanted trash that results from the switching operation. We will see that the filter essentially determines the performance of the system. It is **not** the transmitted-signal bandwidth controller. That function must be performed earlier in the audio signal chain.

With no modulation applied to the modulator, the output from the switch appears as in Fig. 2.

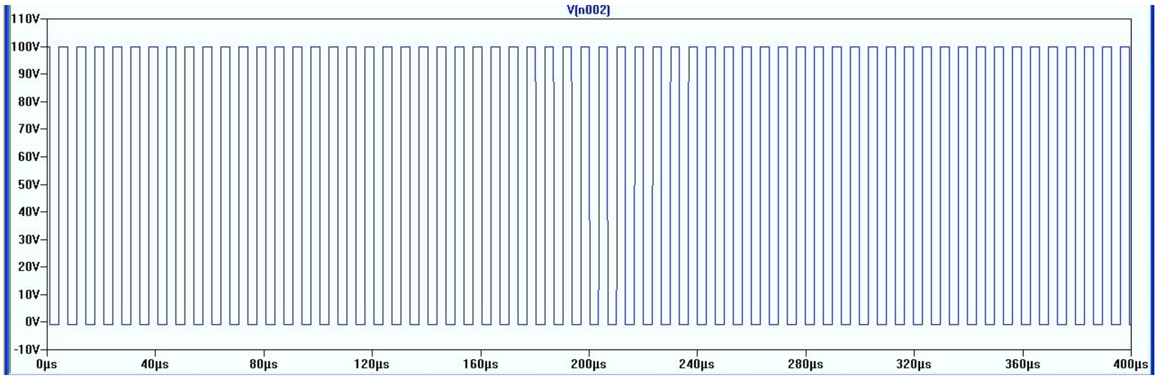

Fig. 2 - Idling (no applied modulation) output from the switcher

That waveform as seen here has a 50% duty cycle. The output from the following lowpass filter will then be half of the maximum. With 100 volts from the power supply, the output from the following filter will be 50 volts.

Some designers adjust the modulator so that the idling (DC) output is somewhat less than half of the power supply voltage. That is an option that may be exercised by the designer.

Now we add a modulating signal as shown in Fig. 3.

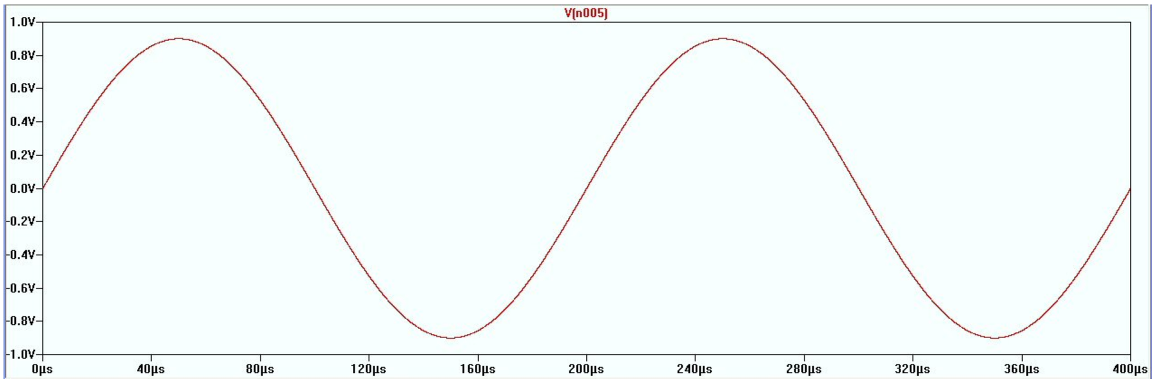

Fig. 3 - The modulating signal

For these illustrations we are using a 1000 Hz sinusoid.

To that signal we add a second signal. It is a triangle wave with a repetition rate typically in the 65 kHz to 160 kHz region.

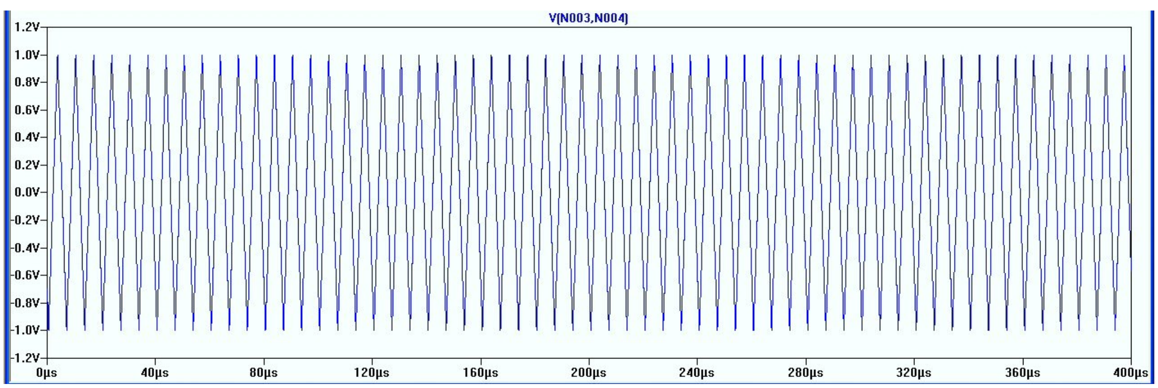

Fig. 4 - The triangle wave

The modulating wave and the triangle wave are summed. The result is shown in Fig. 5 along with a line representing the comparison voltage of zero volts.

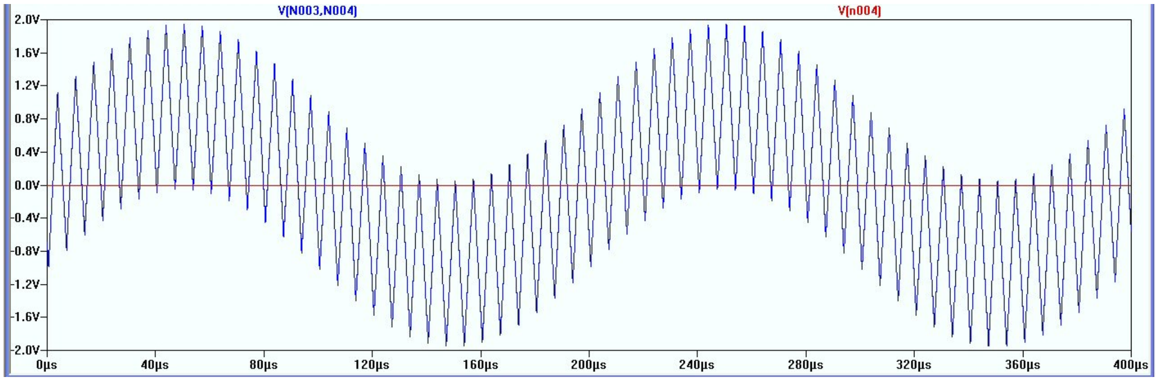

Fig. 5 - The composite waveform: modulating signal plus triangle wave

The sum of those two signals, the composite waveform, is next applied to a comparator. When the instantaneous amplitude of that composite waveform is greater than zero (above the zero-voltage line seen in Fig. 5), the comparator turns on the series switch. That series switch then delivers the full output of the power supply to the following filter. When the instantaneous amplitude of the composite waveform is less than zero (below the zero-voltage line in Fig. 5), the comparator turns off the series switch. At that instant the filter delivers current back to the shunt diode. That diode is called a 'backswing' or 'flyback' diode. In a second version of this switching system a pulldown switch is set to conduct when the series switch is turned off. In this second system only one switch conducts at a time, either the series switch or else the following shunt switch. (If they both conduct at the same time we would have a serious problem !)

The output of the switch and diode (or the pair of switches) now appears as in Fig. 6.

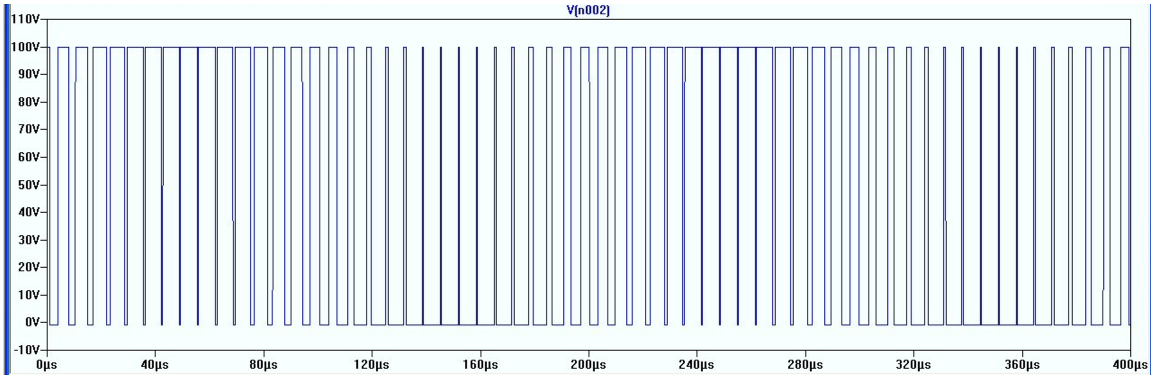

Fig. 6 - Switcher output with modulation

The item to observe here is that the duty cycle of the switch output, the percentage of 'on' time, is changing according to the instantaneous value of the modulating waveform.

This can be examined another way, by showing in one graphic both the switcher output waveform and the modulating waveform. See Fig. 7.

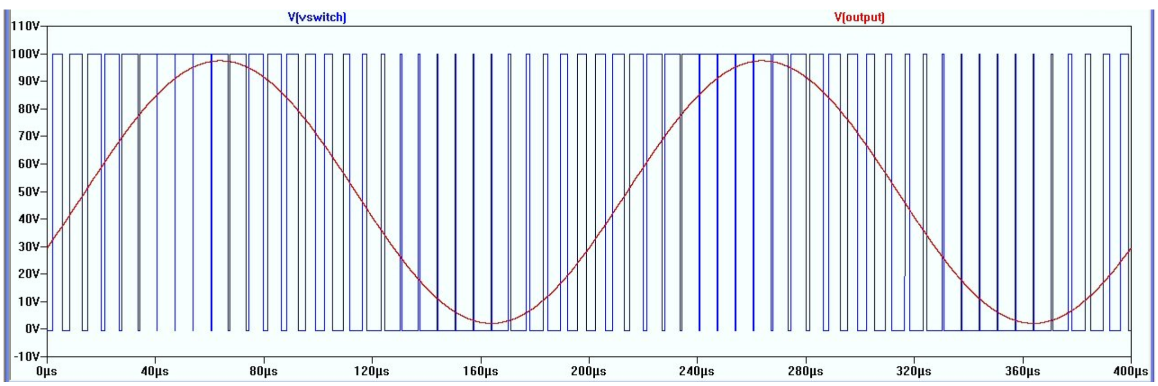

Fig. 7 - The modulating waveform and the switcher output

As the instantaneous value of the modulating waveform (that sinusoid) increases so does the duty cycle of the switcher; the pulses become wider. As it goes down likewise the duty cycle of the switcher goes down; the pulses become narrower.

Two types of switchers

The two versions of this switching modulator are next in our examination. The first, by far the most common, is shown in Fig. 8.

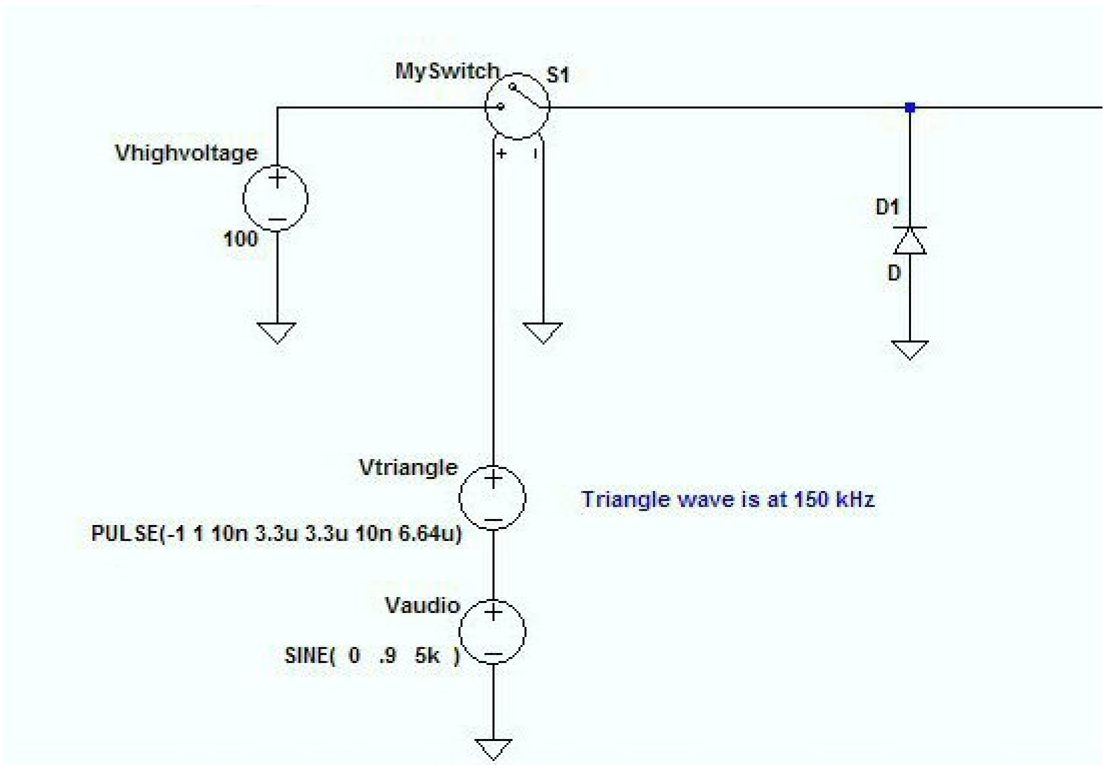

Fig. 8 - Switcher with a series switch and a shunt diode

A second version of the switcher is shown in Fig. 9.

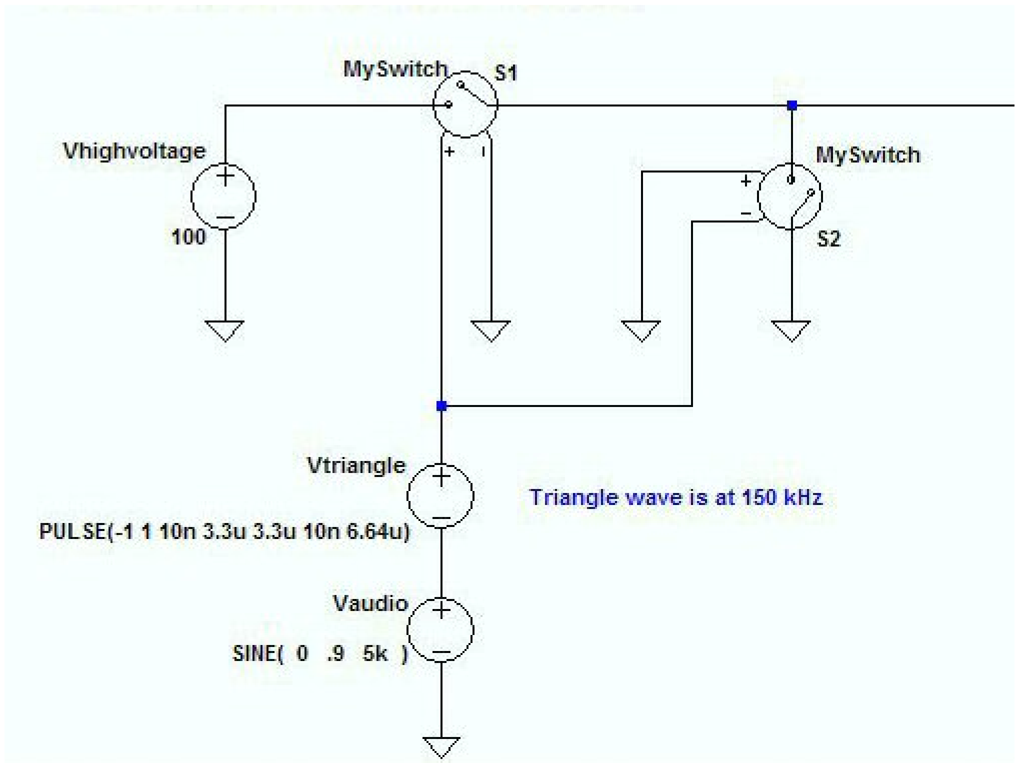

Fig. 9 - Switcher with two switches

This version of the switcher uses a series switch followed by a shunt switch. One switch at a time is "on," first the series switch and then the shunt switch.

Now we look at the two versions along with their respective outputs. First see Fig. 10 which shows the most common configuration. We see the schematic and the filter output.

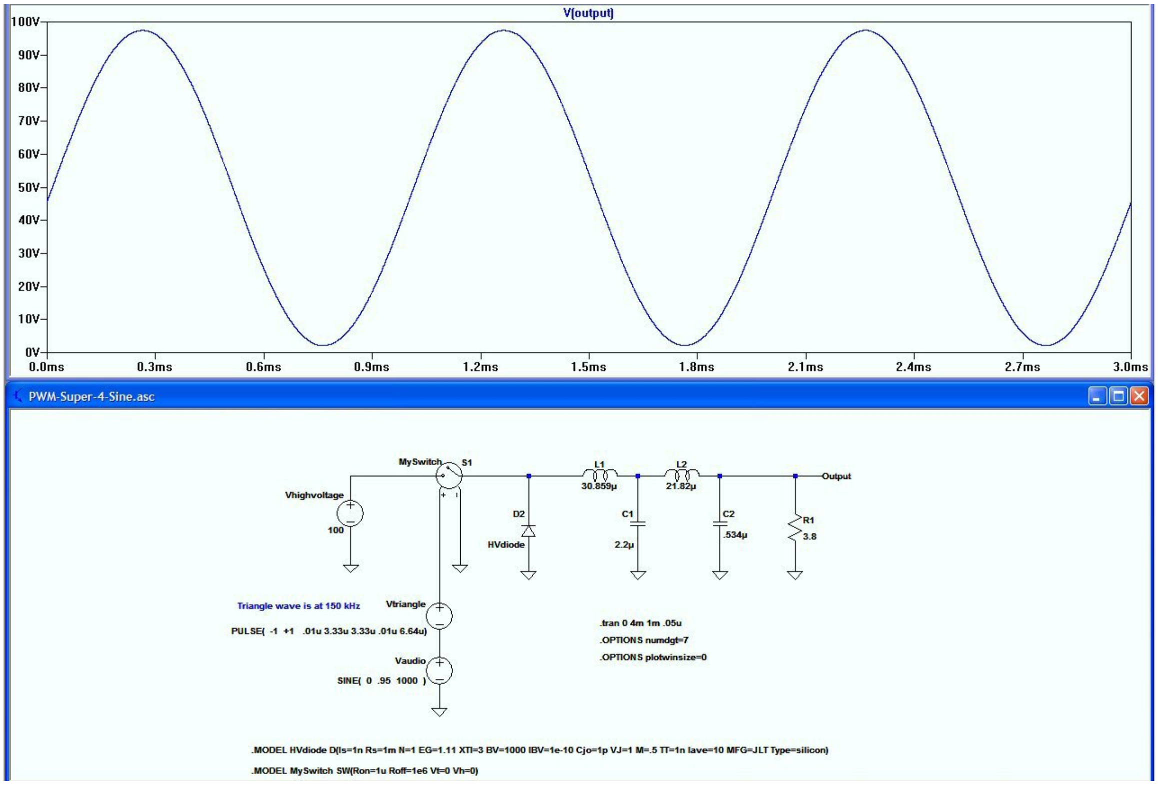

Fig. 10 - The most common system configuration

That graphic shows the first circuit and also its filter output.

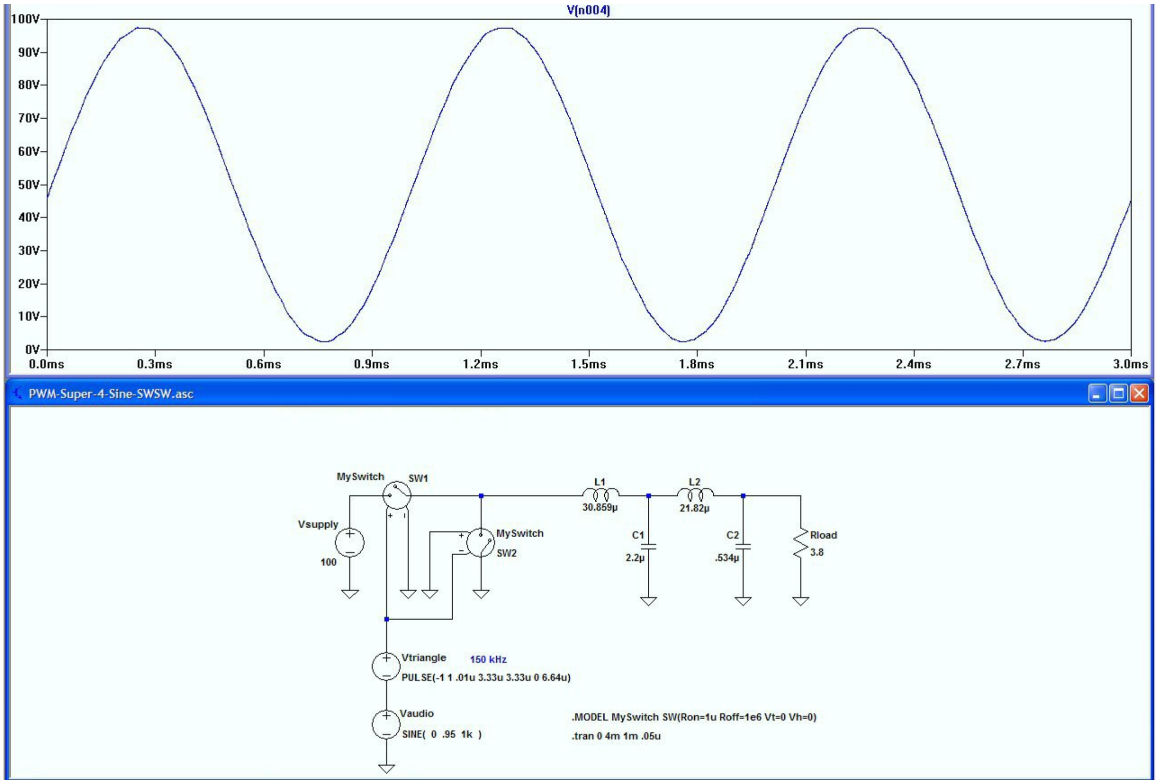

Fig. 11 - The second configuration

This graphic shows the second configuration and its filter output. The performance of this circuit is similar to the circuit shown in Fig. 10 although it seems to be better behaved on complex modulating signals at high modulation levels. This second topology is also called a "half-bridge", sometimes "Class D." The two switches are called "high side driver" (the series switch) and "low side driver" (the shunt switch).

Kindly notice that the 'comparator' is not shown in these schematics as an individual item. Rather it is actually an internal part of each switch.

The lowpass filter

After the switcher, regardless of its topology, we must have a lowpass filter. This filter must pass the DC and audio (modulation) portion of the switcher output and it must severely attenuate the trash that is also present in the switcher output.

One item to choose is how many parts it has. This is called the filter "order" by network designers. A second item that is involved is the topology (how those parts are connected). As will be seen, we must use a series-inductor-input topology.

A third item that is involved is to decide on the degree of attenuation needed at the chosen switching frequency (and its harmonics).

The very first thing that we must do is to examine the spectrum of the switcher output. Fig. 12 shows the output of the switcher in the frequency domain, as might be seen by a spectrum analyzer.

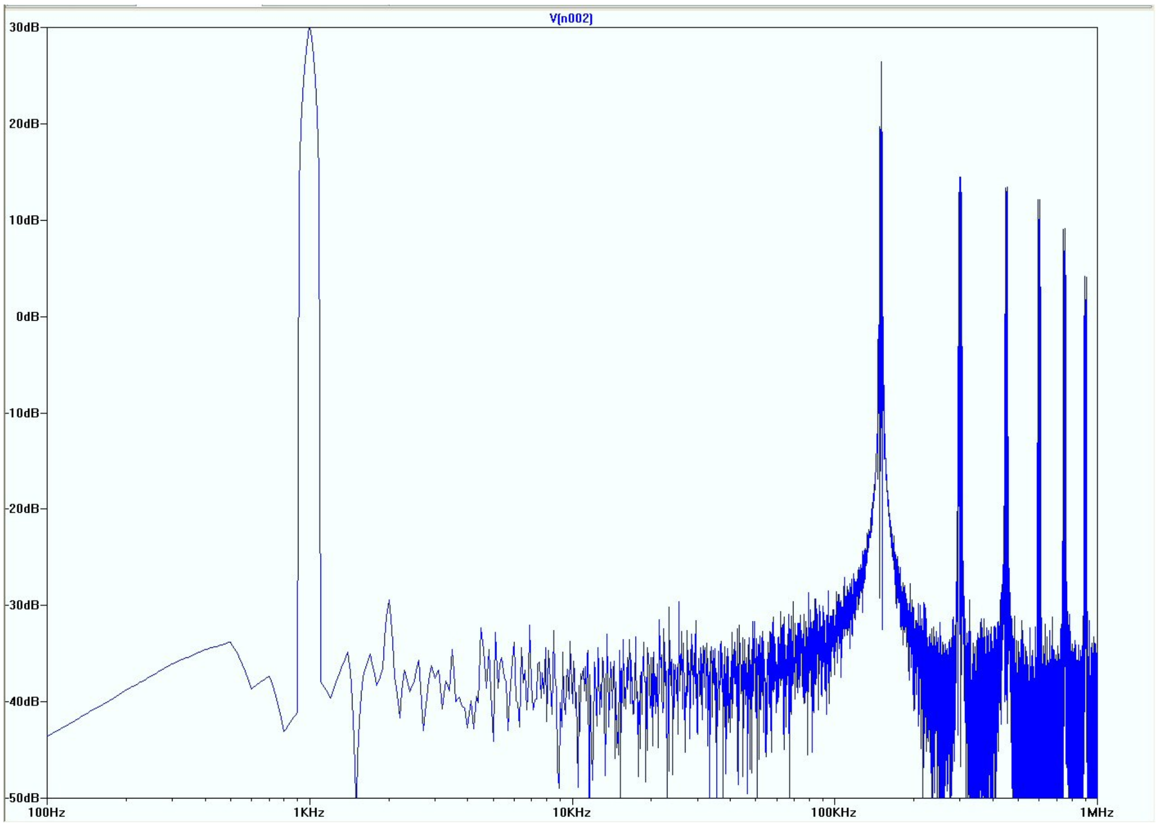

Fig. 12 - Switcher output spectrum

The signal at 'DC' (which sets the carrier level) does not appear in this display. The modulation can be seen at 1 kHz; we are using a 1 kHz sinusoid for modulation in this test. The major components that have to be *rejected* are those of the switcher frequency, set here at 150 kHz. And there are components at harmonics of the switching frequency which also must be rejected. The levels of those harmonics are also significant. Interesting is that the switching frequency and each of its harmonics have what can be called sidebands.

The spectrum in Fig. 12 will appear, after satisfactory lowpass filtering, as in Fig. 13.

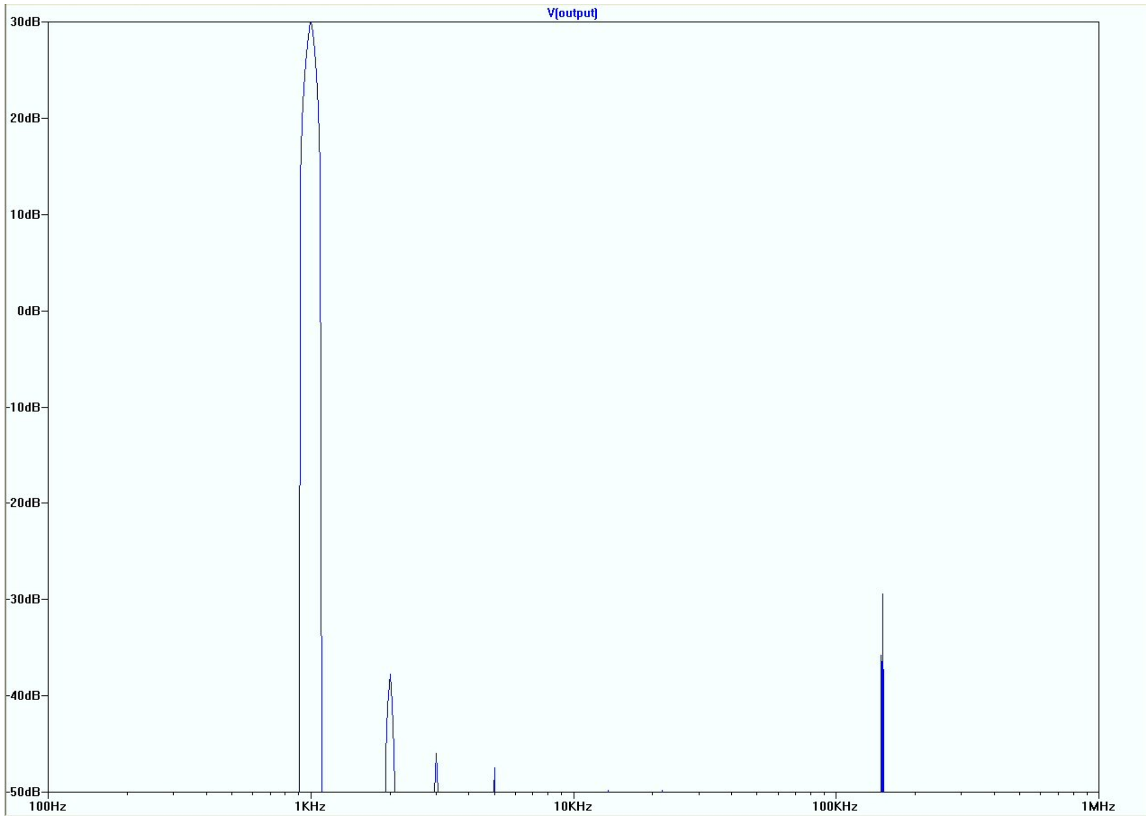

Fig. 13 - Spectrum of signal delivered by the output lowpass filter

That graphic shows that the spectrum of the filter output is satisfactory. We see that the 'trash' has been adequately suppressed. This signal may now be used to supply power to the RF amplifier.

Let's now look at an output lowpass filter design. Its schematic and magnitude response are shown in Fig. 14.

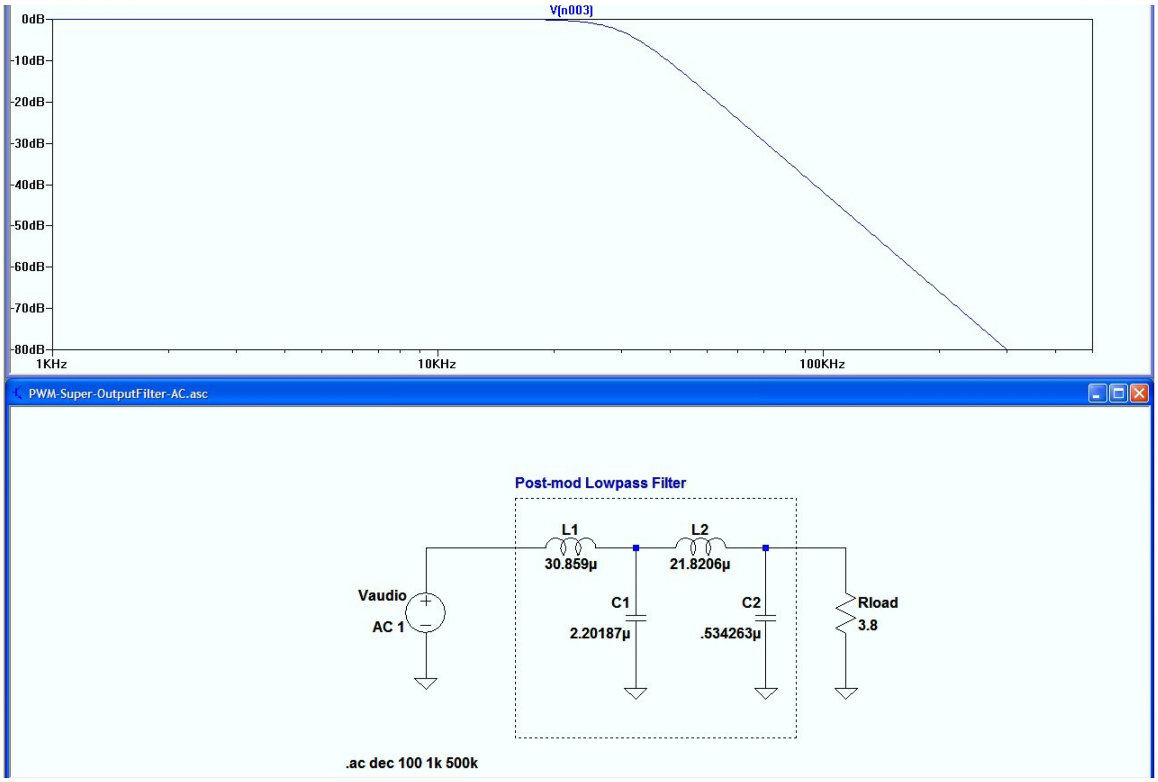

Fig. 14 - A proposed lowpass design and its magnitude response

The bandwidth of this of this post-switch lowpass has been adjusted so that its response is down 60 dB at 150 kHz. As a result its inband response is flat out to about 25 kHz, where it is down 3 dB. It then rolls off at 24 dB per octave. (It rolls off at 6 dB per octave times the order, 4 in this case, yielding a rate of 24 dB per octave.)

The switching waveform trash is about 3 dB or so below the audio level so a filter attenuation of 60 dB will result in a trash level of about 63 or so dB below the modulating waveform, in each sideband. In an AM system, each sideband is down 6 dB from carrier. If we suppress the switching trash 60 dB in the filter, the trash in each sideband will then be down 69 dB from carrier. (60 + 3 + 6) In such a case the trash in a 1000 watt carrier will have a level of 125 microwatts, quite insignificant. How to set the filter bandwidth, and other characteristics, will be discussed in detail later in this document in the filter design section.

Question from *ye scribe*: are your transmitter carrier harmonics also suppressed to that extent ?   ☺

Transient response

It is common in communication applications to use some degree of speech waveform clipping in the low-level stages of the audio chain. The post-switch lowpass has to pass clipped waveforms without introducing overshoots. If the filter we are designing has overshoots ('ringing') in its transient response then the clipper will have lost control of the modulation level. So now we examine what happens when a square wave (not used for modulation; just used for demonstrating overshoot) is applied to our system.

Fig. 15 illustrates a system wherein the modulating waveform is a 1000 Hz squarewave. This is a convenient waveform that can be easily generated and then used for testing purposes. It is a very severe test signal, challenging the modulator clearly more than a sinusoid that has been clipped perhaps 10 dB.

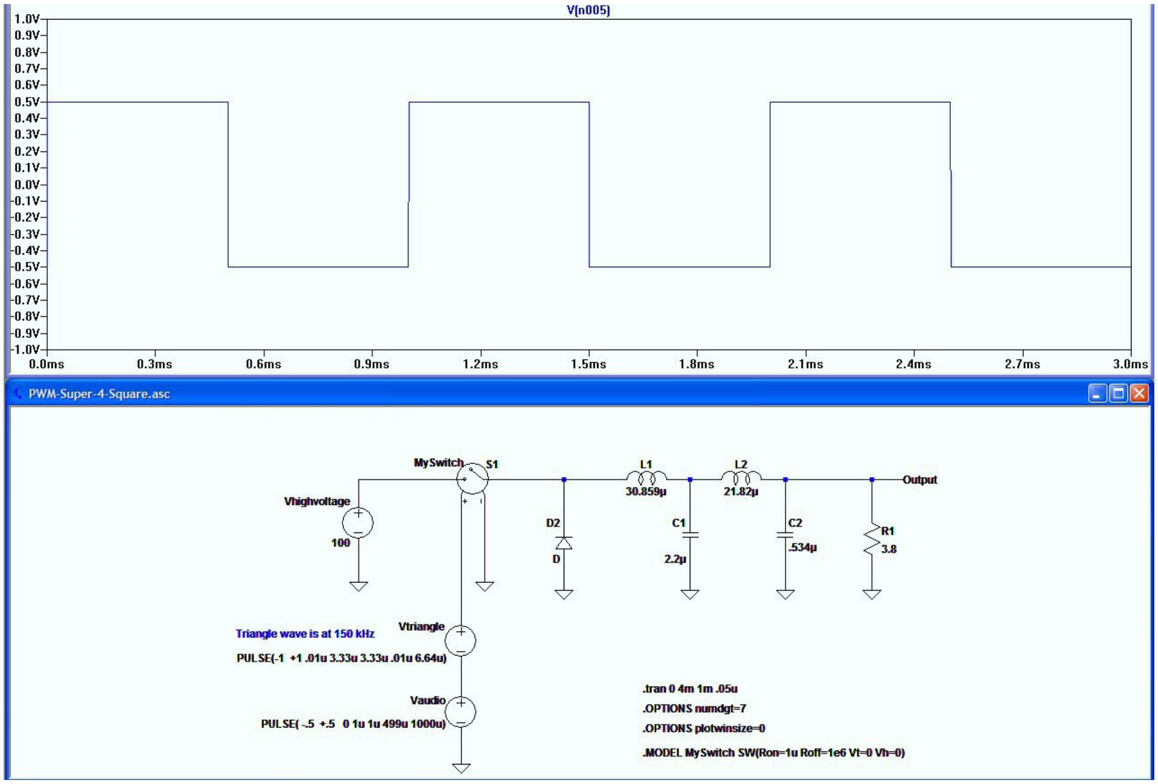

Fig. 15 - Showing a 1000 Hz squarewave used for testing

We will now use that squarewave as a modulating waveform. The output of the system when a recommended lowpass filter design is used is shown in Fig. 16.

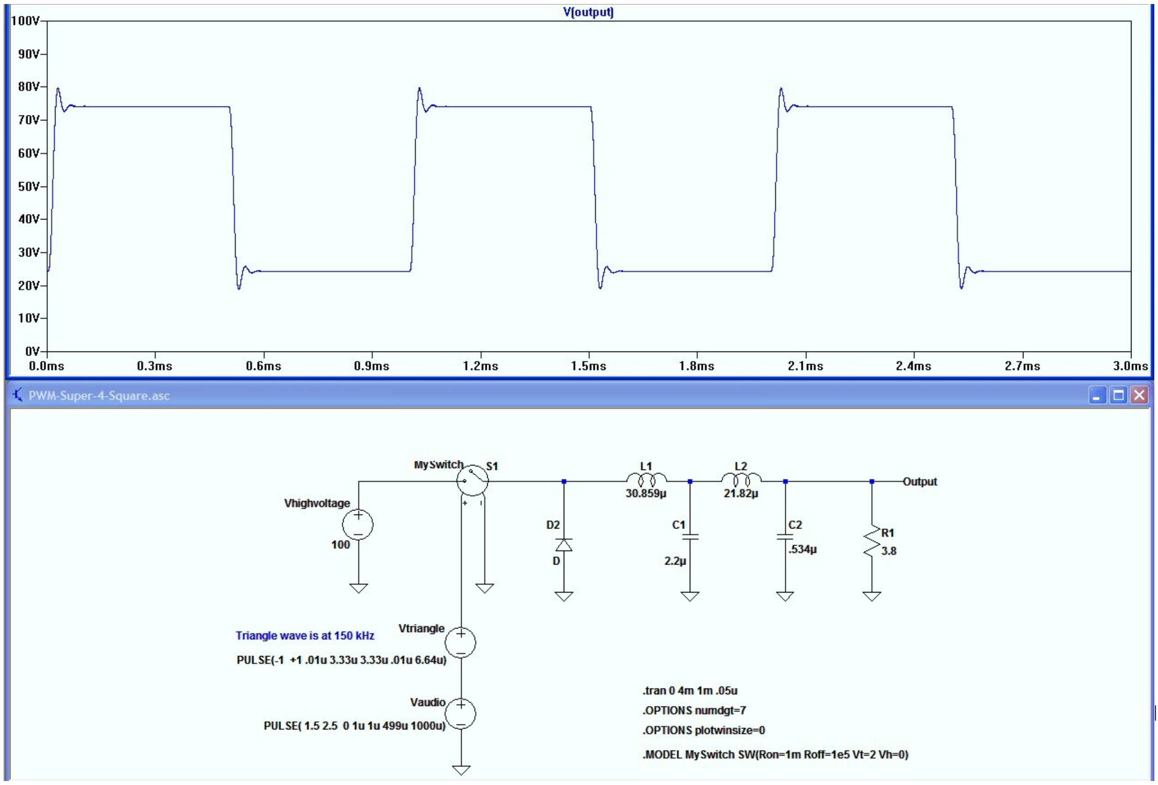

Fig. 16 - Transient response of complete modulator

That graphic shows the system transient response including the switcher and the output lowpass filter. This recommended filter design has a small amount of overshoot on the applied squarewave.

Now let's look at just the output lowpass filter by itself. See Fig. 17.

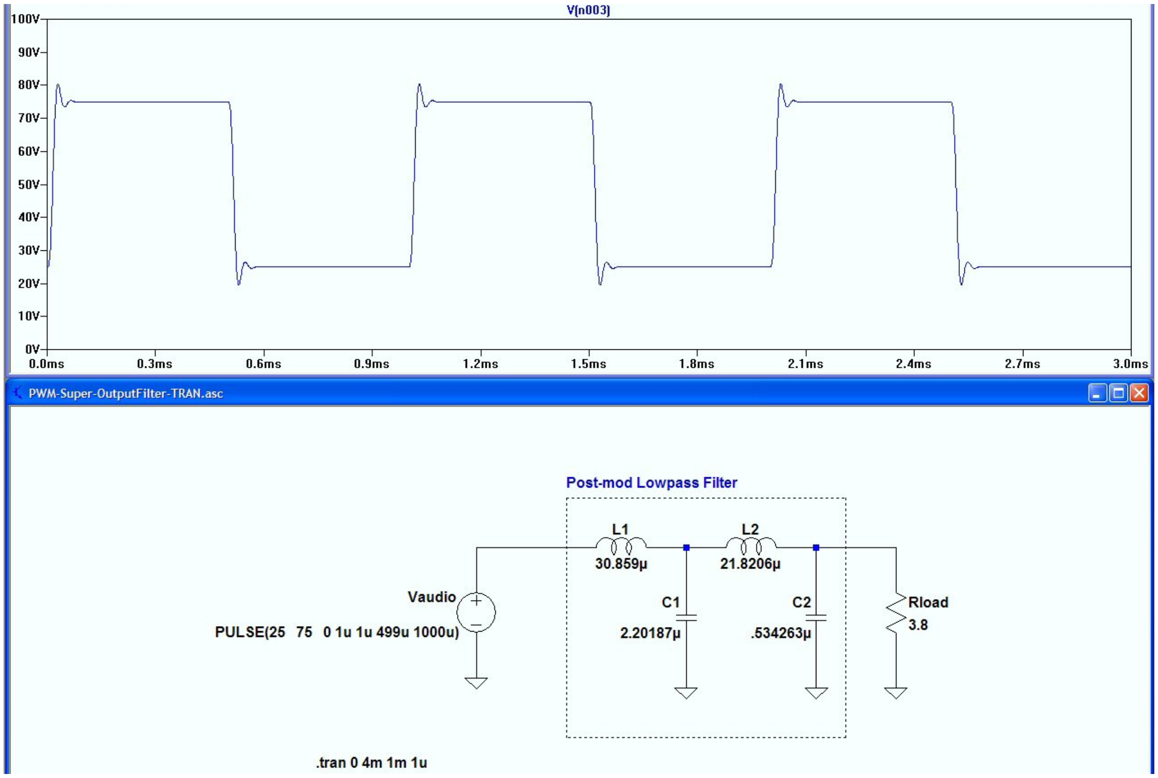

Fig. 17 - Illustrating the transient response of the output lowpass by itself

Notice that the waveshape of the signal delivered from the filter itself is identical to the waveshape from the complete system. It follows that the filter itself determines the modulator performance regarding both frequency (magnitude) response and also transient (squarewave) response. Filter design software can be used to predict system behavior.

But many of the commonly-published lowpass filter designs have a common problem: they have overshoot on applied squarewaves. An example is shown in Fig. 18.

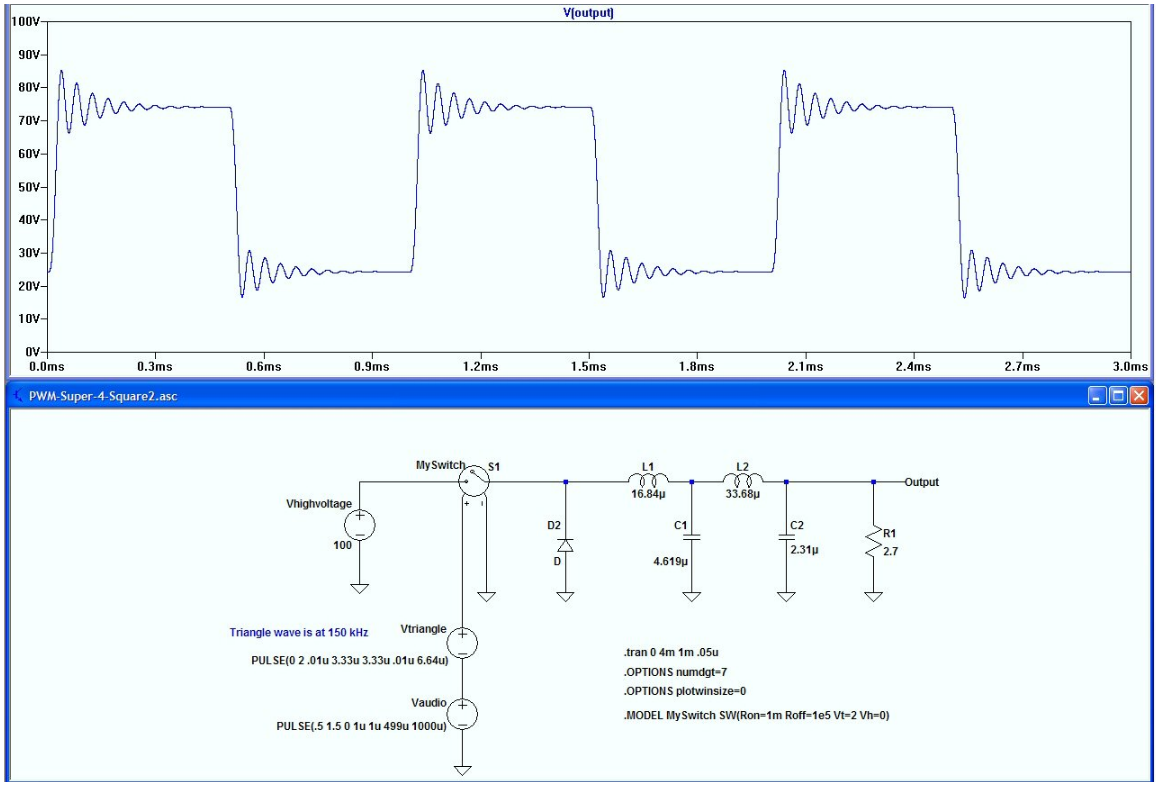

Fig. 18 - Usual design techniques yield this system transient response

That waveform illustrates the signal as delivered from the system. Now let's look at the filter alone. See Fig. 19.

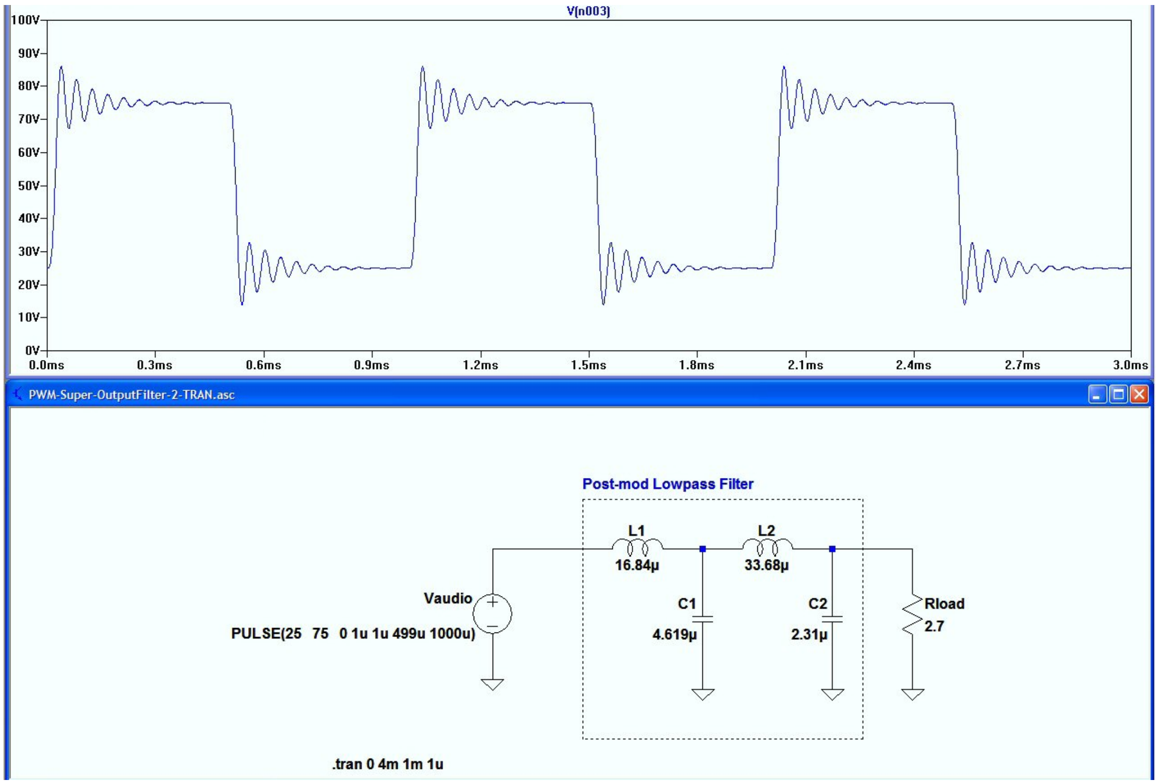

Fig. 19 - Usual design techniques yield this filter transient response

Again we see that the transient response of the system is identical to that of the filter alone. The performance of the filter in a system can be forecast by simply looking at the filter itself. Using a switch instead of a backswing diode (that second modulator scheme) accomplishes nothing to alleviate the problem. See Fig. 20.

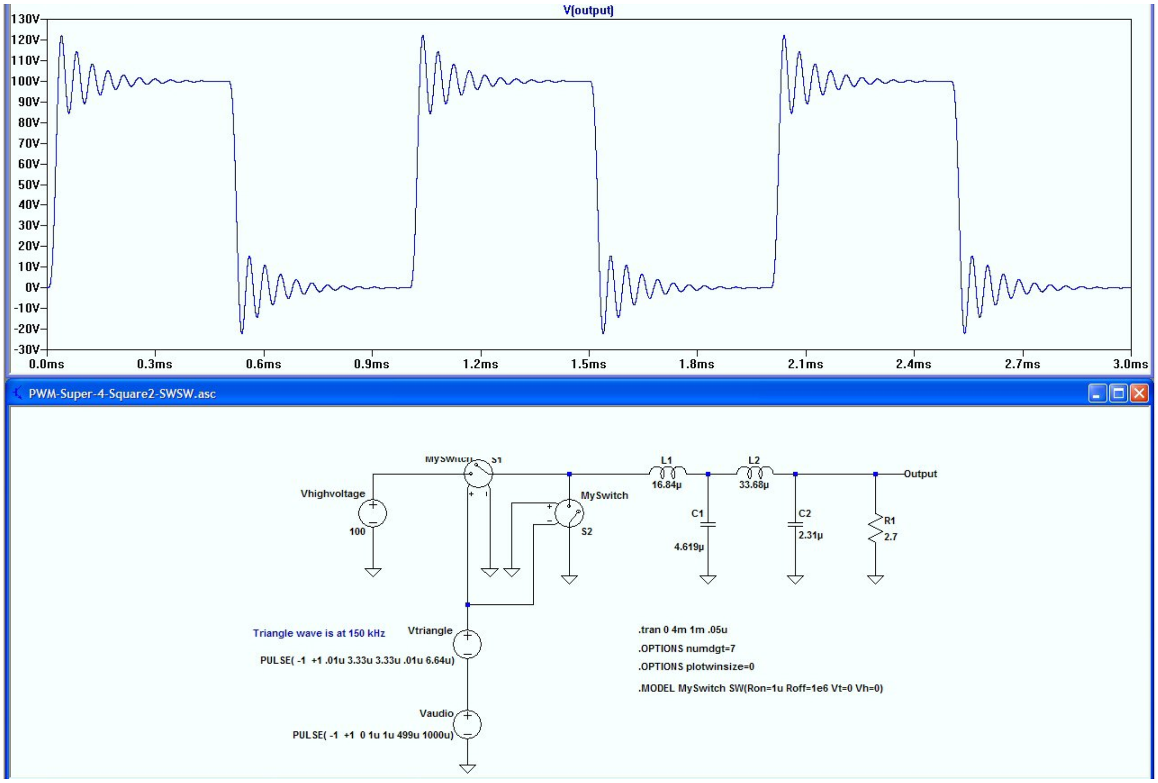

Fig. 20 - The double-switch system has the same problem with transients

That graphic shows the double switch modulator has the same response to a squarewave modulating waveform as does the type of modulator using a backswing diode.

This behavior seems to be due to a common design situation: the Butterworth family was chosen for the filter design (good) but the usual doubly-terminated design was then driven from a near-zero impedance source - the switch (not good). This situation is in part due to the lack of singly-terminated filter design software. A lack of education in these matters may be involved as well.

Polishing using an active lowpass input filter

We have seen that whether the post-switch lowpass is designed by the usual doubly-terminated method or by the highly-recommended singly-terminated method, there can be overshoots resulting from an applied modulating waveform. The doubly-terminated design method - the usual method - has far more serious overshoots when driven from a zero-impedance source, precluding accurate modulation control on transients. Even the recommended design method yields a filter that has small overshoots (*very* small). In each design method those overshoots might be minimized by adding an active lowpass filter to the modulator input. This simple filter may reduce the degree of overshoots caused by our main output lowpass filter. It functions by reducing the magnitude of modulating components near the output lowpass cutoff frequency. If the output lowpass filter is designed by the usual method it will have overshoots on an applied squarewave. Adding the polishing filter can reduce those overshoots. But this active filter must have a gentle rolloff itself so that it does not cause more trouble (by introducing overshoots) itself. This simple 'polishing' filter cannot determine transmitted signal bandwidth. It is in addition to any lowpass filters in the speech processing chain. If this polishing lowpass causes overshoots on application of squarewaves, its purpose is mitigated. This filter must not show overshoots ("ringing") on its output if a squarewave is applied to its input.

The schematic and magnitude response of an active lowpass polishing filter are shown in Fig. 21.

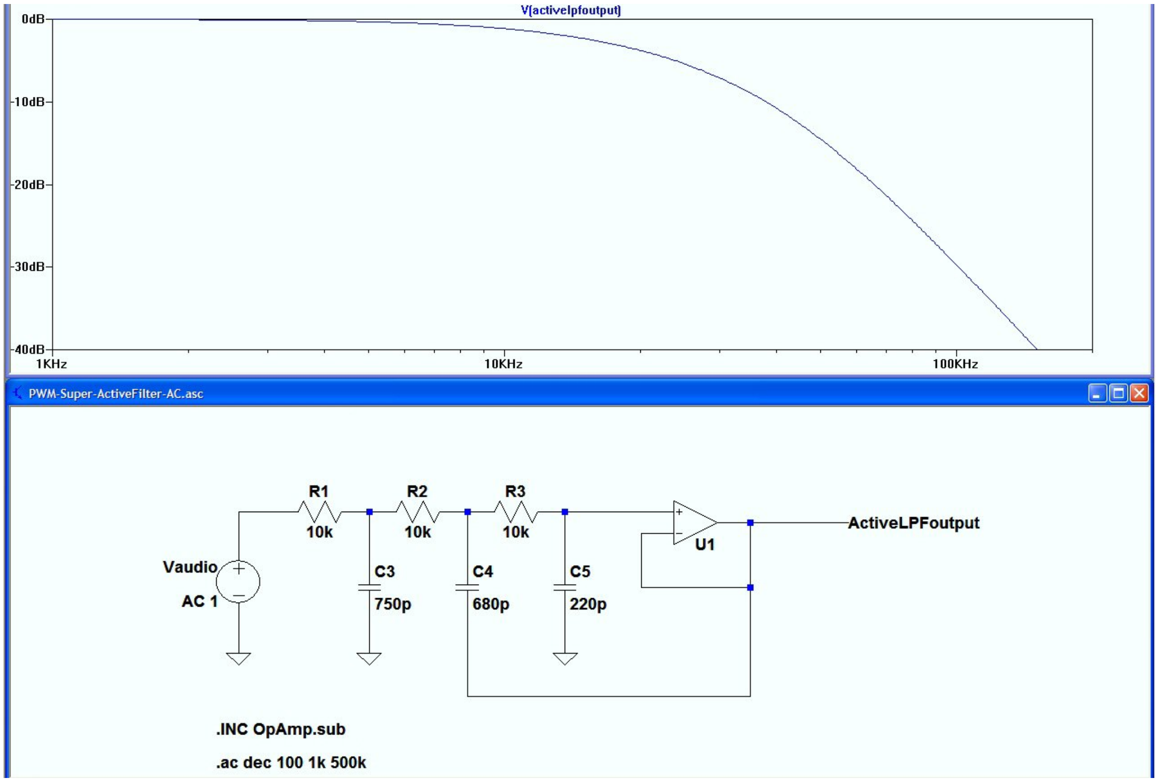

Fig. 21 - Schematic and magnitude response of the simple polishing filter

The transient response of this active lowpass filter when a 1000 Hz squarewave is applied to it is shown in Fig. 22.

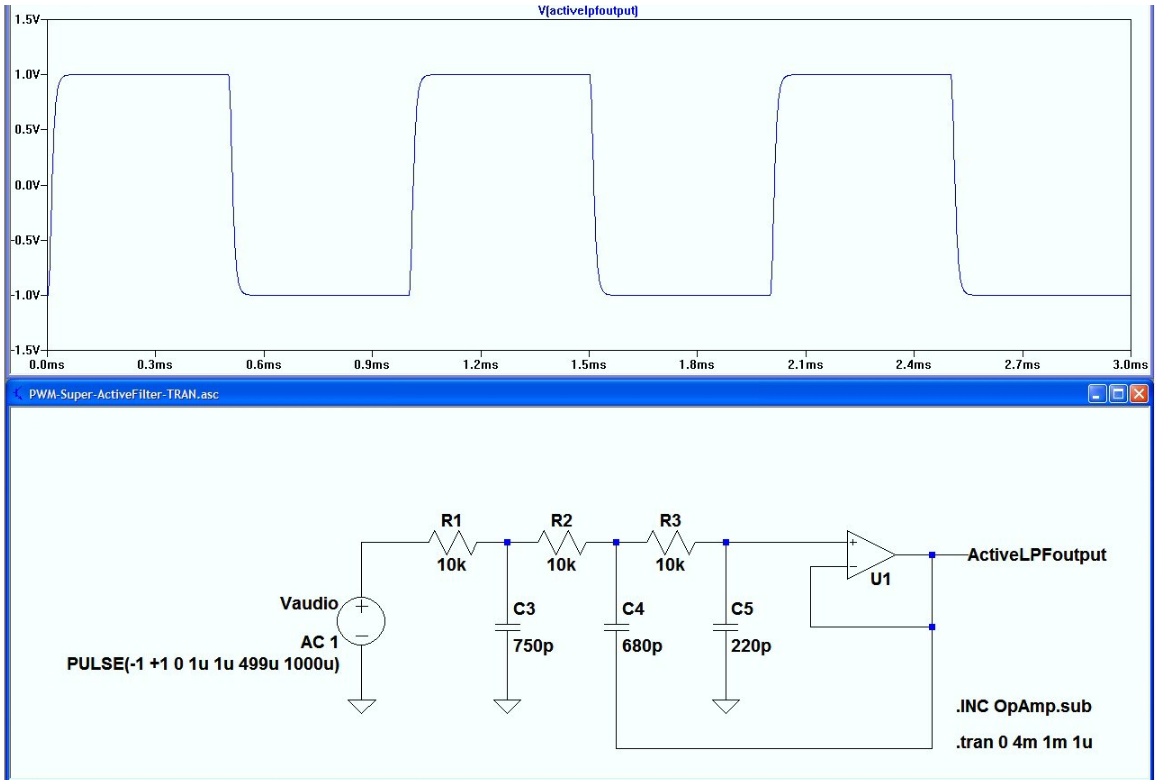

Fig. 22 - Transient response of the simple active lowpass itself

The transient response plot of that active lowpass filter shows that the edges have been slightly rounded. That filter's response is (essentially) that of the Gaussian family, which is characterized by having no overshoot. We will put this filter to use in a moment.

We will now look at some system transient waveforms. In each case we will excite the system with a 1000 Hz squarewave as the 'modulation.' Squarewaves are rather brutal; their main attribute here is to unearth problems.

First we see a waveform from the output of a doubly-terminated lowpass (which was designed by the usual and not recommended method) without this active lowpass in the modulation path.

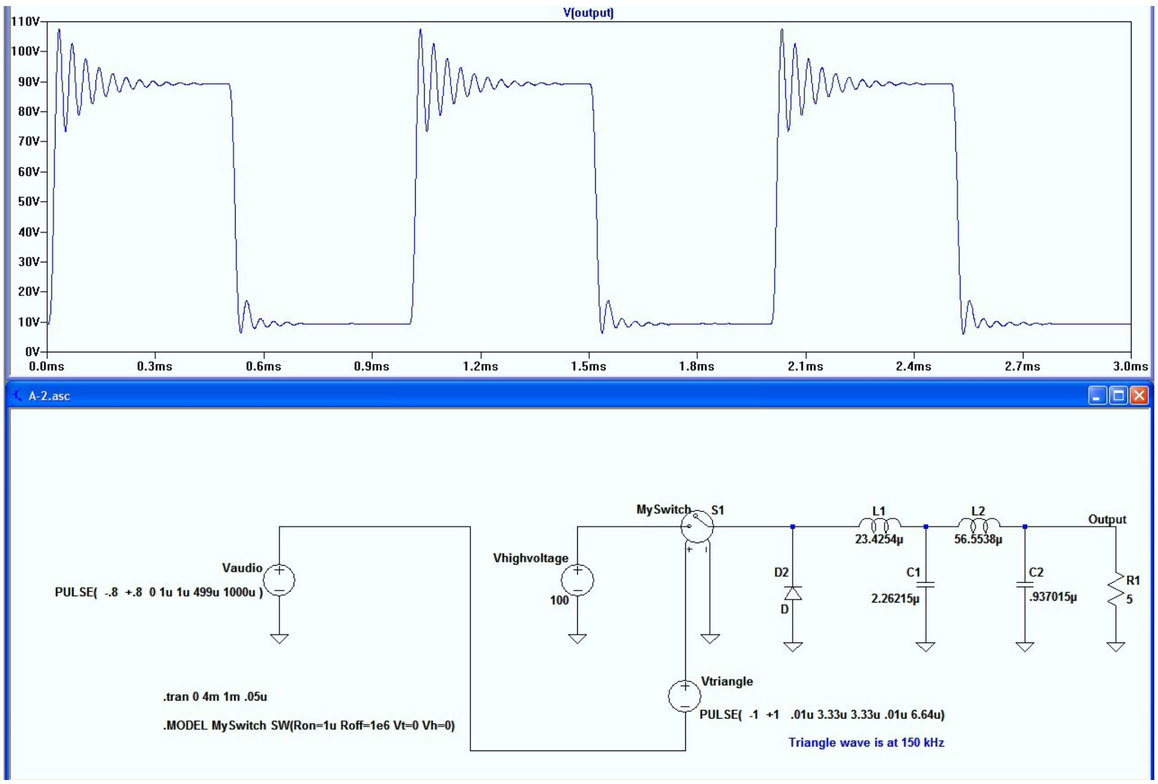

Fig. 23 - Transient response of system - no active lowpass used

To be observed also is that at full modulation levels as seen here the negative modulation peaks are not the mirror image of the positive modulation peaks.

That output filter was designed as the usual doubly-terminated Butterworth filter but instead of the intended input termination (same as the output termination), the filter is actually being driven from the switcher, which has a near-zero output impedance. This causes a severe peak in the magnitude response, along with a huge jump in the envelope delay at that frequency.

Those two items (the peak in the response and the peak in time delay at the cutoff frequency) act in concert to cause a terrible transient response (ringing).

If the squarewave is applied first to that active lowpass at the input to the modulator, the resulting modulator output is changed as shown in Fig. 24.

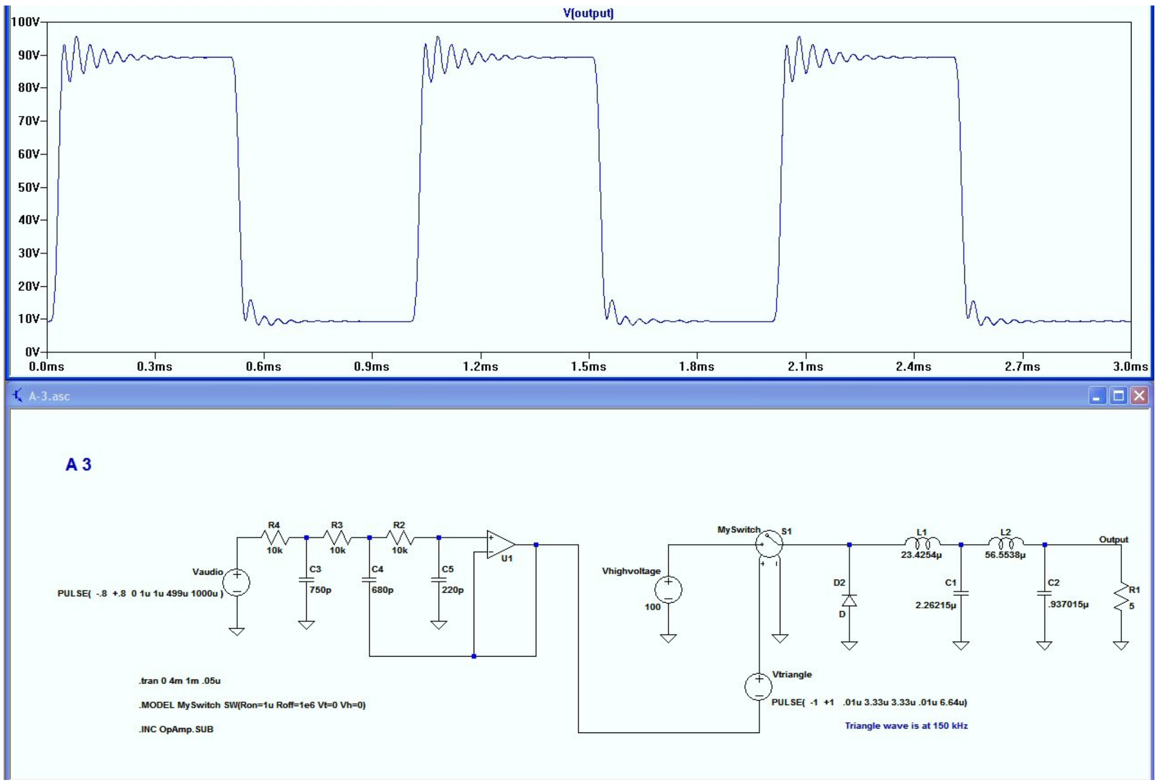

Fig. 24 - Transient response of system - with active lowpass

It can be seen that the serious degree of ringing on the output waveform has been reduced. It functions by reducing the transmission (i.e., dB response) near the output lowpass filter's cutoff frequency. Unfortunately the ringing has not been totally eliminated and also the negative peak problem remains.

Now let us look at the waveform from the output of the singly-terminated lowpass filter (designed by the recommended method), without the added active lowpass at the system input.

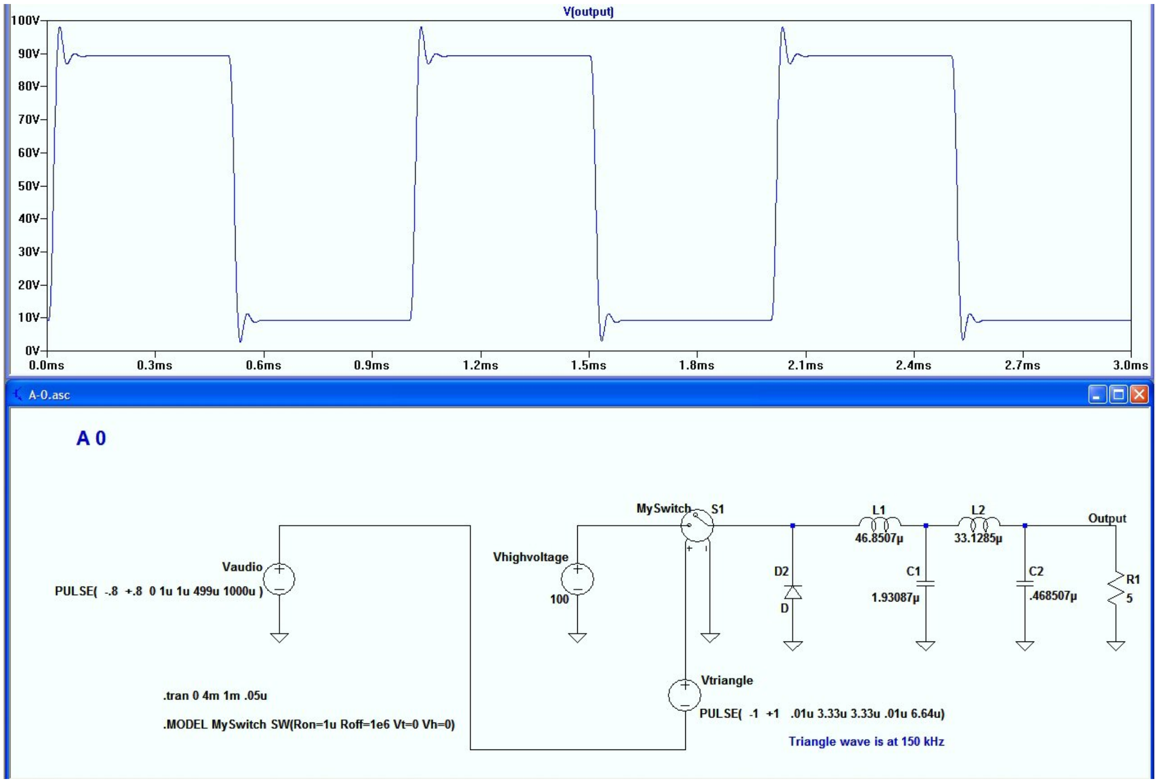

Fig. 25 - Transient response of the system using the recommended design method but no active input polishing filter

This design method results in much improved transient response, allowing clipped waveforms to be handled better.

If the squarewave is applied first to the active lowpass at the input to the system, the output then appears as in Fig. 26.

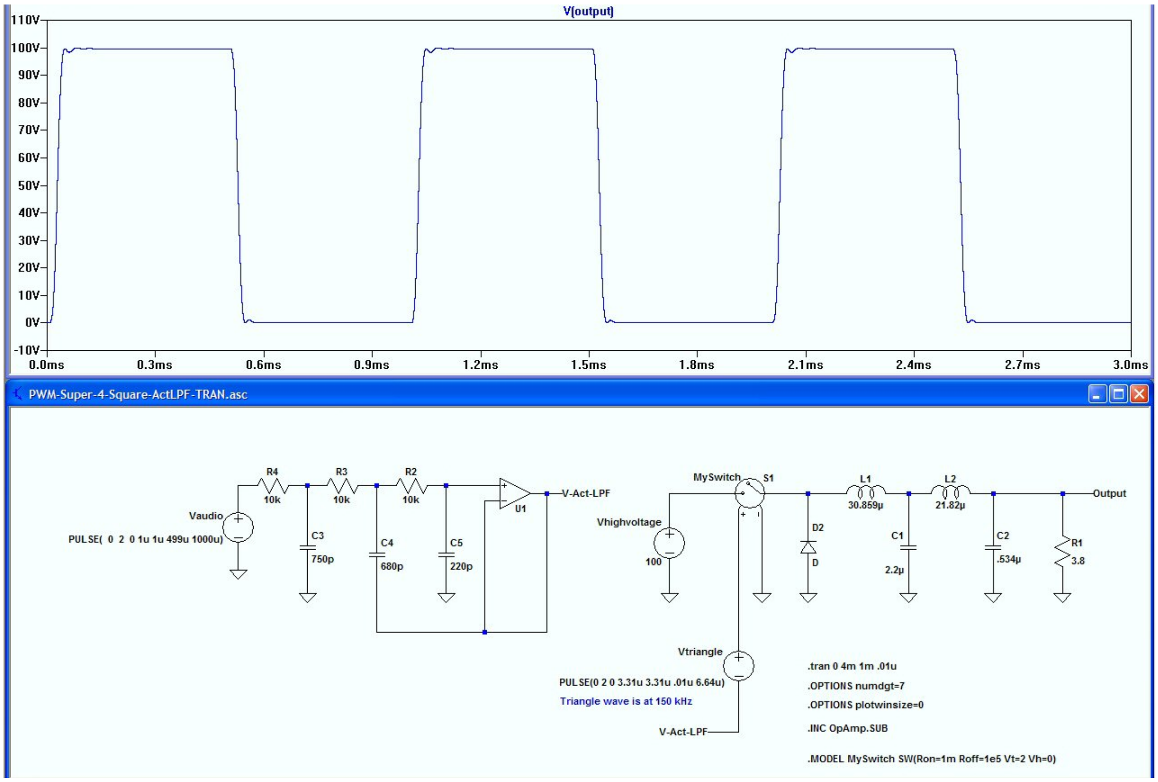

Fig. 26 - System augmented with the simple active lowpass

The small amount of overshoot from the switcher output lowpass has been further reduced. Clipped modulating waveforms will be handled *very* well.

A squarewave test signal is used here as a severe challenge to the modulator. The purpose of using a squarewave as a test signal is simply to uncover problems. A sinusoid that has been clipped several dB is a more realistic test signal and we will discuss that later.

This "polishing" lowpass filter is to be inserted into the system right at the input to the comparator (the PWM generator). Its purpose is *not* to determine the transmitted signal bandwidth. Rather it is to have of the order of 20 dB of attenuation at the main output lowpass bandedge frequency. It will then reduce the ringing caused by the main output lowpass. It will also reduce the possibility of "aliasing" caused by unwanted modulating signal components near the switching frequency. For this reason this simple filter is also called an "anti aliasing" filter.

Filter topology

So far in this document we have shown only inductor-input filters for the post-switcher output lowpass, without any mention of the reasoning behind that choice. But only inductor-input filters can be used at that point. Let us look at what happens when a shunt capacitor is at the input of the filter. This would be the case if an odd-order capacitor-input filter were to be designed and this is illustrated in Fig. 27.

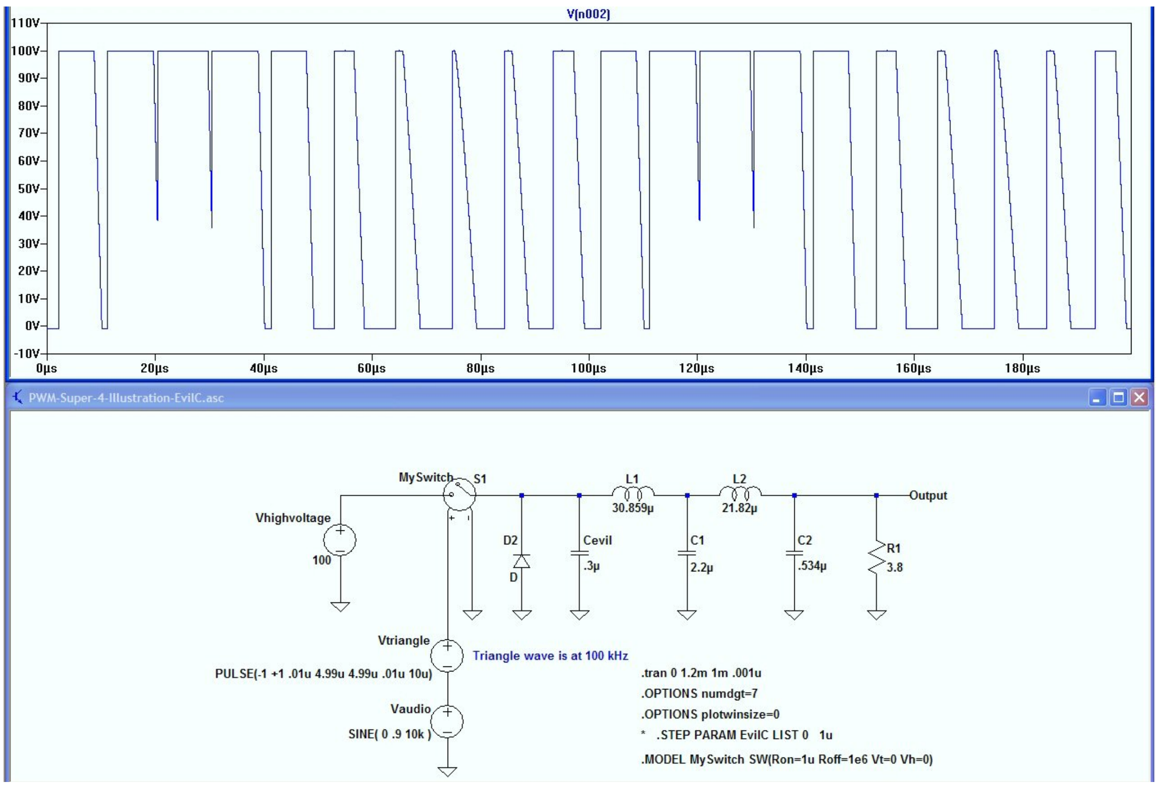

Fig. 27 - Looking at the switcher output to see the effect of a shunt capacitor on filter input

We are looking at the output of the switch here, the junction of the switch and the following lowpass filter to which a shunt capacitor to ground has been added. On each cycle of the switcher its output (the input to the filter) instantaneously rises right up to the power supply voltage. At the end of that cycle the output does not immediately drop to zero as intended but rather it 'ramps down' instead. The times and frequencies involved to generate this graphic have been adjusted to clearly illustrate the problem.

Now we will change the analyzer's time and frequency values to see the problem by showing both the switcher output and also the lowpass filter output in a single graphic. See Fig. 28.

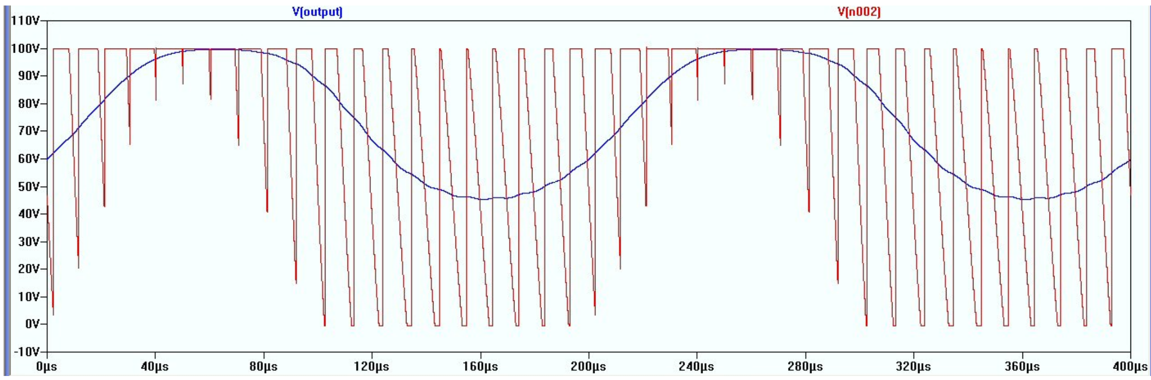

Fig. 28 - Another illustration of the problem with a shunt capacitor on the filter input

This graphic shows both the waveform at the capacitor-input filter's input and also the waveform as delivered at that filter's output.   As can be seen, the output of the lowpass has two problems. One is the obvious distortion and the other is what might be called 'axis shift.'

The high efficiency of the switching system arises from the fact that the switches spend very little time (ideally, none) in their linear region. The graphics above show a system that has the switch spending considerable time in the linear region. The 'sloping' portion of the waveform is the linear portion.

The double-switch system has an even more ominous problem should a capacitor-input output filter be used. As soon as the shunt switch is turned on it is being asked to instantly discharge that shunt input capacitor. There is nothing to limit the current flow and so the shunt switch will probably be destroyed.

The audio itself at that filter's output is shown in Fig. 29.

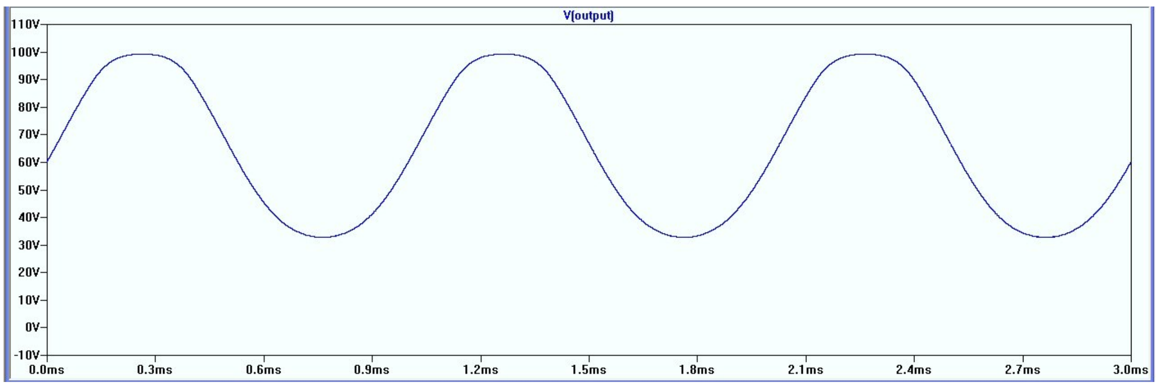

Fig. 29 - Audio from the filter with a shunt capacitor on its input

Fourier analysis reveals that the harmonic distortion of that waveform is right at ten percent.

Triangle wave distortion

We have assumed that the 'triangle wave' in the switching system was ideal. What happens if it is only an approximation in that the rise time and fall times are not linear (not "straight lines") ? That could be the result of accepting a squarewave and passing it through a lowpass R-C network in an attempt to 'approximate' a triangle. A resulting 'sort of' triangle, taken to an extreme for illustration of the problem, is shown in Fig. 30.

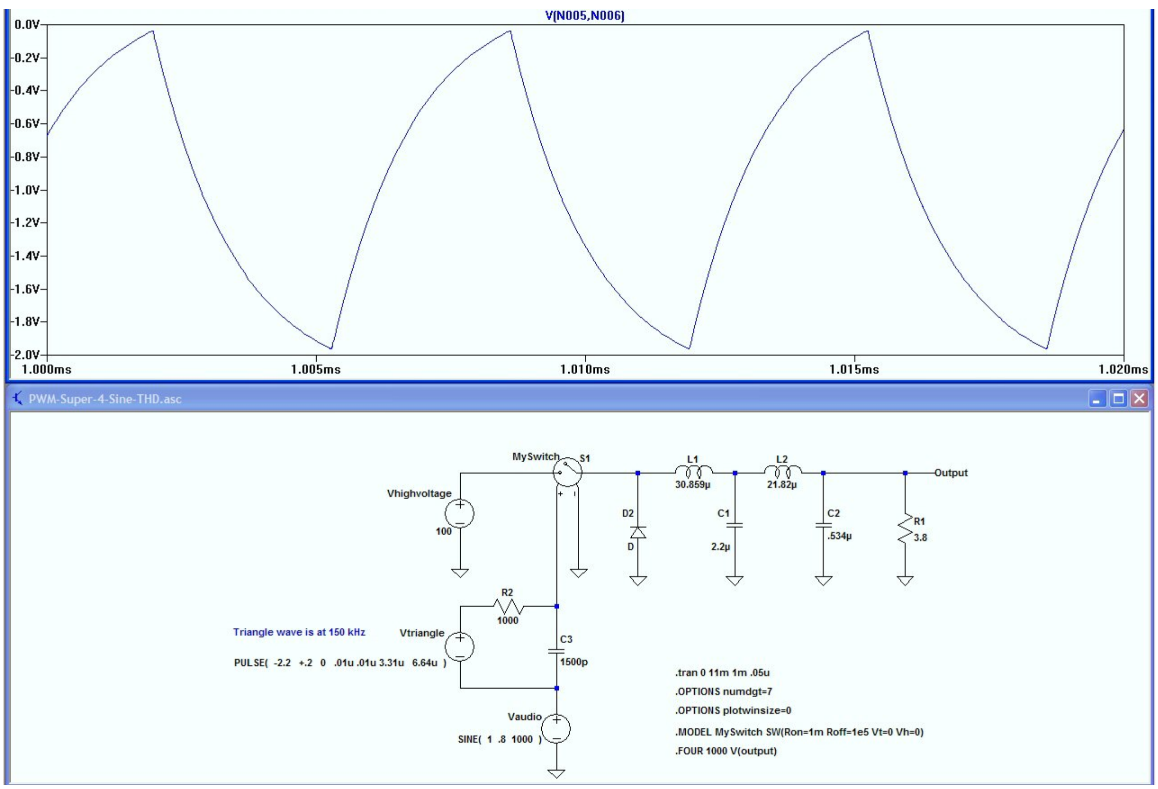

Fig. 30 - Triangle wave approximation

That graphic shows the waveform that results when an ideal squarewave at the switching frequency is simply passed through an R-C network in a crude attempt to make an approximation of a triangle wave.

The resulting output of the system, the audio output from the lowpass filter, is as a result distorted. The test sinusoid has had both the positive-going and the negative-going peaks 'stretched' somewhat. The output signal that results is shown in Fig. 31.

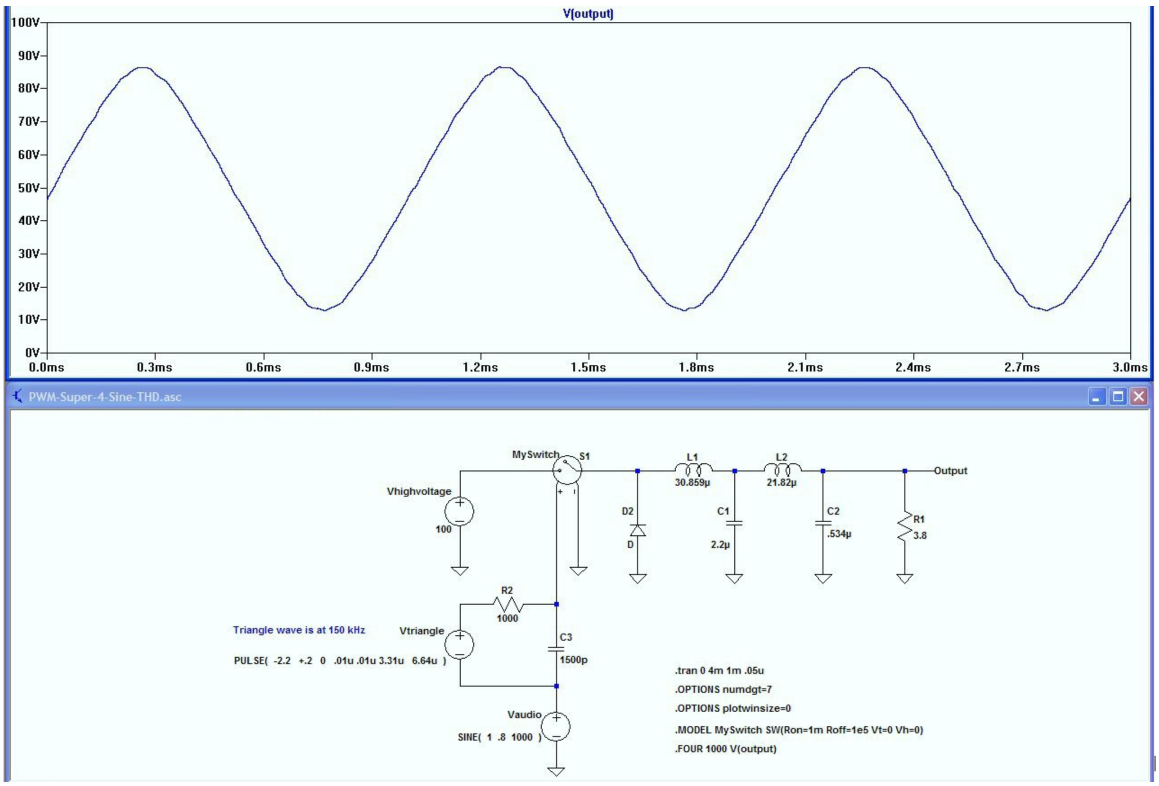

Fig. 31 - Audio out of the filter when a poor triangle wave is used

Odd-harmonic distortion results; this seen as a spectrum in Fig. 32.

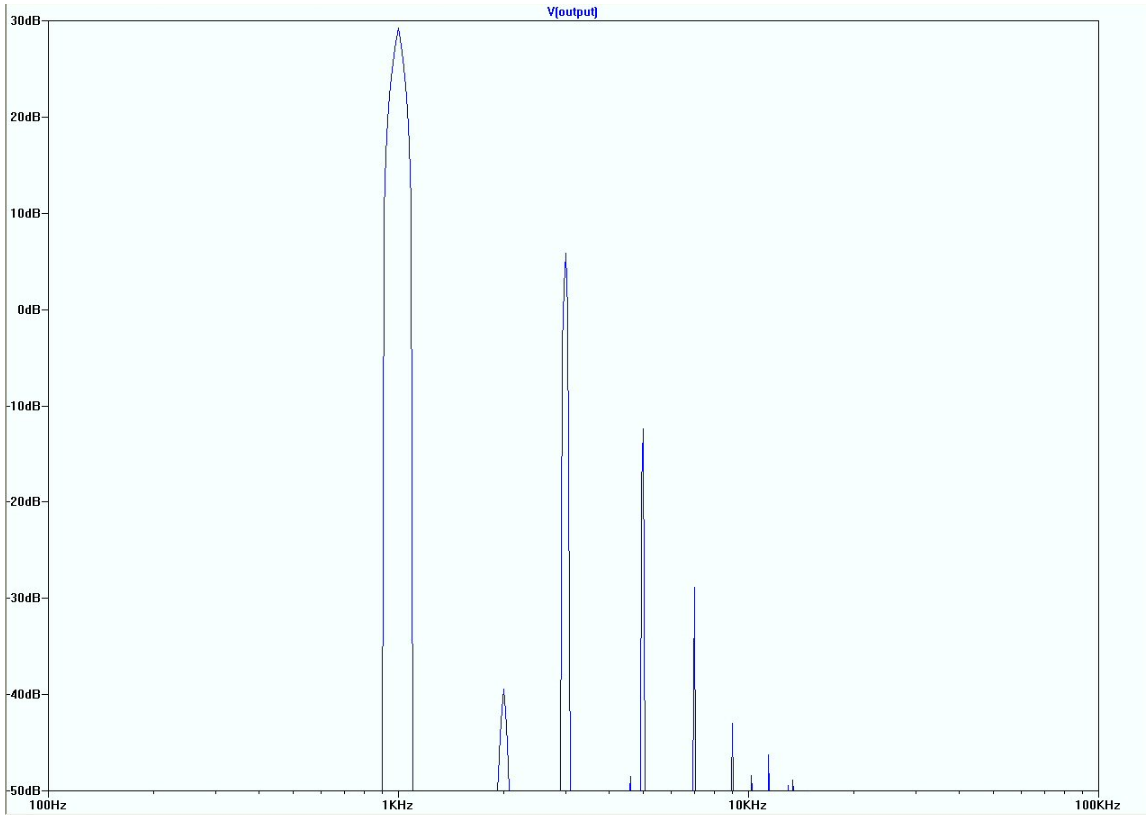

Fig. 32 - Spectrum of audio from a system with a poor triangle wave

This spectrum plot shows a fundamental at 1 kHz and a large amount of odd-order harmonic content (at 3 kHz, 5 kHz, 7 kHz etc.). A bit of second-harmonic component appears about 60 dB below the desired signal, doubtless a bit of 'math dirt' in the analyzer.

Integrated circuits are available that generate nice triangle waves. An example of this kind of IC is the LTC6992. That IC also includes the comparator. Its output is a logic-level PWM signal ready to be applied to an interface IC whose outputs drive the following power MOSFETs. Whether the designer goes that route or generates his own, the best method of triangle wave generation is to use an opamp with feedback to form an integrator. The LTC6992 will be discussed later.

Designing the output lowpass filter

Now we are going to correctly design the output lowpass filter. This detailed procedure involves both design and analysis in an interactive fashion. The software used is *Elsie*, the filter design and analysis program from Tonne Software. Free and capable.

The design family chosen will be Butterworth because that family yields a nice flat passband with only modest overshoots on transient waveforms. Indeed, an alternative name for that family is 'maximally-flat magnitude.' And the transient response (how squarewaves are handled) is quite acceptable for this application. Not perfect, but adequate.

The input and output termination resistances are equal in the usual design methods. Those methods yield what filter designers call a "doubly-terminated" filter. But **for this application a singly-terminated filter design is highly advocated** and that is what we will use. This is because the switcher output has an extremely low (essentially zero) output impedance. Recently an easy-to-use singly-terminated design routine been made available (in *Elsie*). A search on the 'net turned up some singly-terminated designers but the analysis of the design (plotting of response) was missing or poor in some way. In *Elsie* the design and analysis are nicely interactive.

Simply using whatever software might be available to design the usual doubly-terminated filter and then just conveniently forgetting all about that input termination actually being essentially zero will result in the design being clearly less-than-optimal, as we have seen. Easy but absolutely incorrect.

The output termination is set with the value of the load as seen by the filter output. (As an example, an RF amplifier operating at 50 volts with no modulation along with a current of 10 amps would look to the filter like a 5 ohm load.)

In the design process outlined next, the output filter bandwidth is adjusted until the attenuation at the switching frequency is adequate (60 dB of attenuation in the filter is the recommended value). That indirectly defines the passband width of that filter.

Using Elsie to design the output lowpass filter

If you don't have the current version of Elsie, download it from the internet:
http://tonnesoftware.com/elsiedownload.html

1. When you are at that URL, click on "Download Elsie"

2. Install the program. You may tell your antivirus app that Elsie is a nice lady and so it is OK to run with her.

3. Call it up by simply clicking on the Elsie desktop icon.

4. On the opening screen, click on "New design"
You are now on the Design page.

5. On the second row of buttons, click on "Singly-term filter"

6. On the new page, set "Order" to "4"

7. Set "Output termination" to the value your RF amplifier will present to the filter's output. For example if the amplifier idles at 50 volts DC and a current of 10 amps then use a figure of "5" here.

8. Set "Passband ripple" to 0 to generate a Butterworth design.

9. For this first pass at the design, set "Bandwidth" to "30k"

10. In the Topology area, click on "Inductor-input lowpass"

11. We will probably have to adjust the bandwidth to make the attenuation at the switching frequency in spec (60 dB down from the audio region). But first plot the response to see how the just-entered design performs.

12. Click on "Accept and return to Design page." You are now back on the Design page.

13. We may have to define the plotter settings this first time around. On the first row of buttons, click on "Analysis"

14. Set "Analysis start frequency" to "5k"

15. Set "Analysis stop frequency" to "200k"

16. Set "Transmission bottom" to "-100"

17. Set "Number of Y intervals" to "10"

18. Click on "Plot"

19. You can see that the attenuation at the switching frequency of 150 kHz is insufficient. (Of course if you are using a different (lower?) switching frequency, substitute that frequency for "150" in this tutorial.)

20. Let's make it easy to see the responses at a couple of critical frequencies. On the second row of buttons, click on "Markers"

21. In the small window, enter "5k" in the box next to "Frequency 1"
Enter "150k" in the box next to "Frequency 2"

22. Click on "Show"

23. At the bottom of the plot see the signal level at 5 kHz with the computed loss at 5 kHz of 20 dB. This design routine has assigned an input termination value of the output termination divided by 398.04. This is so that due to mismatch the plotted insertion loss is right at 20 dB. It is not a

real-world problem; it is a byproduct of the analysis arithmetic. Far more important is the loss at 150 kHz: seen here to be insufficient.

We are looking for a response at the switching frequency to be 60 dB down from the response at 5 kHz.

24. So we must do a redesign. At the screen top, click on the "Design" button.

25. On the second row of buttons, again click on "Singly-term filter"

26. Adjust the bandwidth downward a few percent.

27. Again select the lowpass topology and then accept this new design.

28. That sends you back to the Design screen; click on "Plot"
See the improved figure at 150 kHz.

29. Return from the plot back to the Design page, adjust the bandwidth and re-plot.

30. **We want the response at the switching frequency to be 60 dB down from the response at 5 kHz.** When you have adjusted the bandwidth to provide that degree of attenuation at the switching frequency, you have finished the design. Call up the schematic and send it to the printer.

Addendum

We are now going to look at some related material and review some of the earlier-discussed material. This section of this paper also includes notes on the "speech amplifier" that drives the modulator circuitry.

Squarewave modulating waveforms

A first item we will look at is to look again at the way that the output filter handles squarewave waveforms. Squarewaves are commonly available in audio signal generators to allow testing of the transient response of systems. The thought here is that if the system can handle squarewaves it can handle any waveform thrown at it. Squarewaves are an extreme kind of waveform whose use here is to flush out problems. Perhaps a clipped sinusoid would be more practical for testing. Or, as will be seen, a wave we are going to call "trapezoidal." We are going to present a wave that is easily duplicated in LTspice so that various scenarios can be tested: "trapezoidal".

First let's look at a squarewave and its spectrum. See Fig. 33.

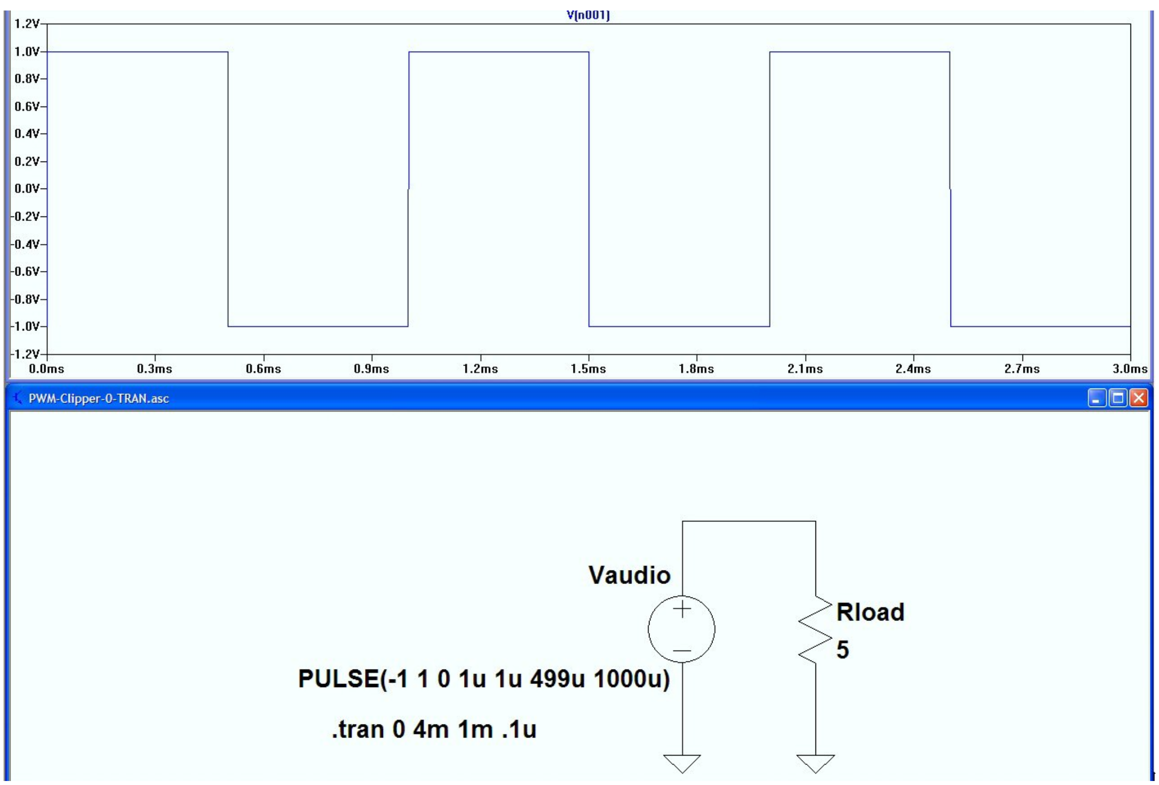

Fig. 33 - Squarewave generation and the resulting squarewaves

The spectrum of that 1000 Hz squarewave is shown in Fig. 34.

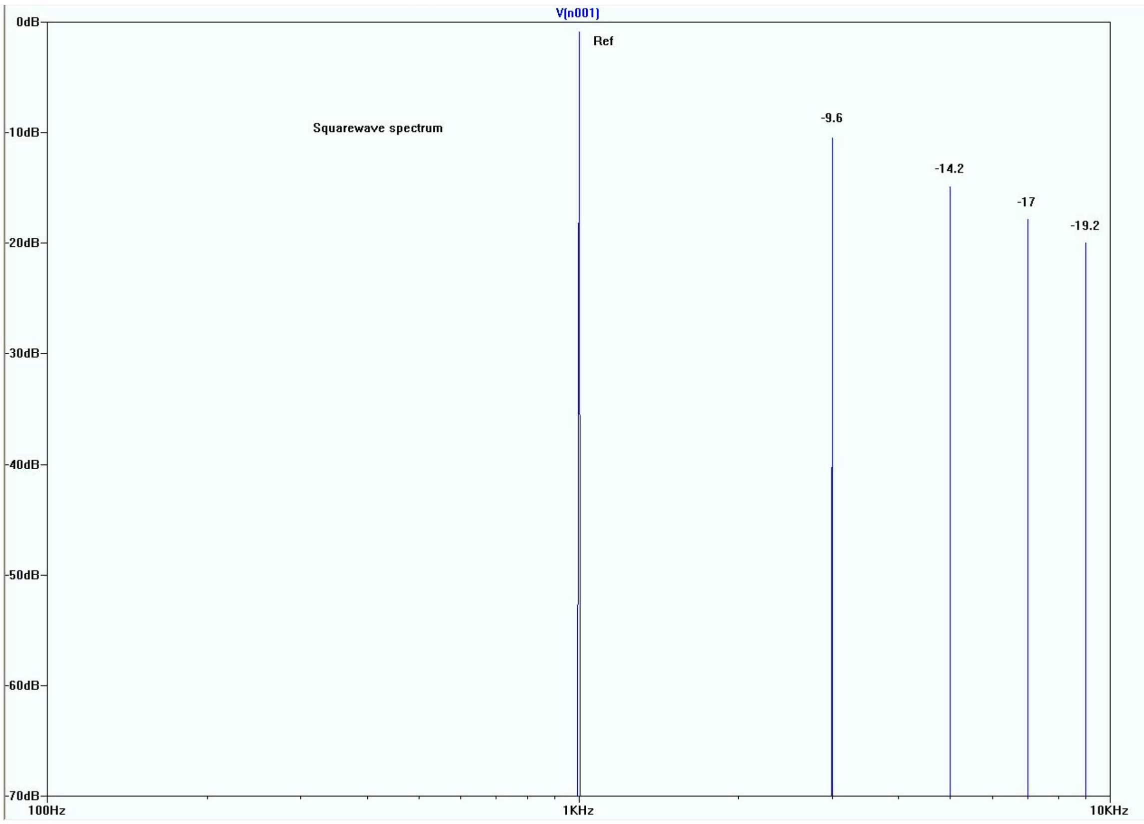

Fig. 34 - Spectrum of the 1000 Hz squarewave

Those harmonics of the waveform continue on without limit. This display, however, stops at 10 kHz and so shows the harmonics of that 1 kHz squarewave up to 9 kHz.

Now we are going to generate a sinusoid that has been clipped 10 dB. This is a believable waveform in amateur radio practice. (A signal that has been clipped 10 dB is not hifi, but it's believable and, I suspect, commonplace.)

The LTspice implementation of the clipped sinusoid is shown in Fig. 35.

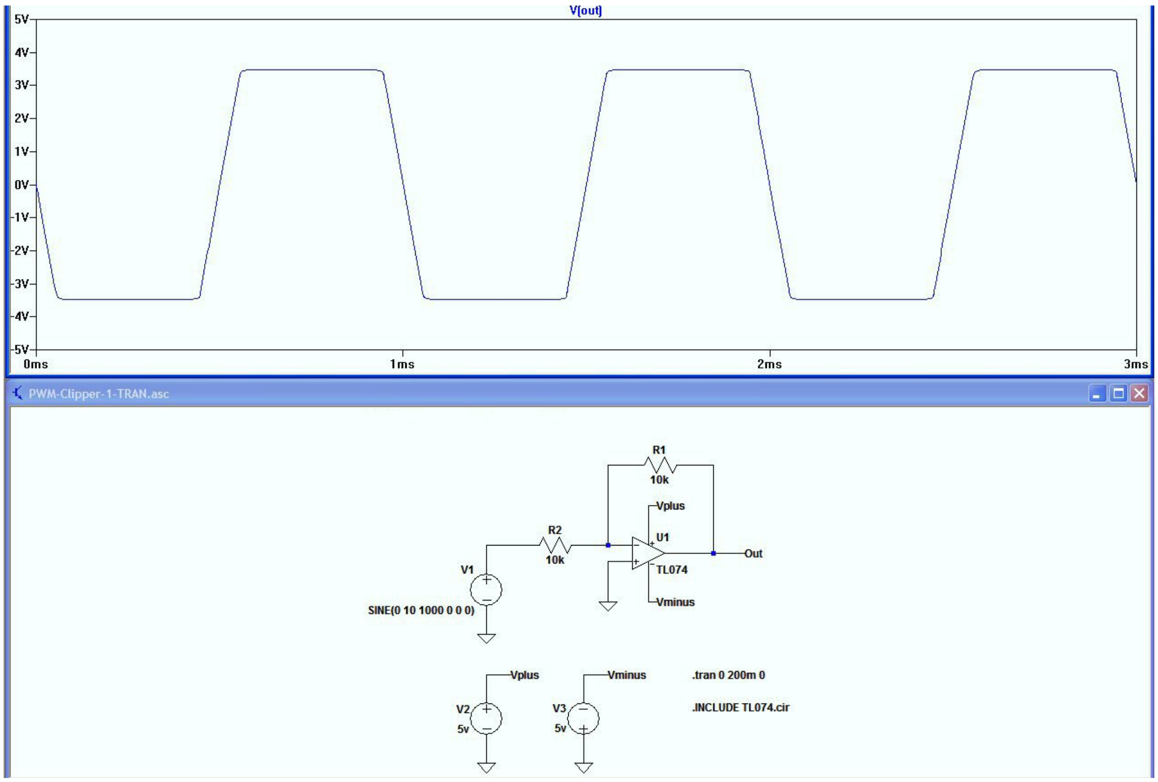

Fig. 35 - Schematic and output of a clipped-sinusoid generator

The spectrum of that circuit's output is shown in Fig. 36.

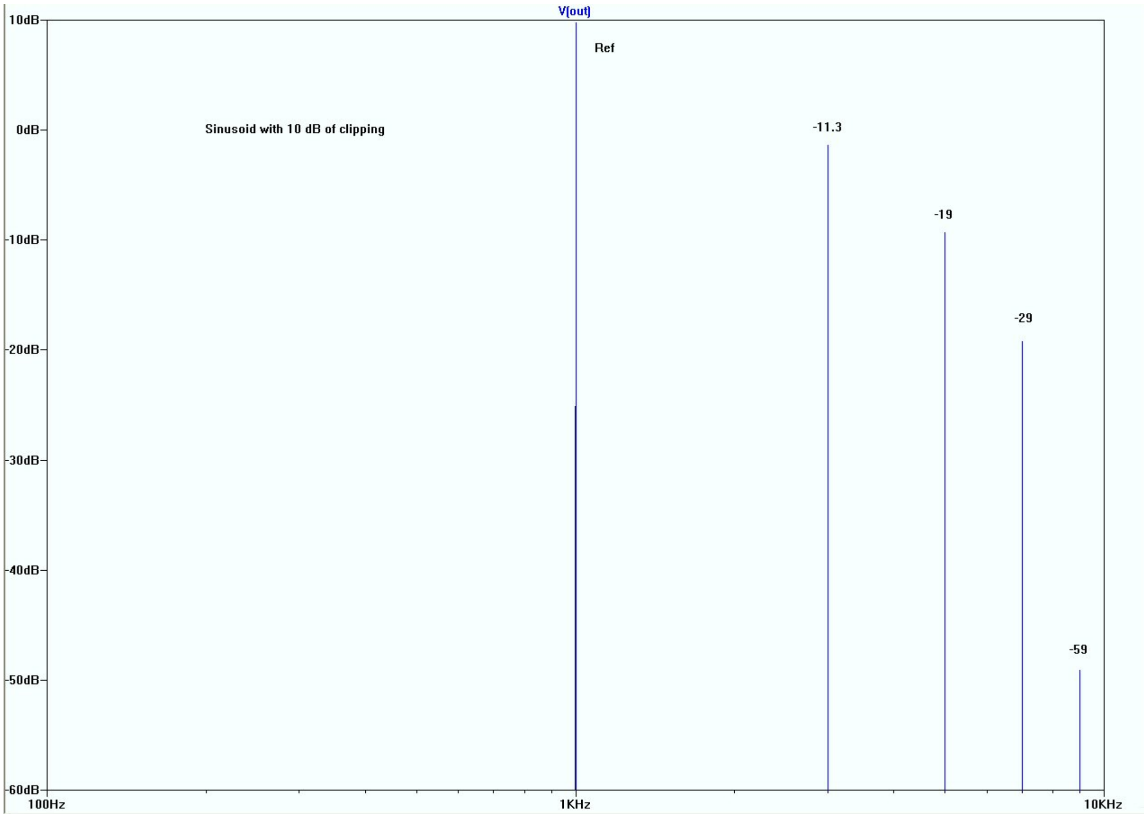

Fig. 36 - Spectrum of the output of the clipped-sinusoid circuit

But for analysis using the computer, this way of generating a clipped waveform is awkward. It requires positive and negative power supplies, an opamp and a pair of resistors.

There is a better - simpler - way (for our purposes) to generate that waveform: a pulse-generator with finite specified rise and fall times. See the resultant circuit and waveform in Fig. 37.

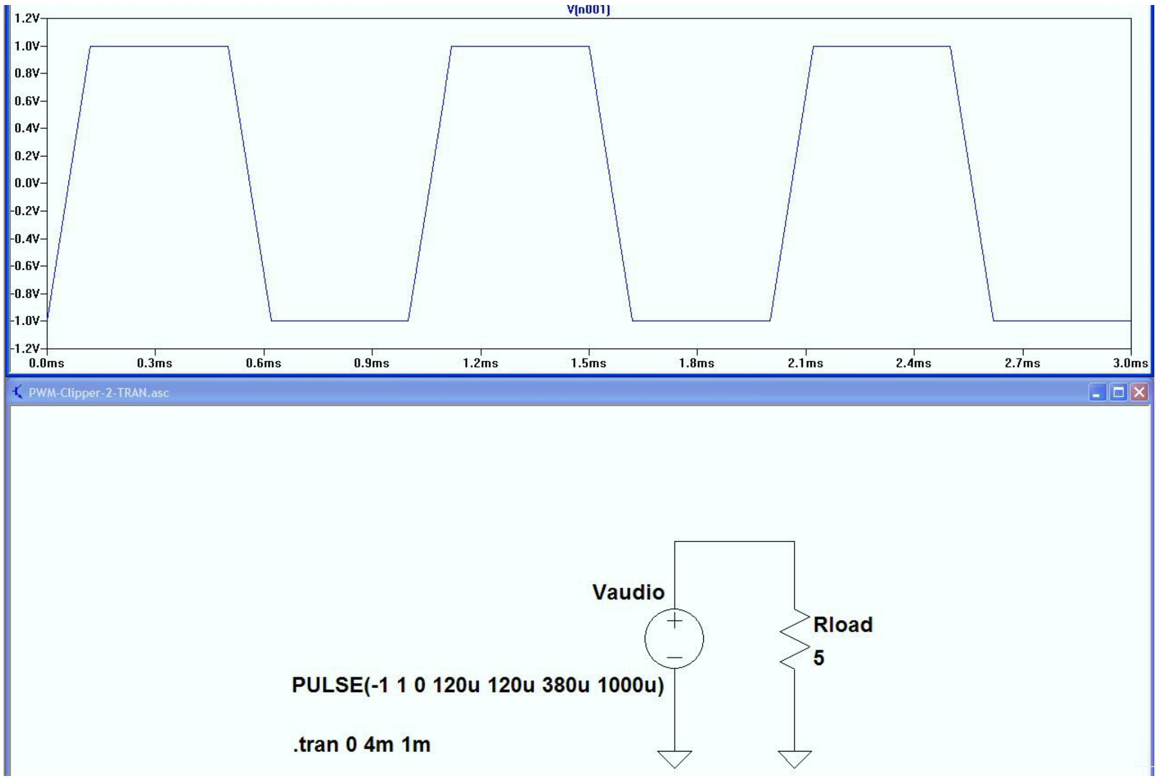

Fig. 37 - A clipped-sinusoid ("trapezoid") as implemented in LTspice.

This is easy to implement: an *LTspice* voltage source (as a pulse). The risetime and falltime are each set to 12% of the pulse's length. The top and bottom are each set to 38% of the total length. Of additional interest to LTspice users is that this generator has zero output impedance, of great value in these discussions.

See the spectrum of that output waveform in Fig. 38.

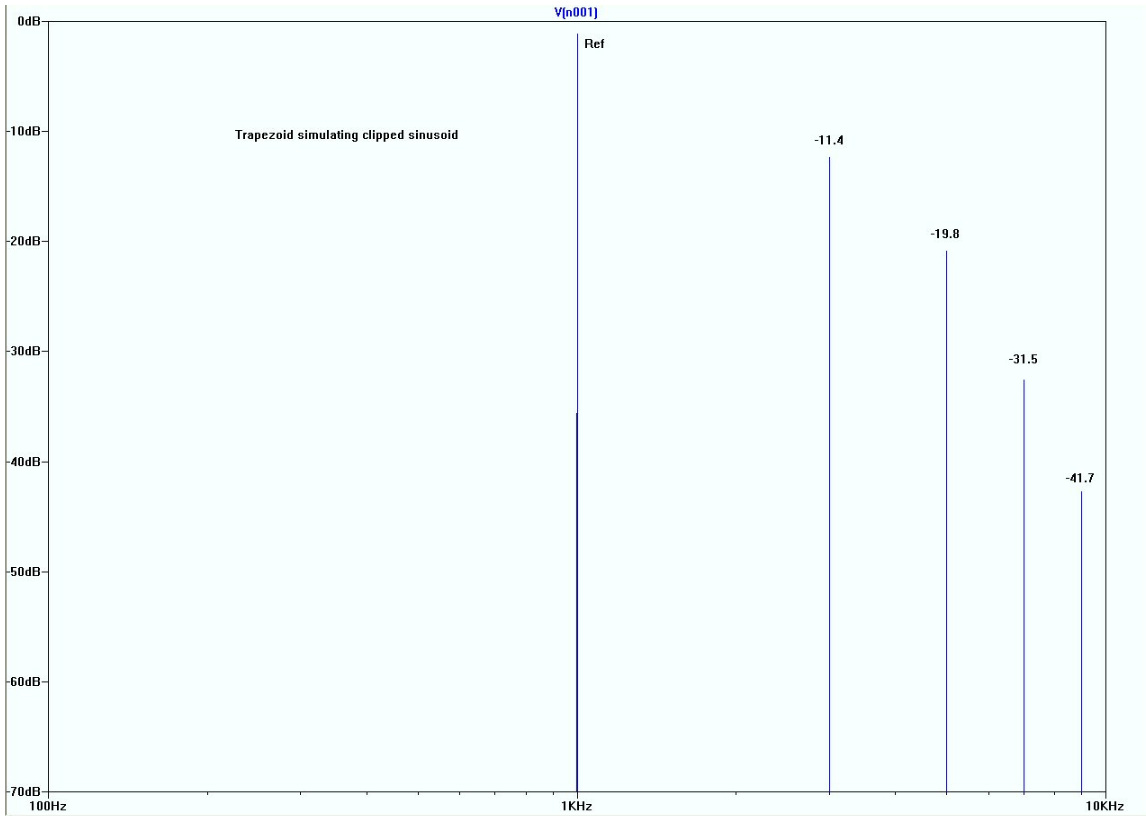

Fig. 38 - Spectrum of LTspice-generated trapezoidal wave simulating the clipped sinusoid

The waveshape and the spectrum of this waveform is strikingly similar to the sinusoid that was clipped using an op-amp. This similarity allows us to use this very simple waveform-generating scheme to look again at the transient response of our output lowpass filter. See Fig. 39.

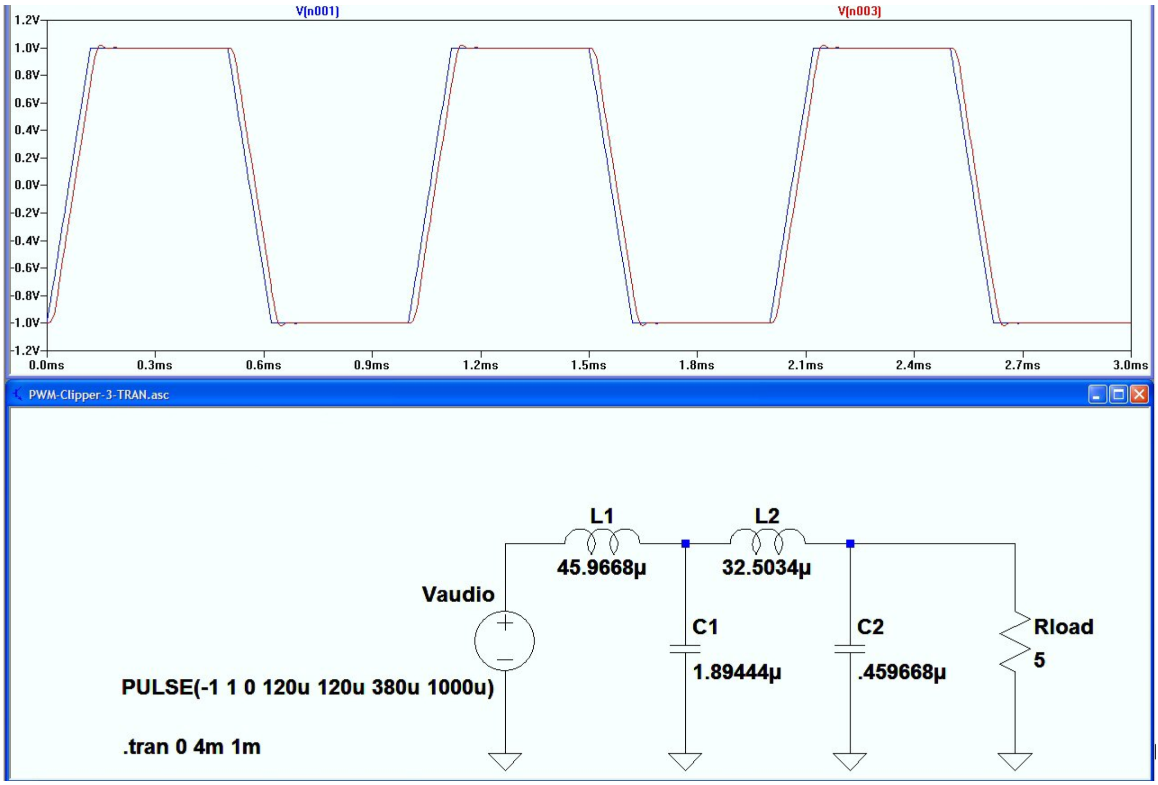

Fig. 39 - The "trapezoidal" waveform along with the waveform as it exits the correctly-designed output lowpass.

The trapezoid is shown here in blue and the signal from the output lowpass filter is in red. We can see that this realistic modulating waveform is hardly modified by the correctly-designed output lowpass filter.

Triangle vs Sawtooth

Let us now look at the spectrum of the switcher output first using a triangular wave to feed the comparator and then a sawtooth wave. This latter is an extreme case of an asymmetric (time-wise) distortion of that wave. The rising and falling portions of this triangle wave are straight-lines as viewed on a 'scope.

First see the switcher output spectrum with a triangular waveform at the switching frequency. For this illustration the switcher has been set to 100 kHz and the modulating sinusoid is at 5 kHz.

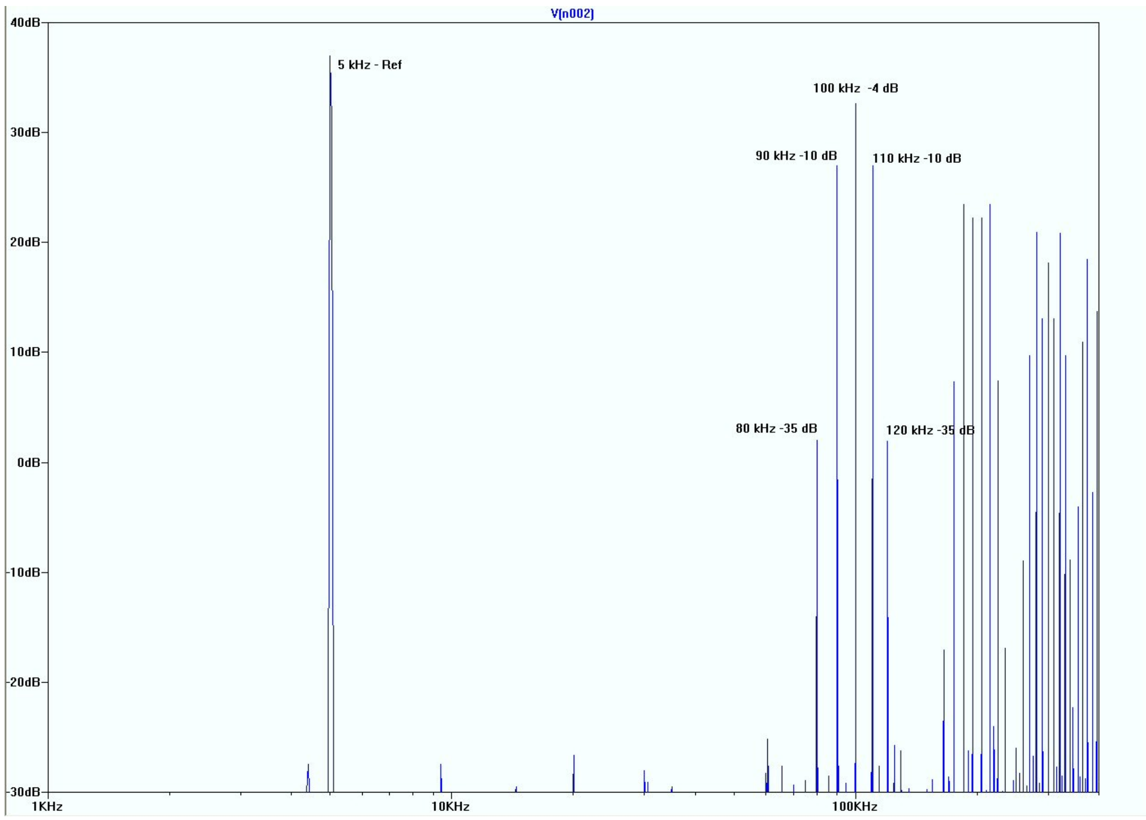

Fig. 40 - Spectrum of switcher output with the comparator being driven by a triangular waveform

Of interest here is that the "sidebands" about the switcher frequency are spaced at twice the modulating rate. But at the harmonics of the switcher signal these sidebands are spaced at the modulating rate.

Now let's look at the switcher output when the risetime and falltime of the triangle are very different: it is a "sawtooth." See Fig. 41.

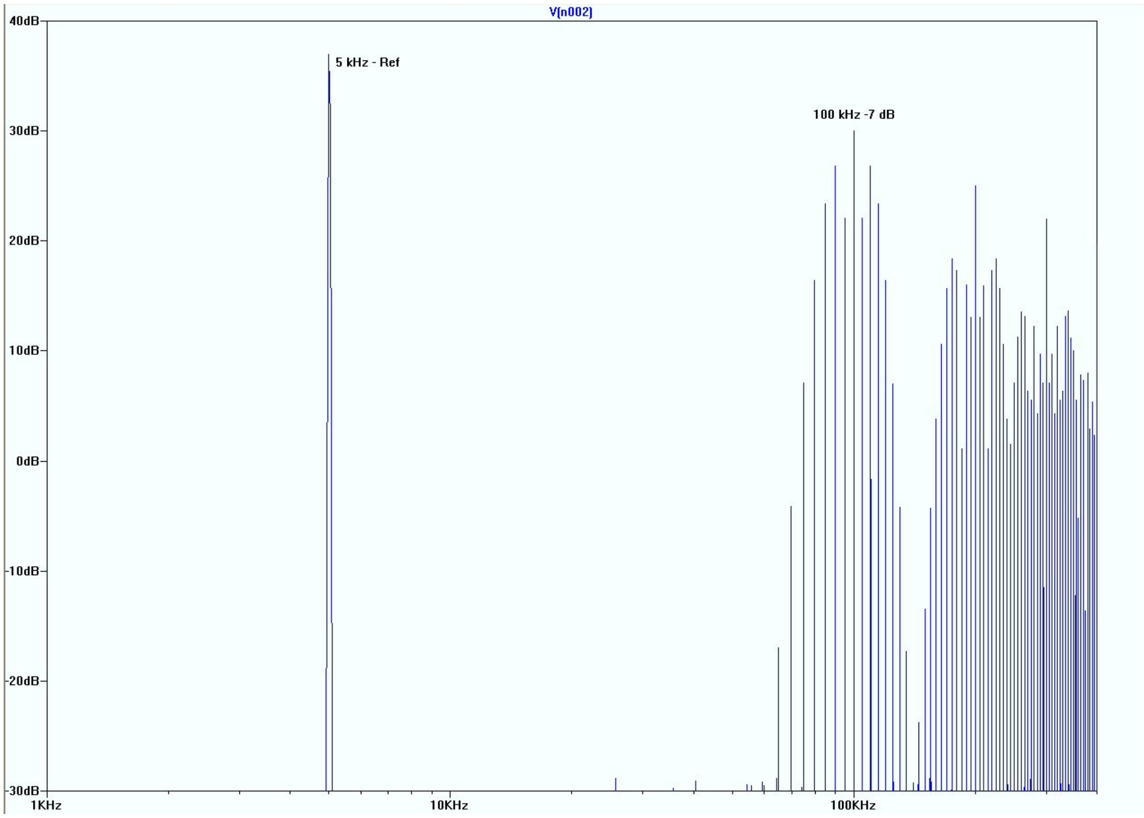

Fig. 41 - Spectrum of switcher output with the comparator being driven by a sawtooth (very asymmetric time-wise triangle) waveform

For this test the triangle has been degenerated into a sawtooth. This forms what might be called a "pulse duration modulator" because only one side of the "triangle" (the sloping side) is involved in the pulse-forming process.

The sidebands of the switching signal and its harmonics are spaced from the switching signal at the modulating rate.

The audio output from the output lowpass filter is unaltered; it remains the same as when an "accurate" triangle is used. In each case we have assumed that these triangle waves (or sawtooth waves) are ideal, not approximations. As seen on an oscilloscope the waveform is made of straight lines.

A two-pole output lowpass filter

It would seem possible to drop the filter order down to two. Such a design will then be simply a single series inductor followed by a single shunt capacitor.

Using the singly-terminated design method in Elsie for a two-poler we see that the 3 dB bandwidth has to be lowered down to 4500 Hz to get 60 dB of attenuation at a switching frequency of 150 kHz. It would have to be pulled in even farther if the switching frequency is lower; a 75 kHz switcher with a 2-pole lowpass would need its 3 dB bandwidth set to 2250 Hz to achieve 60 dB of attenuation at 75 kHz. This is simply unacceptable.

For a 5 ohm load the series inductor would be about 250 uH and the following shunt capacitor would be about 5 uF. The input impedance (seen looking into the filter) is remarkably constant out to the cutoff frequency. This means full modulation capability out to the cutoff frequency is possible. The transient response (how this filter behaves when the modulating waveform is trapezoidal or a squarewave) is quite nice. Such a two-pole design seems impractical for lower switching frequencies because the audio bandwidth would have to be lowered as mentioned above. The parts values, in particular the inductor, will be large. Advice: forget considering a two-pole version of the lowpass unless the switching frequency is at least 150 kHz.

Volts and Amps

It is appropriate to study the recommended output filter design to see the volts across and current through the components. For this study the power supply voltage has been set to 100 and the output load has been set to 5 ohms. The graphics should be easily scaled for other voltage and impedance values. With minimal explanation we now see the values indicated by each figure's label.

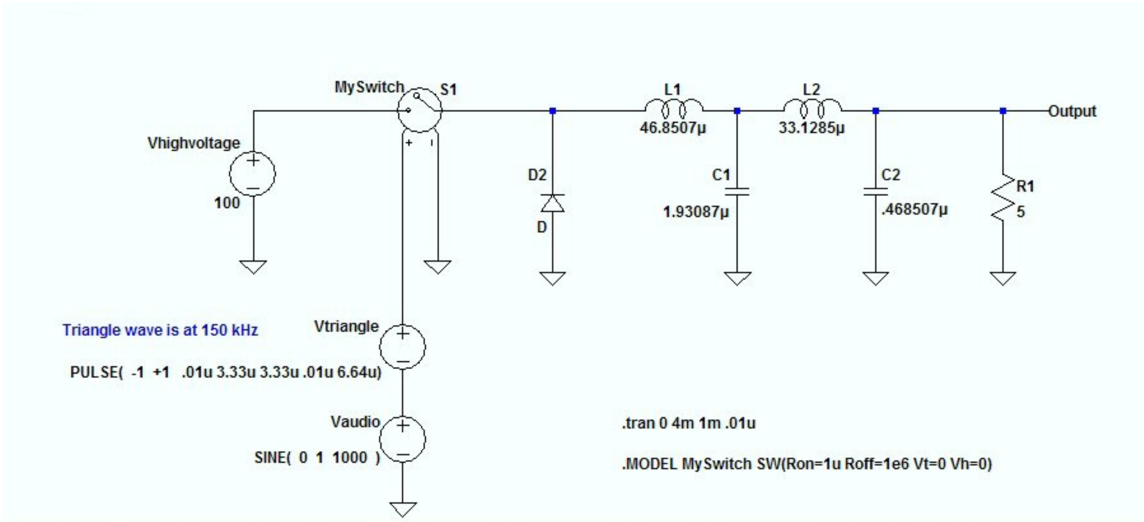

Fig. 42 - LTspice rendering of the modulator used for these tests

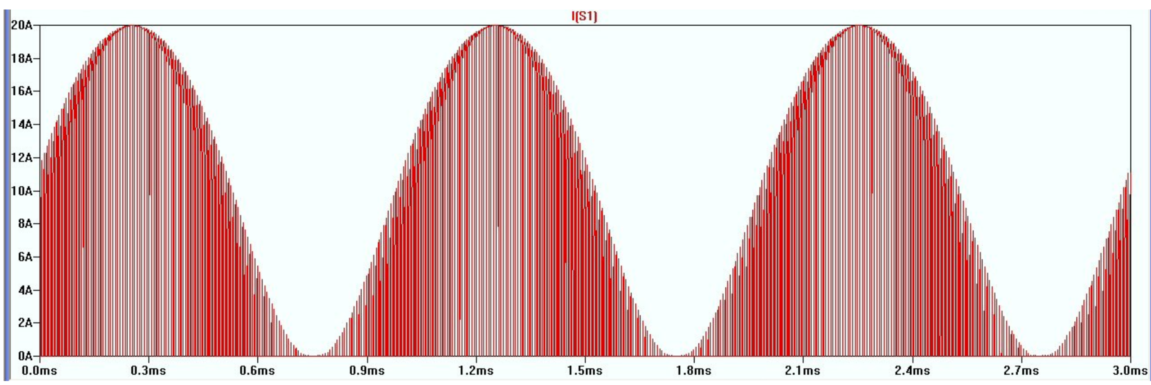

Fig. 43 - Current through the series switch

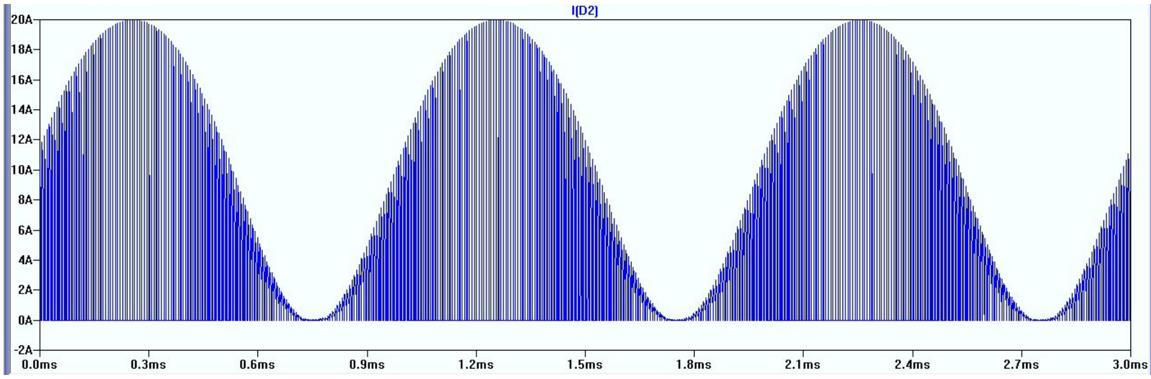

Fig. 44 - Current through the shunt diode

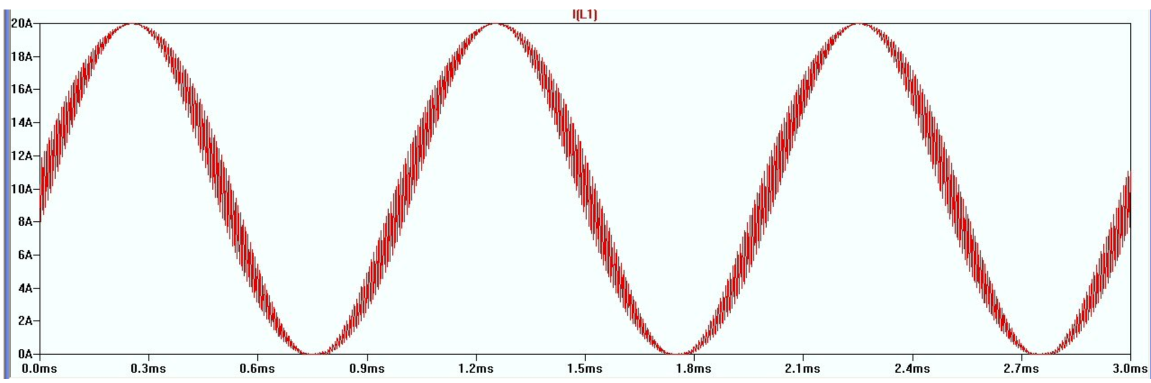

Fig. 45 - Current through L1

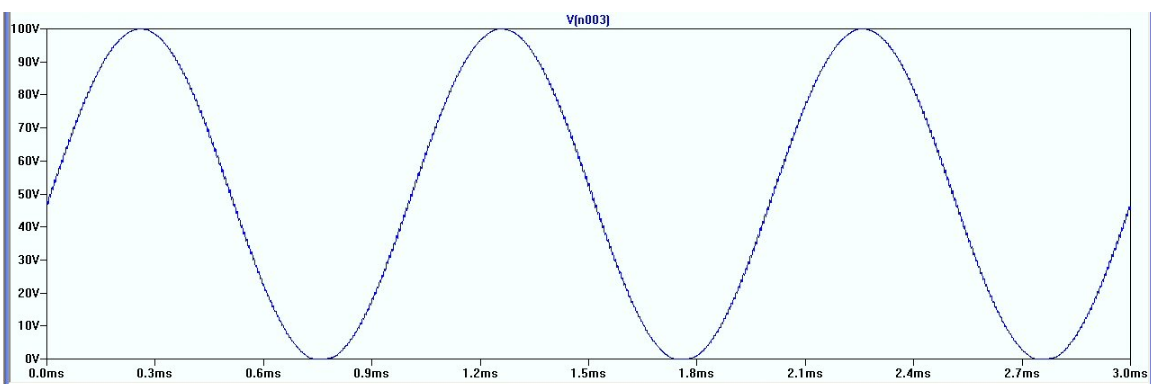

Fig. 46 - Voltage across C1

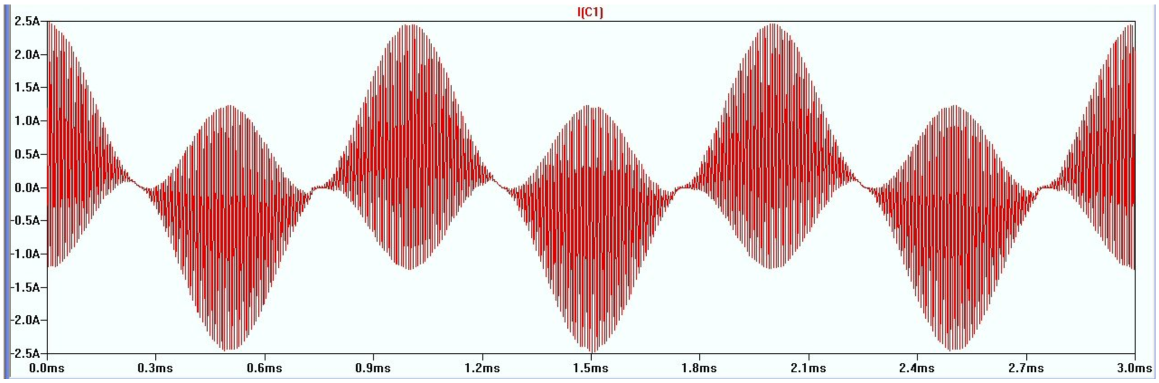

Fig. 47 - Current through C1

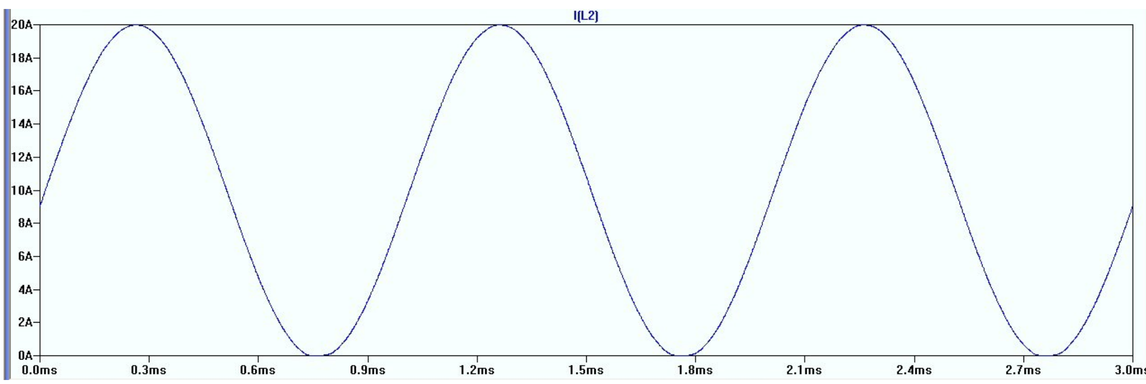

Fig. 48 - Current through L2

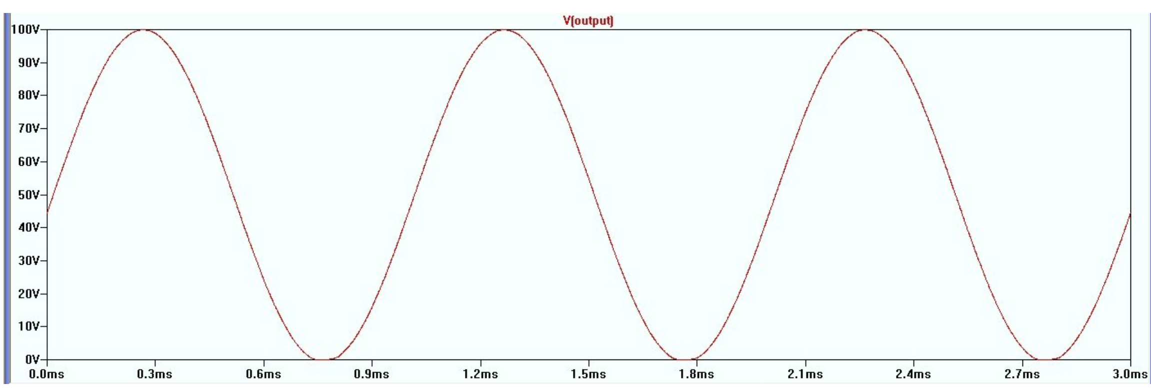

Fig. 49 - Voltage across C2 (the output)

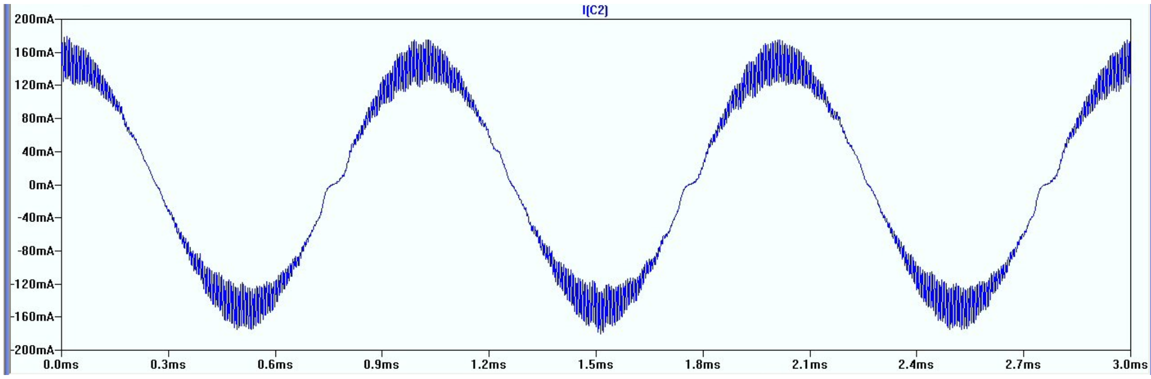

Fig. 50 - Current through C2

Arc Suppression

It is important to protect the RF amplifier should a problem occur somewhere in the system. If an arc occurs, for example, the amplifier should be shut down, the faster the better. Now we explore two ways this can be done. The first is to key off (remove) the drive to the RF amplifier as soon as an arc is sensed. Assuming it is well-behaved in that the amplifier does not go into self-oscillation, this will kill the arc. What happens if the drive is suddenly removed can be seen by examining Fig. 51.

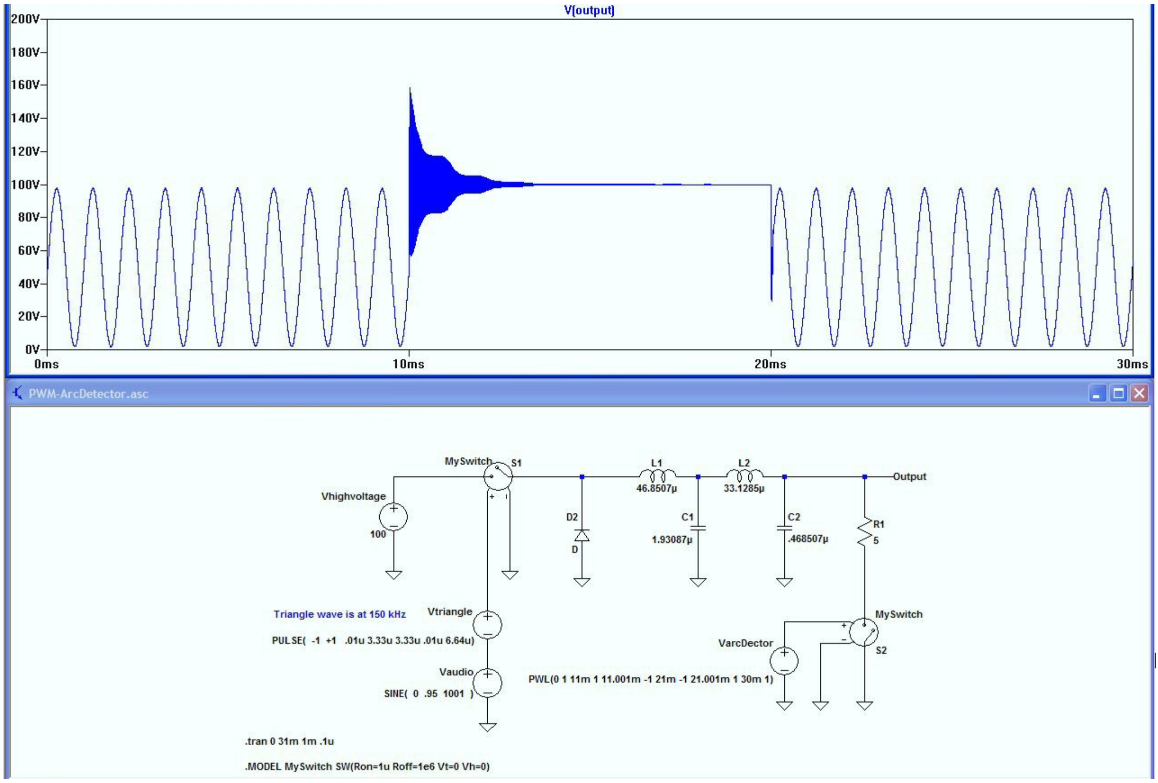

Fig. 51 - Output of filter when the load is suddenly removed

Should the drive to the RF amplifier be removed it stops drawing current and becomes what amounts to an open circuit. Fig. 51 shows the voltage at the filter output, the voltage applied to the RF amplifier. For analysis purposes we have keyed off the RF drive at the 10 millisecond point and immediately the filter output voltage soars up to 160 volts. That odd-looking "oscillation" that continues on for a few milliseconds is a ringing at the frequency determined by inductor L2 in conjunction with capacitors C1 and C2 in series. Those components form a loop and the resultant ringing frequency (in this case) is 45 kHz.

A clamping diode may be added to the output of the filter to stop that serious overshoot. That would also damp the ringing. Such a diode must be fast, its voltage rating must be that of the power supply and its current rating must be that of the load. The resultant output graphic is not shown here.

There is another method of calming down the RF amplifier after an arc occurs and that is to key off the modulator itself. This is illustrated in Fig. 52.

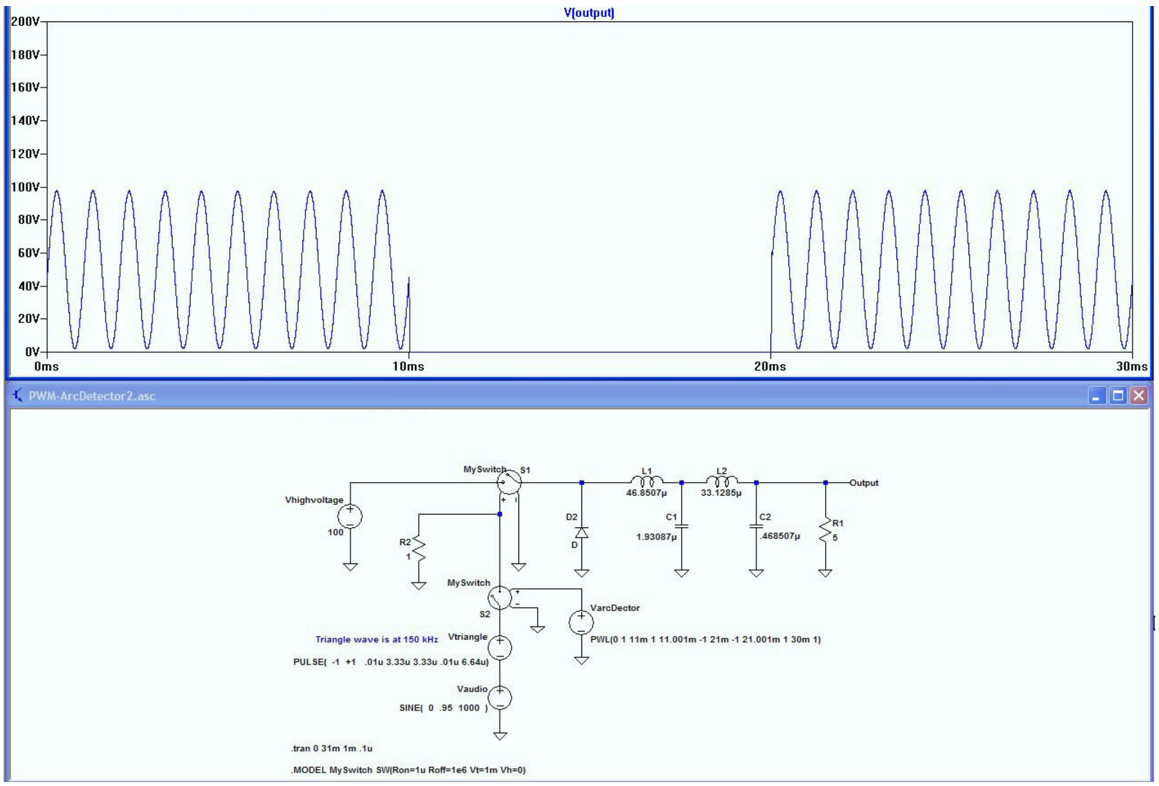

Fig. 52 - The filter output when the modulator is keyed off

We see in that graphic what the output of the filter (the voltage applied to the RF amplifier) looks like when the modulator is shut down. This is certainly a benign behavior.

The fall time of the filter output is about 30 microseconds, doubtless fast enough to quench an arc before it can cause trouble.

Fig. 53 shows the signal applied to the modulator. For circuit analysis purposes it is keyed off at 10 milliseconds and stays off for another 10 milliseconds. Examining that scheme reveals that the system is well-behaved regardless of where in the modulating waveform it is keyed off.

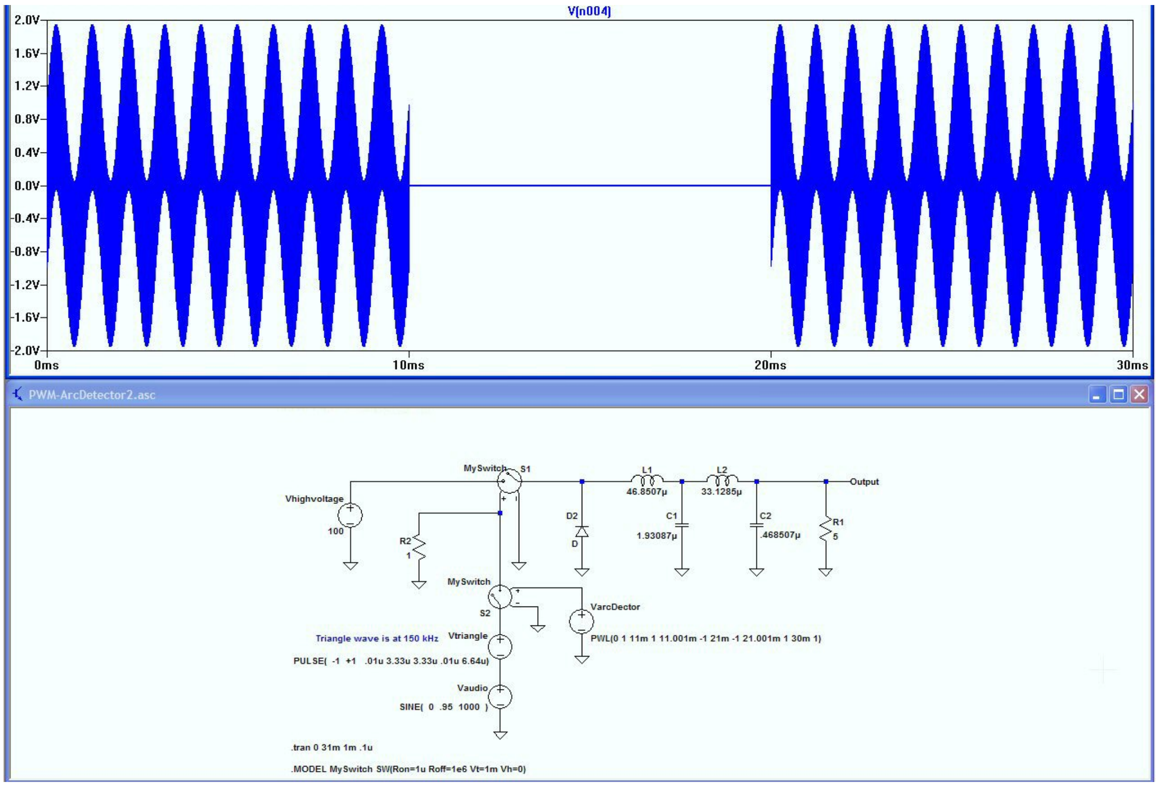

Fig. 53 - Input to the modulator, muted at 10 milliseconds

This should be easy to do and will result in a smooth "muting" of the system. After a problem is detected the modulator should be shut down.

It would be best if a timing circuit is used in the real-world circuit to keep the modulator "down" for some minimum amount of time, perhaps 100 milliseconds.

Other kinds of problems may also cause system shutdown. An example that comes to mind is over-current; another would be over-temperature. Such problems should be detected by appropriate sensors and cause shutdown.

Commercial transmitters have counters so that if a problem occurs too often within some time period the system is shut down and stays off until manually restarted. Consider that in your design.

Exploring a PWM IC

Linear Technology has placed on the market an IC that combines some of the critical items needed to generate the PWM waveform. Their part number is LTC6992. Some useful options are available with various "dash" numbers.

Figure 54 shows the LTspice rendition of a test circuit for this chip.

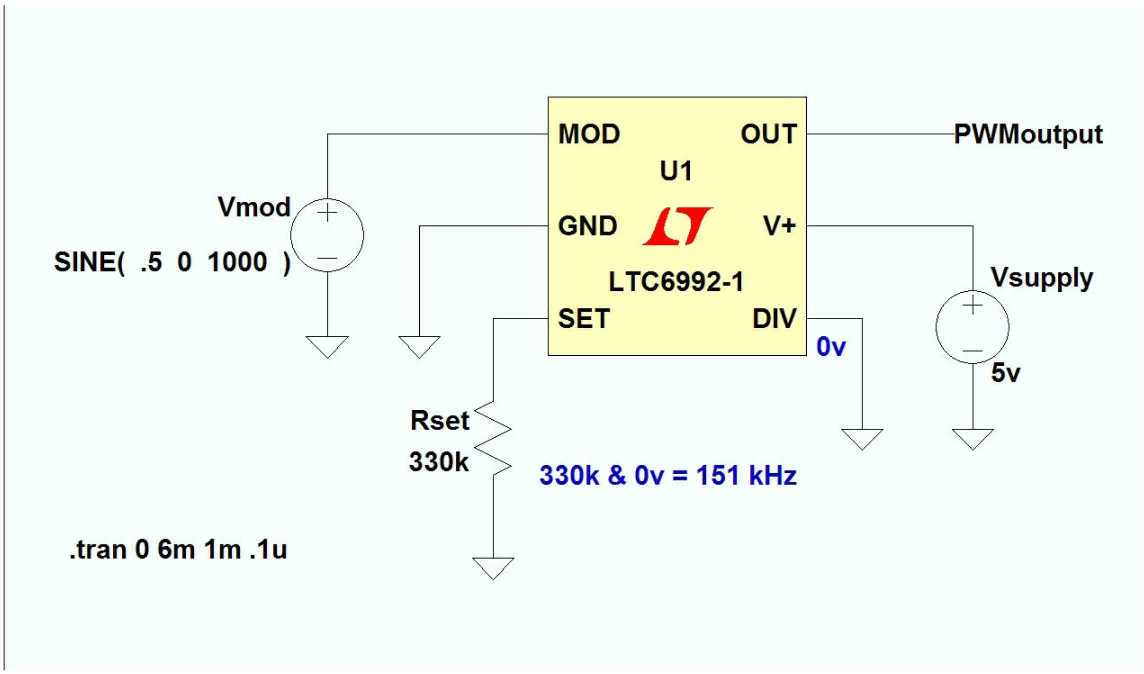

Fig. 54 - The LTC6992 analysis schematic

This IC has an extremely large switching frequency range, set coarsely by the voltage applied to the "DIV" pin and finely by the resistance to ground at the "SET" pin. The resistance from the SET pin to ground must be in the range of 50k minimum to 800k maximum. If those rules are followed the result will be a stable and amazingly high fidelity - excellent linearity - PWM generator.

Fig. 55 illustrates the LTspice schematic showing that chip driving a switch and lowpass filter to see how the LTC6992 handles modulation.

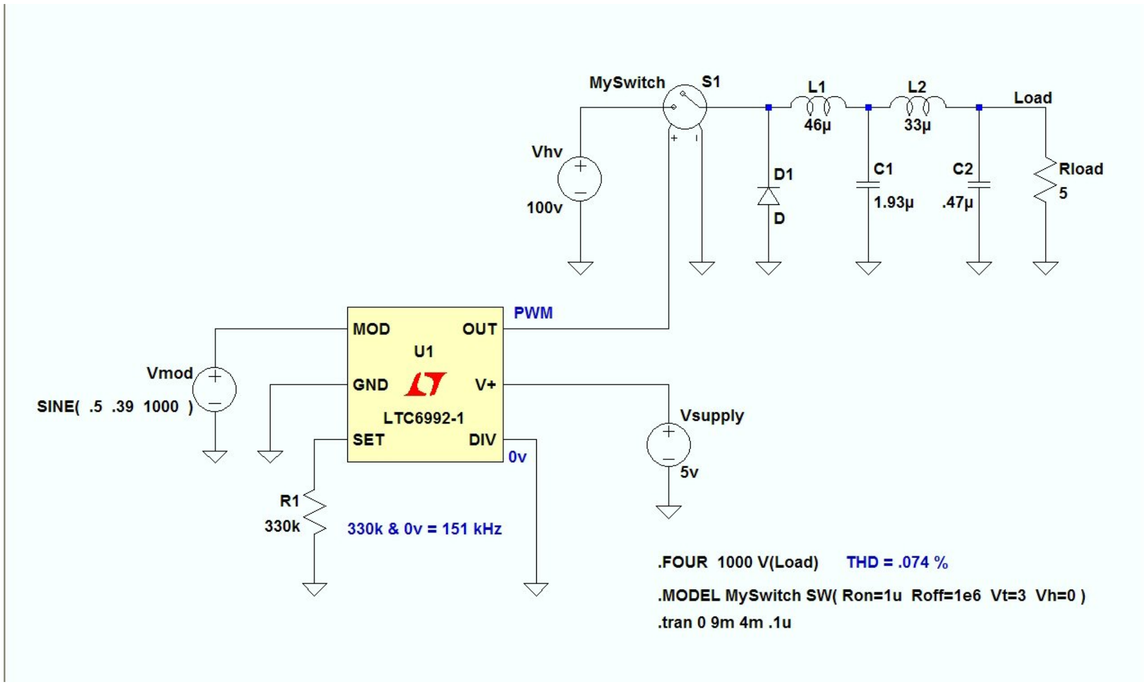

Fig. 55 - The LTC6992 with switch and lowpass filter

The DIV voltage has been taken to 0 volts (ground) and the SET resistor has been set to 330k ohms, resulting in a switching frequency of 151 kHz. With full modulation the total harmonic distortion measured at the load is 0.074%.

Other ICs are available and, as one of my college professors would say, "I will leave it up to you" to explore those chips.

Output Filter Input impedance

When we look into the input of the output filter designed using the old fashioned doubly-terminated design routine we get the input impedance plot as presented in Fig. 56.

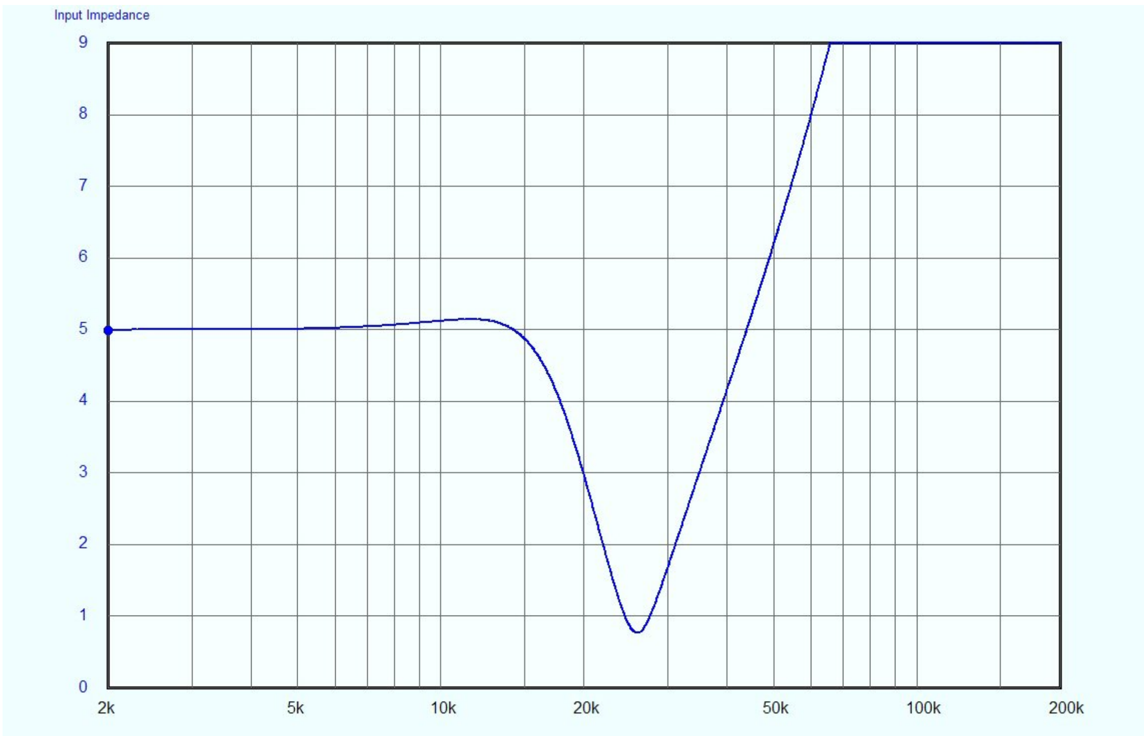

Fig. 56 - Impedance seen looking into the output filter.

Fig. 56 shows the input impedance of a Butterworth lowpass designed using the usual (doubly-terminated, not recommended) design method. The input impedance is rather constant up to about 15 kHz. It drops significantly above that and is nearly a short-circuit in the 25 kHz region (the design cutoff frequency). Attempting to drive this filter with modulation in the 25 kHz region might stress the switch. Fortunately we probably would not apply full-level modulating waveforms in that frequency range, particularly if the "polishing" (anti-aliasing) lowpass filter is used.

Now we will look at the input impedance of the filter designed using the singly-terminated (recommended) design method. See Fig. 57.

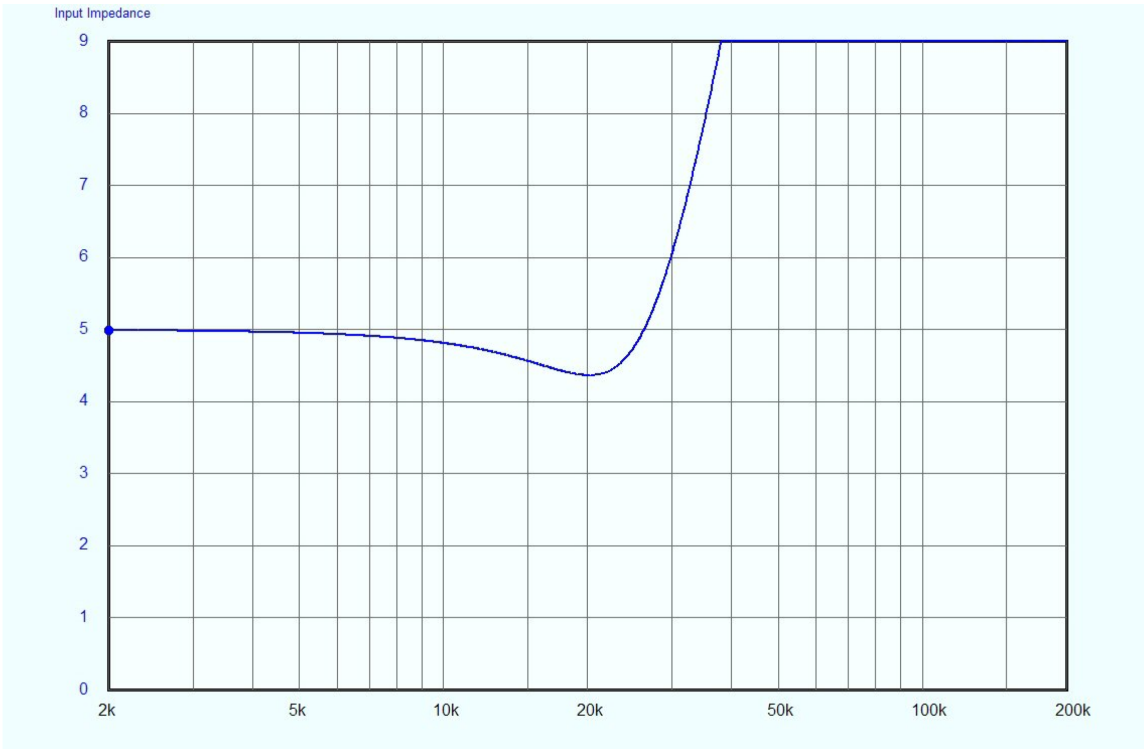

Fig. 57 - Showing the input impedance of the correctly-designed filter

Fig. 57 shows the input impedance of a Butterworth lowpass designed using the recommended (singly-terminated) design method. The impedance is rather constant up to about 15 kHz. Above that range the impedance rises, going in the direction of an open circuit. This essentially unloads the switch at those higher modulating frequencies.

Please note we are talking here about the impedance seen looking into the filter's input, using an "AC ohmmeter." This is a totally different thing from, and unrelated to, the input termination. The input termination is the impedance seen looking back into the source, in this case the switch.

Broadcast Transmitter - 1 kw

Let us now see how the commercial manufacturers of pulse-width modulated transmitters do (or did) things. For our guinea pig we are going to use a Collins 828C1, a.k.a. Continental Electronics 314R1. (The Collins broadcast line was bought out by Continental.) This transmitter is a one kilowatt carrier output unit and is popular in the amateur fraternity for conversion up to the 160 or 75 meter bands. Its post-switcher output lowpass filter schematic along with that filter's response can be seen in Fig. 58.

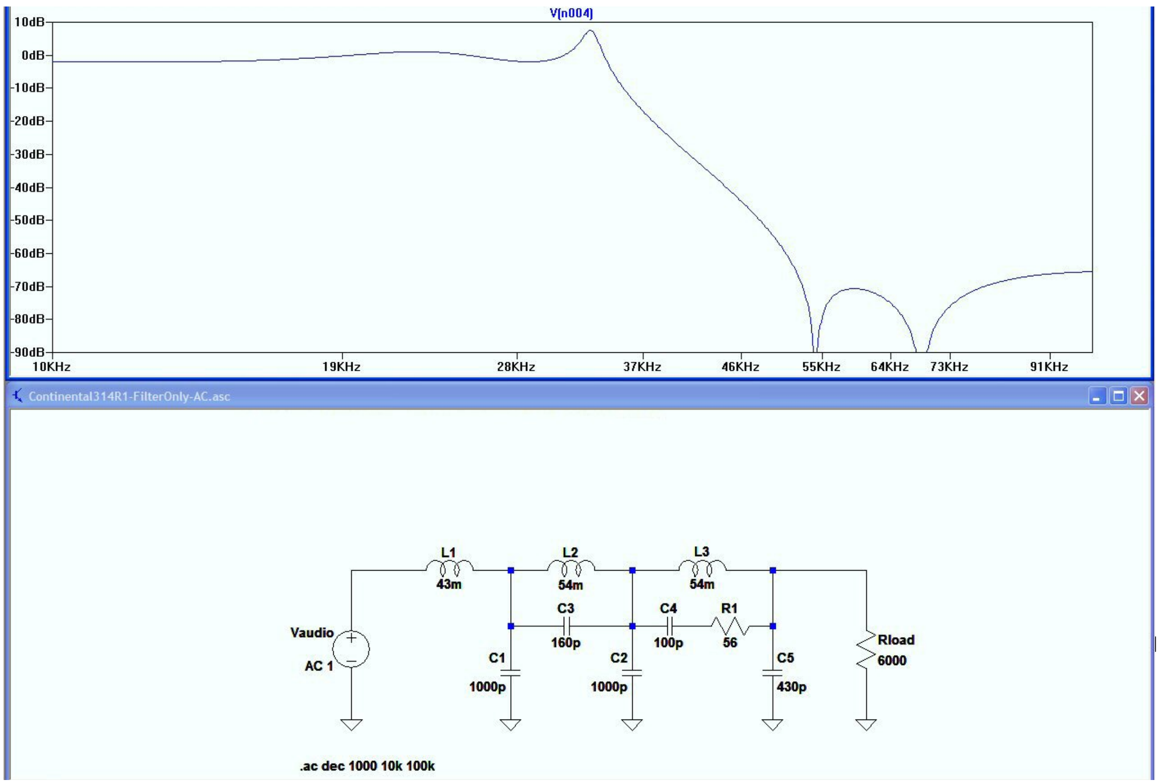

Fig. 58 - The modulator output lowpass for the 314R1

Two transmission zeros ("notches" in the response) have been added by the manufacturer. This is a common addition in commercial implementations of this modulation scheme, the object being to provide extra attenuation in the vicinity of the switching frequency. But notice that the response plot passband has a striking similarity to the earlier-discussed filter plots when the filter is designed as a doubly-terminated type but then it is actually run from a low-impedance source (not good). And as a result this filter has that awful peak in its response at about 30 kHz. And yes, it rings badly as a result. To follow up on a suspicion of how the filter was designed, *ye-scribe* added an input termination resistor to the analysis setup. The result was as seen in Fig. 59.

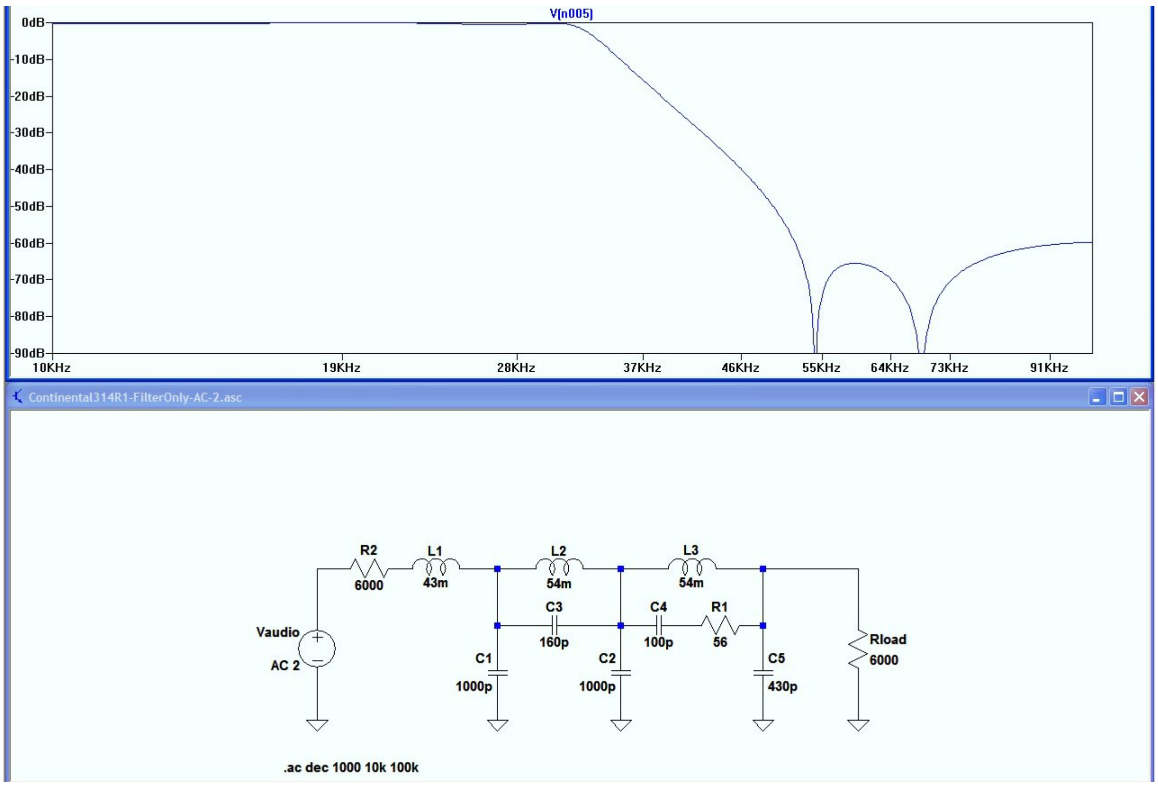

Fig. 59 - That filter with an input termination resistor added

In the above graphic, the response is much better behaved. The passband has a half dB of ripple. The ripple bandwidth is 31 kHz. It certainly appears that this lowpass started out as a sixth-order Cauer lowpass.

Having "a tiger by the tail," a filter was designed from scratch to see if the resulting design matched the above statements. A sixth-order Cauer with a ripple bandwidth of 31 kHz was used. Stopband frequency turned out to be 51 kHz with the stopband depth of 70 dB. The design impedance assigned was 6k on the input end and 10k on the output end. Strange, yes, but the plot and parts values then came close to the Collins/Continental "design."

This little exercise revealed that at least some of the broadcast equipment manufacturers had the same design tendencies (design that filter using a program that did the design using a doubly-terminated method but then operate the filter from a low source impedance). Also examined during this side trip but not shown here were commercial designs for 5 and 50 kw PWM transmitters from two different manufacturers. Each of those had the same design idiosyncrasies. Further, when a new model of transmitter was in the design stages, it appears that its modulator output lowpass filter was based not on design equations or software but instead on the *actual* component values used in an *existing* design, scaled to a new output termination impedance. So over multiple generations the filter "design" could drift. Why did they do this? Because "it worked" and was considered "good nuff."

Broadcast transmitter - 500 kw

But not *every* transmitter designer used odd or incorrect design techniques. From my archive *ye scribe* unearthed the design for a quite high-powered transmitter. This was intended for shortwave transmission, where lots of power seemed (notice past tense !) to be quite common. I gave a casual look at that paperwork and then used the singly-terminated designer in Elsie and was quite pleased to see that the resultant design and the archival designs were astonishingly similar. I had guessed correctly at the 3 dB bandwidth, which did not show up on those archival papers.

This transmitter was a vacuum-tube unit that operated at an idling plate voltage of 13 kv and a plate current of 52 amps. This gave the PWM output lowpass a termination value of 250 ohms. The RF amplifier operated at an efficiency of 74% and delivered a carrier output power of 500 kw.

The filter design from the archive is shown in Fig. 60.

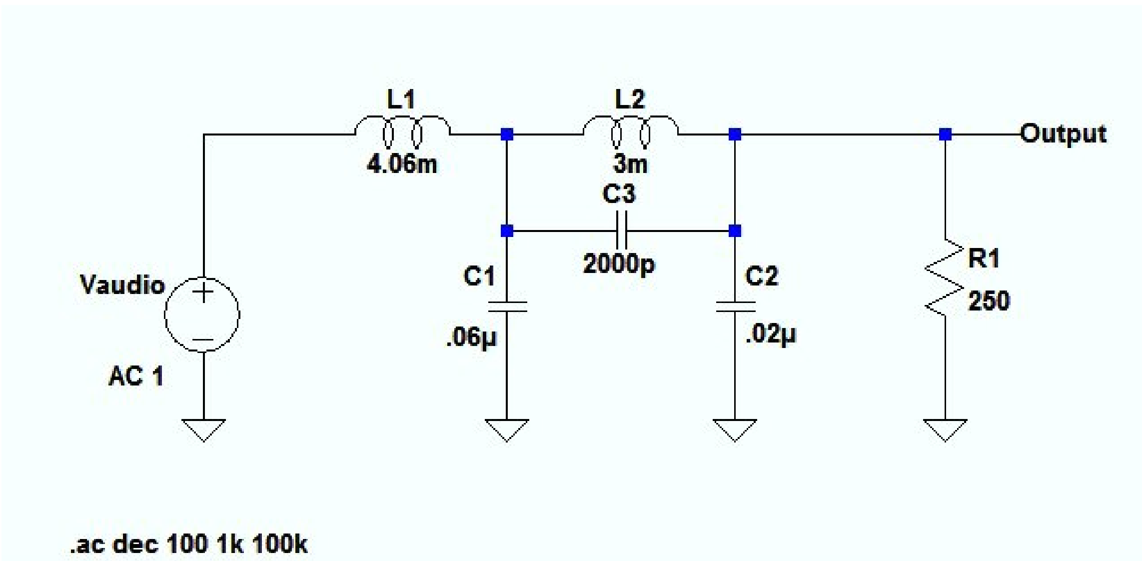

Fig. 60 - Schematic of the big transmitter's lowpass filter

Note the capacitor added in parallel with the second inductor to place a transmission zero ("notch") at 65 kHz, the switcher frequency for this transmitter. Switcher components are thus suppressed additionally in the vicinity of that frequency.

The response of this filter is shown in Fig. 61.

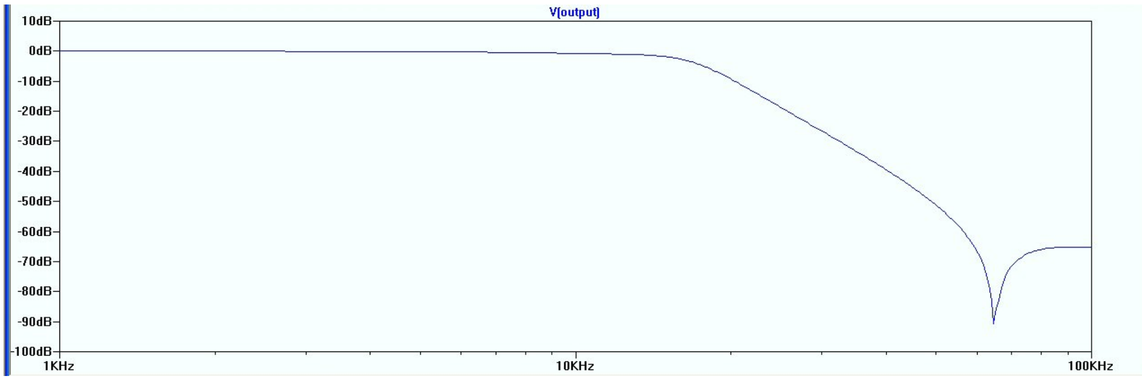

Fig. 61 - Response of that filter showing the added transmission zero

The 3 dB bandwidth of this post-switch filter is right at 15 kHz.

That filter's response to a squarewave is seen in Fig. 62.

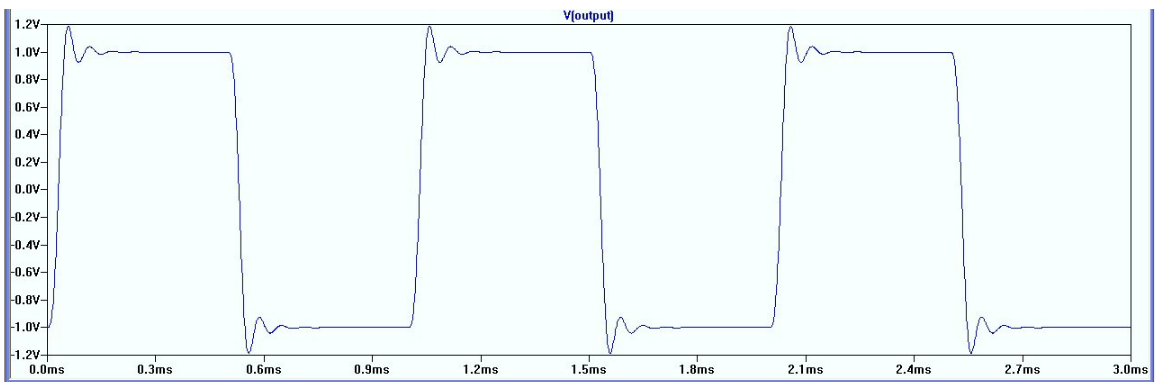

Fig. 62 - The transient response of that filter

This filter is rather well-behaved.

Transmitted signal bandwidth control

Earlier in this writeup it was pointed out that the purpose of the post-switch lowpass filter was to pass DC and modulation with minimal problems while also attenuating the switcher trash significantly. That output lowpass, if designed according to the recommendations posted here, will have a flat magnitude (frequency) response in the audio passband and will also have only a modest degree of overshoot added to a clipped modulating waveform.

But it will not determine the transmitted signal bandwidth; that is not its purpose.

Let's take a look at how the audio frequency response can be (and should be) be controlled prior to the audio signal's application to the modulator proper.

This response-shaping is done early in the audio amplifier chain of circuits, commonly called the "speech amplifier" or "speech processor." It is best done prior to the compressor and absolutely prior to a clipper which follows that compressor. It is intended to minimize those speech components that will cause mischief later on.

A modest low-end rolloff is recommended. Use a simple highpass filter, rolling off at 6 dB per octave with a -3 dB point of 300 Hz. The words "small coupling capacitor" come to mind.

Following that simple highpass element in the speech chain we should install a lowpass filter. That filter should be flat out to 3 kHz (or perhaps 4 kHz) and then roll off at perhaps 30 dB per octave. Some designers may choose a somewhat higher cutoff frequency. This kind of response-shaping is easily accomplished using opamp circuitry.

Fig. 63 shows the schematic of a nice lowpass filter to be installed prior to the speech AGC / clipping circuitry.

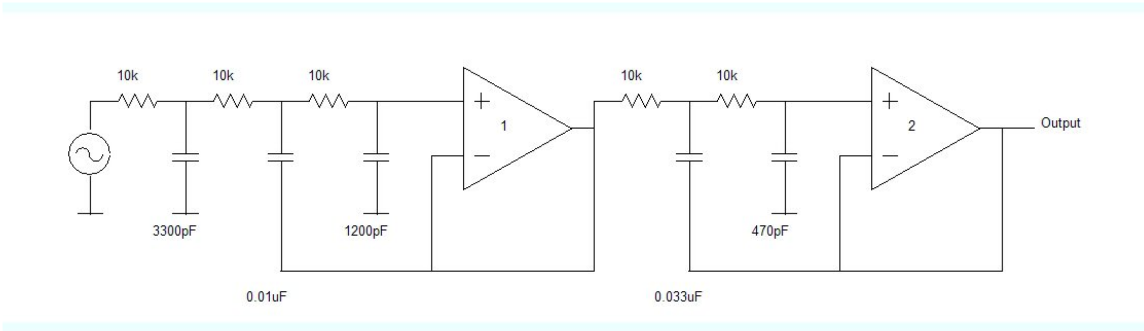

Fig. 63 - The schematic of a lowpass filter for the speech chain

The response of that filter is shown in Fig. 64. As can be seen, it has a modest rise in the response up to about 4 kHz, yielding a slightly "crisp" response. We have, in one rather simple circuit, both a modest treble boost in the area that it is needed along with a lowpass to assist in confining the transmitted signal bandwidth.

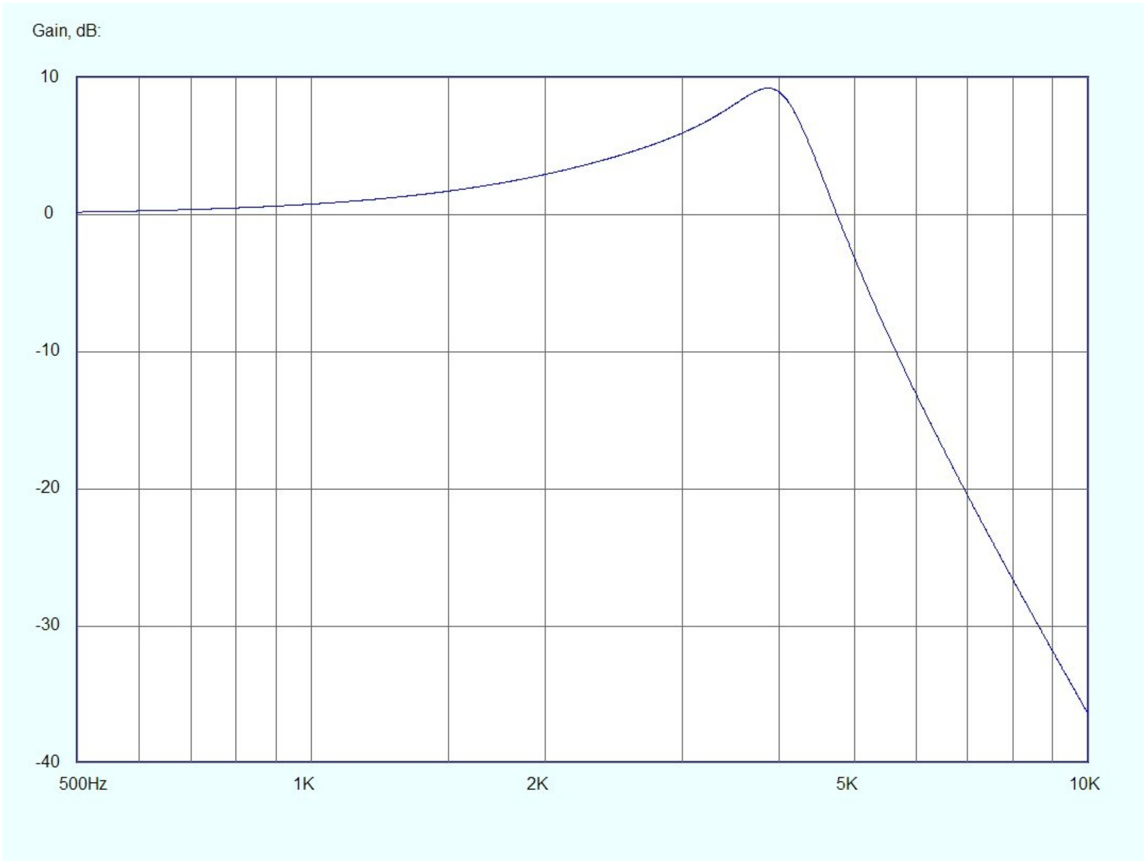

Fig. 64 - The response of the filter shown in Fig. 63.

For those designers who would like the schematic of a "flat" audio bandwidth-limiting filter, see the design in Fig. 65.

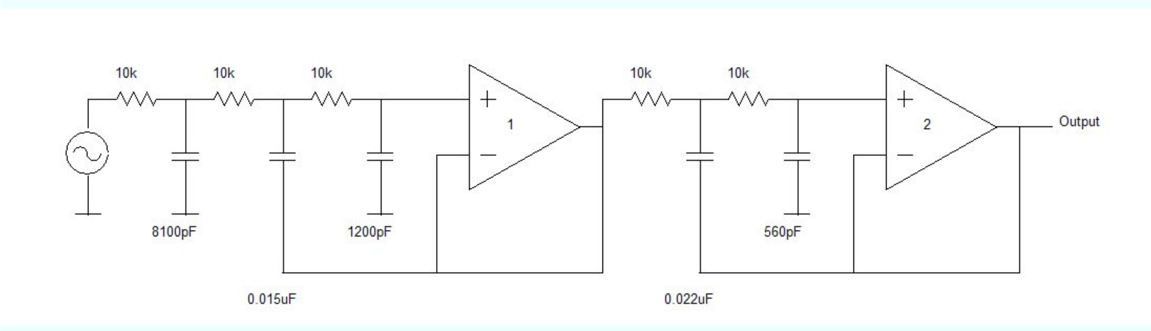

Fig. 65 - Schematic of a "flat" lowpass filter

The design in Fig. 65 is for those who want a more ideal (flatter) frequency response. The response of this filter is shown in Fig. 66.

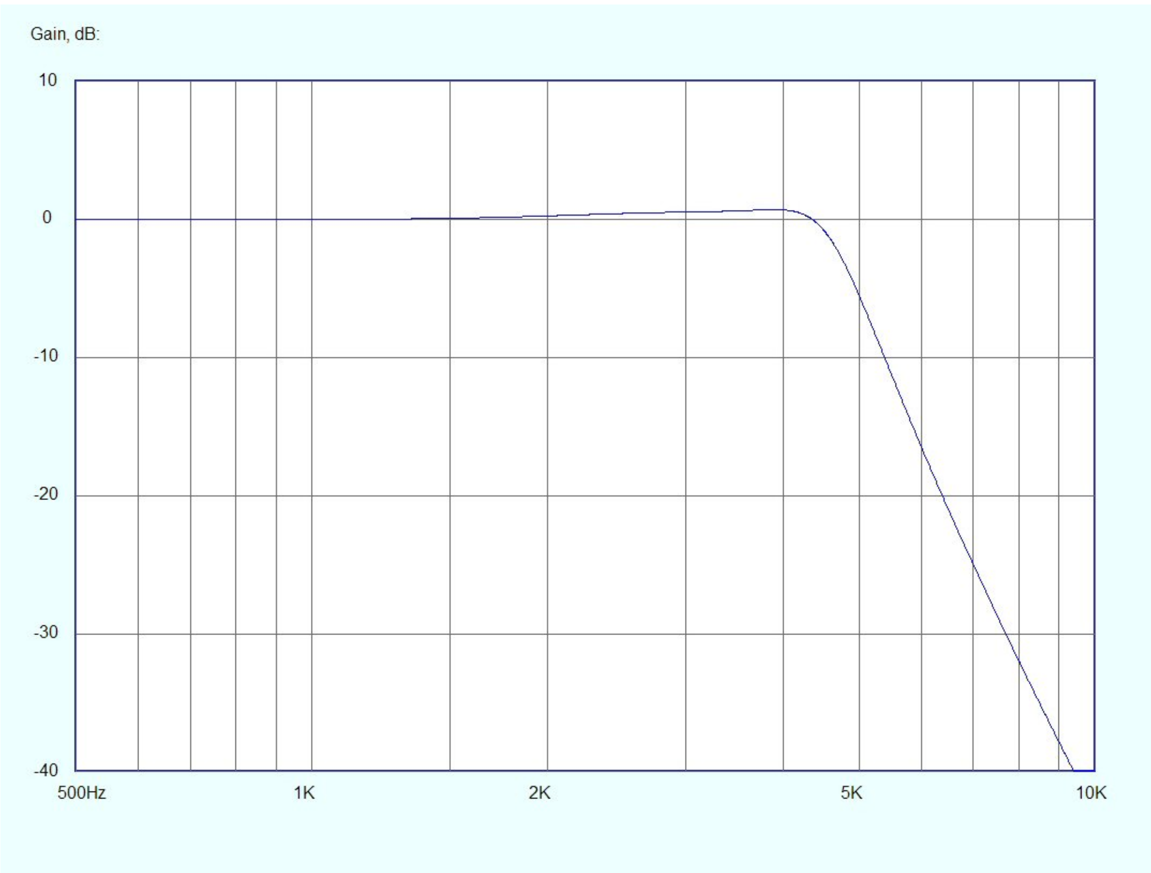

Fig. 66 - Response of the filter in Fig. 65.

As seen, the response of this filter is flat to about 4500 Hz and down over 40 dB at 10 kHz. Beyond that, the ultimate rolloff rate is 30 dB per octave. This lowpass filter can nicely determine the transmitted signal bandwidth.

We are not in the least bit concerned about the transient response of these low-level speech-processor filters. At this point in the audio chain how such a band-limiting filter handles squarewaves, for example, is of absolutely no concern. But **after** the low-level clipper, the transient response of each of the following stages is of **great** concern. In fact, after the clipper, all of the following circuitry should be direct-coupled (or come very close to it) so that low-frequency audio waveforms, if clipped and so have a "flat top," are passed without tilting those flat tops. And **after** the low-level clipper any filtering used must not cause overshoot on applied squarewaves.

These low-level audio filters are intended to be used prior to the speech AGC (compressor) system. They allow rejection of the sibilant frequencies (10 kHz and above, the "s" sounds) which should not be transmitted, certainly can't be received, and serve only to activate the audio AGC (compressor) system, reducing the modulation level. *Ye scribe* considers the sibilant frequencies as nothing more than mischief-makers in communications applications. It is easy to change the bandwidths of the designs as presented: simply scale the values of the resistors that are shown in the schematics. As an example, to move their frequency values up an octave, halve **each** of the resistor values; leave the capacitor values intact.

Or use the software that did these designs in the first place. See the links below.

Antialiasing

In those schemes that involve "switching" of a signal, there is a byproduct of that switching operation that can cause mischief. An example of this is in FM broadcast stereo transmission. That system, which has been globally accepted, is what amounts to a switching scheme which selects for transmission first the left channel and then the right channel, at a rate of 38 thousand times per second. In the receiver the decommutation circuitry then "distributes" the composite (as it is called) first to the left loudspeaker and then to the right. Relatively incidental to that process is a pilot tone transmitted at half the switching frequency and at a low level and whose primary purpose is to keep the two switches (transmit and receive) in synch. [Secondarily this pilot tone turns on a small lamp in the receiver. . .] The pilot tone's frequency is doubled in the receiver and the resultant signal is used to decommutate the composite. But this switching scheme is susceptible to being upset by unwanted components applied to the audio inputs of the generator. Modulation-signal components in the 23 to 53 kHz region would be decommutated back down to the zero to 15 kHz region, quite audible. This annoying process is called "aliasing." To prevent this problem, the stereo signal generators use a lowpass filter in the audio chain to remove components that might cause that aliasing. These "anti-aliasing" lowpass filters, one in each audio channel, pass DC to 15 kHz and reject all audio components at 19 kHz (to protect the pilot tone region) and above (to prevent aliasing).

We should use a similar lowpass filter in the PWM generator to prevent modulation components in the switching frequency region from causing mischief. Supersonic modulation components would be translated back down into the audible range. Such an anti-aliasing lowpass filter is generally installed right at the input to the PWM generator. A study was made to define its requirements in this application.

To make our mental arithmetic easier, and because the readership will be using a variety of switching frequencies, the initial analysis was done using a switching frequency of an even 100 kHz. Test tones, a series of them, were applied to the system at full modulation level and the aliasing components were observed. Then the magnitude of the test tone at each test frequency was adjusted to result in aliasing components 60 dB down from full level. In this test setup, the aliasing components were down 60 dB or better until the test tone was raised up to about 15 kHz. As the frequency was raised above that frequency, its magnitude had to be reduced to keep the aliasing component(s) at -60 dB re full modulation. This is shown in Fig. 67.

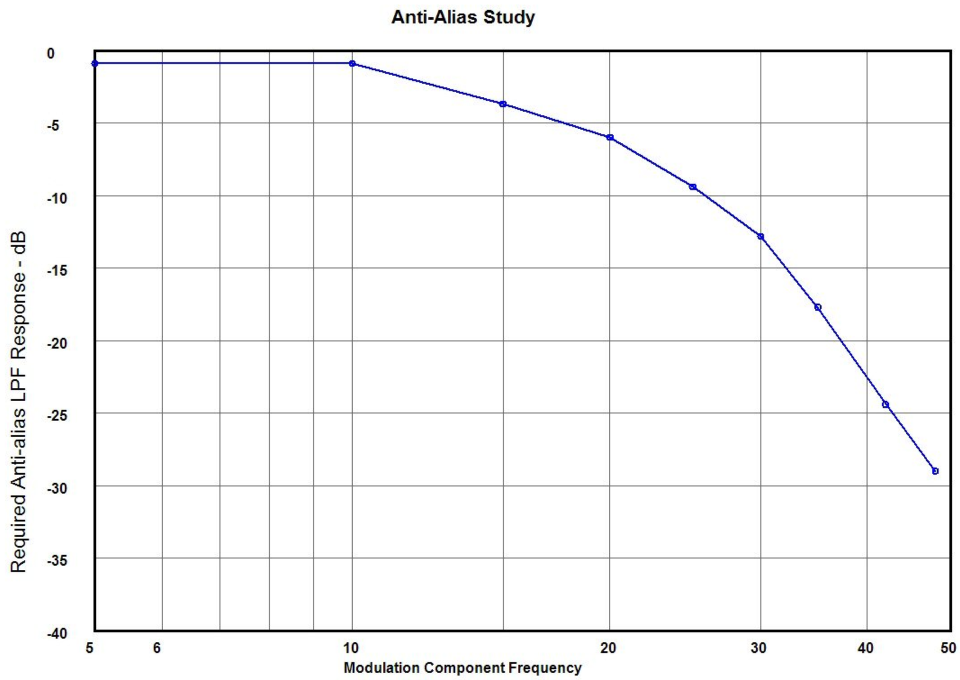

Fig. 67 - Required signal-level reduction needed to prevent aliasing

The modulation components have to be reduced in amplitude as the test frequency goes up if aliasing is to be avoided. But there is no need to reduce their magnitude in the "audio" range. If a lowpass filter is used to reduce their amplitude in the supersonic range (and that is what we are leading up to), that filter can have a flat response in the important audio region.

In Fig. 67, the frequencies shown on the X-axis are percent of the switching frequency because the test run was done using a switching frequency of 100 kHz.

The next step in this excursion is to determine the order required for this anti-aliasing lowpass. Calling on *Elsie* and using its *Limits* routine we generate the graphic shown in Fig. 68.

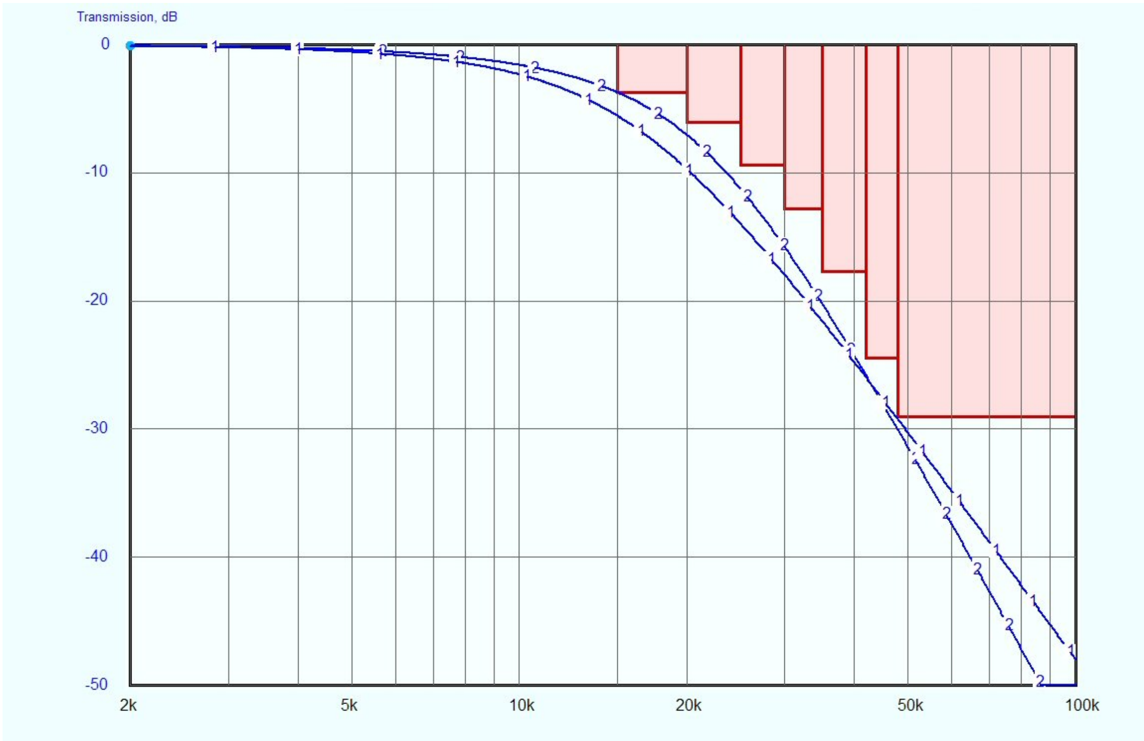

Fig. 68 - Graphic to determine a satisfactory Bessel filter order

In Fig. 68, the red areas indicate the levels to which a component must be reduced in order to keep the aliasing components at 60 dB below full modulation. The "corners" were the frequencies and amplitudes shown in Fig. 67.

To pass squarewaves nicely we have to use a Bessel family filter to avoid overshoot should the modulating signal be complex (for example a squarewave). Both third and fourth-order filters will suffice in this position. In Fig. 68 the plot labeled "1" is a third-order filter with a 3 dB bandwidth set to 11.4% of the switching frequency. The plot labeled "2" is a fourth-order design with its 3 dB bandwidth set to 13.7% of the switching frequency.

Either of those designs will keep the aliasing components at 60 dB below full modulation **when a single sinusoidal modulating tone at full modulation level is applied**. It should be obvious that such an extreme signal level at those supersonic frequencies is quite unlikely in practice. It follows that these anti-aliasing lowpass filters are in effect somewhat over-designed. The aliasing component magnitudes will be very low level indeed in a real-world system.

The above study was done using a switching frequency of 100 kHz so the the results, frequency-wise, are easily normalized to another switching frequency. Then a design was done for a common real-world switching frequency of 150 kHz. The response plot is shown in Fig. 69.

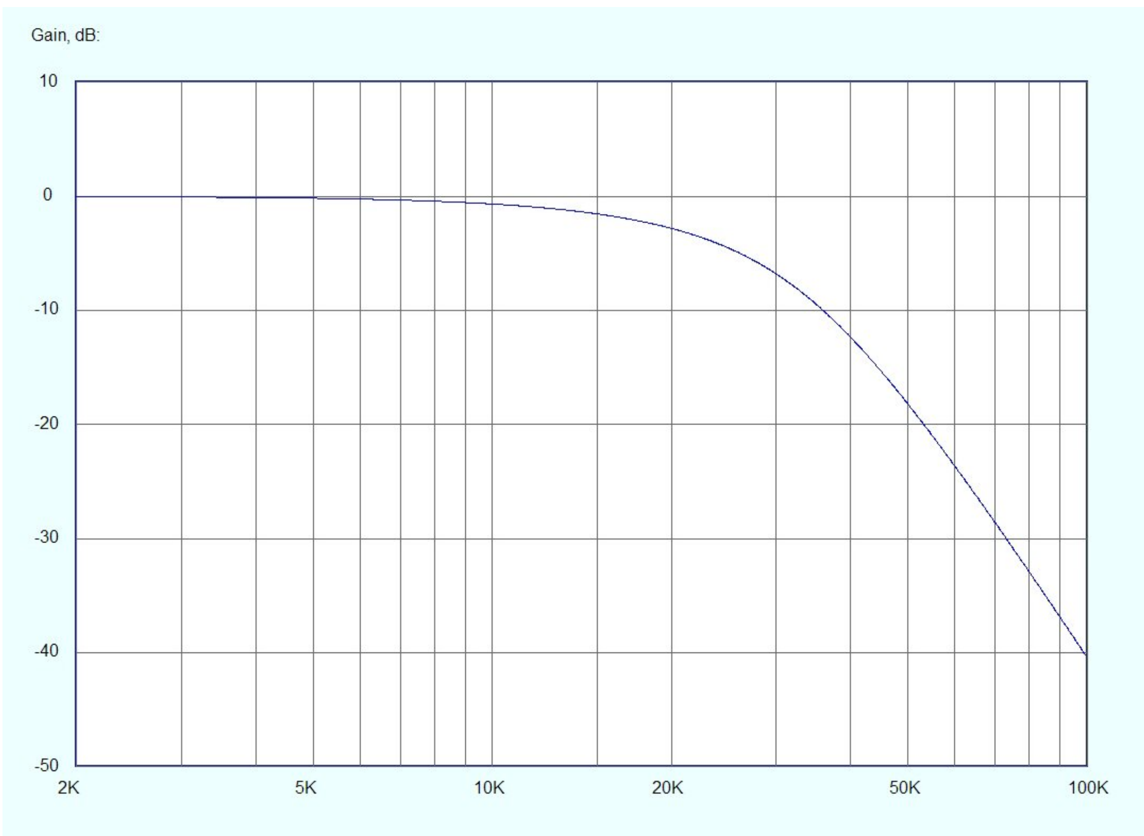

Fig. 69 - Response plot of the anti-aliasing filter for a 150 kHz system

This is the response of a satisfactory anti-aliasing lowpass for a system with a 150 kHz switching frequency. This filter handles transients well. The transient response plot is shown in Fig. 70.

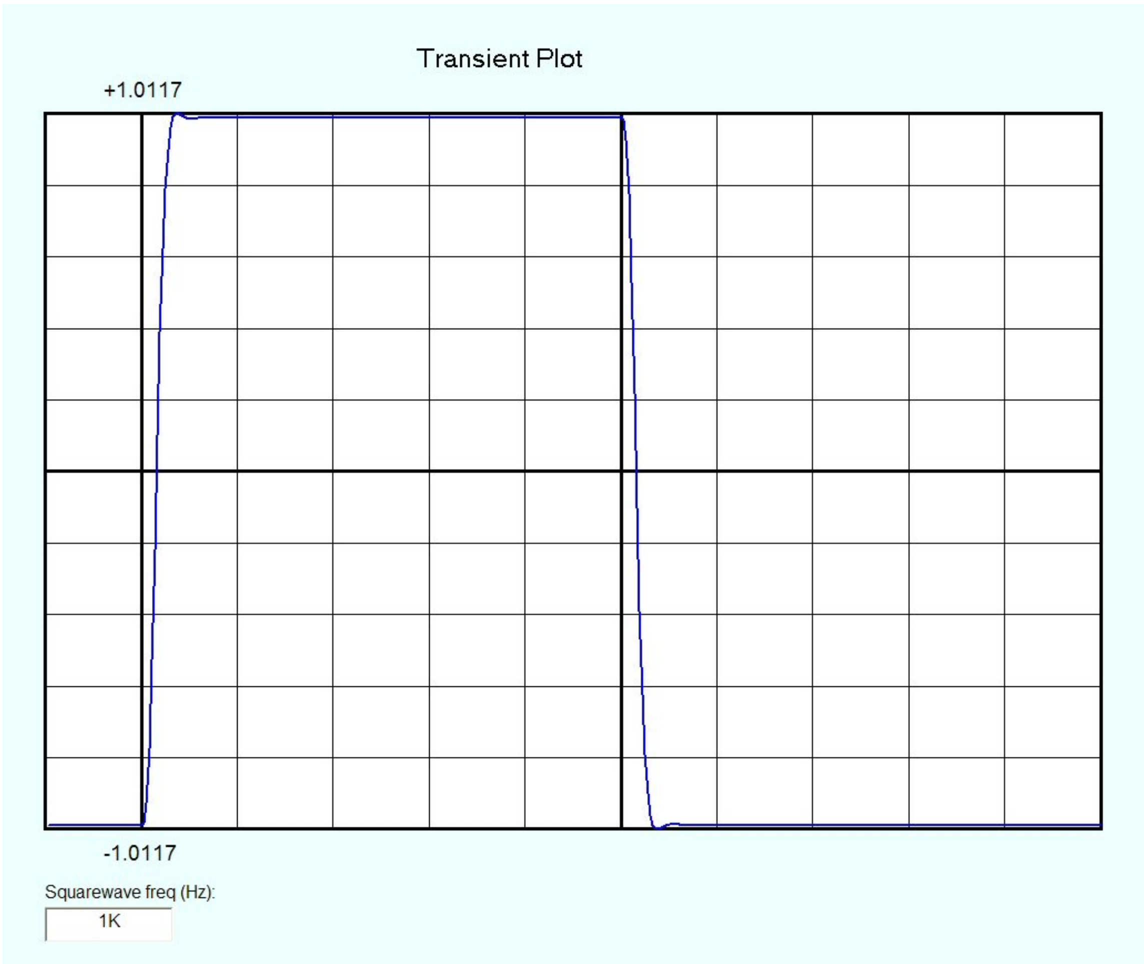

Fig. 70 - Transient response of the antialiasing lowpass

Squarewaves are passed essentially unaltered except for bandwidth limiting.

The schematic of this filter is shown in Fig. 71.

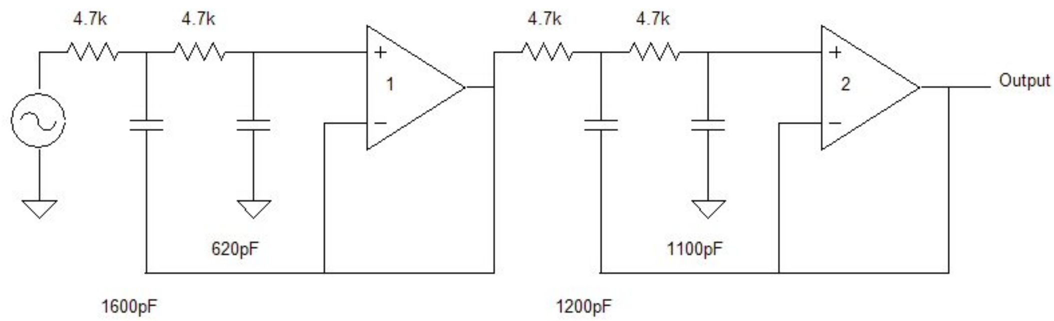

Fig. 71 - Schematic of the anti-aliasing lowpass

This filter is for a system using a 150 kHz switching frequency. For other frequencies simply scale the resistor values. Resistor values of 10k ohms will work nicely for a system using a 75 kHz switching frequency, for example.

Post-Clipping Lowpass

If a clipping circuit is used in the "speech processor" then it really should be followed by a lowpass filter to remove the higher-order harmonics that would cause "splatter" in the transmitted signal. If you don't follow the clipper with a bandwidth-limiting lowpass, you might just as well remove the clipper and go ahead and overmodulate. Get it over with. Simplify your circuitry.

But we have a problem here. That lowpass should have a sharp descent into the stopband; it should have a sharp cutoff. We want to define the bandwidth in no uncertain terms. But those kinds of filters (sharp cutoff) by their very nature normally have significant overshoot (ringing) on applied transients.

There is a solution here but it has a modest peculiarity. The solution is to add a "step" in the lowpass response. This changes the overshoot resulting from a sharp cutoff to an "undershoot" which is harmless. A proposed filter's magnitude response is shown in Fig. 72.

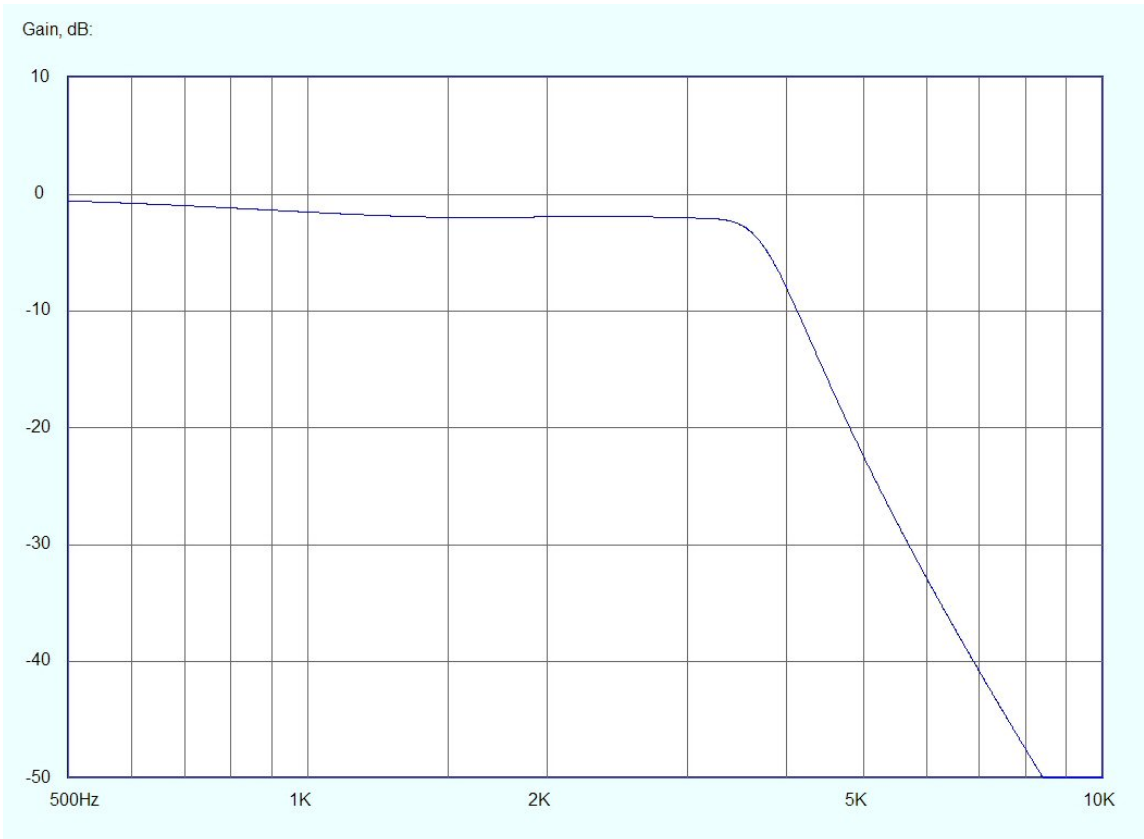

Fig. 72 - Response of proposed post-clipper lowpass filter

This filter, when placed after a clipper, passes the main audio frequencies while attenuating the components that would cause splatter in the transmitted signal.

This lowpass does not introduce overshoots on applied squarewave inputs. And clipping can certainly introduce squared waveforms.

The transient response of this lowpass is shown in Fig. 73.

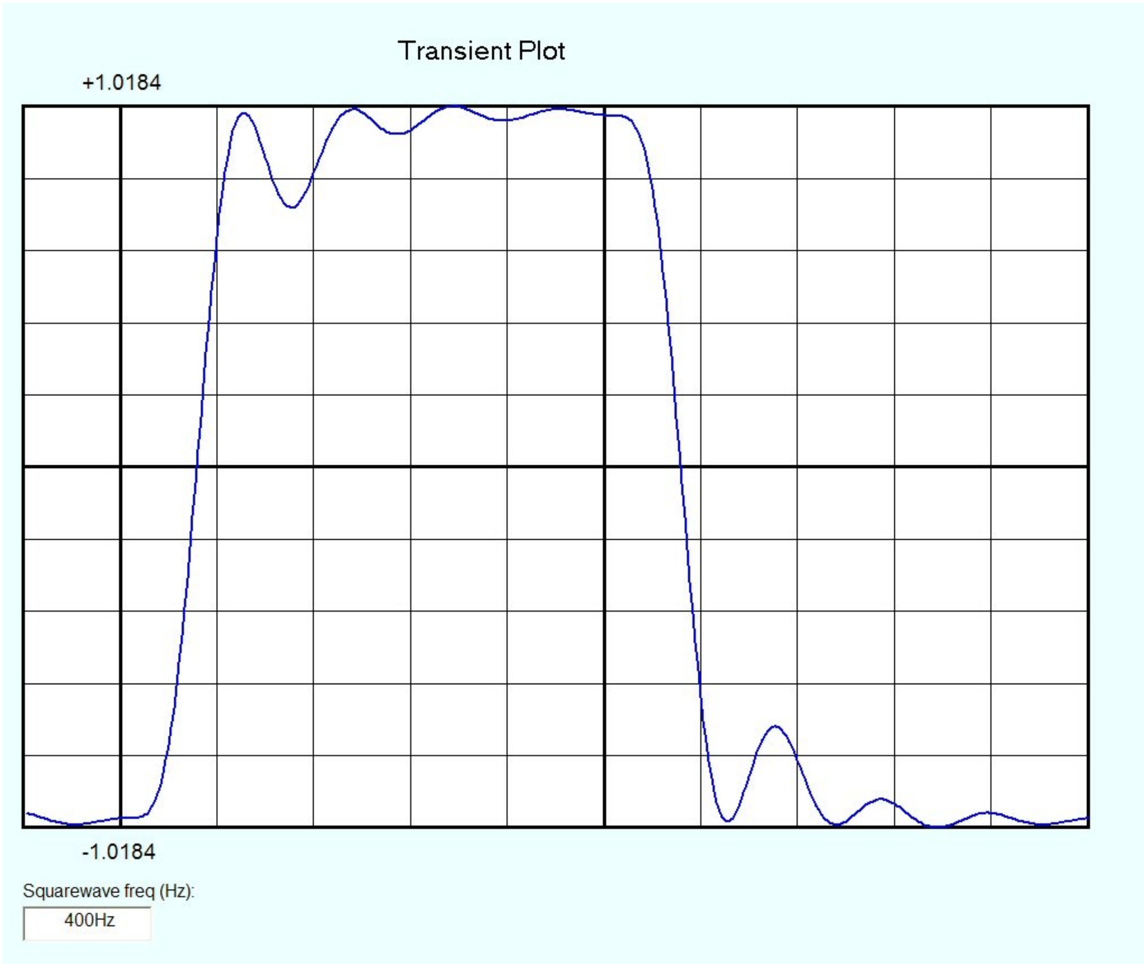

Fig. 73- Transient response of the suggested post-clip lowpass

The overshoot that would have ordinarily resulted from severe clipping has been instead turned into what might be called "undershoot." Looking at Fig. 72 shows the response at the upper audio frequencies to be reduced by about 1.8 dB when a sinusoid is used to measure the response. But when a trapezoidal wave is used for the test signal, full modulation is achieved right up to the filter's bandedge.

If you wish to implement this lowpass the schematic is shown in Fig. 74.

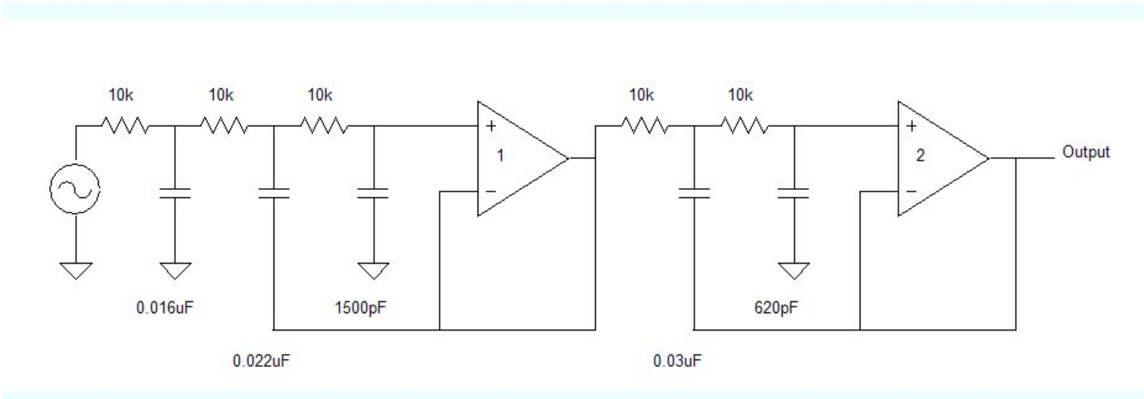

Fig. 74 - Schematic of the "undershoot" lowpass

This post-clipper lowpass filter will determine the transmit bandwidth without question. The assumption here is that all of the following circuitry is linear and does not introduce distortion.

Please notice this: if a clipper is used to set an absolute ceiling on the modulation on an instantaneous basis and then a post-clipping lowpass is used, be advised that there is no need for an anti-aliasing lowpass because there will be no modulating components up in the region where aliasing occurs. The post-clipping lowpass serves as an anti-aliasing lowpass filter.

=====

Where do we go to get a nice active lowpass filter designer with analysis ?
http://www.tonnesoftware.com/downloads/ActiveLowpassInstall208.exe

Where do we go to get LTspice?
http://www.Linear.com/LTspice