Making A Software Defined Radio for the QRP Enthusiast: Part IV

Ward Harriman—AE6TY ae6ty@arrl.net

Those of you who have been following this series of articles know that the description of the "SDR" portion of my project is now complete. At the conclusion of Part III of the series, a fully functioning SDR receiver and transmitter had been described. However, the output amplitude of the transmitted signal was too small even for most QRP users. So, a transmitting amplifier matched to the SDR was in order and will be described here.

I am an active home brew enthusiast and as such, I've read through a large number of articles and built many circuits described by other designers. Many times I've asked myself, "But how did he get here?" Seldom does anyone bother to describe the actual design process. In this paper I hope to show not just the result of the design process but also 'HOW' I got to the endpoint. My goal is not to provide a complete circuit you can just build, but to show you the tools necessary to design your own.

This paper chronicles the design of a Class E, 20 meter, QRP transmitter. This was my very first power amplifier and I got started with a great deal of trepidation. During the design process I found that the class E amplifier is very simple to understand and design and extremely easy to make work correctly. I hope you will find the following enlightening and empowering. There is nothing quite like designing, building and operating your own brainchild.

In the design process, I use four free tools from the internet: "ClassE" from Tonne Software, LTSpice from Linear Technology, Elsie from the ARRL software toolkit and "SuperSmith" from Tonne Software. I highly recommend these tools to anyone actively designing circuits in the HF bands.

The Power Stage

Class E amplifiers are closely related to class C amplifiers. As such, they are generally non-linear and can be used best for CW and FM operations. As with class C amplifiers, class E amplifiers use a transistor as a switching element which is either completely 'on' or completely 'off'. The major advantage of the class E is the higher efficiencies achievable with proper tweaking.

The design of a class E amplifier starts out with a free tool from Tonne Software called "ClassE". This software can be downloaded from www.tonnesoftware. com. This software is extremely easy to use. Figure 1 is a screenshot of the starting point of this design effort:

The six key parameters to be fed into the ClassE tool are quite straightforward. The frequency and supply voltages can be chosen as the designer desires; here 14 MHz and 12 volts. Early in the design process, the other parameters can be guessed. Each of these parameters is now discussed.

The Saturation Voltage is the 'on' resistance of the switching transistor times the current flowing while the transistor is turned on. I knew that I was going to use a MOSFET and that these devices have on resistances of 10s to 100s of milliohms. As can be seen on the screenshot above, the Ipeak current is expected to be around 1.25 amps. Thus, the saturation voltage can be expected to be well below the pessimistic 1 volt I chose.

As expected, the Power parameter can be adjusted at will. I chose 5 watts as a useable power as I generally run QRP.

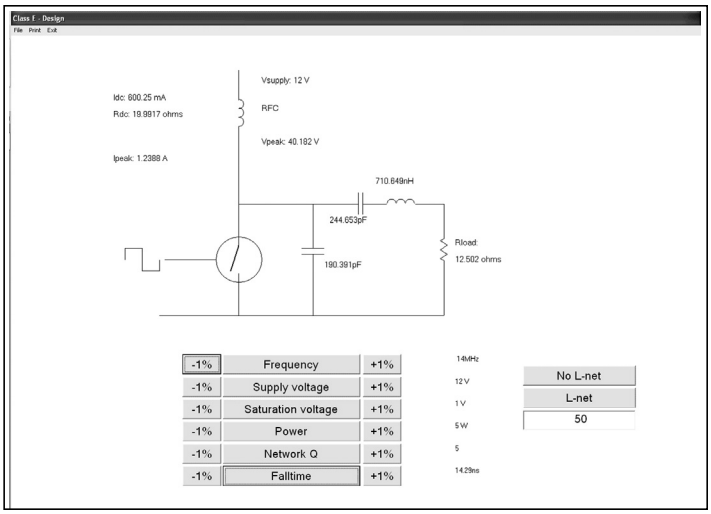

Figure 1—Data entry page for the ClassE software.

Interestingly, changing the power requirements affects only the Rload and the Ipeak values. The Falltime represents the speed of the transistor. I left this parameter alone.

Network Q is the final parameter. This parameter is useful for adjusting the value of the output inductor; larger Q results in larger inductors and smaller capacitors. A Q value of 5 is typical and I left it alone.

Having made a first stab at the power stage of my class E amplifier I decided to go do some simulation. For analog simulation, I use another free tool from Linear Technology called LTSpice. This tool is available to download from www. linear.com. Figure 2 is a screenshot of the circuit entered into LTSpice.

The circuit of Figure 2 may have some unfamiliar symbols, specifically the V1 and V2 voltage sources and the FDS3580 MOSFET transistor. V2 is a simple 12 volt supply voltage. The BSH114 MOSFET is a model provided by Linear Technologies and widely used in switch mode power supplies. It is representative of the transistors I would expect to use in this amplifier. V1 is a little more complicated…

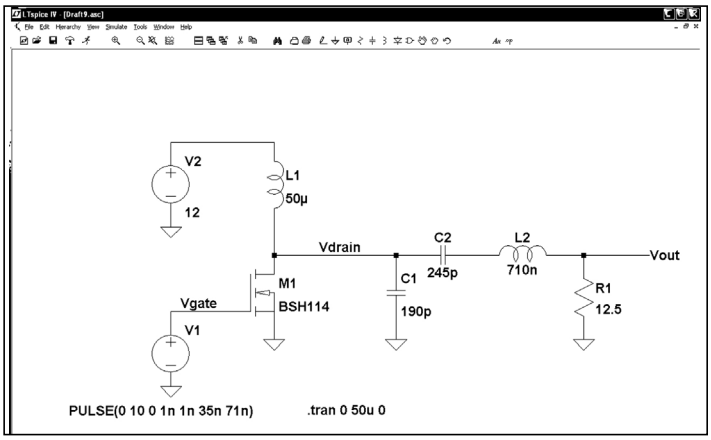

Figure 2—Simulated model of the amplifier using LTSpice.

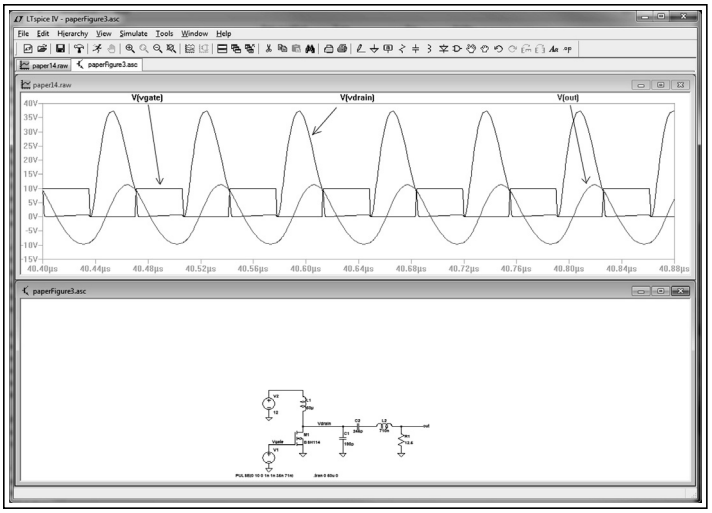

Figure 3—Simulated waveforms in the initial circuit model.

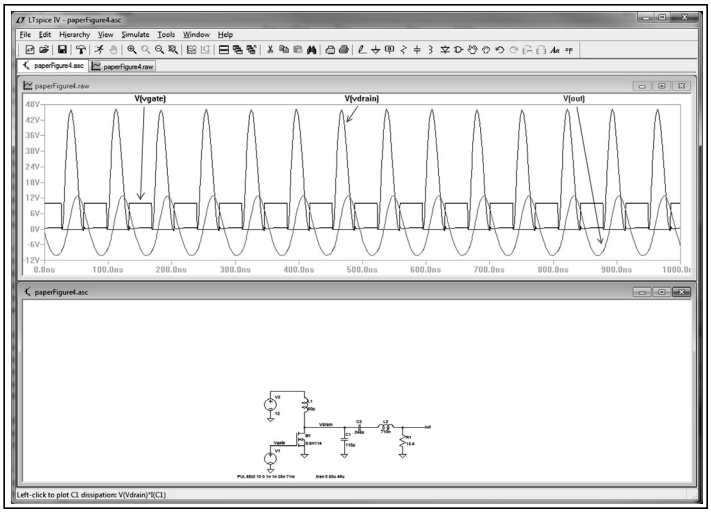

Figure 4—Simulated waveforms taking MOSFET capacitance into account.

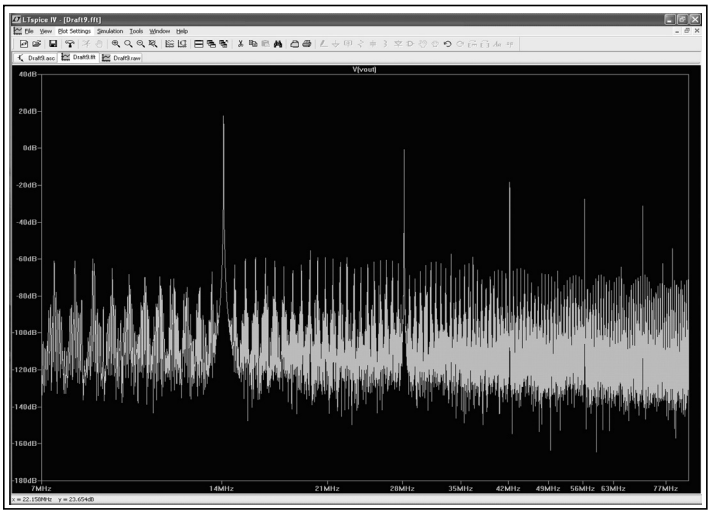

Figure 5—Output spectrum of the modeled amplifier in Figure 4.

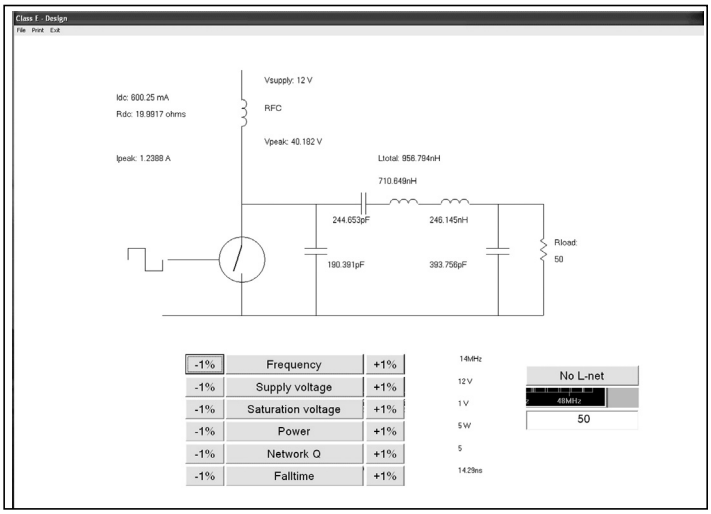

Figure 6—ClassE circuit with L matching network at the amplifier output.

The text under V1 which starts out "PULSE" shows that V1 is not a simple power supply voltage but is, instead, a square wave generator. The key parameters to PULSE are the time the pulse is high (35 nanoseconds here) and the period (71 nanoseconds). The other parameters are the on voltage (10), the off voltage (0) and the rise and fall times of the waveform.

Of course, LTSpice would not be of much use if it could only enter schematics. The real beauty of LTSpice is that it can simulate the circuit as well. Figure 3 is the simulation of the circuit under discussion. The simulation shows the real strength of the class E amplifier. Look closely at what is happening to the V(vdrain) waveform as the V(vgate) signal is going high; the V(vdrain) signal is already going down. Indeed, the goal of the Class E amplifier is to have the V(vdrain) signal be zero just as the V(vgate) signal goes high. Thus, the transistor will have zero voltage whenever it is conducting current; the power dissipated will be (near to) zero.

The waveform of Figure 3 does not have the V(vdrain) signal going to zero just before V(vgate) turns on because we have not yet taken the capacitance of the switching MOSFET into account. This capacitance will result in a lower value for C1. Figure 4 is a simulation with C1 reduced appropriately.

This waveform looks extremely good. A quick calculation is in order. The V(vout) waveform appears to be (largely) sinusoidal so we can compute a rough approximation of the output power: The V(vout) peak is about 12.5 volts so the power out will be Vrms^2 divided by Rload or 12.5*12.5/2/12.5 or 6.25 watts. The current through L1 has an RMS value of about 0.483 amps. Thus, power from the V2 supply is 12*0.483 or 5.8 watts. YES, more power out than in!!! This should raise all sorts of flags. What is wrong? Well, the output power calculation assumed that V(vout) was sinusoidal… is it?

Another great thing about LTSpice is that it can perform an FFT of the waveforms. Figure 5 is the FFT of Vout. Note that the Vout waveform has significant harmonics at 28 MHz (down 17 dB), 42 MHz (down 35 dB), 56 MHz (down 44 dB), 70 MHz (down 48 dB) and higher. Thus, as is expected, the Class E amplifier will need an output filter. Oh… and this explains why power

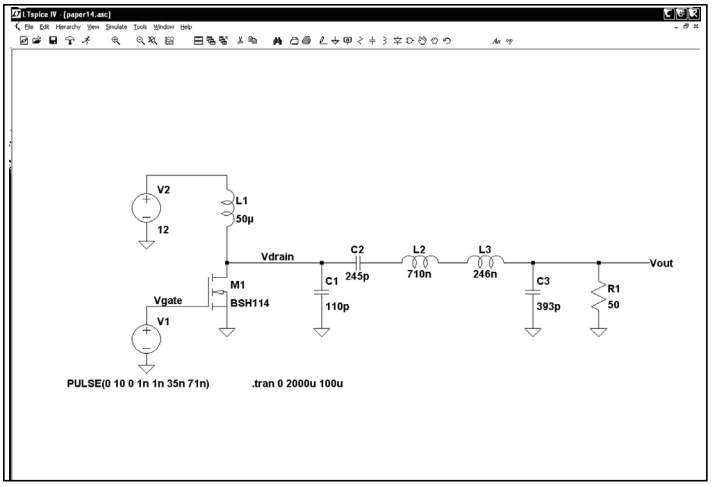

Figure 7—LTSpice simulation for the amplifier in Figure 6.

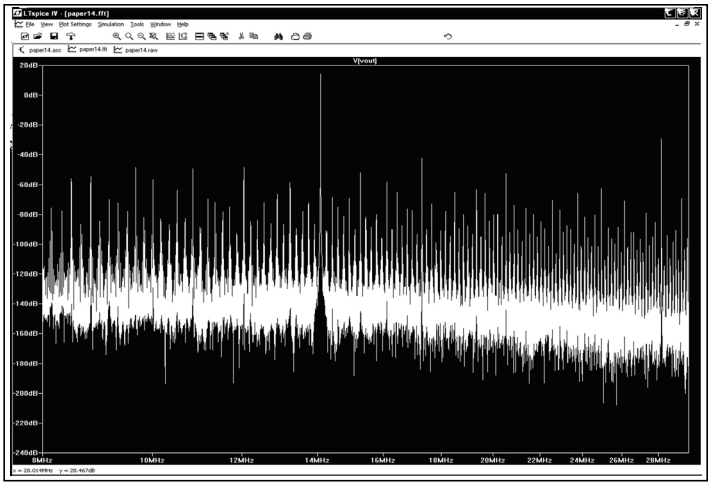

Figure 8—Output spectrum for the amplifier of Figure 6.

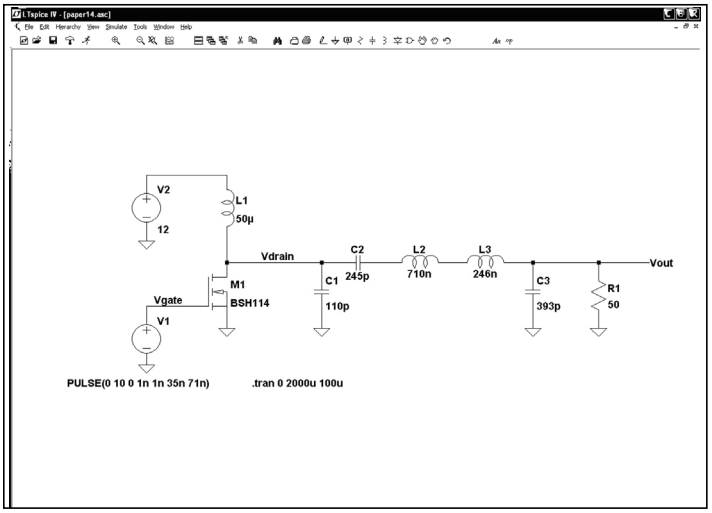

Figure 9—Final design from Elsie showing the output filter.

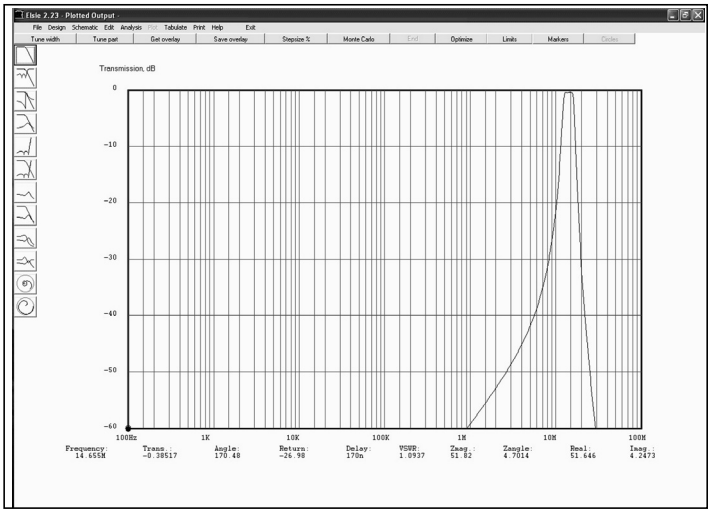

Figure 10—Transfer function of the final output filter.

out appeared larger than power in.

Before proceeding to the output filter, however, there is one issue concerning the power stage of the Class E amplifier. Specifically, Rload in the above design is 12.5 ohms instead of the usual 50 ohms. In the first instantiation of this circuit, I decided that the output filter would have a characteristic impedance of 12.5 ohms and I would merge the L matching network with the last stage of the output filter. This turns out to have been a bad idea. The lower characteristic impedance has the effect of substantially increasing power losses in the inductors of the output filter. Thus, I moved the L matching network to the output of the amplifier and before the output filter. The circuit, shown in Figure 6, is provided by the ClassE tool.

Figure 7 is the resulting LTSpice circuit. And VERY importantly, the output spectrum as shown in Figure 8, now shows greatly reduced harmonic content. The 28 MHz harmonic is now down by 42 dB. The output waveform is much closer to sinusoidal and the output power is now (approximately) 23 volts peak/peak or 23*23/2/50 or 5.3 watts. Much closer to that expected. The difference is probably that the Vsaturated of the transistor is not 1 volt but closer to 0.6 volts. Readjusting the values of the inductors and capacitors should verify this in LTSpice. For this author, that is close enough.

Note now that the output power is about 5.3 watts and the input power was 5.8 watts. This is an efficiency of 5.3/5.8 = 91%, an acceptable value for a Class E amplifier.

Output Filter

For Class E operation, many designers simply design a low pass filter with particular attention to a zero at the second harmonic. I didn't do that. Rather, I chose to design a bandpass filter for 14 MHz. I chose to do a bandpass filter so that I could use the same filter for output and input and not worry about out of band receiver overload.

To design the output filter I employed "Elsie" from the ARRL toolbox. I did extensive playing around with Elsie before settling on the final filter. Elsie is a hugely valuable tool for these explorations. Of key interest to me were the shape of the transfer function and the values of the inductors and capacitors. Figure 9 is the final 'design' page form Elsie. Having decided upon this filter, Elsie is kind enough to plot the transfer function, as shown in Figure 10. As shown, this filter should add another 60 dB of suppression to the second harmonic at 28 MHz. Elsie also can draw the circuit but I won't show that here. Rather, Figure 11 is the

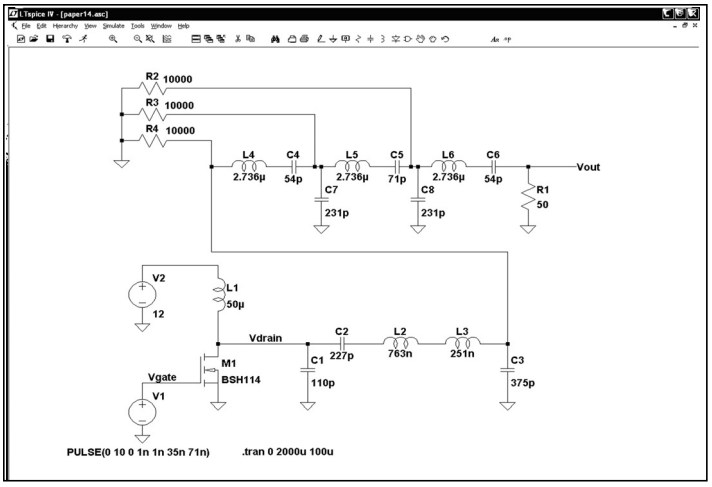

Figure 11—LTSpice circuit; amplifier and final output filter.

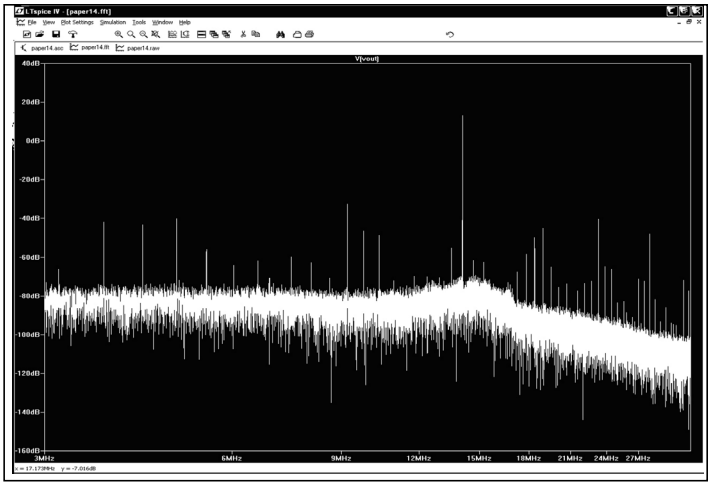

Figure 12—Output FFT of the simulated amplifier and output filter.

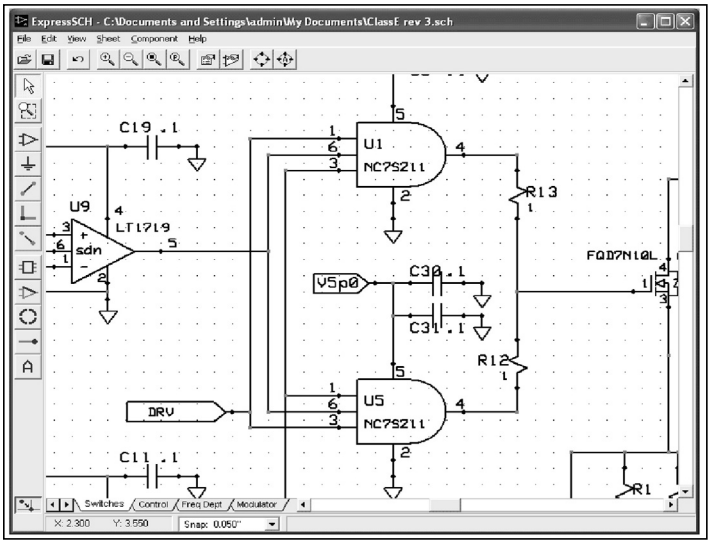

Figure 13—Circuit of the MOSFET driver.

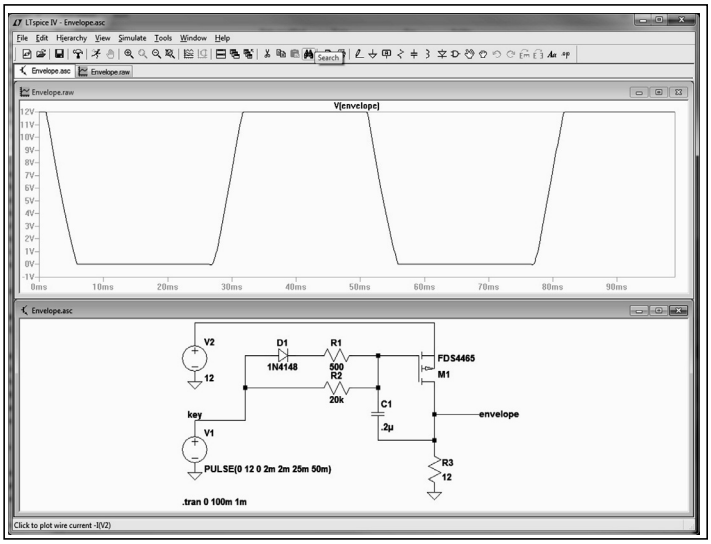

Figure 14—LTSPice simulation showing envelope shaping.

LTSpice circuit. The detail oriented reader will notice the addition of three, 10k resistors. These resistors are needed because LTSpice doesn't work well with undefined voltages on an electrical 'node'. If these resistors were not added then the simulation would run significantly slower.

Figure 12 is the FFT of the output signal. One of the weaknesses of LTSpice is that when computing the FFT it does so over a relatively small subset of time. Even when windowed it can report harmonics which are larger than reality. Still, it can show trends and relative magnitudes. In the above, the 28 MHz harmonic is shown as 60 dB down from the 14 MHz carrier.

Driving the Power Transistor

One of the real niceties of all switching power amplifiers is the simplicity of the drive circuit. Modern MOSFETs can be purchased which have a Vgate on voltage of less than 5 volts. As a result, it is possible to drive the MOSFET gate using standard 5 volt logic components. It is still important to drive the MOSFET gate as hard and fast as possible. Figure 13 is the circuit I used to drive the MOSFET.

Of course, the two NAND gates require digital signals as inputs. There are three inputs to the AND gates shown above. The DRV signal is simply an 'enable' from the SDR processor. The LT1719 is a high speed comparator which converts the incoming analog sine wave (14 MHz here) into a digital square wave. The third signal comes from the protection logic which detects over-current and shuts down the NAND gates when an over-current condition is detected.

Envelope Shaping

In the foregoing circuits, the voltage supply for the power stage is a fixed voltage. However, the polite CW operator tries to reduce key clicks. This is done by 'shaping' the envelope. In this first instance of my power amplifier, I decided to use the traditional shaping approach shown Figure 14.

In Fig. 14, the 'key' signal is driven with 12 volt rail to rail signal from the SDR processor. The values of R1, R2 and C1 are adjusted to fit the operator's desired ramp rates. Slower rates result in less 'key click' noise. Faster rates can improve readability in poor conditions.

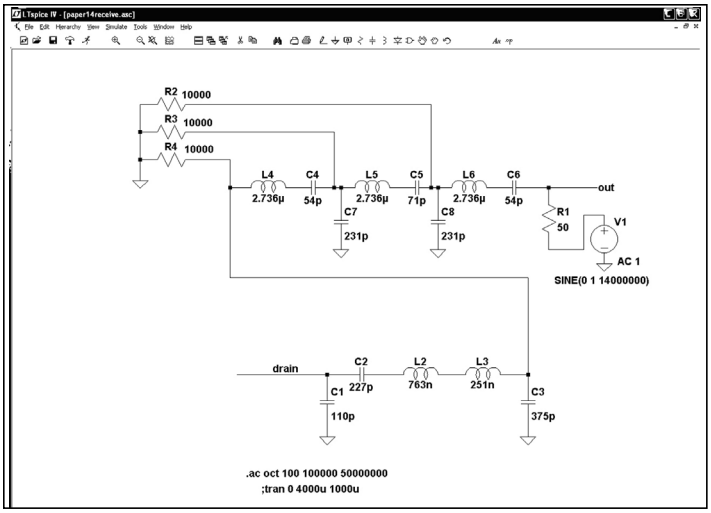

Figure 15—Initial simulation of the amplifier during receive.

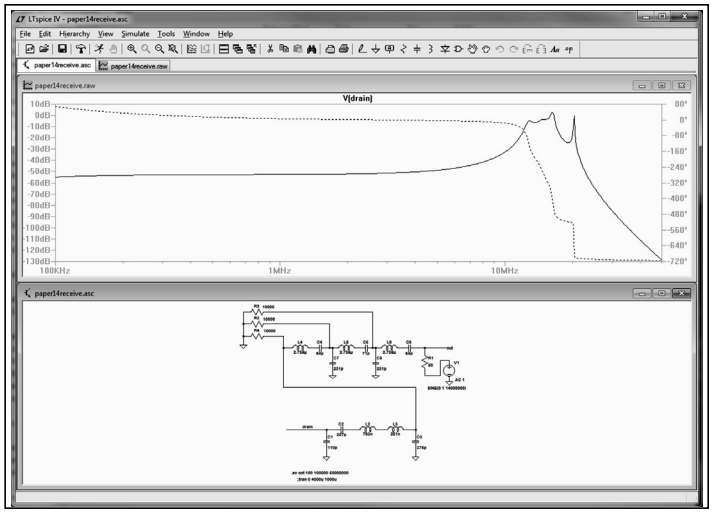

Figure 16—Filter shape for the circuit of Figure 15.

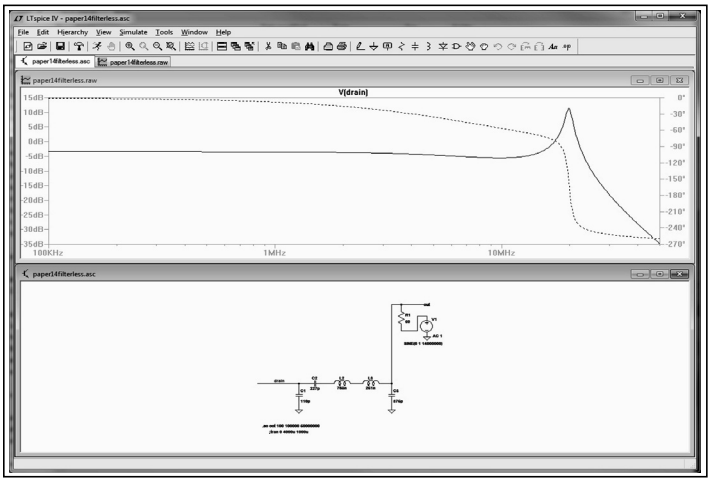

Figure 17—Circuit simulation and transfer function of the matching network only.

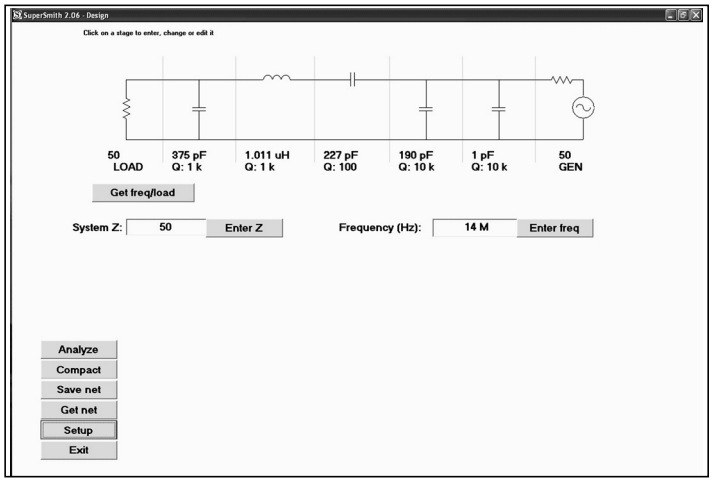

Figure 18—SuperSmith design page for the matching network of Figure 17.

With the addition of an op amp, one can actually modulate the envelope and generate an AM signal. With an 'edge detector' driving the switch FET, one can actually build a linear amplifier using class E as the final stage. There are other ways. The problems involved are brushed over here.

Transmit/Receive Switching

There are many ways to provide transmit/receive switching. In my design, I wanted to use the output bandpass filter as an input filter as well. So, I decided to take the receive signal directly from the MOSFET's drain. It is important to remember that the power transistor and the envelope transistor will be 'off' during reception. Thus, during receive the equivalent circuit is as shown in Figure 15. Note that I have added a voltage source "V1" to the 'out' of the output filter and deleted the MOSFET and envelope voltage sources.

LTSpice has another capability called 'AC' analysis where it can display the filter shape of a circuit. Figure 16 is what LTSpice shows for this circuit. Well, this is pretty ugly, although it really isn't all that bad. Clearly, the problem is that the filter is not terminated properly. Here is where the Smith chart can come in quite handy. Remember that the upper part of the above circuit is just a 50 ohm bandpass filter. Thus, the circuit of Figure 16 can be simplified somewhat to the circuit shown in Figure 17 along with the filter shape. Note the peak at about 20 MHz remains. The goal is to remove this peak and provide a reasonably flat response through 14 MHz.

Moving this circuit to SuperSmith results in the display of Figure 18. Note that I've reversed the circuit and shown the antenna as the load on the left. Also, I've added a capacitor on the right which I'm going to adjust as necessary to get the proper impedance as seen from the GEN side. Clicking the "Analyze" button yields Figure 19.

It is relatively simple to adjust the value of the capacitor on the right to compensate and deliver a 'real' impedance. This has been done in Figure 20 where the 1 pF capacitor has been adjusted to 207 pF. Having adjusted the capacitor value to produce a real value, we can see that this real value is now about 33 ohms. We can put the adjusted capacitor and 33 ohm resistor back into LTSpice just to verify things, and this is shown as Figure 21.

Just for fun, let's look at what happens if we remove the compensating 207 pF capacitor. Figure 22 shows the result. This looks

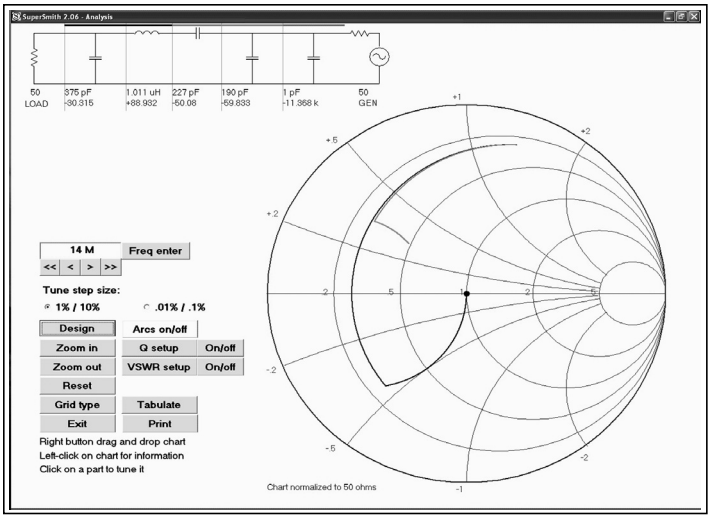

Figure 19—Smith chart analysis of the matching network in Figure 18.

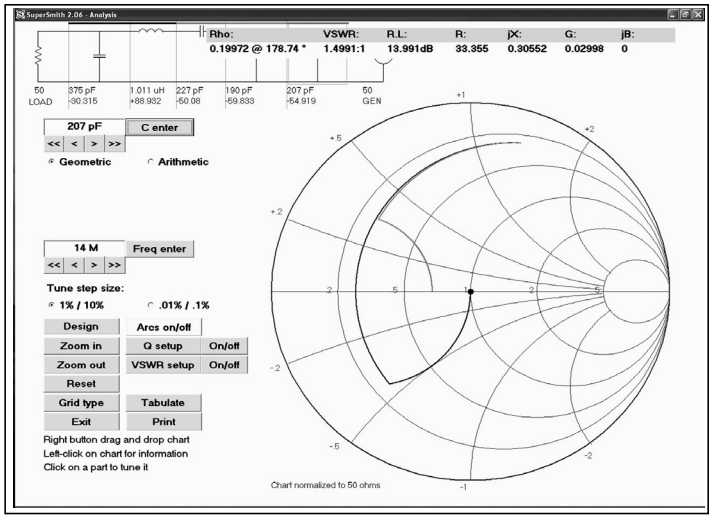

Figure 20—Smith chart analysis after adjusting the right hand capacitor.

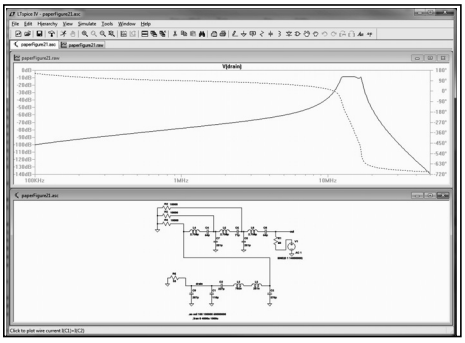

Figure 21—LTSpice simulation after adjusting the right hand capacitor.

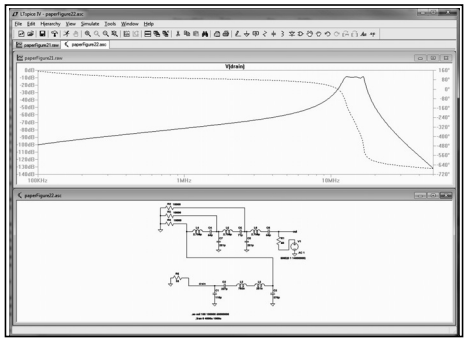

Figure 22—Simulation without the compensating 207 pF capacitor.

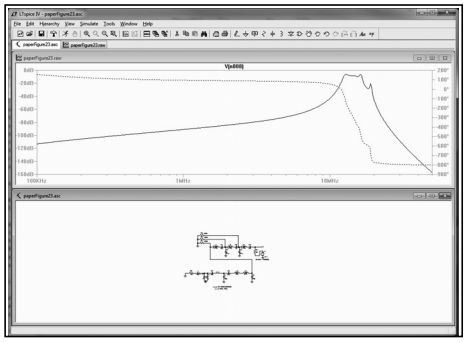

Figure 23—Modifed circuit that includes the 33 ohm resistor only during receive.

pretty darned good (if not better), so let's just leave it out.

The problem now is how to put in the 33 ohm resistor when receiving and how to take it out when not receiving. Figure 23 shows a very common approach. A few notes on the above circuit are appropriate. First, the back/back diodes are generally conducting during transmission. On reception, they don't conduct. Thus, the 25 pF coupling capacitor is in parallel with the MOSFET during transmission and, as a result, C1 has been reduced to compensate. During reception, the 25 pF capacitor has an impedance of roughly X = –500. The inductor cancels this impedance out during receive.

Notice that the peak at around 19 MHz has started to creep back in. If this is a problem, determining the impedance of the diodes and then returning to SuperSmith would help eliminate this peak. I did not think it necessary to take this step. Figure 24 is the more-or-less complete circuit.

Summation

Over the years I've built many circuits designed by other people. Often times I have wondered about the design process itself. How did they make the various decisions? What tools do they use? What compromises do they make along the way? Just how did they get to the end?

In this paper, I have described the step by step process by which I designed my class E, QRP amplifier. It is my hope that the reader has gained an appreciation for the design process and the tools available to help the analysis. These tools allow the budding designer to 'jump in' and start and refine a design with very little experience. Ultimately, when used properly, these tools will let you design, modify and optimize a circuit with confidence. When you actually build it, it will WORK!

●●

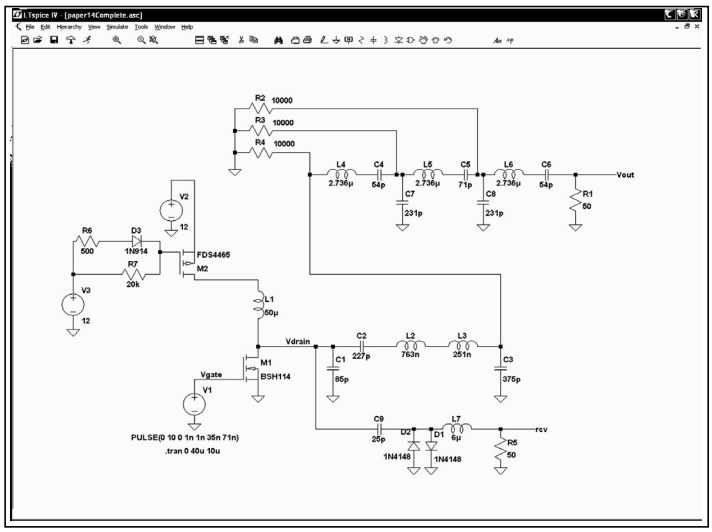

Figure 24—The complete LTSpice circuit of the amplifier.