The Filter Wizard
issue 14: Match Point: why Maximum Power means Minimum Sensitivity
Kendall Castor-Perry

In Filter Wizard #12 (Filter Design using the Million Monkeys Method) I made myself extra work by promising more detail on several filter design and network theory issues that I rushed through at the time in my quest for a pacy, interesting article. So, here's the first of those "deep dives" (sorry, corporate-speak there...): why did I want to set my source and load resistances to the same value?

The filter topology used in the article is known as a doubly-terminated LC ladder filter. The inductors and capacitors form a network with two 'ports' – places where signals can enter and leave the network. In our filter, the input port is fed from a source that has a particular resistance value (that you usually can't change), and the output port is connected to – "loaded by" – another resistor (that's often under your control).

Figure 1 shows an example filter that, like the network from the previous column, is designed to be terminated with a resistor equal in value to the source resistance. Figure 2 shows the nice flat passband response, and figure 3 the rather poorer response we get if we leave off the terminating resistor.

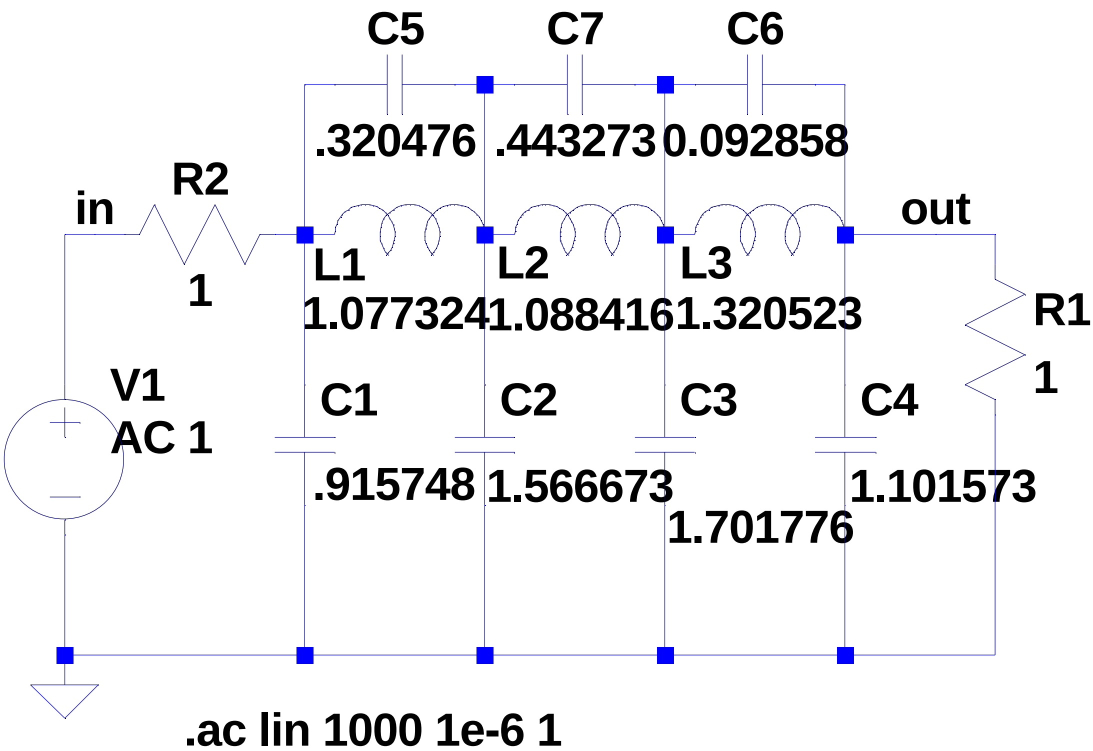

figure 1: an example doubly-terminated lowpass filter network

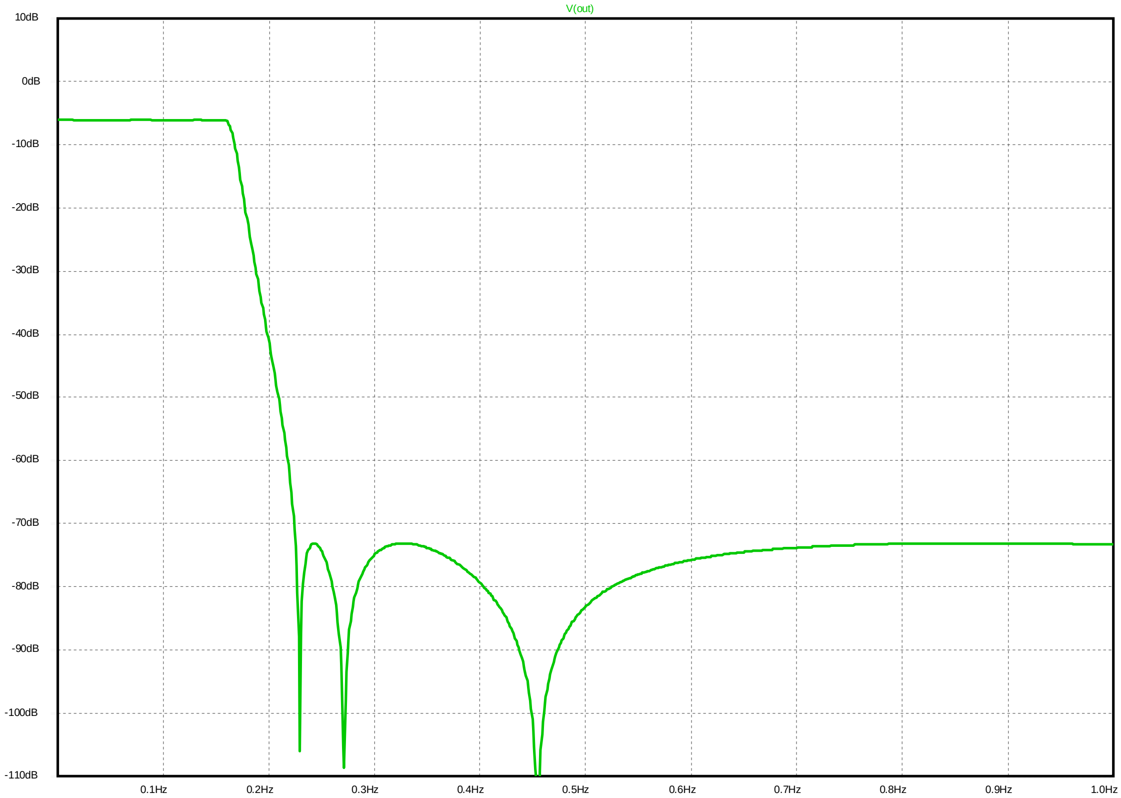

figure 2: the voltage gain of figure 1's values with the load resistor fitted

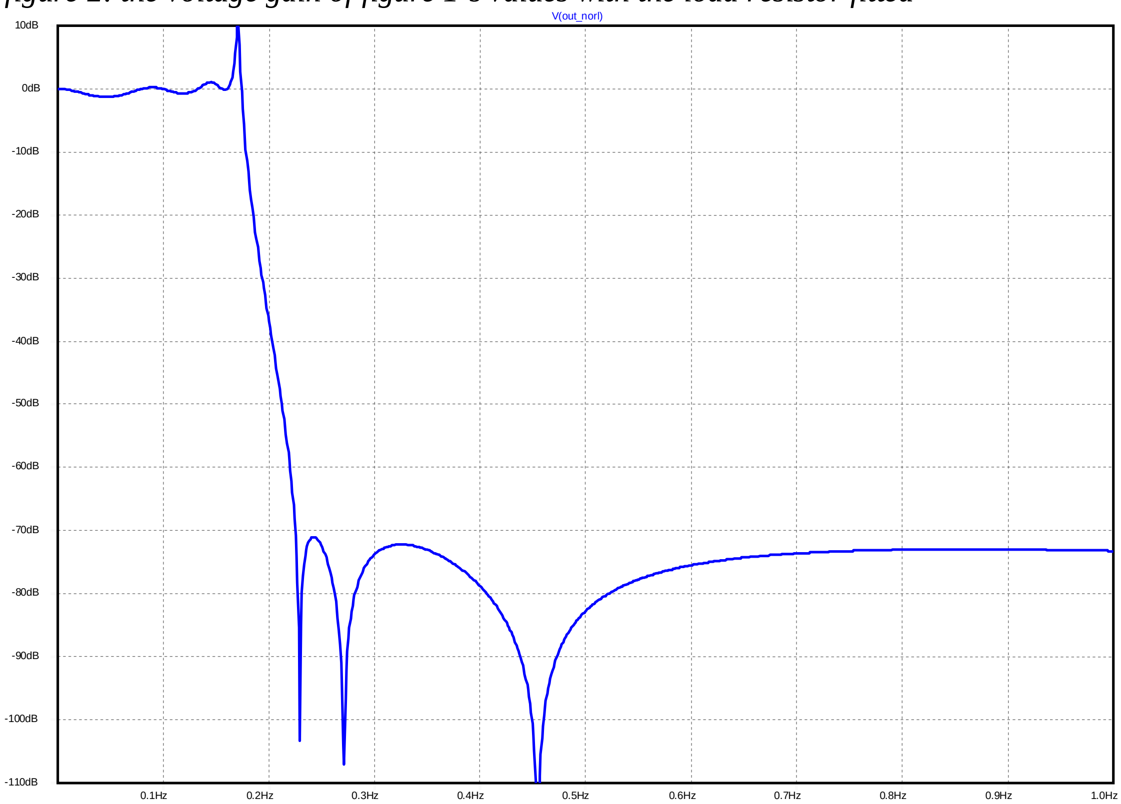

figure 2: the voltage gain of figure 1's values without the load resistor fitted

We can see that for the particular set of values we used here, the low-ripple flat response is only available when we fit the load resistor. But this isn't a *consequence* of having two resistors; figure 4 shows a set of network values that gives us the same amplitude response (figure 5) without a load resistor. It has a voltage gain of 0dB rather than -6dB. Surely you'd use this filter in most practical situations? Same response, more gain and the possibility of saving a component, what's not to love?

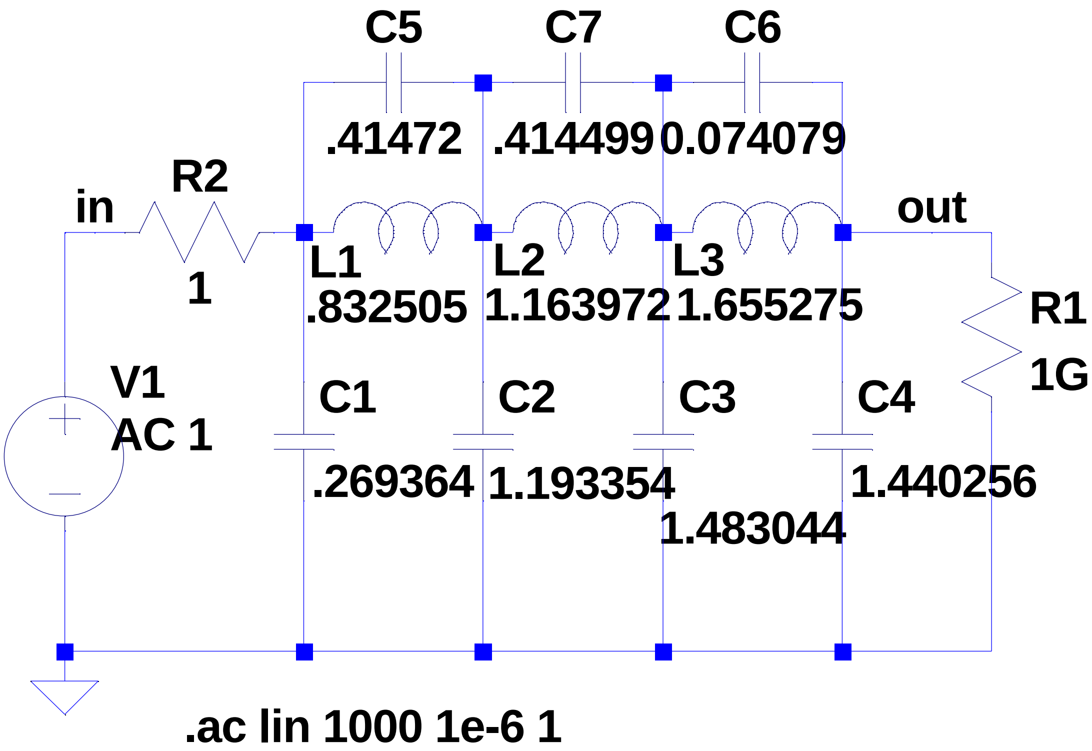

figure 4: a set of component values that doesn't need a load resistor

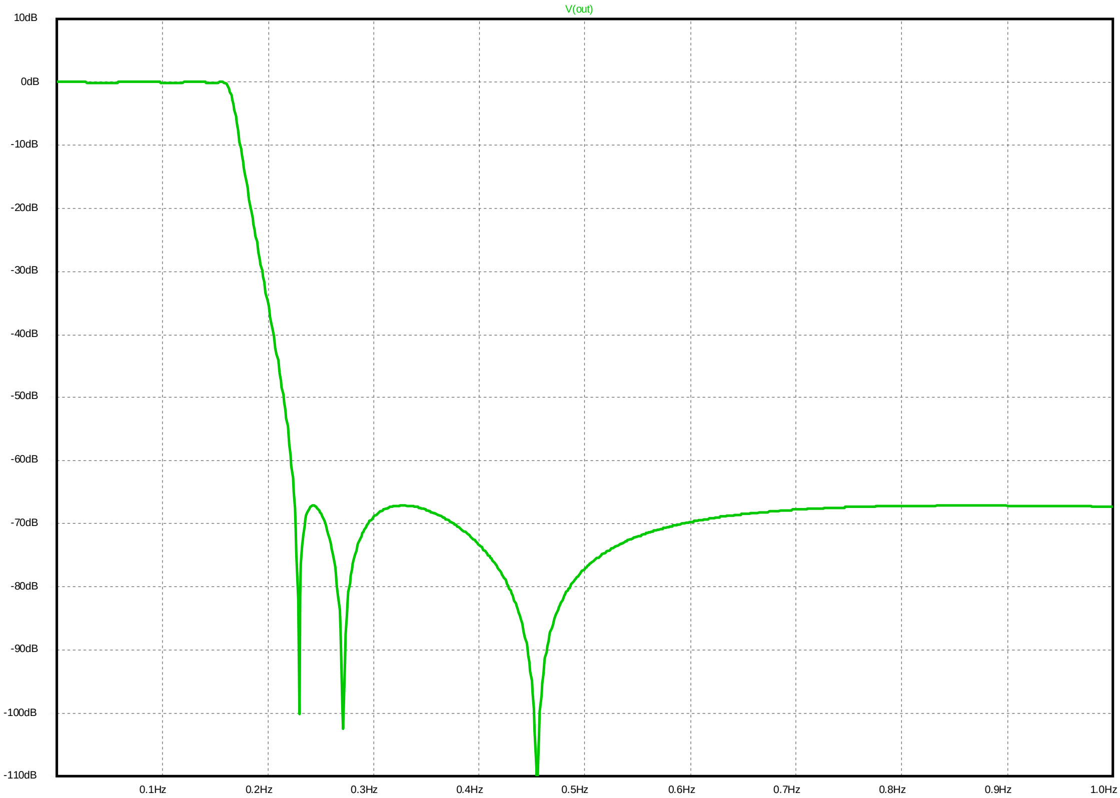

figure 5: look, no load resistor, and still the response is nice and flat

Well, figures 6 and 7 show why we need to resist (sorry) the temptation to go with this "singly-terminated" approach. The figures each show 100 superimposed response plots in which the component values have been permitted to vary over a ±5% range. We can see clearly that the doubly-terminated filter is, literally, much more tolerant to small changes in the value of the components used. We'll see, in the next column, that this behaviour is preserved when we do some scaling tricks that allow us to use these networks as 'prototypes' for active filters without all these inductors.

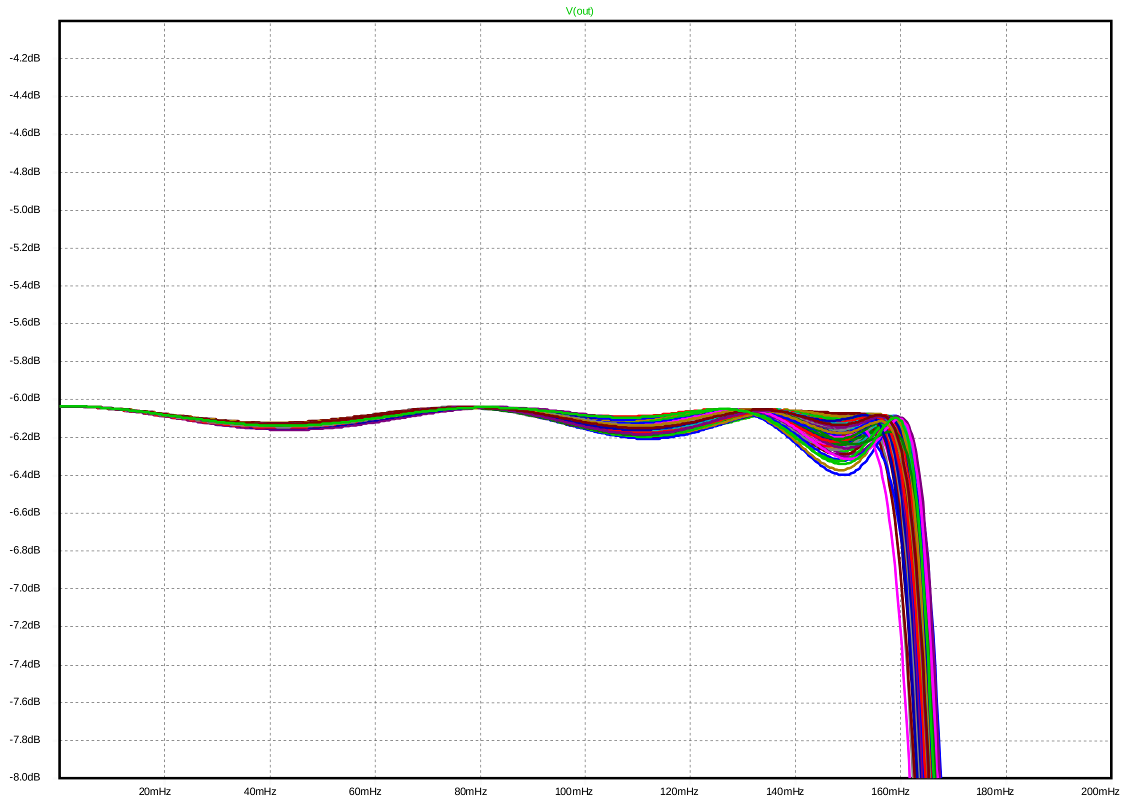

figure 6: Monte Carlo analysis of figure 1's values

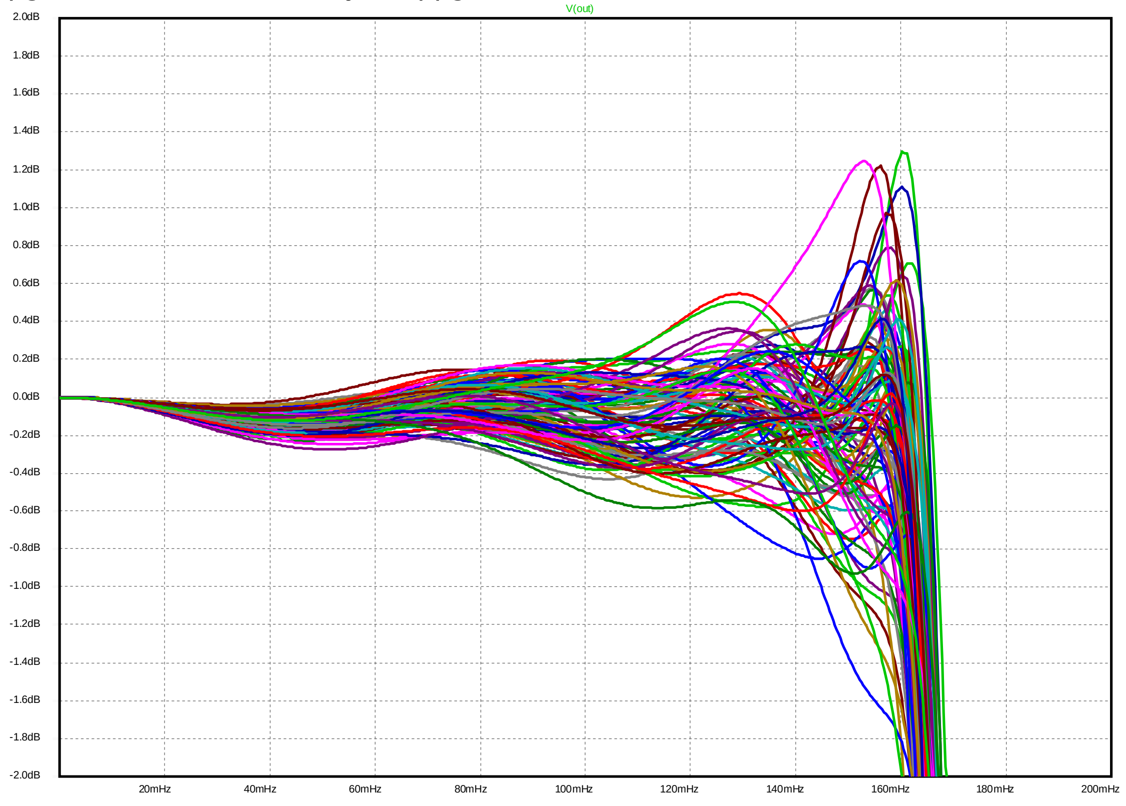

figure 7: Monte Carlo analysis of figure 4's values

What is so special about the doubly-terminated filter? To answer this, we need to consider what is happening when a signal transits the LC network, from the source resistance to the load resistor.

Here's a digression to seed the thinking process. Imagine you're interviewing for an analogue design role, and you're asked the following question:

> "You have a sinewave signal generator with a (resistive) 50ohm output impedance, and it can produce 1Vrms into an open circuit. Your customer has given you a black box with an input impedance of 3.3kohm, and for it to function correctly, it requires an input signal of at least 3Vrms at 50kHz. You must get the system working, but have no power source. The only electronic components you have access to in the lab are passive two-wire components. Explain how you solve this, and why it works."

So, that means no batteries or solar cells, no transistors or ICs, and definitely no transformers. Have a think about it before you read on (especially if you're interviewing soon). Maybe look at figure 2 for a clue. Done that? OK, here's one solution that does the trick:

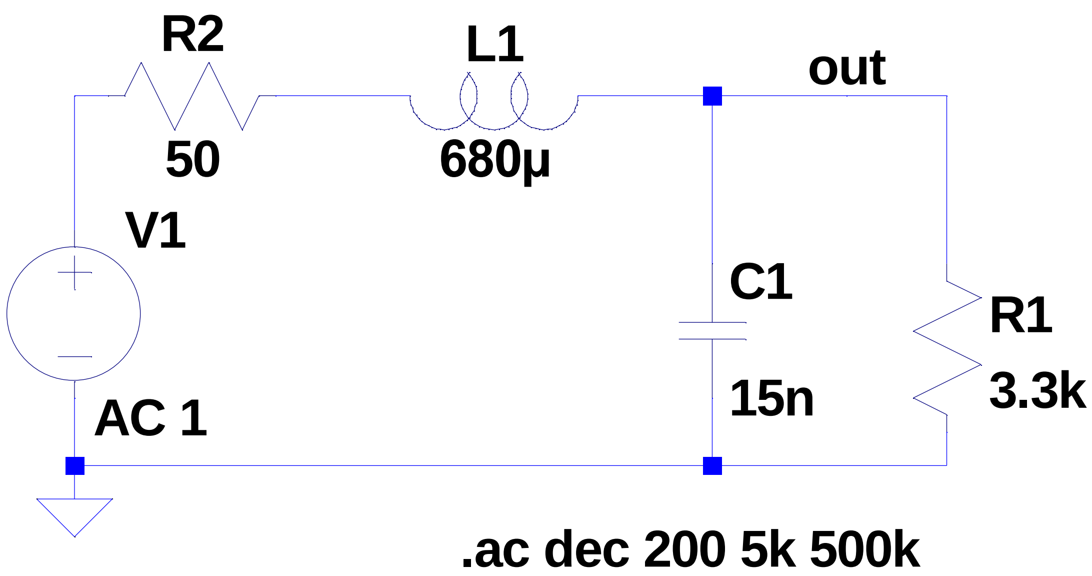

figure 8: a solution for the fiendish interview question

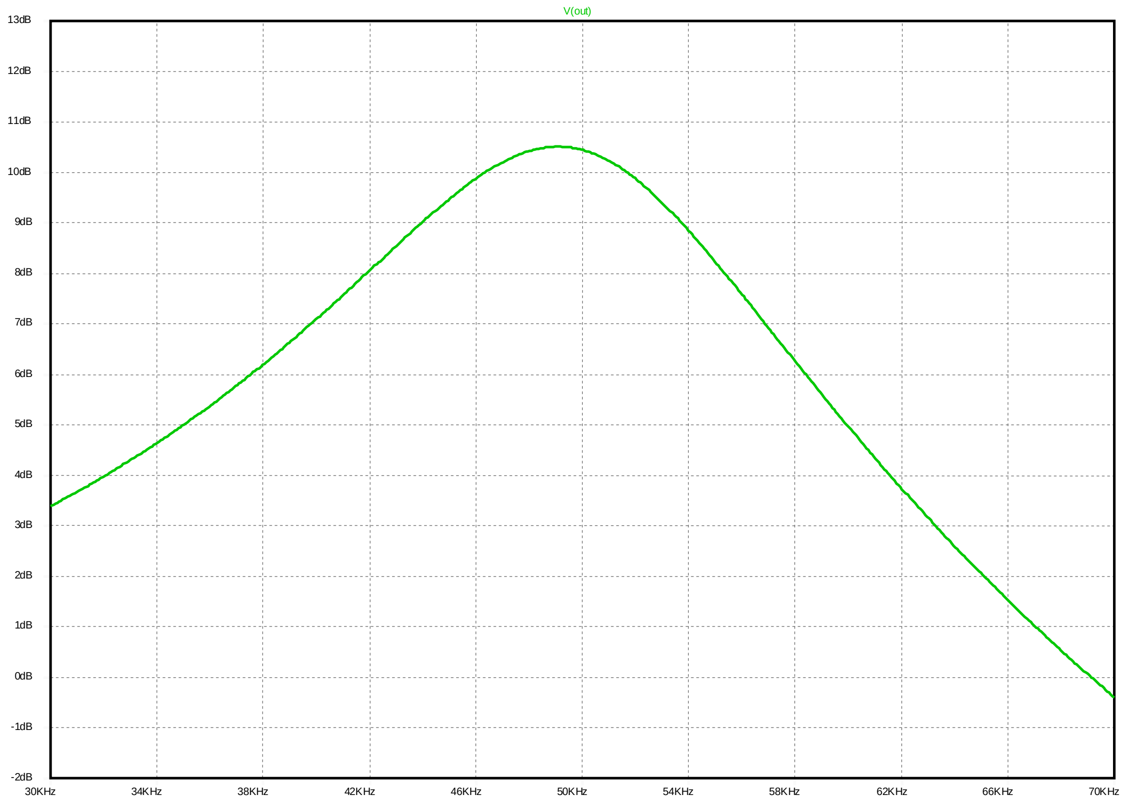

figure 9: frequency response of figure 8's solution

We built a high Q lowpass filter, using only passive two-wire components. Figure 9 shows the voltage on the black box input for a 1Vrms open-circuit setting on the signal generator, swept over frequency. We've clearly achieved the voltage gain necessary to get the system working at the frequency specified. OK, so where's the free lunch?

Well, nowhere, of course. We could have asked, for bonus points: "What's the *maximum* value of voltage gain you could possibly extract from such a solution". To figure this out, the insight is to realize that, while getting *voltage* gain from a collection of components like this is unremarkable, it's impossible to get *power* gain from it. This is key to understanding what happens in these passive filter networks: it's all about power entering and leaving the network. If you have a source that delivers a certain output voltage with a particular output resistance, there's a strict upper bound to the amount of *power* that can be dissipated in a connected load. You may have encountered the Maximum Power Theorem in college, perhaps not paying it much heed. But it's hugely important in passive filter design so, please, heed it forthwith.

The condition for maximum power transfer, you may now be remembering, is that the load resistance is equal to the source resistance. You can prove that by applying some of that dusty old college calculus (there, I said a dirty word, I'm *that* radical). Write down an expression for the power dissipated in the load resistor, calculate the derivative of that power with respect to the load resistor RL, set that derivative equal to zero, and solve for RL. Go on, I dare you.

Is an “aha!” moment hoving into view yet? Your first thought might have been to use a transformer. That’s a classical way to ‘match’ a load resistor unequal to the source resistance. The transformer passes power from the source to the load without (ideally) losing any of it, but it applies it at a different voltage and current level. With the correct turns ratio, you can get the same amount of power dissipated in the load resistor as in the ideal case, whatever the load resistance value.

And that’s what we did with our simple LC filter network in figure 8: we made a transformer. All the power that goes into the LC network comes out again, and with the correct values, we can make it all come out into the load at the desired frequency, whatever the load’s value.

We can calculate what maximum voltage gain would be possible between our 50ohm source and our 3.3kohm load if we’re allowed to use a transformer or its LC circuit equivalent. The voltage transformation needed is equal to the square root of the impedance ratio, so 8.12x in this case. This multiplies the voltage gain we get in the equally terminated case, which is one-half. So the total maximum gain you could get with any LC values here must be 4.06x, or ~12.2dB. Figure 10 shows the result of randomly selecting 1000 different L and C values in our home-made LC transformer; the peak of the response curve never goes above the predicted value The reason why it’s not the same at every frequency is that not all of the power gets dissipated in the load; some of it gets ‘reflected’ back to the source.

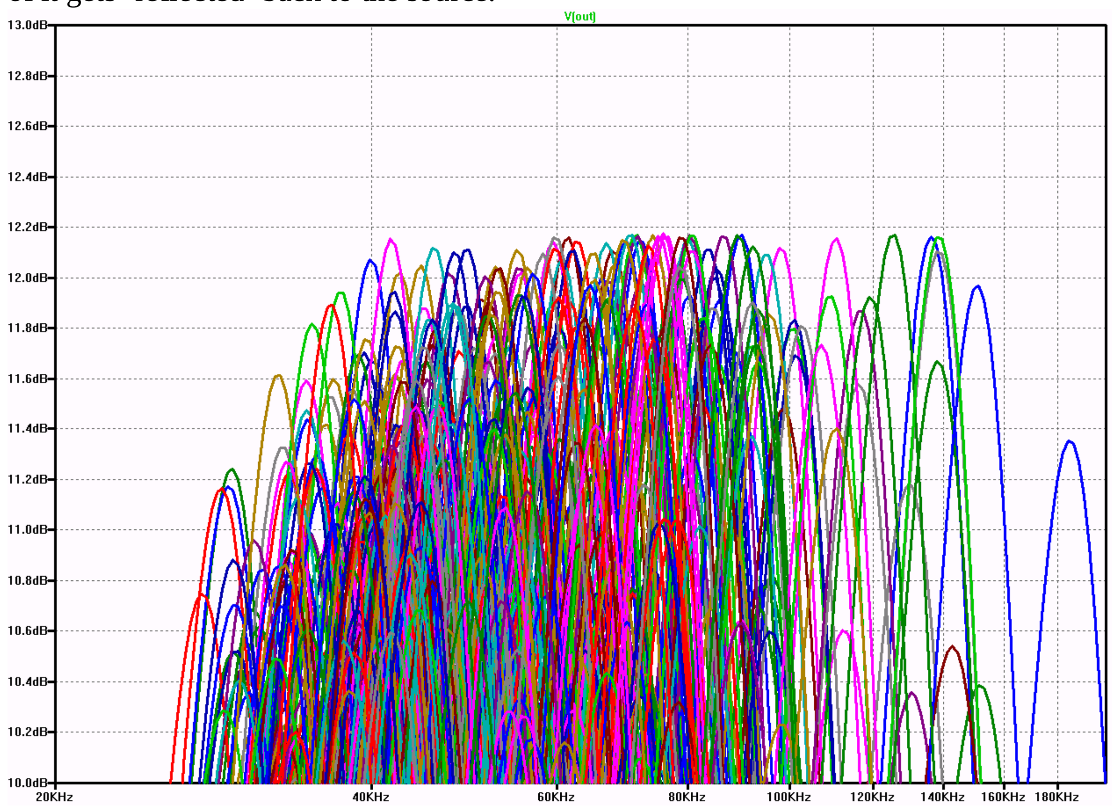

figure 10: nobody beats the Max

All you RF guys out there are ho-humming impatiently, because of course all this is second nature to you. Designing an LC network to ensure that all your precious source power gets transferred to the load is a core RF technique called impedance matching. It's really just like filter design between two unequal resistors (each of which may also have its own accompanying imaginary part that needs to be taken into account). It's typically done with Ls and Cs rather than with bulky, expensive transformers. Except at microwave frequencies, where the transformers aren't bulky or expensive, but spring up all over the place when conductors get too close.

Now, dolly back to our original issue. Why does a (correctly designed) doubly terminated filter have such superior 'sensitivity' properties? It's because you're working at that maximum possible point of power transfer, for one or sometimes several frequencies in the filter passband. Look again at figures 6 and 7. In the case of the doubly terminated filter, any change in component values can only make the power transmission (and therefore the voltage gain) go *down*, not go up. The 'sensitivity' of the filter response at certain key frequencies, called the reflection zeroes, is a downwards parabolic function of every reactive component in the network, and it's very much harder to force the response of the network to be wrong than it is when power transfer constraints don't apply. That's what happens in the singly terminated case, or in any case where the designed response of the filter doesn't actually touch the maximum power gain value, which you get if you throw away the filter and replaced it with an ideal transformer of the correct transformation ratio.

One more thing. We compared two filters way back at the start; one had its load resistor equal to the source, and the other had no load resistor. Can we make successful filters with the low-sensitivity property and the same frequency response for *any* given ratio of source and load resistance, however wide? Well, sometimes. In the case of our particular lowpass filter, though, it's a no. For this particular response, we want a flat response, with the DC gain equal to the value at those 'touching point' maxima. This means that the source and load resistances *must* be equal-valued in the lowest-sensitivity filter. In the more general case where you don't have to pass zero or infinite frequencies with specified attenuation, you can generally design a minimum-sensitivity (i.e. optimally matched) filter between any two impedances. But the how-to for that fills books – lovely books, but not ideal for the faint of heart.

I hope you now sense the magic of Maximum Power Transfer and how it launches us down a path to some high-specification yet manufacturable filter design approaches. Much more on this topic in the next few columns. Happy matching! - Kendall