The Filter Wizard
issue 21: A Fast-Settling Bias Voltage Filter with High Ripple Rejection
Kendall Castor-Perry

Elaborating on a previous article, the Filter Wizard shows how higher order lowpass filters made using a "DC-free" technique can be much more effective in removing ripple from high voltage sources, while responding rapidly to step voltage changes and requiring far lower capacitor values.

In "Filter DC Voltages Outside Your Supply Rails" we saw how a popular active bandpass filter topology could be re-imagined into a building block for a form of higher-order **lowpass** filter that doesn't "see" the DC potential on the signal it is filtering. This lowpass filter is a form of RDC filter, as introduced previously in "Bruton Charisma". Its values are derived from those of an LCR passive filter. The last article begged some questions:

(1) Why can't we just use an old-fashioned passive filter, completely getting around the problem of providing a supply voltage for active components?
(2) What response characteristics do we need from a lowpass filter in order that it has some benefits for filtering practical bias and reference voltages?
(3) Are there any hidden catches in the use of this "DC-free" filter?

Let's set up an example application. Suppose we've got a high bias voltage, say 180 V, derived from an unregulated line-operated power supply running in the US from 60 Hz AC. Let's assume that the destination for our bias voltage can tolerate a 100 ohm source impedance, which sets the maximum permissible value of the series resistor. Also that the supply's 120 Hz ripple level is much too high for the precision system we're driving, and that it must be significantly reduced, from 1 Vpp to 1 mVpp, at this frequency and above.

The simplest lowpass filter is a single pole RC network. The rule of thumb for this type of network is that you get 20 dB of attenuation for each decade increase in signal frequency above cutoff. So we'll need our filter cutoff frequency to be three decades below 120 Hz, i.e. 0.12 Hz. With our 100 ohm series resistor, we'll need a capacitor value of around (fires up handy H-P 15C app on iPhone) 1.33E-2 which is, er, 13300 uF. Dude, I haven't got room for **that** giant on the circuit board. Also, just think of all the energy stored in this capacitor at such a high voltage.

The use of such a large capacitor value raises another issue. That RC time constant is going to be 1.33 seconds. How long is that going to take to settle to 0.1% after we switch the input voltage on? It's the RC value times –ln(0.1%), which comes to 9.2 seconds. Dude, I can't wait **that** long for it to settle!

This is definitely a job for a higher order filter with a more rapid transition from passband to stopband. Because the stopband rejection of a lowpass filter builds up more rapidly with rising frequency as the order increases, we can use a filter with a higher intrinsic

cutoff frequency and still get the desired rejection at 120 Hz. This in turn can deliver lower capacitance values and also a more rapid settling to a stable DC value –if we pick the right shape of response! But what **is** that response? If the filter has a transfer function that exhibits too much ringing when excited with a step, we may still fail to achieve our requirements.

Well, the only frequency response constraint we've specified for our filter is in the stopband – we want it to have 60 dB rejection at 120 Hz and above. When the bias voltage makes its transition from off to 180 V, that's like a step excitation for the filter, and the filter output is going to take time to settle down after that stimulus. The settling behaviour of a filter is determined by the relationships between the frequency components in a step input that **do** get through. This isn't affected by the stopband rejection, only by the **passband** shape and bandwidth. So, we should pick the passband shape that's going to give us the most rapid settling, given our one fixed stopband intercept.

Filters that meet this criterion have a particular, 'soggy' form of passband response. It's called a **Gaussian** response, because its functional form – in both the frequency and the time domain – is a Gaussian function, a scaled and shifted version of y=exp(-x^2). An ideal Gaussian lowpass filter has no overshoot and the fastest monotonic settling for a given filter bandwidth. Unfortunately, Paley and Wiener showed in the 1920s that a filter with a gain function of the form G=exp(-kf^2) for all frequencies **can't** be realized physically; it has an attenuation that just rolls off too quickly. We have to make do with an **approximation** to this functional form, giving us a similarly shaped passband but some realizable level of stopband slope. That's the answer to question (2), by the way.

By far the most familiar approximation to the Gaussian response is the Bessel filter, but there are other candidates. Standard filter tables give values for filters that approximate the Gaussian form down to -12 dB and then hit their final slope value more rapidly than the rather leisurely Bessel. So let's look at the response of the 3-pole and 5-pole "Gaussian to 12 dB" filters in comparison to the single-pole filter, with their -60 dB response points aligned. Why these two choices? Well, the DC-free topology I'll use is most economical in amplifier usage per pole when the order is odd. A 3-pole filter will need one op-amp, and a 5-pole filter will need two. I'm going to assume that I can spare a maximum of two of the op-amps in a nearby Cypress PSoC3 that is already doing some other functions and is running of a local 5 V supply. Figure 1 shows the magnitude responses, with the step responses in figure 2.

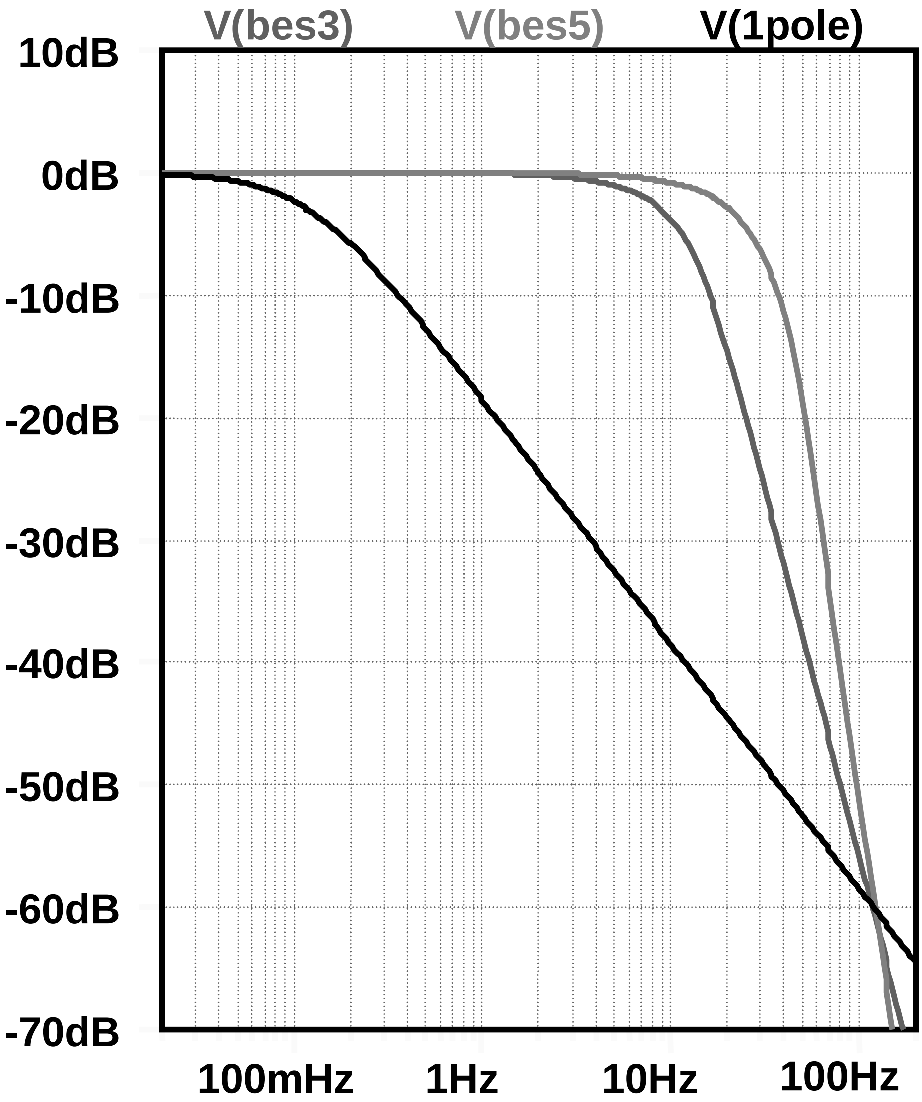

Figure 1: The single-pole filter, and 3- and 5-pole “Gaussian to 12 dB” responses.

We can see on the timescale of figure 2 that the 3-pole and 5-pole filters are **way** faster than the single-pole network, which settles to its final value as predicted by our thumbnail calculation. So let’s ignore the single-pole network’s trace and zoom in on the settling behaviour of the 3-pole and 5-pole “Gaussian to 12 dB” filters, in figure 3.

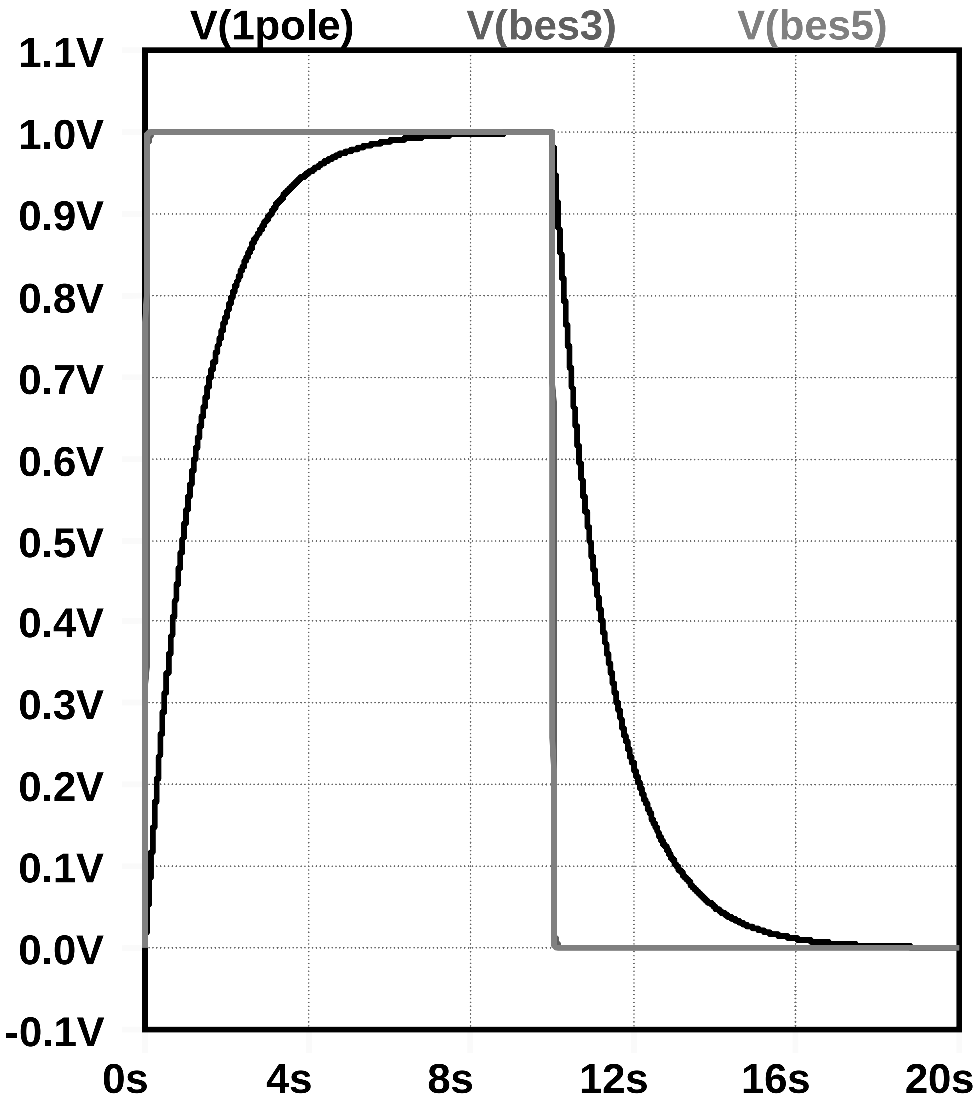

Figure 2: both the 3- and 5-pole filter settle **much** faster than the single pole.

It’s clear that we get the widest passband (-3 dB at 20.6 Hz) and fastest settling time (73 ms to the first 0.1% point) for the 5-pole version. The 3-pole version has about half the bandwidth (-3 dB at 8.94 Hz) and takes just about twice as long to settle down to <0.1% (141 ms), which is still a great improvement on the 9.2 seconds that the single pole network took, when all of them deliver the 120 Hz rejection we need.

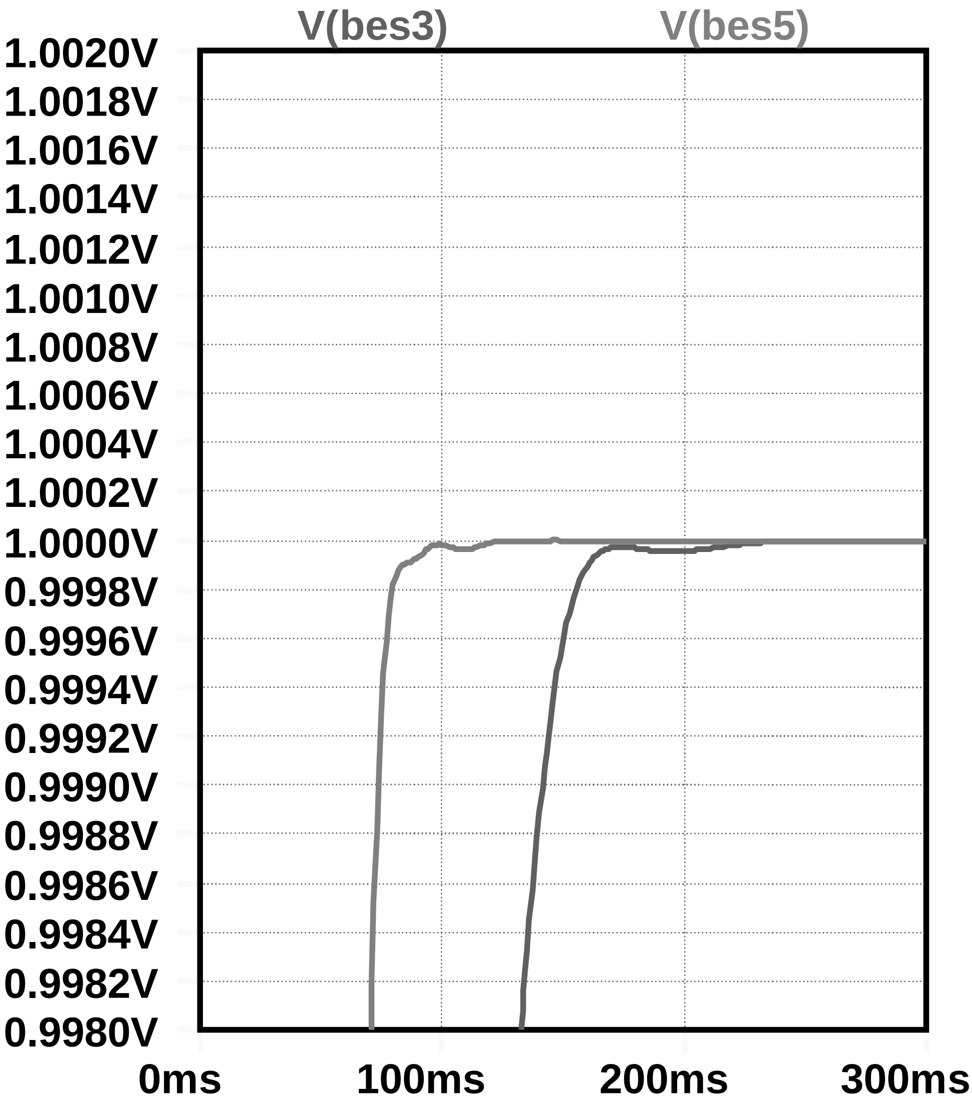

Figure 3: zoom in to the 3-pole (right) and 5-pole (left) "Gaussian to 12 dB" settling.

Realizations of the filters as passive networks at a 100 ohm impedance level with series inductors are shown in figure 4 The initial prototype LC values were taken from standard tables of singly-terminated "Gaussian to 12 dB" filters, collected in Williams & Taylor [1] as well as in older classic books such as [2]. Here we use just a source resistance, and the load is assumed infinite. Look at those inductance values, though; these are completely impractical for modern circuits due to their physical size, cost and also (in this case) susceptibility to hum pickup. So that's question (1) sorted.

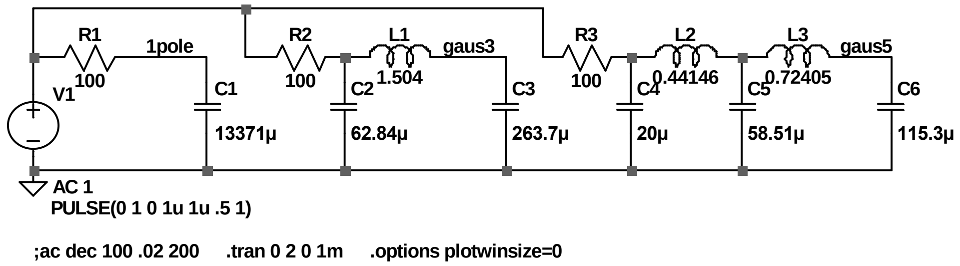

Figure 4: impedance-scaled single-, 3- and 5-pole “Gaussian to 12 dB” LRC networks.

So let’s go ahead and realize the 3-pole and 5-pole filters in the DC-free topology. The constraint that the resistor chain across the top must have a value no higher that 100 ohms dimensions the rest of the network, because the cutoff frequency is set by the stopband requirement. Here we must ‘cheat’ a little for the moment. That’s because the DC-free topology causes additional circuit elements to appear in the equivalent LRC passive filter of this type (see “Filter DC Voltages...”). Working back through the transforms that were applied to derive the DC-free configuration, one can show that the additional elements are the equivalent of a resistor in series with each inductor. There are standard (from the point of view of the 1950s and 1960s, anyway) techniques for designing such filters, but a quick search turned up no tables. For an initial simulation experiment (figure 5), we’ll dimension the component values around the op-amp blocks to render the effects of these components small. It also allows us to use nice low capacitance values, which is obviously attractive. This will cause some problems in a real filter, and we’ll come back to that later.

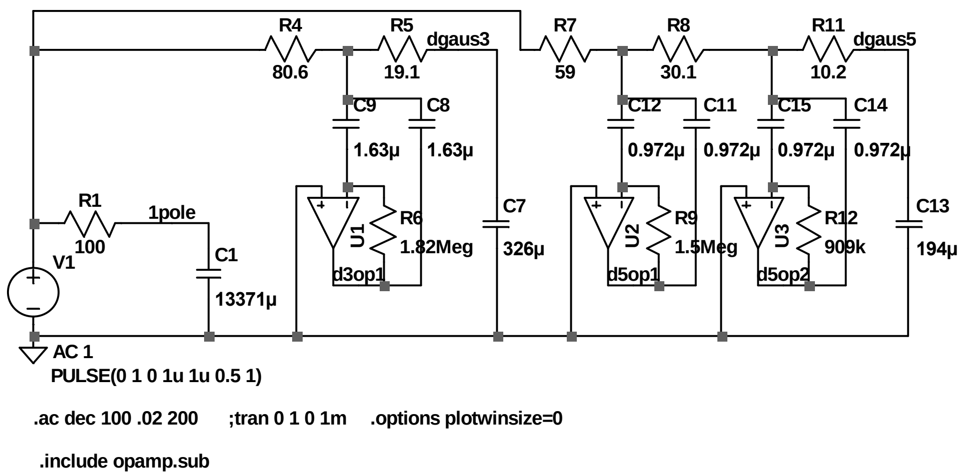

Figure 5: Trial DC-free implementations of the 3-pole and 5-pole filters.

Note that all the capacitors in the circuit have the voltage to be filtered on one terminal, and are at ground potential on the other. This means we can use 200 V rated polarized electrolytic capacitors.

The capacitors in the D-elements have been set quite small (the resistor round the amplifier “takes up the slack” to create the correct D value). This ensures that the response is close to the desired value. But the hidden problem here can be revealed by looking at the gain at the op-amp outputs – and seeing the huge peak (figure 6)! Running with such small capacitor values in the almost-D-element, the gain of the hidden bandpass filter – it’s still there, even though we’re using its **input** properties to give us our lowpass response – becomes large. That’s question (3) sorted out.

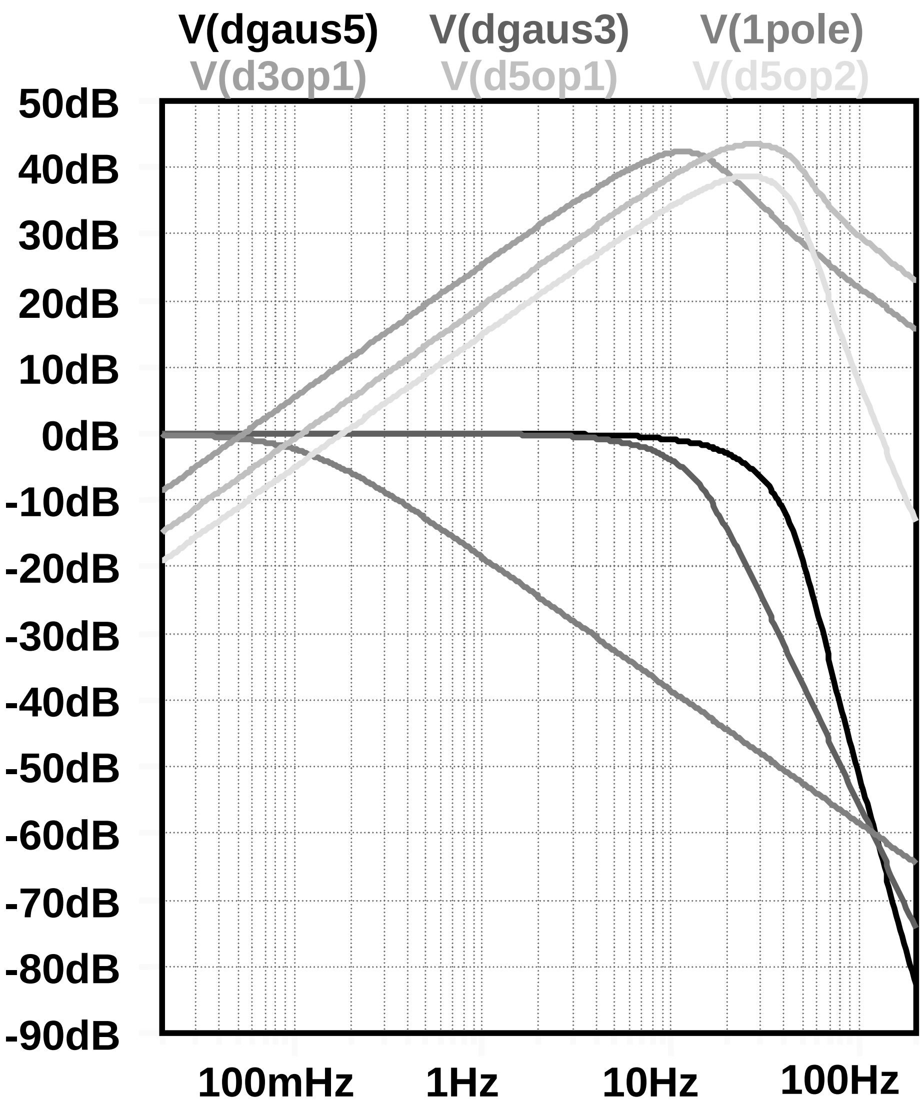

Figure 6: shows huge response peaking at figure 5’s op-amp outputs.

The AC ripple on the voltage we’re filtering will cause the amplifiers to clip on their modest 5 V supplies. To make this circuit viable, we do have to allow these capacitor values to be quite large, indeed of the same order of magnitude as that output capacitor,

to pull the gain of these sections down. This means doing a custom design. Call for the Million Monkeys!

We can use the spreadsheet optimization method introduced in "Million Monkeys" (or, perhaps a PERL optimization script wrapped around our favourite simulator) to generate a new set of component values that retains the exact "Gaussian to 12 dB" response, while reducing the peaking at those op-amp outputs by allowing the capacitor values to be larger. If we're going to do the work anyway, what other tweaks could we make to our response? How could we further increase the effective cutoff frequency, reducing settling time and total capacitance while preserving the precious 120 Hz rejection?

Well, to get the most stopband attenuation for a given passband bandwidth, you need a filter with stopband zeroes. These are the features that look like 'notches' in the response at particular frequencies. Such filters are often loosely called 'elliptic' filters, but that term only actually applies to a specific type of filter having flat passband and stopband response and the maximum possible number of ripples in both bands. Stopband zeroes help to 'pull' the stopband response down quickly. The tradeoff is that the stopband rejection doesn't become indefinitely larger as frequency rises, but ends up bouncing around, ideally staying below some defined minimum level, in our case -60 dB.

Standard tables exist for some filters like this (e.g. in [1]) but in that case, they are not only not 'dissipated' (the term to describe passive filters designed for inductors with loss) but are also only given for doubly-terminated networks, which is no use to us here, as discussed in "Lowpass Filters that Don't".

So I gave an Excel Solver spreadsheet the job of fitting a 5-pole DC-free filter circuit (or rather, the prototype uses to develop the filter) to a response that tracks the Gaussian amplitude profile out to 9.5 dB at the standard 'normalized' cutoff frequency of 0.159 Hz, and that is then down at -60 dB at slightly less than twice that frequency.

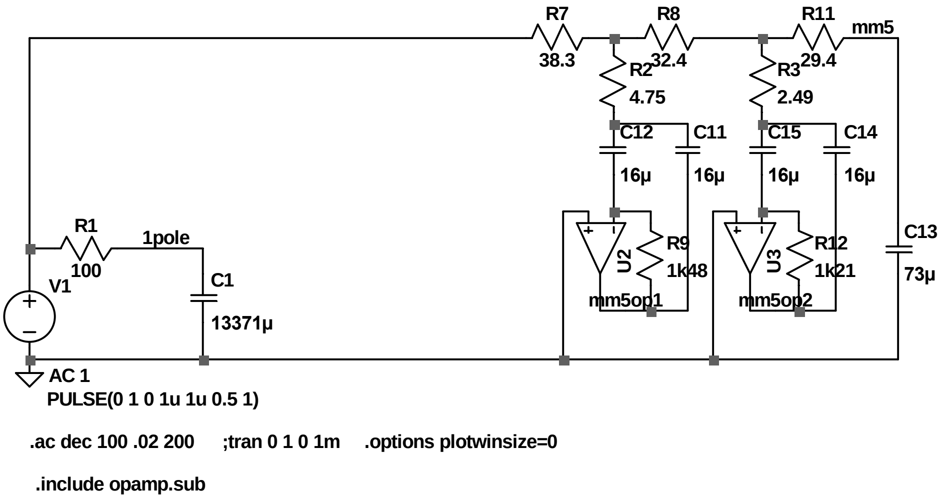

Figure 7: DC-free filter derived from our "Million Monkeys" optimization.

When scaled and transformed into our final filter, we end up with the circuit of figure 7. The schematic doesn't show the surge protection diodes or TVS devices that you really should use in any circuit like this, to prevent the enormous voltage switch-on surges from damaging the low voltage-rated op-amp inputs).

Stopband zeroes are introduced into the prototype with a small capacitor across each inductor, which translates into a small resistor in series with the input to our almost-D-element. Despite the approximations inherent in trying to form stopband zeroes with this filter structure, we easily reach the stopband requirements. I tried several values of 'dissipation factor' and ended up with one that gave nice round numbers in the final filter. The Excel result and its simulated response are shown in figures 8 and 9.

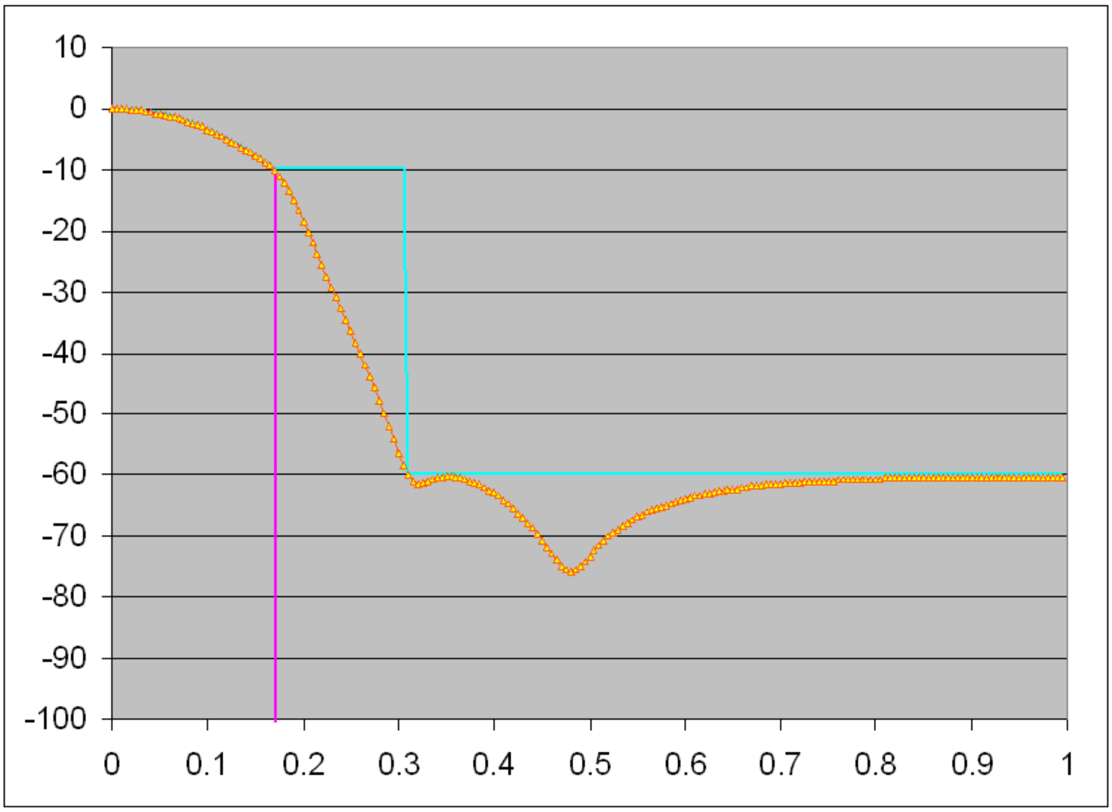

Figure 8: the Excel plot of the fitted prototype filter response.

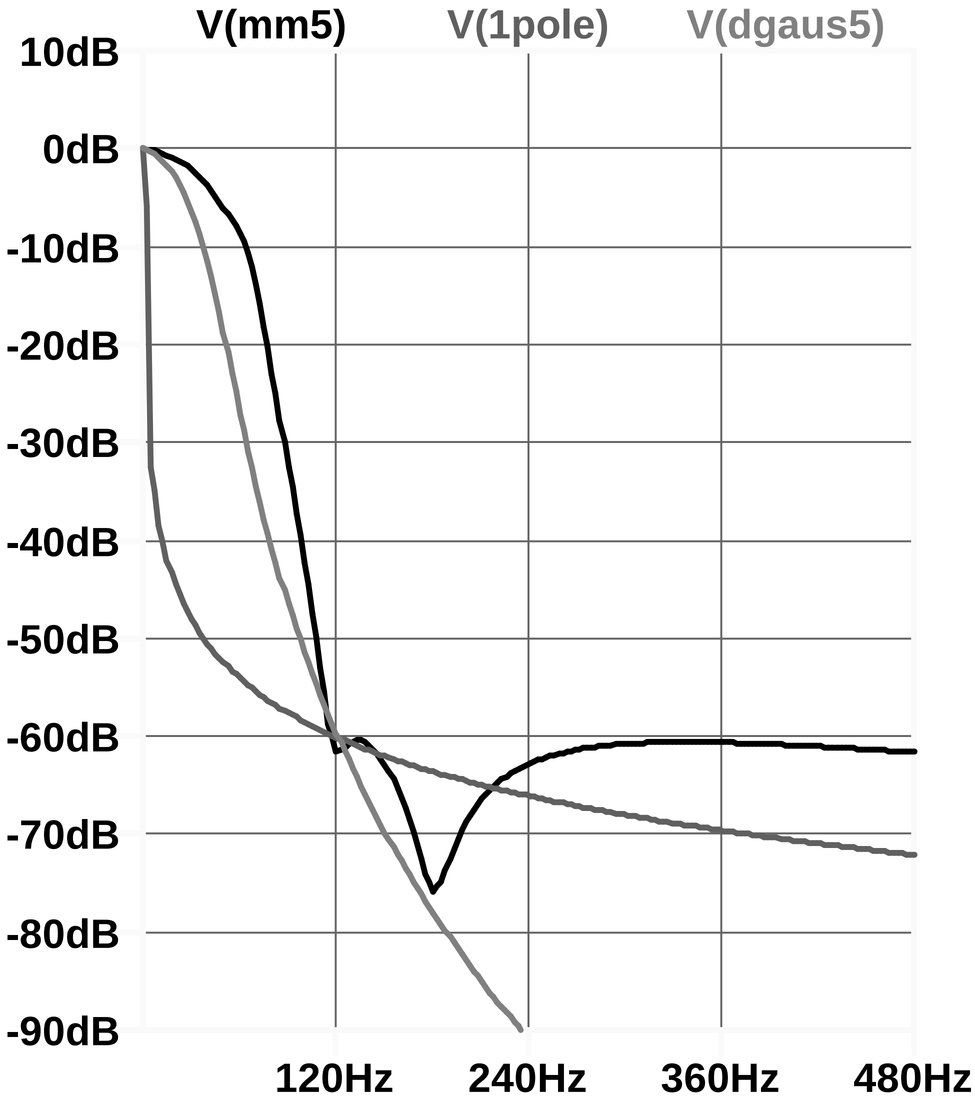

Figure 9: The new filter response, along with single-pole and 5-pole Gaussian.

The settling behaviour, shown in figure 10, is a bit ‘nervous’ in comparison to the Gaussian filter but still easily beats it, at 58 ms to <0.1%. If you’re happy with 1% settling it is **much** faster than the Gaussian version (18 ms versus 46 ms).

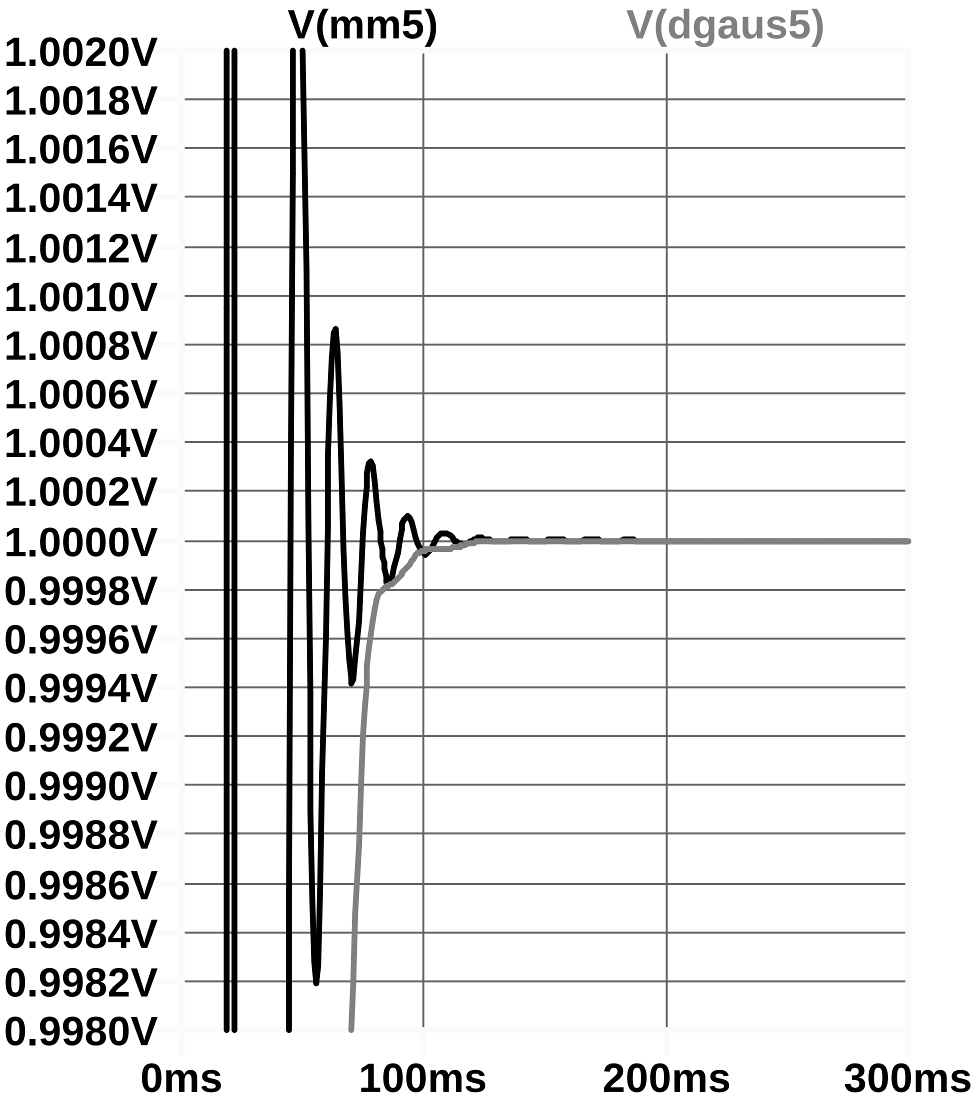

Figure 10: settling behaviour comparison, new filter to 5-pole “Gaussian to 12 dB”.

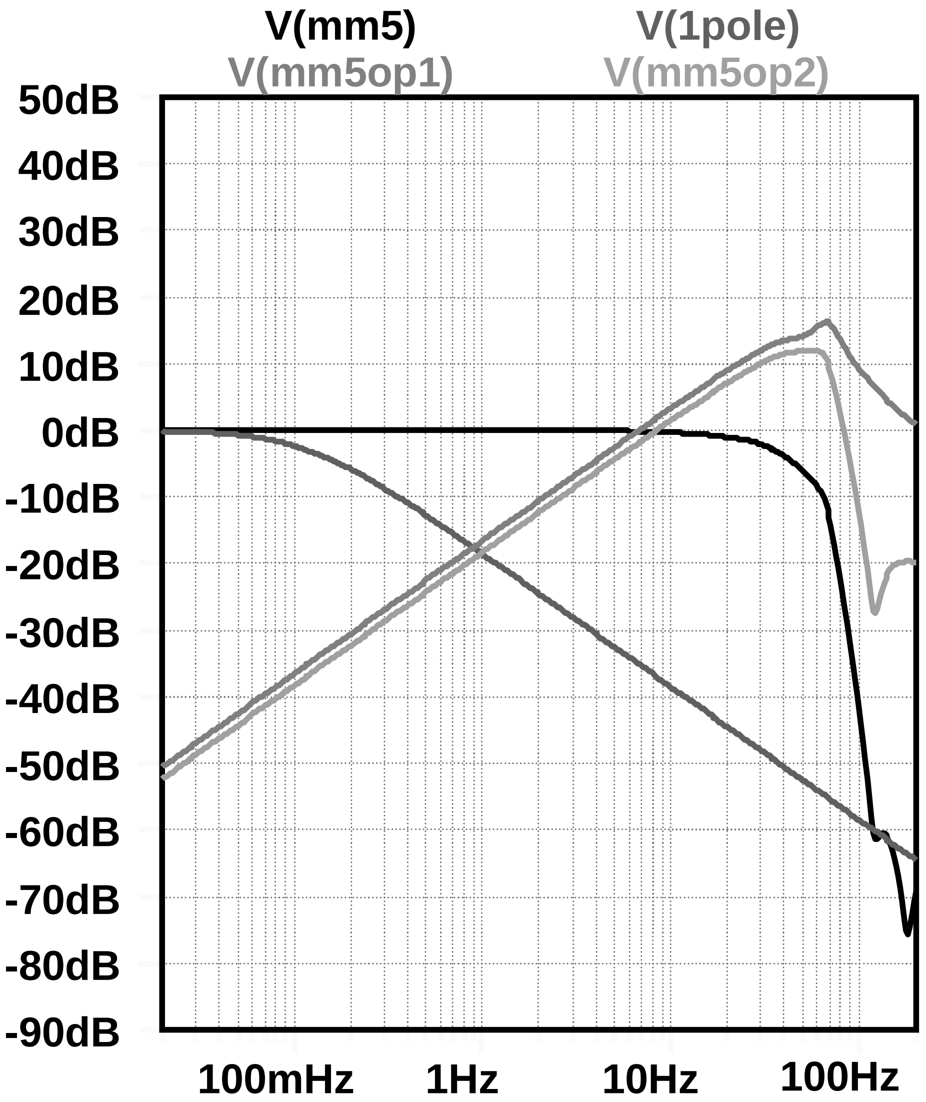

Figure 11: still a little peaking at the op-amps, but not nearly so bad.

We haven't completely conquered that boosting at the amplifier outputs, as figure 11 shows, though it's much better controlled, and we'll be able to cope with our 1 Vpp 120 Hz ripple with no problem. More intervention in the optimization process could probably improve this a little, and perhaps trim a little more total capacitance off (maybe also compensate for any ESR effects). Nevertheless, the capacitors shown add up to only 137 uF, not far short of a 100x reduction in comparison to the single-pole case. Dude, we **might** actually get it to fit!

Will it actually work? Figure 12 shows the op-amp output voltages and currents when running on a ripple waveform on 1 Vpp created by a full-wave rectified power supply feeding a 1000 uF reservoir cap and a 1.37 kohm load. The voltage swings are well inside the 5 V supply rail (the non-inverting inputs are biased at half the supply voltage,

of course) and the required output current is safely inside the 25 mA rating of the op-amps in PSoC3. And the ripple level on the output is too small to show on that scale of graph. So, yes, I think it will!

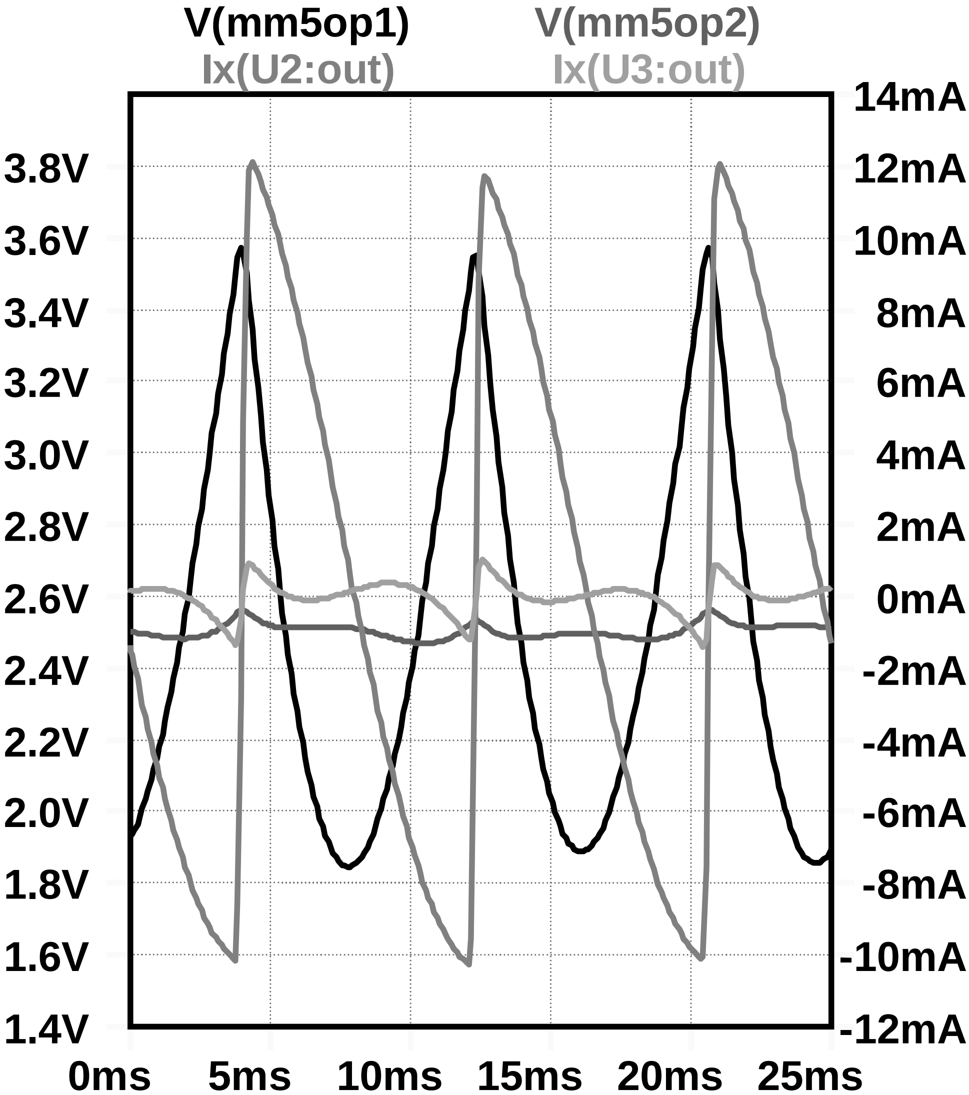

Figure 12: Op-amp voltages and currents in operation.

Are there other ways of doing this? I'm happy to receive all your homework submissions and will publicize the best one. But I don't yet know where I'd start if I couldn't solve the problem this way. Happy (rapid, low-capacitance) bias filtering! / Kendall

[1] Williams & Taylor, "Electronic Filter Design Handbook", McGraw-Hill
[2] Zverev, "Synthesis of Filters", Wiley