The Filter Wizard
issue 22: Active filters that still work even when their amplifiers don't
Kendall Castor-Perry

An addendum to an earlier article in which the Filter Wizard forgot to do a very important simulation. The apparently undesirable behaviour predicted by that simulation turns out to not be a problem after all, and further demonstrates the robustness of the design approach.

In my haste to get "A Fast-Settling Bias Voltage Filter with High Ripple Rejection" out there, I forgot to cover an important aspect, that on the face of it might look to be a show-stopper. It turns out that it's not so big a deal, but does demonstrate something quite interesting about that class of filter design.

To recap, in that article I presented a 'DC-free' filter that not only effectively suppressed ripple on a very high voltage bias line, but also settled very quickly following changes in that voltage. Do check the article out for the full story if you've not already done so.

In that article, Figure 10 proudly showed the settling behaviour of a couple of variants of that filter topology to a step change in the input voltage. As you can see from the units, the particular step I was studying was a unit step, i.e. a 1 V change in input voltage. However, I implied a rapid settling time also to the much larger input step that results when the actual 180 V bias voltage we were filtering was turned fully on and off.

Let's look at the predicted outputs of the op-amps in the circuit when we apply a full-sized 180 V step to the input. First of all, let's use a simple non-powered op-amp macromodel. Whatever simulator you use, it'll have one of these – you can even use the simplest of all amplifier approximations, a voltage-controlled voltage source or SPICE 'E' component, to make an ideal op-amp for this purpose. It's a great way of abstracting the ideal behaviour of a circuit with only the simplest of amplifier characteristics.

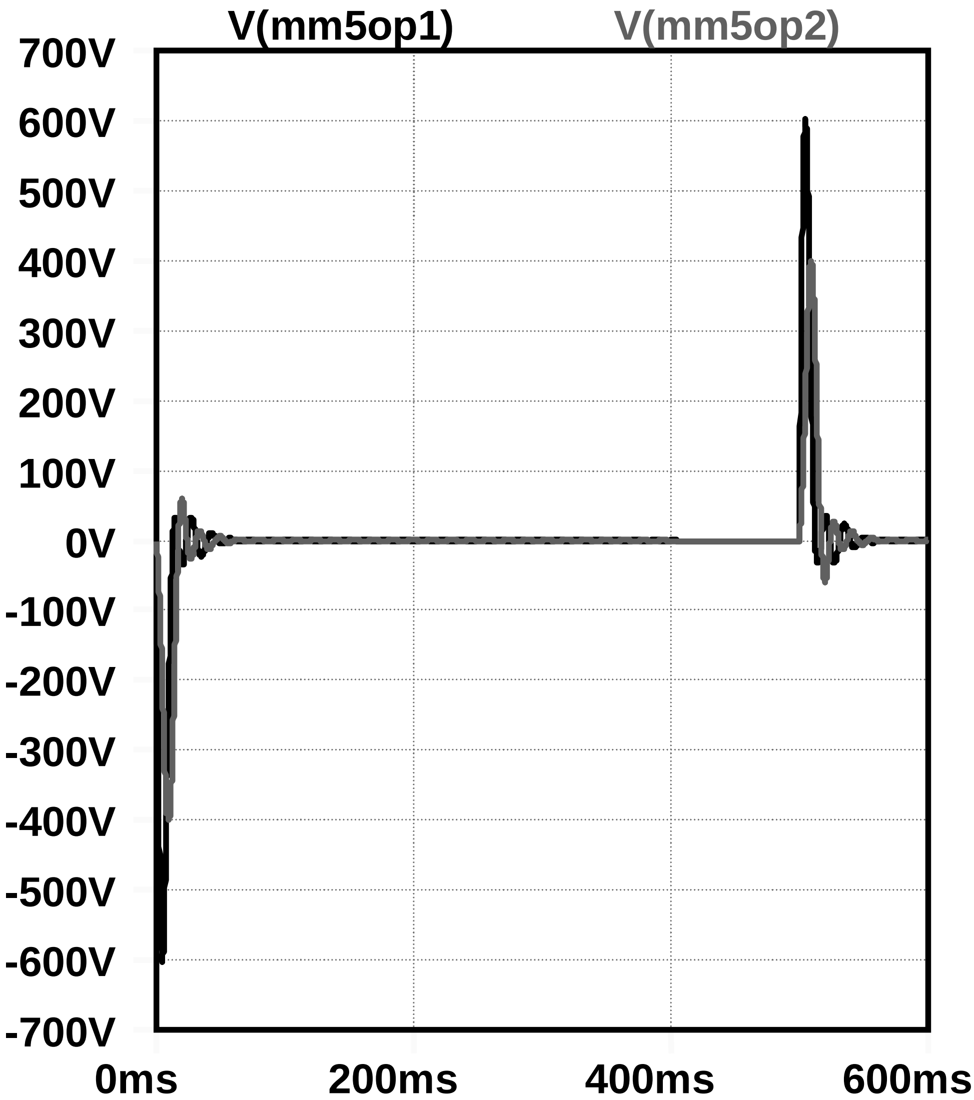

Figure 1: what the amplifier outputs would do with no supply voltage limitations...

Oh, dear. Figure 1 shows the result. Our simulations predict a negative spike of 600 V at the output of each op-amp when the input voltage steps up from zero to 180 V. A same-size positive-going spike follows when the input voltage steps back down to zero again. Now, fellow engineers, it doesn't take class-leading smarts to work out that this kind of voltage excursion isn't available from real amplifiers that are being powered from a 5 V supply! In practice, the outputs can only go close to ground or 5 V, and they will then saturate. As a result, the inverting inputs of the amplifiers, previously held at 'virtual earth' as a consequence of the feedback resistor, will be pulled away from that state of electronic grace by current flowing through the capacitors connected to them. Eventually the inverting input voltage will hit the clamp voltage set by your input protection scheme. Speaking of input protection, in the earlier article I did mention that the op-amp terminals, which are exposed to large voltage changes through the capacitors, need to be protected

with tough devices, such as chunky Zener diodes and surge suppressors – just good basic practice when handling voltage excursions outside your supply rails. These need to withstand quite significant currents during that step, if the source can supply it.

Surely this input and output saturation is fatal for the filter's performance, right? You usually expect that an op-amp-based circuit becomes pretty non-functional if the amplifier outputs hit the rails. Well, once this has happened with our circuit here, the circuit clearly can't work as intended because the amplifiers have an incremental voltage gain of zero and aren't doing anything. The circuit just looks like a ladder of resistors and capacitors, and it won't have either the complex pole pairs or the stopband zeroes. It will longer to settle and it won't have the necessary stopband rejection.

To check the behaviour when we put some clipping diodes in to protect things, I put a 4.7V Zener diode across each op-amp output and each inverting input. My simulator only had a model for a rather weedy 500mW-rated diode, but it soon became clear that you might need quite a high current rating here. Also, to ensure that all the error current does go through the diodes, I put a 40 ohm resistor in series with each op-amp output and gave it a 25mA current limit. You could also usefully put a resistor between each summing node and the actual amplifier inverting input.

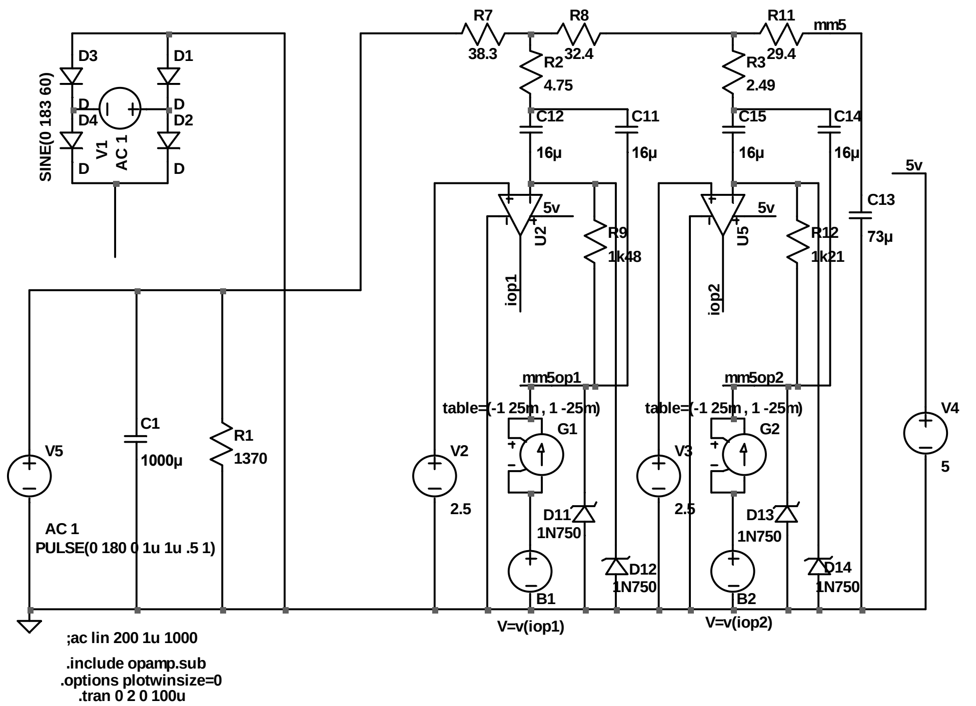

Figure 2: a test schematic with some voltage and current limiting.

Figure 3 shows the voltages we get at the amplifier outputs; much more reasonable. Though the excursion down to -1.4V shows that we must be putting a lot of current through those Zener diode models. And sure enough, when you look, you see predicted spikes of nearly 2A, assuming that your power supply can actually deliver that and that it

ramps up instantaneously (rather unlikely, admittedly). I need some power Zener diode models!

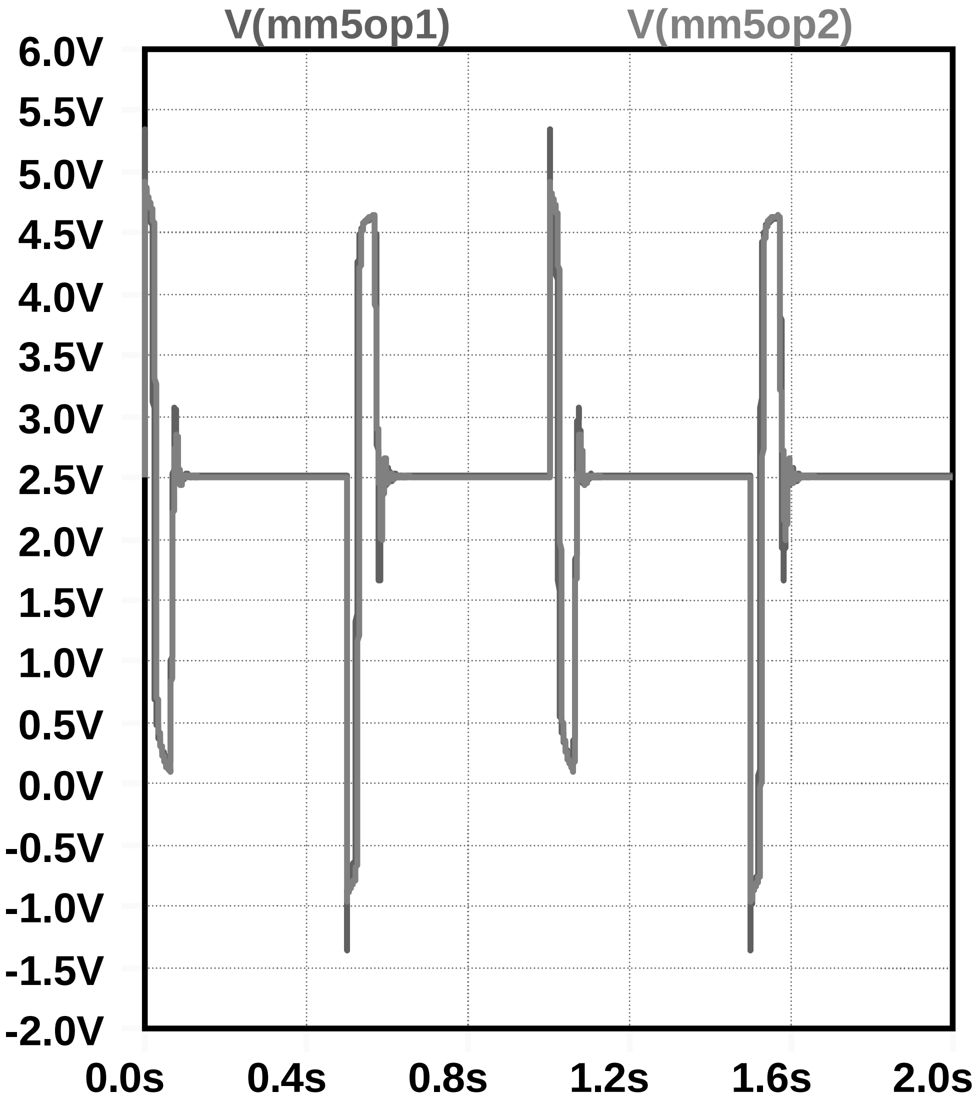

Figure 3: diode-clipping op-amp inputs and outputs to the power supply rail.

However, the actual filtering story is actually not that bad at all. The 'sting' from the huge input step dissipates rapidly, absorbed by the transient protection diodes. The amplifiers come back out of saturation rapidly, because the clamps at the summing nodes hold them close to the supply. When this happens the transfer function, and therefore the residual settling behaviour and stopband rejection, is restored. The entire system still settles to 0.1% in a speedy 67 ms and the behaviour of the circuit is actually now a lot less 'nervous' in terms of the ringing on the input step (figure 4). The settling time to 1% is slower than the 'ideal' circuit, but this wasn't a requirement for our original circuit anyway. The whole thing actually looks nice and clean from the outside, even though we know that terrible things are afoot internally.

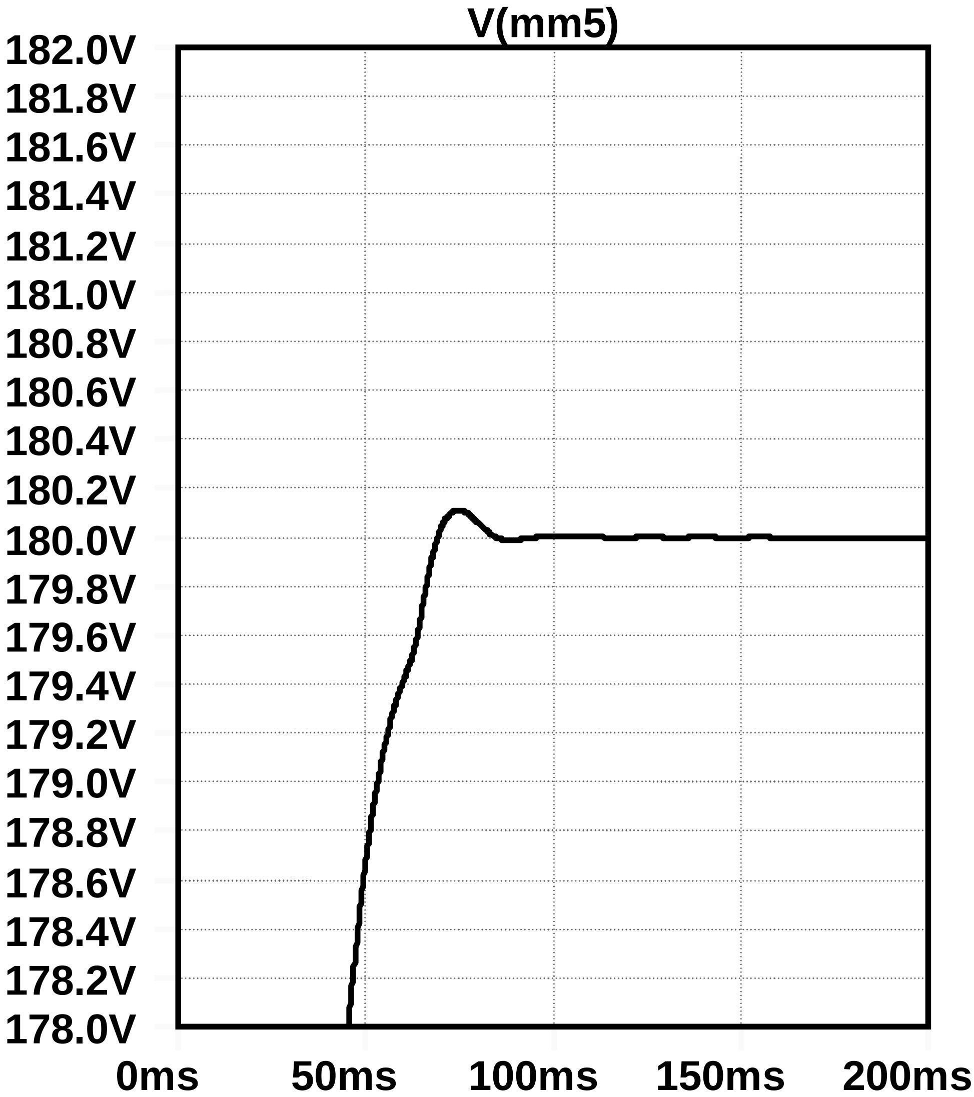

Figure 4: Even with saturating amplifiers, the filter still does a great job.

When our input voltage jumps up instantly from zero to 180 V, the amplifiers spend a certain amount of time in a non-functional state, but the inherent settling time of the RC formed by our fairly small total capacitance still allows an excellent settling time while they are AWOL. Once the step has worked its way out of the system, all the amplifier terminals are operating at their correct values and the promised AC ripple rejection is obtained for steady-state operation. Most active filter topologies fall to pieces when their amplifiers overload – some even go unstable. The robustness of this particular configuration is another aspect that drives me to recommend it so highly.

Other points to note? We're not limited to positive voltages – negative ones are fine too, as long as you remember to connect up any polarized capacitors correctly. And the

voltage doesn’t have to be DC; any AC signal with frequency sufficiently far below the cutoff frequency can be filtered in the same way. The source impedance can be reduced without limit too, as long as you can ensure that the amplifiers can deliver the current necessary for signals near the cutoff frequency. I’ve already had an audiophile amplifier designer get excited at the prospect of using this configuration in front of a voltage regulator, to ensure that input noise falls at a faster rate than the regulator loses supply rejection. Do plenty of simulation to ensure that under normal operating conditions you don’t get an excessive amount of the input signal appearing at the amplifier outputs.

So, my ‘umble apologies if it looked like I wasn’t aware of, or worried about, real-world behaviours with real-world components. It turns out that even though the amplifiers stop working for a short while on that step input, the overall filter can still do the job intended. See if you can come up with some other ideas for where this approach can be used. Ripple begone! best / Kendall