The Filter Wizard
issue 23: Which Filters are Noisier – Analogue or Digital? Part 1
Kendall Castor-Perry

"Give me a lever long enough", Archimedes is said to have observed, "...and I shall move the world". He didn't have the benefit of our modern understanding of materials, which suggests that this boast might be hard to fulfill without a generous supply of Unobtanium. If you drew Archimedes's experimental setup on paper, it might look just like a ball of rock being lifted on the plank of a see-saw (that's a teeter-totter, to US readers) but there's a little matter of scaling. Let's say you double the size of everything – the radius of the rock, and the width and thickness of the plank. The 'strength' of the plank increases by a factor of four (it goes as width*thickness, roughly) but the load on the plank goes up eight-fold. Keep scaling up, and eventually you'll exceed the capabilities of any plank material.

What's the connection with noisy filters? Well, it turns out that things here don't scale quite the way you might expect, either. In this two-part Filter Wizard, we'll look at some fundamental noise mechanisms in filters, using SPICE to illustrate the performance limits you can expect. We'll concentrate on analogue filters in this part, with the biggest bombshell being reserved for part 2's look at digital filters. Yes, digital filters generate noise too! And sometimes in unexpectedly, unacceptably large amounts. We'll see in part 2 how SPICE noise simulation can be used on digital filters, permitting a direct apples-to-apples comparison of analogue and digital filtering solutions.

By the way, if you've got a practical interest in low-noise filters, I recommend that you get Doug Self's new book "The Design of Active Crossovers". While it's focused on a specific audio application, the material has wide applicability, and it's a great new addition to the canon of real-world filter cookbooks. In this article I'll take a slightly more fundamental look at noise issues in analogue filter design, to set the scene for an eventual show-down with digital filters. Which will win? How is it going to end? You'll just have to wait for part 2!

I'm going to concentrate on lowpass filters, whose noise bandwidths are essentially determined by their filter responses. Also true of bandpass filters, this isn't the case with highpass and bandstop filters, whose noise levels are generally determined by the bandwidths of the amplifiers used.

So, where does noise come from in an analogue filter? Well, some of it comes from the opamp(s), for sure. Each active filter topology has its own particular 'noise gain', which causes the inherent input voltage noise of the opamp to be frequency-shaped in a way that is related to – but not identical to – the actual shape of the signal transfer function that the filter creates. This is a fundamental insight useful across all of circuit design – that a circuit doesn't always process its own internal noise in the same way that it processes the inputs you apply to it.

Many opamps have not only an input voltage noise source but a noise current source as well. The consequence is that any finite source impedance attached to an input of the filter's amplifier, and therefore passing this noise current, creates an additional noise voltage that also contributes to the overall filter noise. Each topology again has its own signature here, also dependent on the apparent impedances of the resistor-capacitor networks that are hung on the amplifier.

Last, but definitely not always least, is the fundamental noise contribution from the passive components themselves – resistors and capacitors, since we're not considering any exotic active filters that include inductors. Now, there are two ways of looking at this noise that might seem to be diametrically opposed, but in fact spring from the same well of physical law.

One common perspective holds that Johnson noise in resistors is the only source of noise energy in a circuit, and that capacitors are noise-free and only affect circuit noise behaviour through the interaction of their frequency-dependent impedances with the circuit resistances. Resistance to current flow is a bulk property of matter; anything made out of matter is at a non-zero absolute temperature, and hence the thermal energy in the structure of a resistor 'smacks' free charges around and causes a varying noise potential across any two points in the conductor. This gives answers that match experiment, and is the way that SPICE calculates noise in a circuit. In a SPICE noise simulation, the contribution of the Johnson noise of each resistor in turn at the output is calculated using a linear AC analysis, and all the noises are squared and added. The square root of this sum is the overall rms total noise voltage.

There's another approach, though, used frequently by IC designers. The need for it can be deduced from the thought experiment in which you allow the value of a resistor to tend towards infinity. Since the Johnson noise voltage density of a resistor is proportional to the square root of the resistance value, it's clearly troublesome that the noise voltage of an infinite resistor calculates out as infinite. It may be OK for theoretical physicists to juggle with infinities as part of their day jobs, but in the everyday engineered world we don't see high voltages appear across the terminals of very high value resistors.

In this second approach, we ignore the resistors and concentrate on the capacitors. The usual way of introducing this idea is to consider a resistor R and a capacitor C in parallel, and contemplate the noise voltage of the combination. The total noise voltage is the product of two pieces: the noise voltage density, which is proportional to sqrt(R), and the square root of the measurement bandwidth, which is proportional to 1/sqrt(RC). So the total noise is proportional to 1/sqrt(C) and the value of the resistor doesn't even enter into it! As the resistor value changes, the spectral density of the noise changes, but the total integrated noise value is constant. This applies for arbitrarily high values of resistance, which of course result in arbitrarily low values of bandwidth, and noise voltages that could be moving very slowly indeed.

This is not to say that capacitors actually **generate** noise. The residual uncertainty of the voltage on a capacitor in parallel with an infinitely large resistor (the combination hence having zero bandwidth) is an indication of the Johnson noise that it was exposed to at some time in the past, when the voltage was defined by contact with a circuit in which the resistance was finite. The capacitor has sampled the Johnson noise of the resistor in the circuit that charged it, and now it's holding it for you.

What's interesting is that this behaviour generalizes to **any** RC network, including active ones. If we hold constant the capacitor values in an active filter, and scale the cutoff frequency with those resistors that determine it, the total noise contribution from the passive network stays constant.

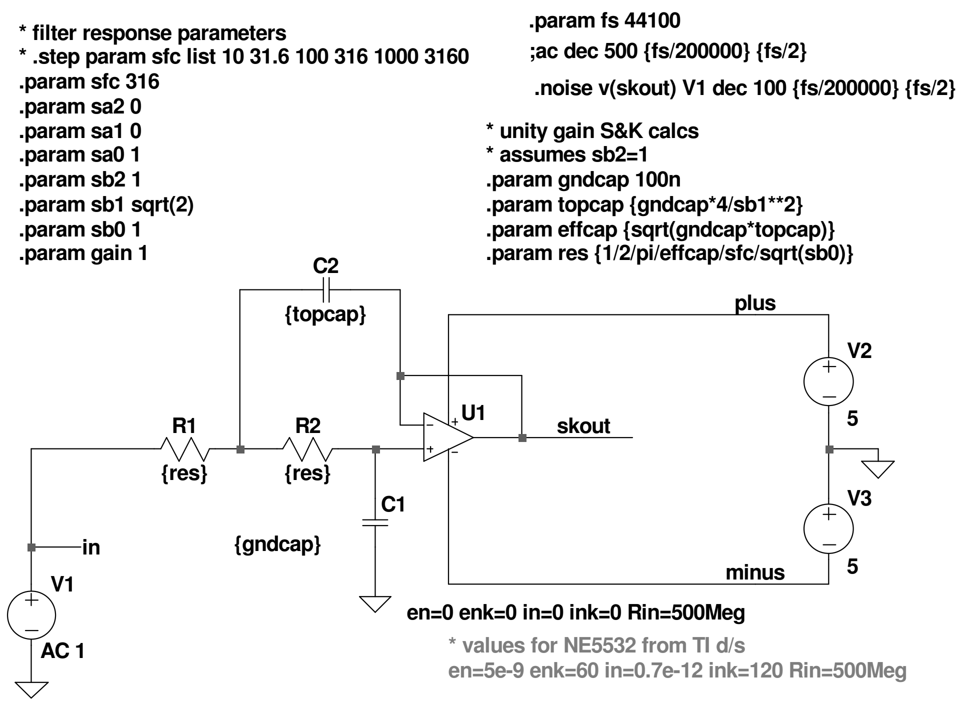

Figure 1: Simple lowpass filter with values calculated on the schematic.

We can show this constancy readily by simulation. Figure 1 shows a straightforward Sallen & Key lowpass filter (2nd order Butterworth) whose values are computed right there on the sheet for a stepped set of cutoff frequencies. The grounded capacitor value is fixed at 100 nF, and the "round the top" capacitor is calculated in the spreadsheet (it will be 200 nF for the Butterworth alignment chosen). These two capacitors are held constant in value as the frequency is stepped, and a new pair of resistors R1 and R2 is calculated each time. Figure 2 shows the frequency response of the filter for cutoff frequencies from 10 Hz to 3160 Hz. Figure 3 shows the noise density versus frequency, employing a noise-free amplifier. They all integrate out to the same value, 287 nV(rms) in the test bandwidth, which runs from 0.2205 Hz to 22050 Hz (audio engineers, you know why).

To me, when thinking about how noise levels can change as cutoff frequency choice is varied, the constant-capacitor approach feels more 'fundamental' than keeping the

resistors constant and changing the capacitors (which **would** change the total noise). Knowing that the total noise is locked to a constant value takes one factor 'out of the equation' in your system design. And, thinking back to the Archimedean scaling metaphor, this is one example where scaling the parameters in an experiment makes **less** of a difference than you'd expect.

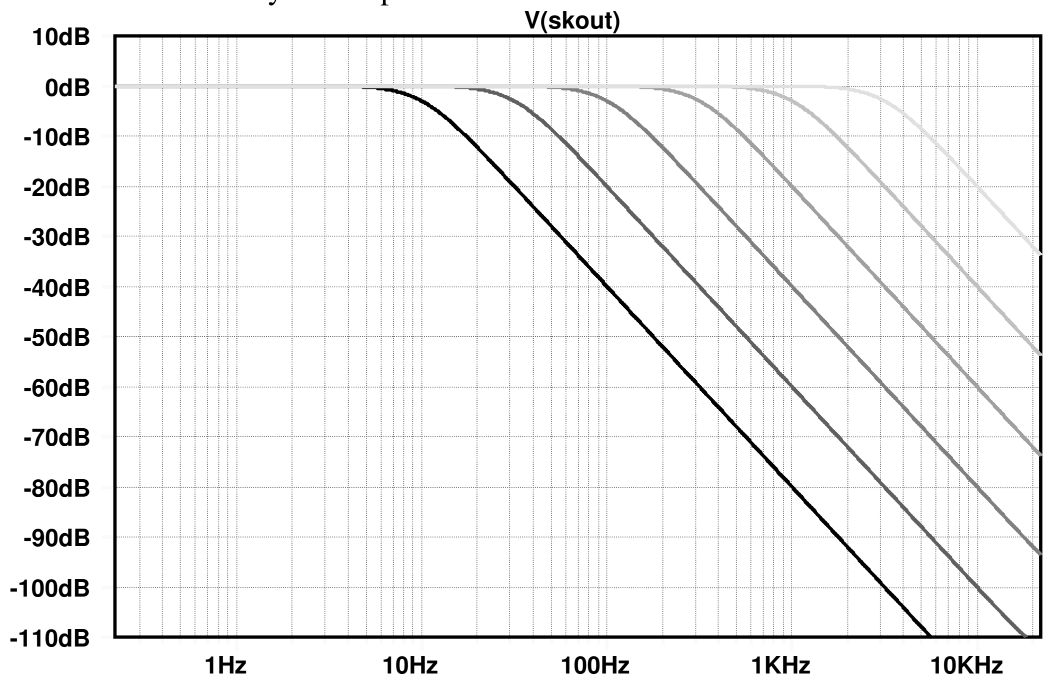

Figure 2: Stepped frequency responses of the lowpass filter.

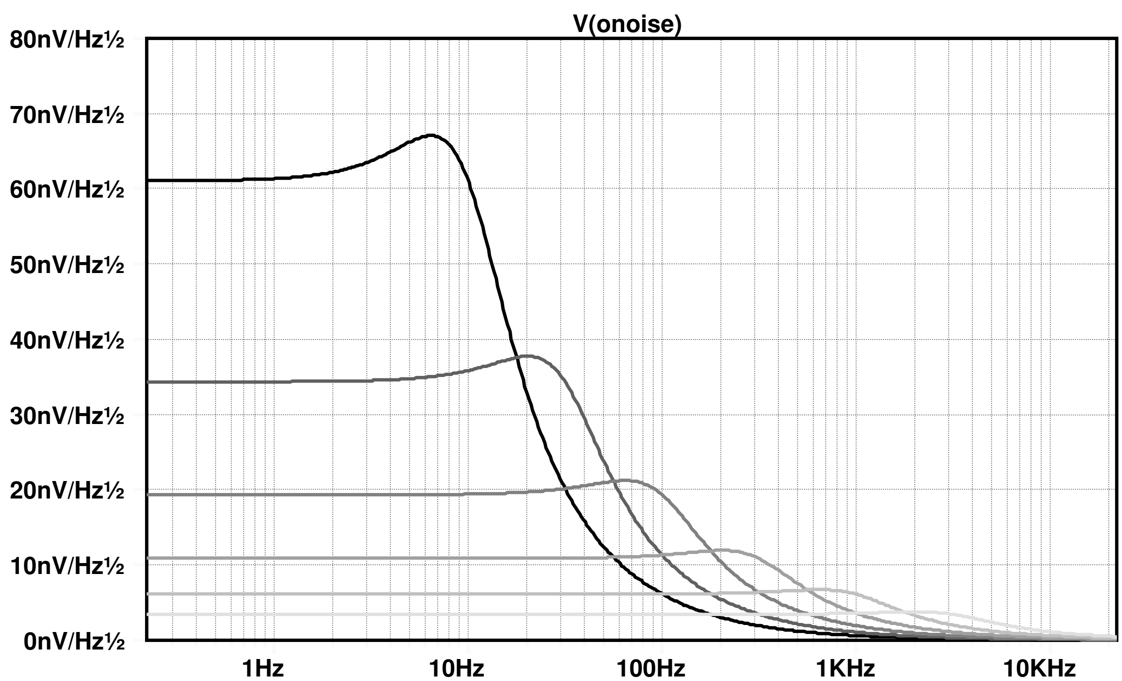

Figure 3: Noise spectral density versus cutoff frequency. Same area under all curves.

Of course, things get more complicated when you use a 'real' amplifier with its voltage and current noise contributions. Each term adds its own 'footprint' to the noise signature. The voltage noise of the amplifier creates a noise contribution whose value rises with the square root of the noise bandwidth it 'sees', which in turn is proportional to the filter's cutoff frequency. Meanwhile the current noise at the amplifier inputs generates an equivalent voltage noise term that's proportional to the impedance level, which is inversely proportional to cutoff frequency. So this contribution **rises** with falling cutoff frequency. The presence of 1/f noise in both the current and voltage terms creates yet more complexity. All these noise terms can be accommodated in the simple 'universal opamp' macromodel included with my preferred SPICE (LTspice). Table 1 shows the integrated noise values simulated at the filter output as progressively more of these mechanisms become non-zero. The values used were interpolated from the datasheet of the NE5532, a deservedly popular audio-grade opamp (and favourite of Doug Self).

cutoff frq	noisefree	en, no 1/f	en & 1/f	in, no 1/f	in & 1/f	all noises
10	286	797	812	700	3607	3686
31.6	287	799	815	461	1346	1546
100	287	807	822	352	564	954
316	287	828	843	309	344	864
1000	285	890	904	292	296	907
3160	278	1043	1054	281	281	1055

Table 1: simulated noise in nV(rms) of figure 1's circuit for various opamp noise terms.

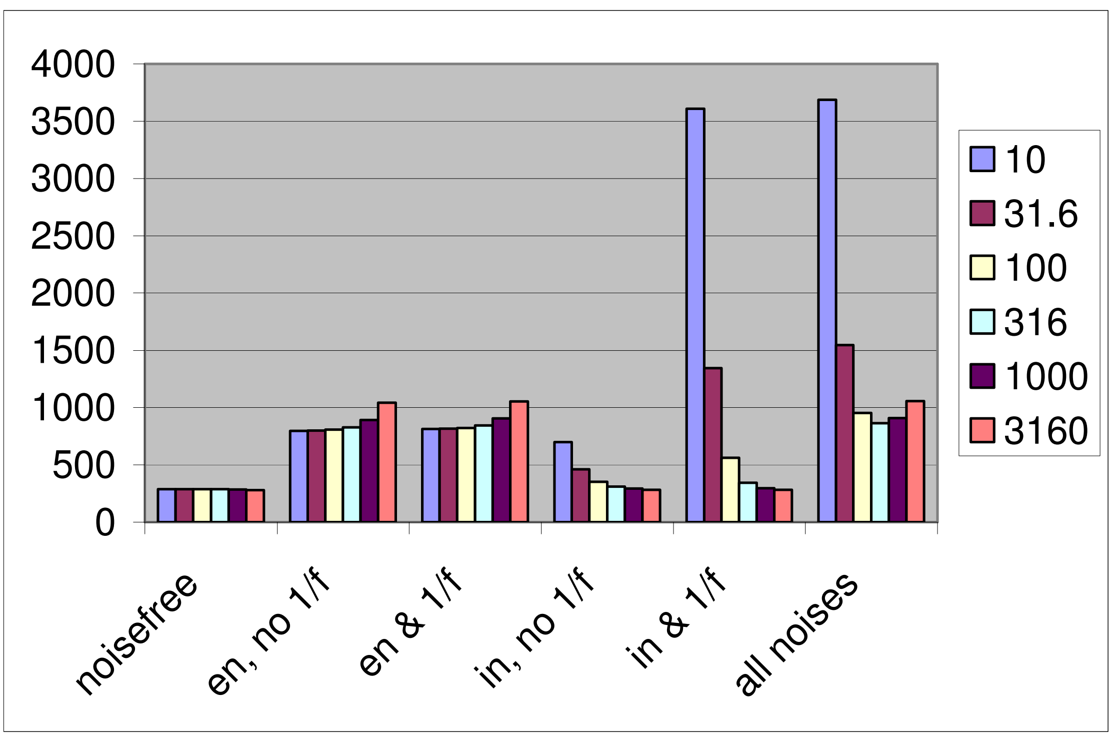

Figure 4: A bar chart of the various noise contributions in nV(rms).

Figure 4 shows table 1's results in a business-like bar graph presentation. It's easy to see the trending of the voltage (en) and current (in) noise terms, plus the dominant contribution of 1/f noise in the current at low cutoff frequencies. There's indeed an 'optimum' cutoff frequency at which the total rms voltage noise is at a minimum.

This approach has given us a baseline case for how 'noisy' a competently designed analogue active filter can be. In the next installment, we'll look at how we can calculate and simulate the noise level of an equivalent digital filter in a meaningful way. We'll see that there's a driving mechanism for noise in a commonly-used digital filter topology, that it can't be ignored, and that scaling effects will bite back in a big way. Meanwhile, keep things quiet! / Kendall