Yet More On Decoupling, Part 4 – Don't Get Into A Macromuddle!
Kendall Castor-Perry

Previously on "Yet More..." OK, that joke is wearing a bit thin. By the end of part 3 we had just got to the point of seeing that when we apply a current excitation to our modelled power supplies, the voltage variations we see on those supplies punch right through to the op-amp output in a quite predictable way. You want to see what that looks like in the time domain, don't you! So do I, but this is the middle third of the story and there has to be an encounter with a scary monster.

Yes, it's time to start using the models of real amplifiers (if that isn't a contradiction!) supplied with LTSpice, and also imported into it from other vendors. It would be disingenuous of me to pretend that I wasn't expecting trouble from these models – and my fears turned out to be justified.

Lies, damn lies and op-amp models

Op-amp vendors don't supply you with the transistor-level circuits of their products; that would give their secrets away to the competition and would demand a great deal from your simulator if you were, for instance, analysing a filter with 12 op-amps, each with 70 or more transistors. They provide 'macromodels' which aim to replicate all the essential behaviours of the real device using as few nodes as possible, for faster running. These models are supposed to be much more accurate than the very simple amplifier archetype which I used in part 3 to demonstrate some home truths about amplifiers. Testing one amplifier after another in the simulator is easy and quick – and initially created great disappointment; the results that came from the simulation were physically impossible.

Figure 4.1 reprises results from the made-up amplifier we looked at in part 3. Closed loop gains of -1 to -1000 are tested, both for the realistic but worst-case in-phase supply modulation (the set of curves asymptotic to 0dB) and the unrealistic (but commonly used by manufacturers) symmetrical modulation (the lower ones). I'm doing both sets for each amplifier to see what this reveals.

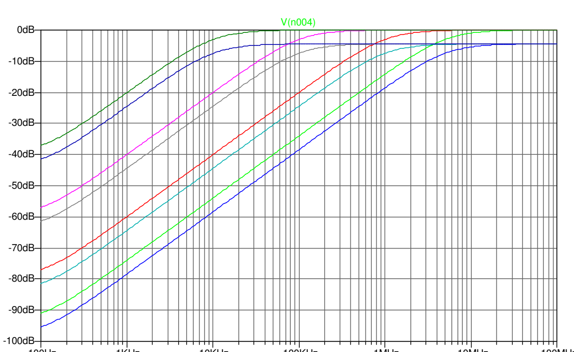

Figure 4.1(L): the closed-loop PSG for part 3's made-up ideal amplifier

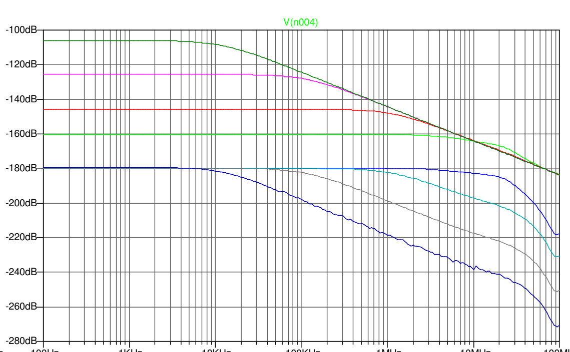

Figure 4.2(R): the open-loop PSG for the LT LT1355. Notice the dB scale

Figure 4.2 shows the results using the model of the LT1355, a nice op-amp I've used in the past, with extraordinary slew-rate (400V/us) for its moderate 12MHz GBW. It looks like a plot of some completely different parameter! These numbers are both physically impossible and inconsistent. The OLPSG is down at around -170dB, a figure so small it is likely not due to attributes of real modelled components but to interaction between SPICE conductance defaults. When feedback is applied, the low frequency supply rejection actually *degrades*. It then moves in the *wrong* direction with frequency, *improving* rather than getting poorer. Something is clearly amiss here. Clearly, this LT1355 model will be no use in my search for what is happening at the output of the amplifier when the supplies are modulated.

Eventually I found a macromodel which did work in my test fixture, and the terrible truth dawned on me: that vendors generally don't see the accurate modelling of power supply sensitivity to be important. In other words, putting most of their op-amp models in my power supply simulation and applying signal to the supplies would *not* result in a realistic change to the op-amp's output voltage. Not until the 14th of 18 devices (fast voltage-mode op-amps) tried from the LTSpice library did I find one whose supply rejection had the right form. Models downloaded from other vendors' websites often fared little better, so it wasn't just a Linear Technology issue.

Here's an example which does seem more reasonable: the LT1723 quoted in part 3, shown in figure 4.3. The traces for in-phase excitation are realistic (though on a close-in look they do show a bit of peaking, this may be a CMRR modeling interaction) and also show the effect of the finite open loop gain at low frequencies. This will let through some of the low frequency stuff on the supply rails. The lower traces for symmetrical excitation are down at a vanishingly small level (~-180dB), implying that the model treats +ve and –ve supply variations identically and they cancel out. We know this is not so from the data sheet quoted in part 3. This was the first device that went into the 'pass' bin.

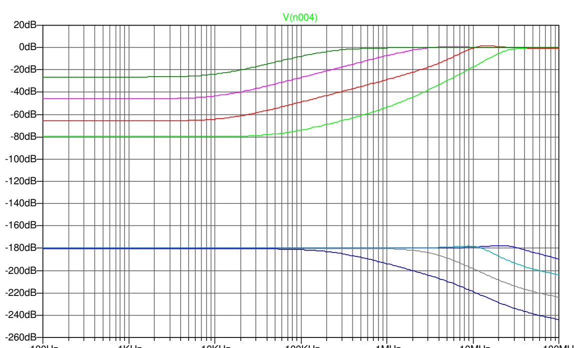

Figure 4.3(L): the closed-loop PSG for the LT LT1723 model, also showing the effect of finite low frequency gain

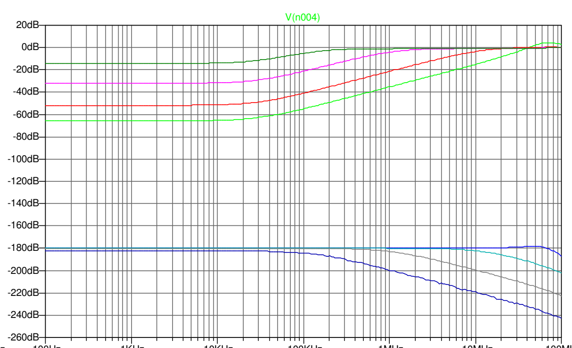

Figure 4.4(R): the closed-loop PSG for the LT LT1812 model.

Another good performer (though also with the unrealistic symmetrical results) was the LT1812 shown in figure 4.4. This has even less low frequency loop gain available, and does go a bit haywire in the upper frequencies. But it also passed this first test.

The next few figures show the variety of results obtained. Most promising 'candidate for perpetual motion' went to a TI-supplied model for the TLE2082 (figure 4.5). At low frequencies the PSG rises, and increased with increasing closed-loop gain as expected. But above 100kHz, in-phase PSG just keeps on rising and rising, reaching +60dB at 100MHz and still rising. In other words, a GBW of 100GHz from a device they rate at 8MHz; I'm very familiar with this device and it doesn't do this!

Again TI, and the model for the popular OPA350 (figure 4.6) showed, for a change, differing sensitivities to the two supplies. One warning though – the OPA350 isn't supposed to be used at a supply voltage of +/-5V, and that's another thing that can go wrong with simulations. Some of these models all seem quite happy to run at supply voltages well in excess of the device's rated value.

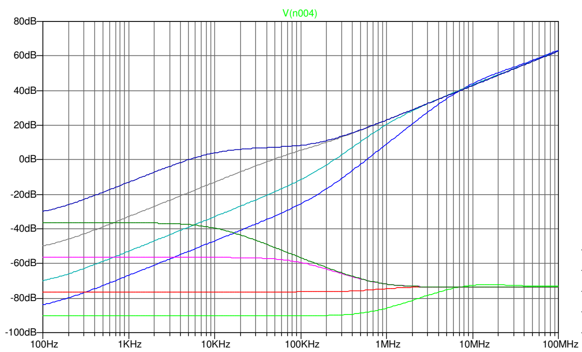

Figure 4.5(L): TI TLE2082 model PSG with closed loop gain of -1 to -1000; record-breaking GBW via supply pins

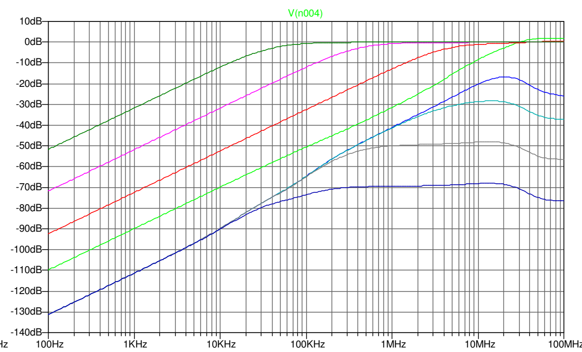

Figure 4.6(R): TI OPA350 model PSG with closed loop gain of -1 to -1000; some thought went into this one

Most disappointing, given the reverence that the device enjoys in the analog world, was the OPA627, the Burr-Brown-birthed Rolls-Royce of op-amps. The behaviour of the model (figure 4.7) as the closed-loop gain was swept combines several competing effects which cause the curve to pivot around as the gain goes up!

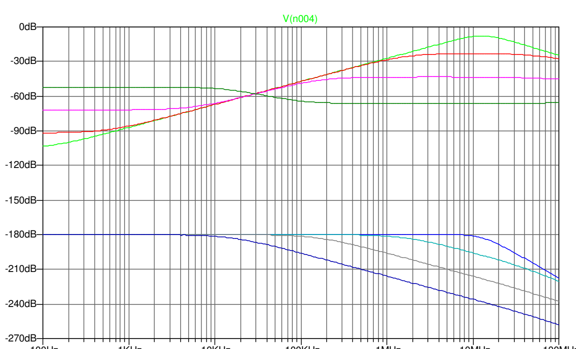

Figure 4.7(L): TI OPA627E model PSG with closed loop gains of -1 to -1000; strange pivoting at higher gains

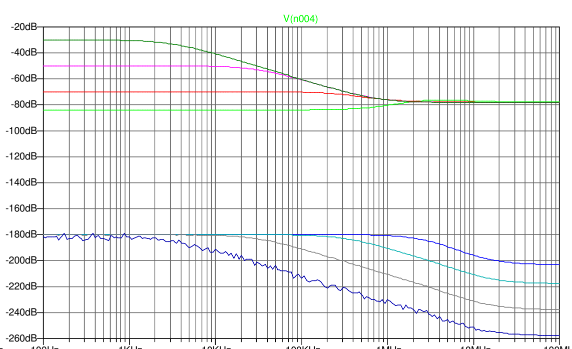

Figure 4.8(R): The TI TL072; behaviour at DC is plausible, but high frequencies are completely wrong

At the other end of the ‘quality’ scale, TI’s humble TL072 – and I have used very many of them so I’m not knocking them – showed behaviour which at DC looked pretty plausible, but was completely erroneous at high frequencies, figure 4.8. I can understand why you would build a model which didn’t implement supply gain at all. But why go to the effort of putting in the components needed to do *something* with the power pins yet do it wrong? Rrrrr...

A similar device from Analog Devices, the ADTL082, initially showed a jumble of behaviour, with huge low-frequency closed-loop PSG. However – good for them – they jumped so quickly on seeing an early draft of this article, accepting the criticism, that it would be unfair to print the first effort. A good model is promised in the very near future. The AD826 (figure 4.9) from the same stable showed impeccable manners, and behaved as expected.

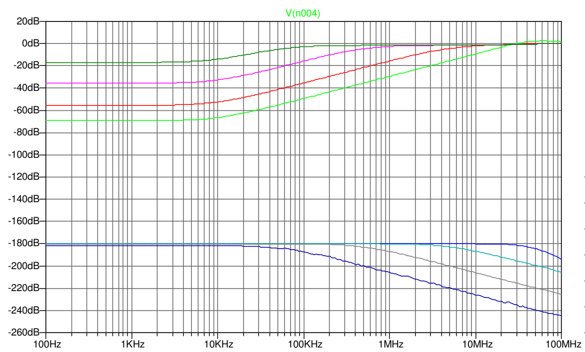

Figure 4.9(L): ADI AD826 model PSG with closed loop gain of -1 to -1000; phew, one that just about works!

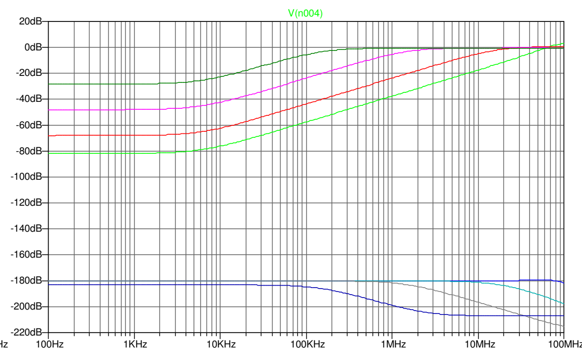

Figure 4.10(R):The Intersil (ex-Elantec) EL2227C model PSG with closed loop gain of -1 to -1000

To round off the manufacturer showcase, Figure 4.10 shows the Intersil EL2227C, good results (though it seemed to run slower than most others on this AC test). Intersil have some great amplifiers, some inherited from past masters such as Harris and Elantec. The difficulties I had with some of their models highlighted two particular issues:

1. not all SPICE syntaxes are quite the same, and some very old Harris models use some non-standard invocations;
2. SPICE model files should be ASCII text files, but some Elantec models had binary characters in. If you open them in Internet Explorer and then save them, the Unicode text file format is selected, which is not compatible with some text-entry simulators.

Hats off to the LTspice team at Linear Technology who rapidly made the necessary changes to SwitcherCad so that you don’t have the same problem if you try this yourselves at home!

Armed with devices (just a handful of them!) which showed plausible supply rejection results, I girded myself up for the next challenge, which was to explore what happens at the output of an op-amp in the time domain when the supply voltage is pinged, first by an external signal and then, finally, by the actual load current the amplifier is being called on to supply. It turned out that the pain inflicted by aberrant device models was not yet over. To be continued...

Takeaways from this part:

- Power supply gain can easily be simulated, if the amplifier's simulation model implements it correctly – but most models from most vendors don't
- Just because the op-amp is good doesn't mean the vendor's simulation model will be
- You'll encounter occasional syntax and file format errors in model files, especially if they are a bit old