Yet More On Decoupling, Part 6 – steering in the right direction
Kendall Castor-Perry

This is the last of the "Yet more..." series. One final plot twist, a moderately satisfying ending, and my work here will be done. Until the sequel, that is...

Let the op-amp do the pinging!

Instead of exciting the supply rail impedance with a test current source, let's drive a signal into the amplifier, so that the load current itself causes the excitation on the supply rails. There must still be supply voltage variations resulting from this, and we know they can affect the amplifier output. This has profound implications for sampled data systems and would explain why, so often, the fine tuning of such a high precision system involves fiddling about with not only amplifier choice but also supply decoupling.

The test fixture stays the same but now we set up the input voltage source to provide a +/-1V signal. Because the amplifier has a gain of -1, we'll invert the input signal when plotting so that we can overlay it on the output and examine the settling behaviour. This signal causes a +/-10mA current to flow in the load, which is a combination of the 111.1 Ohm load resistor and the 1kohm feedback resistor. Now, that feedback resistor doesn't return to 'real' ground, only the nearly-at-ground input of the amplifier, but the error introduced by this is going to be a small fraction of the other errors we're looking for.

To add a pinch more realism, a series RC network, 22pF in series with 100 Ohms, was added to the load. This is typical of the load such an amplifier might have to contend with – ADC input stages and the occasional 'scope probe. There's no little capacitor across the feedback resistor, for two reasons. Firstly, the best value would depend on the amplifier, and secondly I wanted to look at CFAs as well, just out of interest.

Now, it's clear that the very simple macromodel introduced in part 3 isn't going to work in this circuit; the output current isn't supplied from the power pins. In any normal amplifier, when +10mA is being delivered to the load, it is flowing out of the output pin and must be flowing *into* the +ve supply pin. Conversely, when the load current is -10mA, this current flows into the output pin and comes *out* of the –ve supply pin. Each pin therefore should inject a current demand of 0 to 10mA on its supply rail, phased like the way we tested the circuit earlier. So, if the amplifier model obeys these rules, the excitation of the supply impedance should be comparable to the tests in part 5. Let's try some of our models and look at the waveforms. In the busy figure 6.1, the most important traces are the grey and pink traces showing that this model is indeed taking an extra 10mA from the +ve rail when the output voltage is +1V, and -10mA from the negative rail when at -1V. On this scale the supplies aren't moving much, but when we zoom in (figure 6.2) and apply some offset to get both rails on the same scale (400mV top to bottom), we see the expected voltage step and frizz on the supply rails, together with the beginnings of some activity on the amplifier output (green trace).
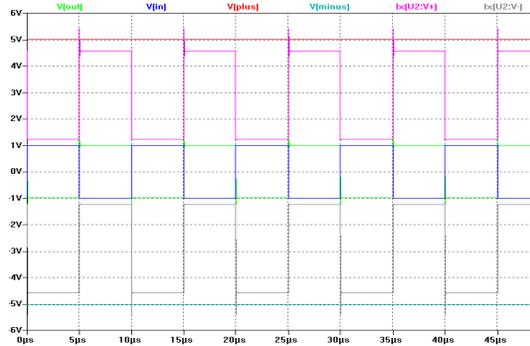
Figure 6.1(L): waveforms from the test rig driving a load, showing op-amp supply currents (grey and pink traces)
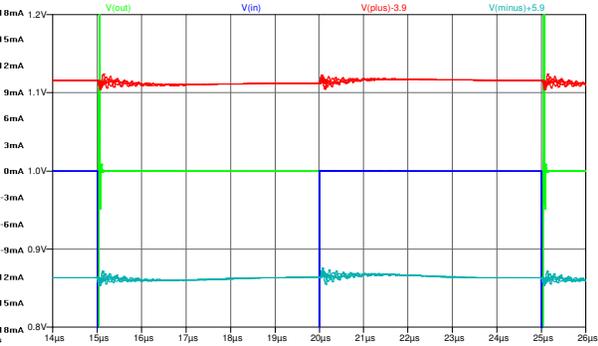
Figure 6.2(R): zooming in on figure 6.1 to show that the supply voltages are moving.

This is with the ceramic caps on the regulators. By the way, if you probe the current in the printed circuit trace, you'll see significant ringing rather than the fairly clean square wave. In other words, the sum of the currents in the trackwork and the local decoupling capacitor is clean, but the transient current demand from the amplifier provokes high frequency circulating currents which are flowing in the power feed, and which could get all over your circuit board., without an obvious voltage cause.

The final plot twist

It's rogue's gallery time again because, you guessed it, many models out there don't let you do this simulation; they don't steer load current from the power pins, but from a magical, everlasting internal battery supply.

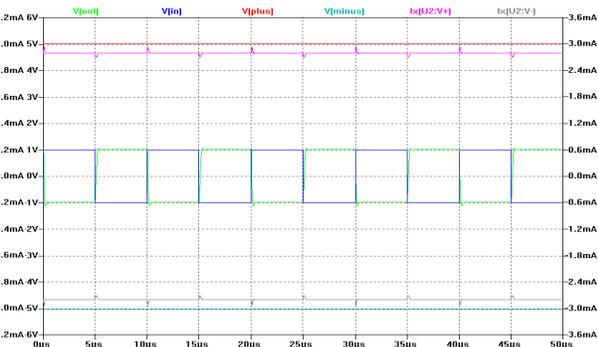*Figure 6.3(L): LT LT1354 model, taking a constant current from each rail (but not the same value)*

Figure 6.4(R): LT LT1055 model; some little high frequency squeaks modulating the supply current

Figure 6.3 shows common behaviour, from the LT1354 macromodel, which didn't model supply rejection either – I included it because it's a good amplifier in the real world! Although the amplifier is certainly delivering current into the load, the supply pins only consume a constant, quiescent current. And the 1mA current that goes into the +ve supply pin doesn't all come out of the –ve supply pin, about 100uA goes missing. Mr. Kirchoff *will* be upset... This behaviour was quite common among models from several

makers. The LT1055 (figure 6.4) is similar, but with a noticeable dynamic twiddle of a current peaking on the signal edges, superimposed on a static current drain.

Several TI models failed this test. The TLE2082 (clear winner of the "impossible PSG" competition earlier in the series), sat there analysing so slowly that it was going to take almost a week to complete (I saw that behaviour with some ex-Elantec devices also). The OPA350 gave very strange results for 22nF decoupler, with the +5V supply voltage collapsing, despite no unusual current being taken from the regulator. At longer timescales, LTspice wasn't able to complete the analysis. Perhaps using excessive supply voltage has an impact in transient SPICE simulations after all.

Winner of the "does it do *that*?" competition this time was TI's OPA2604E model, for a device beloved of audio designers. For some decoupler values, it analysed fine for a few microseconds and then spontaneously *burst into oscillation* on the 'screen' in front of me, as if I had touched it somewhere sensitive with a high capacitance probe! Next thing we know, we'll have thermal feedback from the simulator to the keyboard so you really *do* burn your fingers when the circuit goes wrong! But I digress..

Whoa, *yeah*! Change the decoupling, change the settling time

To see what's actually happening at the op-amp output we need to zoom in further. Let's do that and try some other amplifiers to see what they look like, still with the livelier ceramic regulator output cap. In Figure 6.5, the LT1723 shows some evidence of the ringing with minimum decoupler coming through, at about 1mV p-p. The time delay between the negative-going edge of the blue trace (input) and the bouncing up of the green trace (output) indicates how long it took the device to respond. This is affected by both slew rate and closed-loop bandwidth, which in this configuration we'd expect to be half the GBW. The LT1812 in figure 6.6 has a lower bandwidth but higher slew rate, and the green trace "gets there" more quickly, indicating that at this signal level the slew rate is dominant. Ringing with 22nF is 2mV p-p, which is 0.1% of full scale.
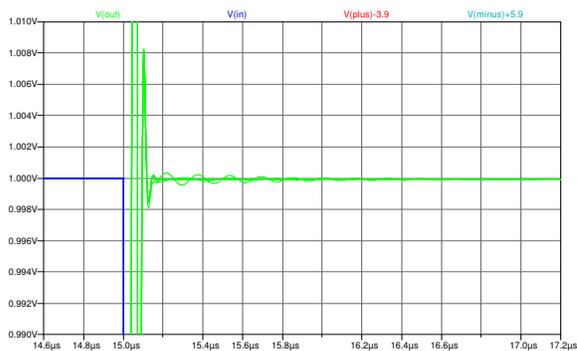

Figure 6.5(L): zooming in on the output of the LT1723 to show the ripple, which is worst for the 22nF decoupler
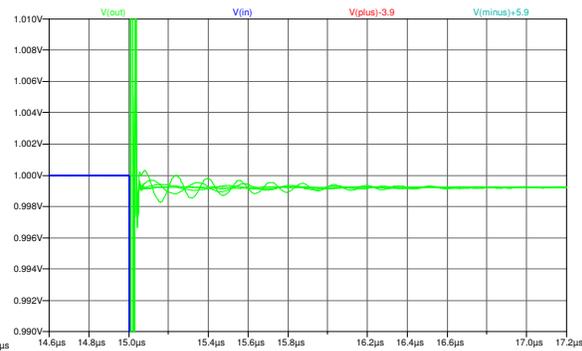
Figure 6.6(R): looking at the higher slew rate but lower GBW LT1812 on the same scale as figure 6.5

What about slower amplifiers? Figure 6.7 shows the LT1632. Note the significantly higher level of ringing now. In fact, if you look at a PSG plot for this model, you can see

3dB peaking before it settles down. If this peaking coincides with a supply impedance peak it will make the ringing look worse than it would do if the model were correct. Figure 6.8 shows the even slower LT1630. It's the slowest one here, yet is not showing quite so much ripple. Why is that?
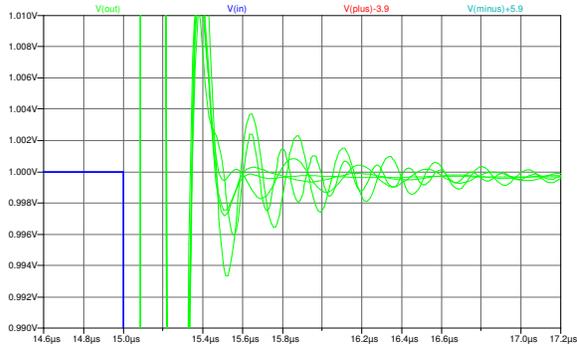

Figure 6.7(L): the output of the LT1632 model, whose slightly peaky PSG may be exaggerating the ringing a little
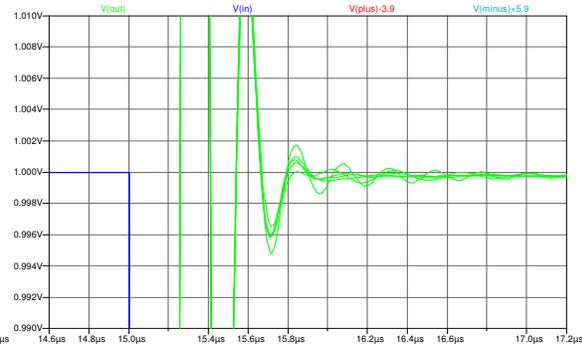
Figure 6.8(R): the output of the LT1630 model. The current demand is so sluggish that the ringing is reduced

Now we start seeing something interesting   The LT1630 takes so long to settle down, and the current excitation's edges are so slowed down, that despite its poor bandwidth, and hence supply rejection, the output ringing is not that severe (once the output has actually settled). The closed-loop small signal response and the large-signal slewing behaviour can combine to 'take the edge off' (literally!) the output voltage, hence the load current, hence the supply excitation. The longer it takes the output voltage to reach the final value, the gentler will be the excitation of the supply resonances.

What about our secret weapon for reducing all this ringing, the tantalum output capacitor for the regulators? Well, it's hardly surprising that it has a marvellous effect here; figure 6.9 shows what's left on the worst-case LT1632 output. Now, the value of the decoupling capacitor really doesn't seem to be having much important effect. The equivalent trace for the LT1723 is now so clean and fast that it's not worth printing; we're going to have to zoom even further in on both amplitude and time.
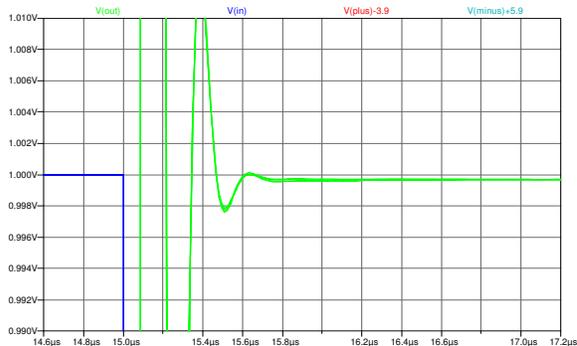

Figure 6.9(L): LT LT1632 settling cleanly with tantalum regulator cap. Compare to figure 6.5.
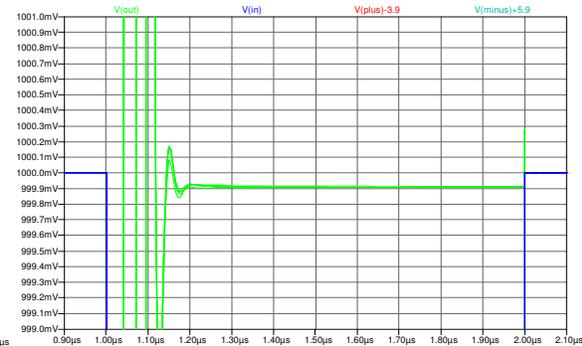
Figure 6.10(R): LT LT1723 settling with tantalum regulator cap, 500kHz squarewave, 100uV/div

Figure 6.10 shows the step response of the LT1723 model with a full scale of +/-0.1%; figure 6.11 shows the LT1812 on the same scale. Now, the LT1812 is an amplifier with a much higher slew-rate than the LT1723 (750V/us versus 70V/us), but a lower gain-bandwidth product, than the LT1723 (100MHz versus 200MHz, so the data sheet says, and the PSG sweeps bear this out). So the LT1723 takes longer to acquire the input signal; it also has a lower initial response ring frequency because on the first few parts of the ring it's still in slew limiting except at the turning points. But the combination of less rail excitation (probe the fixtures yourself!) and more power supply rejection means that the differences between the various decoupling capacitors (swept 22nF to 470nF as always) are much less pronounced than on the LT1812.
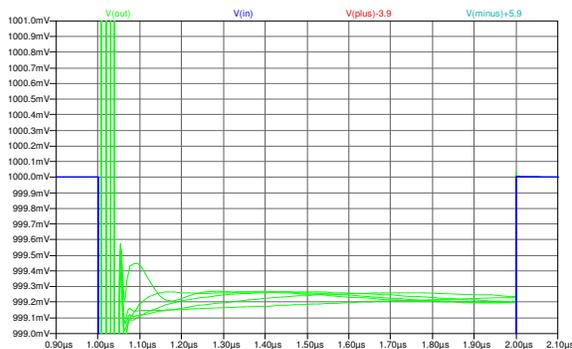

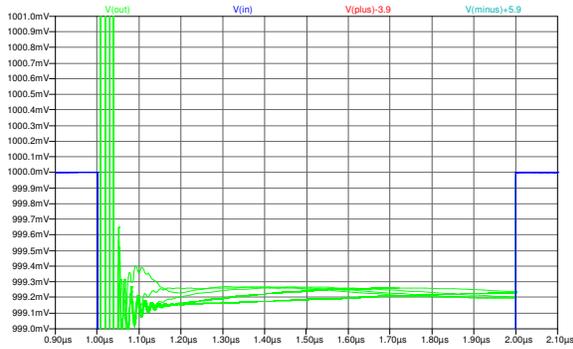*Figure 6.11(L): LT LT1812 settling with tantalum regulator cap, 500kHz squarewave, 100uV/div*

Figure 6.12(R): as figure 6.10 but with an extra 4.7nF capacitor as decoupling

Did you also notice that the actual settled output value is lower with the LT1812 than with the LT1723? The LT1812 has rather a low open-loop gain in comparison to the LT1723, and here's an example of where this can be relevant. The settled value of the output step is dependent on the low frequency error remaining; with a low-frequency open loop gain of only ~66dB, we pick up an extra ~0.05% of error. It's not quite symmetrical when we look at the –ve excursion, due to a little bit of input bias current and offset voltage in the model. Be cautious about offset assumptions in models!

The settling behaviour is always better with the larger values of decoupling, up to 470nF. If we're going to be using such large decoupling capacitors, might it be sensible to put an additional small decoupling capacitor in parallel, to take us up to the three capacitor situation discussed in "Know the..."? Figure 6.12 shows the result of adding a 4.7nF ceramic capacitor (with 0.7nH inductance and 80 milliohms of ESR) in parallel with the existing swept decoupling capacitor. The result is a new, undamped resonance at an even higher frequency, as we'd expect. The extra cap appears to add no value at longer timescales, and just additional ringing in the first few hundred ns. This points to a preference for a single, well-chosen decoupling capacitor rather than several in parallel.

The Big Finish

Finally, therefore, we have arrived: ***the accuracy of the op-amp's output is unmistakeably affected by the choices you make on power supply regulator and***

decoupling components. We've seen that if there's no damping in the regulator output capacitance, an op-amp can self-harm through its interaction with the supplies. Even with damping from tantalum capacitors, the value of local supply decoupling affects the step output of fast amplifiers, delivering an extra error term in even a moderately accurate system. As the decoupling capacitor value is increased, the frequency and amplitude of this resonant tail fall off rapidly. For accuracy to be maintained here, we need as large a decoupling capacitor as possible, with 470nF being clearly superior to the commonplace 100nF in these simulations. Remember, all this work has been done for a gain of -1, i.e. a noise gain of 2. The effect of increasing the noise gain will be the same as choosing a lower GBW amplifier. High gain systems are even more likely to be impacted by these supply disturbances.

Takeaways and conclusions for the series

- Signals on the supply pins of an op-amp *will* find their way to the output especially at higher frequencies, but also at LF if the amplifier has poor low frequency open loop gain.
- Self-inflicted supply variations due to varying load current are *always* present and *can never be* completely rejected by the amplifier. They will affect the amplifier output to some accuracy level, and you can investigate this to a surprising level of accuracy in simulation.
- Op-amp vendors put many obstacles in your way to simulating this effectively – press them (and your regulator and passives suppliers) for proper SPICE models that work! Or just use op-amps with working models, by testing out the models with the test fixtures from this series.
- Don't use low ESR regulator capacitors just because the regulator vendor says they are OK for 'stability'. Your system performance is likely to be significantly better with the additional losses that tantalum capacitors can provide in this role.
- A bit of slew-rate limiting in the op-amp isn't necessarily bad; the resonances on the supplies are excited less strongly the slower the amplifier output is slewing. But if there are some fast-slewing, wide-bandwidth amplifiers affecting the supply rails, then low-bandwidth amplifiers, with their consequently poor supply rejection, will be particularly sensitive to the resulting noise. Careful in a mixed system!
- Use the *highest value* local ceramic decouplers you can get, in the largest package that will fit, and *don't* use Y5V dielectric material unless you've fully allowed for the value change with bias voltage. But the volumetric advantage of the high-k dielectric may be completely lost, with this allowance.
- Adding another small ceramic capacitor in parallel with the main decoupler has a small effect on the amplifier output – but it's often a negative one, so test carefully; most times, it may be unnecessary.