Stabilizing Tube Feedback Amplifiers Without Math: Part 2

By Merlin Blencowe. This article first appeared in AudioXpress, June 2026.

In Part 1 of this article I covered some fundamental feedback theory and how it pertains to a conventional tube amplifier, where feedback is taken from the loudspeaker and is applied to an earlier stage. It will be essential to read Part 1 first, to make sense of what follows. In this part I will cover the most common methods used to stabilise such an amplifier. As before, we will mainly be looking at *loop* gain and phase, but now we interested in how to force the loop gain to drop to less than unity (0dB) before the phase shift reaches ±180°, which is the condition for stability.

Dominant Pole

Fig. 1 shows the loop-gain of a fictitious amplifier, and we see that a very ambitious 40dB of feedback is being attempted. The initial response in red has a pole at 10kHz (P_1) followed by a pair of resonant poles at 40kHz which might be caused by a transformer. The gain-crossover frequency (GCF) is the point where the loop gain drops to 0dB, which is just over 100kHz in this case. Since the GCF lies on a third-order section of slope we can be certain it will oscillate if we close the loop. Indeed, the phase has reversed at just 30kHz where the loop still has oodles of gain. How can we stabilise this amplifier?

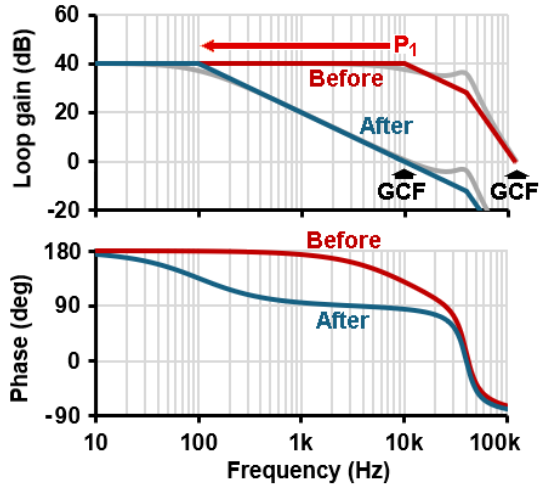

Fig. 1: Dominant pole compensation can stabilise almost any system, at a heavy cost of bandwidth.

Assuming we cannot change the transformer resonance, the simplest way to guarantee stability is to use dominant-pole compensation. This is done by arranging one pole to be much lower than all others, so it dominates the overall amplifier response, making it appear to be first-order over a very wide range. In this case we might try to slug P_1 by adding some capacitance from a preamp stage to ground. By adding enough extra capacitance to push P_1 down to 100Hz, say, we get the blue line shown in the figure. The resonant peak has dropped below 0dB and the response looks first-order all the way up to the new GCF of 10kHz. The phase margin is now excellent, being almost 90°.

But we have paid a heavy price; we had to throw away a *lot* of bandwidth to get this result. Also notice that at the point where the phase does hit the danger figure of 0° the resonant peak in the loop gain is about −3dB, which is the gain margin. This demonstrates that an annoying transformer resonant peak can leave us with a very slim gain margin even if the phase margin is large. We could slug P_1 more, but always at the cost of bandwidth. Moreover, the extra capacitance might place a heavy load on the stage driving it, causing that stage to distort or even hit slew limiting. For these reasons dominant-pole compensation is rarely used in a tube amplifiers; it is more commonly found in transistor amplifiers. Nevertheless, it is an essential concept to keep in your back pocket.

NDFC

An effective stabilising component used in some tube amplifiers is a small capacitance connected from the anode of one power tube back to a convenient point in the preamp. This point might be the input of the phase inverter as in the Rogers Cadet III, or even the same point that global feedback is applied, as in many Dynaco designs; both possibilities are illustrated in fig. 2. This capacitor creates a nested feedback loop for high frequencies. The correct power tube must be selected so this little feedback loop is negative rather than positive, of course.

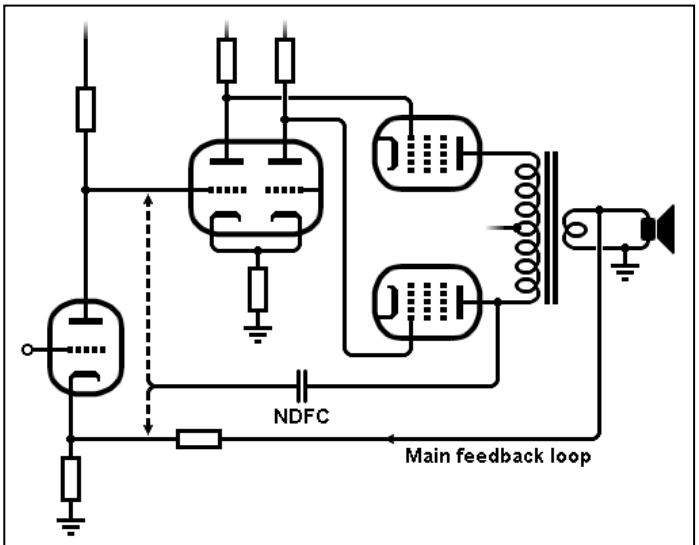

Fig. 2: A nested damping feedback capacitor is an effective stabilising component but needs a very high voltage rating.

This capacitor reduces the effective output impedance of the power tube at high frequencies, which helps to damp transformer resonance, splitting the transformer poles apart to produce a more gentle roll off and less rapid phase change. It also provides some nested distortion reduction as a bonus (unlike a step network which simply throws gain away). It really deserves a name so I have chosen to call it a Nested Damping Feedback Capacitor

(NDFC). A suitable value of capacitance can be found by trial and error on the bench, and a systematic approach was beautifully described by Kauder.[1]

A disadvantage of the NDFC is that it must have a *very* high voltage rating to withstand the anode voltage swing of the power tube, which can exceed twice the supply voltage during clipping. This typically calls for a >1kV-rated component. In an ultra-linear amplifier the NDFC can be taken from the screen grid rather than from the anode, which reduces the voltage swing across it, as seen in the Dynaco ST-70 for example.

Step Network

The most common method of stabilisation found in tube amplifiers is an RC step network, also called a 'lead-lag' network. Fig. 3 shows the basic form of the circuit. At low frequencies where the capacitor appears to be an open circuit, signal passes through without attenuation, but at high frequencies the two resistors form a voltage divider giving attenuation. In practice the series resistor R_1 is not a physical component but is formed by the output resistance of an existing tube stage. The step network is usually applied across the earliest stage of the amplifier since signal levels are smallest there, so they should be least affected by the additional loading.

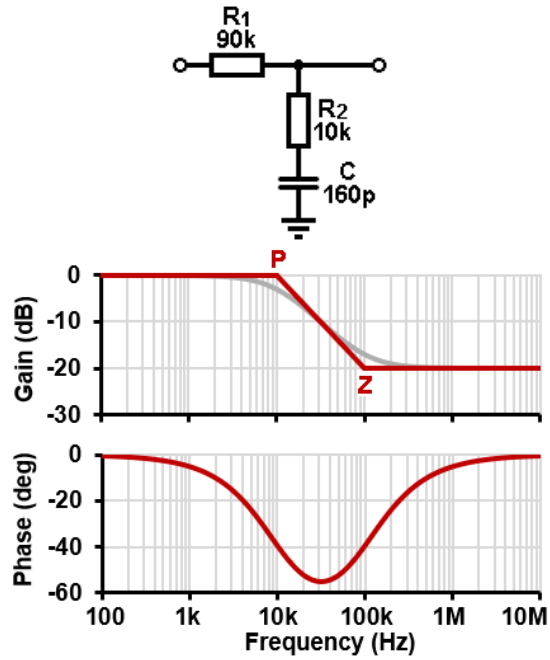

Fig. 3: A step network creates a pole-zero pair. Phase shift is maximum at the centre of the step but returns to near-zero on the flat sections.

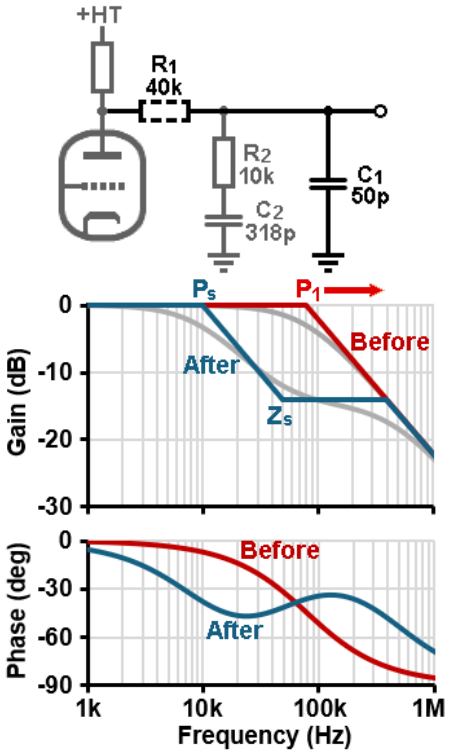

Fig. 4: Model representing a stage with 40kΩ output resistance driving 50pF of input capacitance. Amplitude and phase response before and after adding the step network R_2 and C_2.

The step network forms a 'pole-zero pair'. The zero is given by:

$$f_P = \frac{1}{2\pi C R_2} \qquad (1.1)$$

And the pole frequency is:

$$f_Z = \frac{1}{2\pi C\left(R_1 + R_2\right)} = \frac{f_Z}{\text{Attenuation}} \qquad (1.2)$$

The gain or rather attenuation of the step is:

$$\text{Attenuation} = \frac{R_2}{R_1 + R_2} \qquad (1.3)$$

Which can of course be expressed in decibels:

$$\text{Attenuation in dB} = 20\log\left(\frac{R_2}{R_1 + R_2}\right) \qquad (1.4)$$

In fig. 3 the divider provides 20dB attenuation which is a factor of ten, so the pole is ten times lower than the zero. Similarly, a 6dB step would result in a pole that is half the frequency of the zero. This proportionality makes it very easy to calculate

[1] Kauder, A. J. (1960). Universal Feedback Amplifier Circuit, *Audio*, January, pp19-21.

component values or to sketch a Bode plot. We don't want the network to create a perceptible step within the audio band so we always want the zero to lie above 20kHz (the pole can lie where it may). We can see from fig. 3 that the phase lag has a bell shape; it reaches its maximum at the centre of the step but returns to 0° on either side (the bigger the step, the bigger the bell). The network can therefore provide 'attenuation without phase shift'[2] as the old texts put it, at high frequencies.

However, we must not fall into the trap of thinking a step network provides attenuation forever, from Z_s to infinity. In reality the network will be added to a gain stage that *already has another pole*, created between its output resistance and the input capacitance of the next stage. A more realistic situation is shown in fig. 4. Here we have a gain stage with 40kΩ output resistance represented by the dashed resistor R_1, which is loaded with 50pF of capacitance from the next stage, creating a pole at 80kHz (P_1). Suppose we then add a 14dB step network and select C_2 to give a zero at 50kHz (Z_s). P_s is automatically five times lower, since 14dB is a factor of five. Looking at the gain we can see the flat section of the shelf does not have far to go before it hits the limit of the response we started with. Adding the step network has pushed P_1 five times higher to 400kHz, but the range over which we have useful gain reduction is hardly infinite.

Looking at the phase response we see that in the region of the step itself we have more phase lag than we started with, but this situation reverses exactly half-way (on a log scale) between Z_s and the original P_1; above this point we get a phase improvement, which is the main advantage of a step network in practice. An equivalent step network exists for low frequency stabilisation but it is rarely used in practice since it is simpler to increase coupling capacitances instead, but see GEC[3] for an example.

Speed-up Capacitor

The final way to tackle stability is to apply more feedback at high frequencies, which is perhaps counter-intuitive. By placing a capacitor in parallel with the feedback resistor, as in fig. 5, we create a step network in the feedback divider itself. This allows higher frequencies which were lagging to regain some phase shift or 'speed up', so it is called a speed-up capacitor.[4] The loop-gain curve –which might be diving down faster than we would like– can then be made to transition to a shallower slope, right at the GCF, and phase shift around the whole loop will be pulled back from 180°.

The speed-up capacitor is normally the last component to be added, after the amplifier has already been made conditionally stable by other means and the loop has been closed. The speed-up capacitor is the cherry-on-top; it eliminates residual gain peaking and ringing on square waves. If an amplifier uses only a little feedback, such as most guitar amplifiers, then a speed-up cap may be all that is needed for excellent stability.

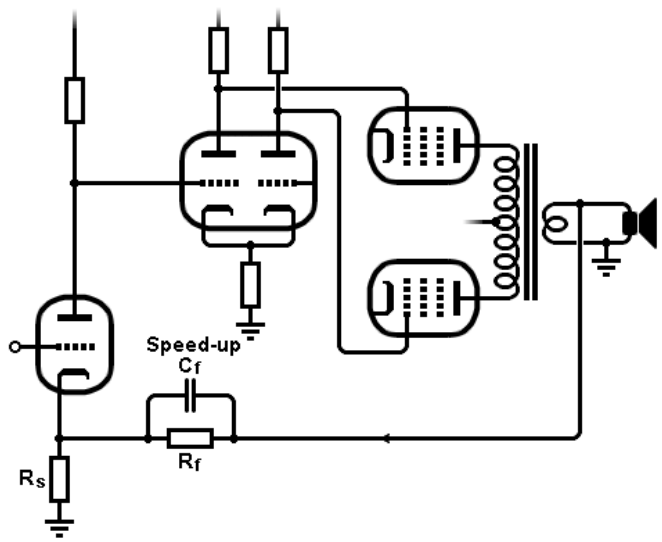

Fig. 5: A Speed-up capacitor is normally added last, to clean up ringing on square waves.

As a starting point try a speed-up capacitance that puts the zero of this step network (f_Z) at the gain-crossover frequency, i.e. the frequency where the closed-loop response is found to be peaking:

$$f_Z = \frac{1}{2\pi C_f R_f} \qquad (1.5)$$

Then feed a 1kHz square wave into the amplifier and tweak the capacitance until visible ringing is entirely suppressed. A clean square wave is a good indicator that phase margin is at least 60°.

A Real-World Example

In is time to put these techniques to the test. As a demonstration I built an amplifier using parts from a scrapped Leak TL/25+; the schematic is shown in fig. 6. I do not intend to give a detailed explanation of the circuit as it is really just a blend between the Leak TL/25+ and the Mullard 5-20, which are well known and documented on the internet. The main modernisations are a solid-state rectifier which allows a much larger reservoir capacitance to be used, so that the extra stage of power-supply filtering used in the original circuits is not required. A series resistor was included to simulate the internal resistance of a GZ34 rectifier, however, to maintain roughly the same HT voltage as in the originals. Protection diodes were added to the output transformer and long-tailed pair. The purpose of the 220kΩ resistor is to keep the voltage across the power supply smoothing capacitors below 450V at start-up.

---

[2] Roddam, T. (1951). Stabilizing Feedback Amplifiers, *Wireless World*, March, pp112-5.
[3] GEC, (1957). *An Approach to Audio Frequency Amplifier Design*, Chapman & Hall Ltd., pp119-121.
[4] Hammond, P. H. (1958). *Feedback Theory and its Applications*, English Universities Press Ltd.

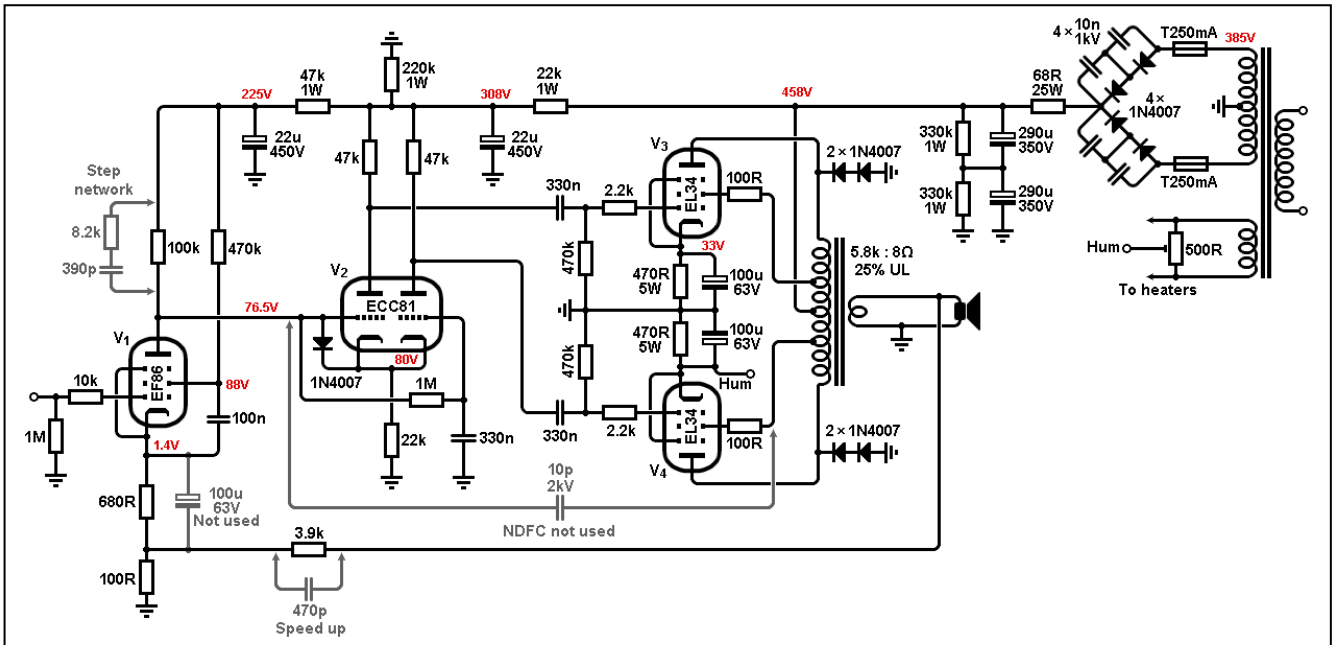

Fig. 6: Test amplifier circuit adapted from a Leak TL25+.

With any design the first job is to get the circuit working open loop; there is no point fiddling with feedback until the core amplifier functions normally. My circuit originally had a cathode bypass capacitor on V_1, but upon first testing the open-loop gain was ridiculously high; just 14mV was needed to drive it to full output power. This was appropriate in the 1950s when signal sources were weak, so that 20 to 30dB of feedback would result in a final sensitivity of a few hundred millivolts, but today we probably want something closer to a volt or so. This would require nearly 40dB of feedback which is simply not feasible around an output transformer without an eyewatering sacrifice in open-loop bandwidth. The cathode bypass capacitor was therefore removed which reduced the input sensitivity to 27mV. With about 30dB feedback (a factor of 31) this should yield an input sensitivity of about 0.8V.

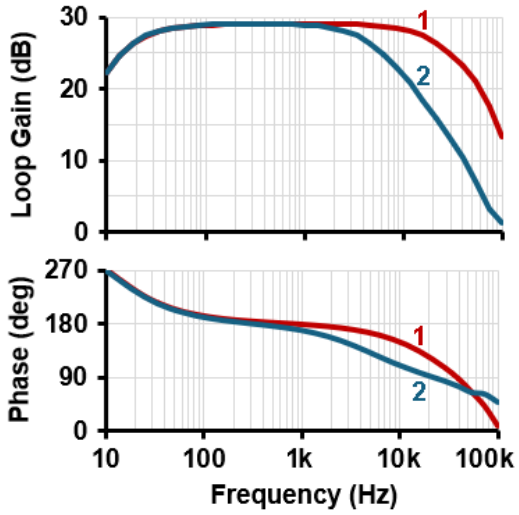

Fig. 7: Initial loop response is unstable (1). Adding a 10pF NDFC makes the loop stable (2).

The open-loop gain between input and 8Ω dummy load was found to be 60.5dB at 1kHz, and we want feedback to reduce this by 30dB. The next job is therefore to set the feedback resistor, either by calculation or by cut-and-try while measuring the loop gain using the method shown in Part 1. In this case a feedback resistor of 3.9kΩ gave a loop gain of 29dB; close enough. We are now in a position to plot the loop gain and phase, and see where we stand. I used an Audio Precision S1 but a signal generator and dual-channel oscilloscope will work too, given enough patience.

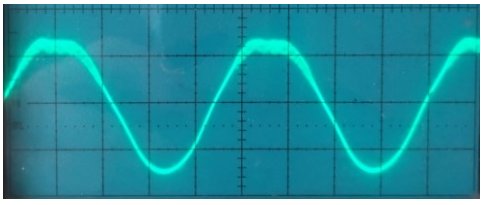

Fig. 8: Parasitic oscillation on a 1kHz output signal at the point of clipping, when using an NDFC.

Stabilizing the High End

Fig. 7 shows the initial loop response in red. The phase drops to 0° at about 100kHz while the gain is still over 13dB, so it will certainly oscillate if we close the loop! Let us try to stabilise the design by adding an NDFC from V_4 back to the input of the phase inverter. A value of 10pF produced the blue curves. This appears to be an excellent result! The gain has dropped to 0dB just a fraction beyond 100kHz while the phase is clearly above 45° –more than the bare-minimum target of 30°– although the gain margin is unclear unless more measurements are made at even higher frequencies. We also do not know what the low-frequency stability margin is, since my instruments don't go below 10Hz. I therefore closed the loop just to see what would happen. Happily, it did not motorboat, but LF stability did need to be improved; this is described separately, later.

I was about to congratulate myself until I discovered a new problem: when the output stage was driven to the point of clipping V_4 would break into parasitic oscillation, visible as a 'furry' waveform in fig. 8. Parasitic oscillation resists systematic bench testing since it is caused by unknown and 'invisible' parasitic elements, and we can hardly open an invisible feedback loop to measure its stability margin. We must resort to cut-and-try fixes. I tried increasing the power tube grid

stoppers, then the screen-grid stoppers, then added anode stoppers, then various inductors, then various Zobel networks across the output transformer primary, then a larger NDFC. None had much effect on this oscillation. Eventually I lost patience and conceded defeat; this particular output transformer apparently won't permit using an NDFC.

Let us abandon the NDFC and try adding a step network across V_1 instead (it does not matter whether it is connected from plate to ground, or across the plate resistor, the result is the same). V_1 has an output resistance of 100kΩ which, in combination with an 18kΩ resistor, will create a 16dB step. Choosing a 100pF capacitor places the zero at 88kHz and its pole therefore 6.5 times lower, around 13kHz. The result is shown in in fig. 9 in blue. Sadly the improvement is not enough. Estimating a little beyond the limits of the graph, the phase is obviously going to hit 0° while the gain is still above 0dB. We need to be more aggressive. An 8.2kΩ resistor will yield a 22dB step, and a 390pF capacitor puts the zero at 50kHz (in truth I tried several other combinations before settling on these values). This yields a good result; the GCF is now a fraction beyond 100kHz and the phase margin is about 30°. Once again I closed the loop to see what would happen.

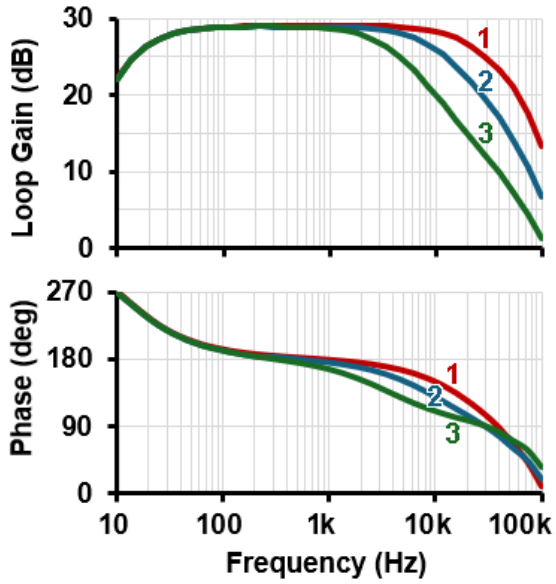

Fig. 9:
1) Initial loop response is unstable.
2) With 18kΩ+100pF step network it is still unstable.
3) With 8.2kΩ+390pF step network the loop is stable.

Fig. 10 shows the forward-path response of the amplifier, starting with its initial open-loop response in red, the effect of adding the step network in green, and the closed-loop response achieved so far in blue. Although it is now conditionally stable there is an obvious peak at the GCF (I took a few extra measurements beyond 100kHz with a different signal generator to show the extent of peaking). This peaking causes ringing on a 1kHz square wave, shown in fig. 11. The final step is therefore to tune out this ringing by adding a speed-up capacitor across the feedback resistor.

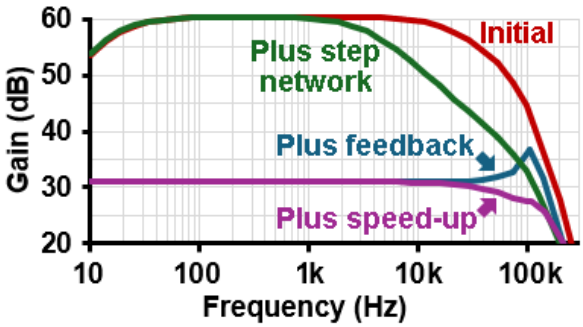

Fig. 10: Forward path response of the amplifier from initial open-loop, to final closed loop (ref. 1W at 1kHz).

Since the peaking occurs close to 100kHz and the feedback resistor is 3.9kΩ, rearranging equation (1.5) suggests a value of 408pF. With a little further tweaking a value of 470pF gave the best visual result, as shown in fig. 11. The final closed-loop response is shown in purple in fig. 10.

The absence of ringing on a square wave implies the gain margin has been improved from the initial figure of 30° to at least 60°. I later confirmed the stability margin by re-opening the loop and measuring the loop gain and phase again, with speed-up capacitor in place. The actual phase margin was about 58°; the gain margin was difficult to determine but exceeded −20dB. Further checks showed that the amplifier remained stable with both open and shorted input, with a real speaker load, and also with a 8Ω||100nF load. Incidentally, distortion at 10 watts was 0.018% at 1kHz but 0.25% at 10kHz. Clipping occurred at exactly 40 watts. Output impedance at 1kHz was 12.5Ω open loop, and 0.17Ω closed loop.

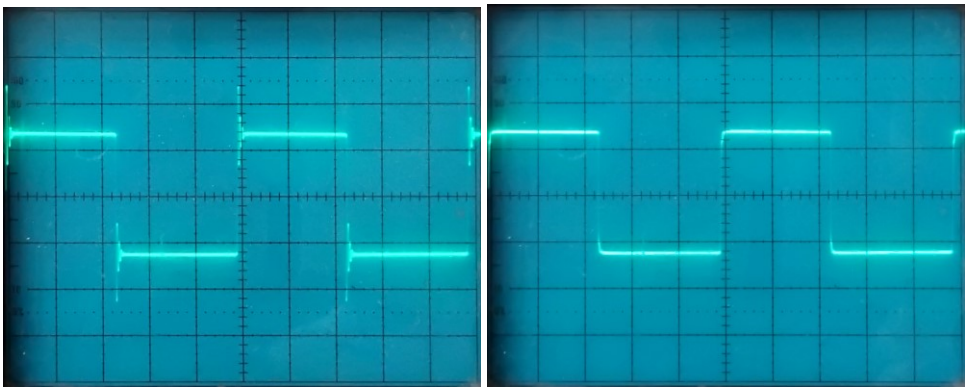

Fig. 11: **Left:** Ringing on a 1kHz square wave after closing the feedback loop.
Right: A 470pF speed-up capacitor eliminates ringing.

Stabilizing the Low End

It was mentioned earlier that upon closing the loop the circuit did not oscillate at low frequency, but this was pure luck since the LF stability margin was not revealed by measurement. If it *had* broken into motorboating I would have needed to increase

coupling capacitances until the loop was conditionally stable. Despite getting lucky, when switching ranges on my signal generator the output signal would visibly 'bounce' for a couple of seconds, indicating the LF stability margin was extremely poor.

LF peaking can be tested by applying a single pulse at the input; about 100ms long will do. With an analog oscilloscope you'll probably want to use a pulse generator, but if you have a storage 'scope that can trigger off a single pulse then you can use something as simple as a DC voltage supply and a momentary push button, which is exactly what I used to produce fig. 12. If there is peaking in the amplifier response at very low frequencies, unsustained oscillation will be apparent after the pulse. Fig. 12 shows the results I got; in each image the top trace is the input pulse viewed at the grid of V_1 (visibly filtered by input capacitance) and the lower trace is the output viewed across an 8Ω dummy load. Note that we are not expecting the amplifier to accurately reproduce the pulse itself, since this is not a DC amplifier; we are only looking for oscillation after the pulse.

Initially I had installed 100nF coupling capacitors to the power tubes, and on the 'spare' grid of the phase inverter. With these values there was atrocious oscillation after a pulse; no wonder switching ranges produced bouncing! Increasing the coupling capacitors to 330nF gave a significant improvement as shown in the figure. Increasing the grid decoupling capacitor on V_2, also to 330nF, gave a further improvement. Increasing the screen-bypass capacitor on V_1 had minimal effect so it was left at 100nF. Exactly what the LF stability margin *is*, is not easy to determine. A total absence of oscillation would indicate a phase margin much better than 60°, but achieving this condition in a high-feedback amplifier may well demand unfeasibly large capacitances, which can result in excessive recovery time if the amplifier ever clips.

The primary inductance out an output transformer varies with excitation level, especially in a push-pull transformer. The associated pole –which is often dominant– will typically be pushed to lower frequencies as amplifier output power increases, eroding the stability margin (until the transformer approaches saturation, when inductance will fall again). For example, fig. 13 shows the difference in response between a 0.2V input pulse and a 0.8V pulse, on the same scale, revealing that LF stability is indeed slightly worse at the higher power level. Happy hours could be spent experimenting with different capacitor combinations but I think the point is made, and my patience with this amplifier is exhausted. I should point out that this is not an amplifier I would recommend anybody copy; it was built only to demonstrate the steps we can take to stabilise a difficult beast. If I were to start again I would design it to have less open-loop gain and use a more modest amount of feedback, e.g. 20dB. This would make it easier to achieve a large stability margin with less sacrifice in open-loop bandwidth, and less worry about what the stability margin might be at low frequencies.

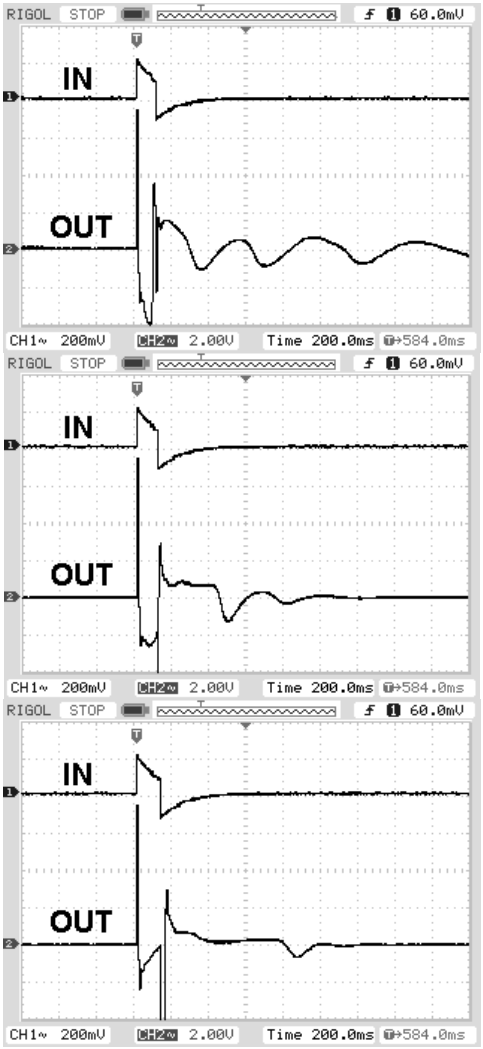

Fig. 12:
Upper: With 100nF coupling capacitors a pulse exhibits atrocious unsustained oscillation.
Middle: With 330nF coupling capacitors the oscillation is improved.
Lower: With 330nF also on the phase inverter second grid, oscillation is suppressed further.

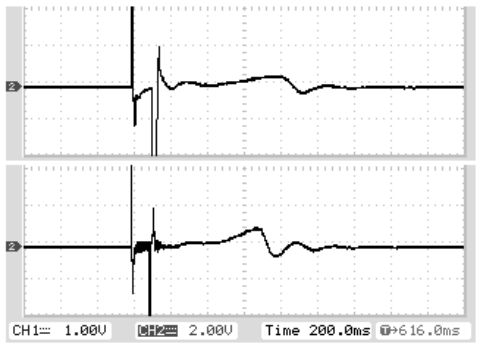

Fig. 13:
Upper: Response to a 0.2V input pulse.
Lower: Response to a 0.8V input pulse.

Closing Remarks

I hope this sheds some light on the mechanisms behind amplifier stability and how you can apply feedback safely, without a lot of mathematics, even when wrapping a relatively large amount around an output transformer. Feedback is an enormous topic in electronics –just look in any good textbook– but almost everything of real-world interest (loop gain, stability margin, bandwidth, signs of instability etc.) can be observed directly on the bench with a signal generator and oscilloscope, though an audio analyser does make the job quicker! More importantly, real-world measurements engender an intuitive understanding of how a feedback design will behave, or misbehave, even before the loop is closed. With that understanding in place, feedback can be applied deliberately and confidently, rather than with fingers crossed.