Stabilizing Tube Feedback Amplifiers Without Math: Part 1

Well, almost no math! By Merlin Blencowe

When global negative feedback is wrapped around a tube amplifier it promises measurable improvements; less distortion, wider bandwidth, lower output impedance etc. When applied thoughtfully it makes a good amplifier better. But these improvements come at the cost of stability; the more feedback we apply, the greater the chance our amplifier will become an oscillator. Judging by the many internet forums I frequent it seems many DIYers apply feedback with no more precautions than crossed fingers. Feedback theory is well covered in tube textbooks and articles[1,2,3] but admittedly it can be difficult to interpret and to apply to new situations. The notation used in older books can be opaque, the techniques were optimised for slide rules, and they focus a lot on Nyquist diagrams which have never been very intuitive. Frustratingly, they often fall short of giving examples that include a transformer inside the loop.

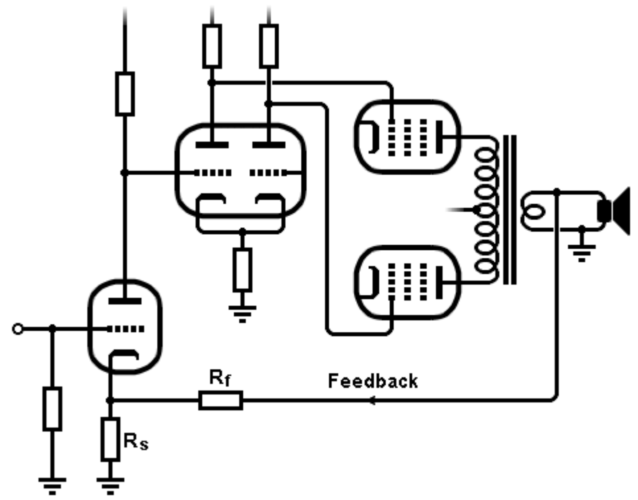

Fig. 1: Typical non-inverting feedback architecture.

In this article I hope to cover the essentials of feedback design for tube power amplifiers in a digestible and practical way using amplitude and phase plots. Bode plots will also be used when it is helpful, i.e. approximating amplitude curves using straight lines of specific slope (phase curves do not conform so nicely to straight line approximations so I will show only true phase responses). I will also show how the stability of a tube amplifier can be assessed by practical bench methods, and how stability can be ensured no matter how much feedback you might want to use –and what is sacrificed in the process. Nevertheless, I expect readers to be familiar with basic principles of amplifier design already, or this article would never end.

There are many ways to implement feedback but in this article I will focus on the familiar situation where it is taken from the output transformer secondary, i.e. across the loudspeaker, and is applied to an earlier stage by a simple resistor divider, something like fig. 1. This is a particularly troublesome arrangement because the output transformer is inside the loop, and transformers add phase shift in a very inconvenient way at high frequencies. Fortunately, one advantage of a tube amplifier is that it can be tested open-loop, before adding any feedback, so we can get reliable information about its gain and phase on the bench. By contrast, most transistor amplifiers have so much gain they will simply oscillate open loop, owing to parasitic feedback.

The Universal Feedback Equation

Readers ought to be familiar with the universal feedback equation, and for the purposes of this article I choose to write it:

$$A_c = \frac{A_o}{1 + A_o B} \tag{1.1}$$

A_o is the open-loop gain; the gain before any feedback is applied;
A_c is the closed-loop gain; the gain after feedback has been applied;
B is feedback fraction; the fraction of the output signal that we feed back.

The denominator $1+A_oB$ is called the *feedback factor* since it is the factor by which the overall gain of the system is reduced; it is the amount of feedback being used. It is normal to express the feedback factor in decibels and to talk about an amplifier as using so-many dBs of feedback.

The term A_oB is called the *loop gain*, since it is the total gain experienced by the feedback signal as it passes around the loop. It is the gain we measure by breaking the feedback loop, injecting a signal into the appropriate side of the break, and measuring the resulting signal appearing on the other side of the break. The loop gain is equal to the amount of feedback being applied, minus one. For example, if we apply 20dB of total feedback, i.e. $1+A_oB = 10$, then the loop gain is 9, or 19.1dB, and we would adjust the feedback divider to produce this much loop gain. It doesn't matter if the open-loop gain of the amplifier is 100 or 100000, if we apply 20dB of feedback the loop gain is always 9, or 19.1dB (in practice we might ignore the '−1' difference, since it is small, and simply set the loop gain to 20dB or so).

[1] Cooper, G. F. (1950). Audio Feedback Design, *Radio Electronics*, October, pp49-50.
[2] Crowhurst, N. (1956). Stabilizing Feedback Amplifiers, *Radio Electronics*, December, pp34-7.
[3] Keroes, H. I. (1959). Stabilizing Feedback Amplifiers, *Radio Electronics*, January, pp45-7.

Stability problems arise at the extreme ends of the amplifier bandwidth –especially the high end– because there are many hidden filters in the circuit, all of which introduce phase shift. At sufficiently high (or very low) frequencies the total phase shift is sure to exceed 180°, at which point negative feedback becomes positive feedback. If the loop gain also happens to equal to unity when this happens, then $A_oB = -1$ and the universal feedback equation blows up to infinity: the circuit will oscillate. Even if outright oscillation does not occur, marginal stability often reveals itself as a resonant peak in the frequency response, and poor transient behaviour, observable as ringing on square-waves. For the best audio we obviously want a flat, well-behaved frequency response, even outside the audible range.

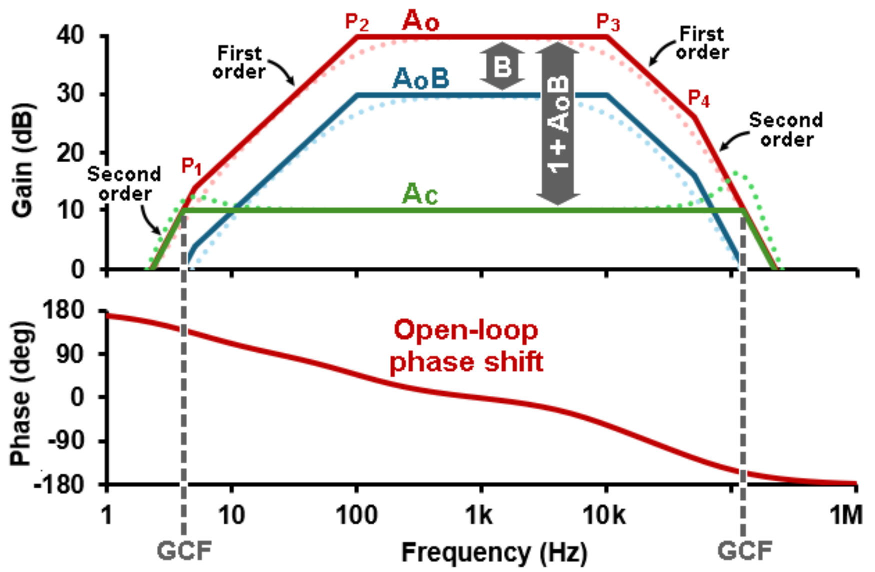

Fig. 2: The universal feedback equation in the form of amplitude and phase plots.

A better grasp of the feedback equation can be acquired by looking at gain and phase plots; fig. 2 shows the response of a fictitious amplifier (Bode approximation in bold, actual response dotted). Starting with the red line (A_o) we see the amplifier has a mid-band gain of 40dB. It also happens to have four cut-off frequencies or 'poles', P_1 to P_4, created by various unavoidable time constants within the imaginary circuit. Each pole contributes 45° phase shift and tends to 90° at infinitely low or high frequencies, leading at the low end, lagging at the high end. Each pole also adds 20dB/decade attenuation. In this case the response therefore ultimately falls at 40dB/decade (second order) at the frequency extremes, and phase shift approaches ±180°.

If 30dB of feedback is applied then the gain of this amplifier is reduced to 10dB and its bandwidth becomes much wider, shown by the green line (A_c). The vertical difference between red and green is therefore the feedback factor $1+A_oB$. It is now obvious that we really only have 30dB of feedback in the mid band; it is less for frequencies outside of P_2 and P_3.

The point where A_c intersects with A_o is called the *gain crossover frequency* (GCF) and defines the closed-loop bandwidth. This amplifier has two GCFs, one for the high end and one for the low end, whereas a DC amplifier would have no lower limit. Beyond the GCF there is essentially no effective feedback so A_c eventually returns to the original A_o curve. It is at the GCF(s) that an amplifier might try to oscillate, so this area gets all our attention when it comes to assessing stability.

The blue line shows the loop gain (A_oB). Assuming feedback is applied with a simple resistor divider with no time constants of its own, the A_oB curve is identical to the A_o curve except for being shifted vertically down by the feedback fraction B. Its phase response is also identical (or shifted bodily down by 180° depending on our frame of reference, as we will see later). Notice that the point where A_oB drops to exactly 0dB is also the GCF. We can therefore assess stability by looking at the GCF point(s) on the A_o curve, or on the A_oB curve; both can give us the same information. However, measuring the latter is both easier and more precise in practice. Once any amplifier has been set up so we can measure its loop gain, we just look for the point where it drops to 0dB and we know we are looking at the GCF, by definition.

The green line shows the closed-loop response of the amplifier, that is, after feedback has been applied. Looking at the true curve (dotted) we find the gain is 'peaking' at each end of the bandwidth, exactly at the GCFs. This is happening because the loop phase shift, shown in the lower graph, is getting rather close to ±180° at these two critical frequencies. In this case peaking at the low end is not as bad as at the high end, owing to the positions of the poles and the phase shift each one contributes. If there were more poles then phase shift would grow more extreme and these peaks would grow higher and sharper. Eventually one of them might shoot up to infinity and we arrive at oscillation, which we obviously don't want. There are some useful guidelines to adhere to: looking at the A_oB curve, if the GCF lies on a section that is falling at a first-order rate then the system will always be stable, because in a linear system the phase will never exceed 180° on a first-order section of slope. If the GCF lies on a second-order section of slope the system may or may not be stable, depending on how close higher-order poles happen to be. If the GCF lies on a third-order (or more) section of slope the system will certainly oscillate. In this example both GCFs lie on second-order sections of slope but the phase shift has not yet reached 180°, so the amplifier will not oscillate. The peaks in the closed-loop amplitude response will, however, produce ringing on square waves.

Measuring Loop Gain and Phase

As mentioned earlier, the best way to assess the stability of a real circuit is to measure its loop gain and phase. Valve feedback amplifiers are invariably built in a non-inverting configuration with feedback returned to the cathode of the input

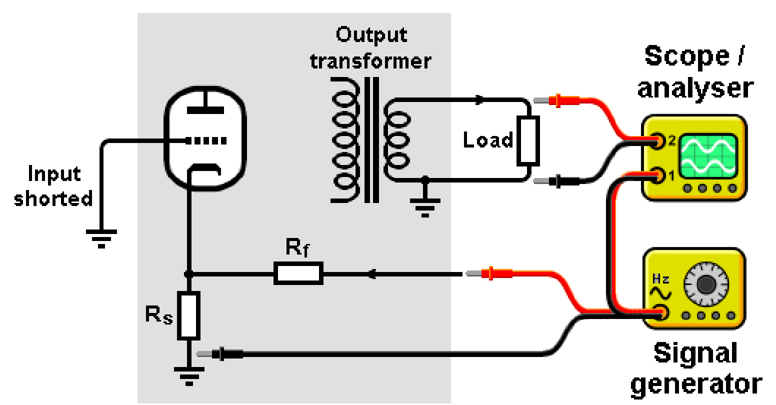

Fig. 3: Measuring loop gain by breaking the feedback path and injecting a signal into the feedback divider.

valve or sometimes to the 'spare' grid of a long-tailed pair phase inverter. In the non-inverting configuration the feedback signal is subtracted from the input signal, so the output (speaker) signal must be *in phase* with the input for feedback to be negative. But when we come to measure *loop gain* we will find it is always *inverting* for feedback to be negative. In other words, loop phase shift is ideally 180° in the mid band and becomes problematic as it approaches 0° or 360°.

To measure loop gain it is important that we break the loop at a point that does not significantly alter the way the amplifier inside the loop functions, or we will get misleading answers. The most practical way to do this is to make the break where the output transformer connects to the feedback divider, illustrated in fig. 3. The test signal is first attenuated by the divider that sets the feedback fraction B, then it passes through the rest of the amplifier with gain A_o, giving the total loop gain A_oB (the main amplifier input should be shorted to ground during measurement).

In a linear system the amplitude response is intimately related to the phase response. Given one, it is theoretically possible to derive the other mathematically and it provides no new information. In practice this is easier said than done; we must measure both the gain and phase. A modern audio analyser should be able to plot gain and phase automatically. If you don't have one then you can use a dual-channel oscilloscope to monitor simultaneously the test signal and the signal retuning from the amplifier; most modern digital 'scopes provide a function to read out the phase difference. An analog 'scope is unlikely to have such a feature, in which case it is possible to measure phase the old-fashioned way using Lissajous figures. However, this is so laborious and prone to mistakes that, in this century, it is honestly worth investing in alternative equipment.

Stability Margin

Fig. 4 shows the loop gain and phase we might measure in another fictitious amplifier, concentrating on the high-frequency end. In this case the loop has a maximum gain of 30dB because that is the amount of feedback being applied (30.3dB if you want to be pedantic), at 100Hz at least. The circuit has, for whatever reason, three poles labelled P_1 to P_3. The total phase lag therefore tends towards 3×90° = 270° which is obviously more than the dangerous figure of 180°. However, loop gain is inverting, so we are actually looking for the point where the phase shift drops *from* 180° *to* 0°.

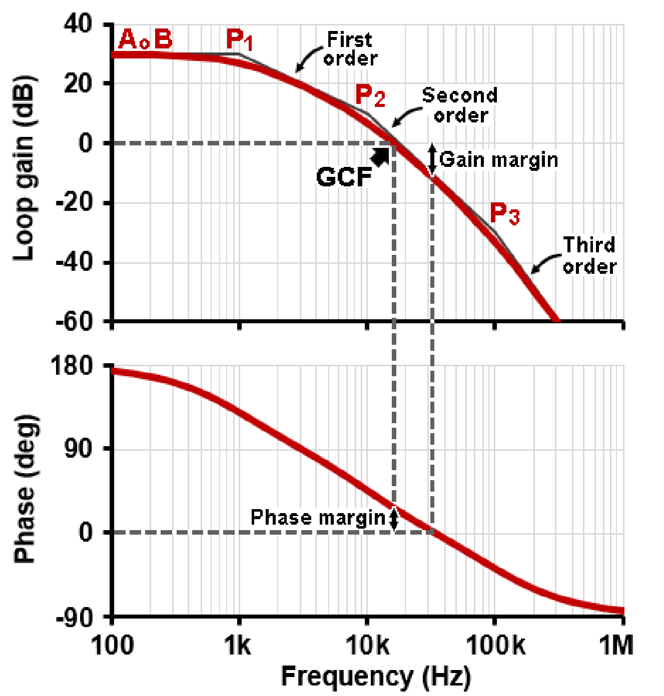

Fig. 4: Loop gain and phase of a fictitious amplifier, illustrating the gain and phase margin.

Will this amplifier oscillate if we close the loop? The GCF is the point where the loop gain falls to 0dB. In this contrived example we can see that it lies on a second-order section of slope, so immediately we know we're somewhere between 'cannot oscillate' and 'will oscillate'. How helpful. Nevertheless, following the dashed line we see that at the GCF the phase shift has only dropped to 30°. The difference between this and 0° is of course 30° and is called the *phase margin*. The phase margin is how much more phase shift would be needed at the GCF to reach the dangerous figure of 0°.

Now look at the point where the phase shift really does hit 0°; here the loop gain is –11dB. The difference between this and 0dB is of course −11dB and is called the *gain margin*. This is how much more amplifier gain would be needed to raise the loop gain to 0dB where the phase shift is 0°. Together, the phase margin and gain margin can be called the stability margin of the amplifier. Exactly the same logic can be applied to the low end of the amplifier's bandwidth, not shown here, except the phase would be leading rather than lagging, i.e. we would be looking for the point where the phase hits 360° rather than 0°. In any case, phase is relative so it doesn't really matter how you choose to plot it as long as you keep track of what you are looking for. In this example the phase shift has not reached 0dB at the GCF, so this amplifier is stable at high frequencies.

Fig. 5 shows the response before and after the loop is closed. The open-loop gain A_o happens to be 40dB, so with 30dB of feedback the closed-loop gain A_c has been reduced to 10dB, but we have some peaking at the GCF. Peaking is an indicator of stability margin; the smaller the margin the worse the peaking. The math of high-order systems is complex but, as a rule of thumb, peaking normally disappears when the phase margin is better than about 60°.

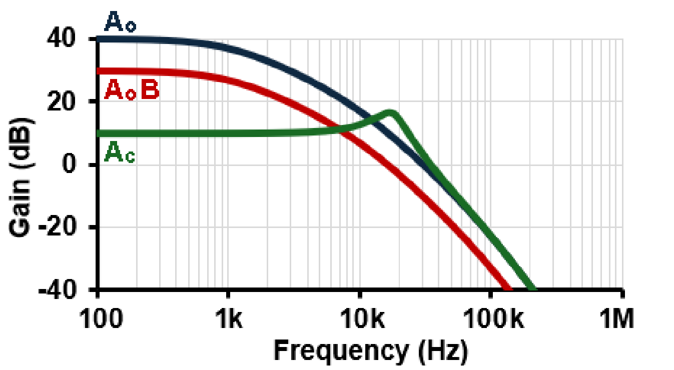

Fig. 5: Open and closed loop response of the fictitious amplifier. Note peaking at the GCF, which will cause ringing on square waves.

Stabilising an amplifier involves altering the loop gain and phase curves until a comfortable stability margin is achieved. How much margin is needed? There is no absolute answer but a phase margin of at least 30° and a gain margin better than −6dB is widely accepted as a good target to aim for, just before the loop is closed.[4,5] Once the loop is closed we may get peaking/ringing but this can finally be tuned-out with a speed-up capacitor as explained in Part 2. The stability margin of the completed design will then be even greater than the figures obtained when the loop was first closed.

Another point to note is that if we apply 100% feedback, i.e. the amplifier's output is connected directly to its input, then the A_o and A_oB curves become one and the same. Putting it another way, if the *open-loop* response drops below 0dB before the phase shift reaches ±180° then *any* amount of feedback could be applied and the system would not oscillate. This is called *unconditionally stable* or *unity-gain stable*. Tube amplifiers are rarely built for unity gain though.

The Problem of the Output Transformer

Wrapping feedback around an output transformer presents a particular problem, owing to its leakage inductance and distributed capacitance, mainly in the primary winding. These stray impedances are difficult to measure and they may vary to some extent with operating conditions, so transformer manufacturers are understandably reluctant to publish such figures. However, since we cannot alter these hidden components (unless you are in the business of winding your own transformers) we don't necessarily need to know their values, we only need to understand why they create difficulties when adding feedback.

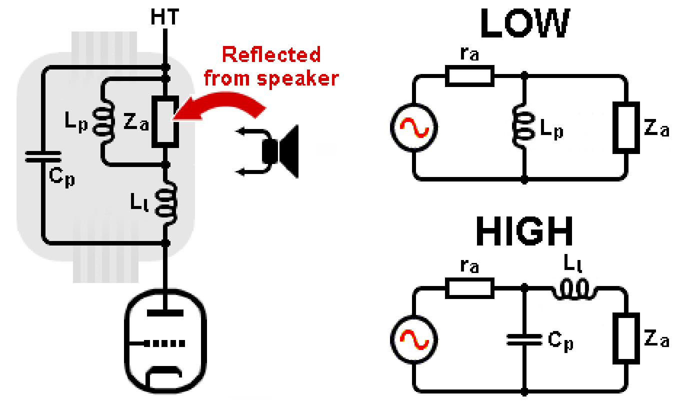

Fig. 6: Simple model of an output transformer, further simplified for low and high frequencies.
r_a = Internal plate resistance of the power tube;
Z_a = Primary reflected load impedance;
L_p = Primary inductance
C_p = Primary winding capacitance;
L_l = Primary leakage inductance.

Fig. 6 shows a simple equivalent circuit of a transformer, further broken down into equivalent circuits for low and high frequencies. Fig. 7 shows the gain and phase response we can expect from a typical tube output stage using such a transformer. At high frequencies the primary leakage inductance L_l, and lumped primary winding capacitance C_p, form a second-order RLC filter. A transformer therefore always introduces *at least two* high-frequency poles (for fuller treatment see Cherry and Hooper[6] or van der Veen[7]). Since the rest of the amplifier also contributes high-frequency poles you can probably anticipate the difficulty this creates if we want to apply feedback around the whole lot.

The larger the load resistance Z_a, or the smaller the internal resistance of the power tube/s r_a, the greater the damping of this RLC circuit. Pentodes have very high internal resistance which often results in under-damping, producing a small resonant peak in the open-loop amplitude response, which then falls off at a second-order rate. This is another way of saying P_1 and P_2 occur at the same frequency (mathematicians would say they are a complex conjugate pair). The phase will snap rapidly to −90° at resonance, and approach −180° equally rapidly.

[4] Hammond, P. H. (1958). *Feedback Theory and its Applications*, English Universities Press Ltd., p109.
[5] Cherry, E. M. & Hooper, D. E. (1968). *Amplifying Devices and Low-Pass Amplifier Design*, John Wiley and Sons Inc., p506.
[6] Ibid, p343.
[7] Van der Veen, M. (2010). *High-End Valve Amplifiers* 2, Wilco Amersfoort.

By contrast, power triodes have much lower internal resistance which helps to damp the circuit and may cause the two transformer poles to 'split', as illustrated in fig. 8 (the same pole-splitting effect will occur if an amplifier is tested with no load). The response therefore starts to fall off at a first-order rate at P_1, then transitions to a second-order rate at P_2. This results in a less rapid phase change, which tends to make triode amplifiers less difficult to stabilise if feedback is applied. The response of an ultra-linear output stage will be somewhere between the pentode and triode extremes.

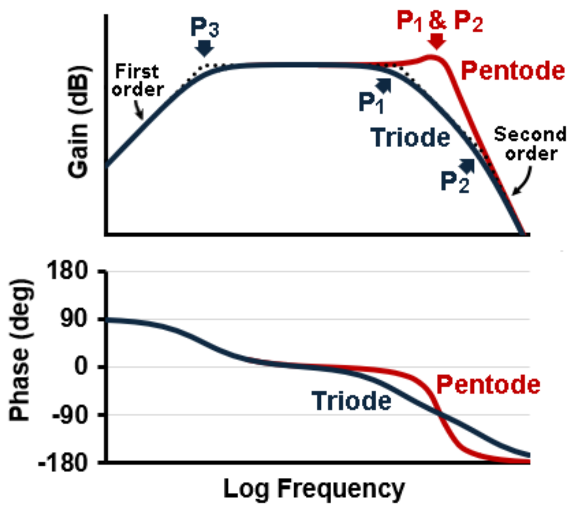

Fig. 7: Typical response of a valve output stage. The transformer introduces two high-frequency poles P_1 and P_2, but only one low-frequency pole P_3.

At low frequencies the situation is much simpler because tiny stray impedances are no longer important, only the primary inductance matters. This creates a simple first-order RL filter with a pole at:

$$f_{P3} = \frac{Z_a \parallel r_a}{2\pi L_p} \qquad (1.2)$$

Phase shift will be +45° at the pole, heading towards +90°. In many designs it is the transformer that creates the first low-frequency pole of the amplifier, often around 20Hz. Many manufacturers publish the primary inductance, at least for single-ended transformers, so this pole can be calculated easily. In a push-pull transformer the inductance will vary more with applied voltage, making it difficult to quote a single useful number. Just another reason to measure true loop gain rather than rely on calculation, then.

The Problem of Miller Capacitance

In textbooks you may find examples where the response of a multi-stage circuit is worked out on paper to assess its stability margin.[8] These examples are sometimes misleading because they often assume the response of each stage considered in isolation can be summed to find the total. In other words, they imply that we can work out the output resistance of each tube stage, and the Miller capacitance of the following stage, and use these figures to find every high-frequency pole in the circuit. This does not work!

For example, fig. 8 shows a three-stage circuit where every stage has a gain of ×29, an output resistance of 10kΩ, and 2pF of anode-grid capacitance. The naïve approach suggests each stage therefore has (29+1)×2pF = 60pF of Miller capacitance. Each stage sees 10kΩ source resistance, so each stage must have a pole at 265kHz. Adding the three together should therefore give an overall third-order response above this frequency, right? Wrong. The actual response is shown in blue; the naïve approach is not even close. The problem is that each local feedback network acts as a reactive load upon its neighbours, changing their behaviour. By sketching in some Bode slopes (dashed) we can probably guess at a couple of poles, but there must be more at very high frequencies since the phase keeps on diving. Nor would it be correct to say any particular stage is responsible for any particular pole, they all interact. Trying to alter one, e.g. by adding more capacitance to C_{ga}, will unavoidably move the others too. However, it is nice to discover that the attenuation and phase change are not as rapid as anticipated, since this works in our favour if we want to wrap feedback around the circuit.

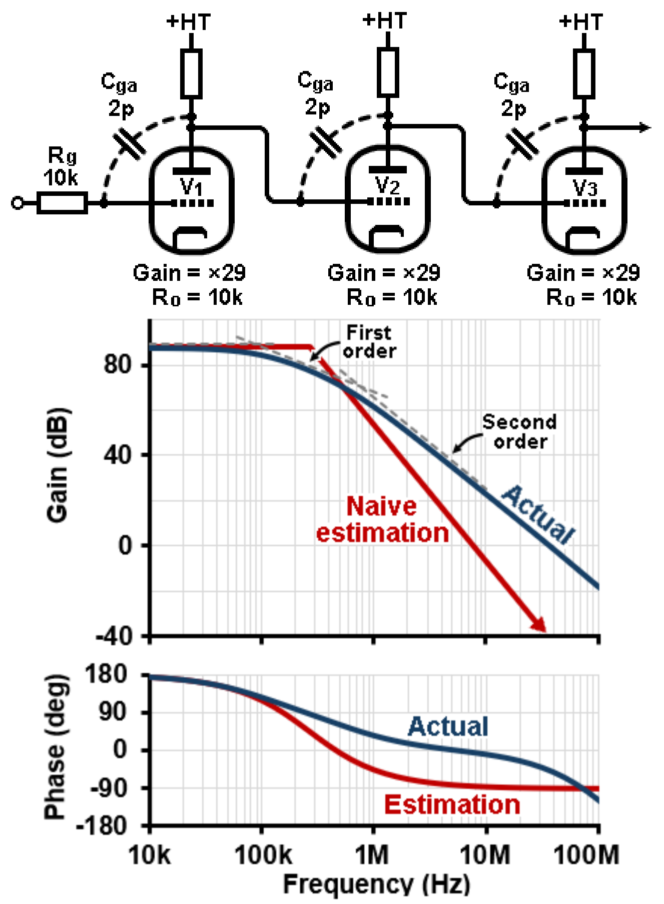

Fig. 8: At high frequencies the response of a multi-stage amplifier does not behave as the simple sum of the responses each stage would have in isolation.

[8] Crowhurst, N. H. & Cooper, G. F. (1957). *High-Fidelity Circuit Design*, Gernsback Library Inc., p35.

I present this example to illustrate how difficult (even futile) it is to predict the amplitude and phase response of an amplifier on paper, at high frequencies. Even Cherry and Hooper, masters of feedback theory, concede that "the complete analysis of a multistage amplifier in which the input impedance of one device forms part of the load for the preceding device is a formidable undertaking."[9] Today we can use computer simulation of course, but when a transformer is included –for which we probably don't have an accurate model– even simulation has its limits. There really is no substitute for measuring loop gain in real life, especially at the high end.

A corollary of this applies when global feedback is returned to the amplifier input stage. Referring to fig. 9, this stage will have an initial pole created between its output resistance and the load capacitance of the following stage, C_1. But if there is a resistance R_g in series with the grid then it will combine with the tube's own Miller capacitance to create a hidden pole and zero, which are pulled to lower frequencies as R_g is made larger. The response of the stage will vary as shown in the figure and, usually, the phase lead introduced by this hidden 'step network' will tend to improve the overall stability margin of the amplifier –more in Part 2.

If there happens to be a volume potentiometer at the input, this pole-zero pair will vary with pot rotation as the source resistance of the pot varies. Consequently, after stabilising the amplifier with the input shorted (e.g. pot turned down to zero) it is important to check that it remains stable at every other pot position too –look for changes in square wave overshoot while rotating the pot. A grid-stopper helps to swamp such changes in source resistance, keeping it more constant regardless of pot position. In other words, an input grid stopper can be a valuable stabilising component even though it looks like it is outside the global feedback loop!

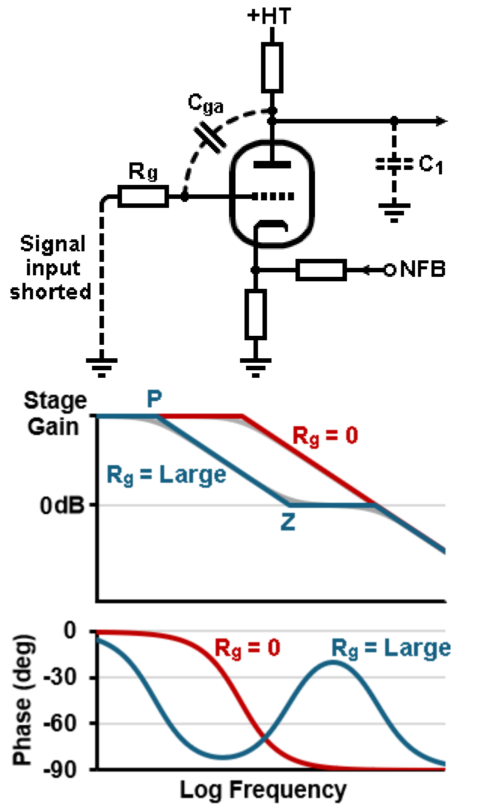

Fig. 9: C_{ga} creates a hidden pole in the stage where feedback is applied, which depends on the resistance in series with the grid.

The Lesser-Problem of Low Frequencies

The low-frequency end of things is much easier to handle because we really can sum the responses of each individual stage to find the total, i.e. low-frequency poles don't interact, they can be moved independently. We are in full control of (nearly) all low-frequency poles and zeros since they are created by coupling capacitors, and to a lesser extent cathode bypass capacitors and power supply decoupling capacitors. The only one outside our command is the pole created by the output transformer, or as Edwin puts it "it is very much easier to increase the size of the coupling capacitances elsewhere in the amplifier than it is to increase the inductance of the transformer".[10]

If you do feel like sketching the low frequency response of an amplifier, the following approach might be used. Fig. 10 shows an example circuit and I assume the reader is familiar with basic amplifier design so should know these formulas already. A cathode bypass capacitor creates a zero at:

$$f_Z = \frac{1}{2\pi C_k R_k} \tag{1.3}$$

And a pole at:

$$f_P = \frac{1}{2\pi C_k r_k} \tag{1.4}$$

Where r_k is the internal cathode impedance of the tube. For the first stage this gives us P_1 = 4.8Hz and Z_1 = 1.6Hz. For the second stage we have P_2 = 23.9Hz and Z_2 = 8Hz.

A coupling capacitor creates a pole at:

$$f_P = \frac{1}{2\pi C_g (R_o + R_g)} \tag{1.5}$$

Which in this case gives P_3 = 6.9Hz. The output transformer creates a pole given by equation (1.2) so we have P_4 = 15.9Hz. The mid-band gain of the circuit is the gain of each stage multiplied together, including the turns ratio of the transformer which we will suppose is 0.04, giving a total gain of 20×50×0.04 = 40, which is 32dB, from input to loudspeaker.

[9]Ibid, p342.
[10] Edwin, G. (1962). Fundamentals of Feedback Design, *Wireless World*, January, pp27-28.

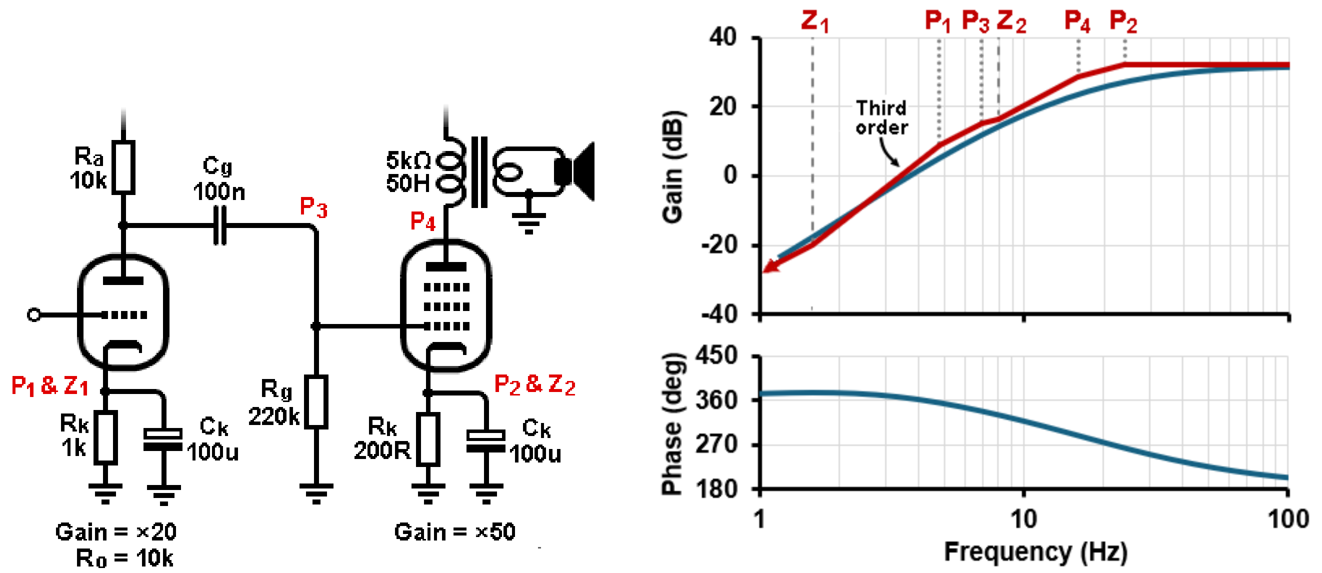

Fig. 10: Estimating low-frequency response is fairly straightforward, and most poles and zeros are fully under our control.

Marking of all the poles and zeros on a graph we can now sketch a Bode plot, starting at 100Hz say, and proceeding from right to left. Any time we encounter a pole the line should bend down by an additional 20dB/decade; any time we encounter a zero the line should bend less steep by 20dB/decade (if a pole and zero lie on top of one another they therefore cancel out). The result is shown in the figure in red, together with the true response in blue. This time the agreement is good. Phase response is not so easy to sketch, so only the true response is shown here. But neither is it really necessary to sketch it, since we can see from the Bode plot that we encounter a third-order slope at P_1, telling us that below this point would certainly be a dangerous place for the GFC to lie. The true phase plot confirms it; phase lag reaches 360° here. It is nearly always coupling capacitors that have the greatest impact on the LF stability margin, so it is desirable to minimise the number of them inside any feedback loop. In this example we could afford to make C_g larger, pushing P_3 lower, but we are free to play will all the capacitor values.

Something we have not accounted for is the power supply. At ultra-low frequencies where power supply decoupling capacitors no longer provide good bypassing, gain will try to rise, but the effect this has on the stability margin is nearly always negligible. Fig. 11 shows an example which is deliberately exaggerated. The power supply capacitor C_p and dropping resistor R_p create a pole at:

$$f_P = \frac{1}{2\pi C_p R_p} \qquad (1.6)$$

And a zero at:

$$f_Z = \frac{1}{2\pi C_p (R_p \| R_a)} \qquad (1.7)$$

Even with the unlikely component values used here, the ‘gain boost’ at low frequencies is less than 4dB and the phase shift less than 12°. Moreover, the phase is *lagging*, which therefore marginally improves rather than worsens the *leading* LF loop phase of the overall amplifier. Calculation is complicated when power supply filters are shared between multiple gain stages, but assuming the capacitors are a few tens of microfarads (why wouldn’t they be?) they are still unlikely to have much influence on the stability margin. With triode stages the effect of power supply filters is smaller still.

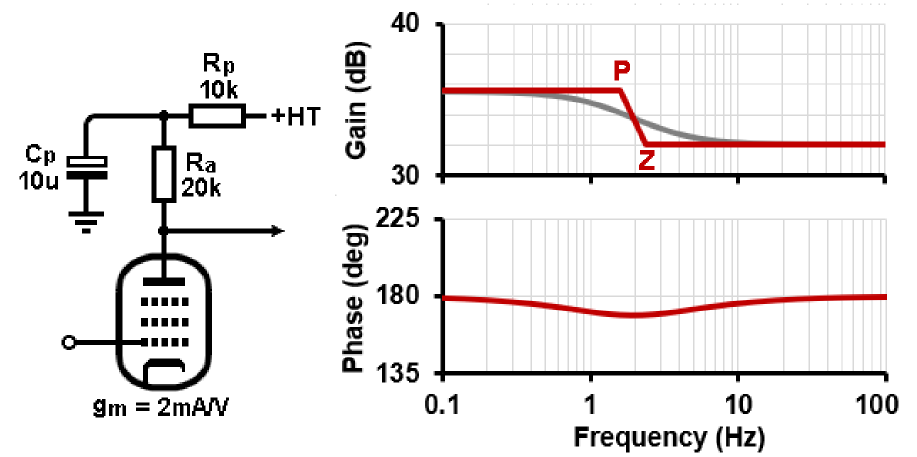

Fig. 11: Power supply decoupling capacitors have minimal effect on the stability margin, and it is in a favourable direction.

Closing Remarks

By working with gain and phase plots it is possible to build an intuitive understanding of how loop gain, bandwidth, and stability margin relate to one another. But so far I have only covered diagnosis of the patient, rather than treatment. We have not yet addressed the practical question every builder eventually faces: how to alter an amplifier that does not have enough stability margin for the amount of feedback desired. That will be the subject of Part 2 where I will cover the most common methods used to stabilise tube amplifiers, together with a real-world example that is taken from entirely unstable to firmly obedient. If you have a head for mathematics and want more quantitative information on feedback and stability, most of the references cited are available for free at: https://www.worldradiohistory.com.

Stabilizing Tube Feedback Amplifiers Without Math: Part 2

Well, almost no math! By Merlin Blencowe

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?

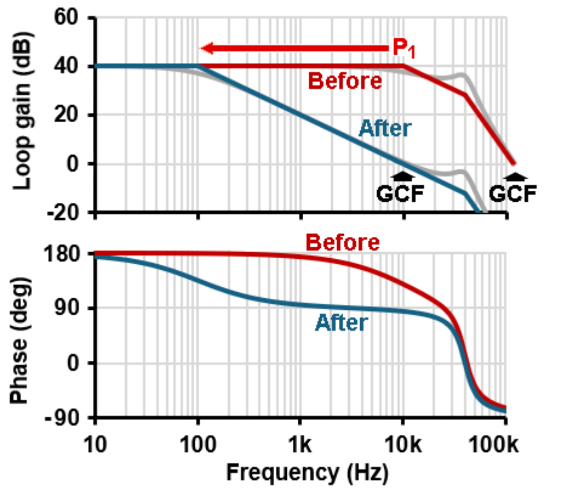

Fig. 1: Dominant pole compensation can stabilise almost any system, at a heavy cost of bandwidth.

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 to make it much lower, perhaps 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.

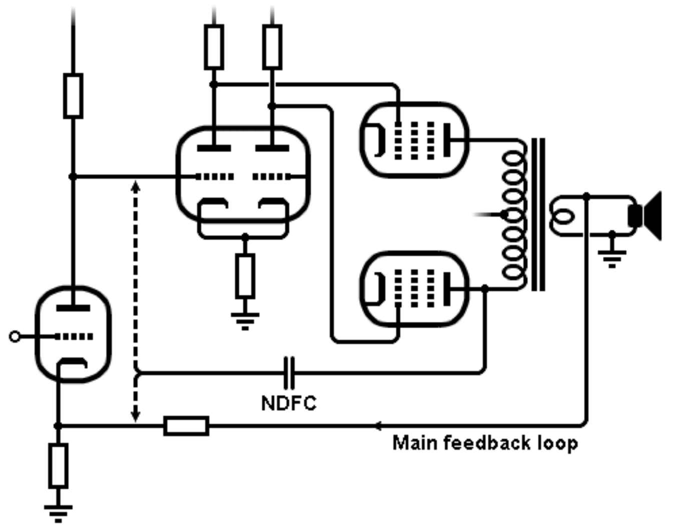

Fig. 2: A nested damping feedback capacitor is an effective stabilising component but needs a very high voltage rating.

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; this can be found 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 an amplifier since signal levels are smallest there, so should be least affected by the additional loading.

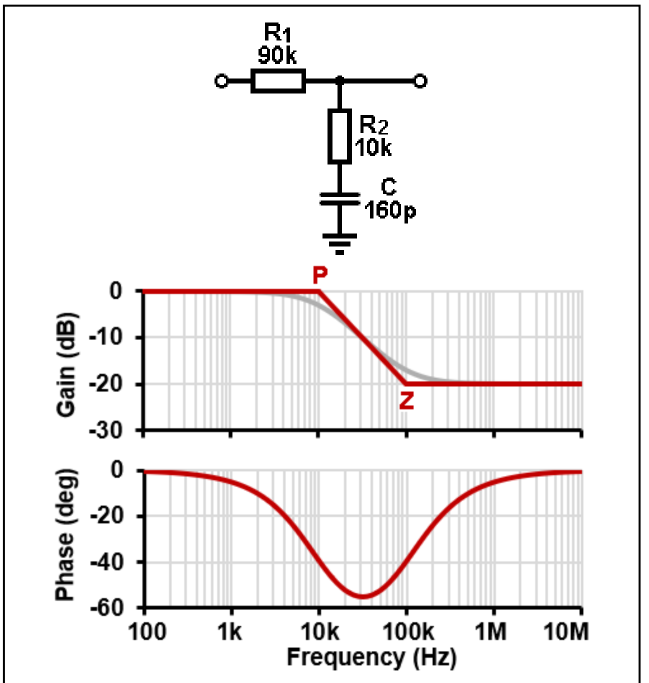

Fig. 3: A step network creates a pole-zero pair. Phase shift is maximum at the centre of the step but returns to near-zero on the flat sections.

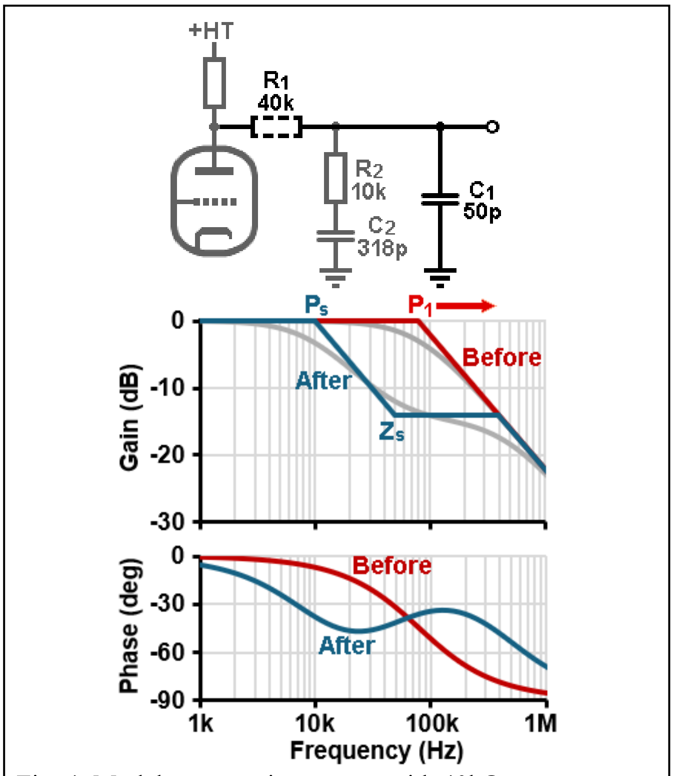

Fig. 4: Model representing a stage with 40kΩ output resistance driving 50pF of input capacitance. Amplitude and phase response before and after adding the step network R_2 and C_2.

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:

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

The gain or rather attenuation, A, of the step is:

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

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.

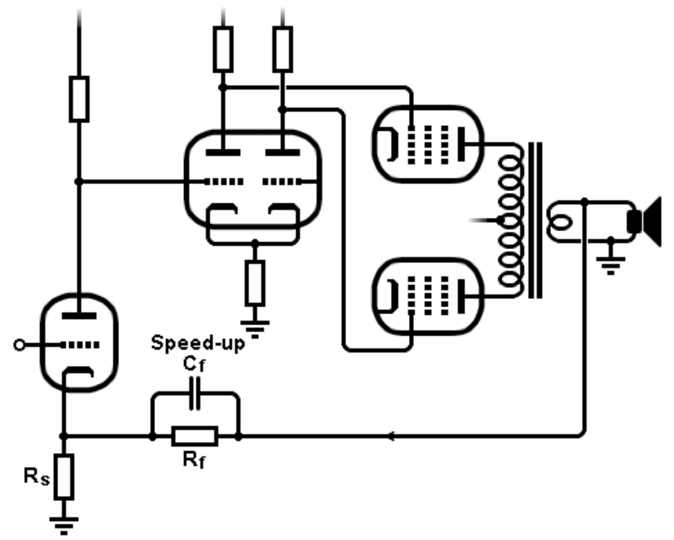

Fig. 5: A Speed-up capacitor is normally added last, to clean up ringing on square waves.

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.

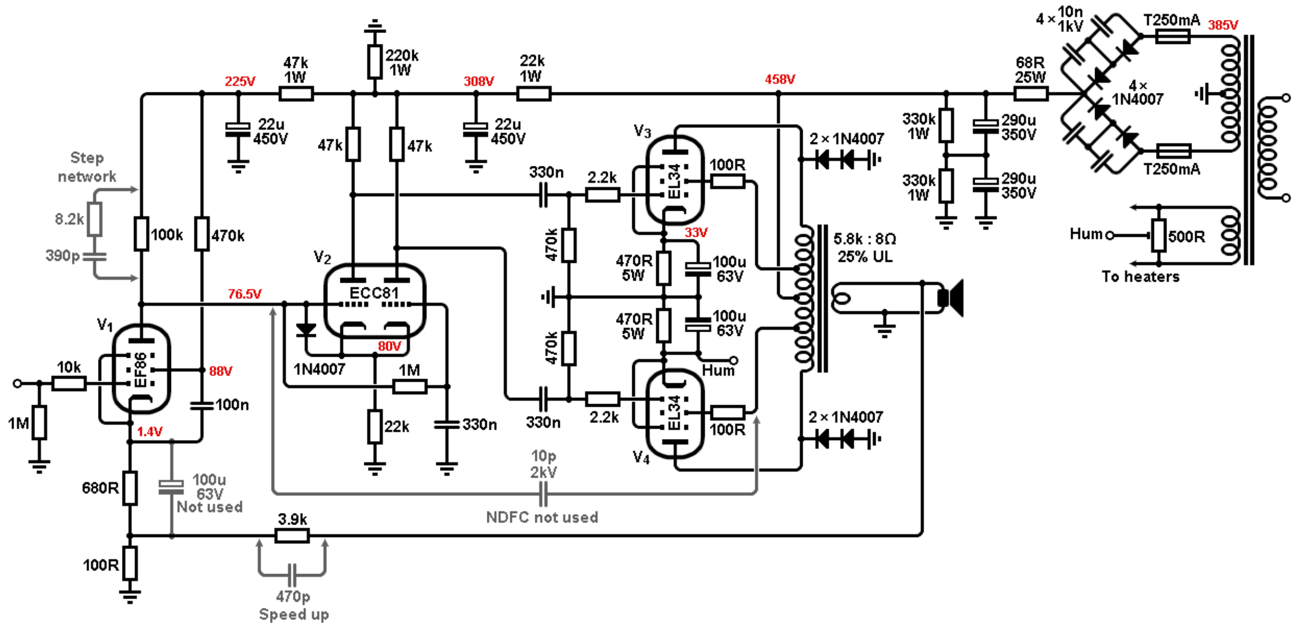

Fig. 6: Test amplifier circuit adapted from a Leak TL25+.

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.

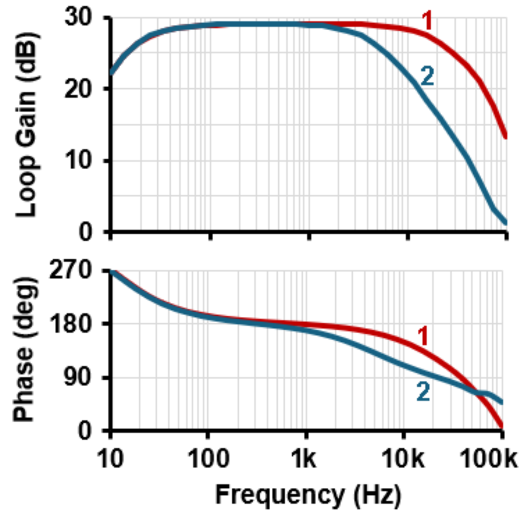

Fig. 7: Initial loop response is unstable (1). Adding a 10pF NDFC makes the loop stable (2).

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.

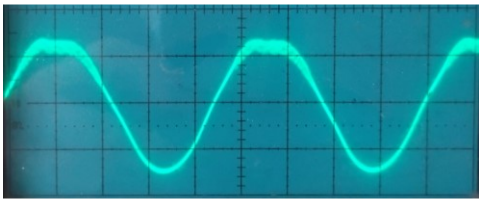

Fig. 8: Parasitic oscillation on a 1kHz output signal at the point of clipping, when using an NDFC.

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. 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.

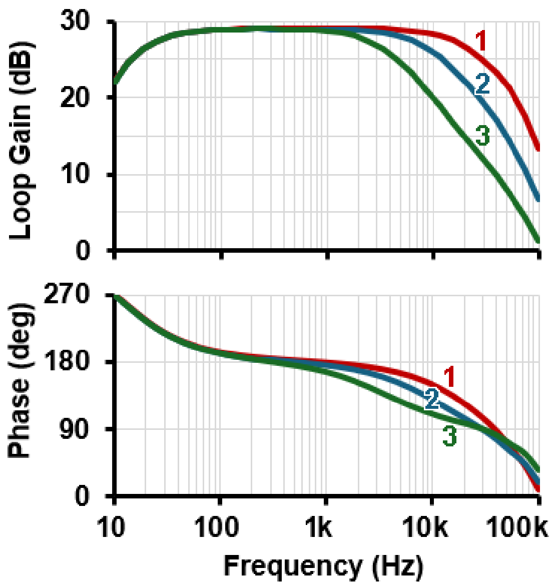

Fig. 9:
1) Initial loop response is unstable.
2) With 18kΩ+100pF step network it is still unstable.
3) With 8.2kΩ+390pF step network the loop is stable.

Fig. 10 shows the forward-path response of the amplifier, starting with its original 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.

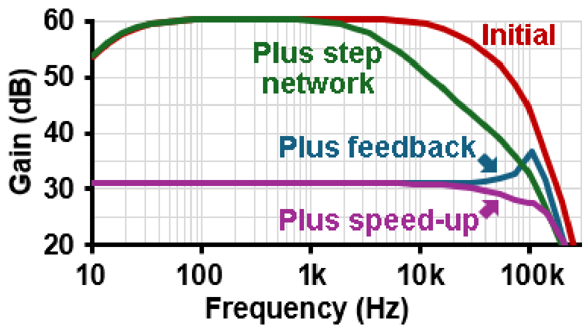

Fig. 10: Forward path response of the amplifier from initial open-loop, to final closed loop (ref. 1W at 1kHz).

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.

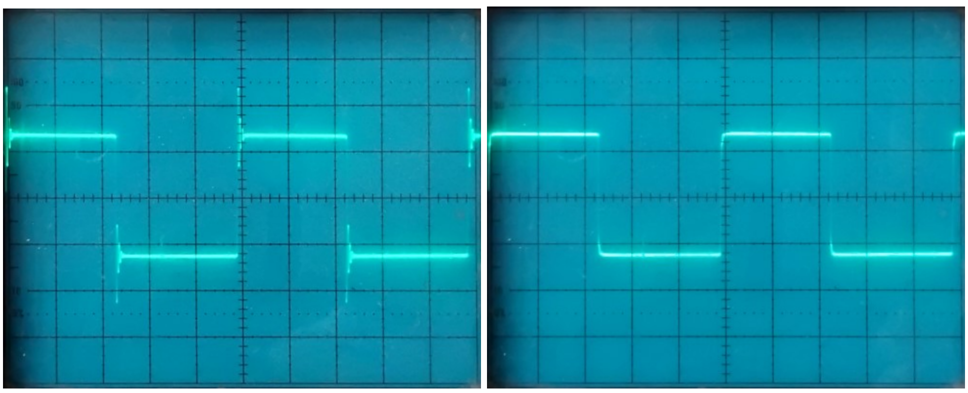

Fig. 11: **Left:** Ringing on a 1kHz square wave after closing the feedback loop.
Right: A 470pF speed-up capacitor eliminates ringing.

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.

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.

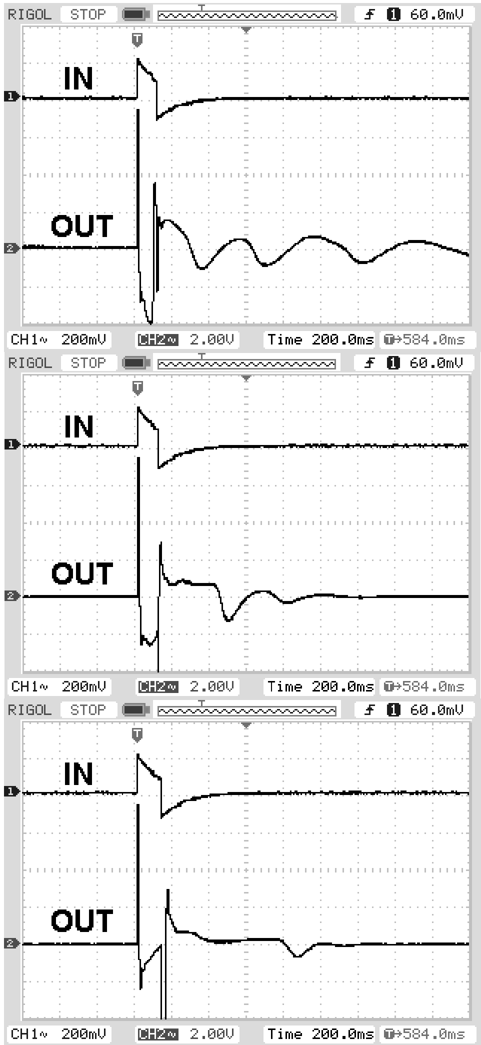

Fig. 12:
Upper: With 100nF coupling capacitors a pulse exhibits atrocious unsustained oscillation.
Middle: With 330nF coupling capacitors the oscillation is improved.
Lower: With 330nF also on the phase inverter second grid, oscillation is suppressed further.

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 was left at 100nF). Exactly what the LF stability margin is, is not easy to determine, although an absence of oscillation would indicate a phase margin much better than 60°. However, 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.

Primary inductance varies with excitation level, especially in push-pull transformers. The associated pole –which is often dominant– will typically be pushed to lower frequencies as the amplifier output power increases, eroding the stability margin (but the opposite will happen if the transformer approaches saturation, where inductance falls). For example, fig. 13 shows the difference in response to 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 principle is well demonstrated, 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.

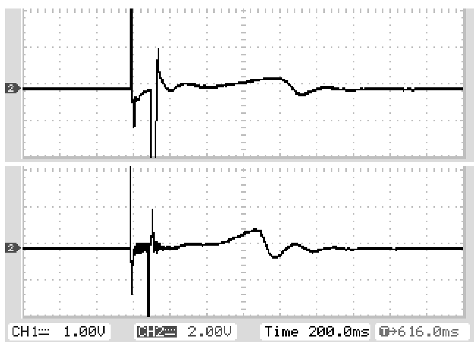

Fig. 13:
Upper: Response to a 0.2V input pulse.
Lower: Response to a 0.8V input pulse.

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.