

HiFi Collective OTL Headphone Amplifier

What is a OTL amplifier?

OTL (Output Transformer Less) amplifiers are simply a type of amplifier that does not use a transformer to couple to the loudspeaker. In general terms, an OTL amplifier refers to a valve amplifier, but to be precise just about all solid state amplifiers could be regarded as OTL's.

Valves, for all their wonderful sonic properties, have a major problem when it comes to driving loudspeakers or headphones. A valve is a high voltage low current device whereas a loudspeaker requires a high current low voltage source; the wrong way around. The conventional way these two opposing devices are married is through the use of a transformer. A transformer is a device that (amongst other things) can convert a low current high voltage signal to a high current low voltage signal suitable for driving a pair of headphones or loudspeakers. Physically, a transformer consists of two (or more) windings of wire arranged so the magnetic field produced by current in one set of wires induces a current in the second set of wires. In the case of output transformers, the wire is wound around a metallic core that increases the amount of power that can be transferred moreover, the number of turns of wire are very large. As a consequence, the resulting transformer is big, heavy and costly to make, especially if a good sound is expected. How many times have you heard it said that the most important part of a valve amplifier is its output transformer? here is much truth in that simple statement.

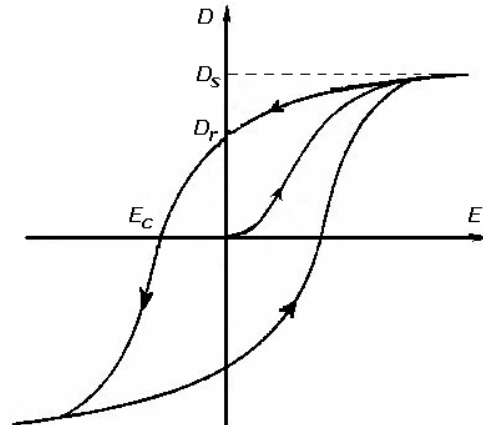
What's wrong with transformers?

Ok, we have now established that the simple way to connect a valve output stage to a loudspeaker is through the use of a transformer, so what's the problem, lets just use a transformer? Unfortunately, it's not that simple! First we have to consider the cost of making a pair of transformers. We can (though we would prefer not to) solve this by just spending more money. However there are other problems that are not so amenable to financial solutions. We want our transformer to pass the delicate audio signal without loss. To achieve this we want it to transfer both high and low frequencies equally well. Typically, audio engineers regard 20 Hz to 20 kHz as the audible range, but to achieve this we normally aim for an octave above and below (10 Hz to 40 kHz). To handle the low frequency end of this range and to transfer enough power to get a good bass response requires a big heavy core. A big heavy core means that many turns of wire are needed. Many turns of wire means that the high frequency response is limited unless complex winding layers are used. So to handle both the high and low frequencies equally as well the transformer will be both large and complex. These two conflicting requirements mean that winding a good transformer gets costly, very quickly.

Assuming we have solved the problem of transforming both high and low frequencies to our satisfaction, we then face another hidden problem. As the signal applied to the transformer changes with the music, there is a trace of the signal left behind. How

does this happen? When the signal level increases the core of the transformer becomes magnetised (like the permanent magnet in a compass), and as the signal drops again a trace of that magnetisation is left behind which then opposes the now dropping signal. The end result? Our transformer will always have an element of memory or smear; this process is called hysteresis. Different core materials have higher and lower levels of hysteresis, so its possible to reduce this with careful choice of the material and by use of expensive additions to the core (such as Nickel or Cobalt), or by altering the crystalline structure of the core (known as Amorphous).

Typical hysteresis curve for ferromagnetic material ¹



Furthermore, if we look at the small signal behaviour of a transformer we find that at very low levels of current a problem called magnetic stiction exists. What this means is that the core of the transformers has a certain 'stickiness' or resistance to change (as opposed to hysteresis which is a reluctance to change over time), so at a very low level the audio signal may not be changing enough to cause the core to change state, so the signal is lost. As a result, there will be loss of very low level information. One of the things that differentiate the high quality HiFi we seek from commodity products is the handling of this low level information.

And finally as if this all weren't enough, the wire that makes up a transformer can be many hundreds of metres in length. We worry about the half metre interconnect between our CD player and amplifier, but inside the transformer the length the signal needs to travel is far greater. We can get our transformer wound with silver wire to help with this, but then the transformer goes from being expensive, to ridiculous.

Why isn't every valve amplifier an OTL?

Having described the assorted problems with transformers, you may be wondering why anyone would ever consider using them? Well, the simple answer is a valve amplifier with all the limitations of its output transformer, still seems to make a sound that many like. If carefully designed and constructed an amplifier with output transformers can sound wonderful. The presence of the transformer also has another

advantage. Inside the valve amplifier you will find high voltages, normally over 400v often closer to 1000v. The transformer provides a simple secure way of keeping that voltage away from our precious loudspeakers, or in the case of headphones from the sides of our head. So any OTL design must also resolve this issue. Many of the valves we associate with audio, 300b, EL34, 2a3 and so on have been designed with an output transformer in mind. So to make a successful OTL amplifier we have to make use of types of valves that traditionally would not have been used for audio purposes and we go well beyond the safe and simple world of "off the shelf" designs. Getting a valve to drive a loudspeaker directly is a hard problem to solve; current production loudspeakers make it even harder. Modern loudspeakers are normally designed expecting to be driven by a solid state amplifier that can produce enough current to weld metal together. It's not uncommon to find loudspeakers that present loads to the amplifier that can drop below 2 ohms at times and, 4 ohms is common even with loudspeakers that are described as 8 ohm loads. It is possible to produce an OTL amplifier that will drive these sorts of loudspeakers, but it requires large numbers of valves all working together in parallel to drive the loudspeaker, so it becomes increasingly impractical to drive these sorts of speakers properly using an OTL. In truth, even valve amplifiers with transformers can have problems with the current demands of some speakers that offer loads below 2 ohms.

Headphones and OTL's

However it's not all bad news! We are considering a headphone amplifier here, and in general most high quality headphones in use present a much more benign load of 32 ohm or more to the amplifier, some models go as high as 300 ohms. The greater efficiency of headphones is also in our favour, we don't have to fill an entire room with sound, just the area around our ears. Using an OTL to drive headphones is a much more practical goal than driving a loudspeaker, and we have to make far fewer costly (both sonically and financially) choices to construct such an amplifier.

How do we make an OTL amplifier?

Well, in principle, it's a simple problem to solve. We need a valve output stage that has low output impedance (at least less than the impedance we are trying to drive, ideally much less) and can produce an undistorted output into the load impedance swinging enough voltage to provide the power required.

If you say it quickly, it's that simple. In practise, however, it needs a little more work. So we will start with a statement of requirements. Looking at a cross section of headphones, and wanting to produce a design that can be simply and cheaply built, the following set of goals seem reasonable.

Goals of the design:

1. Able to drive headphones from 32 ohm upwards.
2. Able to produce 2v RMS output signal into 32 ohm loads (125mW) and 8v RMS into 300 ohm loads (210mW).

3. A Frequency response flat within 0.2dB between 20 Hz and 20 kHz.
4. Distortion less than 0.05% at normal listening levels.
5. Noise and hum level such that it is undetectable in use.
6. Unconditionally stable into any (or no) load.
7. Built using easy to obtain parts.
8. Simple to setup and reliable in use.
9. Relatively cheap to build and run

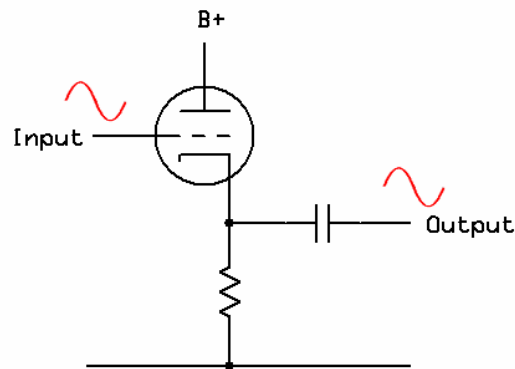
Points 1 to 6 will be addressed by measurements of the prototype amplifier.

Points 7, 8, and 9 will be addressed with a parts list and recommended supplier list and building suggestions at the end of this article.

What this article does not attempt to be, is a set of instructions such that with no prior experience anyone can build this. It will be as simple to build as any single ended valve amplifier, but it is still a valve amplifier, there are dangerous voltages present (less than 250v but still enough to ruin your day) and to build successfully and safely it will need some understanding of electronics, the ability to solder, follow circuit diagrams and the ability to use simple test equipment such as a multimeter. If anyone reading this is unsure of their ability the author will be happy to discuss constructing the amplifier for them, and of course I am happy for other builders to help their less adventurous fellow music lovers by building this amplifier for them (though in this case I would hope that the original design origin was remembered).

First Steps

The simple way of getting low output impedance from a valve is to use it as a cathode follower.



This is a circuit that provides no voltage gain but can drive a low impedance load. The output impedance of a cathode follower is equal to

$$Z_{out} = 1 / g_m$$

Where r_k is the cathode resistance and g_m is the transconductance value for the valve. There are high, medium and low transconductance valves, and the g_m figure is normally given as ma/v (milliamp per volt). So we can look at some common valves and their G_m and the output impedance they would give us when used as a cathode follower

Valve	g_m (ma/V)	Impedance (ohm)
300b	5.5	180
ECC83	1.4	720
6SN7	2.6	380
5687	8	125
6080	6.5	153
E810F	50	20

None of the above looks to be what we would hope for in terms of output impedance, but the E810F looks to be the best choice. However this particular valve fails our list of requirements in part 7. They are rare and getting more and more expensive to obtain. The common choice of valve for this sort of output stage the 6080 is good, but not great and the 5687 would look to be a better choice.

Another successful technique is the use of negative feedback to further reduce the output impedance. In negative feedback, we take a copy of the output voltage, compare it to the input voltage and drive the output valve with the difference signal. This will reduce the output impedance but at the expense of gain, to compensate we could add gain with extra voltage amplifier stages before the output stage. Every 6dB of feedback will reduce the output impedance by half. So in theory we could apply as much as we need to get the output impedance as low as we wanted. If we used a 5687, for example, we start with the 125 ohm output impedance, apply 24dB (16 times) feedback and get an output impedance of around 8 ohm.

However, there is a second requirement, just having a low output impedance on its own is of little use if we can not do work at the output. Doing work (in this case making a sound from the headphone) needs power, and to produce power as well as voltage we need current. To hit our 2v into 32 ohm target we will need 88ma of current; even our 8v into 300 ohm will need 38ma of current. Because the output stage we are using is single ended and running in class A, we need to be passing at idle at least as much current as we expect from the output. We will therefore need to be passing at least 88ma to allow us to produce the required output power into 32 ohm. Assuming we are running 100v or so across the output valve this means it needs to dissipate at least 9W at idle. That removes all the small signal valves such as the 5687 and E810F (assuming we even have some of them) from the running, however, the 6080 is still a candidate.

If one is not enough

Another thing we can do to handle more current is parallel two or more valves. Running two triodes in parallel will double the effective g_m , so half the output impedance, and also half the current that is needed through each triode. So, for example, four 5687 triodes (two valves as there are two triodes in each glass

envelope) would fit the bill. In this case, each valve could be run at 25ma giving 100ma in total, and the combined gm would be 32ma/v giving a output impedance of 30 ohm. This sounds workable doesn't it? Especially for the higher impedance headphones, and if we now add 12dB of feedback and we would have an output impedance 7 ohms. Unfortunately, there are two drawbacks to this. The first is to ensure that each triode passed its fair share of the current. If one triode is slight 'hotter' than the others it might try and hog the current and then it might be damaged. We could prevent this by adding balancing resistors, but doing so would push the output impedance up again. The second problem is the triodes need a load resistor to work into (the resistor below the cathode in the previous diagram). Cathode followers work best with large values of load resistor. The important figure for the minimum load resistor size is the anode resistance of the valve (ra). For a 5687 the ra value is around 2500 ohm. For best results, the minimum load should be 3 times or more the ra, so that gives us a minimum load of 7500 ohm. Because we have four triodes in parallel, we can divide that load by four to give a final value of roughly 1800 ohm . However, recall we want to pass 100ma or so at idle, that means (from ohms law) there will be 180v across the load resistor. So the load resistor will be dissipating a not insignificant 18W. That will need a metal clad resistor on a large heat sink. Now, it is still possible to build this amp, but the more we work through it's becoming less attractive. With 18W from the load resistor and 10W in the parallel triodes, we are releasing 28W of heat to generate a fifth of a watt of sound. Efficiency wasn't one of our stated goals, but even so, 1% is pretty appalling.

Clearly, gm on its own is not the only important factor in the choice of valves, and that the anode resistance of the valve is also important. Let's expand the table of valves previously given to include the ra figure (and add a new valve to the candidates).

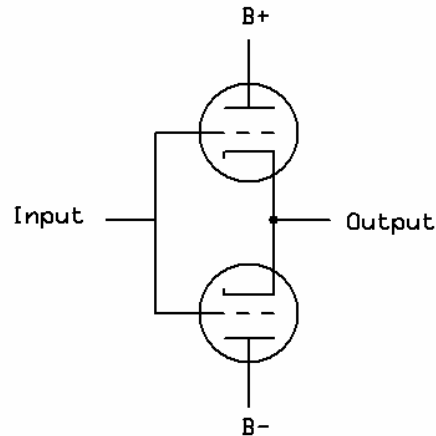
Valve	gm (ma/V)	Impedance (ohm)	ra (ohm)
5687	8	125	2500
6080	6.5	153	300
6c33c	39	25	120

From the above we can see that the standout valve is new newly added Russian 6C33C, both it and the 6080 were originally designed to be used as regulator valves in power supplies where low anode resistance and high current capacity is a requirement.

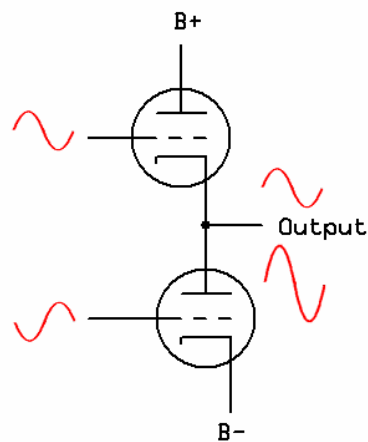
We can now see that by the numbers, a single ended 6C33C running at 100ma into a load resistance of 400 ohm with 6dB or so of feedback ticks all the right technical boxes. But it's not the solution chosen here. Why? It is already widely used for OTL headphone amps. The reason why we don't pick the 6c33c is this valve does a very good approximation of a small electric heater! The heater requirement for a stereo pair is 13 amp at 6.3v; approximately 81 watts. The surface of the valve itself runs at a very very high (read, take your skin off) temperatures. This is hardly in keeping with our goals 8 and 9; simple reliable, cheap and safe.

More than one way of skinning a cat

We saw earlier that paralleling valves can be a good way of solving our problem, and by using the two triodes in a single 6080 envelope we can get close to the performance of the 6c33c but, importantly, with only half the heater requirements. But we can do even better than that. There is another way of using two triodes together and that is to use them in push pull. What we would like to be able to do is this:



Unfortunately this is not possible, we would need to invent a valve that passes current from its anode to its cathode, and to power it we would need a source of positrons, not electrons. This is a bit problematic, so the best we can do with valves is something like this:

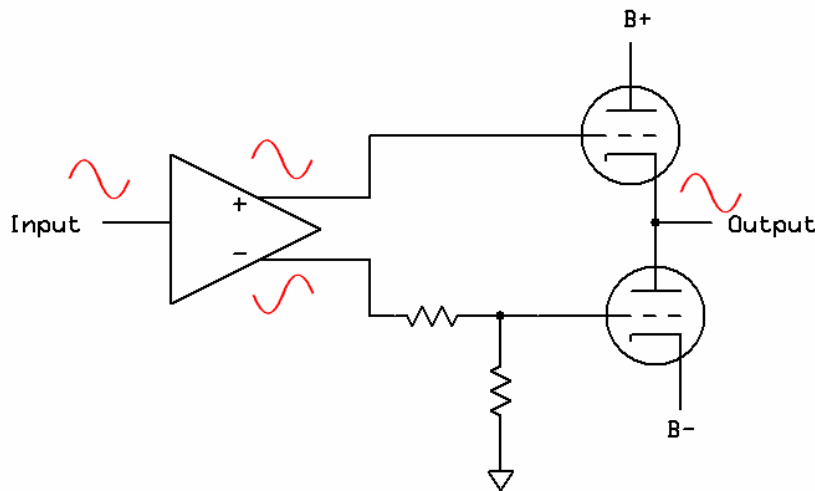


Where one triode drives the output from its cathode and the other drives it from its anode. To make this work the two valves need to be driven with opposite versions of the signal, so somewhere before the output stage we need was is called a phase splitter. We will talk about the phase splitter later. However, because one valve is driving from its cathode, and one from its anode we don't have a balanced pair of impedances. We saw earlier that the impedance at the cathode is given by

$$Z_{out} = 1 / g_m$$

But the impedance at the anode is given by the r_a of the valve. So again we want a valve with a high g_m and low r_a to keep the two impedances as close as possible. Now, returning to our table we find the 6080 has both a high g_m and a very low anode resistance that is only a factor of two different from the cathode impedance (150 and 300 ohms). So again the 6080 this looks like a good candidate. Being a twin triode, we also have the benefit of a single output stage in the one valve.

The slightly more serious hidden problem with the above diagram is that the top valve being a cathode follower as no voltage gain, so the signal on its cathode will be the same as the signal on its grid. But because the signal is taken from the anode of the lower triode, the signal on its anode will be multiplied by the gain of the valve, which for a 6080 is about 2. In other words, to get same signal on its anode as the top valves cathode we need to feed half the signal into its grid. That gives us this workable diagram for a push pull OTL amplifier to drive headphones. Now add 12dB of feedback and we have the output impedance we are looking for.



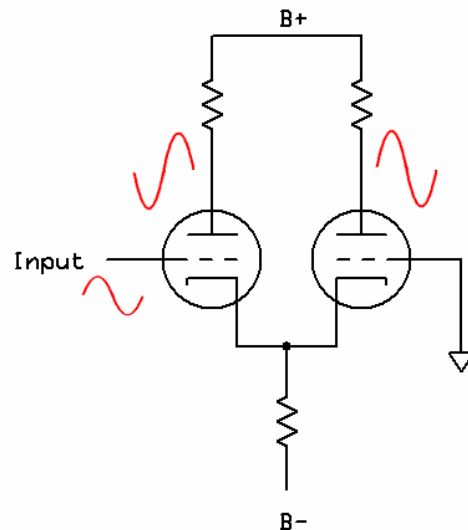
Another advantage of a push pull layout is that we don't have to restrict ourselves to class A operation, but we can let the amp move into B to handle the peaks of the signal, so we don't need to have 100ma flowing through both triodes at all times, and in fact, jumping ahead a little final amp will work well with just 35ma standing current.

Splitting the phase

The history of push pull amps has in many ways been the history of phase splitter. And many texts have been written analysing the merits of various types. This article will avoid talking about phase splitter in any great detail as any interested party may explore the subject further on the internet. However the requirements for our phase splitter are:

1. Good performance
2. Simple and cheap to build
3. Can be used to terminate negative feedback

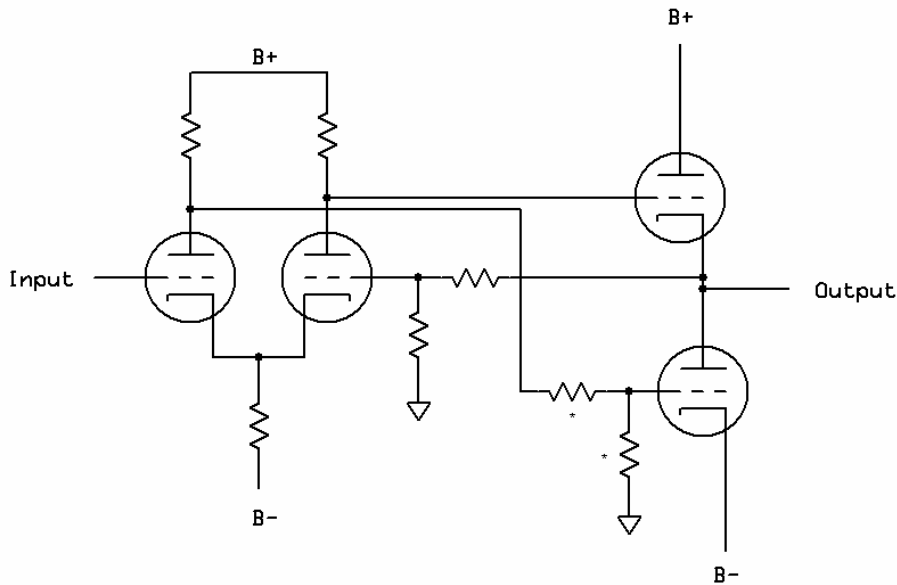
With these goals in mind, the ideal solution is the cathode coupled phase splitter, or to give it its other name the Long Tailed Pair.



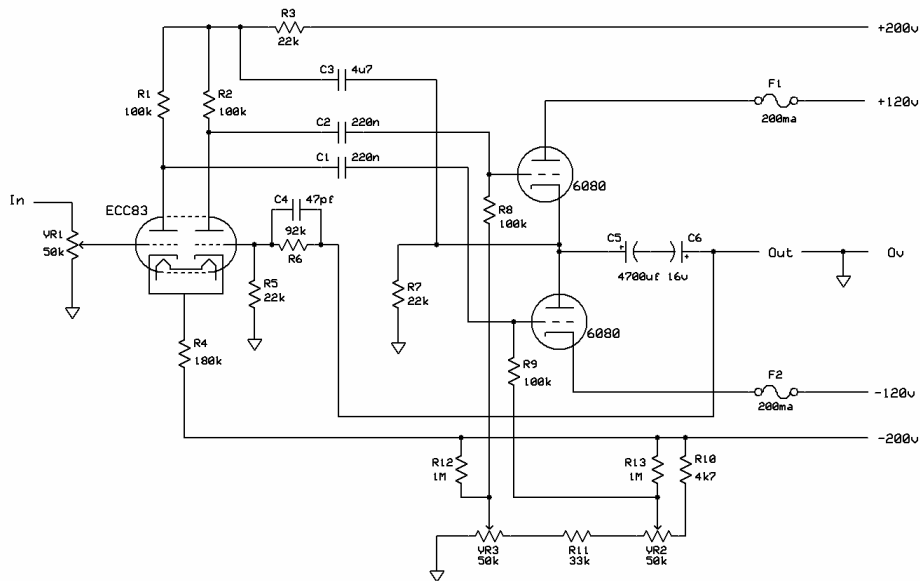
This circuit was invented by the great (and sadly often forgotten) engineer Alan Blumlein in 1936, and works by the first triode amplifying the signal as normal so producing an out of phase signal on its anode. As its cathode is unbypassed, the cathode circuit can be considered as a cathode follower, so will produce an in phase signal on its cathode. Because the second triode grid is fixed at ground potential, the signal on its cathode will produce an amplified signal in phase on its anode. So the signals on the two anodes are equal in amplitude, but of opposing phase; exactly what we need for a phase splitter. In fact, for the two signals to be equal, the signal on the first triodes cathode needs to be as close as possible to the signal on its grid. For that to be the case, the resistor in the tail needs to be as large as possible. Because we want a large value resistor in the "tail" there must be a correspondingly large voltage dropped across it. The simple way of doing this (without requiring a high voltage supply) is to use a split power supply with a negative supply attached to the bottom of the tail. As we already have a negative supply for the output stage (and will need a negative voltage source to bias the output valves) having a split supply is no real disadvantage.

The second advantage of the long tailed pair is it's not just a phase splitter; it's actually a differential amplifier. Why is this a good thing? Both of the triodes grids can have a signal injected into them, and the output on the anodes is two out of phase signals equal to the amplified difference between the two input signals. If we use the first input as our source, we can now use the second input to inject a copy of the output and compare it to the input signal in the differential amplifier; providing our

feedback So now we are almost at the final circuit. This circuit would almost make a perfectly good little headphone amplifier.



The only nagging doubt is the resistor/potential divider (marked with a * in the above schematic) that supplies the lower triode with half the signal of the upper triode. In the real circuit, we need to connect a bias voltage to the two output valve grids, so the resistor to ground is a problem and can not be used as shown. We could remove that potential divider and let the negative feedback do the work for us. Negative feedback can be a good solution to things like reducing the output impedance, but it's less of a good idea when it's correcting a problem that has a more appropriate solution. One of the things that help producing a good sounding amplifier is to get it working and sounding well and then and only then apply our impedance reducing global feedback. We solve the potential divider problem by applying another feedback loop, but this time it's not global, but local to the output stage. If we take the output signal and inject it back into the anode of the phase splitter, it will be in phase with one leg of the phase splitter output (so will add) and be out of phase with the other leg (so will subtract). If we feed 25% of the signal into the top of the phase splitter, one valve will get 25% more signal, the other 25% less, so that gives us our target of the upper valve getting half the signal of the lower one. Now, with no global feedback applied, the current through the two output valves will be equal for any load from 32 ohm to 300 ohm. We need 12dB or so of global feedback to give us the desired output impedance. Combining the gain of nearly 25 (28dB) of the ECC83 stage with the 3dB of local feedback leaves us needing to add an extra 12dB which the 92k/22k divider pairing shown provides.



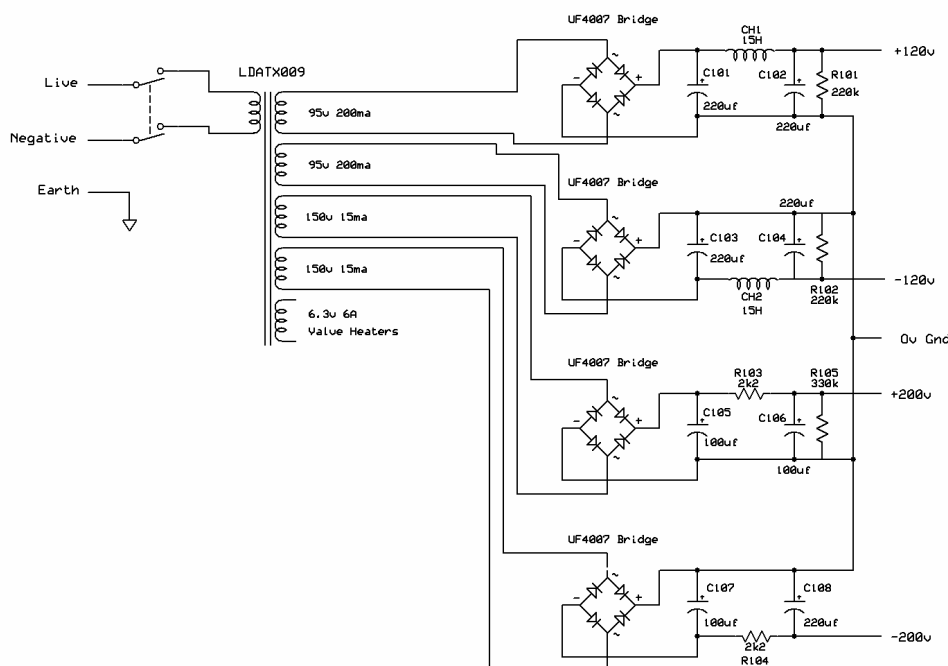
The final circuit has had a couple of extra things added. Firstly, the output has had a pair of back to back electrolytic capacitors inserted between the output and the headphone. The DC voltage at this point can be adjusted to less than 20mv in use, so the blocking caps are not strictly required, But at start up, the two triodes will not warm up exactly in lock step, so a volt or so of offset can develop during this time, clearly we do not wish to apply this DC voltage to your £300 HiFi headphones. Also if one of the triodes fails, a higher than expected voltage can occur at the output, the capacitors provide back up security between the valve's output and your ears. The fuses have been added into the output stage circuit for two reasons. Firstly, they protect the output from failed valves (6080 are hardly long lived devices that are not working hard in this circuit, but better safe than sorry). Secondly, We have to adjust our newly built amp and when first setting up the bias on the output valves, we first install the upper fuse and set the upper valve idle current, insert the lower fuse and set the DC offset to 0v, simple. That's all that is needed to set the amp up. Finally, and its very hard to say that its needed for sonic purposes, the small cap across the feedback resistor is there to compensate for the falling feedback as the input capacitance of the triode reduces the level of feedback. Without this cap, the amp has a slightly rising gain with frequency, so injecting a square wave will show a little overshoot on the output. This is really the icing on the cake as the response only starts to take effect at 50kHz and above and its debatable if there is any real need, but as the amp is still going strong at 100kHz it seems a shame not too. Perhaps also is just the perfectionist nature of the author instead of any real audible reason.

Importantly, because our new amp has only single resistors and capacitors in the signal and feedback path that both with and without feedback applied the amp is unconditionally stable.

The bias arrangements for the two output valves also is also shown giving you the complete circuit and final circuit that has been built and tested and is currently in use on my desk.

Power supply

As is typical in most OTL amplifiers, the power supply is slightly more complex than for a conventional amplifier. The power supply needs to provide two split rail supplies, + and - 200v at 5ma or so, and + and - 100v at 100ma or so. The circuit shown could be simplified by using centre tapped transformer winding to provide both rails from a single supply. However, having the two rails as shown means that during setup, the power rails will not vary as their counterpart is disconnected. This makes setup much simpler, and has only a small impact on the parts count and cost.



The amplifier needs fairly well smoothed supply rails. The amplifier is more sensitive to ripple on the negative rails, but by using the components shown the amplifier is silent in use.

Construction and components

The prototype was constructed using a custom toroidal transformer wound by Canterbury windings and available using the reference number shown in the parts list. The chokes used were by Hammond and are commonly available (also see the parts list). Most other parts used were supplied by HiFi Collective, and a few odds and ends from Rapid Electronics and the parts bin. Takman resistors were used in the signal circuit, and all the coupling caps were Clarity Cap SA types. The amplifier may well benefit by the use of even better capacitors, but the Author has not tried this at the time of writing.

The prototype was constructed in a Hammond Aluminium chassis, using a piece of FR4 board with the valve holders and turret tags to mount the components. The circuit will be relatively insensitive to circuit layout so the constructor can use their preferred method. The prototype was built with a 50k Motorised potentiometer as a volume control and was used in conjunction with a HiFi Collective remote volume control board to provide remote volume adjustment and start-up output muting.

Suggested upgrades

As always it's possible to improve things. The design presented is intended to provide a good performance/cost ratio. But the sound can always be improved on. The first thing that the builder can try is altering the level of feedback. Increasing the value of R6 will reduce the global feedback; this will increase the gain of the amplifier and increase the output impedance. A Lower value will increase feedback, reducing gain and decreasing the output impedance. The gain required will be determined by the builder's source signal levels, their headphone sensitivity and the desired listening levels. Sonically; increasing the level of feedback will make the sound tighter and cleaner, but will also make the sound dryer and more constrained. Reducing the feedback level will open the sound up and increase the subjective warmth of the sound. The design value of R6 (92k) was chosen by the author as the best compromise with his source, headphones, music and ears. The builder may wish to adjust this to his or her own taste.

The output capacitors chosen will have some impact on the final sound, the published value (4700uf) was chosen to allow extended bass response into 32 ohm headphones. If higher impedance headphones are used smaller value of output capacitors could be employed. The prototype was constructed using the Nichicon FG type which was judged to give good performance at low cost. The builder may also wish to experiment with bypassing the electrolytic capacitors with a film capacitor.

The coupling and feedback caps used (C1, C2 and C3) will also have an effect on the sonics of the amplifier. The prototype was built using ClarityCap SA types in these positions. The smaller capacitors (C1 and C2) are the most important ones, so it could be worth experimenting using different caps in these positions.

The volume adjusting potentiometer at the front of the amplifier will have an impact on the overall sound, and as always it's worthwhile to replace the variable resistor with a stepped attenuator. The prototype was built using a Alps Blue 50k unit. The choice of the value of the potentiometer is not important to this amplifier, and it should be chosen with regard to the driving equipment. But increasing the value of this potentiometer above 100k will start to cause the miller capacitance of the first valve to reduce the high frequency roll off point of the amp.

As always with a long tailed pair, the higher the value of the tail resistor the better the stage functions. With this in mind, an obvious modification is to add a constant current source in this position. A solid state CCS can be used here, it needs to be set to pass 1.1ma and will have 200v across it.

Setup and testing

After your build is finished or when changing output valves, it is necessary to set up the bias voltages to the top and bottom 6080 triode.

If you have just finished the build, I would recommend the following check-list.

1. Remove all valves, and all output stage fuses.
2. Apply power and check the voltages on the power rails. Expect all the voltages to be higher than under normal use as they are being run without a load.
3. Power down and insert the two ECC83 valves.
4. Power up again, check the heaters are lighting, leave a moment to warm up and check the voltage at the ECC83 anodes and cathodes. The anodes should be around 120v and the cathodes about 1v. All voltages should be checked with the negative terminal of the multimeter connected to the 0v line.
5. Check the voltages on the grid pins of the output triodes (before inserting them). Adjust the bias pots to give -60v on the top triode and -180v on the bottom triode.
6. Remove power again, wait a few seconds.
7. Switch the multimeter to measure current, and with the negative probe still on the 0v line, connect the positive terminal to top triode cathode. Insert the fuse in the top triode anode (F1) of the left channel. Power up again, and monitor the current as the valves warm up.
8. Once the valves are warm, and the current has stopped changing, adjust VR3 until the measured current is 40ma.
9. Power down, insert the top fuse (F1) into the right channel, and connect the positive probe to the anode of the top triode in right channel. Power up and repeat step 8 to set the current to 40ma again.
10. Power down. Move the multimeter probe to the left channel again, still to the top valve cathode, and insert the lower fuse (F2). Change the multimeter to now measure voltage.
11. Power up again, let the valves warm up, and adjust the bias pot VR2, until the meter reads 0v.
12. Power down, move to the right channel, insert that channel's F2, move the multimeter to that channel, power up and repeat step 11 to again read 0v.
13. As each channel is adjusted, the load on the power supply will cause the voltage on the other channel to change slightly, so repeat steps 10 to 12 several times until both channels read 0v when warmed up. It should be possible to get both channels to have less than 25mv offset at the test point. It is not vital to get this voltage exactly to 0v, and over time as the valves age differently it may drift. Any voltage less than 2v will cause no detriment to the operation of the amplifier, but it will be worth checking every 6 months or so to compensate for valve ageing.

When replacing the output valves, repeat steps 5 to 12.

Measured performance

Voltage gain: 11dB

Output before clipping:

300 ohm load 14.5v RMS (700mW)
32 ohm load 3.75v RMS (440mW)

Frequency Response:

-3dB points 2Hz. - 85kHz
-0.2dB points 17hZ to 22kHz

Output Impedance: 21 ohm

Distortion

1V RMS into 300 ohm at 1kHz

2nd Harmonic -72dB (0.025%)
3rd Harmonic -74dB (0.02%)
Other harmonics below -85dB (less than 0.005%)

0.25V RMS into 32 ohm at 1kHz

2nd Harmonic -72dB (0.025%)
3rd Harmonic -72dB (0.025%)
Other harmonics below -85dB (less than 0.005%)

Parts List

Single channel of amplifier (double for stereo).

Resistors

R1, R2, R8, R9 – 100k 0.5W

R3, R5, R7 – 22k 0.5W

R6 – 92k 0.5W

R10 – 4k7 0.5W

R11 – 33k 0.5W

R12, R13 – 1M 0.5W

VR1 – 50k logarithmic potentiometer, or switched attenuator

VR2, VR3 – 50k 0.5W 10 turn pre-set potentiometer

Capacitors

C1, C2 – 220nf 630v film

C3 – 4u7 630v film

C4 – 47pf polystyrene

C5, C6 – 4700uf 16v electrolytic

Assorted Hardware

F1, F2 – 200ma time delay fuse (with holder)

V1 – ECC83 double triode

V2 – 6080 double triode

B9A Valve Holder

Octal Valve Holder

Headphone jack socket

RCA input socket

Power supply (supplying both channels)

Resistors

R101, R102 – 220k 2W metal film

R103, R104 – 2k2 2W metal film

R105 – 330k 2W metal film

Capacitors

C101, C102, C103, C104 – 220uf 250v electrolytic

C105, C106, C107, C108 – 100uf 350v electrolytic

Transformers

Mains transformer – LDATX009 Custom Mains ²

TX. 95v, 200ma * 2,

150v 15ma * 2,

6.3v 6A

Electrostatic screen

Primary to suit local mains supply

CH1, CH2 – 15H 100ma. Hammond 159M or equivalent ³

Semiconductors

D1 – D16 – UF4006 800v 1A ultra fast rectifier

Assorted Hardware

IEC socket

Main switch DPDT 240v

References

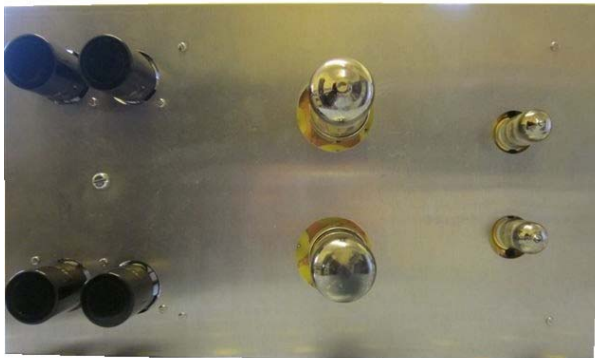
1. Image courtesy of [Ndthe](#) from Wikipedia

2. Custom transformer (<http://www.canterburywindings.co.uk>)
3. Hammond chokes from Bluebell Audio (<http://www.bluebellaudio.com>)

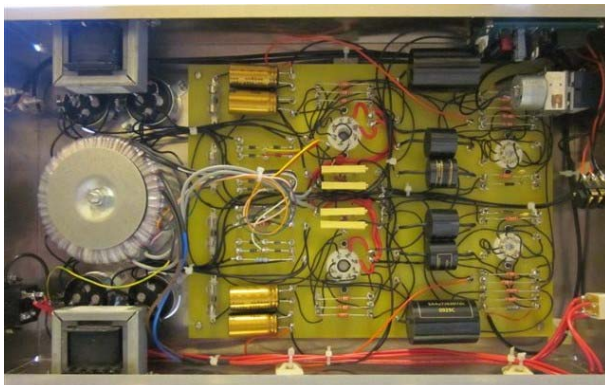
Pictures of the Finished prototype



External view 1.
(note there are 4 exposed electrolytic caps, not visible here as behind the 6080s)



External view 2.



Internal view 1.