Totem-Pole Output Stage PSRR Enhancement

In all the push-pull topologies, the goal is the same: to provide an equal drive signal for the output devices. This goal is important, as it ensures a low distortion output signal by forcing each output device to work equally into the load impedance, a task required for low distortion operation. But even when we have labored to ensure an equal drive signal, we may still find the output signal tainted with noise from the power supply.

Don't all push-pull output stages reject the power supply noise from their output? Some come close, but not all push-pull output stages see a balanced load impedance or an equal share of power supply noise or *are run in true Class-A mode*. This last point is almost never mentioned, but should be, as the push-pull's two output devices must be conducting current in order for the noise to drop out of the output signal. While at idle, both output devices of a Class-AB amplifier certainly do draw some current, so the cancellation of power supply noise is great. But once any one device stops conducting, it can no longer work to cancel the noise. Unfortunately, as amplifiers are only tested for noise at idle, the real dynamic noise rejection figure is never measured.

Therefore do not conclude that if an amplifier is push-pull, it must already be noise rejection optimized. Thus, our added goal is to understand the transmission mechanisms used by power supply noise to make its way to the output.

The techniques outlined in this article are also applicable to pure tube OTL amplifiers or hybrid amplifiers. As a consequence, last issue's use of a schematic symbol made up of a box with a capital "N" in it will be retained to illustrate the universality of the techniques. Once again, this symbol denotes either a triode, or a pentode, or a N-channel MOSFET, or a NPN transistor. All these devices are functionally identical in that they can be used as output devices in the six basic output topologies.

Split-Load Unity Gain Circuits

This topology admits two variations. The first is marked by the cathode side connection, where the output directly (or, somewhat indirectly, via a coupling capacitor or a zener or a voltage regulator) couples to the other side of the phase splitter's cathode resistor. Thus, any perturbation at the output is relayed to the phase splitter's plate resistor. How? Since the cathode resistor sees a fixed voltage at its cathode end, the perturbation at the amplifier's output provokes an increase or decrease of current flowing through the cathode resistor. This increase or decrease of current *finds only one current path to the B+ connection*: the path through the tube to the plate resistor. Thus any perturbation at the output is relayed to the bottom output device's input, which then causes either a increase or decrease in the bottom device's current conduction to counter that perturbation.

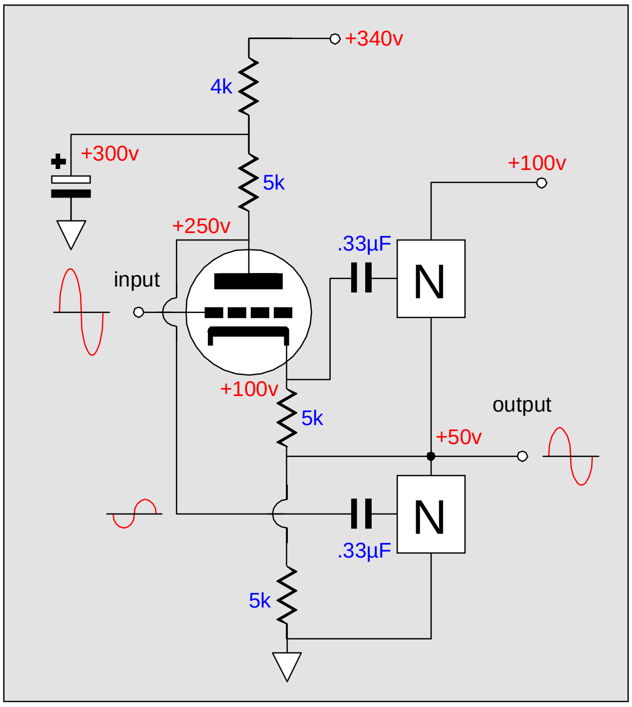

Split-load, cathode-referenced, unity-gain output stage

The second topological variation is marked by a plate side connection, where the output directly (or indirectly via a coupling capacitor or a zener or a voltage regulator) couples to the top of the phase splitter's plate resistor.

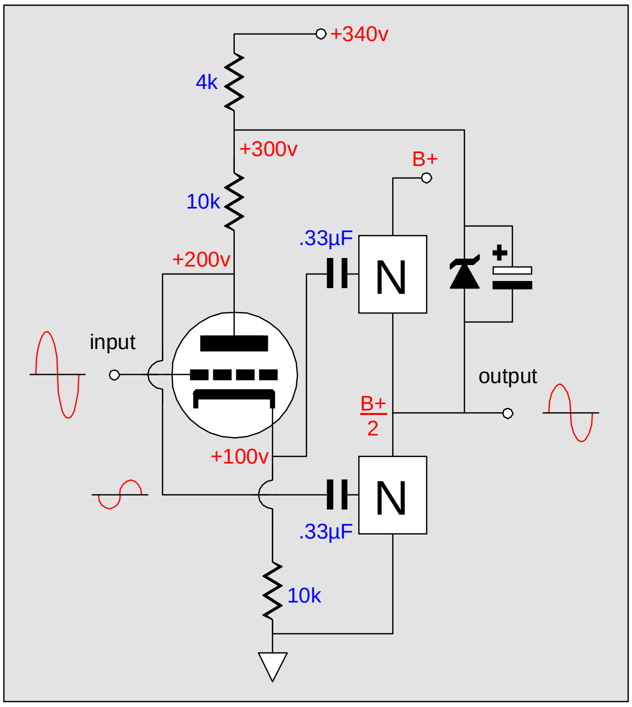

Split-load, plate-referenced, unity-gain output stage

Thus, any perturbation at the output is directly relayed to the phase splitter's plate resistor. Since the triode's effective rp is hugely increased by the unbypassed cathode resistor, the full magnitude of the perturbation at the output is relayed to the bottom output device unattenuated. This error signal provokes either an increase or a decrease in the current conduction of the bottom output device in an effort to cancel the perturbation.

In both topologies, the top output device sees a constant ground referenced input signal, which means that the top device will work as hard as the bottom device to keep the output inline with the intended signal. Furthermore, in both topologies, the bottom output device provides no voltage gain, as all potential gain is entirely fed back to its input via its connection to the split-load phase splitter, which results in 100% degeneration and no gain, just as in the case of a cathode follower. So, in other words, *both output devices function as followers*. When used in a Class-A amplifier, they effectively equal both devices placed in parallel, which means a doubling the transconductance.

For example, a MOSFET based output stage that used two devices with a transconductance of 1 A/V achieves a 2 A/V input voltage to output current ratio. If used, however, in a lean Class-AB or Class-B amplifier, the output stage's effective transconductance only equals that of a single device (or one bank of output devices).

Which topology is better? This was a trick question, of course: there is not a better topology without specifying what the intended use is or what restrictions are placed on the circuit.

For example, if the output devices are pentodes, then the plate referenced variation is preferred, as it provides us a means to driving the top pentode's screen with the an output referenced voltage. This was exemplified in the New York Audio Lab's versions of the Futterman OTL amplifier. Using voltage regulator referenced to the output, these amplifiers fed the output of the regulators to both the top output pentode's screen and to the plate side of the split-load phase splitter. Thus, killing two songbirds with one voltage regulator, if you please. If you pause to think about it, you will realize that a series voltage regulator is nothing more than a voltage shifted follower with the ground as the input source. (In fact, three-pin voltage regulators can be used as output devices in an amplifier. The circuit below shows a single-ended, unity gain, short-circuit protected amplifier made out two LM317s.)

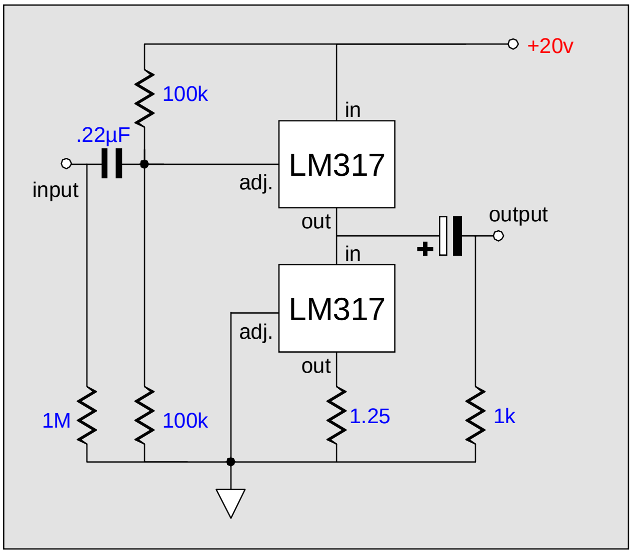

On the other hand, the cathode side connection can eliminate the need or any extra components, as the output is usually at ground level, which allows for DC coupling. Which circuit to use depends largely on which power supply noise canceling trick is used.

As both topologies stand in isolation, the easiest path the power supply noise takes to the output is from the limitations of either the power supply shunting capacitor used in the plate side connection variation or the limitations of the voltage regulator or signal coupling capacitor used in the cathode side variation. Should any of these fail to isolate the power supply noise from the output stage's inputs, noise will leak into the output signal.

But neither topology is used isolation: *thus the input stage often provides a connection to the power supply noise*. Of course, we can strive to eliminate the power supply noise from polluting the first stage's output by using large chokes and filter capacitors or by even using regulation. Or we can use the output stage's common mode rejection ratio to our advantage.

(The following technique works best with a Class-A operation mode for the output stage, but it is still useful with Class-AB output stages.) Remember that push-pull output stages need to see a pair of balanced drive signals, i.e. out-of-phase signals. When the push-pull output stage is presented with in-phase signals, the output stage should ideally provide zero amplification. Thus our goal is to ensure that whatever amount of power supply noise that leaks through must be presented to each output device equally in both phase and amplitude.

The input circuit below realizes an equal power supply noise distribution by halving the power supply noise at the first triode's plate, as one half noise subtracts from one noise to yield one half noise at both outputs. This trick requires that Rk equal (Ra-rp) / (mu + 1), as this is the only ratio that halves the power supply noise. Here is a case where the circuit is device specific, as only a triode will work in the first tube's position, as only the triode has a low rp.

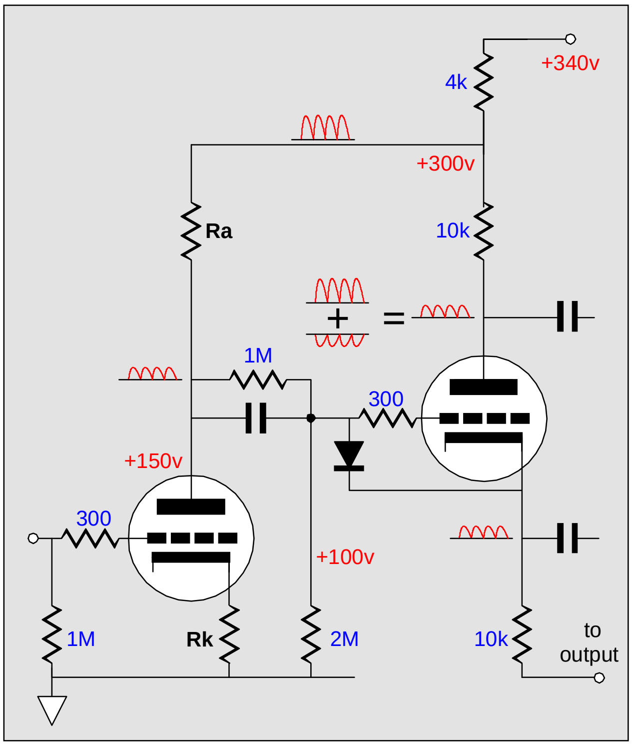

PS noise balancing front-end

If a pentode or FET is used as the input tube, then a 50% voltage divider is needed. Two equal valued resistors will work. The trick will be to find the value for these resistors that is not so low as to excessively load the first stage and not so high as to limit the high frequency response of the signal leaving the split-load phase splitter. Understand that this voltage divider will also divide the audio signal from the first tube. Still eliminating noise often comes at a price.

(A key point here is that both of these front-end circuits rely on the first stage seeing the same amount of noise as does the split-load phase splitter. Paradoxically, this requirement might require the removal of some preexisting power supply filtration circuitry in order to realize a quieter amplifier.)

This noise rejection optimized front-end has a wider scope than just hybrid and OTL amplifiers; it should be used whenever a grounded cathode amplifier cascades into a split-load phase splitter; such as is found in the first half of the Williamson amplifier and even the Dynaco ST-70 and Mark-3.

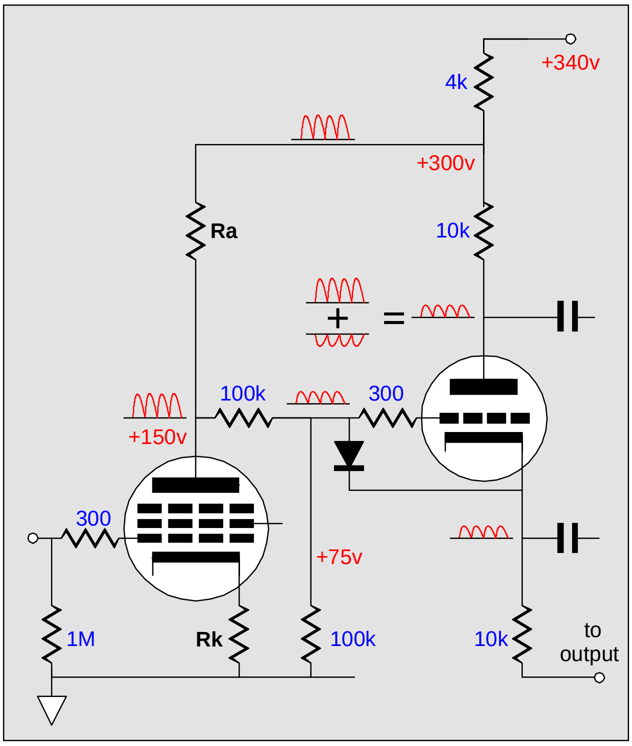

Pentode front-end with a resistor voltage divider

Testing this front-end is easy enough: just ground the input and observe the noise on both outputs of the phase splitter. Should they differ in either phase or amplitude, the output stage will then treat the noise as signal to be amplified.

Split-Load Gain Circuits

Once again we have two variations to choose from: plate referenced or cathode referenced. In both cases we have an output stage that provides gain at the cost of a much higher output impedance and a greater distortion figure. In the previous examples the output devices functioned as followers, but as they are configured here, they function as either grounded cathode or source or emitter amplifiers. But more gain means more potential feedback, so what we really have is a zero-sum game, wherein the feedback can be applied locally or globally. Which is better? The global feedback network will provide a lower distortion figure. This is so because all the feedback can be used to iron out the most guilty device; whereas its own gain would set a lower limit to how much distortion 100% degenerative feedback could eliminate. (Still, my bet is on local feedback. I just do not trust a global feedback loop that envelopes several gain stages. And I worry about the clipping behavior of the global feedback based amplifier. Yet I have heard such amplifiers sound good. More experimentation is needed.)

The circuit bellow is cathode referenced. The difference between this circuit and the unity-gain one is that inputs to the output device have been switched: the top device is driven from the plate and the bottom device is driven from the cathode.

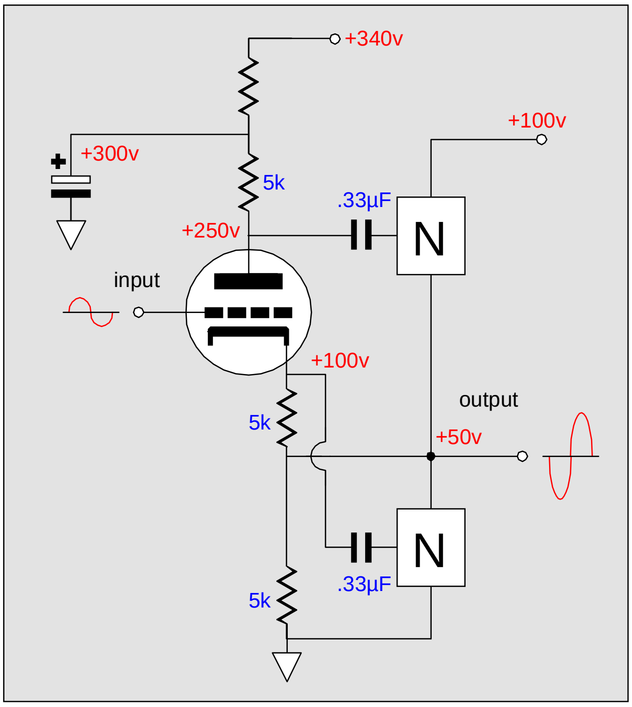

Split-load, cathode-referenced, gain output stage

The bottom device is obviously a grounded "something" amplifier. This is easy to see. But the top device also looks like it is functioning as a follower, as the output is taken at its cathode or source or emitter. The twist lies in the top device's input being referenced to the output and not ground. What difference does this make? Considered the what happens when the output is externally forced up 1 volt by a quick going pulse. When the top device is ground referenced, the pulse moves the cathode or source or emitter to become 1 volt less positive than its grid or gate or base.

Thus the device decreases it conduction equal to its transconductance, if the idle current was high enough to absorb that amount that is; or it ceases to connection altogether, if the idle current is not high enough. And a negative going pulse will provoke an increase in conduction equal to the transconductance of the device. This is classic follower action in a nutshell.

On the other hand, when the top device is output referenced, the positive going pulse is relayed in its entirety to the input of the top device. In other words, the cathode or source or emitter may have been bumped up 1 volt, but then so has the grid or gate or base. And as it is the voltage relationship between grid and cathode or gate and source or base and emitter that controls the flow of current through the output device, the output device effectively does not see the pulse. This explains why the output impedance is so high. (The triode has an advantage here, as its rp will buck the pulse, which will offer some output resistance.)

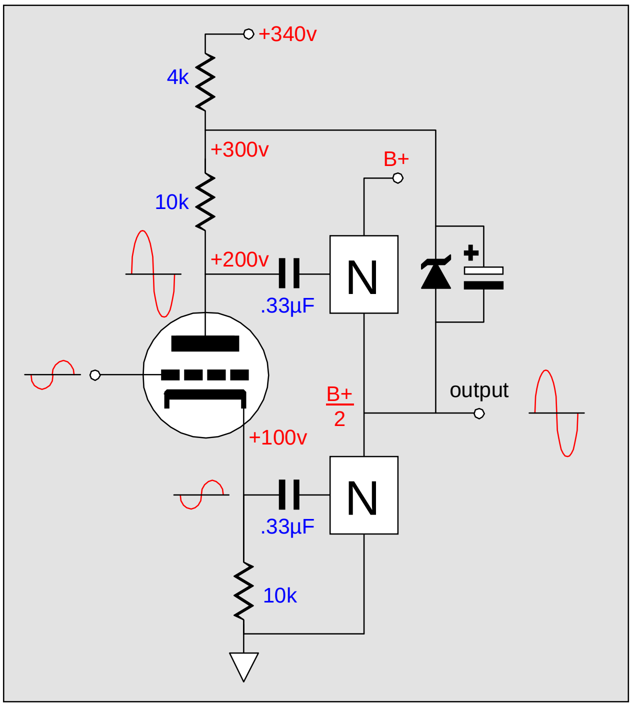

Split-load, plate-referenced, gain output stage

OK, the high output impedance makes sense, but where does the gain come from when the load attaches at the cathode or source or emitter?

The answer is the same place the gain came from when the load attached to the plate or drain or collector: the transconductance of the device gives rise to a change in current flow, which when this flow also travels through the load, creates a voltage across the load's resistance or impedance. Thus, the rough equation for gain is gm against the load: gain = g_mR_L. The top device sees the same magnitude of input signal from its input to its output s the bottom device sees. The only difference is that the bottom device's voltage reference is ground and the top device's reference is the output of the amplifier. If the load is a dead short to ground, then both top and bottom devices become ground referenced and the drive voltage relationships become obvious.

Admittedly, when the load impedance is some value greater than zero, the drive relations become more difficult to see. Still, as long as we remember that the reference has moved from ground to the output, we stand a better chance of mentally unraveling the drive voltage relationships. Here is a different example of a change in reference, in point of view.

A man stands in a large basket-like structure. In front of him his a small vertical lever that is labeled "Up" at the top and "Down" at the bottom. Finding the lever at the mid-position, he pulls the lever completely up and then completely down. As he stares at it, the lever has only traveled half a foot from center to either extreme. But to the onlookers on the street, who see a man standing in the container at the end of fire engine's crane-ladder, the lever (along with the man) actually travels 20 ft from one extreme to the other. Who is right? What is your reference? If it is that of the man in the basket, then the lever only moved half a foot from its center position. If it is that of the onlookers on the street, then the lever actually moved 10 feet up (plus the half a foot) and then 10 feet (plus the half a foot) down from its center position.

After that long recapitulation of how this circuit works, let's turn to how the power supply noise leaks into output in this circuit.

Will the halving of the power supply noise trick from the circuit work here? Yes, if we carefully choose the value of capacitor C1, then the amount of noise at the top of C1 will match that at the top of 33µF capacitor and the balancing of the power supply noise will be assured. These two capacitors cannot share the same value, as they do not see the same shunt impedance. The first stage represents a much lower impedance than the split-load phase splitter. If the value of resistor Rk is correctly chosen, then the total effective impedance of this first stage must equal twice the value of resistor Ra. Whereas, the split-load phase splitter's effective impedance equals rp + (mu + 2)Ra, a truly huge number. In short, capacitor C1's value will have to be somewhat greater than 33µF.

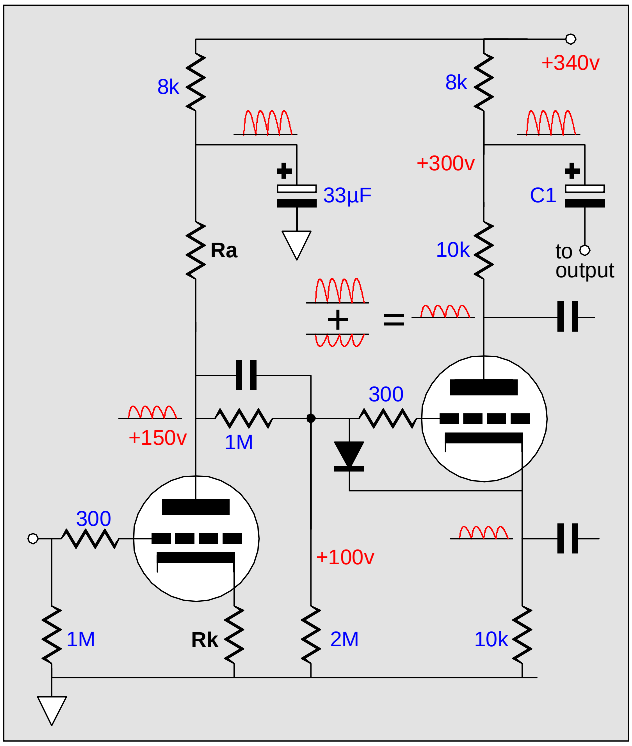

PS noise balancing front-end

When the connection is made at the cathode side, the previous circuit works perfectly, as both the first stage and the phase splitter share a common B+ connection. In other words, in this circuit it does not matter that the output devices are providing gain.

Split-Load and Class-AB

What if we are not willing to pay the price of high current draw and heat generation that Class-A demands? The alternative is what 99.9% of power amplifiers run in: Class-AB operation. (In fact, many (most?) advertised Class-A amplifiers actually run in Class-AB. If the amplifier's output stage does not at least draw half of the peak output current swing, it is not Class-A.) Class-AB provides a compromise between Class-A and Class-B. It is not as distorting as Class-B and not as inefficient as Class-A operation.

In fact, many (most?) advertised Class-A amplifiers actually run in Class-AB. If the amplifier's output stage does not draw at idle at least half of the peak output current swing, it is not Class-A.

Eliminating power supply noise is still possible within a Class-AB amplifier; it just requires a little more thought. The circuit below is a case in point. A bipolar power supply is used and the output is DC coupled. The bottom output device must see the same power supply noise at its input as it does at its connection to the –B connection.

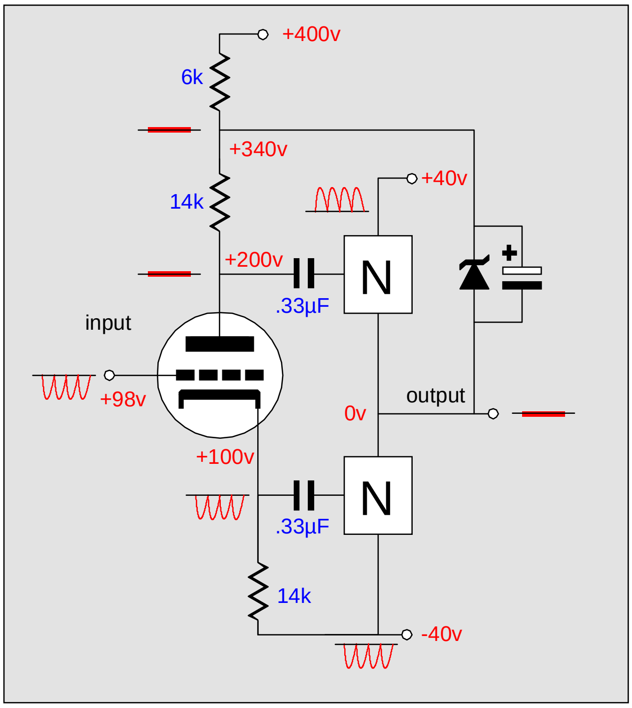

Introducing this noise at the split-load phase splitter input ensures that its cathode will relay the noise faithfully to the bottom device. Amazingly, the interjection of noise also ensures that the top device does not see the noise. How can this be so? The phase splitter's cathode resistor sees the same noise signal at both its ends, thus it never experiences any current fluctuations due to that noise signal, and thus the plate resistor cannot reproduce the noise signal. As it sees a constant current draw from the plate.

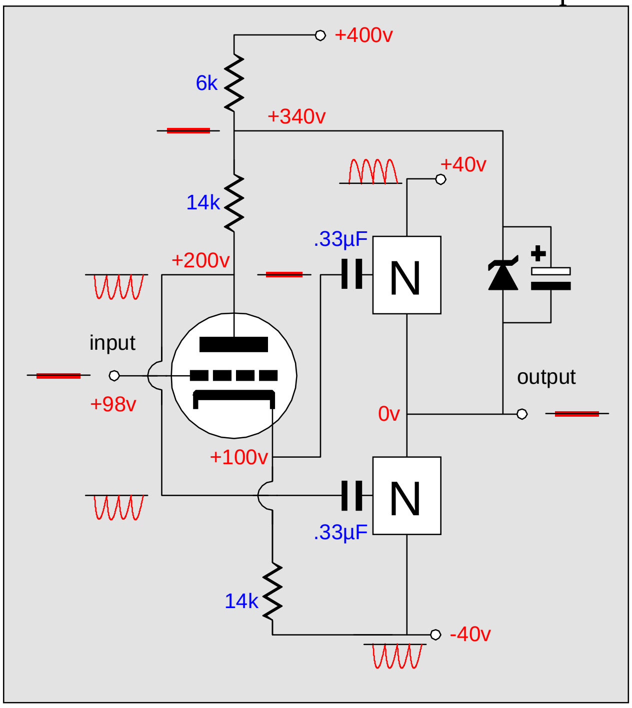

If we reverse the connections from the phase splitter, as in the circuit above, then a different strategy is required. The bottom device still needs to see the full negative power supply rail noise and the top device needs to see no noise. Both conditions can be satisfied by virtue of the bipolar power supply. If the phase splitter's cathode is noise free, the cathode resistor will see the full negative power supply noise signal, which will in turn create a current signal that the plate resistor will experience. This current signal Creates an in phase replica of the negative power supply noise at the plate, which is then relayed to the bottom output device via the coupling capacitor. (*This trick only works because the cathode resistor is not grounded.*)

If the direct connection layout is used instead, then the path to the negative power supply noise will be effectively severed. In the circuit below, a new path is created by connecting a power supply filter capacitor not output ground but to the negative power supply rail.

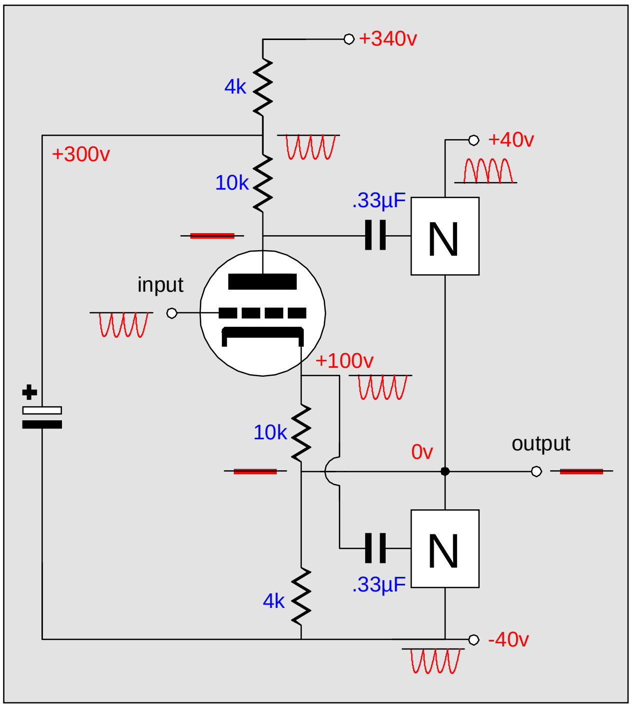

The split-load phase splitter is marked by an extremely high plate impedance due to the presence of the large unbypassed cathode resistor. Whatever signal appears at the top of the plate resistor will also appear at the bottom of the resistor. If this signal meets an inverted version of its self, both signals will cancel. Thus if we superimpose the negative power supply noise on the input signal, the bottom output device will see the its required noise signal and the top device will see the null that results from the input signal being inverted at the plate canceling the in phase noise signal.

Had the filter capacitor been connected to ground, the negative power supply noise would not be canceled at the plate and the output would polluted by it. The lesson here is that filter capacitor should not always be terminated into ground. (Some are not mentally flexible enough to grasp this and, unfortunately, they design most amplifiers.)

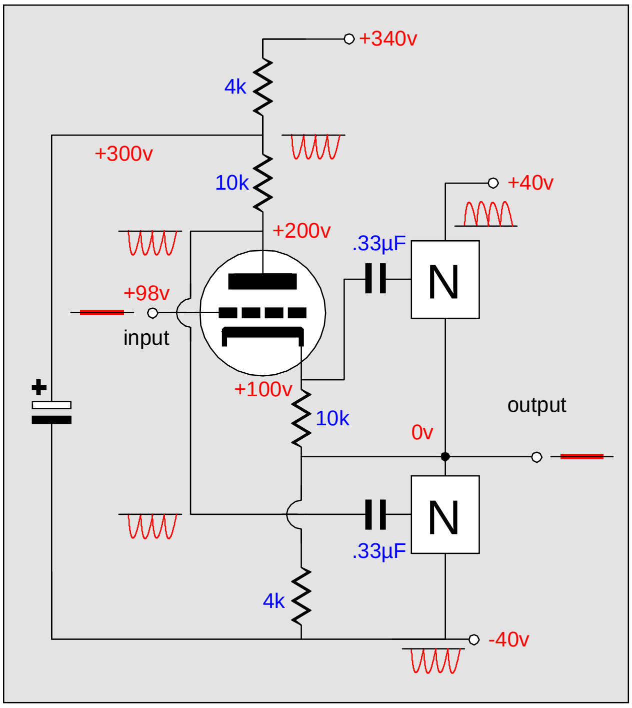

Reversing the connections to the phase splitter requires reversing our strategy. Now the phase splitter's cathode needs to see no noise and the plate needs to see the full negative power supply rail noise. The filter capacitor's connection to the negative power supply rail ensures the coupling of the rail's noise to the phase splitter's plate and, then via the coupling capacitor, to the bottom device's input.

(The 4k resistor that spans the output to the negative power supply rail only serves to prevent the phase splitter's current from being drawn through the output stage. It could be eliminated, if desired. It is, however, useful for allowing testing of the driver stage in isolation of the output stage.)

I can imagine that some readers are troubled by the task of superimposing the negative power supply rail noise on the phase splitter' input signal. Just how do you do that? Depending on the input circuit, several techniques are possible. For example, a common-cathode amplifier whose cathode resistor is connected to the negative power supply rail will find an easy path to passing the noise to its output, as the signal present on the cathode resistor will be relayed to the plate resistor.

Differential Unity Gain Circuit

With the addition of the negative power supply rail comes a new source of noise. Of course, a negative power supply is not strictly required; the long-tail phase splitter only requires a large valued cathode resistor. The negative power supply allows for a greater resistor value; but if the B+ voltage is sufficiently high, or if a constant current source used, then need for a negative power supply falls away. Or does it? If the output stage uses only a mono-polar power supply, then certainly the large valued resistor or current source will prove adequate.

On the other hand, if output stage uses a bipolar power supply, then the task of noise elimination becomes more difficult. The problem lies with the bottom device being referenced to the negative power supply rail, not to ground. Thus if the bottom device's input is grounded, its cathode or source or emitter will see the full negative power supply rail noise and amplify it at its output. On the other hand, if the bottom device's input is referenced to the negative rail, then both its input and its cathode or source or emitter will see the same noise signal. But no amplification of the noise ensues.

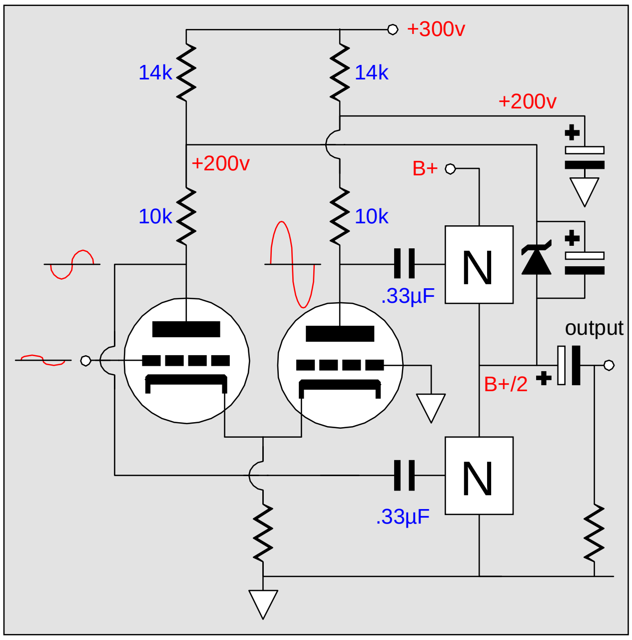

Diff-amp, plate-referenced, unity-gain output stage

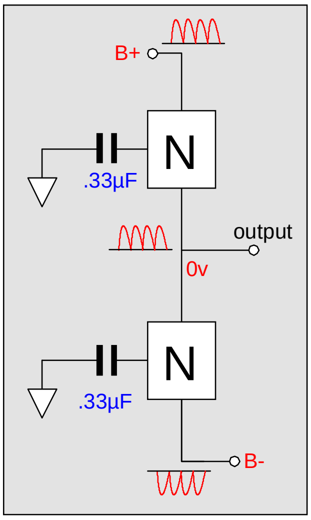

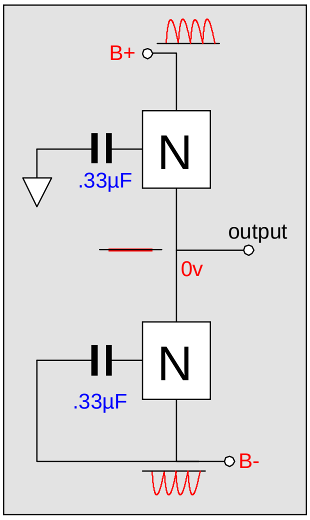

Now if the power supply rails were truly noise free, grounding the input to the bottom device would not result in any noise at the output, as the ground and the negative rail would both be noise free. So wouldn't our efforts be better spent in trying to cancel the power supply at its source rather than trying to cleverly cancel it through circuitry? If the amplifier is run in pure Class-A mode, the brute force approach has some merit. But hybrid power amplifier are seldom run in anything close to pure Class-A and pure tube OTL amplifier never are, in spite of advertising copy. What does mode of operation have to do with noise transmission? More than most tube gurus know it turns out.

Hybrid and OTL output stages usually run an idle current of 50 to 200 mA, which translates into .8 to 3.2 volts into an 8 ohm load. Beyond these voltages, only one output device delivers the current into the load. Thus any scheme that seeks to eliminate power supply noise at the output by having the noise currents through bottom output devices cancel ceases to work beyond this point. Fortunately, the music signal will tend to mask the noise contribution to the mix. Now, let's imagine a power supply that had no hum. Such a power supply would be greatly valued when no music was playing, but it would still give rise to a different type of "noise," one that would not show up at idle, but would manifest itself when play music.

Power Supply Sag

This "noise" is the collapsing and expanding of the rail voltage in response to the demands made on the power supply during the playing of music. The sag is caused by the greatly varying current demands the output stage makes on the power supply, which when multiplied against the effective series resistance of the power supply, create a varying power supply rail voltage. At full output a 100 watt amplifier draws a peak current of 5A, which results in a 5 volt sag in power supply voltage, if the power supply's ESR is 1 ohm; and 25 volts, if the power supply's ESR is 5 ohms. (The latter figure is more likely in an OTL amplifier.) The ESR is the sum of the diode and transformer winding resistance; and the higher the power supply voltage, the greater the ESR.

So even a perfectly quiet power supply can cause trouble. If the fluctuations in power supply voltage are not matched in the drive signal for the bottom device, gross distortions occur. Imagine a loud passage collapsing the negative power supply rail, but the DC value of the drive remains at the idle value. In this scenario, the cathode or source or emitter has effectively been made much more positive than the grid or gate or base, which in turn turns off the bottom device altogether. Yes feedback will work to correct the problem, but this means that the feedback must ease up on its correction of bandwidth limitations and general distortion cancellation.

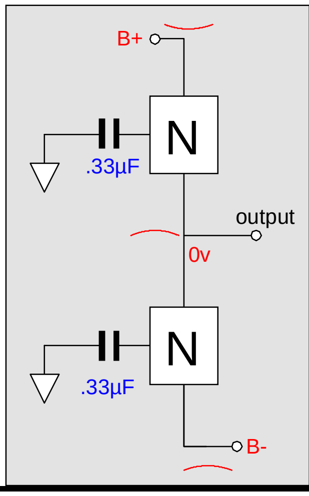

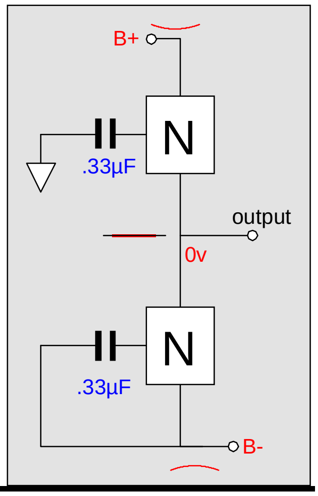

And if the amplifier does not use a global feedback loop, the sound will suffer. Unless, of course, we match the sagging negative rail voltage with a negative drive voltage. The circuit below shows the top device seeing the negative power supply rail noise reflected to its input, which will make it to the output. Is the answer to replace the cathode resistor with a constant current source?

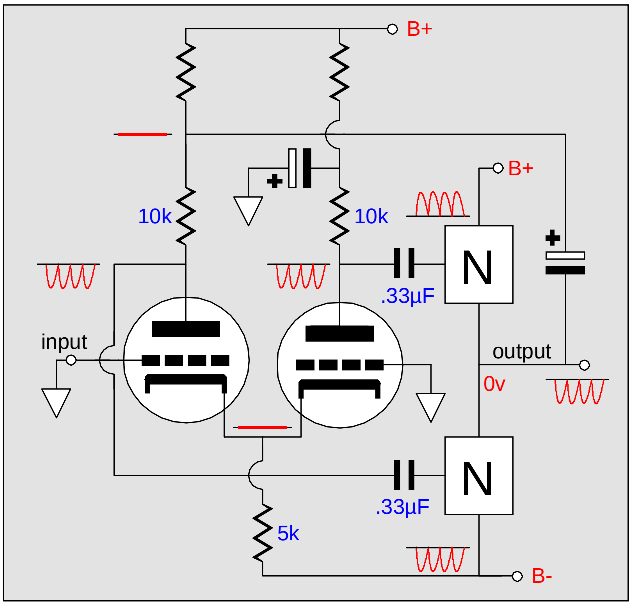

The circuit below shows the addition of the constant current source and the inversion of the noise at the output. What happened? The constant current source only shifted the source of the noise to the bottom device.

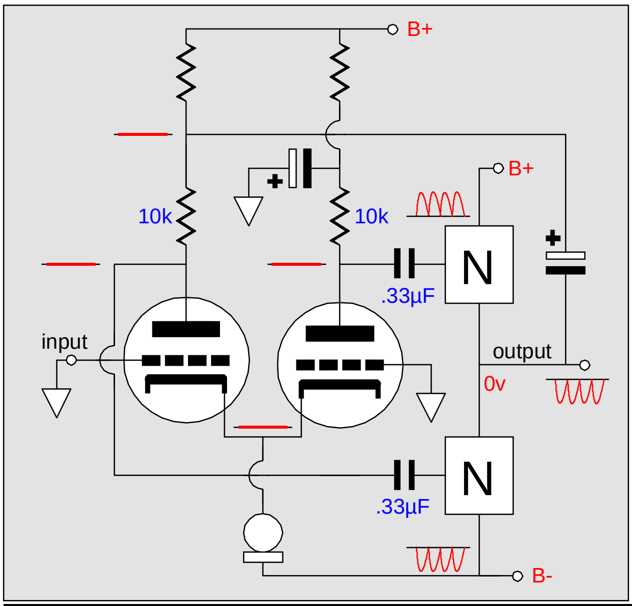

The solution is to shift the power supply filtering capacitor connection from ground to the output stage's positive rail. A small change in layout, a large reduction in noise. Now both the top and bottom device's see the their required noise referenced drive voltages. The bottom device sees the same noise signal as its cathode or source or emitter sees. Thus without a difference there cannot be any amplification. The top device sees the sum of both the positive and the negative rail noises, which cancel to null. Once again, nothing to amplify.

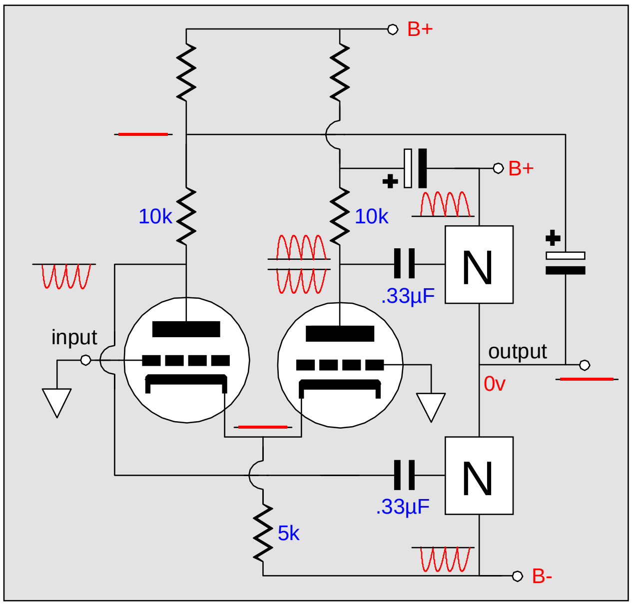

The cathode resistor's value must equal half the plate resistors value or the noise cancellation will not occur. Thus *the cathode resistor should not be replaced by a constant current source*. The same holds true for replacing the plate resistors with current sources. So the answer to the question, "Are current source a good idea in tube equipment?" is sometimes yes, sometimes no, and sometimes it makes no difference.

Look carefully at this last circuit and you will see that the circuit compensates for power supply rail sagging and expansion. Should the negative rail voltage creep upwards, the current flowing through the cathode resistor will decrease, which in turn will increase the plate voltages, as the voltage drop across the plate resistor will collapse.

In fact, the rail voltage fluctuations are a type of noise that is induced from dynamic operation of the amplifier. Unfortunately, most amplifier tests only evaluate the output at idle and under steady state output signal conditions, which does not reveal the dynamic failings of an amplifier. For example, the sagging of the power supply rail voltages under a heavy steady current draw will only initially causes a blocking distortion, as the bias voltages will eventually catch up with sagging by on the time constants of the coupling circuits. Constantly shifting in frequency and in volume, very little music contains a steady RMS signal. Thus if the amplifier is always seeing a sagging and expanding rail voltage and if not designed to compensate for this fact, it will always be faltering when faced with music.

Differential Gain Circuit

The output stage can be configured to provide gain. This is accomplished by bootstrapping the triode that drives the top output device. This allows the top output device to realize gain at the cost of a high output impedance.

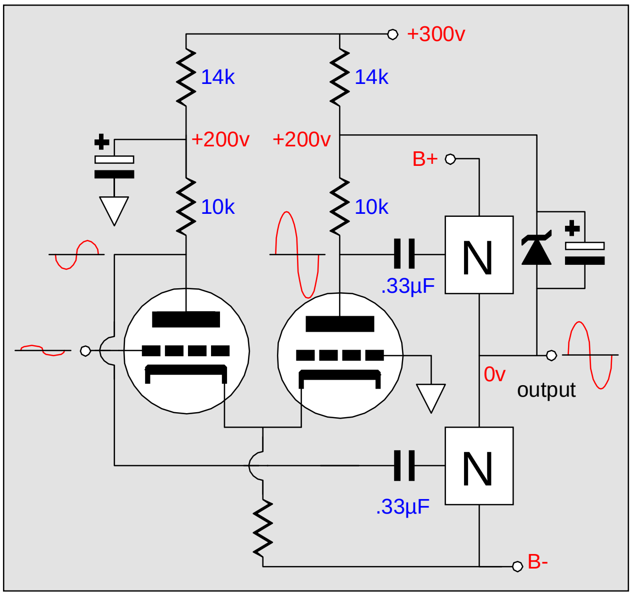

Eliminating the noise is a bit different in this circuit compared to the no-gain version. As shown below, the power supply noise grossly amplified by the top output device. Here the cathode resistor relays the negative rail noise to top device's input where it is greatly amplified.

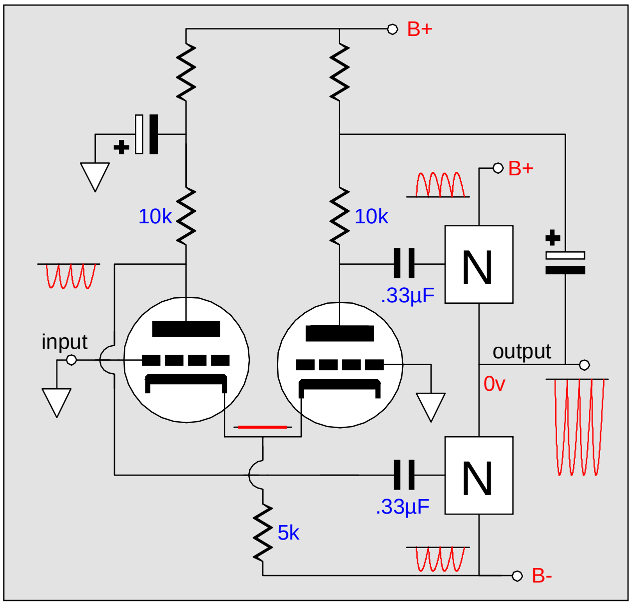

Replacing the cathode resistor with a constant current source, once again, does not eliminate the noise, but interestingly it does not invert the phase of the noise at the output.

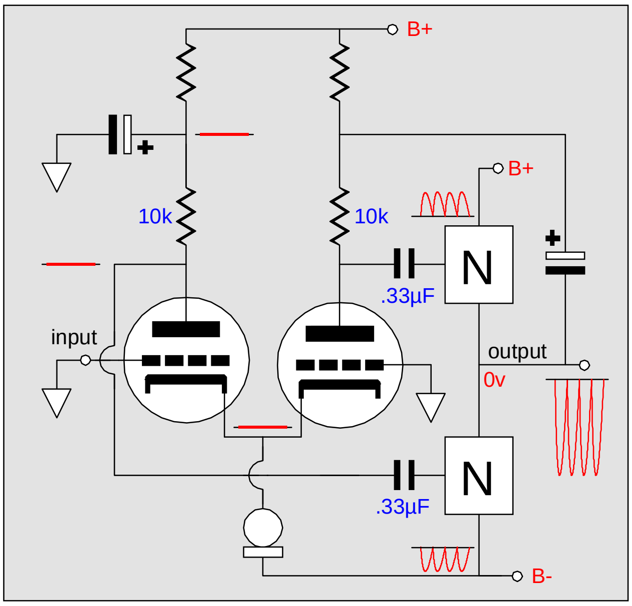

Keeping the current source and shifting the power supply filtering capacitor's termination from ground to the negative power supply rail (as shown below) is the trick to eliminating the noise. This is another example of how the whole circuit, the driver circuit and the output stage and the power supply, must be analyzed to fruitfully design a power amplifier.

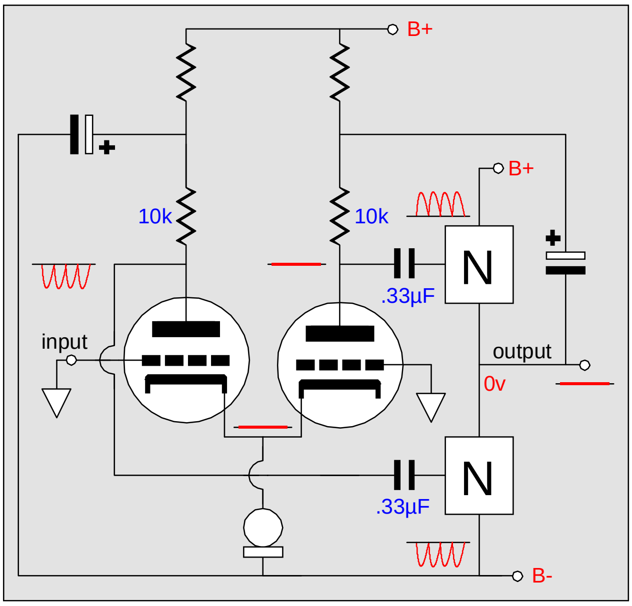

Correcting the Long Tail Imbalance

This last circuit holds the answer to a problem not yet mentioned: the inequality of gain from each leg of the long-tail phase splitter. In a nutshell, the first triode realizes more gain than the second triode, which requires using a larger valued plate resistor for the second triode.

Why is there an imbalance? The quick answer is that the common cathode resistor is not infinite in value. The longer answer is that long-tail phase splitter must present an equal drive voltage to each triode in order to develop an equal output gain. But as the first triode's grid sees all of the input signal and as the second grid only sees ground, how is this condition attainable? First of all, the input grid does not see all of the signal.

If we look at the first triode carefully we see that it works as both a cathode follower and a grounded-cathode amplifier. Of course, if the cathode followed the grid perfectly, there would be no grid-to-cathode voltage variation, and thus no variation in current conduction, and thus no extra gain at the plate. However, the load that the cathode works into is not just the common cathode resistor, but that resistor in parallel with output impedance of the second triode. In other words, we have a cathode follower that works into another cathode follower's output.

We begin by making the common cathode resistor infinitely large. In other words, if it is infinitely large, we can ignore the cathode resistor. As both triodes share the same rp , mu, and plate resistor values, we should expect that they would also share the same output impedances. And two equal impedance in series define a voltage divider that halves any signal. So if the first grid sees a +2 volt input signal, the first cathode will only be able to follow half way and thus see a +1 volt signal. Being grounded, the second grid sees 0 volts, but the second cathode sees the same +1 that the first cathode see, as they are connected. Therefore, the first triode effectively sees its grid move up 1 volt, and the second triode effectively sees it grid move down –1 volt.

So if the common cathode resistor were infinitely large in value, the gain would match (actually, they wouldn't, as no cathode follower has a gain of 1: an imbalance of 1/mu remains). If it is not, it decreases the voltage division and thus decreases the input signal the second triode receives. This explains where the name "long tail" came from; the larger (the longer) the common cathode resistor value, the less the imbalance. Of course, an infinitely large cathode resistor implies an infinite negative power supply voltage.

The constant current source frees us from the need to provide an infinite voltage. Independent of the voltage across it, it draws a steady current. Making a constant current source out of pentodes, FETs, MOSFETs, and transistors is easier than with a triodes. In this application, the triode's chief advantage, rp, becomes a liability, particularly when high current is needed, as high current triodes, such as the 6BX7 or 5687, have very low plate resistances. Still, by using a large valued cathode resistor, the triode's effective rp can be greatly increased. For example, a 6SN7's rp is 8k and its mu is 20, which when combined with a 10k cathode resistor yields an effective rp of 218k. Put mathematically,

$rp' = (mu + 1)Rk + rp.$

Feedback and Gain Imbalance

If a constant current source is not used, the imbalance can be partially corrected through the use of a feedback loop. This is the lazy engineer's answer to virtually any problem: noise, limited bandwidth, high output impedance, and distortion. Design the sloppiest, cheapest circuit you want, no worry, the feedback will take care of it. Imagine if a politician were to argue that every social ill had one remedy: more laws and more police, one cop in each business, school, playground, church, home, and bedroom, if need be. Sad to say, but it does not require much effort to imagine such a fool.

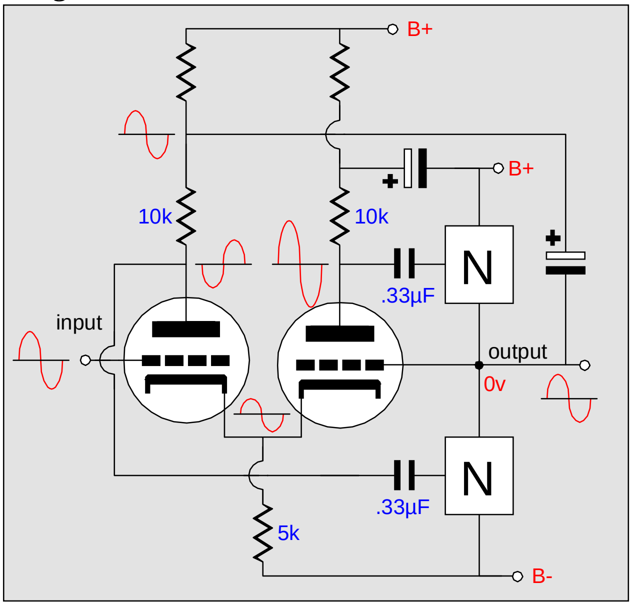

Here we have an interesting example of how a feedback loop can exacerbate a problem. In the usual arrangement of the long-tail phase splitter, the second grid goes to ground. If instead it were to connect to the output of the totem pole output stage, then a feedback loop is created. The total gain of the circuit falls bellow unity, the output impedance decreases, and the distortion lessens, but not as much we might expect. Because of the 100% feedback ratio, a much bigger input signal is needed to bring the output stage to full output. This means that rather than 4 volts of input signal, 40 volts of signal might be needed to drive a tube output stage to 36 volts of output signal.

In both cases, with feedback and without, the output devices see the same drive voltage. In the feedback case, however, the common cathode resistor sees not the 2 volts of signal, but 38 volts. This huge increase results in a huge imbalance in drive signals, which the feedback must work to eliminate.

The feedback must compensate for its use in this circuit, which means that the feedback cannot give its all to extending bandwidth or lowering noise. Remember having 40 dB of feedback is like having $40,000 in the bank: you can buy a BMW or a high-end sound system or a diamond necklace, but not all of them at once. So it is with feedback, it gets used up.

One solution is not to tie the second grid to the output and use global feedback that terminates into the first stage, usually a grounded-cathode amplifier's cathode. An alternative is to use the differential, gain realizing circuit instead. This topology demanded the use of a constant current source for power supply noise elimination and the current source would also serve to eliminate the unbalanced drive signals. The circuit below illustrates the similar drive signals, but what cannot be seen is the gain imbalance elimination.

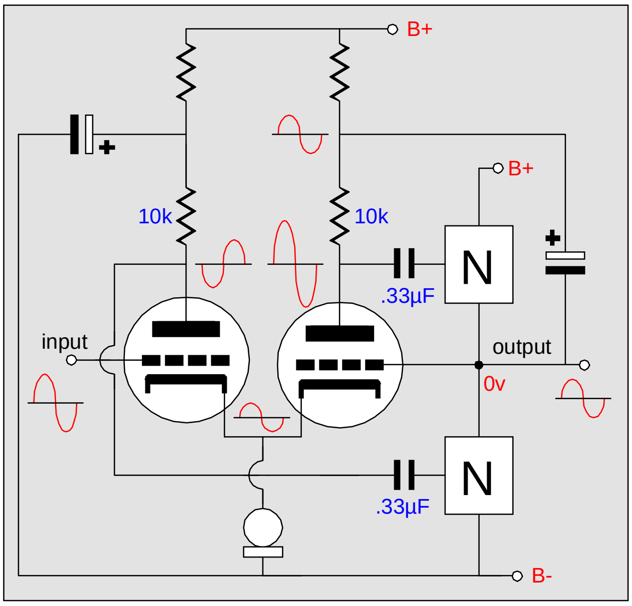

Transcendental OTL Remake

As was promised in the issue, we take another look at the Transcendental OTL amplifier. In light of what we have covered in both this and the last issue, we definitively see opportunities for improvement.

Some readers might be stunned by the hubris of the last sentence. Do I not realize that the Transcendental OTL is a *Stereophile* "Class A" amplifier? How dare anyone presume to improve upon perfection?

(I am reminded of when I was in elementary school my sixth grade teacher told us the brain teaser about the man who hears a noise outside is home and fears it might be a bear and runs outdoors with his rifle. He follows the tracks and proceeds exactly one mile south, then one mile east, and then one north, where he finds himself back home and sees the bear clawing on his front door. "What color is the bear?" was the question. The answer is white, which a few of us gave. The teacher went on to explain that from *only* one point on the earth's surface could one return home after following the man's course and that point was the North Pole. I disagreed by claiming that an infinite number of points existed that met that condition and that these points define a circle roughly 4 miles in circumference. I told him that any point that fell 1 + 1/(2pi) miles north of the South Pole qualified, as one mile south, then one mile east, and then one mile north from any of these points would bring the man back to his starting point. The teacher was not amused; the rest of the class was not amused. Obviously, I did not understand that the brain teaser was perfect.)

North Pole **South Pole**

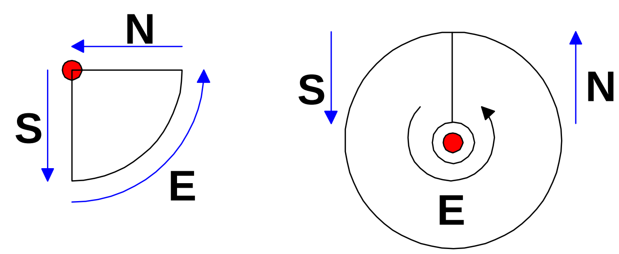

In the world of hi-fi, many believe the notion that all that can be done in terms of circuit topology has been done; and thus, the widely held belief that only improvements in part quality can lead to improvements in sound quality. So while the Transcendental OTL is highly regarded, it certainly is not perfect; change is possible. Of course, if a recommended set of changes are great enough, the original topology disappears. So let's try to retain as much as possible of the original amplifier, while bearing in mind what we have covered so far.

From the last issue, we saw that the patented drive equalization technique, the bootstrapping of one of the cathode follower's B+, used in the Transcendental OTL was a bust. But we have seen how to balance the drive signals and configure the long-tail phase splitter for power supply noise canceling. Applying these techniques will require an additional adjustment, however. The examples of the long-tail phase splitter based output stages that we have covered have all assumed that the first grid does not see any power supply noise. If the grid does see any, this noise will be amplified along with the signal. We know that no mater how carefully we filter the power supply, some power supply will be superimposed on the input signal to the output stage. So what we must do is configure the circuit so that the noise is largely rejected.

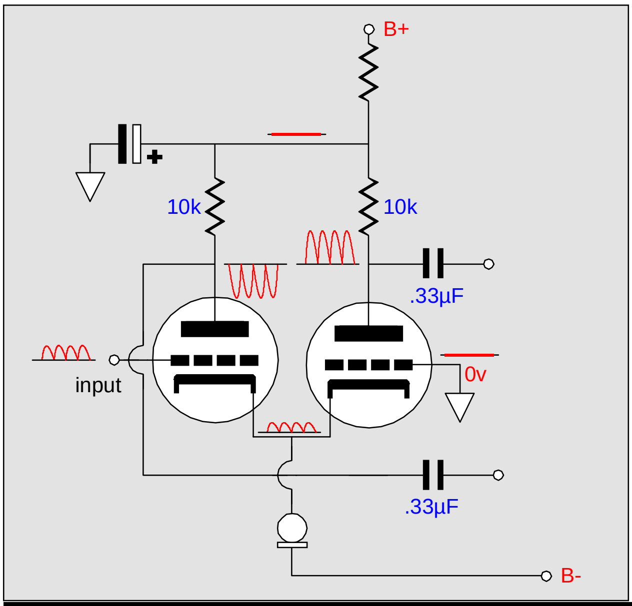

Here again is an example of how by purposely interjecting noise into a circuit we can lower the noise at the output. Although seldom mentioned, the long-tail phase splitter, much like the differential amplifier, contains a good CMRR mechanism. If both grids see the same signal, the common signal is barely amplified compared to the same signal applied to only one grid. So what is needed (in order to lower the noise at the output) is to ensure that the second grid see the same noise that the first stage sees.

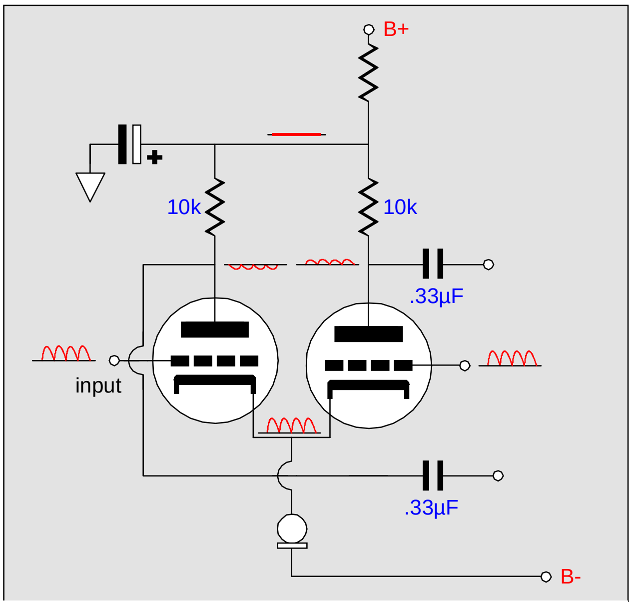

How do we balance the noise at both grids in this amplifier? The 12AX7 based first stage functions as a voltage divider to the power supply noise at the top of its plate resistor. In this instance, the division equals about 30%. In other words, 10 mV of noise reduces to 3 mV at the 12AX7's plate. The same division can be achieved by using a 70k and a 30k resistors in series. The problem with using resistors is that they divide AC and DC voltages equally, whereas the triode and plate resistor do not, the DC division being closer to 70%.

Remember we need to ensure that phase splitter's both grids see the same DC voltage or the long-tail phase splitter will not work well if at all. The solution is found in using two capacitors in series and a 1M biasing resistor. The resistor forces the capacitor junction to the same voltage as the first grid sees.

In the circuit below, the bottom capacitor equals .1 µF and the top capacitor equals .047µF. If the values seem reversed, remember that the reactance of a capacitor is inversely related to its capacitance. The higher the capacitance, the lower the reactance. Of course, the exact values required will probably differ, but these will be close. Some fiddling is always needed when canceling noise.

What if the first stage had used a pentode or a cascode amplifier and the voltage division equal virtually zero, i.e. all the power supply noise is present at the plate? In this case, we dispense with the bottom capacitor altogether.

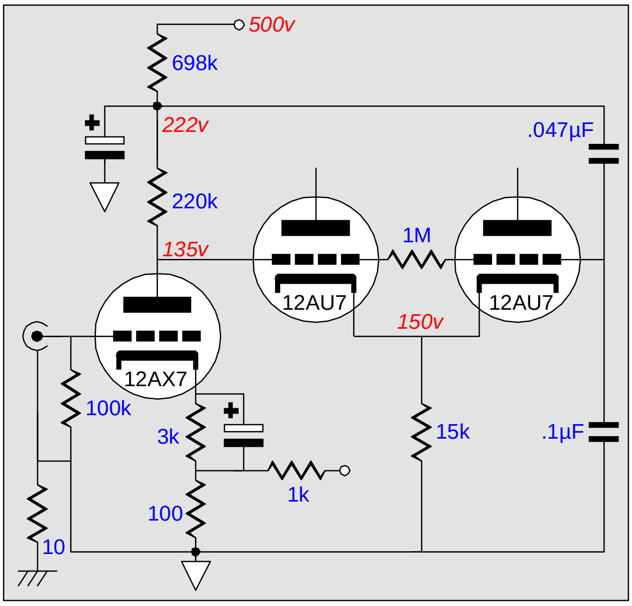

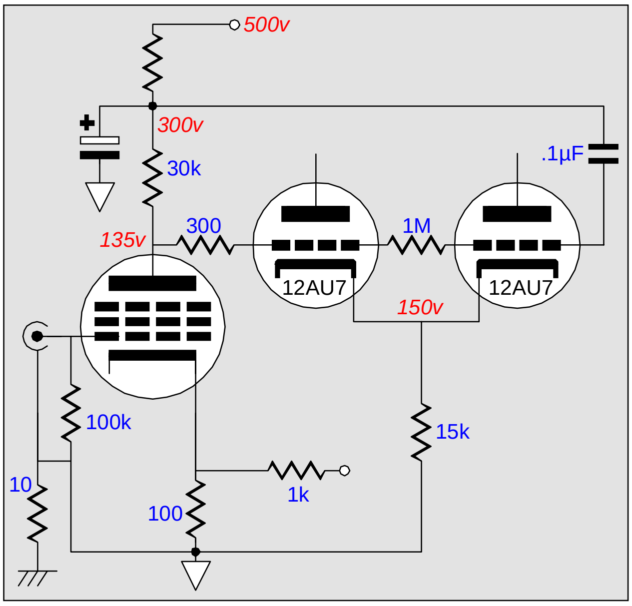

The observant reader will have noticed that in the last schematic that the plate resistor's value is much lower than in the original circuit. One problem with the stock Transcendental OTL is the ridiculously low plate current for the 12AX7 based input stage: 0.4 mA. Why was this trickle current chosen? The answer is that the designer hoped to develop all gain he could out of this stage. Higher current would have meant a smaller valued plate resistor and thus less gain. Why was the extra gain desired? Probably to lower some of the distortion and to extend the smaller bandwidth that resulted from using so little current. Huh? Much like a dog chasing his tale, we create distortion in the attempt to lower distortion. Starved of current, a triode produces more distortion, as it must operate in the bottom, gooey region of its plate curves. Additionally, the wimpy current is insufficient to charge the following stage's input capacitance (the Miller effect capacitance) quickly enough to ensure a wide bandwidth. Our solid-state brothers figured this out long ago. If an amplifier cannot charge its internal capacitances quickly enough, the feedback is going to send the amplifier into a slew-limiting distortion, e.g. SID.

What's the solution? Replacing the 12AX7 and 12AU7 with different tubes and increasing all the idle currents are a good start. Tube myopia is a debilitating disease that inflects more audio designers than can be imagined. (I remember John Atwood and myself promoting the 5687, 5965, 6072, and 7062 back in the 80s and meeting blank stares. I remember being told, "Don't you understand, only the 12AU7 and 12AU7 can be used in audio equipment; the tubes which you like best are meant for industrial uses and cannot be used in audio equipment." Today, I hear, "Only the 300B or 2A3 can be used in a single-ended amplifier.")

A better choice might be 6N1Ps or 7062s throughout. These triodes could easily handle sufficient current to make the amplifier sing. But I am also troubled by the first stage's power supply filtering: just a large valued resistor and a large valued (22µF) capacitor.

Improving this situation can be easily accomplished. A shunt regulator can be made out of the input tube's unused triode. Or a constant current draw stage can be made by adding a cathode follower to the input tube's output. For as long as input stage and the cathode follower shared the same valued load resistor, their combined current draw will always equals a constant, as the cathode follower's current conduction is in anti-phase to the grounded-cathode amplifier.

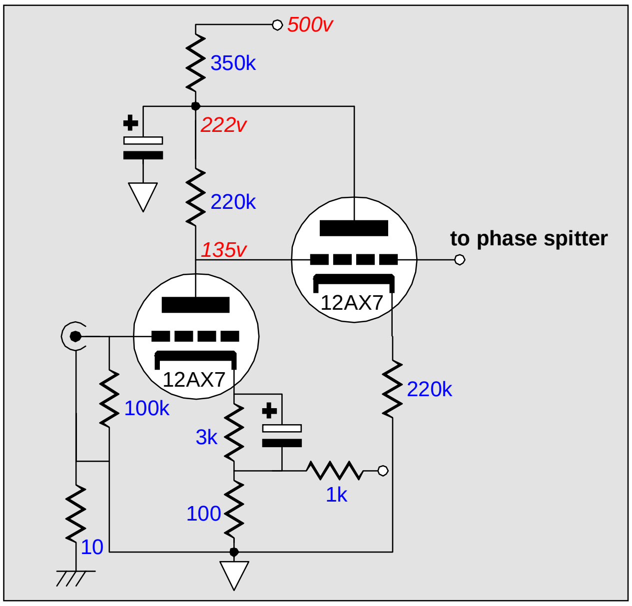

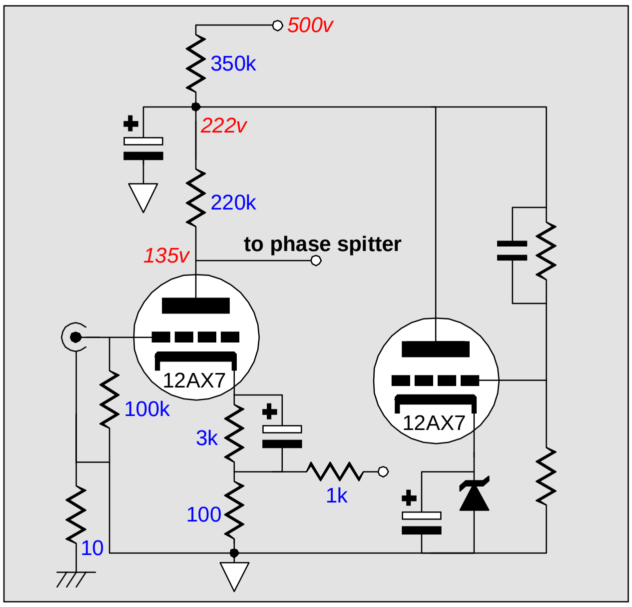

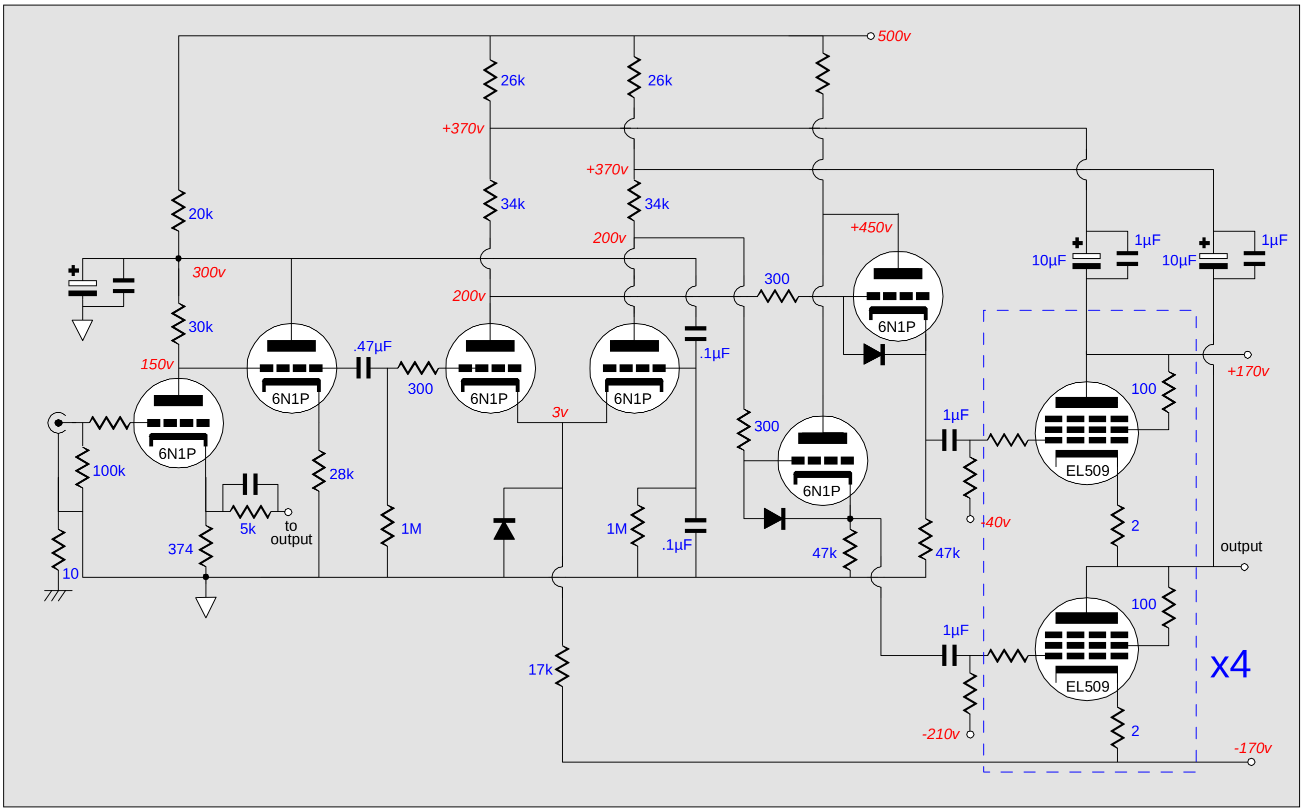

Click on image to see enlargement

The circuit above shows the enhancements made to the amplifier. As I now look at it, I see that the 5965 might be a substantially better choice for the long-tail phase splitter, as it would provide both more gain and a larger voltage swing. But then this remake was not meant to definitive (as I wanted to retain as much of the original design as possible).

Which of the above modifications yields the most results? Balancing the drive signals comes first and is followed by the noise canceling techniques; and although it is last, using a different input tube is needed to run a higher idle current and would have to be implemented.

Conclusions

I remember seeing a circuit that purposely added hum and limited the bandwidth of solid-state amplifiers in order to reproduce the golden sound of tubes! That tube circuits often hum is as unfortunate as some great wines often smelling of sulfur dioxide, but only a fool imagines that sulfates make a grand cru.

Rarely is any topology fully optimized, since few designers consider noise reduction possible without the aid of heavy feedback or without brute force techniques, such as chokes or regulators. So how can we fairly judge a topology unless it is completely and competently designed? Sadly, we can't.

Do not forget that the techniques outlined here can be applied to a wholly solid-state amplifier with a few modifications. So is there anymore to be squeezed from this topic? I am sure that there is, but we will wait to see what the e-mail brings before returning to it.

(I had hope to include an explanation on how to design a 4 watt OTL amplifier in this article, but it will have to wait until next time.)

//JRB

Reference

"Analyses of Drivers for Single-Ended Push-Pull Stage," Hiroshi Amemiya, *IRE Transactions-Audio*, Sept-Oct 1955

