

The Sweet Spot

Nelson Pass 10/10/09

Introduction

Every audiophile understands the concept of the “sweet spot”, that happy balance of conditions that elicits the best possible sound. It can be a particular listening spot, the best positioning of loudspeakers, or the fortuitous combination of components which complement each other perfectly. It should be no surprise that the desire for the best performance takes the search for the sweet spot into the interiors of the components themselves.

This article will concern itself with finding the sweet spot for each gain device in audio amplifiers. It is a commonly held belief in audio that the best amplifiers are composed of one or more active gain stages, each made as intrinsically linear as possible before negative feedback is applied to further improve performance.

Taken to the extreme, linear stage = *good* / negative feedback = *bad*. The latter half of that statement is the subject of controversy, but there is no argument about linear stages being good.

Achieving intrinsically low distortion in active devices (tubes or transistors) is usually not easy. We can use good parts, see to it that they have adequate voltage and current available, and use them in topologies that play to their strengths. These are the basics, and this is what you generally see when you examine schematics of popular amplifiers.

However each individual device has a sweet spot where its performance is the best, often dramatically better, and finding that sweet spot is a powerful technique for maximizing the performance.

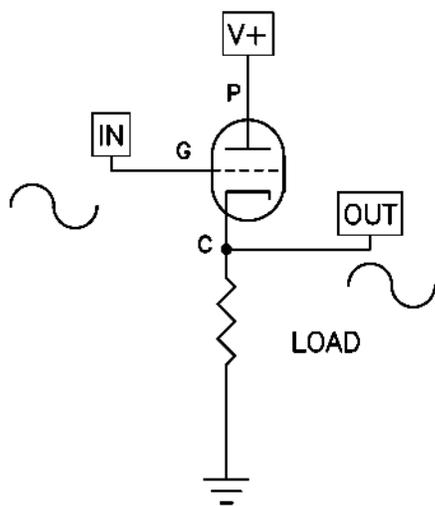
I am writing about this because over the years of communicating with DIYers and many professionals I have discovered that this is a poorly understood concept, and I have seen many circuits which have failed to take advantage of it.

For every gain device, there is a number which characterizes the amount of gain. When we talk about distortion we are talking about the alteration of the amount of gain in the device as the conditions change. The gain of a tube or transistor is altered primarily by changes in voltage across the device, current through the device and the temperature of the device. Changes in any of these conditions changes the gain, and all gain changes make the distortion. More precisely, they *are* the distortion.

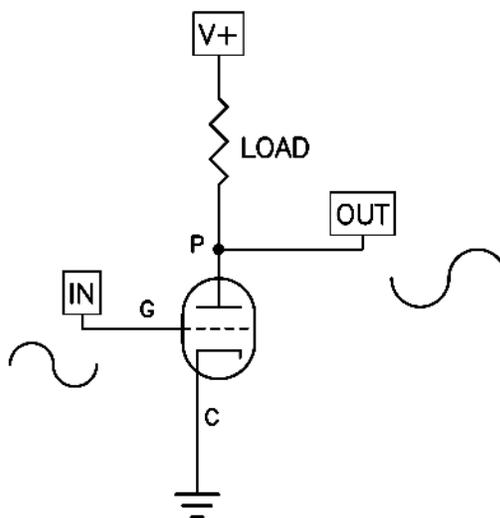
Of course the concept of a sweet spot depends on having an idea what of what constitutes the best performance. It could be that you want the lowest measured distortion, a particular mix or phase of harmonics in the distortion waveform, the greatest efficiency, greatest power or simply the best subjective experience when you listen to it. The sweet spot is whatever you want – after all, **you** are the designer.

I Brake For Tutorials

If you don't mind a little basics, let's look at examples of gain devices in generic circuits, starting with a triode:



COMMON PLATE MODE
(CURRENT GAIN ONLY)



COMMON CATHODE MODE
(VOLTAGE + CURRENT GAIN)

TRIODE SINGLE-ENDED CLASS A

This illustration shows simplified examples of the two most common of gain stages using any kind of gain device, and the device pins are the Grid, Cathode, and Plate. In both cases the input signal appears at the Grid and variation of voltage between the Grid and Cathode causes a change in the current flow from the Plate to the Cathode. (All the examples you will see are simplified, and omit DC bias that may be present on the pins except for the V+ supply).

In Common-Cathode mode, the tube develops both current and voltage gain, and the output is taken off the plate, inverted in phase and usually increased in amplitude.

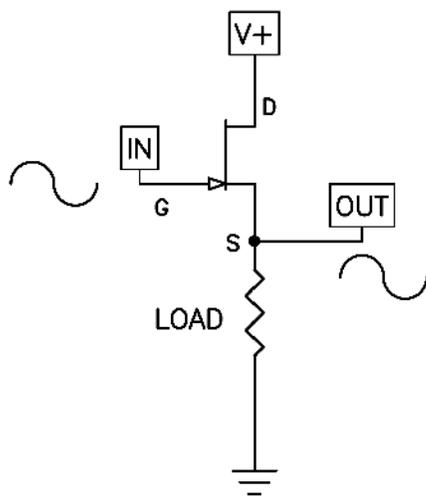
In Common-Plate mode, the tube seeks merely to have the Cathode voltage follow the Grid voltage, so the output voltage is ideally the same as the input, but with some current gain. That is why it's called a follower.

For both modes we would normally look for the sweet spot, the conditions giving the best performance, by varying the values of the supply $V+$, the value of the load resistance, and the bias current through the device. Inevitably there will be some combination of these which is the best. Sometimes one of the circuit conditions is fixed – often it is the load value, in which case we vary the other conditions – the supply voltage and the bias.

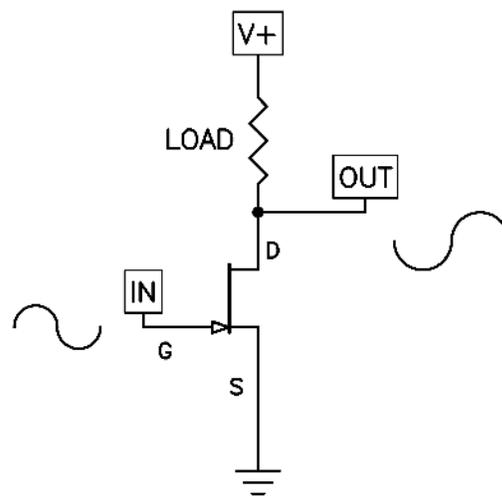
In all of the examples we will look at later, I have chosen to simply look at Common-Plate (or Common-Drain / Common-Collector) performance, but the same ideas apply to Common-Cathode (or Source or Emitter) operation.

Actually, Common-Drain circuits would seem to offer the least opportunity for improvement - as followers enjoying 100% degenerative feedback, they already have a much lower distortion character than Common-Drain circuits. But we will see that even they can be improved, often dramatically, if you find the sweet spot.

Here's the example of the JFET counterparts to the tube circuits seen above:



COMMON DRAIN MODE
(CURRENT GAIN ONLY)

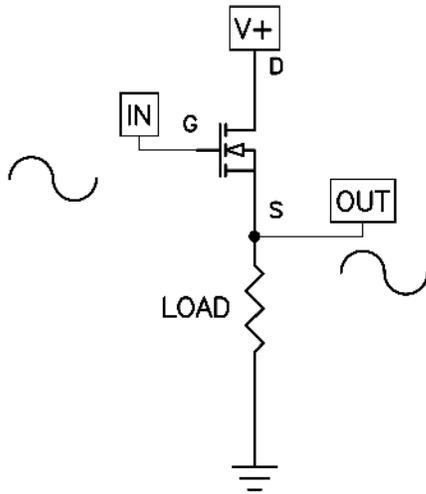


COMMON SOURCE MODE
(VOLTAGE + CURRENT GAIN)

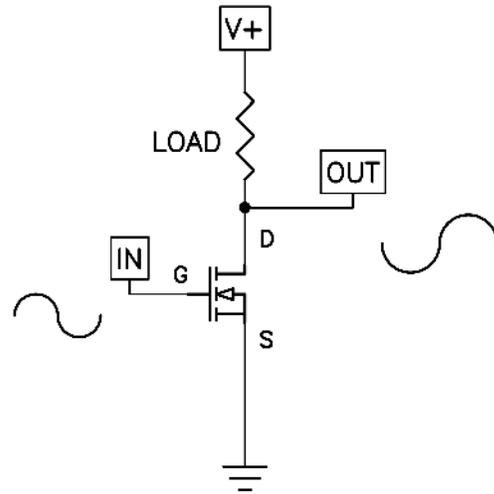
JFET SINGLE-ENDED CLASS A

Here the pins are Gate, Source, and Drain. In the Common-Drain follower, the input voltage pretty much follows the output voltage, and in the Common-Source amplifier, the output is inverted and has both voltage and current gain available.

Here's the MOSFET example, and it has the same pin designations as the JFET.



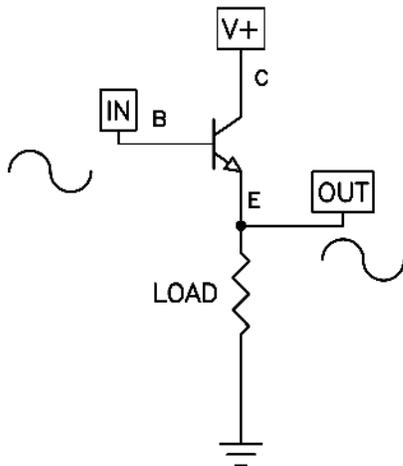
COMMON DRAIN MODE
(CURRENT GAIN ONLY)



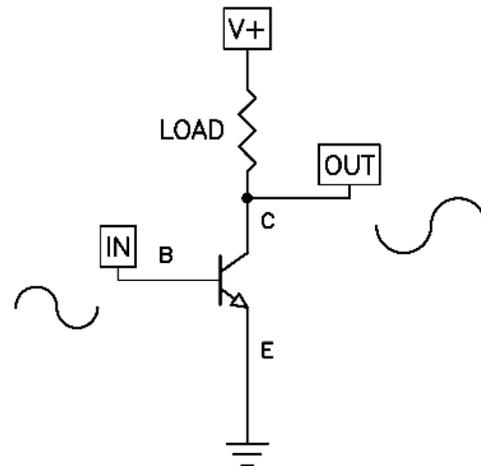
COMMON SOURCE MODE
(VOLTAGE + CURRENT GAIN)

MOSFET SINGLE-ENDED CLASS A

And last, here is the example of a bipolar transistor, where the pins are Base, Emitter and Collector:



COMMON COLLECTOR MODE
(FOLLOWER)



COMMON EMITTER MODE
(VOLTAGE + CURRENT)

BIPOLAR SINGLE-ENDED CLASS A

The choice of a particular type of gain device - tube or transistor - is often arbitrary, but ideally it plays to the particular strength of the device for a given application. In principle you can substitute one type for another, but as a practical matter, there is usually a good reason for the selection of one over the others.

It's doesn't matter in terms of the idea of finding the sweet spot – the terrain changes at bit, but the game remains the same.

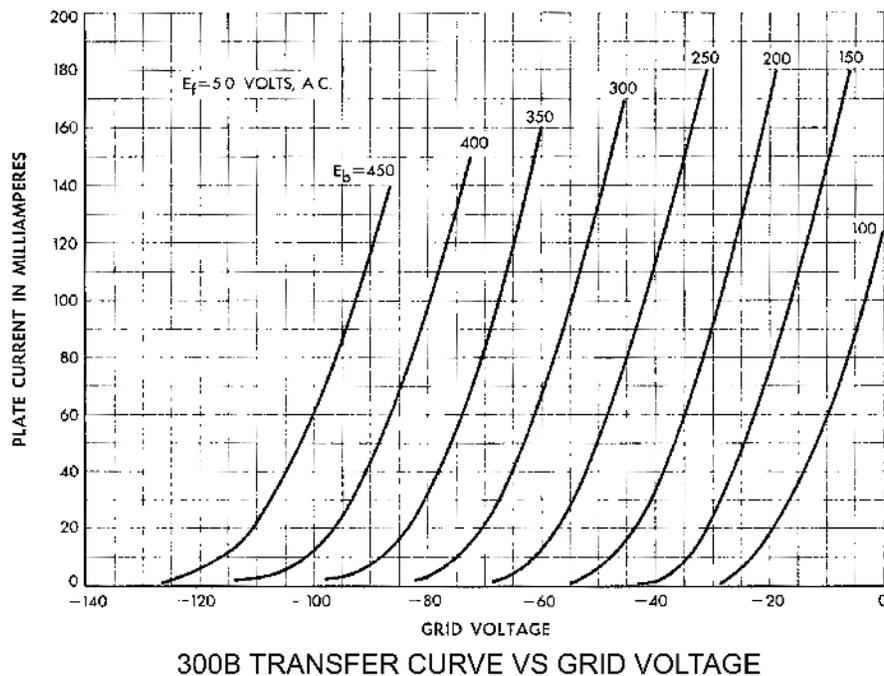
Let's Get to the Sweet Spot

For any of the above circuit types and devices, if you start with reasonable textbook values for the circuit, you will get reasonable performance. If you start to play with the values a little bit, you will find that the performance changes for better or worse. Often you will find a combination of voltage and resistor values which give a lot better than the generic performance. That's the sweet spot.

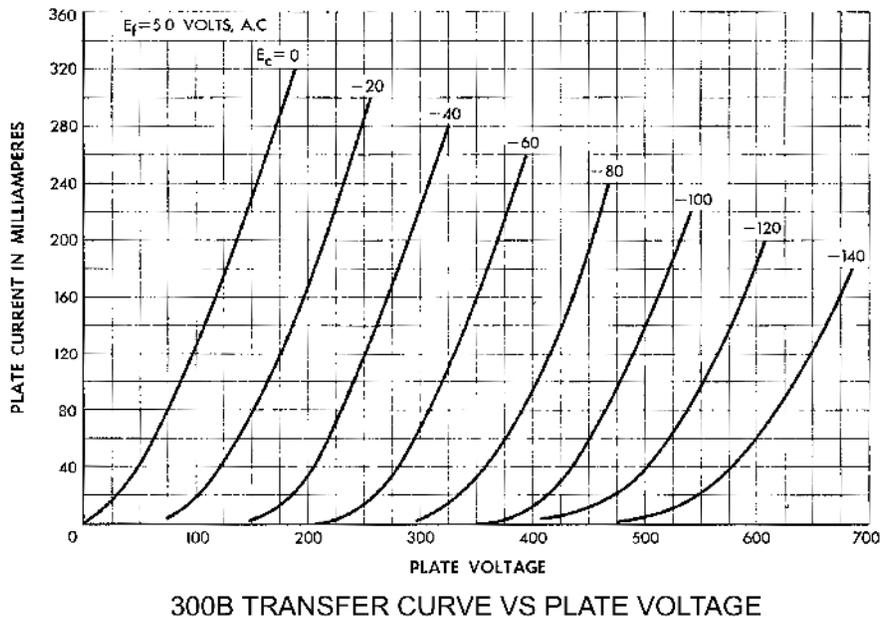
If you have a distortion analyzer, you could simply run through the range of combination of values and select the result you like best. If you simply want good measurements, you might be able to stop there. If you are looking for better subjective performance, you might find this a good place to start your listening.

What creates the existence of the sweet spot in gain devices? To see where this comes from, we will look at the characteristics of some tubes and transistors.

Let's start with a simple big triode, the 300B. This curve shows the Plate to Cathode current as a function of Grid voltage, with each curve representing a fixed value of plate voltage.



Looking at this curve, we see that the lines are curved – the more current, the faster they rise. Below is a different view of the same character, but with the current plotted as a function of the Plate voltage with each line representing a fixed value for Grid voltage.



These lines are also curved with a shape very similar to the previous graph. A perfectly linear tube will have these lines in both graphs perfectly straight and equally spaced. Any deviation from that would be distortion. Looking at the curved lines, we can see that there is quite a bit of distortion.

For a given Plate voltage, the current increases exponentially with the Grid voltage, in what is known as a square law characteristic, resulting in second harmonic distortion. For a given Grid voltage, the current increases exponentially with the Plate voltage, also a square law and also creating second harmonic distortion.

In this curve the Y axis is still the plate current, but the X axis is now the Plate voltage, and each line represents the Grid voltage. It looks familiar – just as the plate current is an exponential function of Plate voltage, it is also an exponential function of the Grid voltage. This is a particularly important observation with respect to finding the sweet spot.

So Here's the Deal:

Ordinarily in a tube amplifying an AC signal, a positive change in Grid voltage causes an *increase* in Plate current and is accompanied by a *decrease* in Plate

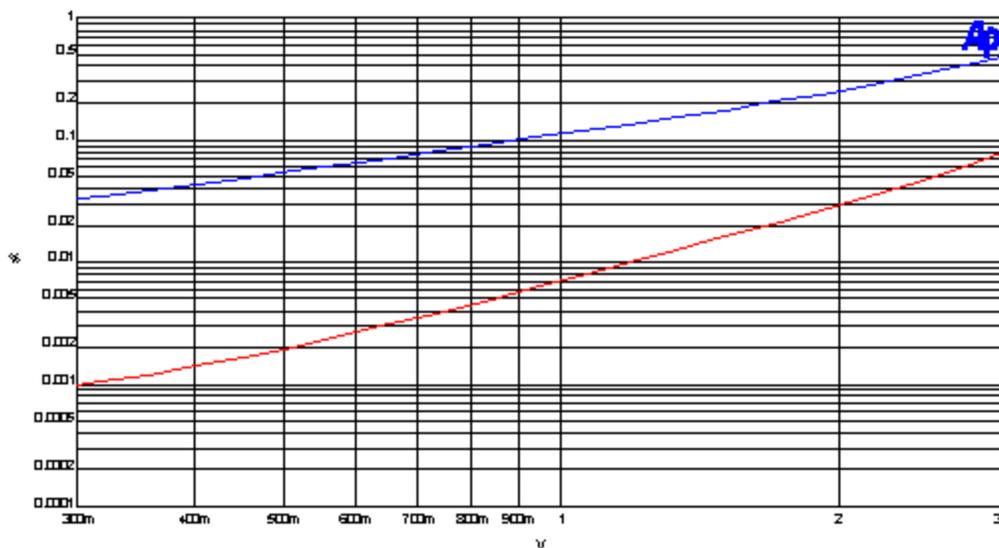
voltage. The gain due to Plate current *increases*, and the gain due to Plate voltage *decreases*. The opposite happens when the Grid voltage goes negative.

These two gain variations tend to cancel each other, resulting in more constant gain and lower distortion. If you choose your conditions so that the simultaneous gain increases and decreases are equal, you can lower the distortion a lot.

This technique is often referred to as *load-line cancellation*. It is called that because the range of the device's operation sits on a line in the transfer curve, and the particular shape and position of that line results in minimum distortion.

By some criteria, this would be the sweet spot. At this point you will probably find that your second harmonic has largely disappeared and you are left with some third harmonic. This is because you can't completely cancel two square law distortions without leaving a cubic trace – the third harmonic.

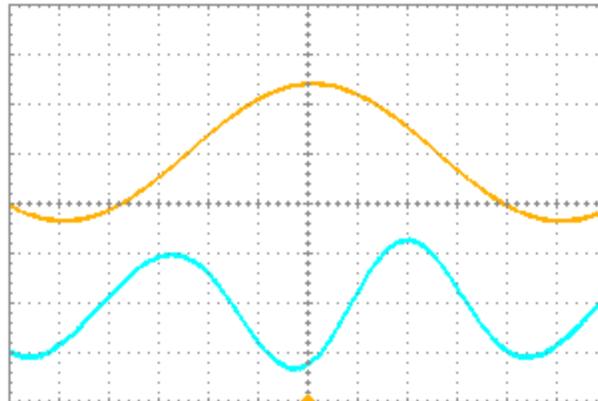
So how much distortion reduction are we talking about? On a clear day, we can think in terms of a 90% reduction, as seen in the following curve, where a triode follower is operated at a high plate voltage giving the upper blue distortion curve. Without changing the load or the current, I simply reduced the Plate voltage until the gain dependence on Plate voltage matched the gain dependence on Plate current (as determined by lowest distortion at 1 volt) and ran the red curve.



TRIODE DISTORTION

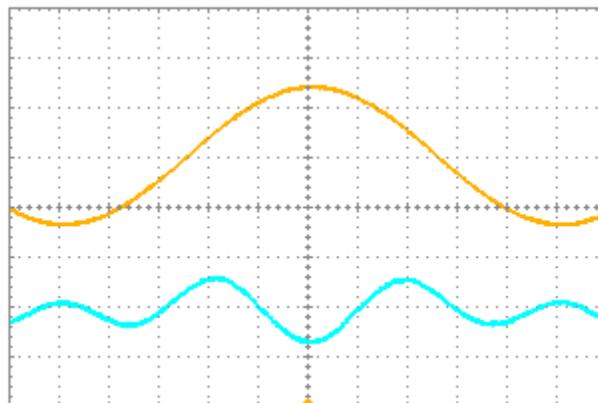
I think we can agree that this is a significant reduction, roughly equivalent to what we would see with 20 dB of negative feedback, except that we have not attached a negative feedback loop and the output gain and impedance of the circuit remain about the same.

Because of the square law characteristics of the devices, the primary distortion is second harmonic in content, and when we cancel two second harmonics we find that the remainder is a third harmonic. The original second harmonic looks like this on an oscilloscope viewing the output of a distortion analyzer:



SECOND HARMONIC DISTORTION

After cancellation, the harmonic content looks like this:



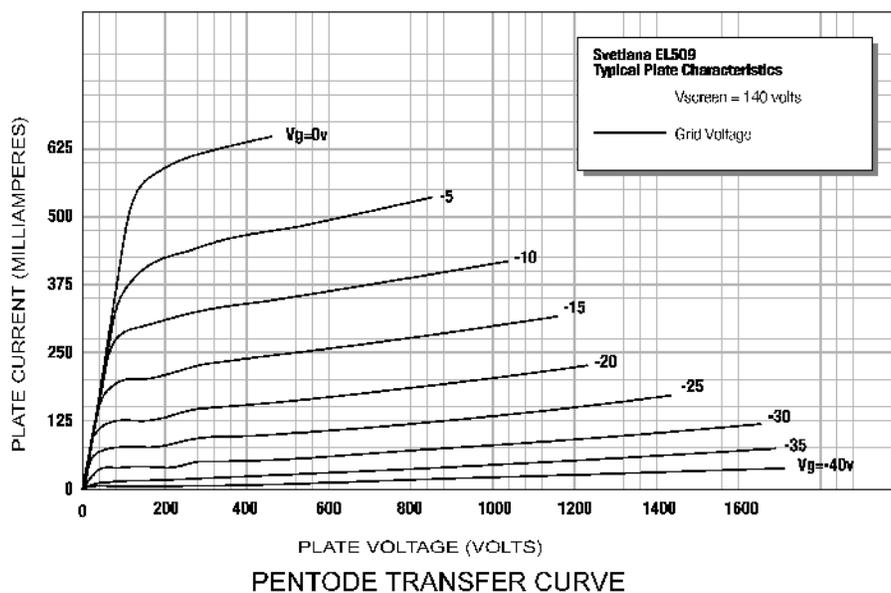
3RD HARMONIC REMAINING AFTER
LOAD LINE CANCELLATION

If your criteria instead is a particular harmonic content and/or harmonic phase, a deviation from this point in either direction will give you second harmonic back with its phase depending on which way you went, and you can tweak this in relation to the amplitude of the third harmonic if you like.

There is a reason why triodes are popular in single-ended applications – their gain dependence on Plate voltage is quite strong and makes for easy load-line cancellation.

Load line cancellation is available in the other types of gain devices, including pentodes, JFETs, MOSFETs and bipolar transistors, but they take more subtle advantage of load-line cancellation. In many cases it is easier to find the sweet spot with push-pull operation. We will see real world examples of both single-ended and push-pull later, but first we will consider the characteristics of the other devices.

This technique also applies to pentode tubes, although it will not be quite the same. The transfer curve of the pentode is much the same with regard to grid voltage, but the current variation due to Plate voltage is different. Here is the transfer curve of a pentode:

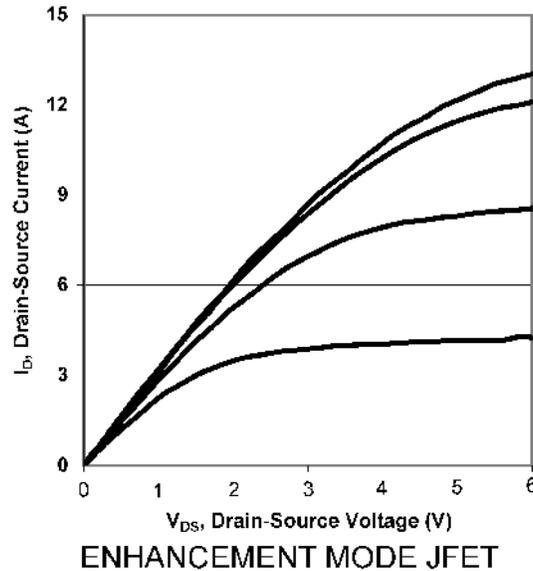


Here we see a family of curves which indicate that the gain (transconductance) increases exponentially with Plate current, but has less variation with respect to Plate voltage as compared to the Triode. You can still vary the voltage, current and load parameters to find the sweet spot, but it will be in a different place.

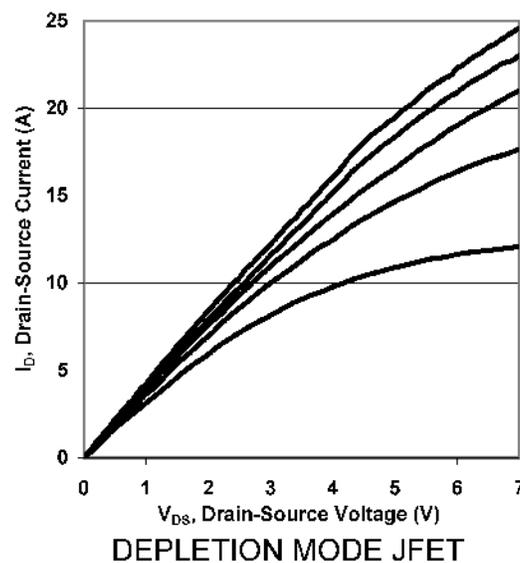
JFETs have the exponential current vs Gate voltage but enjoy a wider variety of transfer characters situated between triode and pentode and are generally well suited to exploit the concept of a sweet spot in single-ended Class A circuits.

JFETs come in both enhancement mode and depletion mode, the primary difference being the DC voltage required on the Gate to turn them on. Depending on the dimensions of the Gate structure on the chip can exhibit pentode or triode characteristics, or an in-between “mixed mode”.

Following is the curve of an enhancement mode power JFET:



And here is the curve of a depletion mode power JFET:



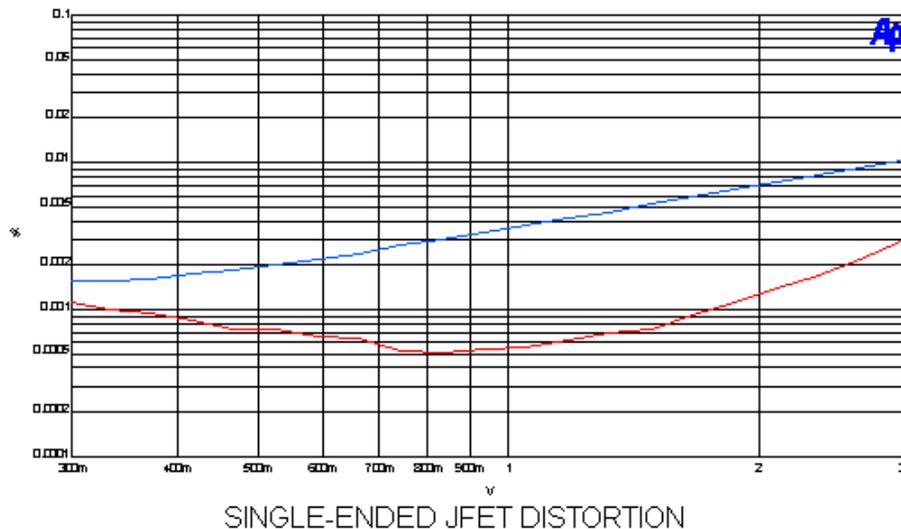
You will note that these aren't the tiny little signal JFETs that everyone is accustomed to. These are examples of the new generation of high power JFETs coming out of the labs. In the 1970's when Yamaha and Sony produced power JFET devices for use in their own amplifiers it created some excitement among audiophiles. Unfortunately it seems that they were ahead of their time, and now those parts are quite rare.

Fortunately totally new devices show promise in modern power amplifiers both as switches and linear amplifying parts – tomorrow's transistors (*at tomorrow's prices*).

For power JFETs the I_{ds} dependence on V_{ds} is partly a function of the Gate channel depth on the chip. The pentode-like character is seen in high current power “Vertical” JFETs and the triode-like character is seen in SITs, (Static Induction Transistor).

The enhancement mode power JFETs do well in single-ended Class A applications and can take advantage of load-line cancellation, and my samples of depletion mode JFETs have a more pronounced voltage dependence, and are thus even more interesting.

Here is an example of a JFET follower whose characteristic is mostly “pentode-like”. The upper curve reflects the distortion at $V+$ voltage near the device rating, and the lower curve shows the distortion cancellation available by searching for the “sweet spot” at a lower supply voltage. You can see that at 1 volt, the distortion has been reduced by about 85%.

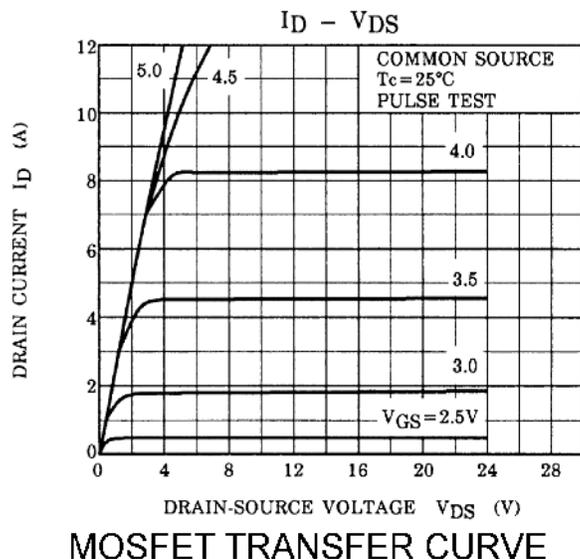


Alternatively we could have adjusted the bias current or the load resistance, and we also could have tried for best performance at some higher voltage than 1 volt.

This example was of one of those “little” JFET devices, but it works with parts of all sizes. The thing to remember is that this cancellation is there somewhere, and you might want to go looking for it. Also keep in mind that the sweet spot is a bit different for every part, even within the same part types, and you will want to consider expending the effort for each individual part if you want the absolute in performance.

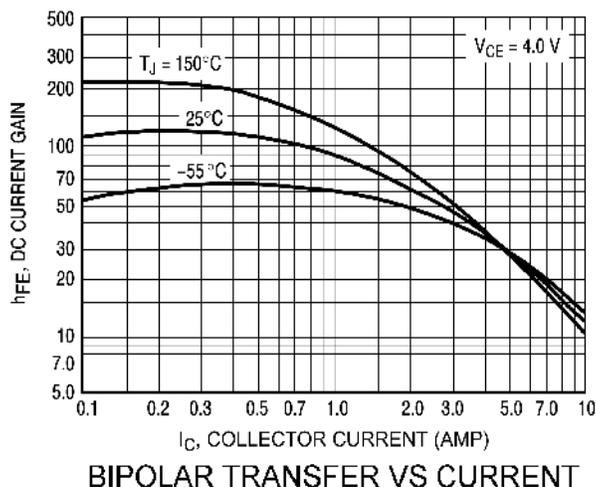
This is important when you change out tubes in an amplifier, as every tube-o-phile knows they are all different. Setting them to a generic bias current does not guarantee the sweet spot for a given tube. You might want to consider how you can vary the supply voltage and the bias while evaluating the performance.

Below is an example of a power MOSFET's transfer curve:

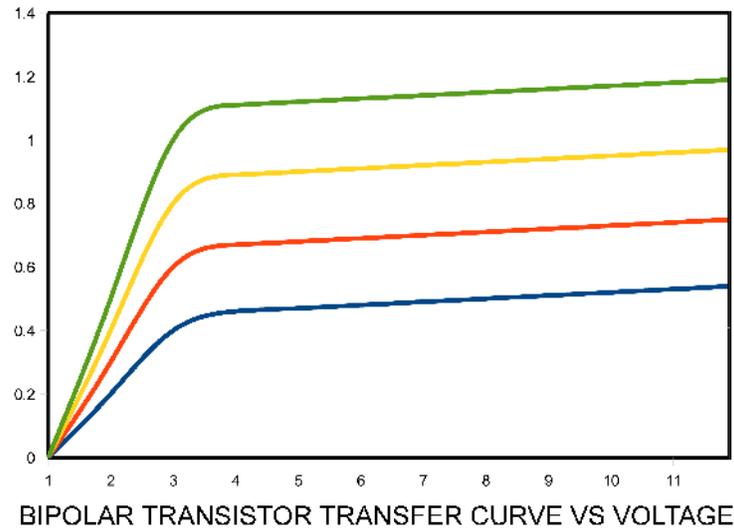


Power MOSFETs have a character somewhat similar to a pentode. The Plate has been replaced by the Drain, the Grid has become the Gate, and the Cathode is now the Source. We see from this diagram that the gain increases with I_{ds} but there is not a lot of dependence on V_{ds} above a few volts. At just a couple volts we operate in what is known as the “linear” region of the MOSFET, where there is strong dependence. You can work the load-line in this region effectively, but it is not a popular technique for a number of good practical reasons. Fortunately, for many MOSFETs operating in Class A, the distortion of the current transfer curve is low enough that you can still find the cancellation for a pentode type voltage characteristic – just go looking for it.

And of course, someone is going to ask, “What about bipolar transistors?” It so happens that bipolars have a sweet spot also. Here is an example of the current gain figure of an ordinary NPN transistor vs Collector current.

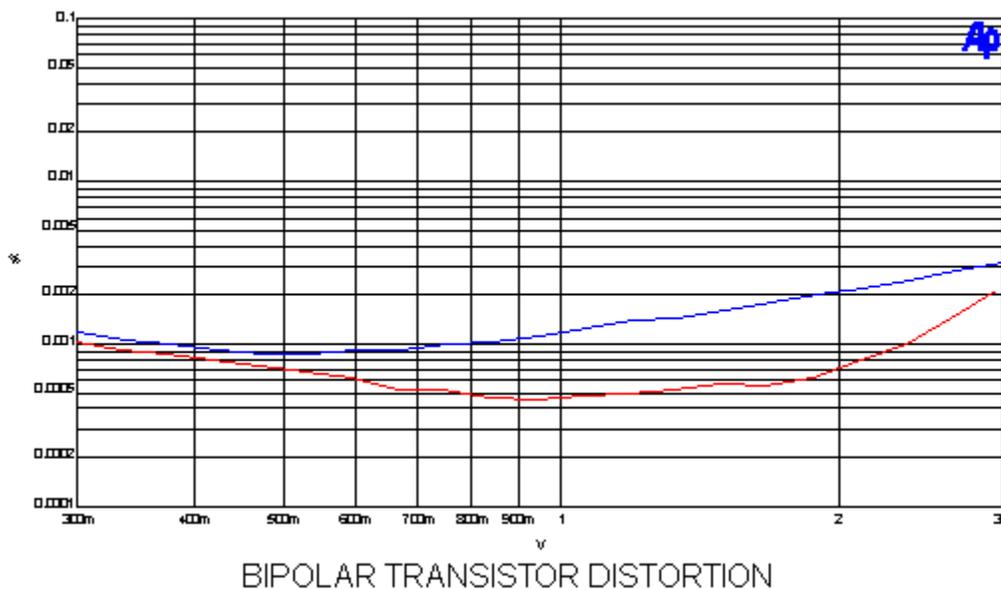


There is also a dependence on V_{ce} known as the Early effect, a slightly straight line dependency increasing with voltage. As with other devices, finding the sweet spot for a bipolar device means locating the point where these variations result in distortion cancellation.



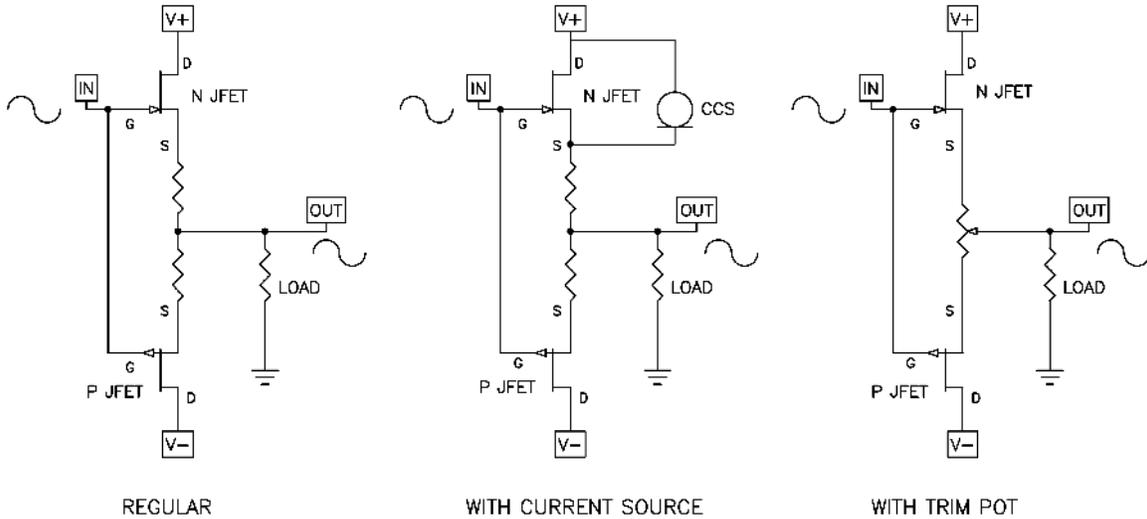
Here we see the gain vs current curve of an NPN power transistor showing some gain dependence vs Collector current. Above a couple of volts, the variation is much less and forms a fairly linear straight line with voltage.

Of course you can work these two curves. The following distortion curve shows a follower circuit like the others, but with a bipolar transistor. The blue line is with a high supply voltage, and the red is the reduced voltage. The reduction in distortion is less dramatic, but it's still about 7 or 8 dB better:



Push-Pull Operation

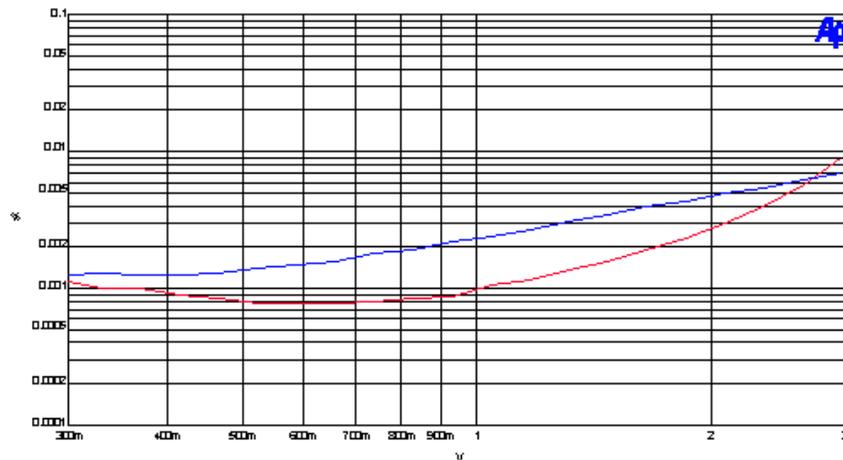
Of course, single-ended Class A is not the only way to operate gain devices. Large reductions in distortion can be had by operating gain devices in push-pull, where two similar devices operate in balanced opposition such that their second harmonic characteristic is naturally cancelled. Below on the left, you see a pair of JFET followers operated push-pull:



PUSH PULL JFET FOLLOWERS

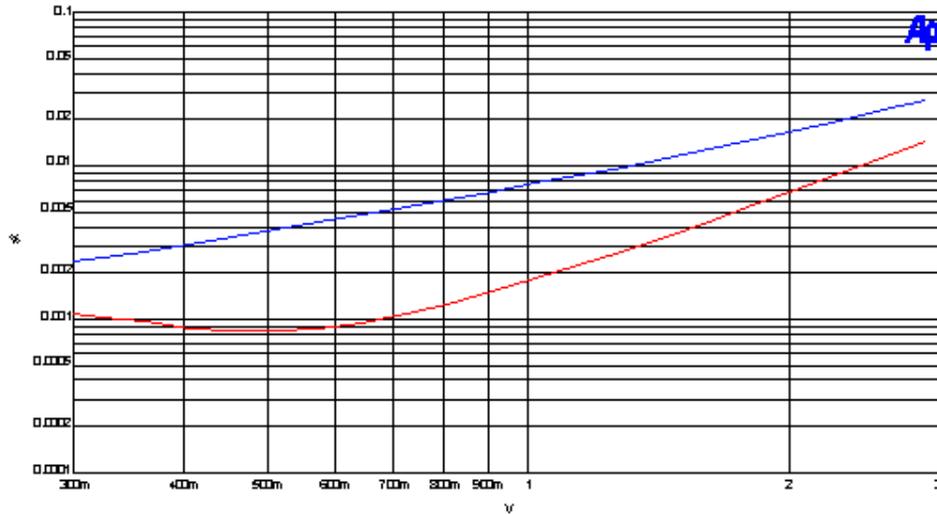
So where's the sweet spot for this sort of circuit? It's there if you go looking for it. The following sets of curves show three techniques for getting a pair of these devices into the sweet spot.

First off you can take the “regular” follower shown above and simply vary the supply voltages a bit. The curve below shows distortion for two examples, the upper blue curve being with two equal supply voltage values. The lower red curve shows what can happen if you vary only one of the supply voltages:



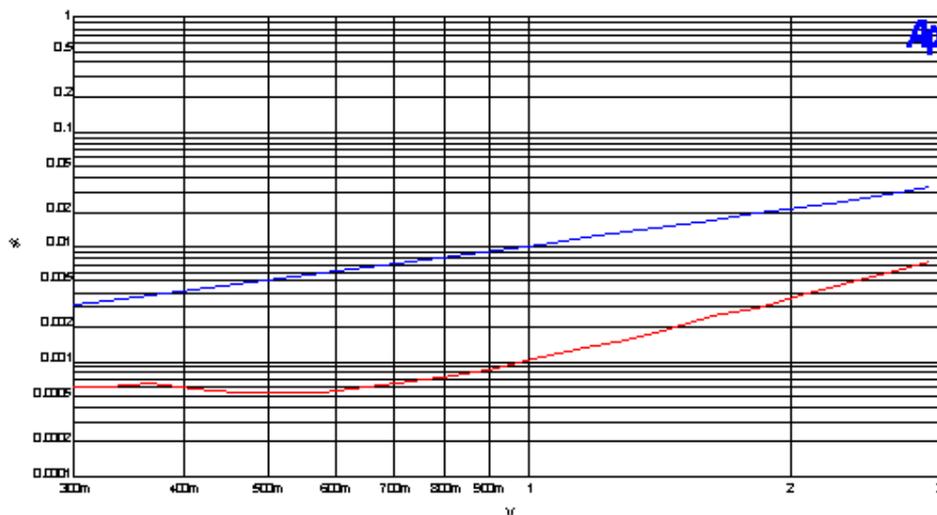
PUSH-PULL JFET DISTORTION WITH DIFFERENT V+

Here's a technique we use at Pass Labs – the use of some single-ended bias in a push-pull output stage to reduce distortion. The distortion curve below shows what you can get if you simply put a small value current source in parallel with one of the JFETs as shown on the middle circuit. The current source displaces the bias current values on the JFETs until they achieve a more perfect distortion cancellation:



PUSH-PULL JFET DISTORTION - WITH AND WITHOUT CCS

There is yet another way, shown by the circuit on the right in the diagram, which involves the use of a trim potentiometer instead of the Source resistors of the “regular” circuit.



PUSH-PULL JFET WITH VARIABLE SOURCE RESISTANCES

Here we see that a simple adjustment of the pot can drop the distortion by 90% without having significant effect on the other aspects of performance.

Temperature

There is one last important variable affecting performance and the location of the sweet spot – temperature. All of these gain devices have transfer curves which are some function of the temperature of the device. All those transfer curves I have shown you will be different if you vary the temperature.

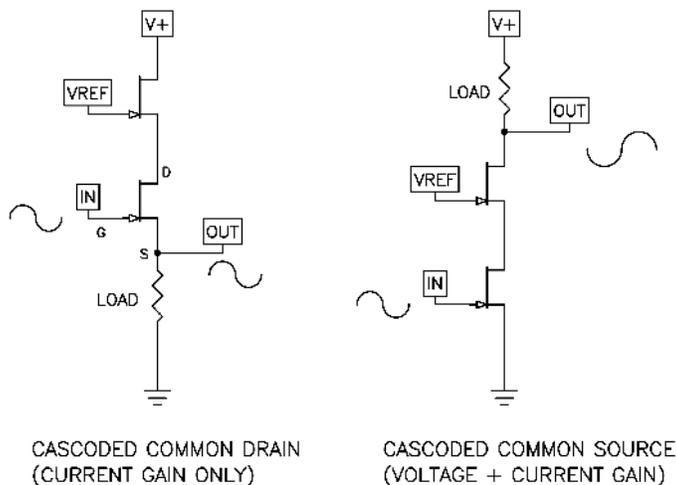
Conceivably you could adjust for the sweet spot by varying the temperature. More practically, you will want to see to it that our circuits have reached their ordinary operating temperature before adjusting for the last bit of performance.

Cascoding, Paralleling and Transforming

So what do you do if it's impractical to find the optimal load-line in a given circuit? Occasionally the sweet spot occurs at voltage values that are impractically low, or at currents that are higher than a device can handle for a given voltage. Here are three things that a designer can do to get into the zone.

The first is cascoding, where the gain device is coupled with a Common-Gate / Common-Grid / Common-Base (depending on the type of device!) tube or transistor which adds practically no influence of its own but which allows a more arbitrary DC and AC voltage across the gain device. Here are a couple of examples, using JFETs in both Common-Drain and Common-Source circuits.

The top JFET is the cascode device, and its Source voltage, which will be seen by the Drain of the JFET below it, is set by V_{ref} . The idea is that the cascoding device provides a “voltage umbrella” for the gain JFET, and all manner of voltages can appear at the output of the circuit while the gain device sees all, a portion, or none of it.



CASCODE OPERATION

It's pretty easy to get into the sweet spot using a cascode, and if you want to explore this in some more detail, I recommend Zen Variations 9, which you can download from www.passlabs.com. There we explore getting the lowest possible distortion from an amplifier using a single stage with a power JFET having a high current rating, but a low voltage and dissipation rating.

Another thing you can do to get into the zone of the sweet spot is to parallel devices. There are times when the load is too low (and you can't change it) or the voltage is too high, but often you can mitigate this situation by operating devices in parallel. This way you can limit the dissipation of each device, and the load appears as a multiple of the number of devices.

A third thing you can do is to use a transformer, which can allow you a flexible range of voltages and currents through the device while delivering the appropriate values to the load. Tube amplifiers do it all the time, and it gives the extra flexibility that then allows you adjust for the sweet spot.

There's a Small Catch

All the loads we have examined are resistors. Many times that is all we need to work with anyway, but when it comes to loudspeaker drivers and passive crossovers we have to contend with loads which are also reactive and which vary in impedance. This can be a problem when the load-line that gives best performance depends on resistance. A reactive load will take it out of the sweet spot.

This is particularly true for tube amplifiers operating single-ended Class A without feedback – much of the performance depends on triode load-line cancellation. These amplifiers are particularly appreciative of resistive loads, and if you have this situation, it's worth considering ways to flatten the load impedance, giving a more resistive load for the amplifier.

There is another solution to this problem if you employ cascoding. Referring again to Zen Variations 9, we see there is a technique, which I have dubbed “Cascode Modulation”, in which the gain device can be made to see controlled load-line voltage variations which are only a function of the output current, and not related to the output voltage. In this way, the sweet spot is preserved into any load reactance or impedance, and it is very easily dialed in by adjusting a couple of resistor values. In the example of ZV9, we get a 90% improvement in distortion by simply attaching the cascode's voltage to a different spot, without changing any values or any other characteristics.

Conclusion

So there we have it. As an audiophile you want the best performance, and you probably aren't above spending money and trying tweaks of various sorts to help you get it. But there are real improvements that can be had without extra accessories or emptying your wallet.

Amplifying circuitry can be made better without more complexity and without more feedback. The best part of this for do-it-yourselfers is that these are things that can be done on your own bench, costing mostly just the time it takes to tweak the circuit and evaluate the results. Moreover, this approach is not seen on the factory floor – for most manufacturers it's simply too time consuming to do, and the rest probably never heard of it.

So here's your chance to make simple high quality audio amplifiers better than ever. Buy yourself a cheap used distortion analyzer and go for it.

Copyright 2009 Nelson Pass