L'Amp: A Simple SIT Amp

by Michael Rothacher

Introduction & Plausible Deniability

I'm not a real amplifier designer, and this is not a real article about a wicked-good sounding, super-simple amplifier anyone can build using light bulbs and alien technology. Now that the pressure's off, we can have a little fun.

All Hail Unobtanium

Okay, it's no secret that I'm a Pass fanboy. So sue me. Some people dress up like Spock and go to comic book conventions. I build amplifiers.

Like many a Pass fanatic I check the Pass sites every hour or two to see if a new project has been posted. (I'm not the only one that does that, right?)

Sometimes months go by, but I am undaunted.

If only Pass would get a Twitter account, then I could get minute-by-minute tweets like:

"*Tested awesome speakers you'll never own.*"

"*Tested awesome speakers you could own, but I bought up all the stock.*"

"*Partied with Thievery Corporation.*"

"*Developed new transistor.*"

Whoa, wait, what? New transistor?

Yep, Srajan Ebaen first broke the news over at www.sixmoons.com. It seems fearless leader struck-up a little skunkworks project with the guys at SemiSouth to develop a super-cool new version of the Static Induction Transistor, or SIT. Pass has pointed out elsewhere that active gain devices are generally the weak link in amplification, contributing a lot more distortion than other components like resistors or capacitors. Until now, little could be done about this, and designers were at the mercy of the big

labs, which primarily make transistors for non-audio applications. But, leave it to Pass to shift another paradigm and crank out some highly specialized gain devices to advance the audio art. Their official part number is PASS-SIT-1. They're awesome, expensive, and completely unobtainable.

Now, it's a well-known fact that audio hobbyists love unobtanium, and any part we can't get (or can't have) usually goes right to the top of our wish list, becoming the part that stands between us and sonic nirvana. So when I first heard about Pass' new super-secret device, I did what any self-respecting audio nerd would do.

I began planning a heist.

As I was searching Google satellite images for possible entry and exit routes in preparation for my moonlight raid on Foresthill, I stumbled upon a cheaper, easier and far more legal way to get a taste of the sweet audible nectar of these intriguing new devices.

It turns out that the people who make Playstation once made audio equipment back when Led Zeppelin still made records, and they even made some very special transistors just for audio. One of those special transistors was the Sony 2SK82, a Static Induction Transistor.

And, you can still buy them (for now).

This article won't be a complete primer on Static Induction Transistors and is mainly intended for beginners with a curiosity for what Pass is working on

in his underground lair, but there's a little something for everyone willing to try these exotic parts in a crazy-simple and delightfully rewarding build.

I bought a quartet of them online and when they arrived I wasted no time plugging them into a circuit to see what would happen.

Curves Ahead

I wasn't able to find the official manufacturer's datasheet for these parts and the data that's out there is sketchy, but that's all part of the adventure. When I first wired-up the 2SK82 it worked, but it wasn't behaving like the characteristic curves from the internet*. Lacking a fancy curve tracer, I tried fiddling with the drain and gate voltage supplies to capture some data points, but this proved to be a little clunky and inaccurate as things tended to drift, so I built up the apparatus in Figure 1. If you have an oscilloscope with X-Y capability, you can use this setup to capture some current vs. voltage (IV) curves for yourself. This can also be handy for times when you'd like to look at a particular sample, or zoom in on a certain part of the graph.
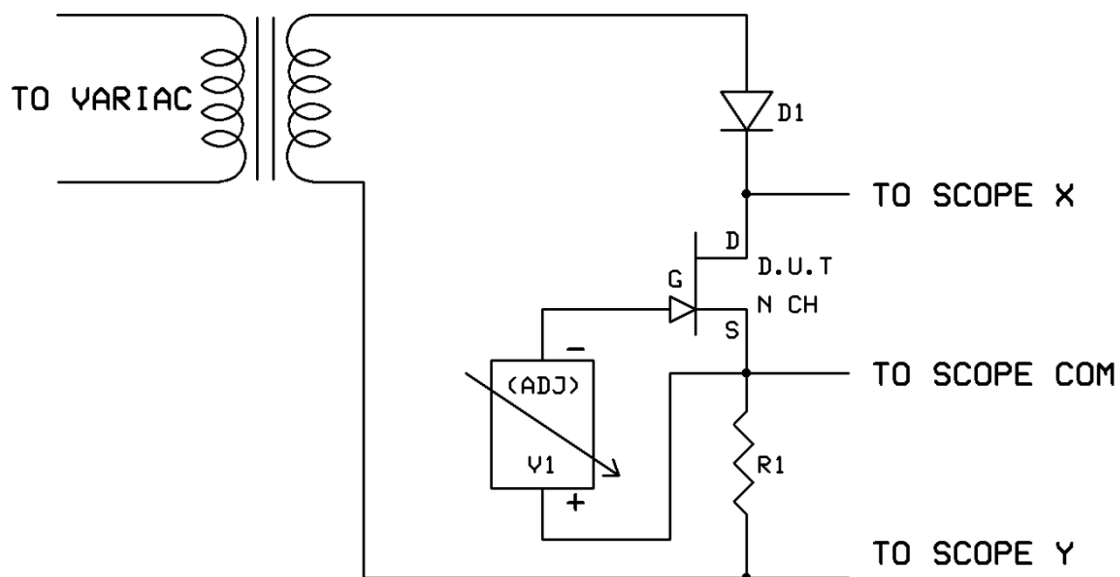

FIGURE 1

I used a toroid transformer with two 18 volt secondaries in series to supply the AC voltage and provide isolation (If you build Pass amps you have lots of these on hand). An 8 amp rectifier supplies a half-wave rectified voltage to the device under test, and a 1 ohm resistor provides a place for us to monitor the current, in this case 1 volt = 1 amp. A bigger resistor would provide some current limiting so we don't accidentally destroy our device, but that would eat up voltage we need across the device for our curves, so you have to be careful, and it helps to know a little about the behavior of your device beforehand. The whole thing is connected to a Variac, so you can bring the voltage up slowly and adjust it as needed. Remember, we're not testing little devices; we're working with bigger voltages and currents, so safety first. Don't touch the apparatus while it's powered up and shield yourself from the device under test. I wear safety glasses and keep a fire extinguisher right next to my workbench just in case.

For the 2SK82, the bias supply is negative, and for enhancement mode devices you would reverse this. Also, if you intend to test a P channel device, you'll need to flip the diode polarity in order to provide the negative half of the wave to your device. You can test other devices as well, even light bulbs.

Mounted to a proper heat sink, I connected the 2SK82 to the makeshift curve tracer, dialed-up -12V, and adjusted the supply. Then I repeated this for the remaining Vgs values (-10, -8, -6, -4, -2, 0). I took a screen shot of each trace and composited them together to get Figure 2.
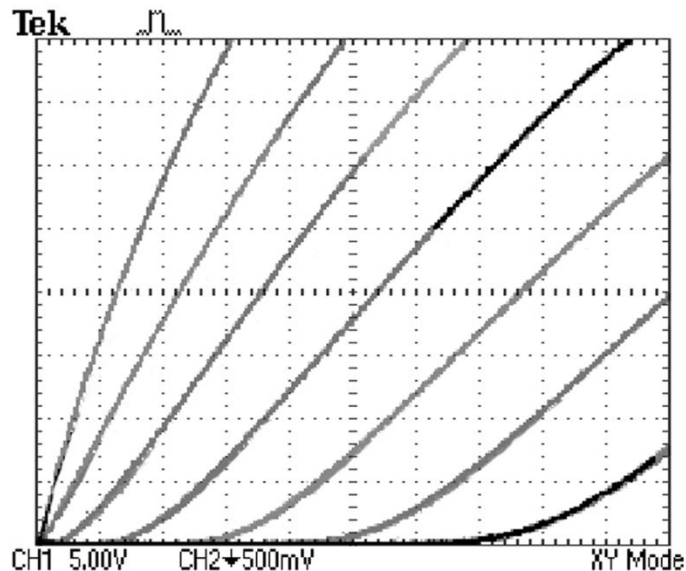
FIGURE 2

A number of inexpensive USB scopes would do fine here, and many of them allow you to export the waveform data right into Microsoft Excel.

I'll also point out that real curve tracers use a staircase generator on the bias to get all of the curves on the screen at once, and you could do the same using an arbitrary waveform generator. I just do it the cheap and simple way.

Warning: This Next Part Is Pretty Graphic

Now we have some curves, so what can we do with them?

Even if you don't intend to do any curve tracing of your own, understanding how they're made should help you to understand how they graphically represent the operation of a particular device. If you're just starting out, all those datasheet curves might seem a little esoteric and you might wonder how others divine so much information from them.

MOSFET and JFET datasheets will usually include a graph like the one we made in Figure 2. We know from our curve tracing experiment that the horizontal axis represents the voltage across the device from Drain to Source (Vd). The vertical axis is current flowing through the device from Drain to Source (Id). And, each curve represents a different value of Gate to Source voltage (Vgs).
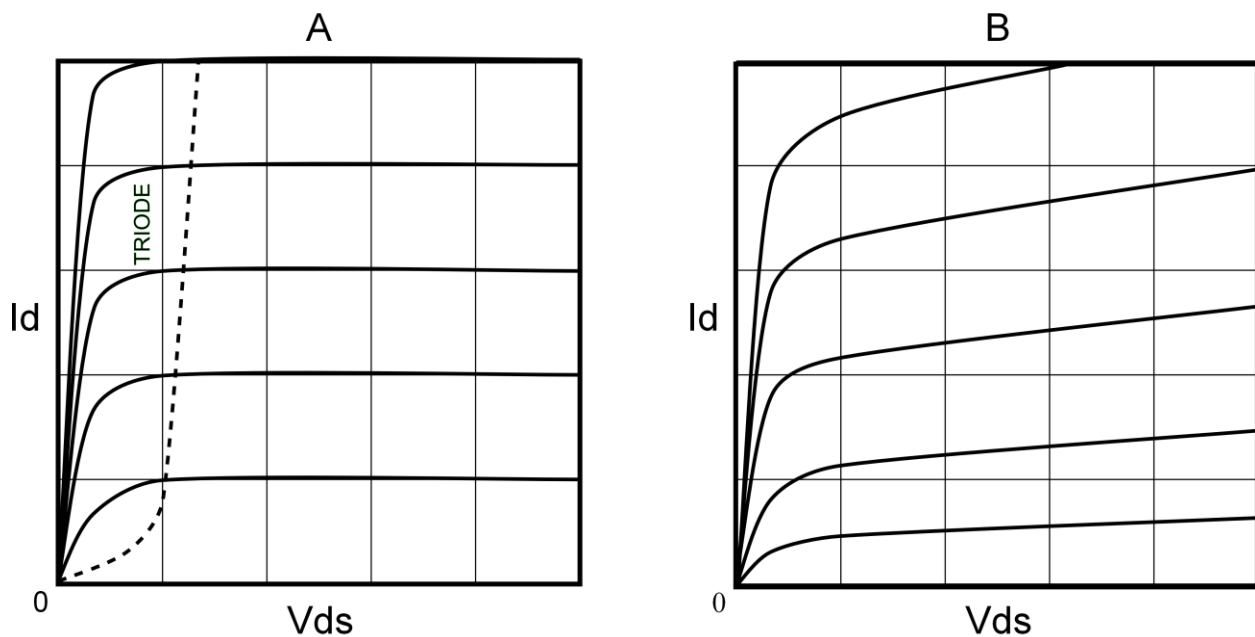
A

B

TRIODE

Id

Vds

Id

Vds

FIGURE 3

Figure 3(a) is a set of curves for an imaginary MOSFET transistor. Starting at 0 volts from Drain to Source we see that the current (Id) rises sharply for the first few volts in what is called the triode region, then levels off in the so-called saturation region. This sort of curve is very similar to a Pentode vacuum tube.

We notice that the lines in the saturation region are nice and evenly spaced. The spacing here represents the gain. If we want lots of gain, we'd like to see a big change in Id for a small change in Vgs. Change in Id over the change in Vgs is called Transconductance ($\Delta I / \Delta V$). Variations in gain will cause distortion, and this would be a low distortion device indeed, since the gain is perfectly uniform throughout. We could select an operating voltage and current (*Operating Point)* just about anywhere and get a good low-distortion result.

However, we're more likely to encounter something like Figure 3(b). Here, the spacing between lines isn't even and we see that they tilt upwards and diverge in the saturation region. In a case like this, you might move your operating point closer to the top of the graph where the spacing is wider and more even (Higher gain, less distortion). If you've built the MOSFET Zen amps, it's possible you've developed a penchant for turning up the current; this is because the MOSFETs we use look something like this.

Now let's turn to Figure 4(a). Again, the horizontal axis is the voltage across our imaginary device, the vertical axis is the current flowing through it, and each line represents a different Gate to Source voltage. These curves clearly look different. We see that Id rises with Vds throughout the graph much like the triode region of the MOSFET, so Id is dependent on Vgs *and* Vds. In this ideal example, the gain is nicely uniform, so selecting a low-distortion operating point would not be difficult. However, we're more likely to see something like Figure 4(b). In this graph the lines aren't evenly spaced or parallel. As before, we'll want to select our operating point accordingly, but as you will see in a moment, there is a *twist.* If you're thinking this graph looks a bit like the 2SK82 curves we made earlier, you're right; this is what Static Induction Transistors look like, and they closely resemble Triode vacuum tubes.
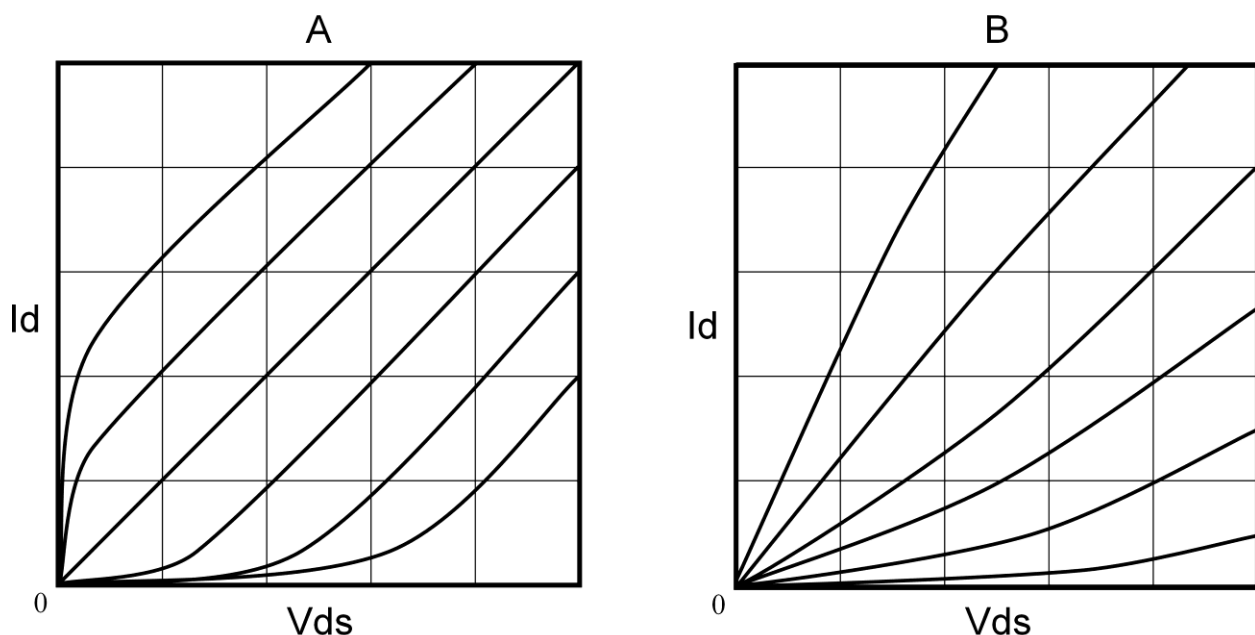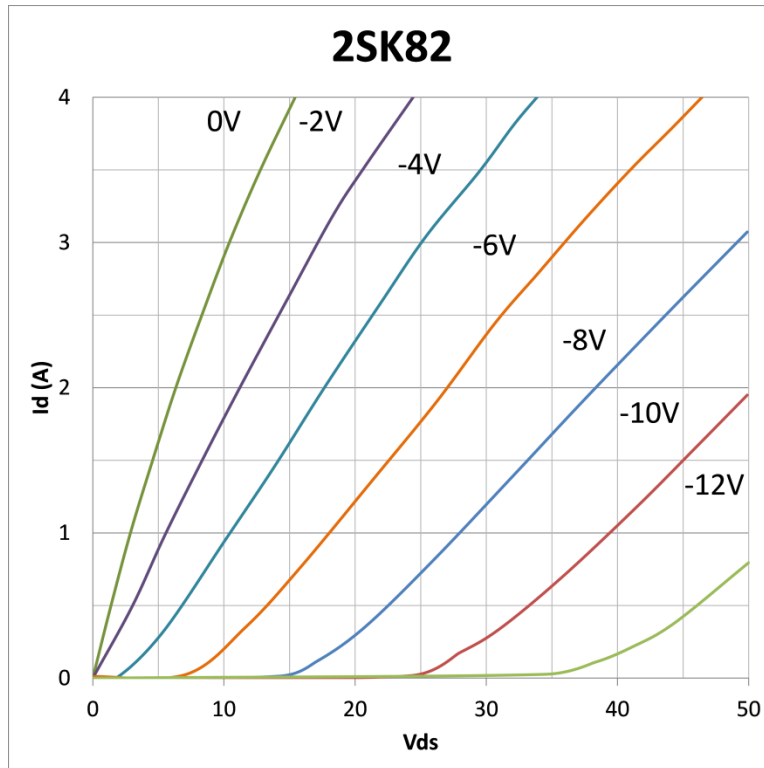

FIGURE 4

FIGURE 5

Now, let's return to our real curves for the 2SK82 (Figure 5). Our Vgs values in this graph are negative because the 2SK82 is an N channel depletion mode device, and we'll need to interpolate (by calibrated eyeball is fine) for values of Vgs not drawn.

One thing that catches our attention is that the 2SK82 will conduct when the Gate and Source are at the same voltage, which would be nice for self-biasing. However, Id increases very sharply here so we need to keep an eye on our dissipation as we increase Vds. The 2SK82 is rated for 95 watts and, ideally, we might want to see a 4:1 safety margin to ensure our parts will last and perform well in the hands of regular consumers, but this is DIY; we have soldering irons and we know how to use them so, we'll push it to 45 watts or so. It will be helpful to see this maximum power on the graph, so for each value of Vgs, let's draw a point where the dissipation equals 45 watts and connect them with a curve as in Figure 6. When we choose our operating point we'll want to stay below this dashed line.
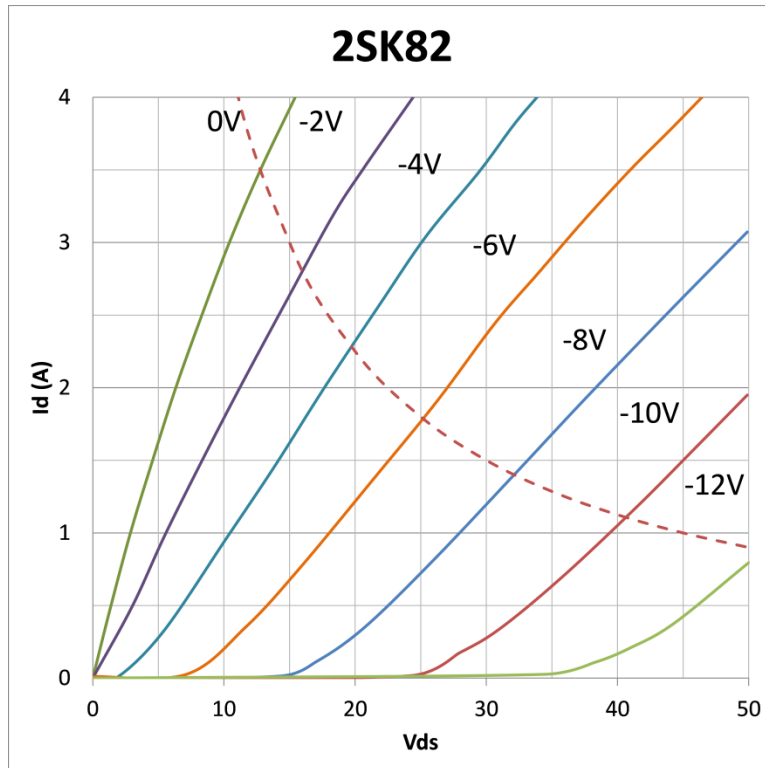
FIGURE 6

Now, what about all of this load line and sweet spot business? First let's talk about what a load line is. Let's draw a line from 32 volts (Vds) to 4 amps (Id) as in Figure 7.
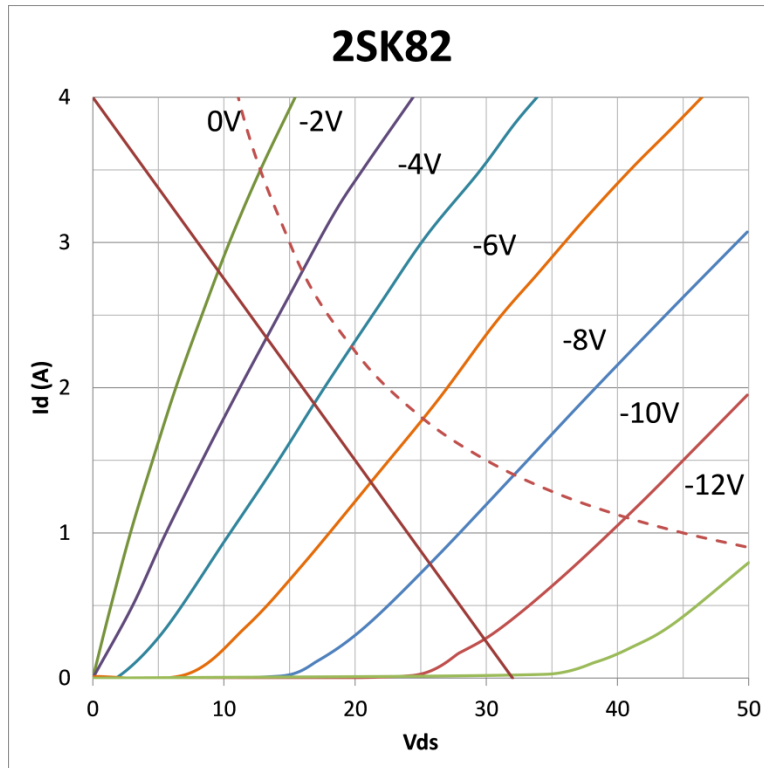
FIGURE 7

There you go; that's a load line. In this case, it represents an 8 ohm load (32 volts divided by 4 amps). Not so bad, right? Of course, there are an infinite number of 8 ohm load lines we could draw, but they will all be parallel, having the same slope. And, the load may be something other than eight ohms; it might be the impedance of a transformer primary, the impedance of light bulbs in parallel with your speaker, or something else altogether.

Now, let's choose an arbitrary operating point along our 8 Ohm load line and place a dot there as in Figure 8.
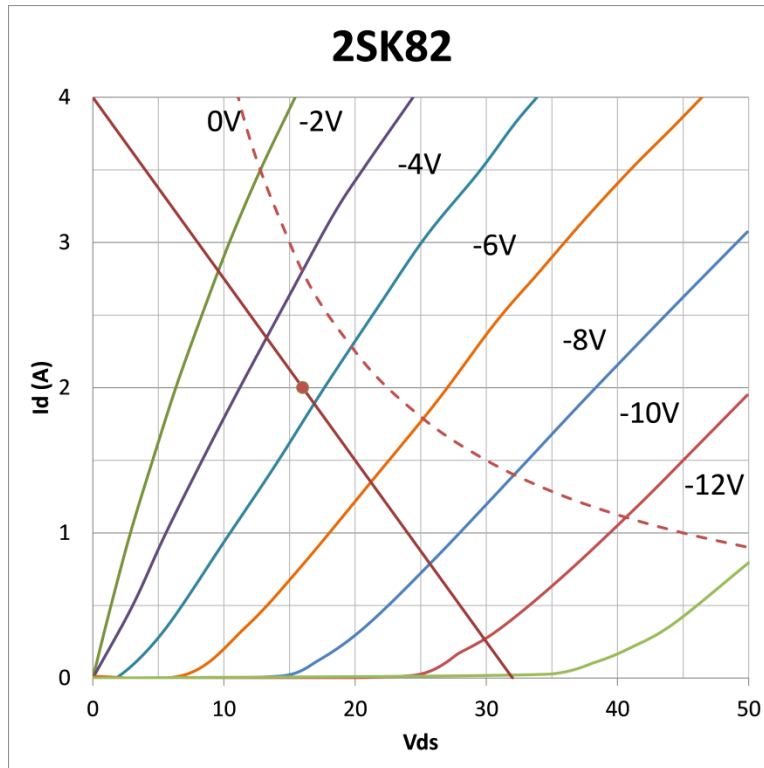
FIGURE 8

Here I've chosen 16 volts, 2 amps. Now imagine rotating our load line around the operating point.

You may want to print this graph and try it with a straightedge.

If we twist the line clockwise, the slope gets steeper, so the load decreases. Conversely, when we twist the line counterclockwise, the slope gets shallower, so the load increases. As we twist our load line back and forth, watch the spacing between curves along the straightedge; the spacing changes, and when they're the most evenly spaced, we've found the point of minimum distortion for this particular operating point. Keep twisting until the light bulb above your head comes on. Eureka! You've got it.

Well, almost. Changing the slope of the load line is only one of the variables and sometimes you have little or no control over the load value. In that case, we can move the operating point to find the sweet spot. Here, the load line (or our straightedge) stays at the same angle as we work

around the area of interest. Of course, we have to choose an operating point that provides adequate voltage swing, and the sweet spot isn't necessarily the point of lowest distortion; maybe you're after a special blend of odd and even harmonics, or maybe you have other design constraints you need to achieve. In any case, the sweet spot is there for the finding.

Hopefully, this little thought experiment has helped you to understand what happens when we vary certain parameters in search of the sweet spot. I'd like to point out, however, that finding the sweet spot on paper is really only useful as a *rough* starting point. In real life, speakers aren't resistors, and gain devices never look exactly like the datasheet. For that matter, they seldom look like one another, and if you've ever matched MOSFETs for Vgs, you already understand this to be true.

The Amplifier Already

Much of what follows is expanded upon in thirty-plus articles published by Nelson Pass and I hope you'll permit me to refer to them here in a general way rather than flood you with citations.

The articles of Pass are the *I Ching* of do-it-yourself amplifiers, particularly the less-is-more Zen variety. They can all be found at www.passdiy.com or www.firstwatt.com . If you have not already done so, print them out, or load them on your Kindle, then retreat to a nearby mountaintop, meditate on them until you have reached enlightenment, then return to your workbench and use real parts with real electricity to reinforce what you have learned. As I have discovered for myself, you may need to repeat this process many times.
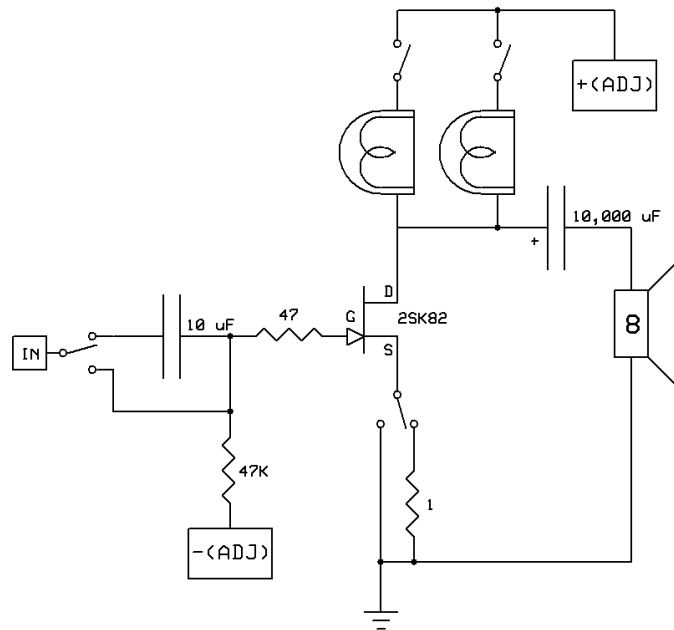
FIGURE 9

Figure 9 is my "throw science at the wall and see what sticks" test setup. It has adjustable power supplies for the bias and V+, and selector switches (alligator clips, actually) for the various components, which allowed me to try a few different configurations quickly and without a lot of fuss. You'll recognize this as your basic, common source amplifier similar to Zen-Lite, De-Lite, et al. The beauty of Nelson's light bulb amplifier projects is that they're easy to build and learn from, while being fun to listen to and watch in the dark.

The first configuration I tried is seen if Figure 10. There's a cap on the input and no source resistor. The input impedance is the value of R1, and the output impedance is approximately equal to the total resistance of the light bulbs. I used two 300 Watt light bulbs whose parallel resistance is in the 11-13 Ohm range. Make sure the bulbs you buy are 300W @120V; the first bulbs I tried were marked 300W@130V, so they're really 266W bulbs. In any case, you're looking for 11-13 Ohms and you can always use a big power resistor.
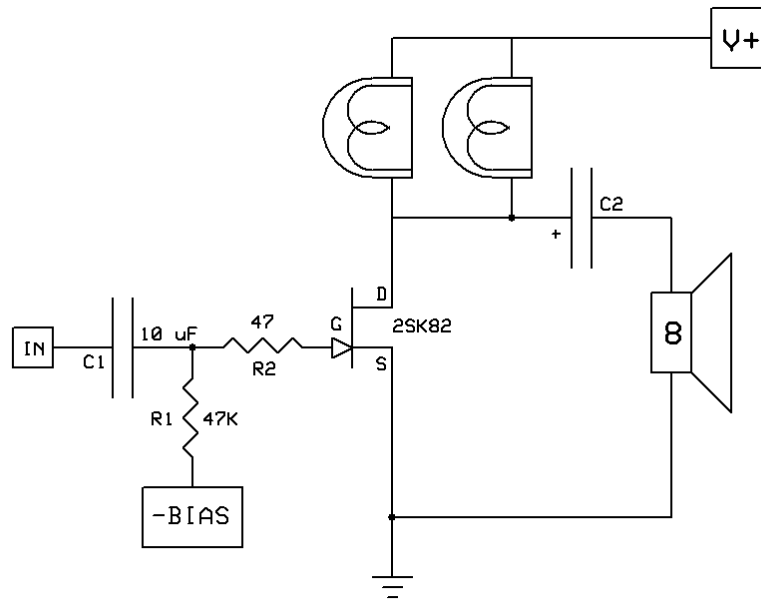
FIGURE 10

In trying to keep to a logical workflow, I moved the operating point around the area of interest by stepping through several power supply values, and for each one I adjusted the bias for lowest distortion at 1 watt. I also checked the voltages across the JFETs and the bulbs.

If you plan to try this for yourself, you'll need a distortion meter of some sort. There are many used distortion analyzers available that will get you down to .01% or so. You could also try one of the soundcard + software solutions, and I'm sure you will find a discussion about them somewhere in the Pass Labs forums. I use an HP8903, which isn't hard to come by, and measures down to .003%.

Here are some of my findings:

L'AMP NO SOURCE RESISTOR & NEG. BIAS SUPPLY (FIGURE 10)					
	Q POINT				
V(supply)	V(ds)	I(d)	V(bias)	V(bulbs)	1W THD%
40	13.6V	2.31A	3.04V	25.83V	.168%
45	13.93V	2.55A	2.69V	31.17V	.106%
50	13.37V	2.74A	2.27V	35.81V	.143%
55	15.3V	2.91A	2.91V	39.7V	.575%

You can see from the data that the resistance of the bulbs changes a little with the voltage across them, so the load line moves slightly as you adjust the supply. This wasn't a big problem, but noteworthy in the context of our discussion about load lines.

The measurements were different from my intuition, which told me the middle of the graph had potential, but I tried one, two, and three light bulbs and the results weren't very interesting. Perhaps that area could be exploited with some other load.

The gain of this circuit is the effective load divided by the inverse transconductance of the JFET. Returning to our characteristic curve, we see that the transconductance in the operating area is around .4A per input volt, or .4 Siemens. The effective load is our light bulb resistance in parallel with our 8 Ohm speaker, or about 5 Ohms.

$$5 / (1 / .4) = 2, \text{ or } 6 \text{ dB}$$

That's not a lot of gain, but the distortion performance is looking pretty nice and this is a flea-watt amplifier anyway, so we press on.

You could build this one using a suitable negative supply for the bias, but maybe you don't have the equipment for dialing in the sweet spot, and you just want something you can build and listen to without a lot of knob-twisting. Figure 11 doesn't require any adjustment, or a separate bias supply and it'll get you pretty close.
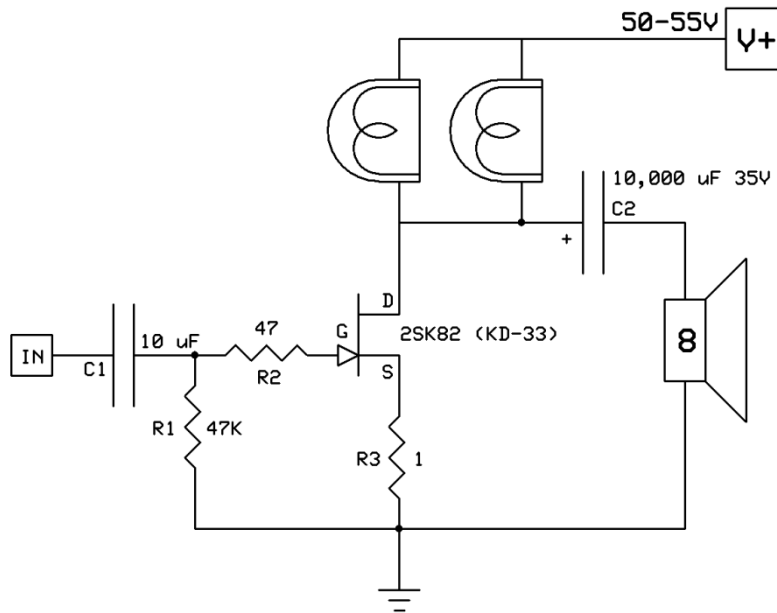
FIGURE 11

It has a 1 Ohm source resistor, which lowers the gain (and distortion) even further. Here, the gain is the effective load divided by the inverse transconductance of the JFET plus the 1 ohm resistor.

$$5 / ((1 / .4) + 1) = 1.4, \text{ or 3dB}$$

R1 is attached to ground and the JFET is said to be self-biased because the Gate is kept at 0V and the Source is raised by the voltage dropped across R1, so the voltage from Gate to Source becomes $-V_{R1}$. If you build this one, you'll want to use SITs graded KD-33, which Acronman sells in N Channel singles.

Having captured a few interesting data points at the bench, I assembled a pair of plywood prototypes. Then I lugged the prototypes, three regulated power supplies, a bunch of clip leads, and my notebook two flights upstairs to my listening room to see if they liked music. You can build your amps this way too, and if you plan on experimenting a lot, it will afford you a lot of flexibility. Look for the nice dual *linear* regulated supplies rated at the voltages and currents you see in the various Pass projects. Here's an action shot of the prototypes:

The Build

Having listened to the plywood prototypes for a while, and needing my regulated power supplies back in the lab, I built the more portable and domestically acceptable version you see here. They're monoblocks.

Figure 12 is the power supply schematic for one channel, and there's nothing unusual here. I used 400VA transformers from Antek, which have two 40V secondaries, which are wired in parallel. The Anteks also have two 15V secondaries, convenient for adding a little regulated supply for the bias if you decide you'd like to try that later.
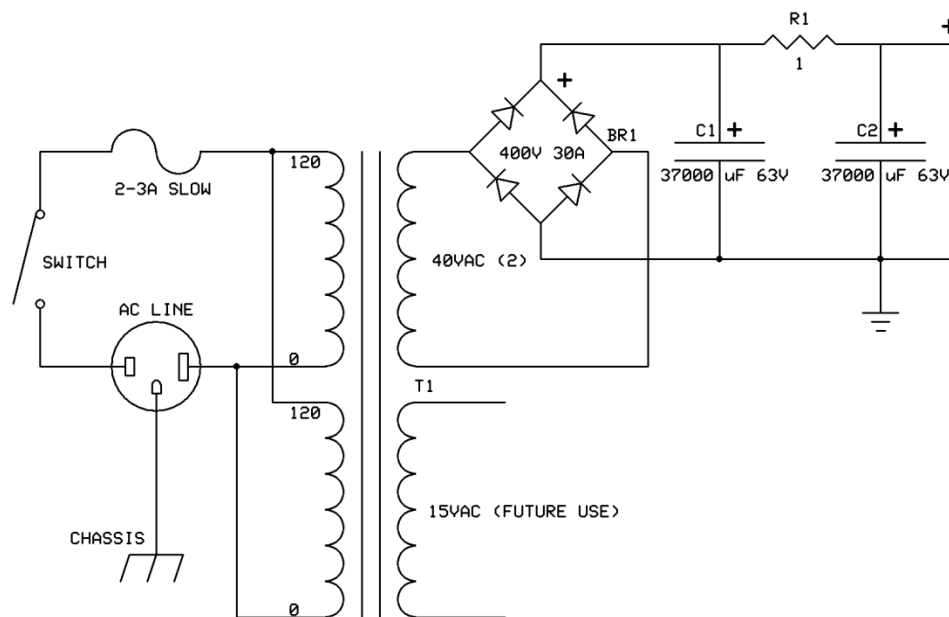

FIGURE 12

I wired everything point-to-point using solder lug terminal strips, keeping my wiring as short and neat as I could. I know the thought of wiring point-to-point makes some do-it-yourselfers tremble, but a PCB would just be silly here.

The SITs dissipate around 45 Watts, and I attached them to big heat sinks from Wakefield, drilled for TO-3 transistors. Their thermal resistance is around .5C/W. I used sockets and thermal pads on the transistors, but mica and goop would probably be better. Here's a pinout diagram for the 2SK82.
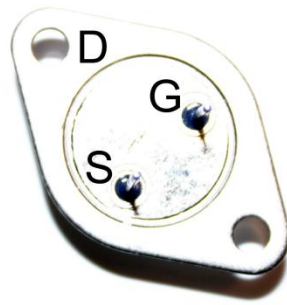
I attached R3 to a small heat sink, but you could attach it to the same heat sink as the JFET. I included a few niceties like two input jacks; one is direct coupled and the other has a cap. Then there's the Sweet Spot Meter, which is mostly a gimmick, but fun. Here's what it looks like close up.
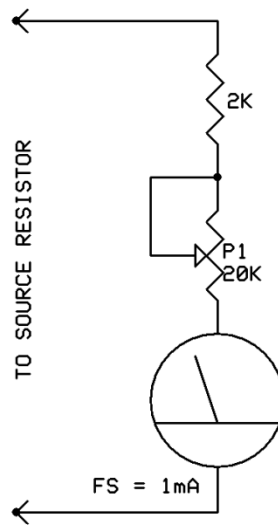
Figure 13 is the wiring diagram.

FIGURE 13

The idea is to adjust P1 so that the meter goes to ½ full-scale (the sweet spot mark). Later, you may decide to try the separate bias supply and/or plug your amps into a Variac to fine-tune them with a distortion meter, after that, you can readjust P1 for the sweet spot, making it easy to return to the same setting later.

Finally, I should add that I simply addressed turn-on thump the manual way by installing one of those banana plug shorting bars during power up.

Measurements & Stuff

This little amp measures pretty well, especially when you consider it within its class. Figure 14 shows distortion vs. power for the Figure 11 circuit.
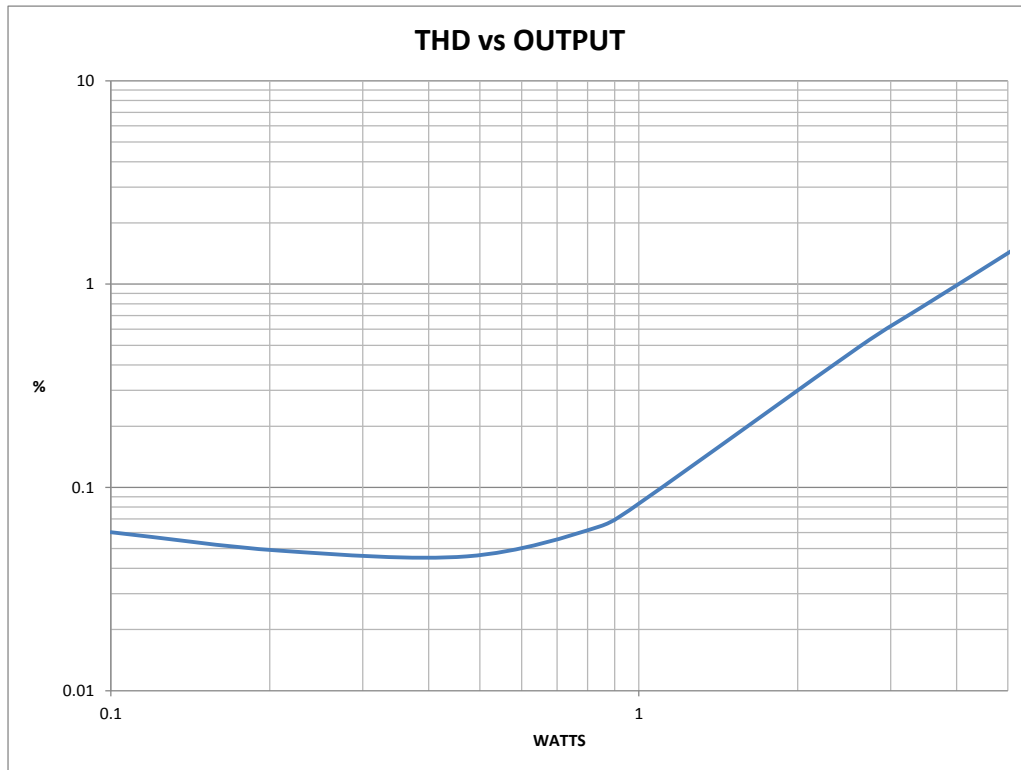

FIGURE 14

We're looking at 1 watt distortion below .1% in a no-feedback amplifier. Of course, the gain is pretty low, and we don't want to obsess on numbers too much, but it's interesting to note how far we've advanced along the performance vs. simplicity curve, and I can only imagine what the PASS-SIT-1, with higher transconductance and lower distortion, might yield.

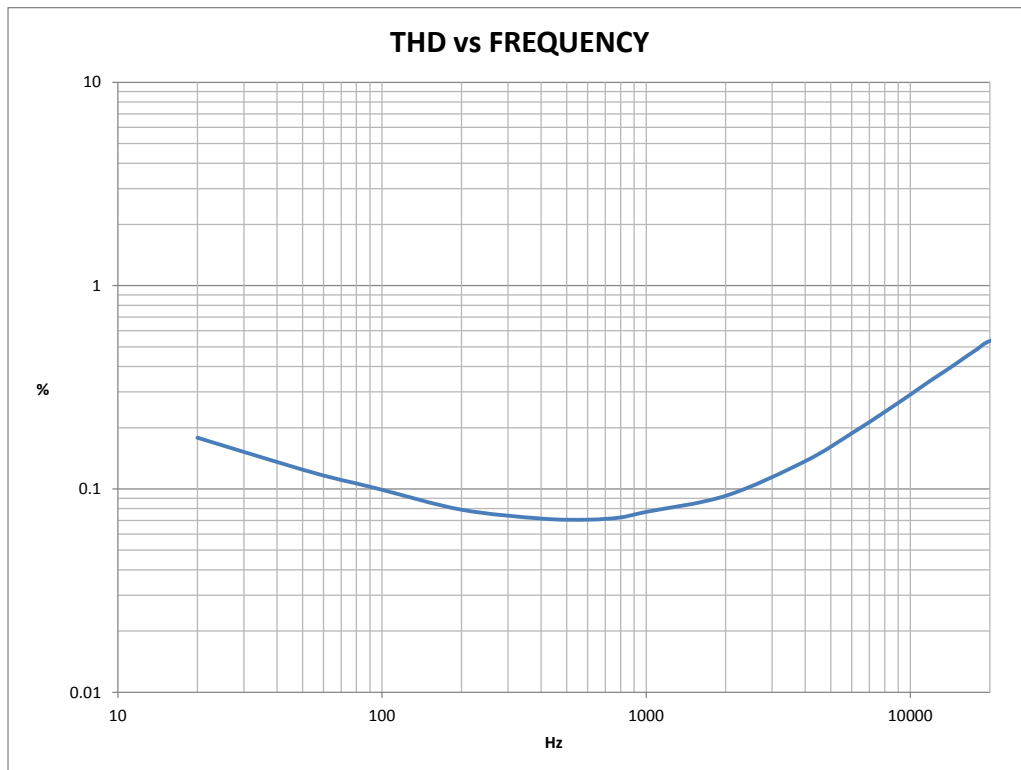Figure 16 is the distortion vs. frequency at 1 watt.

FIGURE 16

The frequency response is -.04 dB at 20 Hz, and -1.5 dB at 60 kHz. Figure 17 is our square wave at 40 kHz.
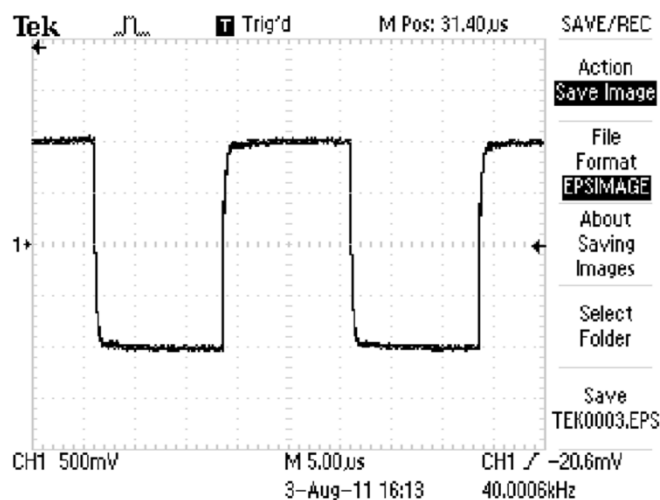
FIGURE 17

The Sound

I really wish there was an independent review outlet for DIY projects. Authors don't really like to self-review their stuff, but potential builders want to know about the sound before they commit to a project. Perhaps someday, some clever online publisher will figure out a way to make such a service viable. Until then, I'll do my best to describe this thing for you.

First, I should tell you that I'm in agreement with the notion that this is the entertainment business and I think your audio equipment should make *you* happy. If you like to wear really uncomfortable shoes because they look good, I respect that, but I'm more of a tennis shoe guy. Next, let's put things into context; five watts is five watts, and it's hard to know if you'll like the flea-watt game until you've tried it, but clearly it has a sizable and very dedicated following, so on with the review.

One thing little zero-feedback tube amps often get right, and transistor amps seldom get right, is what reviewers call "palpable presence", which gives the music that magical "you are there" quality. When this is really

good, the vocalist and musicians are in the room with you, each occupies a distinct location in three-dimensional space, the walls disappear, and the next thing you know it's 2:00 AM and you're running out of records.

With the right speakers, in a smaller room, this little SIT amp can do that.

Of course, it doesn't have a lot of power and it has a low damping factor, but if you use it with the same music and equipment you might choose for a little 2A3 amp, you should be fine. You'll need an active preamplifier or other source capable of 5 or more volts output, and high-efficiency speakers. I did my listening with a pair of Tekton Lore's, and I have found these 10-inch two-way floorstanders to sound consistently good with a number of low power DIY projects.

So, what about the inevitable question: does it sound tubey?

Yes and no. It is certainly one of the tubiest solid state amplifiers I've listened to in terms of the aforementioned "palpable presence". It does have a nice warm midrange too, but it bests a number of tube amps I've heard in terms of detail. The highs are sweet, but not syrupy, and the bass was nicely punchy and full, but not too tubby. As the reviewers might say, with the right music, speakers, and room, it's great. Of course, what matters most is what *you* think, so if you build one, I hope you'll drop me a line and let me know.

The Pleasure of Finding Things Out

Well, that's our amp. Many improvements are possible, and perhaps you're already thinking about current sources, buffers, and so forth. You could certainly improve the performance and efficiency of this circuit, but I recommend you build and enjoy it in this simple form before adding parts and complexity. This will give you a chance to get to know the sound and learn what happens when you tweak the various parameters. Get yourself a distortion meter that allows you to monitor the distortion waveform, or a spectrum analyzer, so you can try different harmonic seasonings and see how they affect the sound.

You can learn a lot from books, articles, and the forums, but the best way to reinforce these concepts is to build and listen, then repeat as needed. Admittedly, I fall into the "plug it in and see what happens" camp of audio enthusiasts. I encourage you to study much, but be adventurous too, try things for yourself, and avoid analysis paralysis. Discovery and optimization generally happen at the workbench, not at the drawing table.

So Long and Thanks for All the FETs

Well, I hope you've enjoyed this little entertainment piece. If you're new to DIY and not sure where to start, perhaps you'll give this amp , or one like it a try. The idea is to hook you with a little sample so there will be more of us in this addicting but mostly harmless hobby.

I'd like to thank the gang at the Pass Labs Forum for their help and, of course, Nelson Pass for tirelessly answering all my questions and never once throwing an object at me.

Now go build something!

© Michael Rothacher 2011

---

*The curves at www.circuitdiy.com are fine but 0V starts with the third curve.