Nelson Pass Talks at Burning Amp Festival 2012

The Burning Amp Festival is a DIY audio event held in San Francisco in October for the past few years. This is a loose transcription of Pass talking without notes but with pictures. Edited for clarity and brevity and to make Pass sound more articulate than he really is.

(We open with smattering of applause)

OK, well once again, as always what we do here is uh..., I string a bunch of pictures together and then talk about each one. And the subject of this talk is the First Watt F6 design which those of you who hang out on the forums at DIY Audio know......

*(**loud brown note** from the PA system - "Too much feedback!" audience laughter at clumsy microphone manipulations)*

OK! The F6 is a project where I looked around for something to do and saw that we already had an F5, and I did want to do something a little different. What drives this design is a couple of things: One is that I wanted to do more with these power Jfets, the ones that are no longer being made. The bad news being that SemiSouth apparently has closed its doors. The only thing I can say is that I bought a big pile of them and so have every intention of marching ahead. Also I wanted to do a little more work with audio transformers.

I announced the F6 on the Pass forum at DIYAudio with a little teaser schematic of what I had in mind:

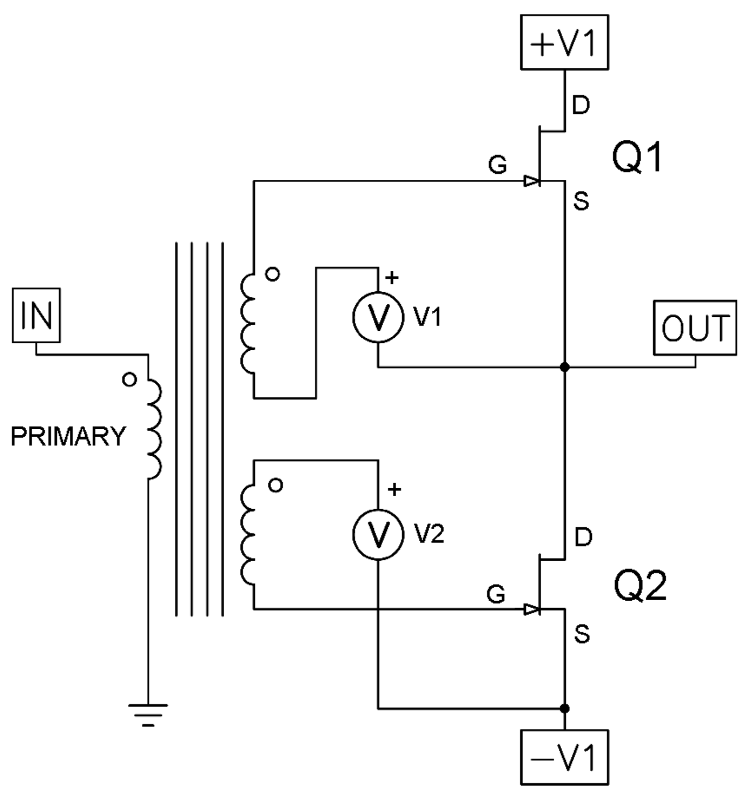

F6 - MOST SIMPLIFIED

Actually, I originally put up an erroneous version of this schematic, but the corrected version you see is a push-pull Class A power Jfet amplifier. The power Jfets that I have are only available as N channel versions – there aren't complementary P channel parts.

We don't have the ease of operating them in push-pull stages that we get with complementary parts, so usually some accommodation has to be made to get push-pull operation. A lot of amplifiers use what is known as quasi-complementary, stuff that dates back to RCA, the Harmon Kardon Citation 12, Phase Linear and other things that came out of the late '60's, and they inevitably involve some additional circuitry that allows the device on the negative side to behave as if it were a P channel part.

Even when you have P channel parts, there's some flaw, being that there's really nobody making truly symmetric P's and N's. That hasn't stopped anybody - it certainly hasn't stopped me, but I found an opportunity here to maybe play around a little bit using a transformer similar to the one used in the M2 amplifier, where the transformer produced all the voltage gain. In the M2 I drove the primary of the transformer and the secondaries had more windings on them so that all your voltage gain can be generated by that, and then a complementary output stage with pair of N and P channel Mosfet power transistors as voltage followers, and *Voila*! You've got a single-stage amplifier with no feedback. It's a nice little amplifier, and I am still selling them.

As I said, I wanted to so something a little different and I still had some stock of the transformers I had been playing with, and I decided first off to build an amplifier which was push-pull complementary but with only N channel devices and really good decent symmetry between the plus and minus halves, and one way to achieve that was how they used to do it back in the old days - with transformers.

And like I said, I had some laying around...

Looking at the above schematic, the primary winding drives two secondary windings with a 1:1 ratio. The top transistor Q1 is driven in phase and the bottom negative transistor Q2 is driven out of phase, using only N channel devices with good symmetry and biased into Class A . You can build this amplifier as presented.

Finding this particular transformer is kind of a rarity and so you look through the catalogs and most of what you find are two primaries and two secondaries. That's OK, we'll adapt to that.

The little circles with a “V” in them are voltage sources. The Fets we use have to be given positive Gate voltages relative to the Source pins (note the S for Source, G for Gate and D for Drain on the transistors). We put a little positive voltage on these things to get them to turn on - in this case about 1.2 volts will get them running. You have to provide that 1.2 volts, and in this circuit I have specified voltage sources. In each case the voltage source is relative to the Source pin of the device.

You bias the Fets up at about 1.2 volts and you have an actual amplifier. There is a small problem – The Jfets would have to be perfectly matched and all the other conditions would have to be equal because there's nothing in here to control the DC offset voltage at the output – it might be free to drift around a little bit.

Here is the less simplified version:

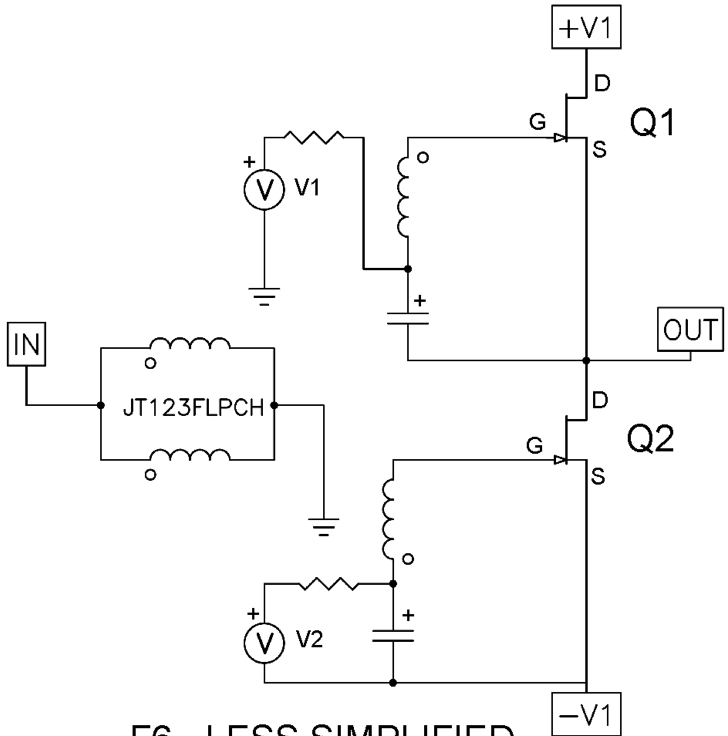

F6 - LESS SIMPLIFIED

It shows a real part, a Jensen JT-123flpch, a nice little $30 transformer which mounts on the pc board and has four windings which are all identical. In case you don't know this about transformers, the little dots indicate the polarity of the windings, so if I put a positive voltage on the dotted end of one winding, a positive voltage will appear on the other dotted end of the other windings. The primaries are running in parallel, and of the two output windings, the one driving the negative transistor Q2 is flipped in polarity. When a positive voltage appears at the input, a positive occurs at the Gate of the positive transistor Q1 and the output will follow that. A negative appears at the Gate of the negative transistor Q2, causing it to carry less current, supporting the positive output of Q1.

And that's a push-pull circuit.

There are a couple other differences; Instead of referencing the bias voltage to the Source of the output transistor Q1 as in the previous circuit, I now reference it to Ground. This solves the DC offset problem I mentioned, and the transistors don't have to be perfectly matched and you don't have to tweak them so that the temperature is always the same and so on. As the output DC starts to drift this will drive that transistor in such a way as to compensate.

Now we have a circuit where the voltage V2 determines the bias current of the amplifier and V1 determines the output DC offset, and they are independently adjustable. A little bit later you will see that we put potentiometers to do that. You can use anything to make the bias voltages – you can use a battery, in fact a 1.5V battery is just about the right voltage for that.

I did build exactly this circuit and the distortion curve looks like this:

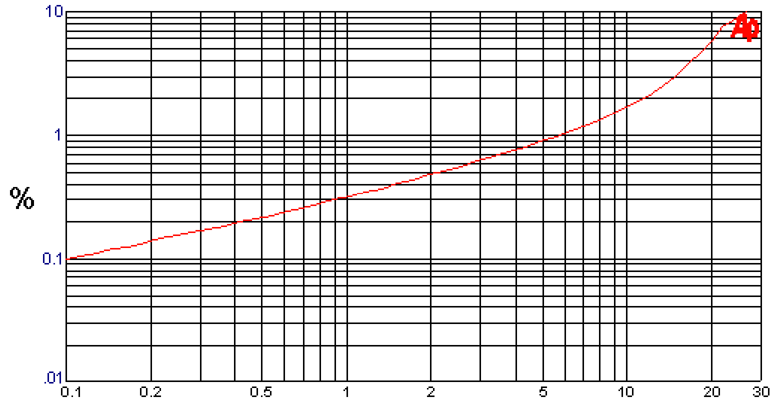

F6 OPEN LOOP DISTORTION VS WATTS

This is a “no-feedback” circuit. Down around a tenth of a watt it has 0.1% distortion, mostly second harmonic because the two transistors are not perfectly matched. As you can see, up at 5 watts it's running about 1%. That's fairly respectable as compared to “no-feedback” single-ended triode (SET) amplifiers.

This performance is *open-loop* (no loop feedback), and the output impedance of the amplifier is high – it's a current source amplifier and has a very small damping factor. It works fine – you can listen to it and blood won't come out of your ears. There are many speakers that will probably not appreciate this performance, but there some that will sound quite good.

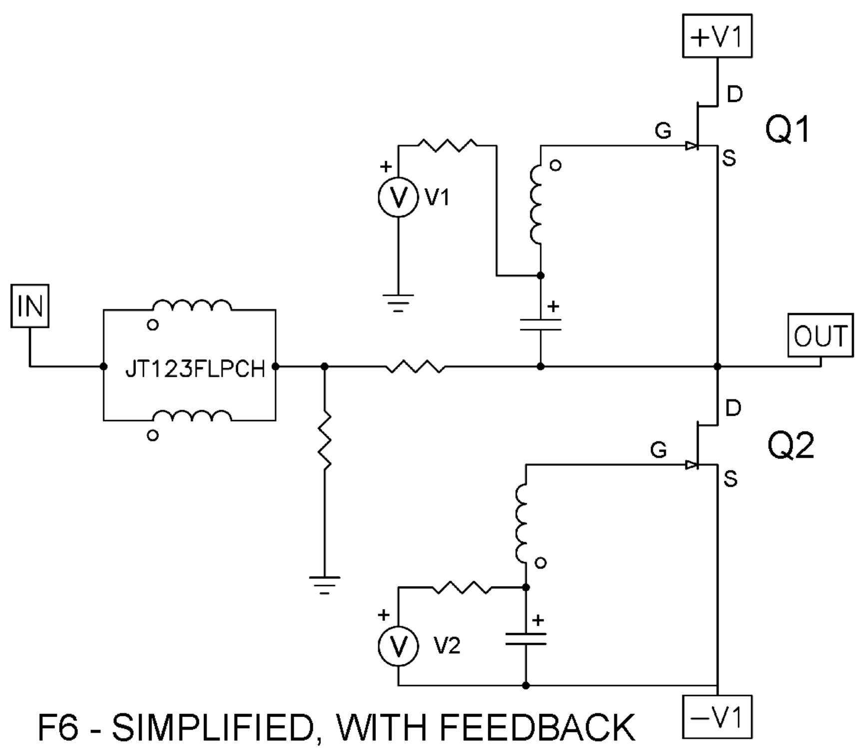

F6 - SIMPLIFIED, WITH FEEDBACK

Here's that same circuit, but I've done something new – I've put a feedback loop in it and I'm using the primary of the transformer as the feedback element. When I input a positive signal, the upper secondary drives the positive transistor Q1 to produce positive output, the other secondary drives the negative transistor Q2 oppositely, also producing a positive output, and this amplified output comes back to the primary windings through the feedback loop, where a resistive divider delivers the positive voltage at the minus side on the input coils of the transformer. This reduces the amount of gain, provides a damping factor for the amplifier, lowers the distortion and raises the input impedance of the amplifier.

So we've morphed from a more simple transformer with no feedback driving a not very DC-stable output stage into something where we have some DC control and we can set the gain with feedback and lower the distortion and so on. The open-loop figure I was getting for the previous circuit with no feedback was about 38 dB of gain, with a frequency response that rolled off at about 5 Khz or so. With feedback the gain is 15 dB with a bandwidth of 50 Khz.

Now for a little more detail on how those bias voltages get developed. Resistors bleed a small amount of current through LED1 and LED2. I used blue LEDs which develop about 2.7 volts. I need about 1.2 volts bias, so I put potentiometer dividers P1 and P2 for adjustment. This result goes through 10 Kohm resistors R4 and R5 to the drive coils, which pass the bias voltage to the Gates of the transistors. The 1.2 volts or so biases the amplifier at about 1.5 amps for Class A operation. To see to it that the secondary still presents a low AC source impedance to the Gate of the transistor, we couple it to the Source pin through large value capacitors C1 and C2. As a result, when we turn the amp on, it takes a little while for them charge up and the amp gets going.

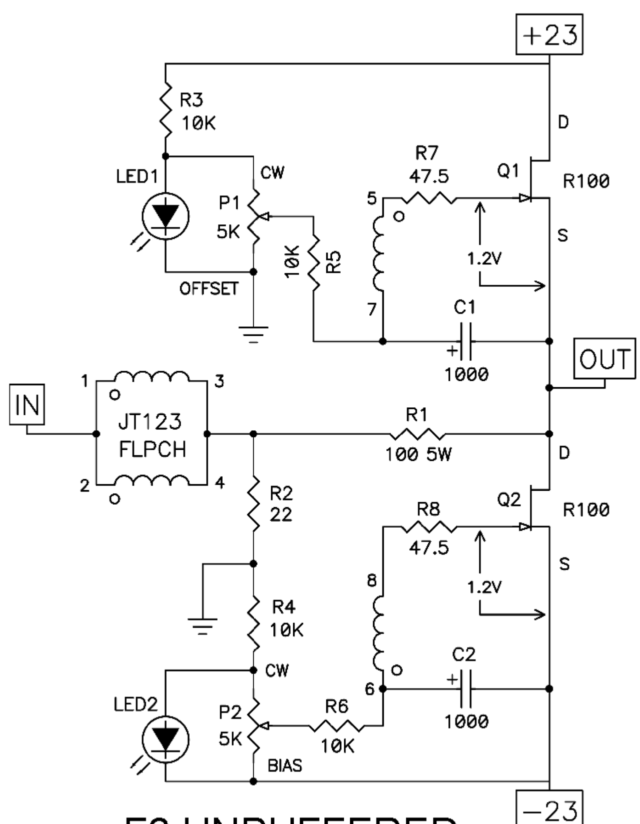

F6 UNBUFFERED

On the thread on DIYAudio, in advance of releasing the design, I had people speculating on all the things I might be doing, and so on. At some point they got extremely close (at the end here I have a version that gives a nod to a feature one of them came up with).

We've done some fill-in here, but it's still the same basic circuit – we are driving the input primary of the transformer as the input. With the feedback, the input impedance is about 20 Kohm, at lower frequencies – but it has a fair amount of capacitance between the windings, in particular the one seeing the output voltage of the amplifier, and so it's rather limited unless you have a low source impedance. If your preamp has an output impedance of 100 ohms or less, it's no problem, but if it's 600 or 1 Kohm, then it's going to be real soft on the top end.

This has a feedback figure of about 24 dB, that lowers the "open-loop" distortion by more than a factor of 10. It also delivers a damping factor which is almost identical to the amount of feedback. That 24 dB of feedback is equal to 16 times, and when I measured the output impedance of the amp I got 0.5 ohms, which is a damping factor of 16.

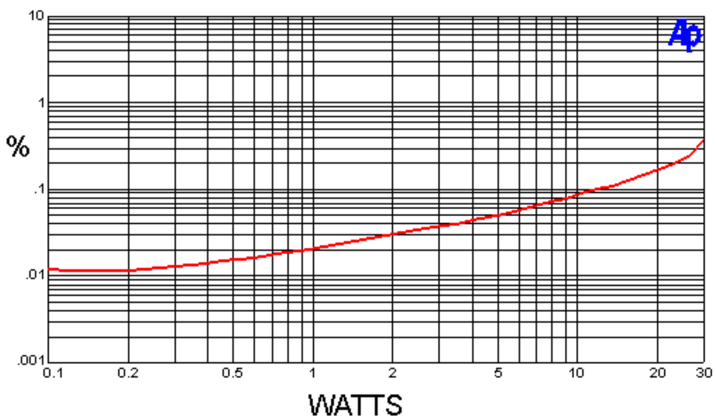

DISTORTION VS POWER @ 8 OHMS

The amplifier maintains a second harmonic character through the lower part of this region. You can usually tell a second harmonic character – the distortion rises as the square root of the power, so that if it's 0.01% at 0.1 watt, then it will be .1% at 10 watts.

Here is the distortion waveform as seen through the Audio Precision analyzer on a scope. You can see the two traces, the output signal and the remnants of distortion and noise left over if you subtract the pure part of the sine wave output signal, magnified for easy viewing.

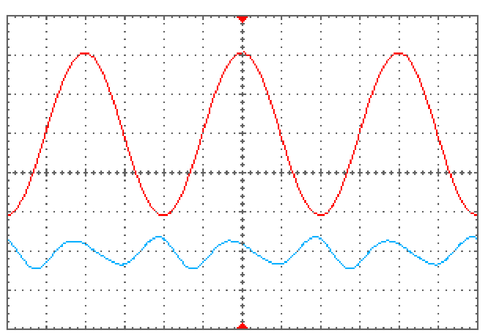

DISTORTION WAVEFORM @ 2 WATTS

In the waveform of the distortion component, what you really want is a nice smooth character and you want it as a low order harmonic. Looking at this you see that it is dominantly second harmonic (twice the frequency of the undistorted sine wave). There's a little expansion added to the top (positive) of the original waveform and a little compression to the bottom (negative) of the original. If you look at it carefully you can also detect a little bit of third harmonic.

The interesting thing about this is that second and third harmonic character correlate to a lot of people's listening preferences. I recall Jean Hiraga's comments about liking amplifiers which have a particular amplitude relationship between second and third harmonic (and no higher harmonics, probably), and it appears that he preferred it over purely second harmonic.

I don't think anyone has particularly improved upon that observation – if you're going to have some distortion, then this is likely how you would want it. I can tell you that many of my amplifiers that have done well in the marketplace start out at low levels being dominantly second harmonic in character, and with increasing the power you start seeing some more third harmonic, and somewhere below clipping the distortion is dominantly third.

We are talking about human ears here, and it's useful to remember that the ear is not a microphone and the brain is not a tape recorder. We have very complex neural networks that in many ways defy our efforts at simple analysis. A lot of my thinking on this subject is merely the result of observation - “this is what I like, this is what I perceive” and is difficult to describe further.

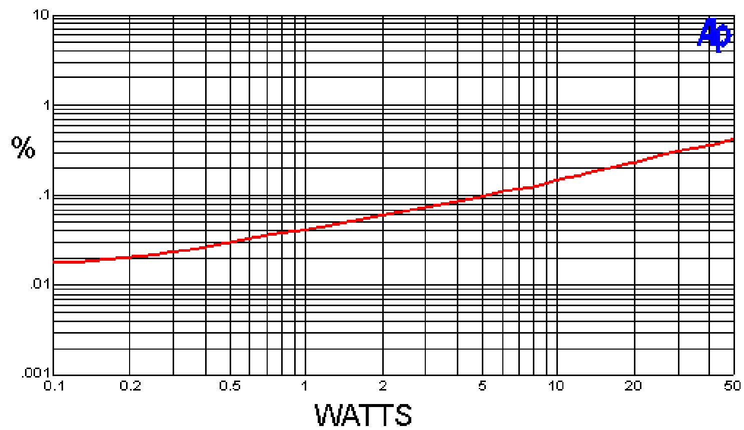

DISTORTION VS POWER @ 4 OHMS

Here's that same amplifier's distortion into 4 ohms. It actually does quite well – into 8 ohms it does about 30 watts, and into 4 ohms it gets up to 50 watts, making it better at driving lower impedances than many of my other little amplifiers.

Question from the audience: “Just to clarify, that's 50 watts out of one pair of SITs?”

These aren't SITs, they're regular power Jfets - SemiSouth SJEP 120100. The SITs are a more exotic form of a Jfet, with a character more like a Triode than a Pentode.

There are frequency limitations to transformers at the low frequencies where they lose some inductive coupling as well as losses at the high frequencies due to winding capacitance.

Here's the distortion vs frequency:

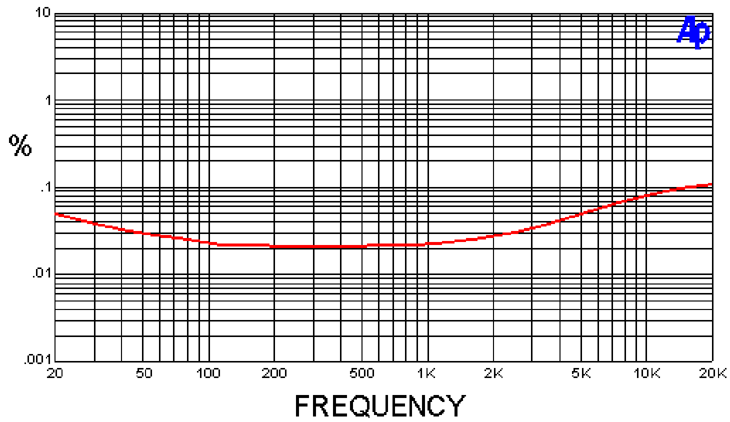

DISTORTION VS FREQUENCY @ 1 WATT

Here's the open loop and closed loop gain vs frequency of the amplifier. 38 or 39 dB of open loop gain, rolling off at 5 Khz. With feedback we get 15 dB gain, with about 24 dB of feedback, rolling off at about 50 Khz.

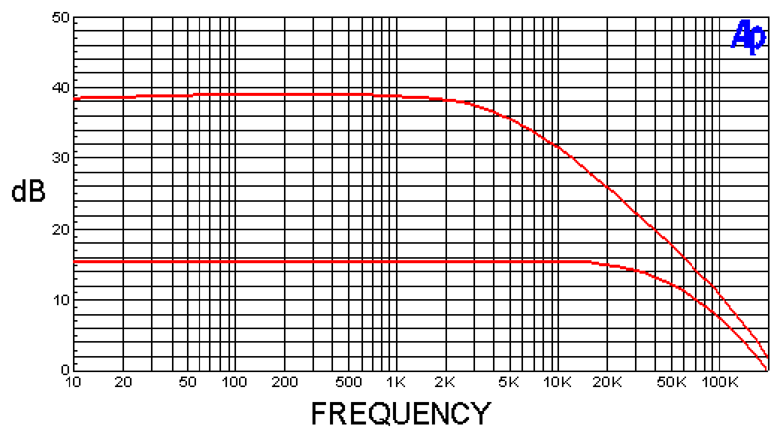

OPEN AND CLOSED LOOP GAIN

Here's the 5 Khz square wave at 1 watt. It's not real fast, but it's nice and clean, no ringing or anything nasty. It's not exceptional in this regard, but it is better than your CD player.

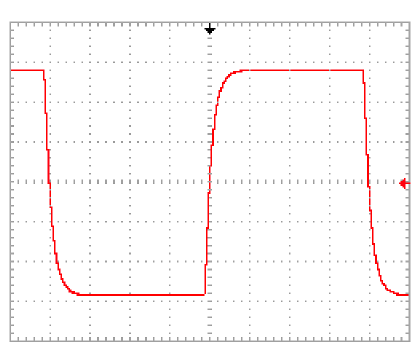

SQUARE WAVE @ 5 KHZ 1 WATT

A problem that we will now address is that the input capacitance of this amplifier is relatively high – it requires a low source impedance in order to get acceptable frequency bandwidth. 100 ohms is fine, and not much improvement is available if you provide lower than that. Some sources don't provide this, particularly "passive" preamps and a number of tube solid state and tube preamps.

Here is one solution to this issue (familiar to many by now), a buffer formed by a pair of Jfets Q3 and Q4 arranged as followers. They self-bias conveniently – if you tie their Gates and Sources together a known amount of current (their Idss) will run through them. We put a resistor to ground at the input so that there is no confusion without a source connection, and a 1 Kohm resistor in series with the Gates because that's just smart. Without an input resistor you would inevitably find some system or cable interaction or something that would create stability problems for you. The buffer provides a 47 Kohm input impedance for the amplifier and a 25 ohm or so source impedance to drive the transformer.

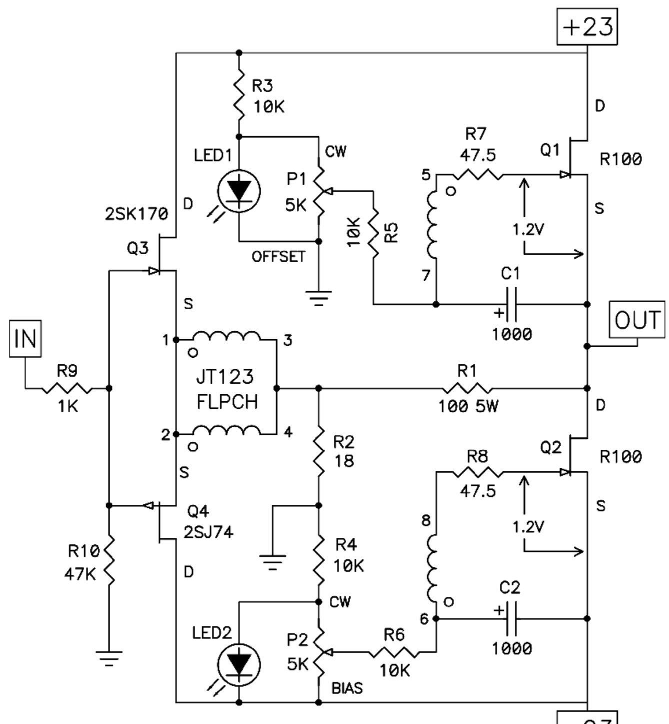

F6 BUFFERED

This is the circuit that we have in the working display down the hall. The buffer didn't make any significant difference in the measurements or the sound as compared to a low impedance source, but here is the option for when you need it.

There are those who would correctly point out that you can run lower distortion if you put some resistance on the Sources of the input Jfets because they would not then be running near their Idss figure. What happens there is that the Gates can start drawing a little current on peak inputs. I tried it both ways, and I didn't have any issues either way, so I left those resistors off. However, if you have Jfets with high Idss values then some resistances on the Sources would lower the bias current and thus the heat dissipation of these transistors.

In this circuit, the input Jfets were chosen for Idss at 8 mA or so, but you can get them with Idss up to 20 mA, which times the 23 volt supply will exceed their 0.4 watt rating. Or you might try some other Jfets with higher Idss or higher supply voltages. If you find that you need to *degenerate* the input Jfets with some resistance, you will find that there is little or no performance penalty for small resistance values.

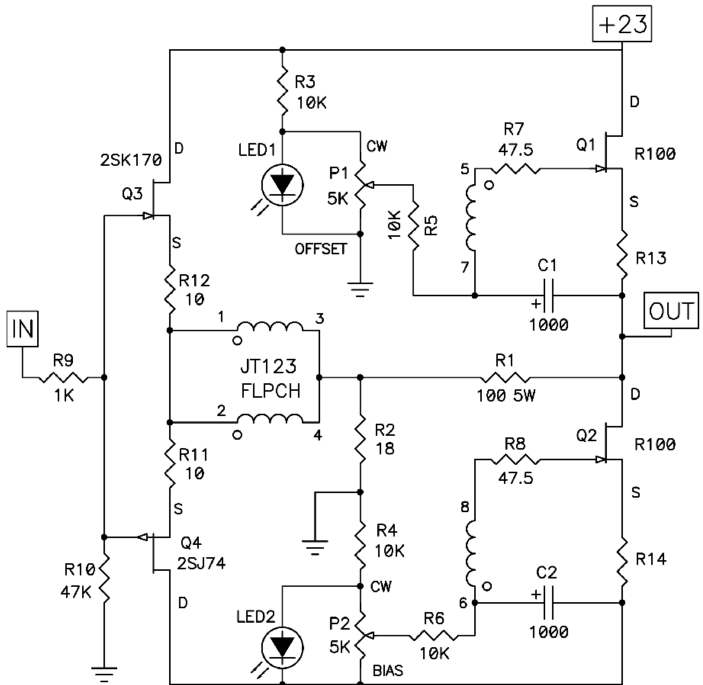

F6 - OUTPUT DEGENERATION

Here is the “second harmonic” version, showing the input degeneration I just mentioned (R11 and R12) and also a provision for output degeneration (R13 and R14).

It happens that the second harmonic character we have talked about is kind of random due to the imperfect matching of the parts. Typically you will see some measure of second harmonic in the output, and that is ordinarily what you would want. You might find when you build it that you didn't get the amount of second harmonic that you were looking for – maybe it's not in the proportion you want, or maybe you want to null it out altogether to get the lowest possible distortion figures, leaving only third harmonic.

These are all decisions you can make, and I always encourage people to try it – you might like some particular result. There are people who like purely third harmonic, and presumably there are people who prefer purely second, and people who like a slight mixture.

Recall that even if you design for second harmonic, you will tend to get more and more third as the power goes up. In an amplifier with dominant second harmonic rated at 20 watts, you will start seeing lots of third above 10 watts.

I have to add, down the hall we have a system with this amplifier and we are running an oscilloscope watching the output with what I call the “1 watt window”, where the top and bottom of the screen show instantaneous peaks at 1 watt, and the image freezes momentarily on peaks above a small fraction of a watt so that you can the waveform. An amplifier rated at one-half watt rms can just make it to the top or bottom of the screen. With music, this 1 watt peak will occur with average music levels of 0.1 watt or less. The remarkable thing is that with efficient speakers (these are 94 dB/watt) 1 watt peaks are really pretty loud, and sometimes you want to turn it down. Where I come from, 1 watt is usually enough.

If you are interested in adjusting the harmonic structure in this amplifier, there is a nice easy place where you can do it. I have added low value power resistors (R13 and R14) in series with the Source pins of the output transistors. Typical values for these would be 0.025 ohm to 0.1 ohm. If you leave R13 at 0 ohms and put 0.05 ohms for R14, you will find yourself with a positive-phase second harmonic. By this I mean there will be expansion as the wave goes positive and compression as the wave goes negative. If you have a negative-phase second harmonic, you could use this approach to null it out. If you leave R14 at 0 ohms but set R13 at a small value, you do the opposite, either creating a negative-phase second harmonic or nulling an existing positive-phase second harmonic.

The two phases of second harmonic do sound different, in fact the most consistent observation people have reported is that positive-phase has a little more projection to it, that it's a little more in-your-face and immediate. Negative phase tends to add more depth.

This is something you can play with, and moving on to the next slide, for this one we credit Patrick (EUVL) on DIYAudio. What he did was come up with a variable version – Why not put a pot on it? So here you can adjust P3 and P4 up and down, trimming the relative perfection of the symmetry, creating or nulling the the second harmonic.

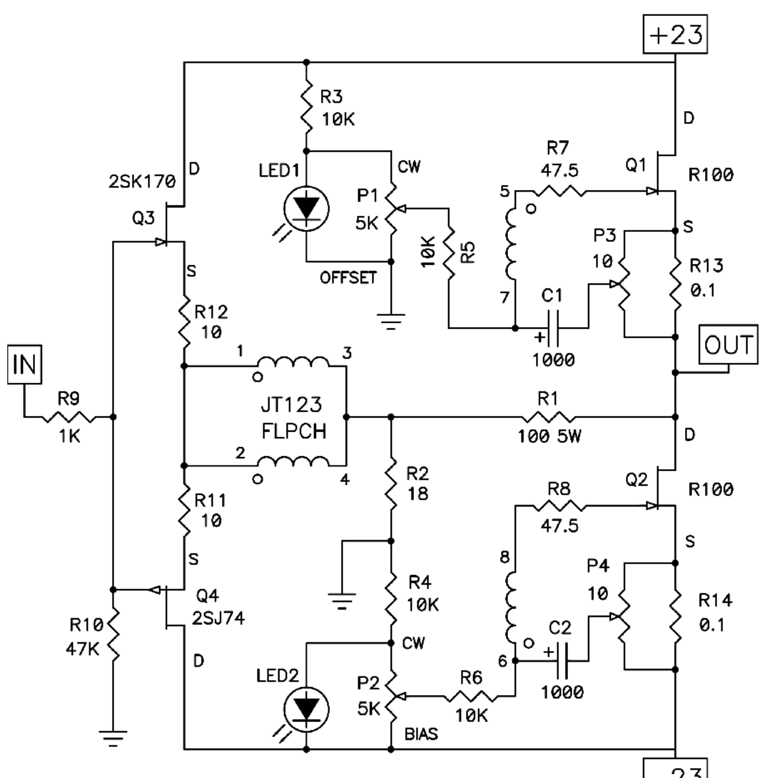

F6 - VARIABLE 2ND HARMONIC

By the way I numbers I showed you before were without either form of trimming, where there was generous second harmonic content. You can trim this to make the numbers drop dramatically, and if you are chasing the double-0 distortion numbers, this is a perfectly fine way to do it.

You can look at the popular F5 project at see a version where we did the same thing but on the input stage. It does, however, take a distortion analyzer to see these things. If you just want to listen to different settings, that's OK, but if you want it calibrated, a distortion analyzer is the way to go.

As an aside, whenever I use Fets I put a Gate-stopper resistor in the design to prevent parasitic oscillation. Here R7 and R8 at 47.5 ohms are used. You can often run these amplifiers without them, but it will not be reliable. On this amplifier I originally took the Gate-stoppers off because I like to remove all the parts I can, and only grudgingly put them back. Without Gate resistors it seemed to work fine, but I discovered parasitic oscillation during turn-on. For a few seconds while the output stage was biasing up the output stage showed oscillation, and then behaved itself when full bias was achieved.

While we are talking about adding resistors to pins, when you add resistance at the Source pin you alter the apparent characteristic of the transistor, and this is known as degeneration.

Degeneration is generally thought of as a form of feedback, but I always make a point of distinguishing between degeneration and loop feedback (which by the way we are also doing here – we have a feedback loop). You can degenerate this circuit at the input and output devices to control the bias current or adjust the gain, as we have seen above.

As a push-pull Class A amplifier we traditionally expect it to operate in Class A to peaks of twice the value of the bias current. With degeneration a circuit designed to operate at 30 watts rms into 8 ohms in Class A, would by definition leave Class A at 60 watt peaks. Here, the actual bias of 1.55 amps would be expected to give 3.1 amps peak output, a 76 watt peak into 8 ohms (with adequate supply voltage), and half of that (38 watt peak) into 4 ohms.

This calculation of the bias and the “Class A-ness” of the circuit depends on the output devices having significant Source resistance, which is generally the case. Fets and Tubes have what is known as a “square-law” characteristic, which means that as you raise the input voltage to the Gate, the Drain-to-Source current increases disproportionately. In a single-ended design this produces the familiar second harmonic. Using degeneration resistance tends to remove this effect by adding a linear component to this characteristic.

You can't imagine the number of emails and phone calls I get, wanting to know exactly at what wattage any given amplifier leaves Class A. Now I have to take the blame for some of this, having written an article called “Leaving Class A”.

(audience snickers)

Some people have become overly concerned about where an amplifier leaves Class A as if there's going to be a *Klunk* or some other noise that accompanies it, or something. However, in a highly-biased square-law circuit, this point is approached asymptotically, and there is no big discontinuity to talk about. Often you can't see it on a distortion analyzer waveform.

On the other hand, Class A is *Class A*, and we like as much as we can afford. In a circuit like this, we have square-law components that are working in high-bias complementary push-pull.

Jan Didden (here in the audience and publisher of Linear Audio) recently had an article about a square-law amplifier, where the major thrust of the design was “I can get you more Class A with less bias current”. In other words, you can make a Class A amplifier more efficient by exploiting the square-law characteristics of the devices. I've done some similar things – forty years ago my first commercial amplifier was based on the premise that everyone wants the Class A, but nobody wants the heat.

If you don't degenerate the output devices, that is if you don't use Source resistors, the square-law characteristic sees to it that you are going to get a wider band of Class A operation for a given bias. I took a measurement of the F6 down the hall (with no degeneration) which was was biased at 1.55 amps. With degenerated devices I would expect a bit more than 38 watts peak into 4 ohms. Without degeneration, it left Class A at 64 watts.

One thing you have to watch out for when you run devices without degeneration is the thermal coefficient of the transistors – generally they conduct different amounts of current at different temperatures, with the result that they move to a different bias point, and many, particularly bipolar transistors, will simply run away. They conduct more and get hotter and then they conduct more...

The R100's have a “*0 temperature coefficient*” in the region just above 1 amp, which is one of the reasons I can easily run this amplifier without degeneration. If you bias it at an amp and a half at operating temperature it will not drift enough to worry about. An additional benefit of operating at such a point is that it doesn't suffer the so-called “thermal distortion” that some designers have talked about. Not all devices out there are so well behaved.

There are people building (and selling) amplifiers (most of them bipolar it seems) who swear by the idea of not having emitter resistors (the equivalent of Source resistors on a Fet or Cathode resistors on a tube) and they're saying things like “The dynamics are incredible! It's the best sound I ever heard!” for all of ten minutes....

(audience laughs)

If you want to perform harmonic content adjustment without degeneration there are a couple options. First you can carefully select the parts for transconductance to favor the desired harmonic content. You can also take some value of resistance and place it across the secondary coils of the transformer so as to load one slightly. They can take as low as a few hundred ohms without too much of a performance hit.

For me, it's just a matter of picking the slightly higher-transconductance part, putting it on the positive side for positive-phase second harmonic, or on the negative side for negative-phase.

Lastly of course, we can note that degeneration values at 0.05 ohms or so are not so high that you don't get a good deal of square law action anyway, so you can still do this and enjoy an expanded Class A region, though to a lesser extent.

So, That's our circuit.

It has a power supply, and here is the mono version:

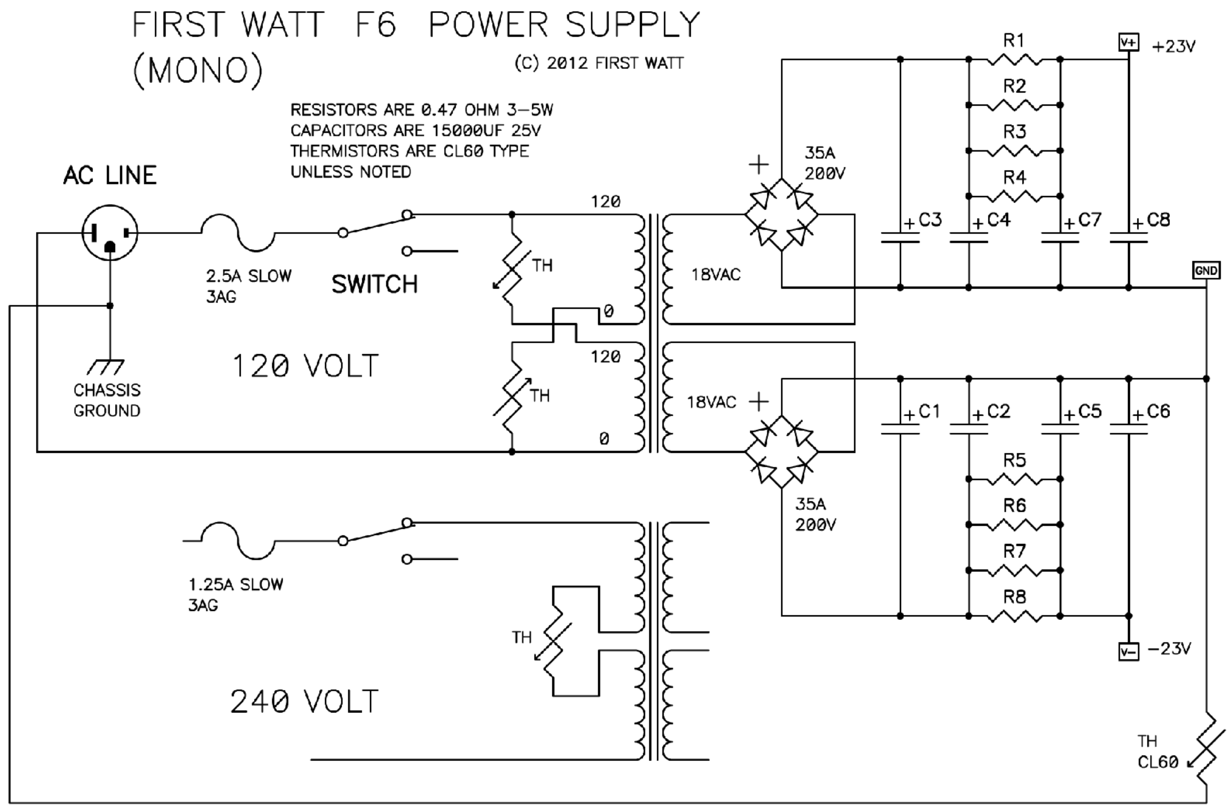

You've seen this supply many times, and here it is again. It's got 18 volt secondaries that you can get from Antek or Plitron. You can wire it for 240V or 120, its got a thermistor isolating the analog ground from the hard chassis (earth) ground, giving it a safety connection but at the same time having enough resistance to discourage ground loops. The secondary uses CRC networks to reduce the ripple noise.

You don't have to do this stuff, this is just a sample circuit. You can put coils in there for a CLC filter if you want, or nothing at all. With a little decoupling I'm getting noise figures on the order of a hundred microvolts or so. I consider that to be quite good. I have seen better, but I rarely achieve it in amplifiers with more than 25 watts or so.

People often ask about how much capacitance should be on either side of the filter, and the answer is that there is no hard rule. I usually do half and half (the Even-Steven rule) but you don't want to be stingy on the first half, as the ripple currents are much bigger than you think – like 10 times bigger. I see examples of amplifiers which still work fine, but the heat shrink plastic sleeve on the capacitor bodies has continued to shrink from high temperatures until it's a belt. That tells me that it's been running too hard because it's taking all the ripple current - short-duration high-current pulses. The energy dissipated is proportional to the square of the pulse size, divided by the pulse length, so you want to make a serious allowance for it.

By the way, you can use the voltage drop across the resistors to measure the bias current of the amplifier.

Here is the stereo version:

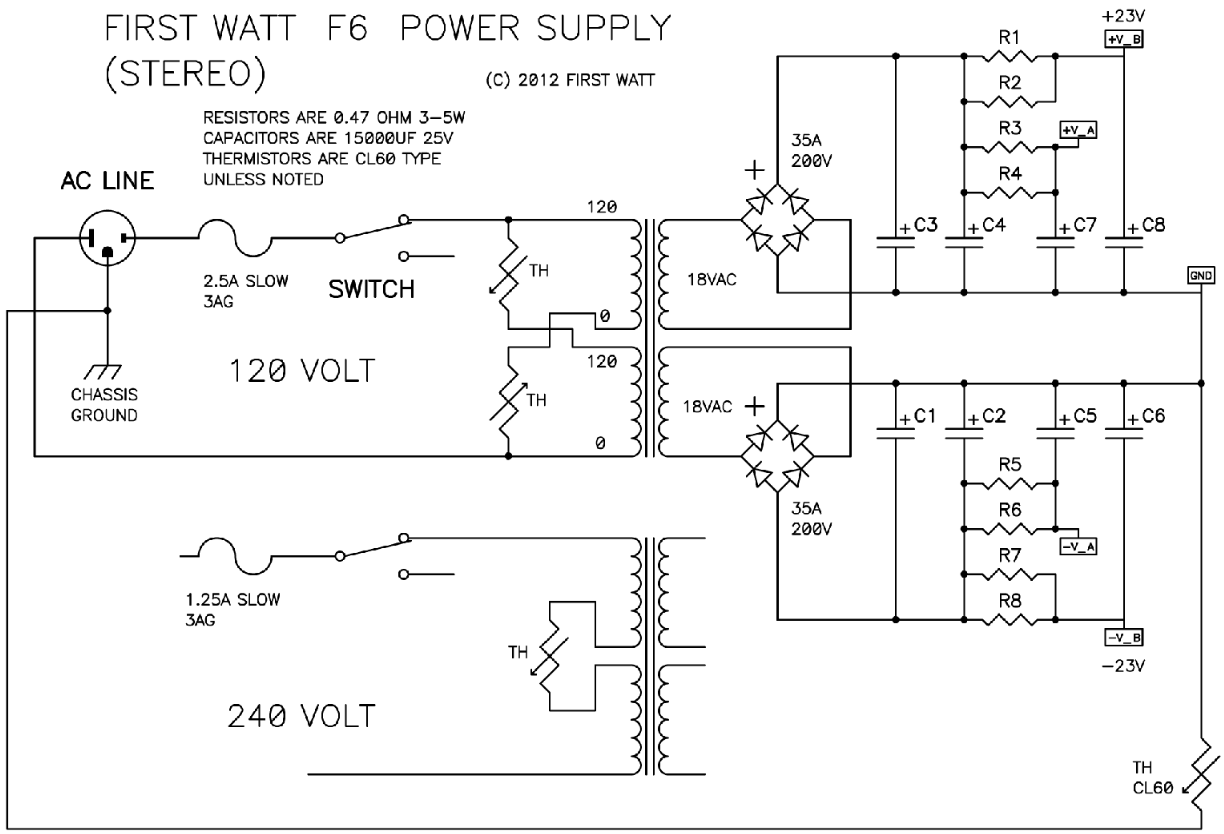

It's the same thing except that you can split the RC networks out to each channel separately. These sorts of filters are very effective in lowering the noise, particularly the higher order harmonics of the ripple noise. You're free to put large values of capacitance here, it doesn't bother me at all.

Here is the actual amplifier interior *before* I got it to work properly.

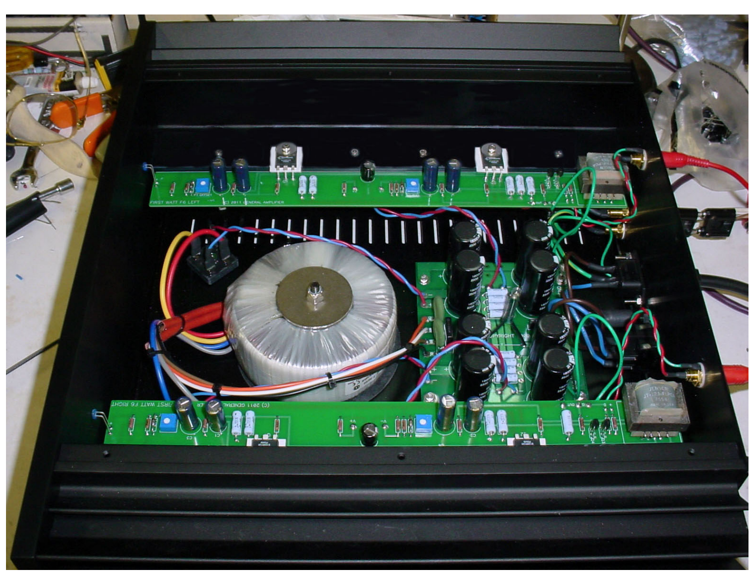

Unless I'm mistaken, this is end of the presentation. Any questions?

Q: "The degeneration resistors, the most common kind would be the sand filled block things. I think they use some nichrome wire. I think I've measured distortion in those before, do you just recommend those or any other material?"

A: I use metal oxide film types. You can go get very snazzy resistors, and at some point you're spending a lot of money but getting diminishing returns. The whole thing of eye-candy parts is great. I have no objections to people using gold-plated anything, and there's some really nice parts out there. And I'm the last person to argue that they're not better. What they are for sure is more expensive. I'll be honest: One thing I can tell you for sure about DIYers is that they're really cheap guys.

(audience laughs)

"But Mr. Pass, that transformer costs $30! I can't afford that sort of thing!"

OK, you can understand why I don't automatically point people to expensive parts. As a default I pick cheap/available parts (read: crummy) for these projects. When I'm asked what calculation led me to a 220 uF value for a capacitor, I explain that I have thousands of them on the shelf, and when you look at the wire on my display today, you will see nice quality clean copper wire from Fry's. It works fine, and I should say that I also have some very nice wires that I think work a little better.

The way I look at it: If you build one of my amps from the cheap parts and it sounds good, the design gets all the credit   :)

If you want take it up a notch, that's great and I have no argument with that. I don't even consider it a waste of money, and I personally spend a lot on my own toys.

Q: "About the power supply, it turns out that Duncan Amplifiers makes something called PSUV and its originally intended for tube amplifiers, it turns out it does a wonderful simulation of these kind of voltages and it shows everything, it shows the current pulses,whatever you want, and its a wonderful simulation tool and its highly recommended."

A. The other thing you can do is download a copy of LT Spice which is a little bit of work (I'm really bad at Spice, but I can get what I want by beating on it) and you can simulate anything, and it will tell you a lot about power supplies.   Perhaps the easiest of the them is the free MicroCap student edition - it's nice and dumb and easy to use. And the last time I looked John Curl was still using it, so I'm with him...

(more audience laughter)

Q: "The specs on British and American power amplifiers are almost identical, very similar, but the tonal character are differentiated because the Brits have the windings start on the opposite end of the coil that flips the polarity. You say that if we do a little bit more second order harmonic on the top or bottom side it changes the character. Have you tried flipping the polarity on the speaker and see how that...."

A: Absolutely. One of the most fascinating things is the whole thing about absolute phase – how much can you hear absolute phase? Not whether your speakers are in or out of phase with each other, but absolute phase. It has become quite apparent that it matters. It matters in the context of how you look at the second harmonic structure of the amplifier as it relates to the speakers and to what comes out of the recording process. I mentioned that positive-going phase for second harmonic has a particular sound. If you flip that characteristic you've got a different sound.

I'm not here to tell you what you like, I've noticed that when you get reasonably experienced listening to that effect you can then go through your record collection, often deciding which recordings are in-phase and out-of-phase. I find that totally fascinating.

It relates to something I can talk about briefly, one of my favorite soap-box subjects. If you go into the literature of psychoacoustic perception, there is a very good book by Diana Deutsch at UC San Diego called “The Psychology of Music” and in it there are several chapters talking about how the low level neural networks of the brain take the data from the ear and what they do with it like the bureaucracy at the DMV, and you have an army of these things and for each of them the job it to make a decision – what goes with what, and these are called grouping mechanisms. Each bit of the network takes disparate bits of audio information and decides whether they go together or not. The system is sensitive to such things as loudness, timing, pitch, harmonic structure and phase. Decisions are made at very low levels and then get passed upwards for increasingly more abstract decisions until the final result, the “executive summary” is handed to the guy who sits behind your eyes at the control panel and imagines he's in charge.

So what are we doing when we play with the distortions of an amplifier? Well, we're just fooling ourselves, fooling the ear and the brain. And sometimes that's a good thing. It's plausible to me that if you tag the sound with a particular characteristic (I'm not claiming that expertise) it seems to slip more easily through these neural systems like poop through a goose, and the decision-making process is easier. There is a lot of *work* going on in the brain when we are talking about listening – a vast army of neurons working this thing, and if you make their life easier, they aren't working as hard.

We are talking about listener fatigue, talking about people who get tired after a half hour and shut the music off versus guys who go through their entire record collection all night long. We are literally talking about fatigue – the brain gets tired.

So why do we try to fool the ear? It makes people happy. It helps them to relax while they listen to music and try to forget all the terrible problems in the world. I'm not here to deliberately create distortion, but if my simple little circuits are going to have some distortion anyway, I can at least try to organize it the way I want.

Perhaps you say that it's not accurate? I say it's entertainment!

(thunderous applause)