KISS — Keep It Simple Stupid

Designing a 2A3 Push-Pull / Interstage amplifier

The goal of this project was to find out how far you can get without leaning on perfect measurements, complex topologies, or feedback. What does that restraint do to the sound? And the question lurking behind it: are a lot of other designs simply over-engineered?

“Keep It Stupid Simple” became both the working title and the design philosophy. One voltage amplifier, one interstage transformer, a pair of output tubes, and nothing else in the signal path. As you’ll see, “simple” turned out to be a slippery word.

The finished amplifier: proudly on the living room’s cabinet

Inspired by…

I built this amplifier around my own take on the Klipsch La Scala. It's a speaker I love, not least because it is mercilessly revealing: every small change anywhere in the system comes straight back at me, loud and clear. That makes it an ideal development tool. And a harsh judge.

Its designer, Paul W. Klipsch, famously drove his speakers with a Brook 12A: a push-pull 2A3 amplifier. I had used the 2A3 in an earlier project and was very happy with it, so following Klipsch's lead felt natural. The Brook doesn't use an interstage transformer; it drives the 2A3 grids through anode chokes. Not quite the same approach, but it also puts iron in that part of the circuit, and that idea stuck with me.

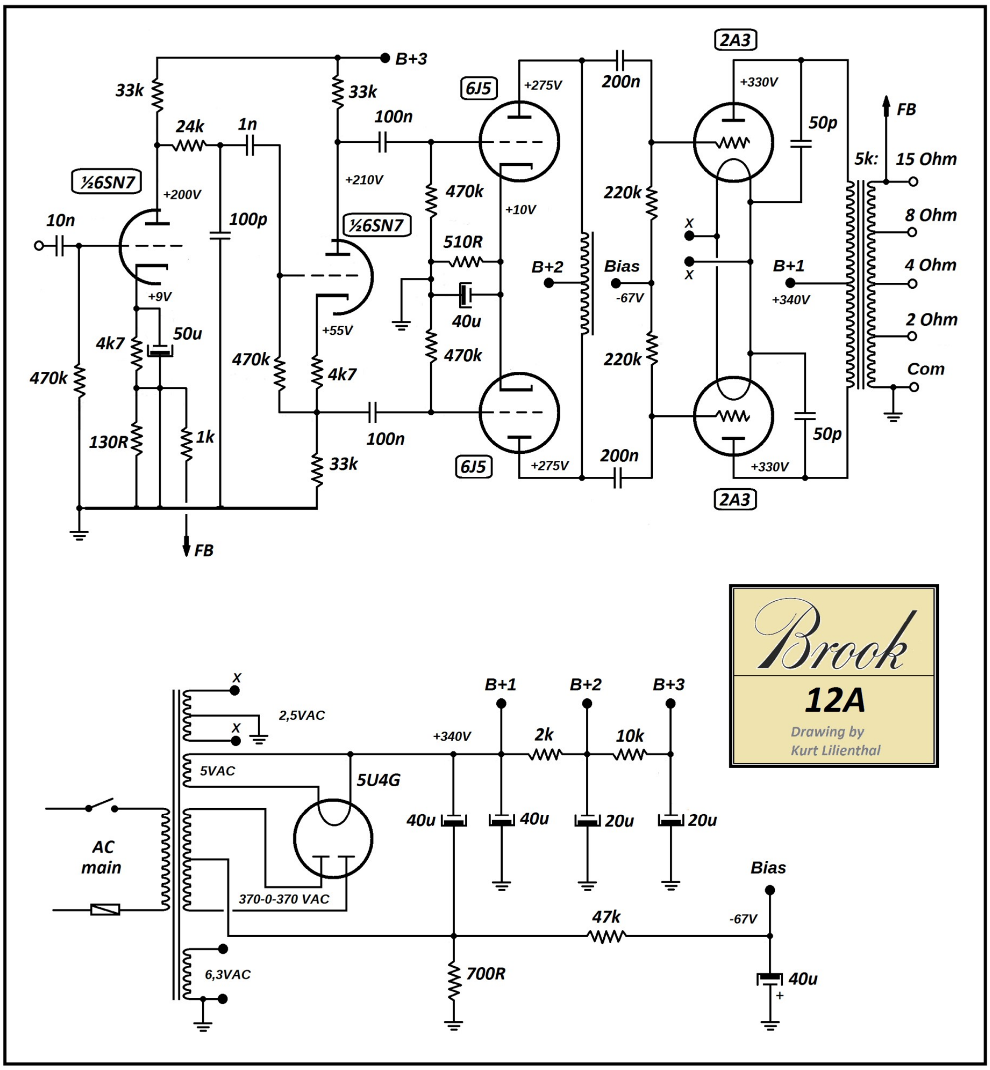

The Brook 12A Amplifier - 2A3 Push-Pull

Why interstage?

I had never built an amplifier around an interstage transformer before, and the idea appealed to me: with a single component I could solve several problems at once. It handles phase inversion for the push-pull stage, the coupling between the driver and the output tubes, and the bias path, and it does this without a coupling capacitor in the signal path. One well-chosen component doing the work of half a schematic fits the KISS idea well.

Why 2A3?

The 2A3 is a directly heated triode with a deserved reputation for a beautiful, natural sound. In push-pull, a pair of them yields a little more power and the usual benefit of cancelling even-order artefacts, while keeping that direct-heated character. Together with my earlier good experience and the Klipsch/Brook heritage, it was the obvious output tube for this build.

Why D3a?

With the topology fixed, the driver becomes the critical choice. To swing the interstage transformer properly it needs both enough gain and real current drive, and there aren’t many tubes that offer both. The D3a (and its equivalents) does. It is also well regarded in the online community, and I always try to weigh those subjective impressions alongside the datasheet numbers when making a design decision.

The schematic

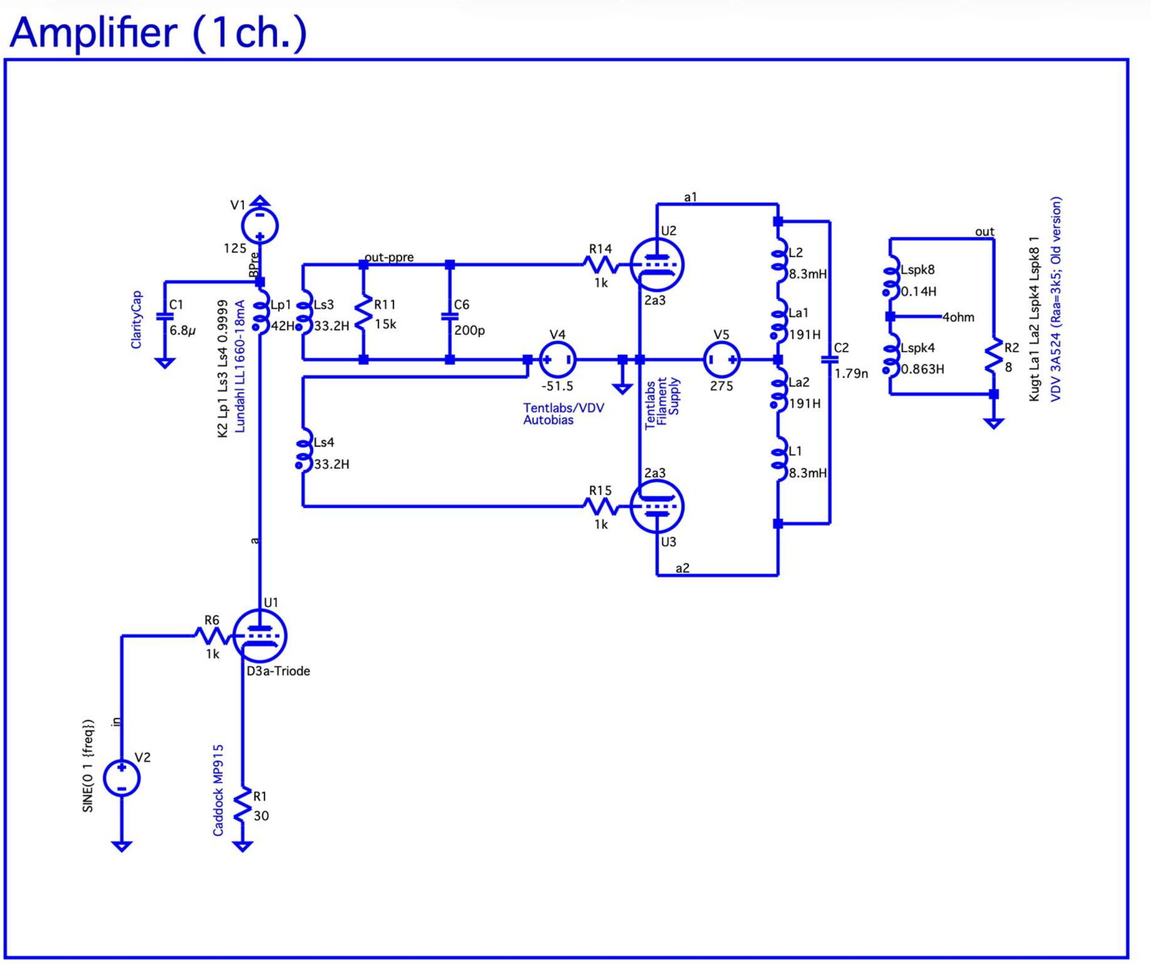

- D3A biased at B+ = 125V. Ia = 26.5mA, Vgc = 800mV
- Lundahl LL1660 interstage in modified Alt-V configuration (each winding having its own bias supply)
- One secondary loaded with 15K + 200p (silver mica) for optimal balance
- Tentlabs/VDV autobias module on each of the four 2A3 tubes
- Tentlabs Filament supply on each of the four 2A3 tubes
- 2A3 biased at 50mA, B+ = 275V
- Output transformer: VDV 3A524 Raa = 3.5K

Simple, but not so simple

On paper the concept could hardly be simpler: driver tube, interstage, output tube. But it pays off enormously to make everything *around* those stages as close to ideal as possible, and that is where the simplicity quietly disappears.

Power supply and rectification

A push-pull amplifier may, and probably should, have a robust supply. As the stage moves from class A into class B, the current draw can swing hard, and the supply has to keep up without complaint. On top of that, the driver and the output tubes each want their own supply voltage. I chose Jan Didden's T-Reg for the job.

Earlier in the year I had organized a demo in which we switched between a rectifier tube, ordinary 1N4007 diodes, and UF4007 diodes. The difference in sound was startling. The rectifier tube sounded calm and musical; the 1N4007 sounded hard and unpleasant; the UF4007 sat somewhere in between. It made painfully clear how large (and how audible) the influence of diode recovery time really is. So I applied a trick from Blencowe's book: place one good diode with a favorable recovery time in series after the rectifier bridge, so that it dominates the recovery behavior of the whole supply. The T-Reg itself is a very sophisticated, thoroughly worked-out regulator. Yes, the amplifier schematic is KISS. The supply modules? Not in the slightest.

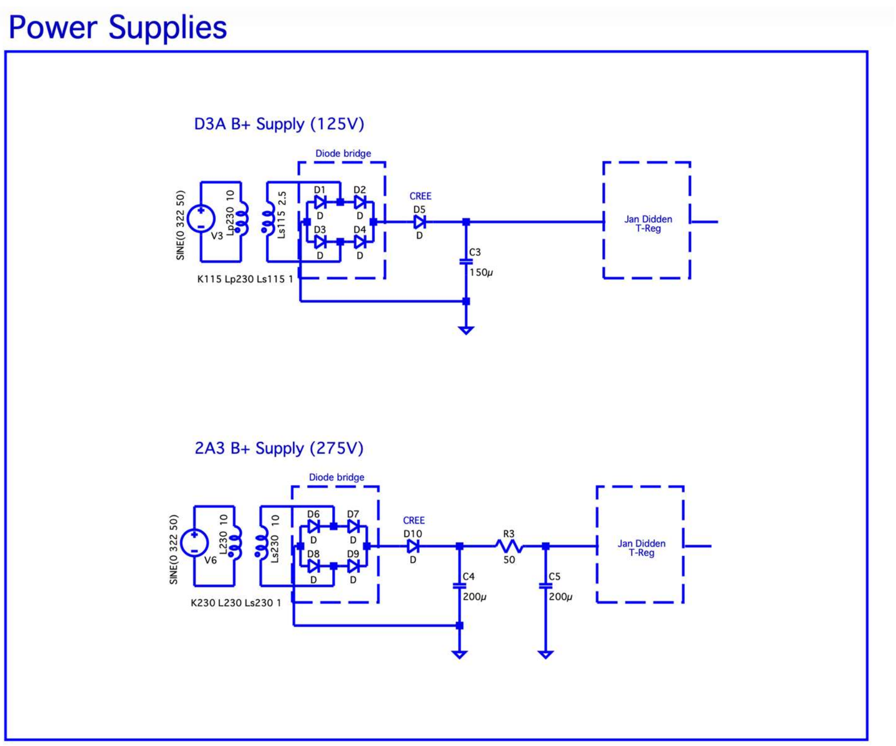

Dual power supply using Jan Didden's T-Reg

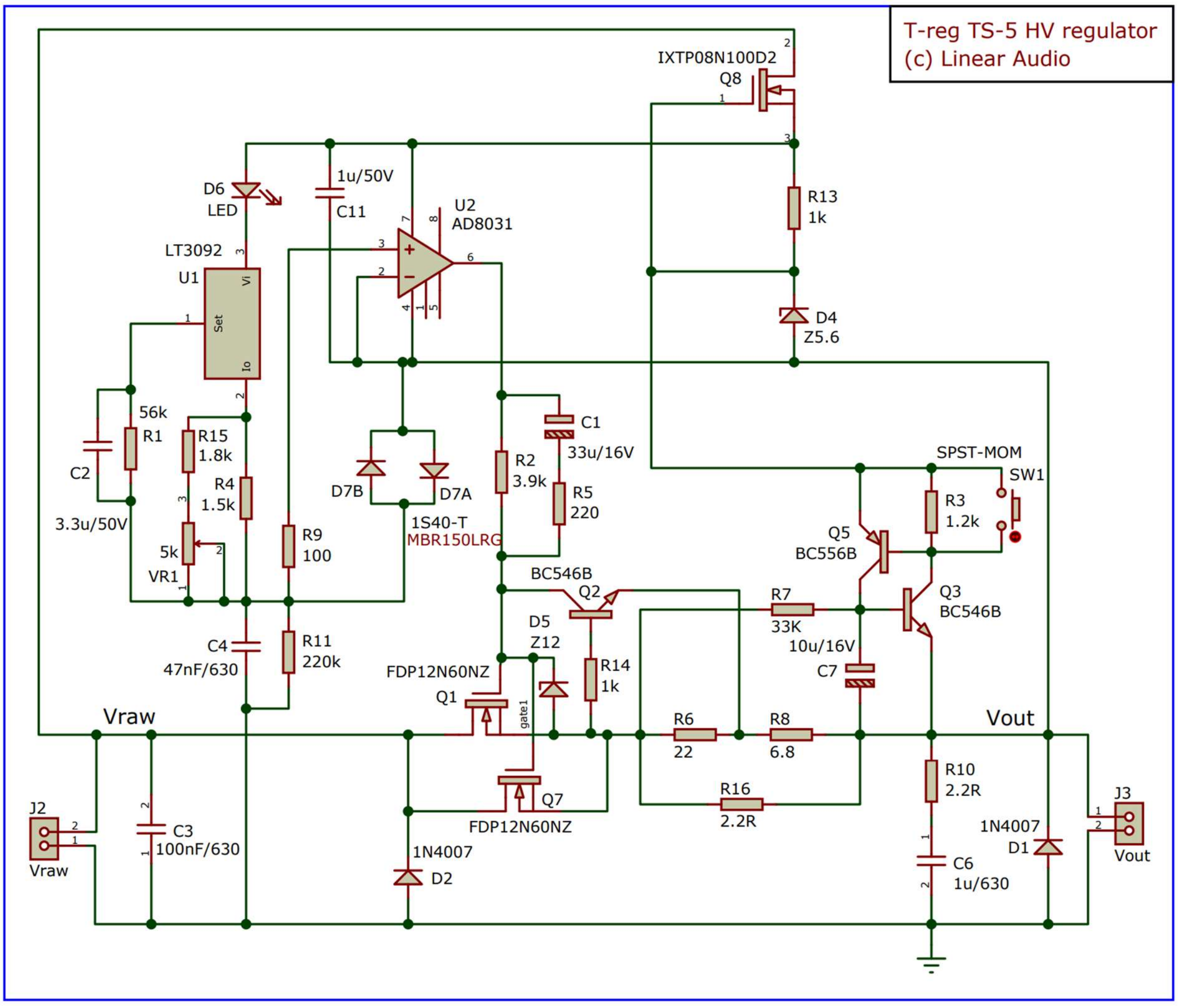

T-Reg schematic

Current-source filament supply

Because the 2A3 is directly heated, the way you power its filament matters a great deal. Guido Tent investigated this and concluded that a (voltage-controlled) current-source approach works best (see Tentlabs). I decided to use his regulators. The catch: this required four separate 4 V windings, which you simply don't find on a standard transformer. With some advice, I wound the extra filament windings myself onto the power transformer. Still KISS? In spirit, perhaps, but a dedicated winding per tube, the considerable board space the regulators take up, and the design effort to make it all work are not exactly a picture of simplicity.

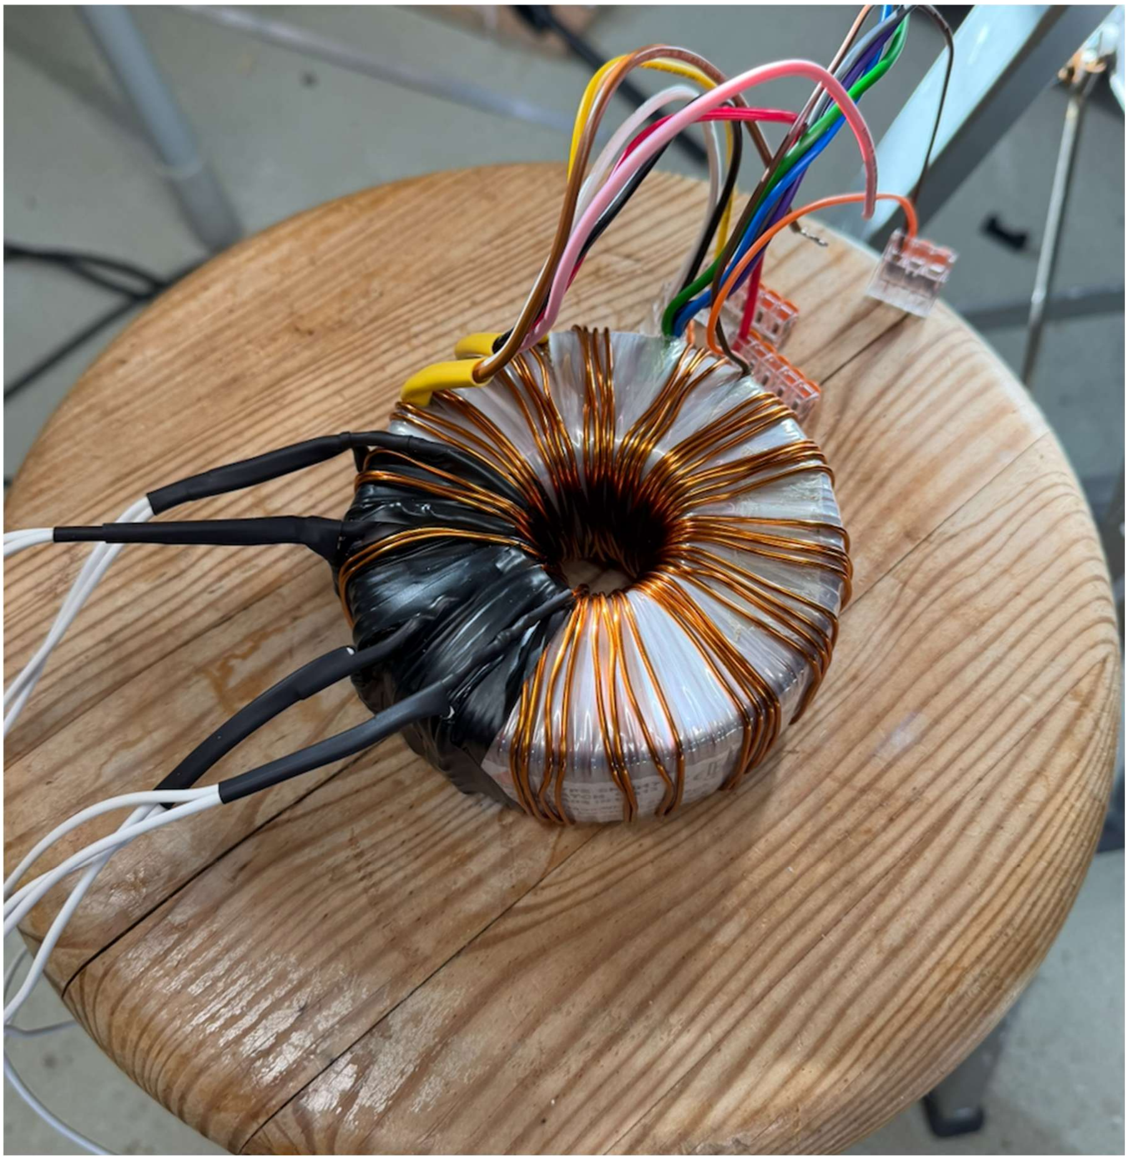

Adding windings to the power transformer

Autobias

In a push-pull design, a near-perfect balance of the quiescent currents is essential for a clean distortion signature. To get this right I used the autobias module developed by Menno van der Veen and Guido Tent. It senses each tube's current through the cathode and corrects continuously. Importantly, it is designed with class B operation in mind: when one tube is driven into cut-off its current is, by definition, different, and you don't want the autobias to “correct” that asymmetry and slowly drift away. Once again, a module with a great deal of thought and complexity behind it. Not exactly KISS.

Add it all up: the regulator boards, a filament supply per tube, the autobias electronics, and the wiring loom that ties them together. The space all of this takes up inside the chassis tells the story. A KISS schematic can still turn into a build that is anything but simple.

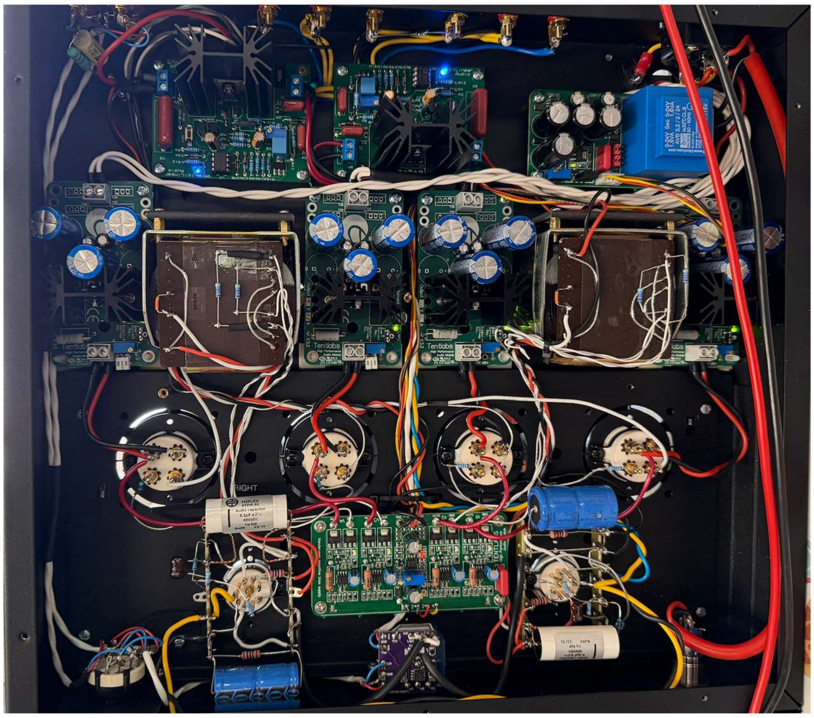

Picture of all parts in early stage development. KISS?

The interstage transformer

Now for the heart of the project: the interstage transformer. Before getting to what I built, it helps to set out what this transformer has to do, and what each requirement brings with it.

What it has to do

The first requirement is a gapped core, because the drive is single-ended. Not every interstage transformer has an air gap, and here that gap is essential. The D3a sits with the primary hanging directly off its anode, so the tube’s DC current runs straight through the winding. On top of that, the 15 kΩ secondary loading resistor (more on that below) reflects a relatively low Ra back to the D3a, which means the tube has to be biased at a fair amount of current to drive it. My first estimate put that around 20 mA, and the transformer has to carry it without saturating. That is exactly what the air gap allows.

The second requirement is two fully separate secondaries, for phase splitting. For a push-pull output stage I want two identical secondary windings that can be driven in opposite phase; that is what turns the transformer into the phase inverter. Just as important, the two windings have to be galvanically separate, because each one feeds its own 2A3 grid and therefore has to sit at its own bias voltage.

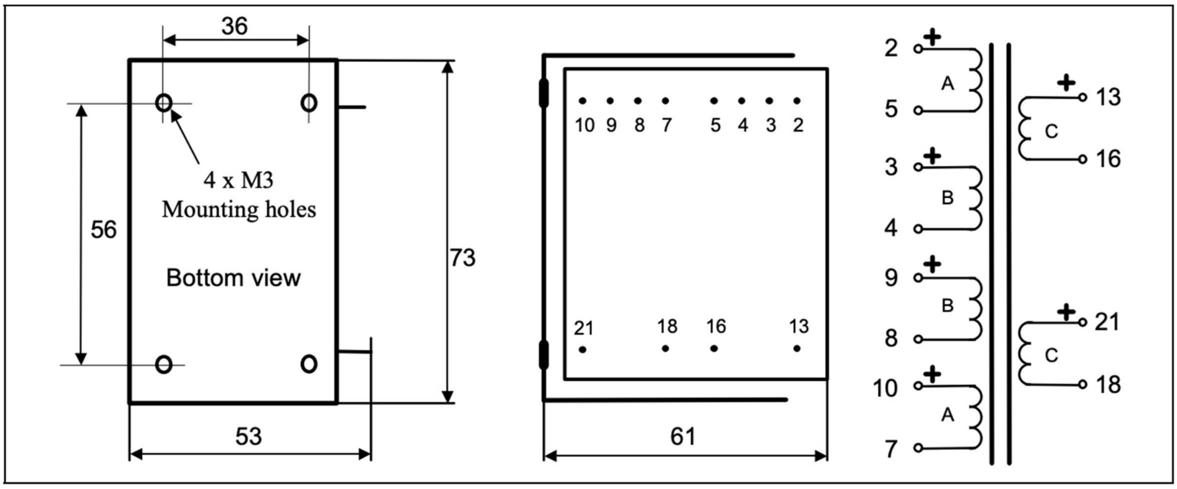

Weight	Turns ratio	Static resistance, winding A	Static resistance, winding B	Static resistance, winding C
0.75 Kg	1+1+1+1 : 2.25+2.25	315 Ω	240 Ω	625 Ω

Lundahl LL1660 interstage transformer

Loading, resonance and the balance problem

A real transformer is never just an inductance. It also has winding resistance and parasitic capacitance, and left unloaded those combine into a peak (a resonance) in the frequency response. A well-chosen loading resistor on the secondary side damps that peak. The catch is that the parasitic values aren't something you can simply read off the datasheet, so finding the right load is a matter of measuring and finetuning.

Ideally the two secondaries are as identical as possible. An equal frequency response, and equal phase behavior, is what guarantees that the second-harmonic cancellation stays optimal across the whole band. Even a small mismatch already changes the distortion behavior at the higher frequencies. And this is where an important lesson hides. In the Lundahl LL1660 the windings are physically stacked on top of one another, which means one secondary inevitably sits closer to the primary than the other. By definition, then, each winding has slightly different parasitic properties. That asymmetry showed up clearly when I measured the frequency response of the two windings separately, loaded with the 100 kΩ input impedance (plus input capacitance) of the measurement gear:

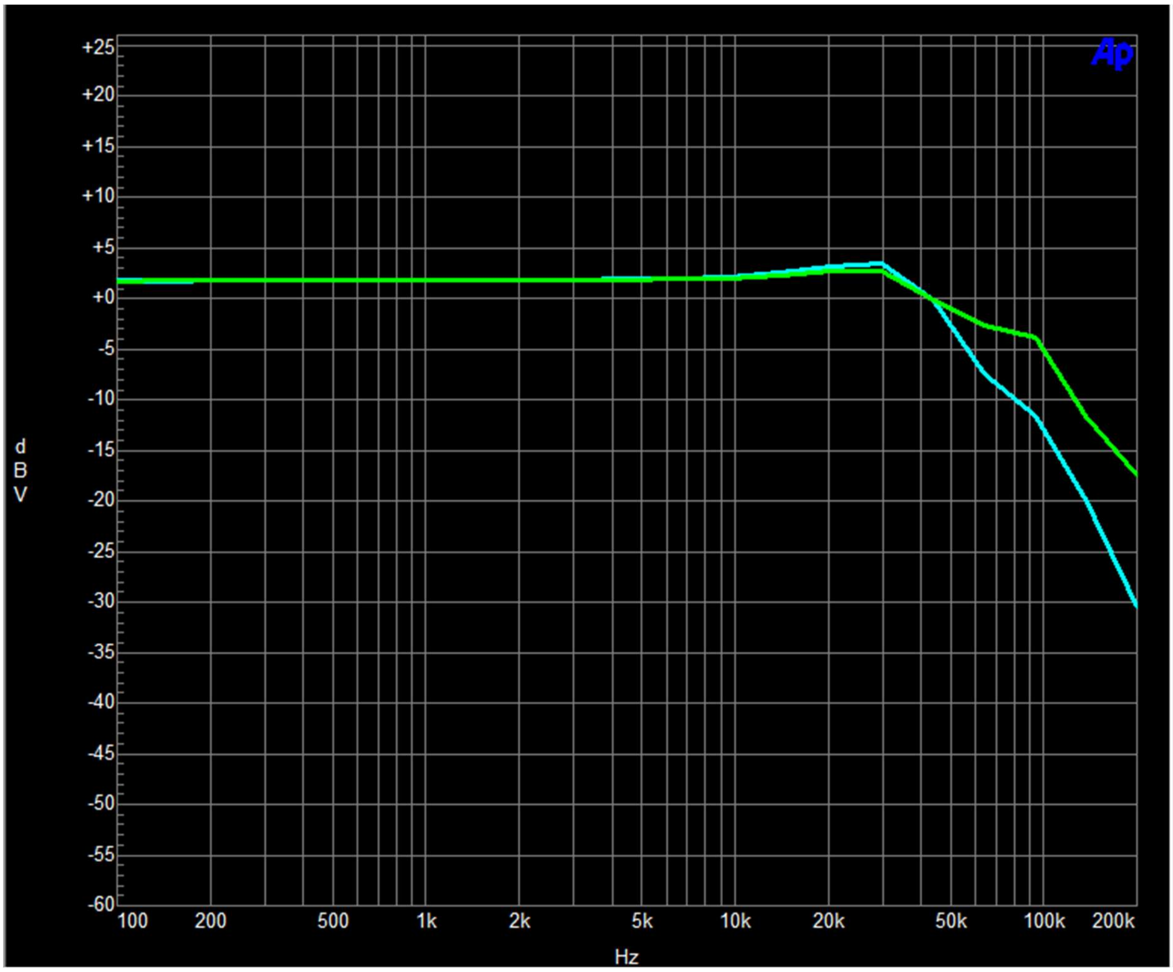

LL1660 response. Driven by d3a. Cyan = inverting winding; Green = non-inverting winding

After a good deal of email back and forth within the study group, a lot of reading, and plenty of experimenting, I found that the best balance comes from loading only one of the two secondary windings with 15 kΩ plus a 200 pF silver-mica compensation cap. The other secondary is left floating. The measurements below show clearly how much that compensation improves the distortion. Both measurements are at 1 Watt into 8 ohms, connected to the 4 ohms tap of the output transformer. The first measurement uses a 12 kΩ loading resistor on each of the 2 secondaries. The second measurement uses 15 kΩ + 200pF only on one secondary. Notice that H2 distortion at 10 kHz has decreased from a problematic number around 1.15% to an acceptable number around 0.25%.

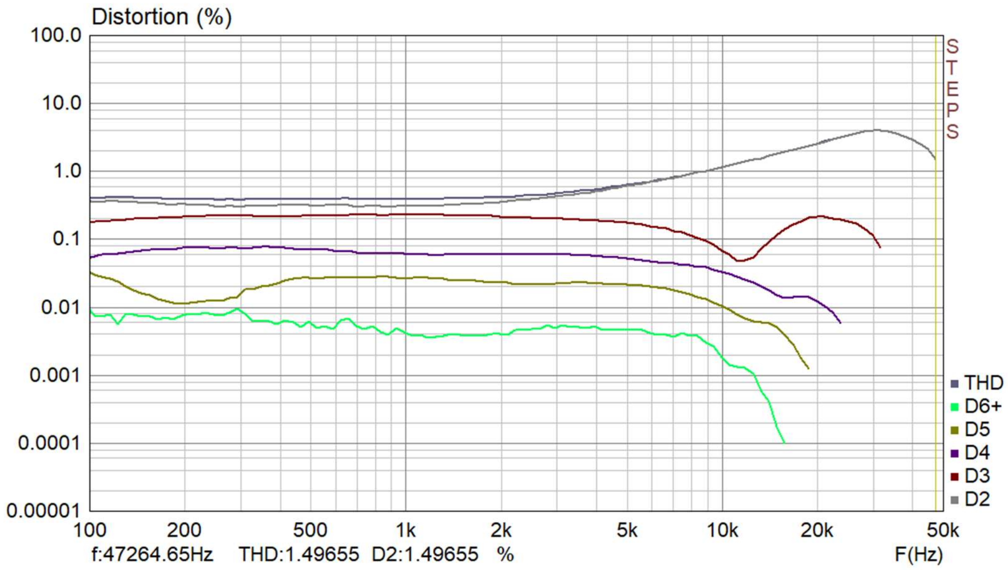

Loading: both sides 12K

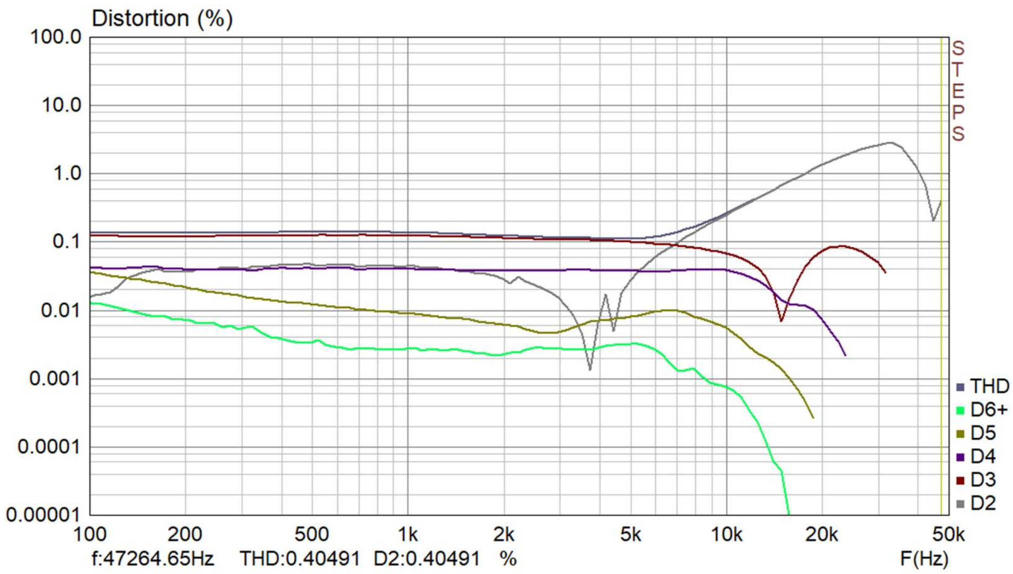

Loading: 1 side 15K + 200p

The LL1660S version

You learn by doing. Lundahl also makes an “S” version of this transformer, developed specifically for this kind of job: getting the balance between the windings right. The S stands for shielded. It places a grounded shield between the primary and secondary windings, which reduces the parasitic capacitance. Unfortunately the S versions didn’t arrive in time for the demo; in the 18 mA version they are made to order. Once they come in, I want to measure the parasitic capacitances of each winding properly beforehand and use LTSpice to work out a smarter secondary load.

Tuning the D3a’s operating point

The secondary loading resistor sets the Ra that the driver tube sees, and that in turn sets its ideal operating point. But the tube’s internal resistance Ri also shapes the transformer’s frequency behavior. Ri is itself then partly determined by the operating point. So this is a subtle balancing act: pull on one end of the rope and the other end moves too. In the end I settled on the final operating point using trimmers, measurements and my ears.

Low-impedance drive: the pentode experiment

In an early version I also tried running the D3a as a pentode. Its high transconductance (35 mA/V), together with the transformer’s Ra, gives a relatively high gain, and with plenty of feedback you can build a clean stage with a fairly low output impedance. That low Zout is attractive for driving an interstage transformer: the lower the source impedance, the less the winding’s inductance and hysteresis get in the way. But our group listening tests pointed the other way: this gave a smeared, unmusical result. And on closer inspection the calculated Zout with feedback wasn’t substantially lower than what the D3a already gives in plain triode mode.

Solving the remaining technical problems

With the transformer now behaving well on the bench and the operating point settled, a handful of practical problems still had to be ironed out.

Supply instability — and a deliberate torture test

When I build a tube amplifier I always put it through one rigorous test: I feed it an input signal that goes *far* beyond its maximum input voltage. Here the amplifier already delivers full power at roughly 1 V, so I decided to throw 5 V at it. For a few seconds at a time, then quickly back to a normal level. In practice I switched repeatedly between 5 V and 100 mV, in bursts. This test matters all the more with the D3a, which is a frame-grid tube and has a reputation for oscillation.

The problem showed up immediately: the supply sagged and never climbed back to the right level. I fixed it in steps. First I increased the reservoir capacitance and lowered the B+, which gave the regulator more headroom to work with. I also placed a resistor between the two electrolytics, to keep the T-Reg from being overloaded thermally.

That already gave a much healthier picture, but the amplifier still tipped into oscillation now and then. The cure was a 6.8 μF film capacitor per channel, soldered as close to the tube socket as I could get it. After that the supply was indestructible and rock-solid.

Heater supplies

Tentlabs recommends a 4 V AC supply for the 2A3. In practice, though, the extra transformer windings I had wound came out at around 4.1 V under real load. That was just enough to push the filament modules into thermal protection in warm weather, after about six hours of playing. Adding a 0.1 Ω resistor to each module made it stable.

The finishing touch

Everything up to this point was about reaching a healthy baseline: the power supplies, the operating points and the interstage all in good shape. This is the stage where testing and measurement were crucial. Without them I would never have discovered these pain points, let alone fixed them with any confidence. And most of it is hard to predict or find in LTSpice.

I am a firm believer in measurement as a way to discover, fix and understand both problems and sound flavors. But the ear should always have the final say. So from here on I deliberately stopped measuring for its own sake and shifted to listening and finetuning. That meant swapping tubes and components.

Tube rolling

2A3 vs 45

With the B+ now at 275 V, the bias adjustable on a trimmer, and a filament module that delivers a steady 2.5 V regardless of current draw, an opportunity opened up: trying the 45. The 2A3 is already praised for its lovely sound, but the 45 is praised even more. It has something of a cult status on the forums. So I wanted to give it a go. I still had two lying around, and through eBay I found a matched pair of Philco/Sylvania tubes. With a few small adjustments I got it working. The result? A very soft, friendly sound: wonderfully forgiving. But I missed the power and body of the 2A3. As much as I wanted this to be the final answer, I kept drifting back to the 2A3.

Sovtek vs RCA

In my treasure chest, a big steel drawer in my workshop, I had four Sovteks (new production, mono plate) as well as two RCA and two Hytron 2A3s, all picked up cheaply over the years and all still measuring comfortably above 100%. That gave me plenty to swap between. This choice was harder. The Sovtek has a noticeably cleaner, more analytical sound: lovely detail and a fresh top end. The RCA/Hytron combination (RCA on the left, Hytron on the right) wasn’t the last word in detail or analytical behavior, but in return it gave a very soft, musical, forgiving sound. I found it genuinely hard to choose, but I noticed that with the RCA/Hytron pair I listened for longer and switched the system on more often, whereas the Sovtek impresses more in a short listening test. I decided to stay with the RCA/Hytron combination.

D3a vs E280F

This difference was much like Sovtek vs RCA: the D3a is analytical and fresh, the E280F soft and musical. Here too the choice fell on the more forgiving option: the E280F. Bear in mind that my Klipsch speakers are quite revealing and put down a fairly sharp top end; I wouldn’t rule out that with other speakers the more analytical D3a-and-Sovtek combination sounds better.

Distortion pattern

In between the tube rolling I occasionally went back and measured the distortion. What stood out was that with the analytical-sounding combinations (D3a and Sovtek) the H2 sat below the H3, while the more forgiving setup often had H2 just above H3, or at about the same level. That makes sense, because I had originally tuned for minimum H2 and THD using the Sovtek and D3a. The E280F-and-RCA combination lands at a slightly different operating point, and with it a little more H2.

This contrast, H2 below H3 versus H2 above H3, is something we have discussed often within TubeSociety. Menno van der Veen calls it the “H2 ghost” and clearly prefers H2 below H3, which he argues gives a better rendering of microdetail. Nelson Pass has written about it too (see Pass: Audio, Distortion and Feedback). By adjusting operating points, or by rolling tubes, you can shift between these regimes. In this case, and also in an earlier project with the Nelson Pass ACA amplifier, I personally prefer an H2 that sits slightly above the H3. It is a very subtle optimum. At least, that is what I think at the moment.

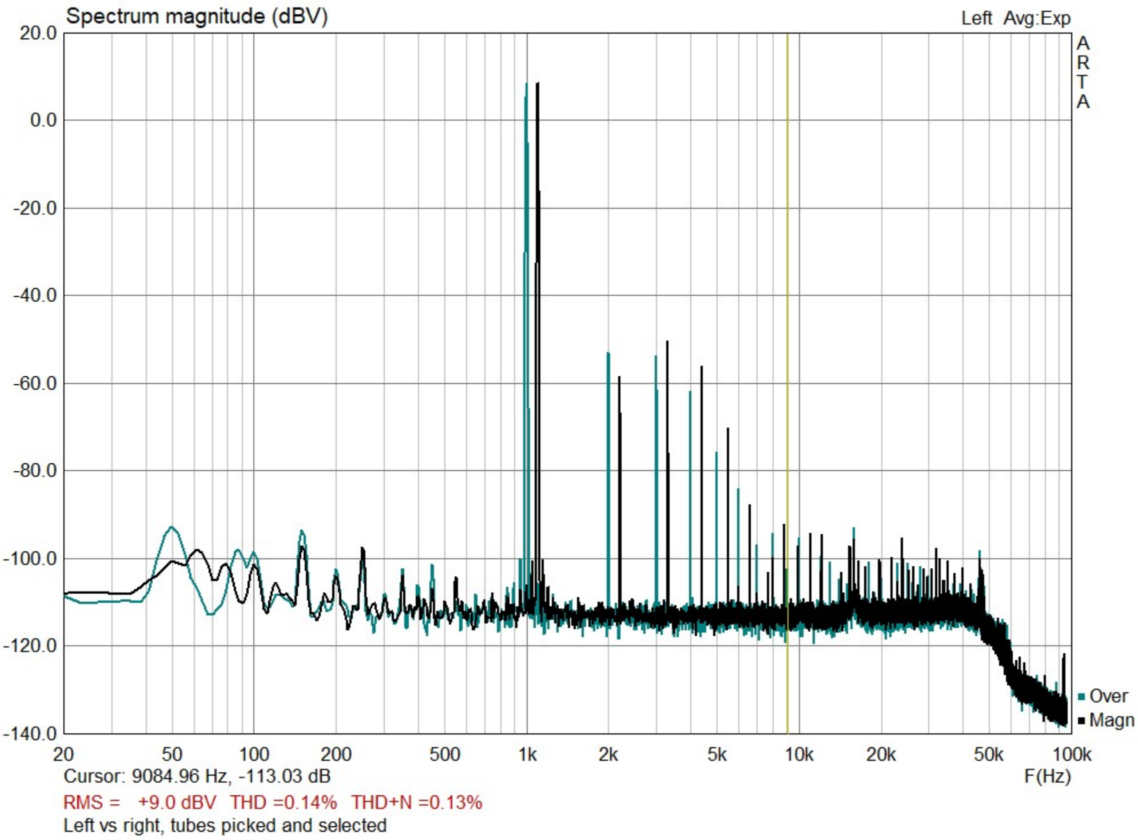

Distortion measurement (d3a + sovtek) at 1Khz into 8ohms.

The cathode bypass capacitor

In an early version I used a 2200 μF bypass cap across the cathode resistor, and I also experimented with an 82 μF tantalum cap that someone once handed me (“you have to try this one: it sounds good!”). In the end the amplifier sounded best with no bypass cap at all. That cost a touch of gain, but the effect was limited, since the cathode resistor is only 30 Ω.

Resistor types

At first the amplifier sounded a little less beautiful in the highs, even after fixing the transformer loading. Several people confirmed this in a group listening test, without me prompting them. It kept nagging at me, so I decided to experiment with resistors. I had initially used cheap parts from an experimenting bin. Reading online about D3a oscillation and resistor types, I decided to replace the cathode resistor and the g2 stopper with Caddock resistors. After that change the sound improved considerably. As a final tweak I replaced the g1 stoppers with carbon composite. A place where this type of resistor, with its low parasitics, can genuinely make sense. That last change, however, made no audible difference.

Thoughts and potential improvements

I have been listening to this amplifier with a lot of pleasure for quite a while now. By design it isn't the last word in detail retrieval, but what you get in return is worth it. It's soft and forgiving, and it invites you to just put on music. The imaging, the stereo picture, sibilance control and the bass (not too much, but well defined) are all pleasant and in balance.

It has since replaced my Line Magnetic with the 845 tube, which, in hindsight, sounds a good deal messier than this well-balanced 2A3 push-pull. I would also pick this amp over a Hypex Nilai, which sounds very good and measures very well, but lacks the musicality of the 2A3.

I'm still waiting on the S version of the Lundahl LL1660, and that is something I definitely want to experiment with further. I have also been reading about bifilar interstage transformers, where the windings run right alongside one another instead of stacked on top, which reduces the imbalance considerably as well.

And I came across the E810F: an even higher transconductance (50 mA/V!) and lower internal resistances, in theory an even better match.

Conclusion

Back to the question I started with: KISS.

The first thing that has become even clearer to me is that the quality of an amplifier, or rather its musicality, can't be read from absolute measurements. A certain group of amplifiers, and their fans, focus on a single figure like SINAD. This KISS amplifier shows that the distortion pattern is really a matter of taste and careful adjustment. On top of that, even though the pentode-plus-feedback version had lower distortion, it didn't sound any better than the straightforward triode version.

My current view is this. Distortion does have an effect. But whether less distortion sounds better depends on the route you take to get there, and quite often the cure is worse than the disease. A good example is applying a big dose of global feedback. Once you let go of the demand for extremely low THD, you find that a lot of the components and circuitry simply become unnecessary. Exactly when the cure becomes worse than the disease, I'm not entirely sure yet. But feedback, as I learned this year and last, is something to handle with care.

And THD is only part of the story. There are plenty of other aspects that deserve at least as much attention: damping factor, where higher isn't necessarily better; resistor types, which I've now learned do make a difference; capacitors; tube types and brands; diode recovery time; the winding technique of the transformer. The list goes on.

So: start simple and stupid, and use your ears. Skip either one, and you risk masking what really matters, never learning a thing.