(The following article contains information on building high-voltage equipment, which unavoidably involves a high risk of electric shock. Please be extremely cautious if you attempt to build something based on the information you find here.)

Mini Valve Curve Tracer

I built my first valve curve tracer a couple of years ago. It was a design by Merlin Blencowe. He published it on https://www.valvewizard.co.uk/, where you can still find it. His goal with this design was to create a decent curve tracer with a minimal amount of components. It works great and has never let me down. But not so long ago I built my own curve tracer for transistors. It is low-voltage (32V) and tests all common kinds of transistors (see picture). It works very well. I was so thrilled with it that I immediately got the idea of tweaking the valve curve tracer I had. In many ways, a transistor curve tracer is harder to design, because of the different kinds of devices: BJT (npn and pnp), JFET and MOSFET (n-channel and p-channel, enhancement, depletion). You have to take into account base-emitter voltage drops, gate bias voltages,… Compared to that, valves are pretty straightforward devices. Of course there are many different kinds as well, but the last half century only a few types are still common. However… they operate on high voltage, which creates its own design issues.

There are some reasons why I wanted to change anything about the old curve tracer. First of all, the tracing frequency is rather low (rectified mains, 100Hz). It takes 10ms to trace out one curve. It draws 8 curves, so it takes 80ms to trace out an entire picture. Although that is fine for reading the curves, the flickering became a little annoying when I knew I had designed a transistor curve tracer without that problem. The second reason was that the X-axis pointed to the left. Again, not a real problem, but it got under my skin. Thirdly, there were some double traces due to each trace being drawn twice (once on the upgoing half, again on the downgoing half of the rectified AC wave). This pulsating DC is not superclean either. So although it worked fine, I was curious whether I could design something that addressed all these items. Of course, that meant letting go of the original design goal: minimalism. The device I created is still mini, but it is not minimalist.

I also tried to use components that I had lying around already. You might think some choices I made along the way are a bit strange, in that you might see a more efficient solution for some problems here and there, but some choices were dictated by what I had in my parts boxes.

In any case, I am heavily indebted to Merlin's design. I do not consider this design an improvement of his design. His design was a starting point to create my own. Check out the valve wizard website. It contains so much valuable, well-explained information on valve and solid-state circuitry.

Power Supply

I used four separate transformers to produce the necessary voltages for the different sections of the circuit.

- A 30VA toroidal transformer (meant for use in a low-power valve amp) with two secondaries of 0-230Vac and 0-6.3Vac. The high voltage secondary supplies current for the anode and screen grid. After rectification and smoothing it produces around 340Vdc. I strived for 300V for the anode supply, so 340V was perfect to start from, with enough headroom in case momentary larger current draw should pull the tranformer's output voltage down a bit (e.g. when testing larger power valves). The low voltage winding supplies current for the PIC16F1765 and surrounding circuitry, which operates on 5Vdc.
- a 50VA toroidal transformer with two secondaries (each 0-12Vac), which I connected in parallel. This is used to supply current to the filaments. I used the same arrangement as in Merlin Blencowe's minimalist curve tracer: an LM2596-based DC-DC buck converter. Its output voltage can be adjusted with a simple potmeter and it can deliver a fairly large amount of DC-current at a very stable voltage.
- A small PCB mount transformer with two isolated secondaries: 0-22Vac and 0-8Vac. The former is used to produce 30Vdc for the LM358 opamps, and the latter to produce the negative voltages that drive the control grid. I actually needed a higher (more negative) dc voltage than the 8Vac could be used to produce, so I used a voltage doubler to get about -18Vdc. This negative-voltage circuit doesn't draw much current, since it just needs to power one opamp, which drives the control grid (which, under normal circumstances, has a very high input impedance and demands very little current). Under these low-current conditions a voltage doubler is a good solution.
- Another small PCB-mount transformer (with two separate 0-9Vac secondaries) just to supply the panel-mount voltmeters (one to monitor the filament voltage and one for the screen grid voltage).

I actually used a simple EI transformer to power the LM2596 module first. However, I always keep the curve tracer on top of my analog oscilloscope. It's basically the only thing I use the analog scope for, so I want to keep them together. The EI transformer creates a considerable EM-field around it, though, and it proved strong enough to disturb the beam of electrons emitted from the scope's CRT. It made the traces wobble around quite a bit. Finally, the 50 VA toroidal transformer I had lying around for years, had found a purpose in life. The EM-field around toroids is much weaker and it made a big difference. I also put three layers of 0.8mm steel plate between the CT and the scope. The traces are very stable now.

All the low-voltage sections have linear regulators, since they're easy to use, dependable and they have internal current limiting and thermal protection in case anything should go wrong.

There is a common ground for all secondary windings, except the 9V windings that power the voltmeters – they need to be isolated from the rest of the circuit, or funny things happen -, and the filament circuit. However, I made sure to avoid any 'overlap' between sections in the PCB layout. The different circuit sections' supply traces are kept separate on the entire main board. The grounds meet at one point on the power supply board, to avoid currents from one section creating unwanted noise and such in other sections. That turned out well.

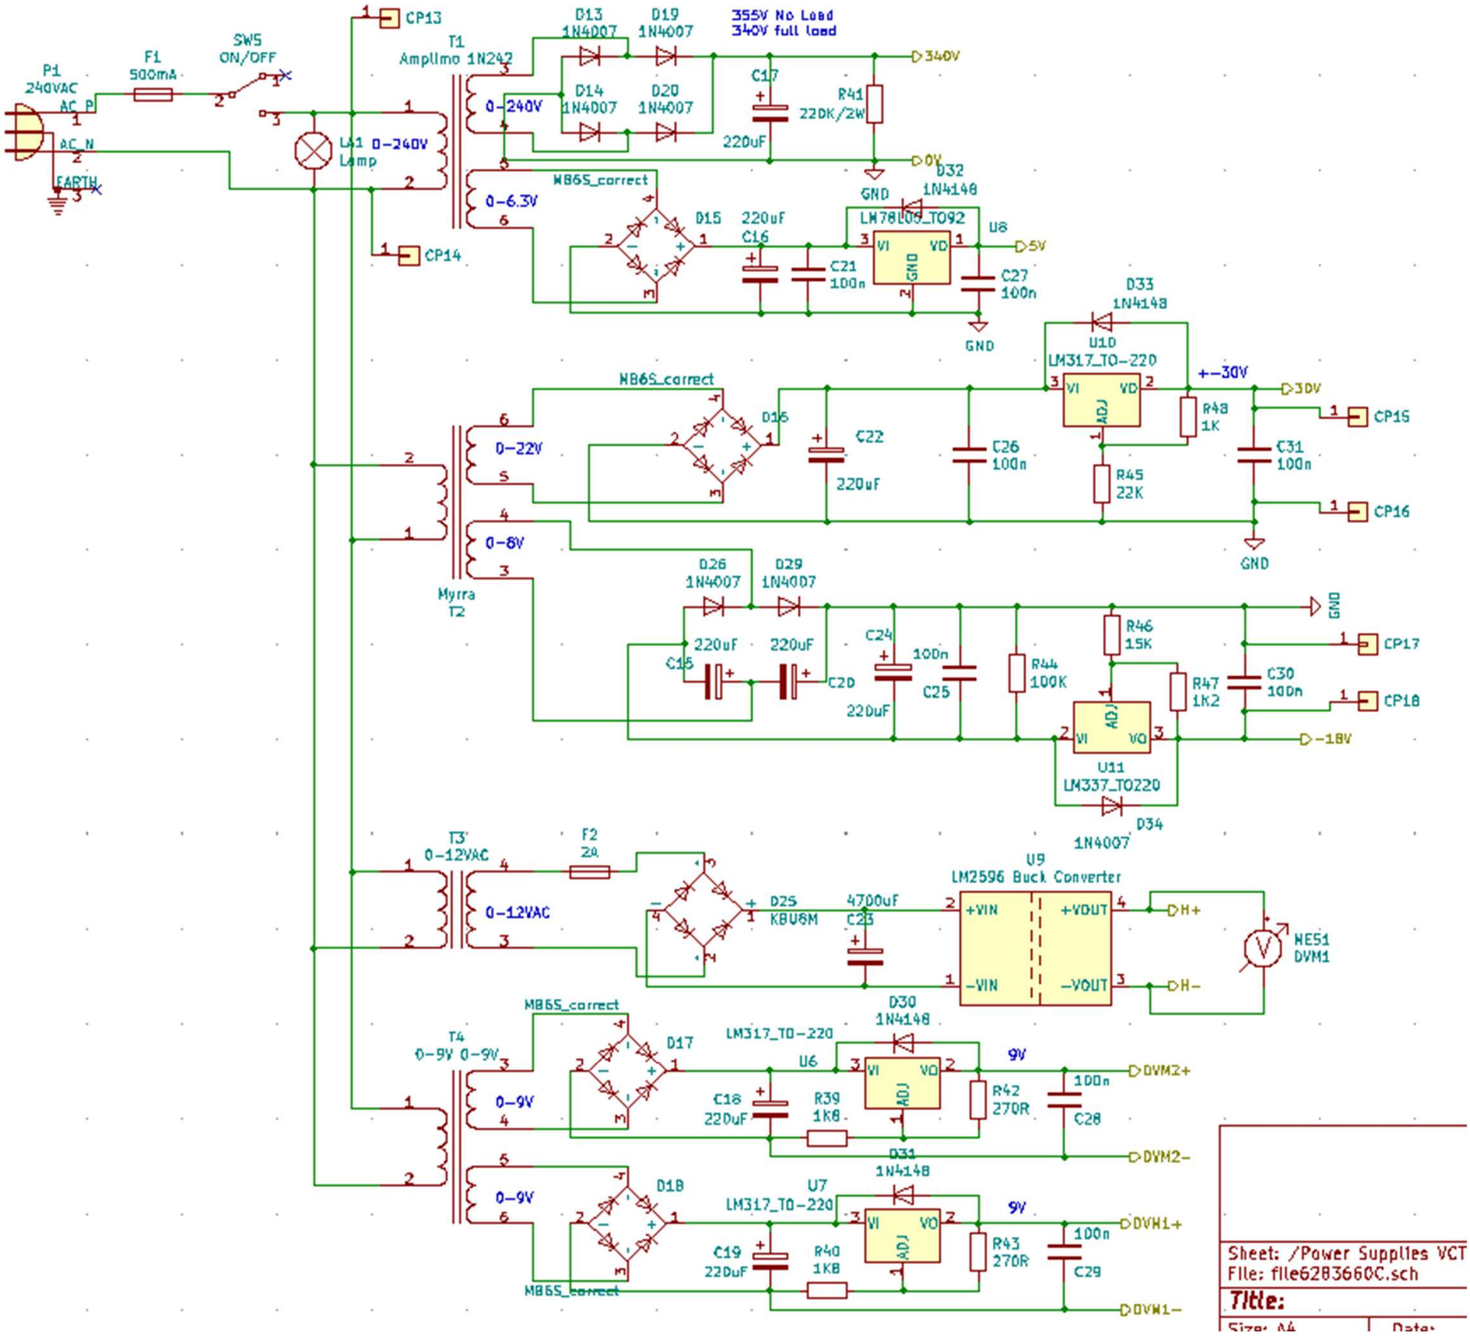

P1 240VAC
F1 500mA
SW5 ON/OFF
T1 Amplimo 1N242
CP13
CP14
LA1 Lamp
0-240V
0-6.3V
D13 1N4007
D19 1N4007
D14 1N4007
D20 1N4007
355V No Load
340V full load
340V
0V
C17 220uF
R41 220K/2W
GND
D32 1N4148
MB6S_correct
D15
C16 220uF
C21 100n
LM78L05_TO92
U8
5V
C27 100n
D33 1N4148
U10 LM317_TO-220
+-30V
30V
R48 1K
R45 22K
C31 100n
CP15
CP16
D16
C22 220uF
C26 100n
0-22V
0-8V
Myrra T2
D26 1N4007
D29 1N4007
C15 220uF
C20 220uF
C24 220uF
C25 100n
R44 100K
R46 15K
R47 1K2
C30 100n
CP17
CP18
-18V
U11 LM337_TO220
D34 1N4007
T3 0-12VAC
F2 2A
D25 KBU8M
C23 4700uF
U9 LM2596 Buck Converter
+VIN
-VIN
+VOUT
-VOUT
H+
H-
ME51 DVM1
T4 0-9V 0-9V
D17
LM317_TO-220
U6
C18 220uF
R39 1K8
D30 1N4148
9V
R42 270R
C28 100n
DVM2+
DVM2-
D18
U7
C19 220uF
R40 1K8
D31 1N4148
R43 270R
C29 100n
DVM1+
DVM1-
Sheet: /Power Supplies VCT
File: file6283660C.sch
Title:
Size: A4
Date:

Anode supply

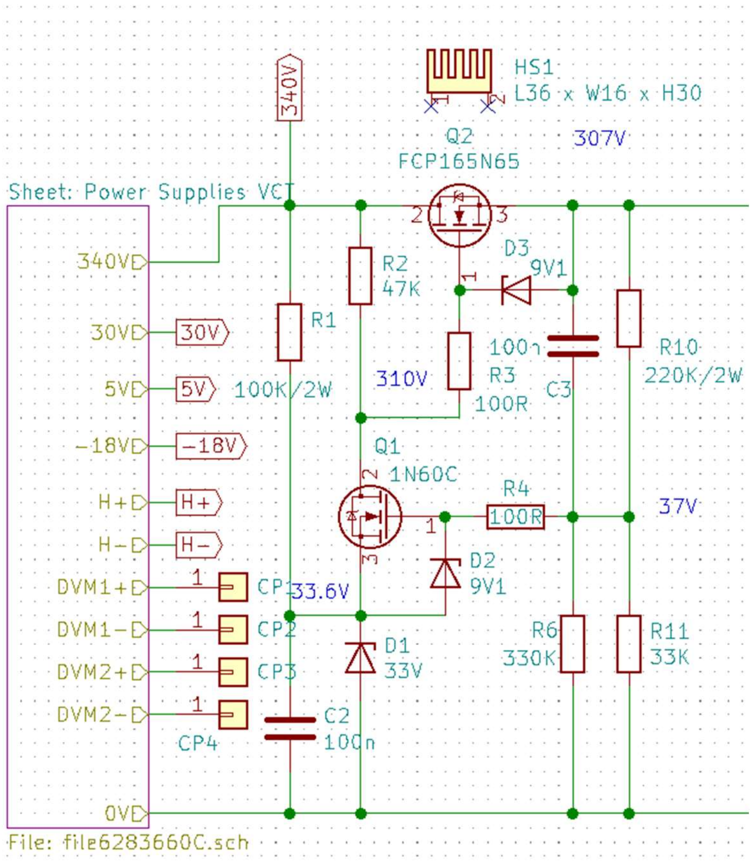

The smoothed 340Vdc is fed into a series voltage stabiliser arrangement. The pass device is mosfet Q2, mounted to a heat sink: the voltage drop across it is 30-40V, so it will dissipate a few watts if it passes even moderate currents of 50mA or so. Low-power mosfet Q1 regulates the voltage at Q2's gate by means of a 33V zener diode D1. R1 supplies current to D1 to keep it active and C2 removes any noise. So the voltage at Q1's source is held constant at about 33.5V (these zeners are not high-precision). Vgs_{on} for this mosfet is around 3.5V, so the gate is held stable at about 37V, as is the voltage across R6||R11, which means a constant current of 1.2mA flows through the parallel combination. This current must also flow through R10, resulting in a voltage of about 270V across it. So in effect, Q1 will try to keep the voltage at Q2's source constant at around 307V. Another way to look at this is to imagine that Q1 corrects any 'error' that appears at the output (source) of Q2: higher voltage at Q1's gate => Q1 conducts more => pulls more current through R2 => lower voltage at Q2's gate => lower output voltage at Q2's source. Error corrected.

The 9V1 zener diodes protect the gate-source junction of the mosfets – a very necessary protection measure (the mosfets will fail at some point without the zeners, as I experienced). There are internal zeners to protect the drain-source channel from reverse voltage. The 100ohm gate resistors prevent current spikes (dis)charging the gate capacitance and unwanted oscillations; for good measure.

R2 has around 30V across it, allowing a current of about 0.6mA to flow through Q1 and D1. Since Q1 has a voltage across it of about 280V, it will dissipate on average 280V * 0.0006A = 180mW, which should be safe for this device.

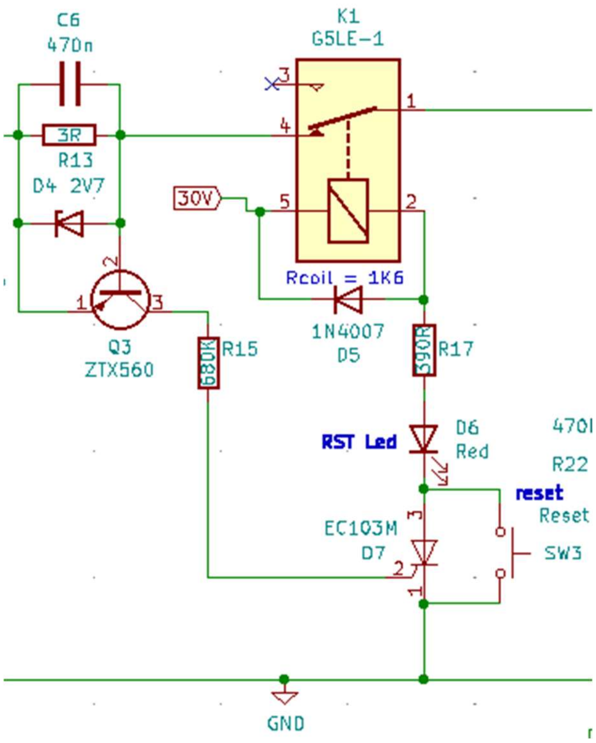

Next in line is a protection circuit. Relay K1, in its de-energized state, passes the supply current from Q2 to the rest of the anode supply circuit, but when this current exceeds 200mA, the relay coil is energized and the anode circuit is disconnected from power.

When 200mA flows through R13, a 0.6V voltage drop results, which activates Q3. Q3 now passes current to thyristor D7's (sensitive) gate and triggers it into a conducting state. Current flows through K1's coil and energizes it. K1 is actually a 24V relay. I don't have a 24V supply, so I measured the coil resistance (1K6 ohms) and added a series resistor R17 to get around 24V across the coil when it is active. The exact values aren't critical, as long as the current through the coil is enough to keep it energized.

Thyristor D7 remains active until the SPST reset switch is pushed (this bypasses D7, reducing the current through it below the holding current level, switching it off), provided the low impedance causing the 200mA of current is removed. Both SW3 and LED D6 are mounted on the front panel of the CT.

C6 smoothes out any spikes from the switching contacts of the relay (which could engage the sensitive-gate thyristor again, as I found out) and D4 protects the base-emitter junction. Q3 actually failed in my experimental setup without D4, though I did abuse it somewhat in testing its limits. With D4 in place, it survived all abuse.

You might wonder why I added this to the circuit. See below under 'current-limiting'. I could also have chosen to somehow pull Q2's gate down to ground or short the gate-source junction to turn it off fully, but I felt more for the idea of using a physical hard switch to disconnect the whole anode circuit from the HV supply.

We now have a stable supply voltage of around 300V to create the triangular voltage sweep with.

Screen grid supply

This part of the curve tracer is almost entirely copied from Merlin Blencowe's minimalist curve tracer. It is a simple, but efficient and practically failsafe circuit.

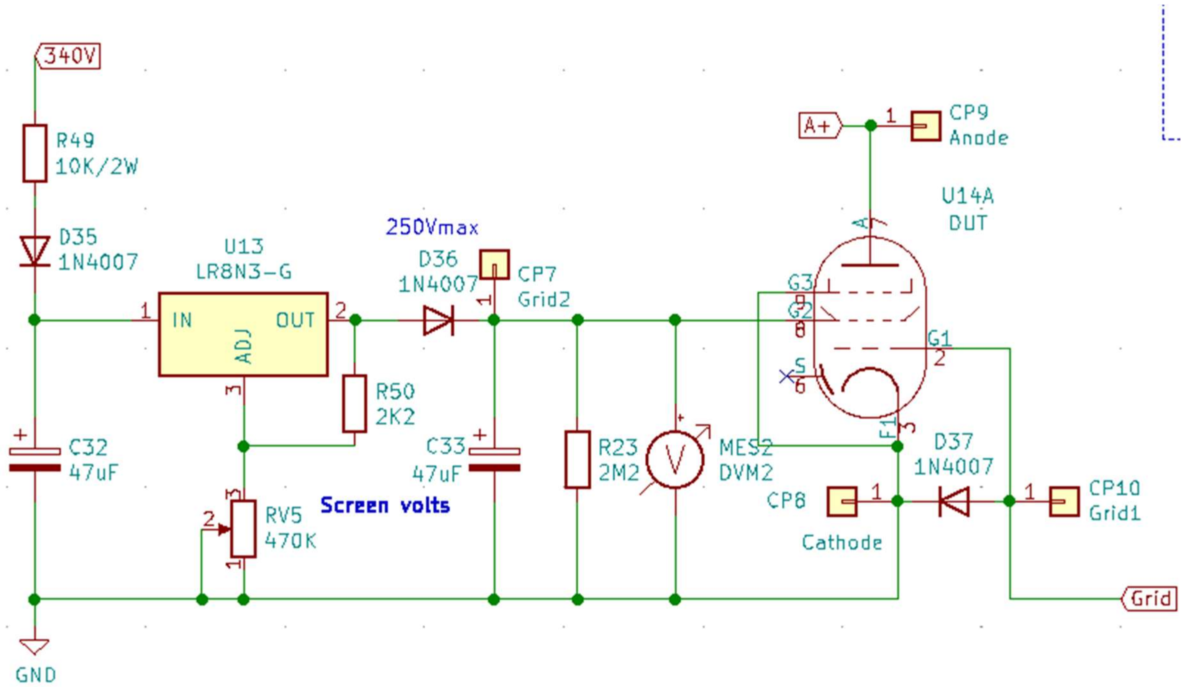

Current is supplied from the main smoothing capacitor C17 (340Vdc). It is fed into a high-voltage regulator, LR8, smoothed again by R49 in conjunction with C32. According to the LR8 datasheet, the voltage drop between Out and Adj pin is 1.20V, so the current flowing through R50 is 0.54mA, which then also flows through RV5, along with 10uA of adj current, but we can disregard the latter without sacrificing too much accuracy. The voltage across RV5 then ranges from 0 to around 250V. D36 protects the LR8 and prevents an internal short between anode and screen to push current into its output. C33 further smooths the output. R23 is there to help discharge C33 when power is switched off. A voltmeter MES2 shows you the voltage on the screen grid. D35 is not strictly necessary in this design, but I left it there, as in the minimalist curve tracer original design.

The LR8 cannot deliver much current. It is internally limited to about 20mA. This makes the curve tracer unsuited for very large valves that can draw substantial amounts of screen current. But this does not seem to be much of a problem, since the valves I need to test are nearly all small-signal triodes (12AX7 and such) and pentodes (EF86), or output valves 6L6, EL34, EL84, and 6V6. The most common ones in valve equipment today (and the last 70 years, I would think). I mostly test output valves at a screen voltage of 100V, since it is a convenient value and a lot of datasheets provide specs at that value of screen voltage. They can also be extrapolated from specs at higher screen voltages, if necessary.

A curve tracer built to handle big bottles would require a separate transformer to supply the necessary current. I did not think it was worth the extra cost and circuit space. I can still test them at low screen voltages.

The regulator is current-limited and has internal junction thermal limiting as well. At maximum current output, this design limits the power dissipated in the LR8 further through R49. LR8 comes in a TO-92 package (and some surface-mount packages) and will heat up rather fast. R49 takes up a large chunk of the power delivered to the regulator to prevent it from shutting down too fast. Again at the cost of available screen voltage. But the design is practically failsafe, and that outweighs any drawbacks, to me at least.

Generating the anode sweep voltage and the control grid step voltage

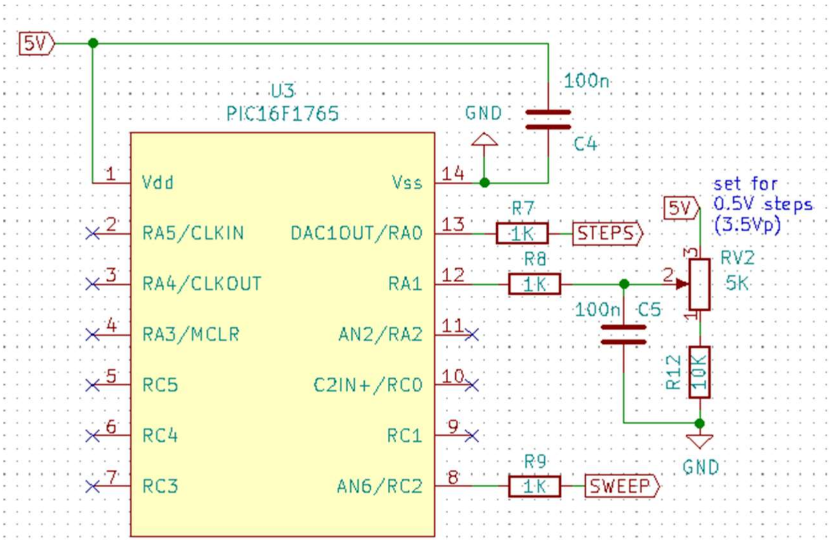

In order to save space on the board I programmed a microcontroller to produce the necessary waveforms that drive the anode and control grid. These can be produced with analog components, of course, but a microcontroller is perfect for this job and saves a tremendous amount of work and space. I used a PIC16F1765 because I already had them for a previous project and they have 2 DAC's: a 10-bit and a 5-bit DAC. The 10-bit DAC is programmed to produce the sweep voltage or triangle wave. The 5-bit DAC produces the step waveform, synchronised with the triangle waveform. 10 bits equals 1023 discrete steps to produce one half of a triangle wave with amplitude 5V. This results in a nicely even triangle wave with a maximum frequency of 125Hz. Of course each slope of the triangle is used to trace out the current through the DUT at one specific grid voltage step, so the 'tracing frequency' is actually twice that: 250Hz. This cures the flickering you get on an analog scope when you use the 100Hz rectified mains waveform to sweep the anode voltage and you get rid of some double traces resulting from tracing at a specific grid voltage while the anode voltage swings up ánd down.

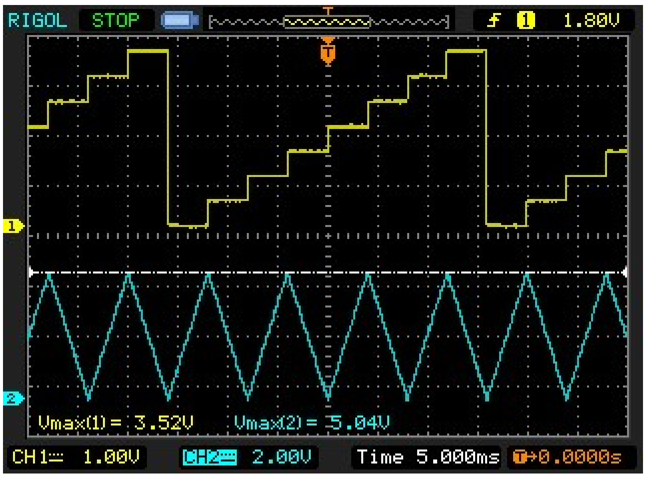

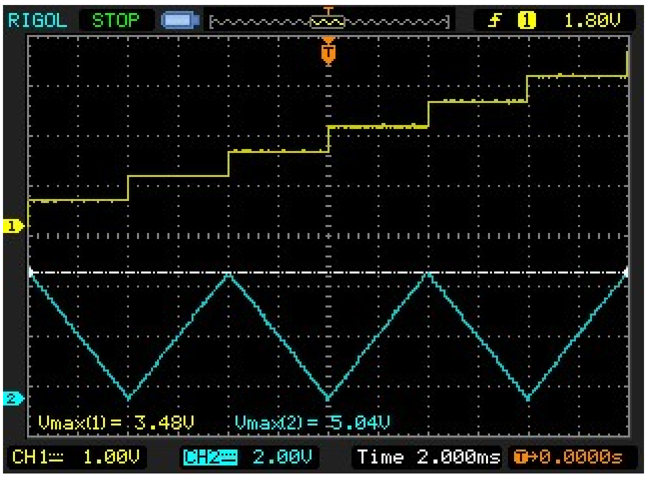

The figures show the waveforms at the outputs of the PIC.

The step voltages can be adjusted to 0.5V by RV2. You get 8 different levels for the grid voltage, from 0 to 3.5V, which is enough to get a clear picture of how a valve performs.

Generating negative grid voltages

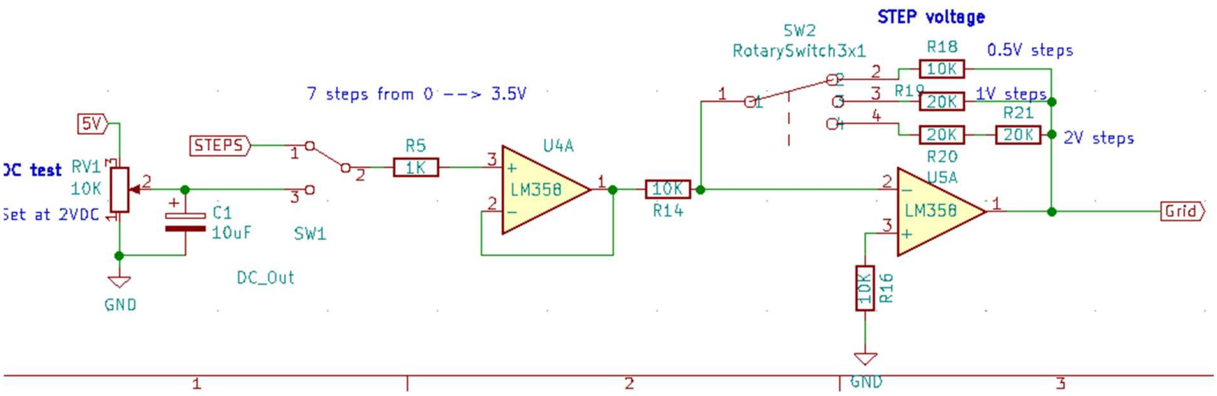

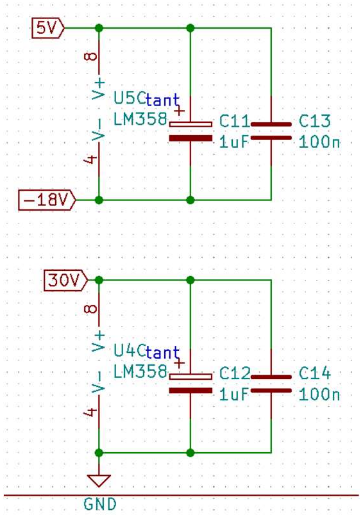

The ‘steps’ output from the PIC is buffered. The DAC output from the PIC is relatively high impedance, so buffering is necessary to avoid signal loss. The signal is then fed into an inverting amplifier with switchable gain to allow larger step voltages for power valves. 0.5V steps are ideal for high-gain triodes like 12AX7, while I use 2V steps for 6L6’s and such.

U4 is powered from the 30V circuit. U5 is powered from 5V at V+ and -18V at V-. The reason for this is that an LM358 can swing its output down to V-, but not up to V+. It can only output up to 1.5V below V+. If V+ were tied to GND, the LM358 would not be able to reliably output the 0V step for instance. So I tied V+ to 5V. It may have been ok to use 1 opamp for this arrangement, powered from +5V and -18V, but the buffer’s output goes to 3.5V, which is a bit close to the maximum output swing. For the cost of one extra opamp, I didn’t have to worry about that. Rail-to-rail output opamps are relatively expensive (and prone to failure, in my experience), so this was a better option.

U5A produces the negative going steps (0V → -14V max) with reference to GND (+ input pin).

In my first version I used a 10K resistor between U5A’s output and the DUT control grid. That didn’t work very well, because at grid voltages at or close to 0V, some grid current starts to flow. That caused some traces to wobble due to a reverse voltage drop across the 10K resistor. I just jumpered it in the end.

I also added a switch to bypass the step voltages and select one DC voltage instead. I don’t know if this is going to be useful in the future, but it was easy to add. It can be set to any DC voltage between 0 and 5V. I set it at 2V. This results in selectable grid voltages of -2V, -4V or -8V.

I used precision resistors (0.1%) around the inverting amplifier, for peace of mind. It probably wouldn't affect performance very much if I had used 5% devices (or select the ones with values closest to each other).

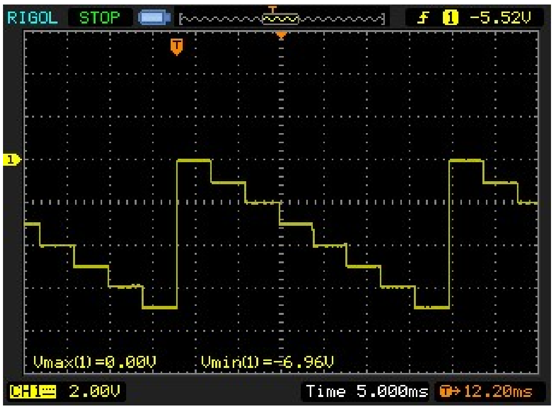

The figure shows the waveform at the output of U5A, the step selector switch set at 1V/step. This goes straight to the DUT control grid.

Anode HV sweep

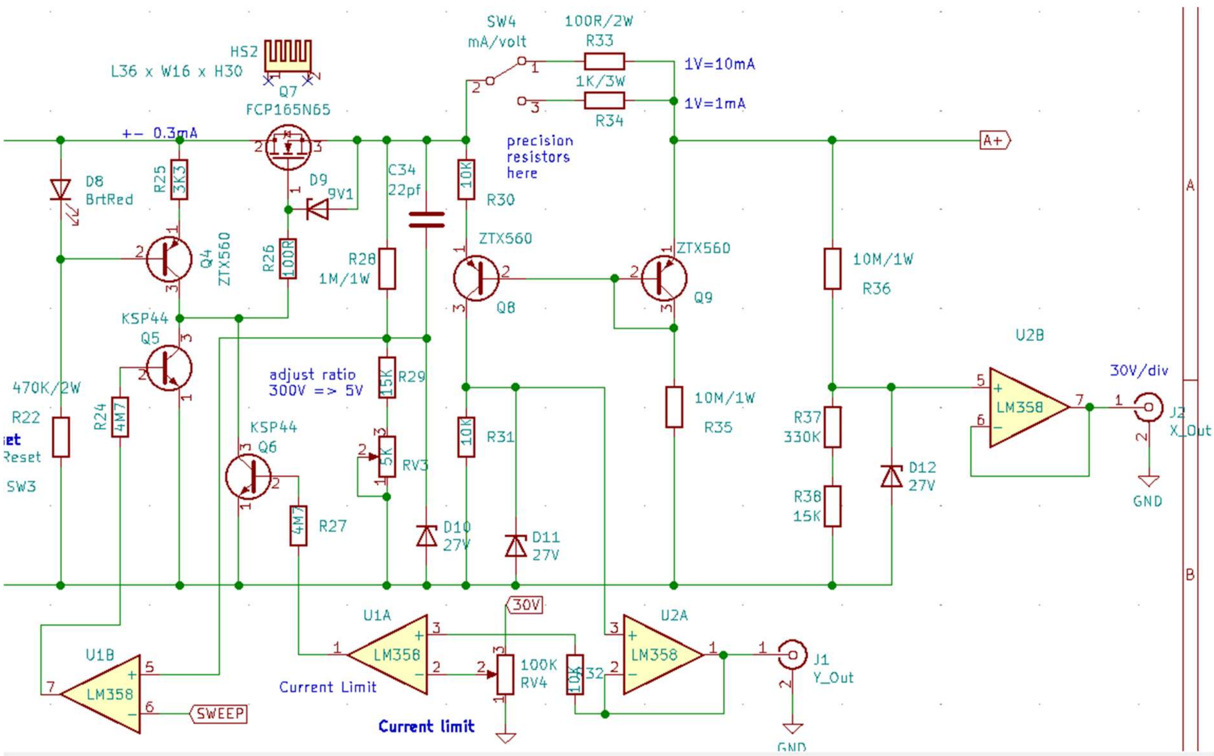

This was by far the trickiest part of the circuit to get right…

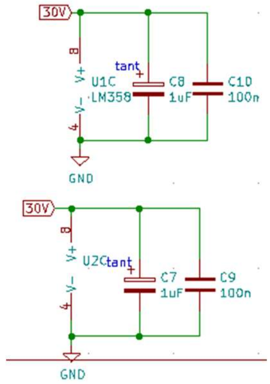

The idea is to apply the sweep waveform from the PIC to the gate of a mosfet (Q7) so that it generates a 0-300V sweep voltage across the DUT. I also wanted some sort of current-limiting and I wanted to keep the Y and X outputs both referenced to ground so that the X-axis points to the right and the Y-axis up on the oscilloscope screen. In the original simple curve tracer, a 1K anode resistor served the dual purpose of converting current to voltage (Y-output) and limiting the current through the valve. At low current levels (below 10mA) that works fine, in that the anode voltage still swings close to the full supply voltage. But at for instance 80mA of current, the anode voltage drops by 80V, which appears on the scope screen as a shorter trace. I wanted to avoid that by adding an SPDT switch: 1K or 100ohm anode resistor. The 1K resistor for low-power valves, the 100ohm resistor for higher power valves. See below for further explanation.

- *anode sweep*

The basic principle was to apply the 5V sweep to the base/gate of a common emitter/source amplifier, which in turn drives the gate of the pass device mosfet Q7, which acts as a source follower. In my first design I used a HV small-signal mosfet for Q5. Problem is that a mosfet gate is very sensitive. I needed a very small amplitude voltage at a specific DC offset to produce an output vaguely resembling a triangle wave. Using a drain resistor as the load didn't work well. I then replaced the resistor by a simple constant current source: D8, R22, R25, Q4. Since the voltage drop across D8 is more or less constant, the current through Q4 is constant. From Q5's viewpoint, the load is now very high impedance. I managed to get decent looking triangle waves now, but it was very fiddly to get the input voltage at exactly the right amplitude and DC offset (using 2 potmeters). Even then, the mosfet's Vgs_{on} would start to drift a little as the device's temperature drifted and the output would start to look uglier. Another solution was needed.

Eventually I came up with a feedback system. The full 5V sweep waveform from the PIC is applied to U1B's negative input. A sample of Q7's output voltage is applied to U1B's positive input. The opamp compares both inputs and will output whatever is necessary to make the voltage at both inputs the same. Suppose the sampled output voltage is higher than the voltage at the negative input. This will make the opamp's output go positive and cause a current to flow through R24 into the base of Q5. This will then make Q5 turn on harder, lowering the voltage on the gate of Q7, lowering the output voltage at its source.

The output is sampled by a series combination of R28, R29 and RV3. If the output at Q7's source reaches 300V, the sample should reach 5V (the same amplitude as the PIC produces). Taking a 1M resistor as a convenient value to start from, the calculated 'tail' resistance needed would be 16949 ohms. I used a 15K resistor and a 5K trimpot to dial in the exact value needed to produce a clean triangle wave.

Q5 is a BJT, not a mosfet. As I had experienced before, a mosfet is a very high gain device, which gives the opamp only a small amount of headroom to work with. A BJT has less gain (KSP44 current gain is around 150), and it is current-driven, so the device operation can be further controlled via a single base resistor. By experimentation, I found that a 4M7 base resistor gave the opamp ample headroom to perform its duty. The exact value is not critical, though, as long as the opamp gets that headroom (even 1M worked fine, as I recall, but with less headroom).

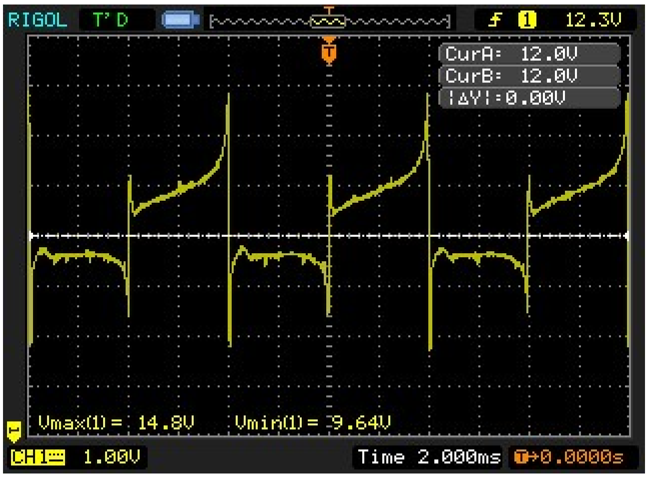

The figure shows the output of U1B. You can see that its output is nicely centred around 12V (U1 is powered from V+ = 30V, V- = 0V). The output swings around 5Vpp.

C34 is parallel to the 1M feedback resistor. Without it, oscillations occur on the output of U1B, resulting in curly traces on the scope screen. 22pf is enough to get rid of them and get a clean output.

R28 is a higher wattage device, since they can handle high voltages better. 1/2W devices are only rated for 300V. For peace of mind.

D10 protects the opamp's input from overvoltage. This would only happen if RV3 or R29 fails open, but still. It's a small effort.

Q7 also needs to be heat-sunk. I guess if you would test high-current devices or even EL34's at elevated screen voltages for extended periods you would need a bigger heat sink than the one I used, but usually I keep the screen voltage at or below 100V so the current does not exceed 100mA (e.g. for an EL34 at Vgk = 0V). I also wouldn't leave the valve in circuit for more than a couple of minutes. Enough to let it warm up nicely and look at the curves. It would be too involved to calculate exact dissipation figures, since the current through Q7 varies with each grid voltage step. If it gets destroyed at some point, I'll update this section.

This is how the anode voltage looks with a 150K resistor across the anode and cathode terminals (no valve inserted of course).

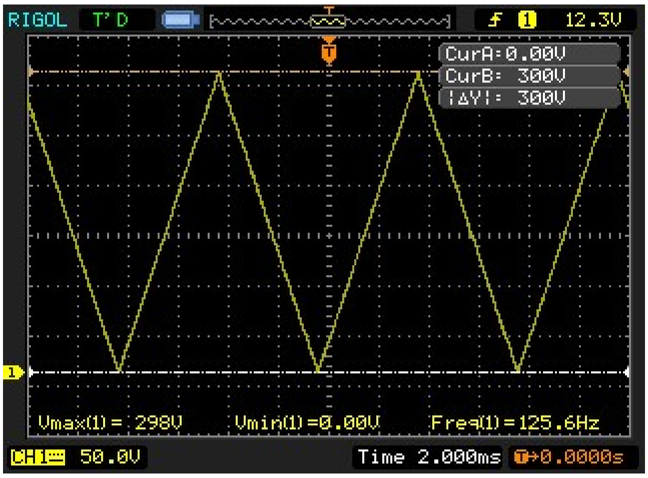

You can see how nicely the feedback arrangement works. It is a near-perfect high voltage copy of the PIC's 5V sweep output (inverted, but that is of no importance here).

- *sampling the current*

I was looking to find a way to convert the anode current into a manageable voltage that was referenced to ground. Taking the voltage straight from the resistors in series with the anode was not an option in that respect. Putting a low-value resistor between cathode and ground was not an option either, since it affects the grid-to-cathode voltage (the step voltages are referenced to ground) and in the case of pentodes, screen current must be added to the equation. I didn't like the idea.

So after quite a bit of experimentation, I came up with the following circuit. Q8 and Q9 are HV small-signal BJT's (ZTX560 has an hfe of 300 max, which is quite a lot for a high-voltage PNP device) in an arrangement that is reminiscent of a current mirror. If we go around a CCW loop and add the voltage drops from Q9's base back to Q8's base, we get $0 = (Vbe_{Q9} + V_{Rs} - V_{R30} - Vbe_{Q8})$. V_{Rs} is the voltage drop across R33 or R34. Since both Q8 and Q9 are ZTX560 (ideally these are a matched pair, but it seems the error is really negligible), their Vbes cancel and $V_{R30} = V_{Rs}$. So, in effect, V_{Rs} appears across R30. Hence a current flows through Q8, equal to V_{R30}/R30. Since R30 = R31 = 10K, $V_{R30} = V_{R31}$. So with V_{R31} we have a copy of the voltage across the anode series resistors, and it is referenced to ground.

There are a couple of factors that cause errors in the copy: 1) the Vbe's don't exactly match, 2) there is some base current flowing out of Q8 that flowed through R30, but not R31, and 3) there is some current flowing through Q9 that has to flow through the series resistors and cause a slightly higher voltage drop there. However, the errors are very small and errors 2) and 3) work in opposite directions, possibly cancelling each other out at least partly. Q9's tail resistor is 10M and draws only 30uA at 300V. I didn't measure everything to the microamp, but if I connect a 47K resistor across the anode and cathode terminals, the following voltage appears across R31:

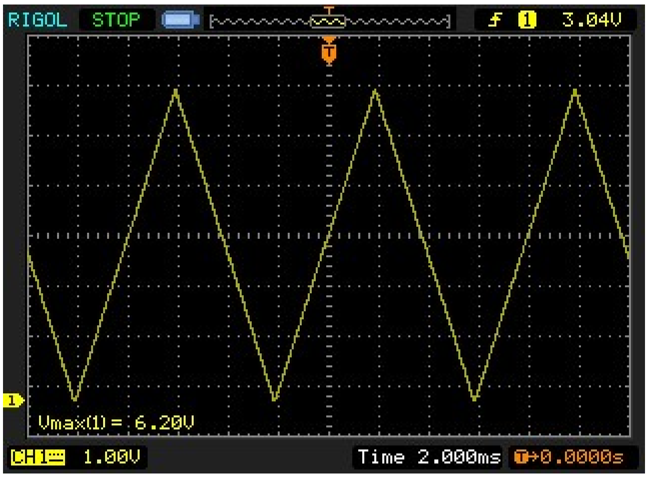

We would expect a peak current of (300 – V_{Rs})/47K. V_{Rs} in this case would be about 6V if we use the 1K resistor. So the expected peak current is 294/47K = 6.26mA and the peak voltage 6.26V across R31, which is very close to what we can read off the oscilloscope.

Regarding the resistor values: if the series resistor is 1K, then 1mA through it causes a 1V drop. This 1V appears across R31, so we have a 1mA/V scale. If we switch to the 100ohm resistor, we have a 10mA/V scale. This turned out to be a very convenient feature. I also used precision resistors here, although you could debate whether it is worth it since there is some margin of error anyway, albeit small.

The voltage across R31 also appears on the positive input of U2A, which acts as a high input impedance non-inverting buffer. Its output goes straight to the Y-input of the oscilloscope.

Suppose I forgot to switch to the 10mA/V scale and I put in a higher current valve, the voltage across R31 could reach levels that damage the opamp reading the voltage. So I added a 27V zener diode to protect the opamp's input. In addition, I added a current-limiting arrangement.

- *current-limiting*

In the 10mA/V setting, there is not much to limit the current through the DUT. Pentodes have screen grids that keep the anode current in check, but triodes don't. An ECC99 could draw a lot of current if it were left running free.

I created another feedback arrangement to accomplish the current-limiting. U1A's positive input is connected to the output of U1B (the Y-channel voltage). U1A's negative input can be varied from 0V to 30Vdc. Basically a simple comparator. If the voltage across R31, i.e. the current through the valve, exceeds the voltage set by current limit pot RV4, U1A's output switches high, which feeds current through R27 into the base of Q6. Q6 turns on fully, pulling the gate of pass device Q7 low, effectively switching it off.

Here is the same 47K resistor across anode and cathode, but with current limited to 3mA:

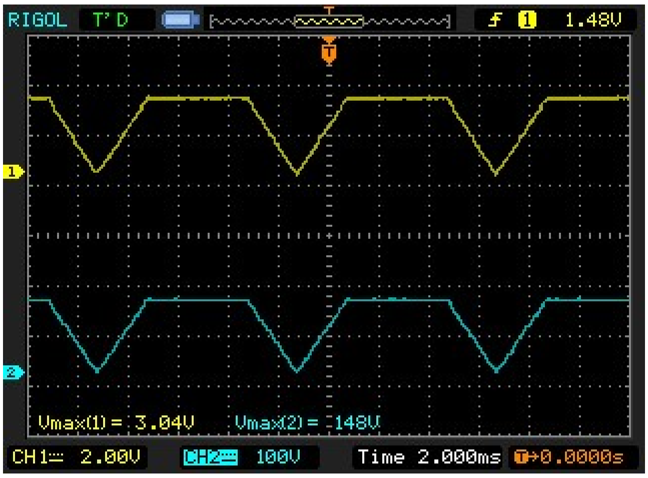

We can see the Y-output levelling off at 3V = 3mA on the top trace. The bottom trace shows how Q7's gate is controlled by U1A and Q6, giving an anode voltage that levels off at whatever voltage gives a current of around 3mA through the DUT.

R27 is a 4M7 resistor. I experimented with several values and again, anything between 1M and 10M would work. Lower values like 10K will result in strange behaviour from the opamp, because it would be struggling to find the right amount of voltage to output. Suppose Q6 is fully conducting, then only 0.34mA of collector current would be flowing through it via the constant current source (I measured the voltage across R25 to be 1.13V). KSP44 has a fairly high hfe (200 max). That means the base current needed is tiny.

This is the actual response of U1A vs. Y-output:

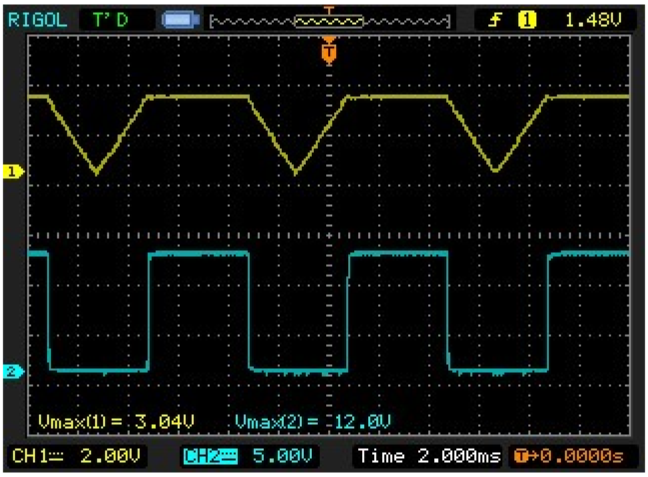

We can see the opamp has to output 12V to turn Q6 on. Base current is then (12-0.6)/4M7 = 2.4 uA. 0.34/0.0024 = 142. This transistor shows an hfe of 142 in this circuit. Point being it doesn't require much base current. If the base resistor were 10K, the opamp would only need to output 0.024V to turn on Q6. That is no headroom at all.

The current-limiting action is derived from the Y-output, so if we switch to the 10mA/V position, current-limiting will set in at 10x the previous value as well.

The absence of a larger impedance in series with the anode will make it possible for the curve tracer to push the DUT out of its safe operating area. That is not likely to be a problem, since the time spent in the danger zone is short (only at low Vgk), and you wouldn't leave the DUT in circuit for more than a few minutes. In addition, in the case of power pentodes, I can't run them on high screen voltages anyway.

The current-limiting mechanism doesn't work in case there is a very low impedance across anode and cathode (like a dead short). Q9's emitter would sit too close to GND at that instant to perform its duty well. That is why I added the relay and surrounding circuitry (see above). I actually made a mistake when I had been testing power pentodes in triode mode and inserted a 12AX7 without changing the connections on the front panel. The relay tripped and protected the triode from damage, and a number of components in the curve tracer from destruction.

Outputs to the oscilloscope

As explained before, the Y-output is taken from R31 and buffered. This goes to the Y-channel of the oscilloscope.

To sample the anode voltage, I simply connected a voltage divider between anode and cathode or GND. In order to minimize the error due to the extra current draw, I used large values. Starting with the 10M resistor, if I wanted to show the 300V as 10V on the oscilloscope (corresponding to the 1V/div. setting on the X-axis with 10 divisions), I would need a 'tail' resistor of 344 827 ohms. That became 330K + 15K in this circuit. This voltage is buffered and sent to the X-channel of the oscilloscope. Again a zener as a protection diode for the input.

I needed to calibrate my oscilloscope to get a reading as accurate as possible. I simply attached a couple of different resistors across the anode and cathode connections so I could calculate the current and voltage drops and compare that to the reading on the oscilloscope.

The insides of the mini curve tracer:

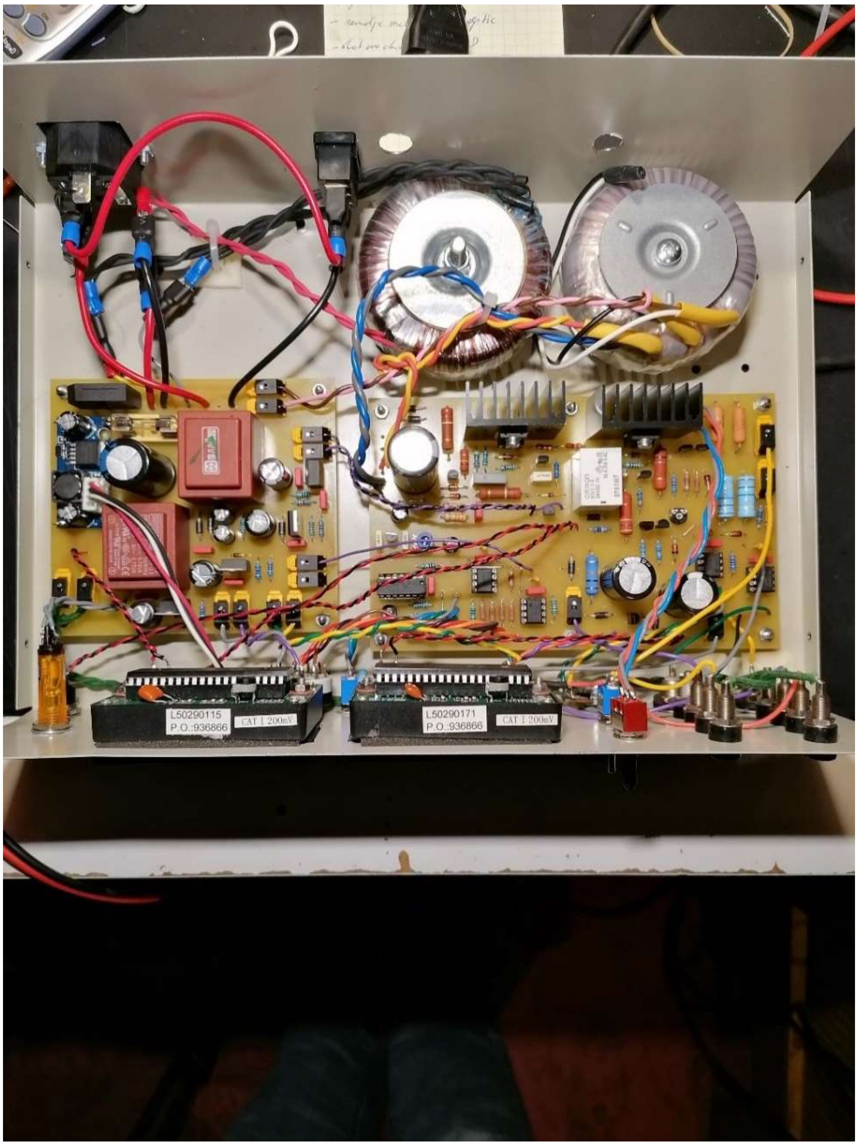

The breadboarded version on which the experiment unfolded. I had to keep my head clear working on this thing, with 300Vdc sitting naked on the bench… Labelling the supply traces is a good idea.

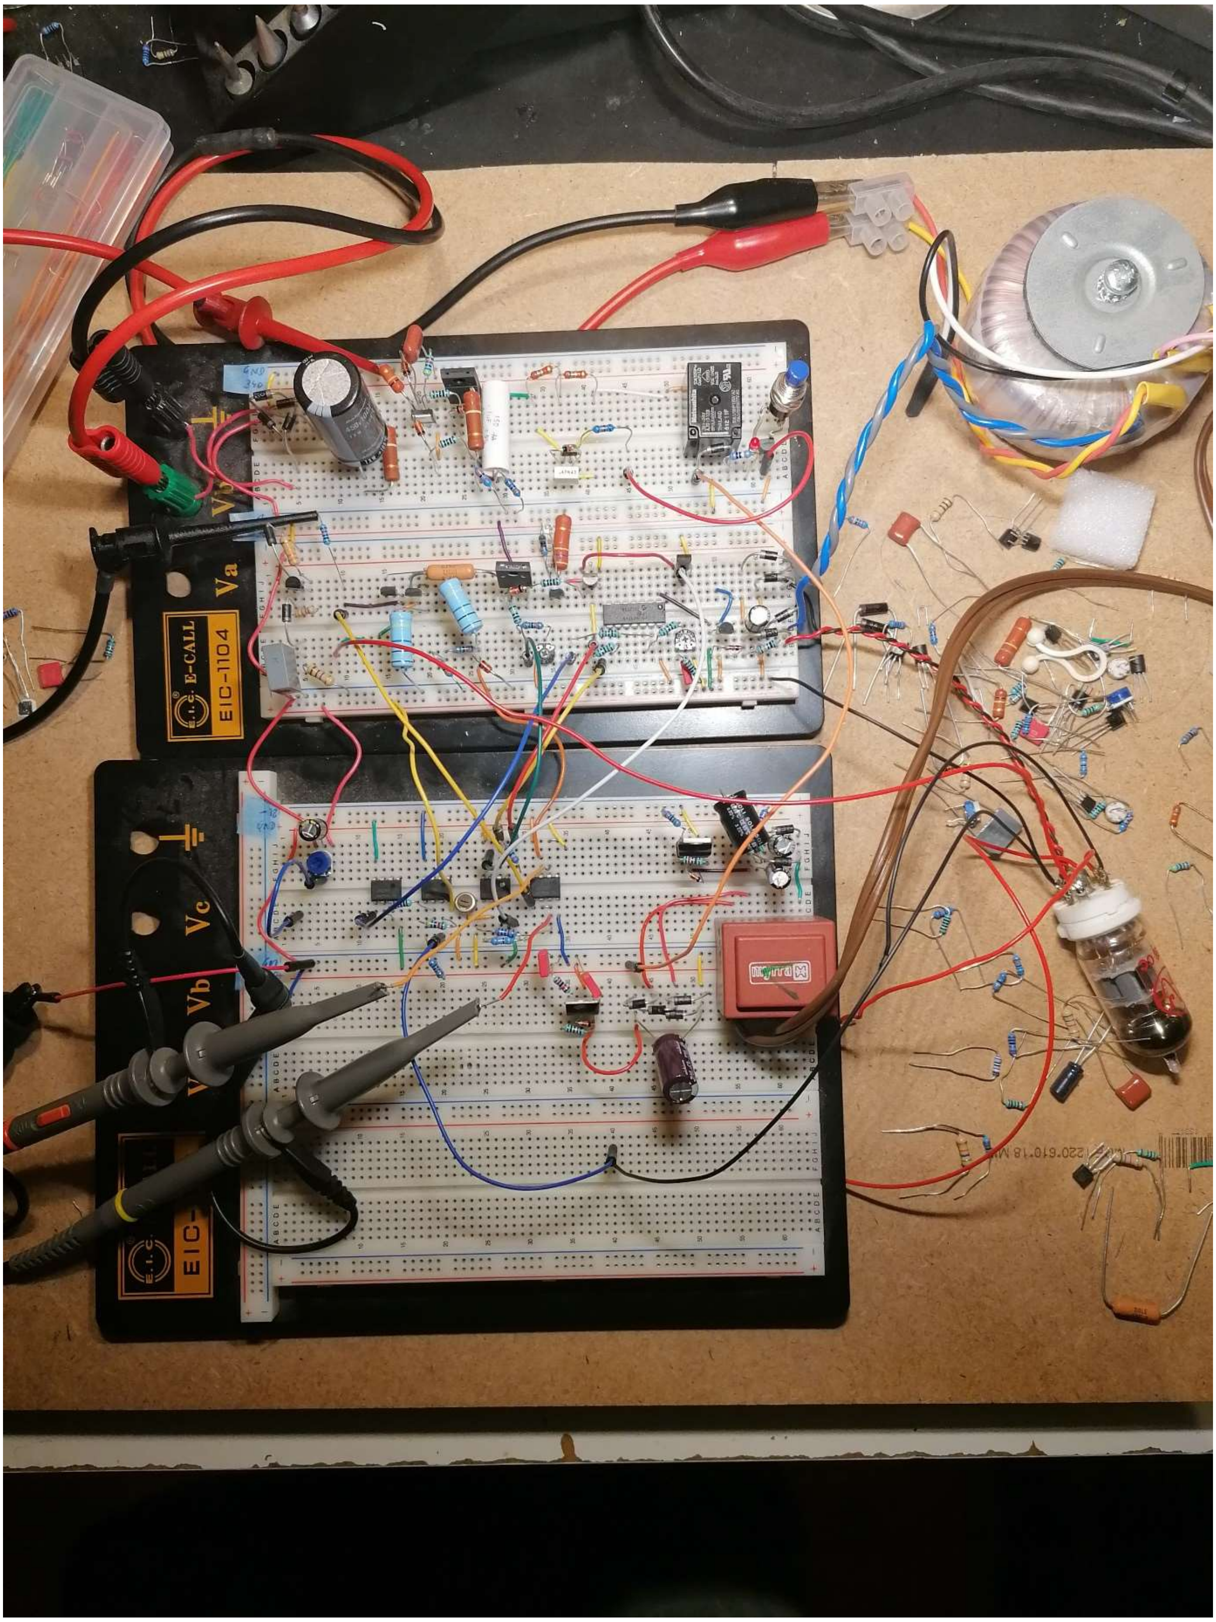

A healthy 12AX7. If you wonder why the volts/div's are set on 0.5: I used fixed resistors between anode and cathode to calibrate the oscilloscope in the XY mode. The centre button can be used to calibrate it. This scope read a little low on the 1V/div setting. I use this scope only for curve tracing, so I can leave it all as it is.

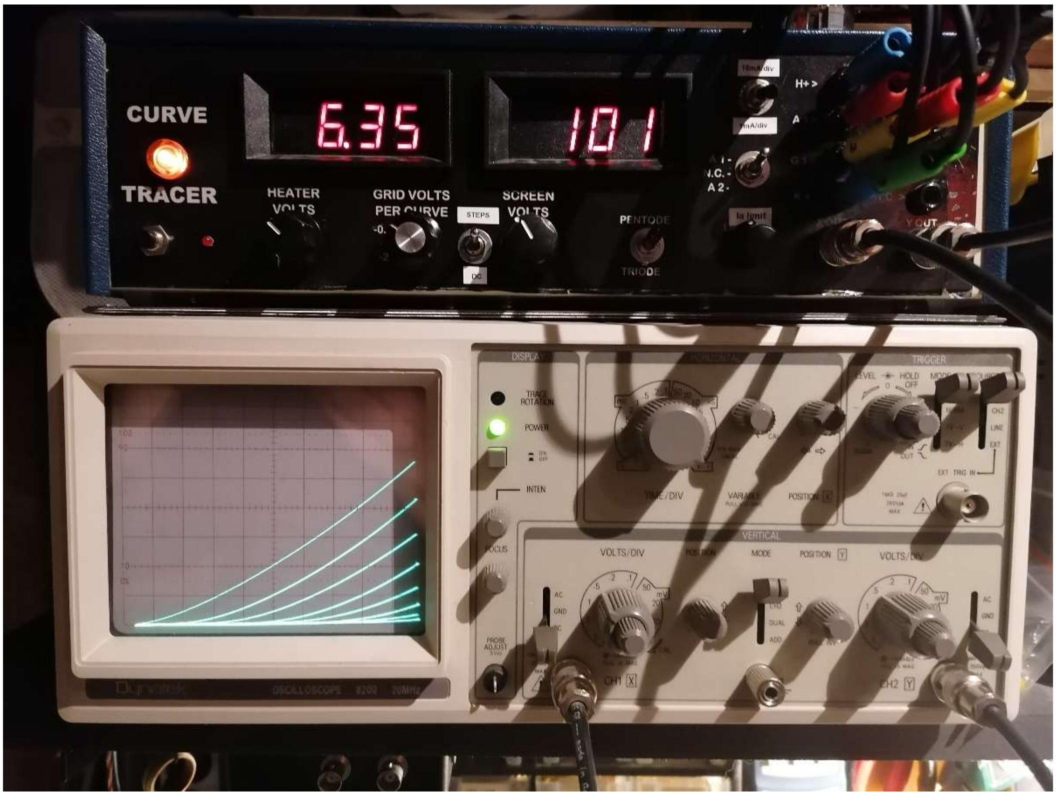

A less healthy 12AX7

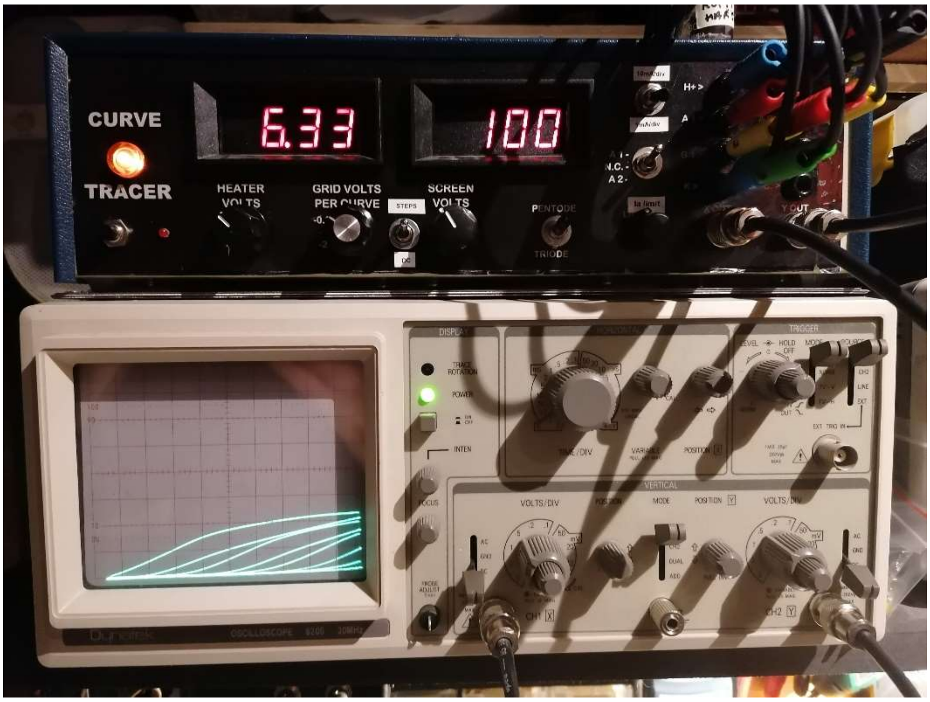

A 6L6 that has been used quite a lot (emission on the low side)

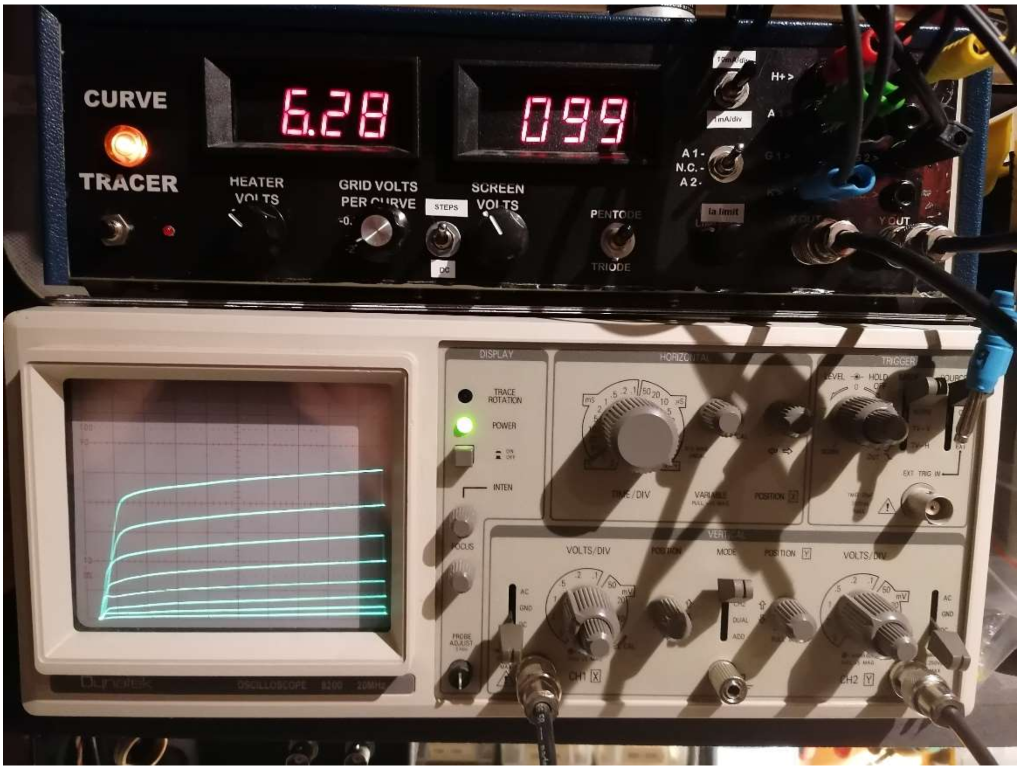

And this is the same thing, but sampled on a digital oscilloscope. It should be clear why I prefer analog scopes for this.

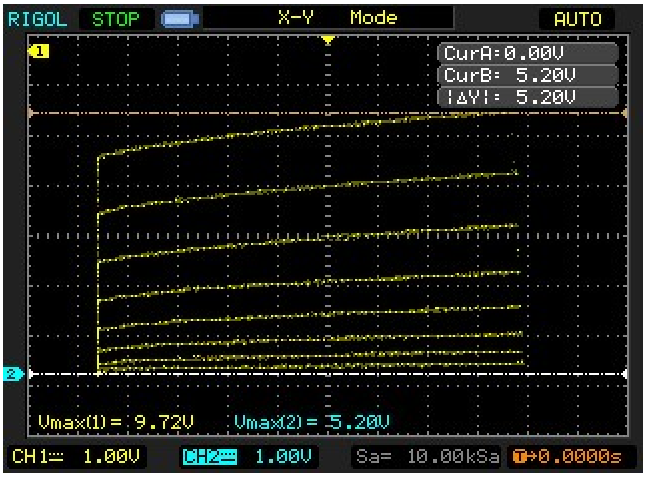

The same 6L6 with current limited to 30mA

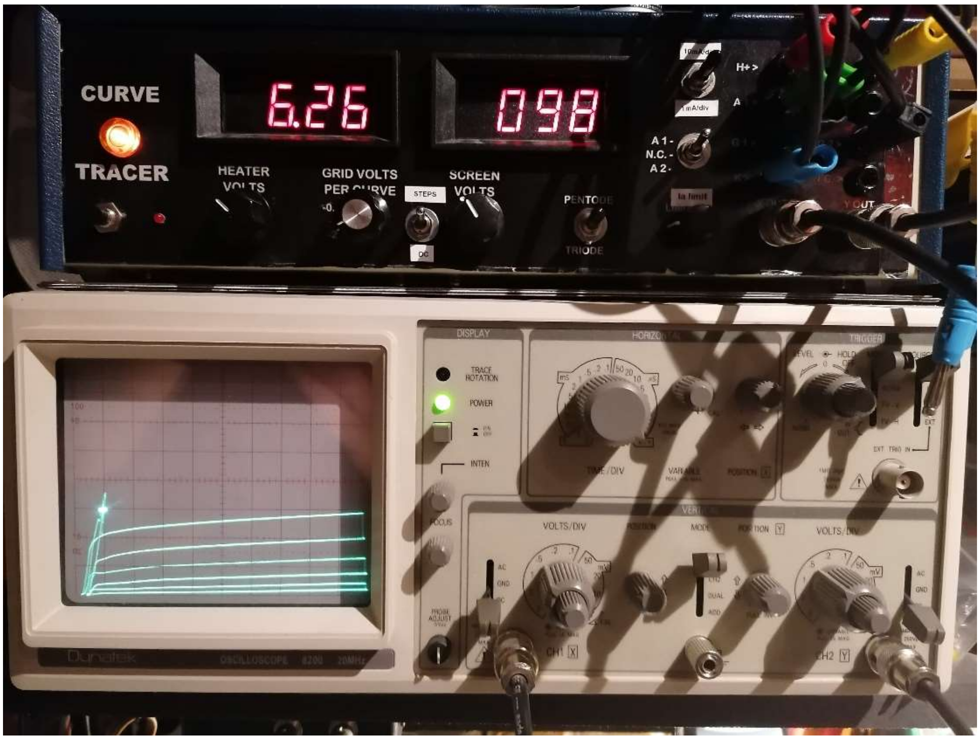

The same 6L6 connected as a triode, with current limiting

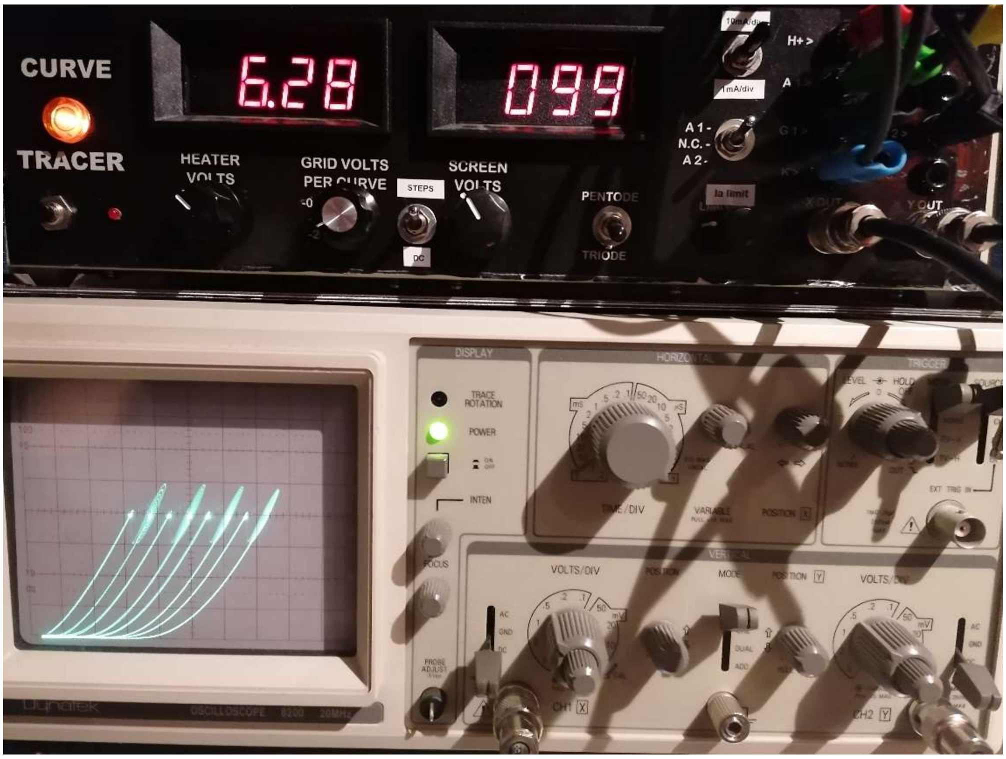

A new 6L6 with healthy emission (65mA at Vgk = 0V, Va = 300V, Vg2 = 100V):

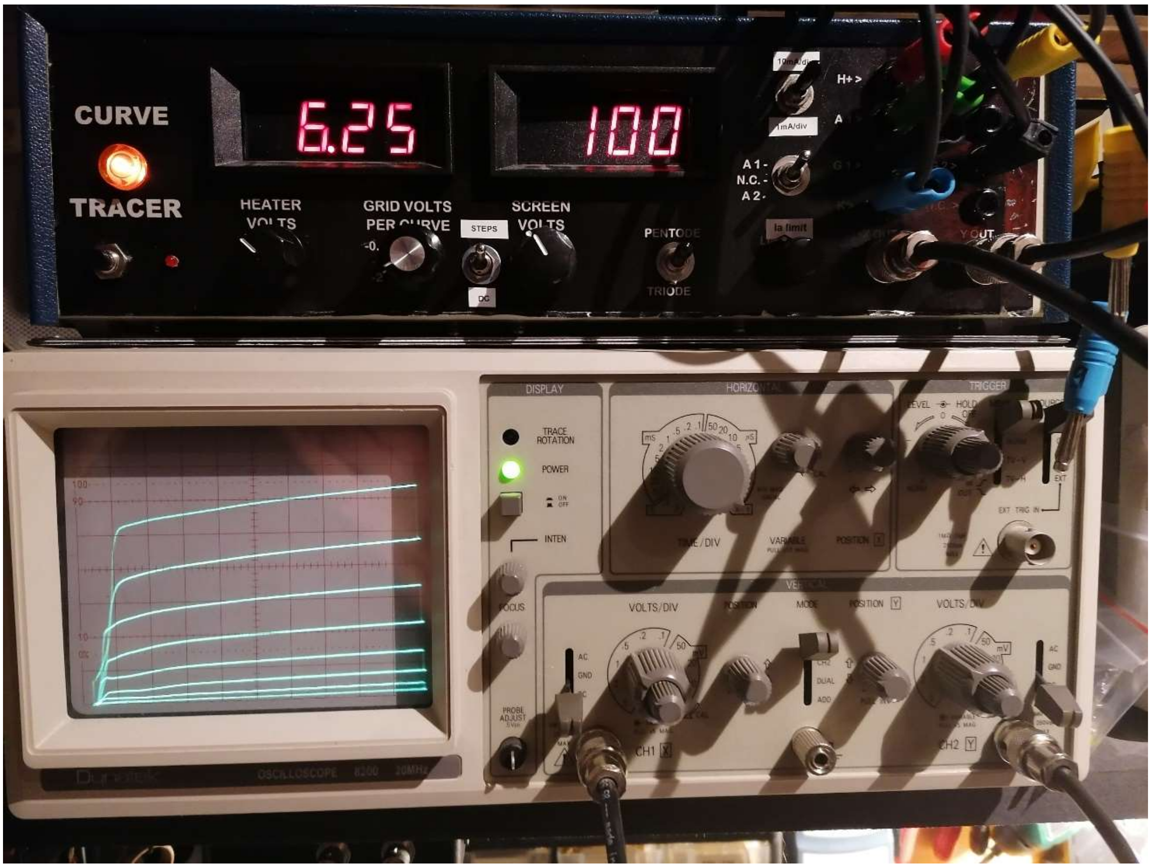

An EL34. You can clearly see the higher gm compared to a 6L6:

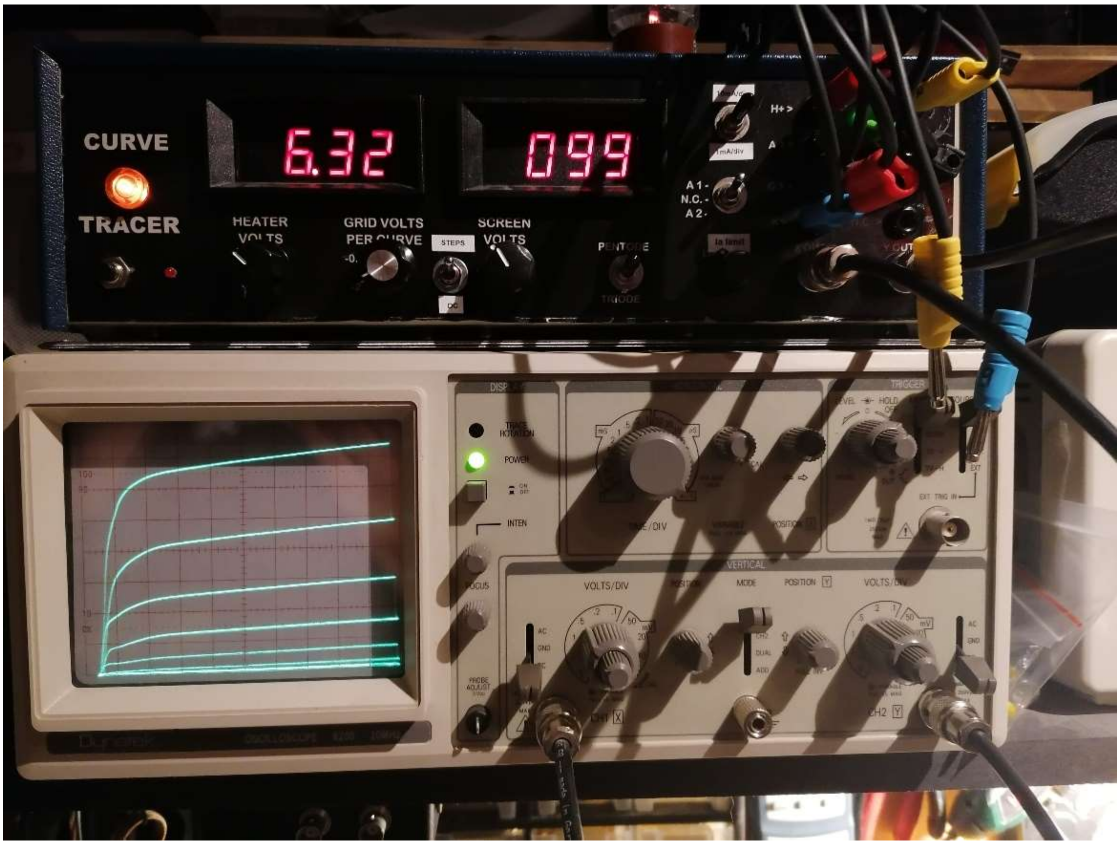

The curve tracer can also be used to test other high-voltage components. Here is what you see when you connect a 100V zener diode across anode and cathode terminals. I put a 47K resistor in series with it, since above the breakdown voltage, the zener is a very low impedance and current-limiting does not work (trips the protection relay). You can see the knee at around 100V (1 division corresponds to 30V). Current limited to 2mA:

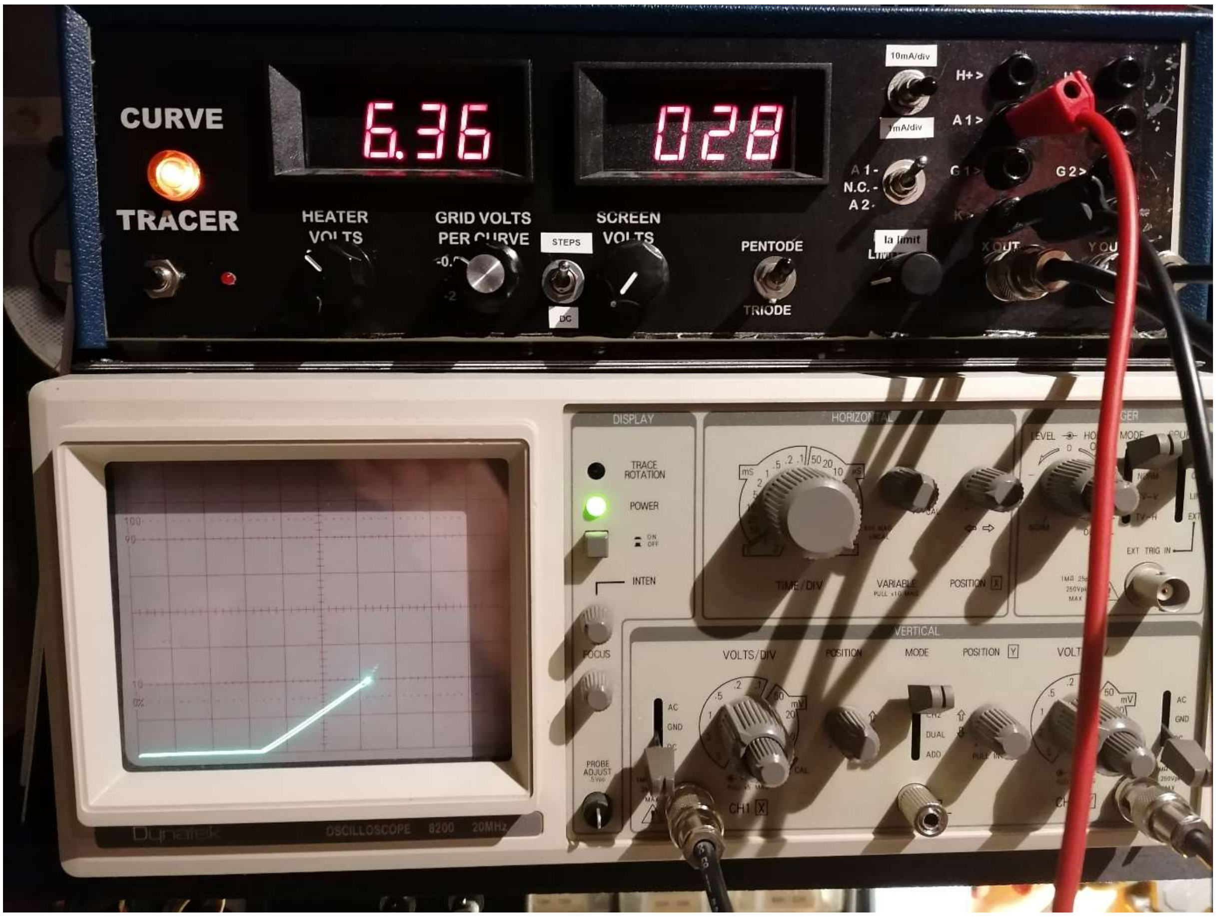

Top side. You can see I installed a couple of different sockets (standard noval, octal, a 7-pin socket, two sockets for very early tubes of a radio I restored a while ago).

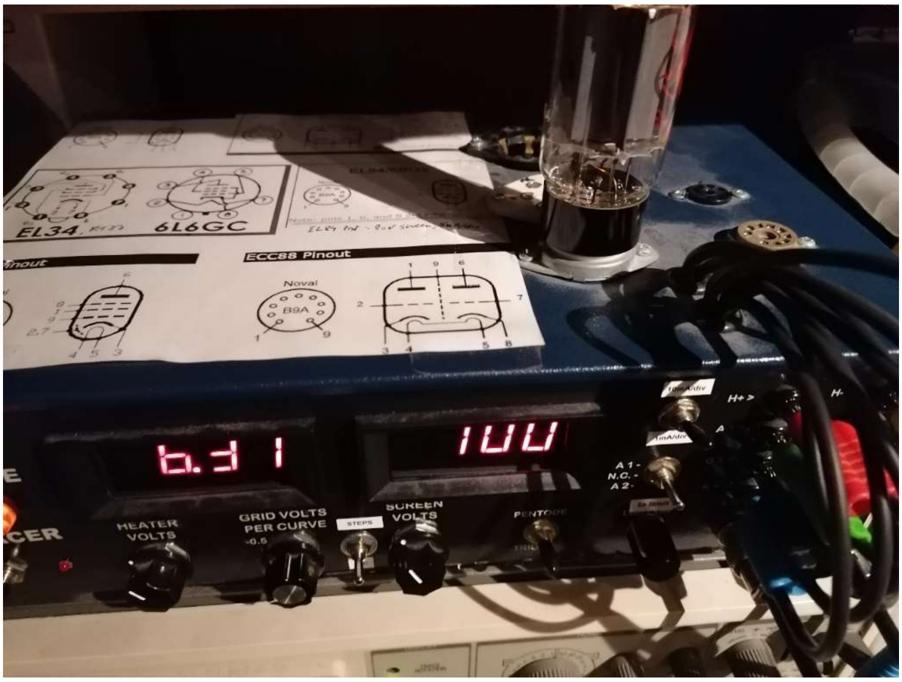

Front view: I have not yet made a new sticker. Used the old one and added some dymo prints where needed.