Valve stability tester

Designed and built December 2024-January 2025. Updates April/September 2026.

This text describes a mains-operated high-voltage device. Be extremely careful if it gives you ideas. It can kill you, especially when you think you know what you're doing, but even if you actually know what you're doing. So: I hereby disclaim any liability for damages should anyone be killed or injured while working on this project, or projects based on this design, or any other project or design presented on these web pages and any associated web pages.

Vacuum valves. They used to be ubiquitous, but are nowadays condemned to a life of abuse in guitar amps and rather specialised hi-fi audio gear. The outlet market for vacuum valves is quite small and only three factories still make them. The JJ valves are made in Slovakia, the TAD valves in China, and the rest of them (Sovtek, Svetlana, Mullard, Tung-Sol,…) are made in Russia. That's it. Much has already been said about the quality issues with some of these current-production devices. It seems that the rigourous quality control from 'back in the day' is very much a thing of the past. You can easily find several old documentaries from the 1940's and 1950's on Youtube, showing how valves were made then and how every device had to pass several stages of testing before it could leave the factory. I can imagine the modern standards of quality control are inferior to the standards held back then. It probably simply does not generate more profit; it costs more to keep the standards high and not enough people care.

Are current-production valves by themselves really inferior to the new-old-stock or NOS valves you can find on for instance Ebay? Not necessarily, if you ask me. There are EU directives that prohibit or limit the use of certain substances, which might explain a difference between devices now and before. More importantly, however, the economy dictates that profit should be maximised and quality control after production just costs too much. The production of vacuum valves is still quite artisanal, in that a considerable human effort is still required in the process, and it doesn't lend itself well to full automation (for the time being). The different electrodes in the valves are very delicate, tolerances (spacing, etc.) are tight, and fine assembly is done by human hands. This could explain how some valves from the same batch work fine for years, and others fail after a couple of hours or less. The human uncertainty factor.

The production process of valves and the design requirements are very well-known by now. I believe a correctly functioning current-production valve is just as good as any old device. I certainly do not believe that NOS valves are intrinsically better sounding. Many of the so-called NOS are probably used, or B- or C-stock from the past. The NOS label generates interest, creates a halo around the device, and thus inflates prices. The price does not reflect some inherent superior mystical property. I have had NOS valves that are very microphonic, or have low emission out of the box, … In addition, it is very hard to prove that this or that 12AX7 sounds better than the other, because experiencing that something sounds better is a subjective process by nature, and because the operating tolerances of one type of valve (even from the same manufacturer) are quite broad. So I can imagine someone installing a NOS 12AX7 that is somewhat depleted, and finding

©Jygsawaudio 2026

that it sounds better, while the device itself is objectively actually worse than any good current-production device. Because it has low emission, the bias point in the circuit will shift somewhat and that may well change the texture of the sound eventually coming out of the speaker.

So, assuming that I can't do anything about the internal construction of valves and the real problem is the lack of quality control, I made my own testing device or valve quality control department. It is true that there are a number of vacuum valve testers on the market, some of them really good, but many also not very reliable. Old testers will very likely require thorough servicing before they can be made to work and produce reliable results. Modern testers will give you a detailed account of how the device performs (static characteristic curves, emission, etc.). But still, a valve's life in a guitar amp is different from the couple of minutes it spends in the tester. Although the performance information is certainly valuable, what I really want to know is how the device performs at the conditions it lives in in a guitar amplifier, and that for hours on end. Many old testers also use AC mains-derived voltages (pulsating DC) to test valves, but that is not what the valves are subjected to in a guitar amp. So I made what I call a 'valve stability tester': put in a quad of 6L6's, set the quiescent current to the desired level and let them sit there for as long as I need or want, while I can read off the cathode current at any instant (and check if the valves are and remain well-matched).

The design goals.

- mimic the electronic circumstances of a generic guitar amplifier:
    - octal valves (6V6, 6L6, EL34, 6550,…): anode voltage V_a around 450Vdc, screen voltage around $V_a - 50V$. Mostly however the screen voltage equals V_a more or less, but that is a heavy burden for some valves and exceeds the datasheet maximum limits. So to keep things safe enough for every kind of octal and not damage any valves, I respected the datasheet limits. The exceedingly high screen voltage combined with too low-value screen stoppers has killed many valves.
    - noval (only EL84): anode voltage around 300Vdc, screen voltage around 300Vdc.
    - adjustable negative bias on the grids
    - 6.3Vac heaters
- keep it simple. No extensive range of controls, but only basic things:
    - switch between octal and noval operating conditions
    - adjustable negative grid bias
- ability to test four valves at once. Most guitar amps have four output valves or less. Old SVT's have six 6550's, but that is an exception. And SVT's have crazily high anode voltages on the valves (750Vdc). Scary stuff.
- measurement: a system for easy monitoring of the cathode current for each separate valve. In that way it can also serve as a matching device.

©Jygsawaudio 2026

- protection: if the cathode current exceeds a certain value, I need the circuit to cut power to the valves immediately and show me which valve drew too much current.
- allow for prolonged use (cooling, construction ruggedness,…).

I made an overview of some current-production valve types and their ratings and characteristics according to the datasheets:

Power tube data	**$Pa_{(max)}$**	**Vf_{nom}**	**If @Vf_{nom}**	**$Va_{(max)}$**	**$Vg2_{(max)}$**	**$Rg1_{(max)}$**	**$Ra\text{-}a_{(typ)}$** (2 tubes in push-pull)		**Ig/Ia**
EL84 (JJ)	12W	6,3V	0,76A	300V	300V	300kΩ/1Meg	8kΩ	*pentode*	10%
EL84-STR (TAD)	12W	6,3V	0,76A	420V	300V	300kΩ		*pentode*	
EL84-CZ (TAD)	12W	6,3V	1,5A	400V	300V	300kΩ		*pentode*	
6V6S (JJ)	14W	6,3V	0,5A	500V	450V	100kΩ/500kΩ	10kΩ	*pentode*	7%
6V6GT-CZ (TAD)	14W	6,3V	0,5A	500V	450V			*pentode*	
6V6GT-STR (TAD)	14W	6,3V	0,45A	450V	400V			*pentode*	
6L6GC (JJ)	30W	6,3V	0,9A	500V	450V	100kΩ/500kΩ	5kΩ	*pentode*	4%
6L6GCM-STR (TAD)	30W	6,3V	0,9A	550V	500V			*(long bottle)*	
6L6WGC-STR (TAD)	30W	6,3V	0,9A	550V	450V			*(short bottle)*	
5581 (JJ)	23W	6,3V	0,9A	500V	450V	100kΩ/500kΩ	5kΩ	*pentode*	8%
5881-Cz (TAD)	23W	6,3V	0,9A	500V	450V			*pentode*	
EL34 (JJ)	25W	6,3V	1,5A	800V	450V	700kΩ	3,5kΩ	*pentode*	13%
EL34-STR redbase (TAD)	25W	6,3V	1,5A	800V	500V			*pentode*	
EL34B-STR (TAD)	25W	6,3V	1,5A	800V	500V			*pentode*	
KT88 JJ	42W	6,3V	1,6A	800V	600V	100kΩ/270kΩ	4,5kΩ	*pentode*	5%
KT88-STR (blackplate TAD)	42W	6,3V	1,6A	800V	600V			*tetrode*	
KT88-STR (redbase TAD)	42W	6,3V	1,65A	800V	600V			*tetrode*	
6550 (JJ)	35W	6,3V	1,6A	600V	400V	50kΩ/250kΩ	4,5kΩ	*pentode*	9%
6550A-STR redbase (TAD)	42W	6,3V	1,7A	680V	450V			*pentode*	
KT77 (JJ)	25W	6,3V	1,4A	800V	600V	250kΩ/500kΩ	5kΩ	*pentode*	10%
KT66 (JJ)	25W	6,3V	1,3A	550V	550V	100kΩ/500kΩ	8kΩ	*pentode*	6%
KT66 (TAD)	25W	6,3V	1,3A	500V	400V			*tetrode*	

It can be seen that in order to accommodate all the octal types without more extensive controls (anode voltage, screen voltage, etc.) the safe maximum is an anode voltage of round 450Vdc and a screen voltage of around 400Vdc. So that is what I decided to go for. It is close to the conditions you will find in most guitar amps.

HT Power supply

1. *Transformer, rectifier, smoothing*

The choice of transformer was determined by the worst case conditions: a quad of 6550 valves under test. A typical 6550 valve is rated at 35W P_{max}. If I aim for an anode voltage of around 450V, then at a fairly typical 55% idle dissipation one 6550 would draw around $\frac{0.55 \cdot 35}{450} \approx 43mA$ of anode current. Add to that an extra 10% for the screen current.

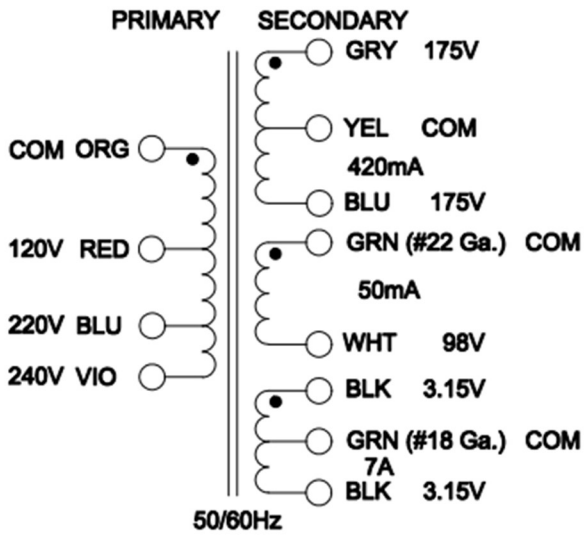
©Jygsawaudio 2026

Let's call that 50mA of cathode current for simplicity. Four of them then draw 200mA DC. As a rough (but reliable) rule, if 200mA DC is drawn by the load, then twice as much RMS ripple current will flow through the transformer-rectifier-smoothing cap combo. That means the transformer HT secondary must be able to handle 400mA rms current at $450V_{peak}$ or around $320V_{rms}$. That gives us a VA rating of at least 0.4 x 320 = 128VA.

I would also need a decent heater supply of course. Again, worst case scenario: four 6550's drawing 1.6A of heater current each, at $6.3V_{rms}$. That is 6.4A of constant heater current. Quite a lot.

Next, I would need a high (or low) enough voltage, preferably on a separate winding, to generate the negative bias voltage from. I would like to get at least -60V negative bias.

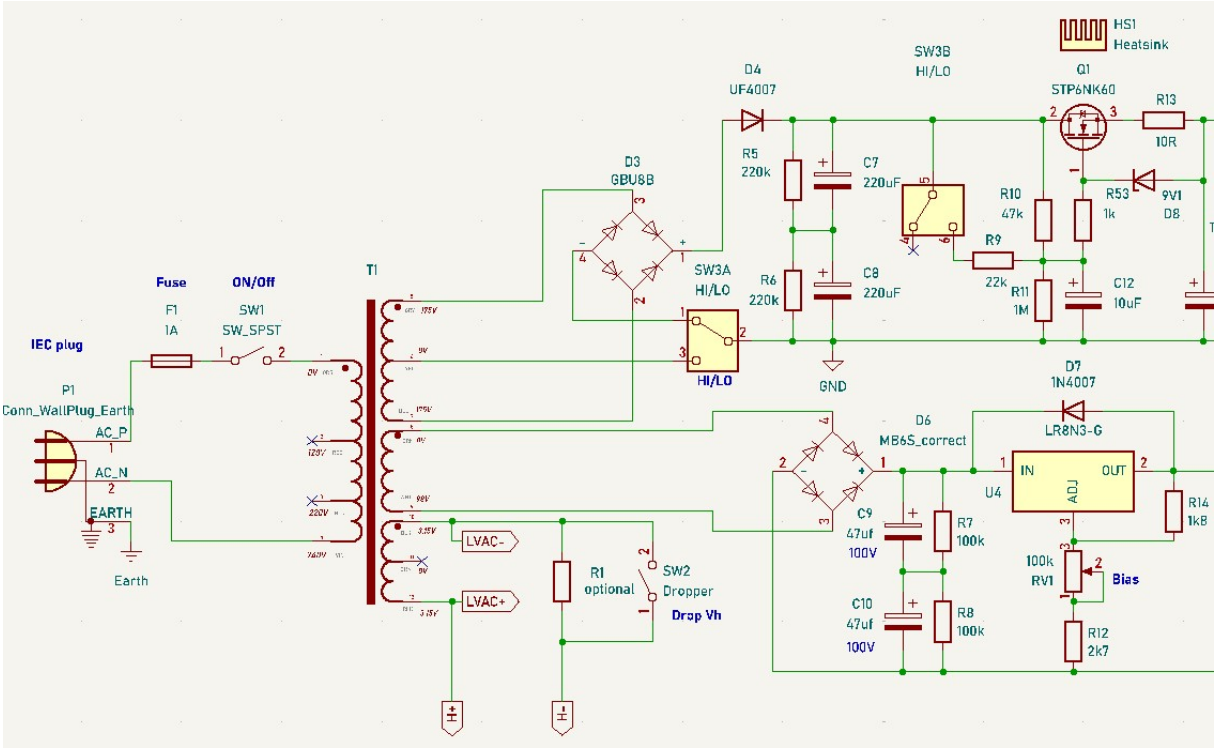
Browsing through the Hammond catalogue of high-voltage transformers, I found the 290HX. The primary can be wired for 120V/220V/240V mains. I will be using the 240V tap, since we have a nominal 230V mains here.



The transformer is rated at 150VA, which is enough for the purpose. The fact that the HT secondary is centre-tapped creates possibilities for making the DC supply switchable between hi and lo. At about 150€ it is not cheap, though. But transformers have become increasingly expensive the last couple of years.

The rectifier is a single package bridge rectifier RS1005M, rated at $600V_{RRM}$ and 10A continuous rms current. I always put a fast UF4007 diode after the rectifier. It minimizes switching noise as the rectifier switches on and off each cycle. Not that it matters much in this application, but it's standard practice for me. It also assists in protecting the bridge rectifier against the peak reverse voltage, which would be around 500V in this application. Smoothing is handled by two 220uF 350V radial caps in series. Balance resistors are used in parallel with each cap to ensure the voltage is divided equally between the caps. They are 220k 3W devices.

©Jygsawaudio 2026



In order to get two different voltages from this power supply I added a DPDT switch. In the position shown above, the centre tap is out of circuit and the rectifier acts as a bridge rectifier. The resulting RMS voltage across the secondary winding is therefore 175 x 2 = $350V_{rms}$. That gives a theoretical maximum of $350V \cdot \sqrt{2} = 495V_p$. It also needs to be considered that under no or light loads, this voltage will be even higher (since the 175 – 0 – 175 figure is the full load figure). More on that in a minute.

Switching SW3 to the other position, the centre tap is switched into the circuit and the leftmost two diodes of the bridge rectifier do nothing. The rectifier now functions as a full-wave rectifier instead of a bridge rectifier. Each half-winding conducts one half-cycle. The 175-0-175 winding and following rectifier should now produce a voltage around $250V_{dc}$ or thereabout. That is adequate for EL84's. It is not the desired $300V_{dc}$ I had in mind, but to avoid adding any more complexity I left it at that.

It would be a bad idea to switch between modes when the power supply is live. During the couple of milliseconds both switch contacts are open (and bouncing on and off!) the secondary return current is cut off. It's a transformer and has substantial inductance. High voltage spikes are possible between windings. Under normal use however, it would not be necessary to switch when the device is live. Just don't.

©Jygsawaudio 2026

As can be seen in the figure below, the estimated 400mA rms ripple current through the transformer's HT secondary at 200mA DC load current is quite accurate. The peak current even reaches 1A.

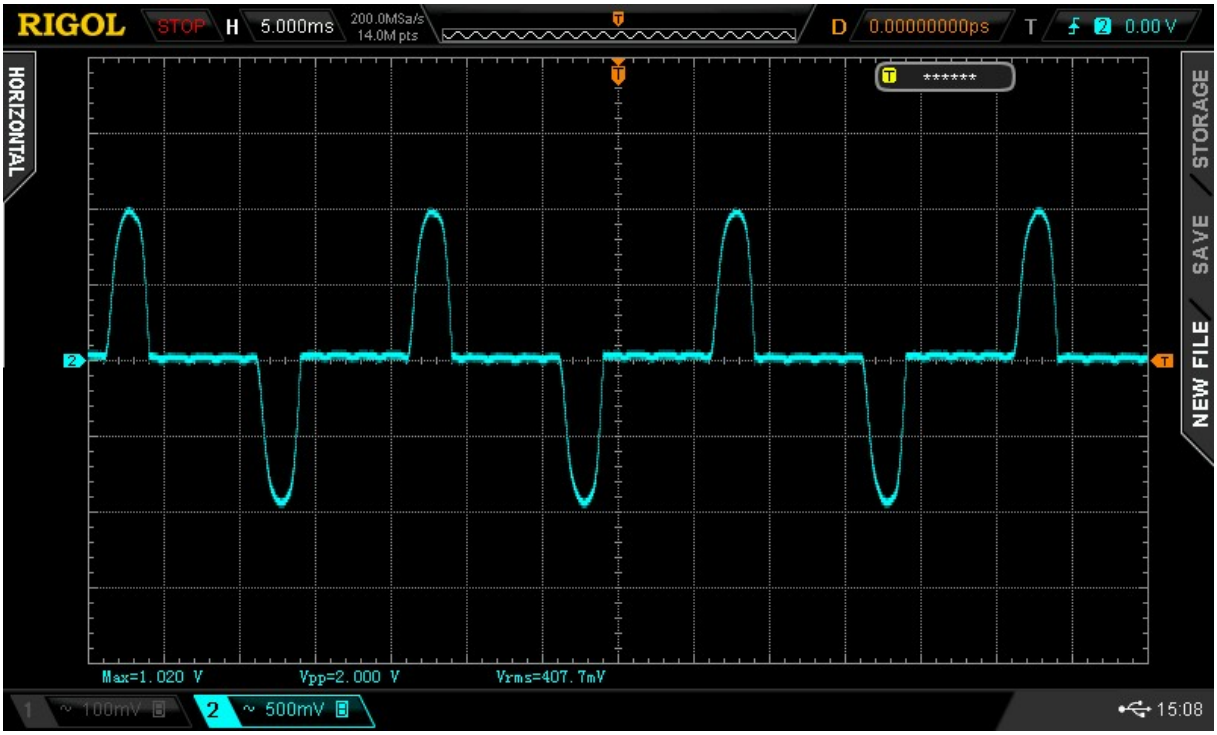


Figure 1: ripple current in HT secondary at 200mA DC load. 1mV equals 1mA on this scale.

2. Voltage stabilisation

At this point I could choose to add a voltage stabilising circuit to produce a fixed HT, but there are problems to consider. First, the switchable HT complicates matters. I would need two different voltages, adding complexity to the circuit. It can be done and is not so hard (zener diodes as a reference, switching between two values, etc.), but I wanted it to be simpler. Second, if a fixed reference is used, what exactly should it be? Too high and the stabiliser may drop out of regulation at a high current draw from a quad of 6550's, or too low and the pass transistor will get very hot at a relatively high voltage drop and higher current draw. So I decided to let go of the idea to fix the output to one value, and let the output of the stabiliser follow the input. The circuitry around Q1 accomplishes that.

When SW3 is switched to the HI position, as shown in the figure, raw DC is applied across R10 and R11. With R10 = 47kΩ and R11 = 1MΩ the voltage at the gate of Q1 is 95.5% of the raw DC at the drain of Q1. I decided to drop some HT since, under no or light loads, the HT will exceed $500V_{dc}$. Suppose the raw DC is 500V, then 477.5V is applied to the gate of Q1. That puts 22.5V + V_{gs} across Q1. V_{gs} will typically be about 5V, so 27.5V is the voltage drop across Q1. At the maximum current draw of $200mA_{dc}$ Q1 then dissipates $27.5V \cdot 0.2A = 5.5W$. A heat sink will be needed to help the transistor dissipate that power.

©Jygsawaudio 2026

These figures are all a bit overrated, because the transformer primary will be wired for 240V use, and the nominal mains voltage is 230V here. Excursions to 240V are certainly possible, but not constant. Additionally, current draw will cause voltages to sag a bit.

In the LO position SW3B puts resistor R9 in parallel with R10. That is because the raw DC voltage is now only 250V and there is no need to drop much of it. Together they make 15kΩ of resistance to filter out the 100Hz in conjuction with C12. Very little voltage is dropped.

Capacitor C12 is a 10uF 500V cap (another reason to drop some HT in the HI mode: avoid having to use 2 HT caps in series there). It smoothes out the ripple at the gate of Q1. Thanks to the high value of R10 = 47kΩ, this is an efficient ripple filter. At 100Hz (rectified DC pulses) C12's reactance is only $\frac{1}{2\cdot\pi\cdot100\cdot10\cdot10^{-6}} = 159\Omega$. That makes this a very efficient filter.

R53 is a gate stopper. Without it Q1 will oscillate at some point.

C13 is another high voltage cap to stabilize Q1's output.

R13 and D8 form a basic current-limiting circuit. As soon as the current through R13 exceeds a certain value, D8 will start conducting, stealing current from C12, lowering the voltage at Q1's gate, thus limiting current through the load. D8 is a 9V1 zener diode rated for 1W. It is put across Q1's gate and the output of R13. That means the voltage across D8 under normal circumstances is Q1's V_{gs}(≈5V) and the voltage drop across R13. So as soon as this voltage drop exceeds around 4V, D8 conducts and current-limiting starts. R13 is 10Ω, so maximum current is $\frac{4V}{10} = 0.4A$. At 400mA Q1 starts to limit the current.

There is a possible snag here: if you connect D8 straight to C12, before the gate stopper, in case of a dead short at the output C12 will discharge straight through D8, probably destroying it. So I put D8 after the gate stopper, so in case of a shorted valve and the ensuing current surge, D8 has a chance of survival since C12 can only discharge through the gate stopper before D8. 10uF is not that much capacitance, so I don't really expect the gate stopper to burn. The protection circuitry should kick in anyway. *Edit: the zener failed when a short happened in a DUT (EL84). In the first version the diode was a 1/2W device. I then replaced it with a more rugged 1W device and haven't had any failures since then*.

The current-limiting feature is only intended to protect all the following circuitry from possible shorts in the valves under test. I have had mosfets fail by sudden current bursts caused by faulty valves before, even when there is short-circuit protection elsewhere in the circuit. There is always some response time of maybe only a couple of milliseconds, but that is enough to destroy a sensitive device. This simple current-limiting requires only two cheap components and works very well for this purpose.

©Jygsawaudio 2026

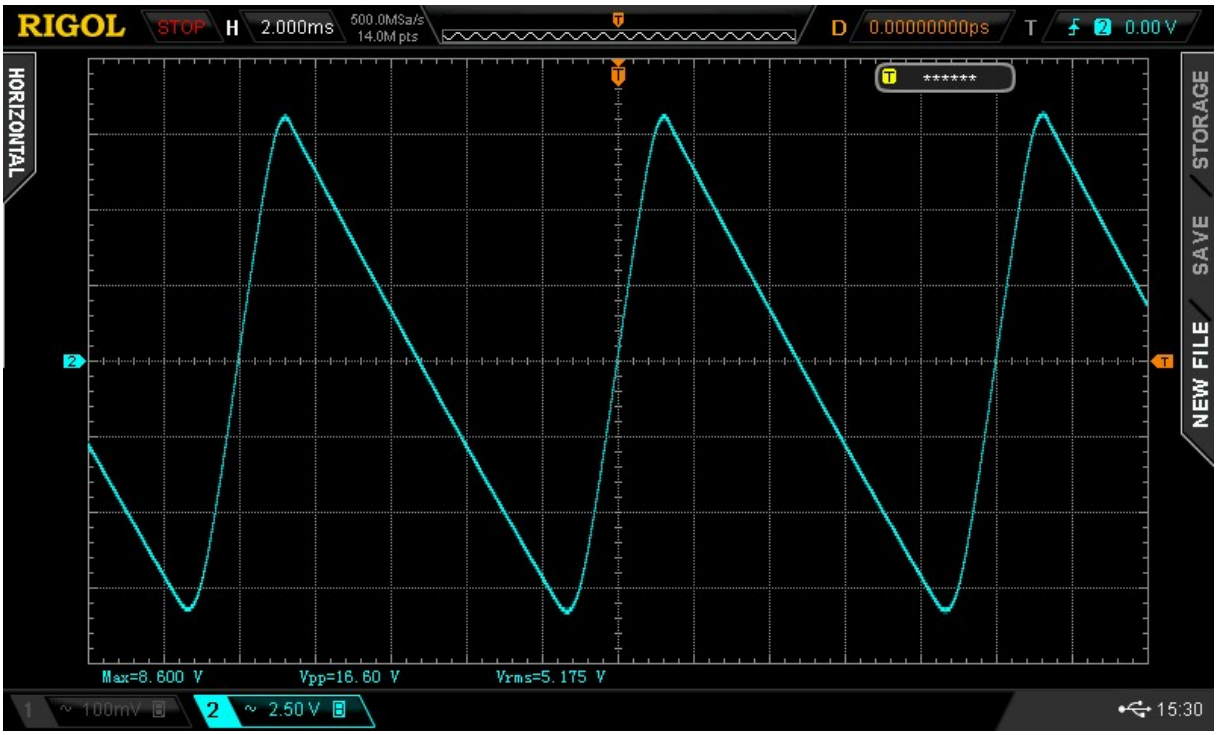
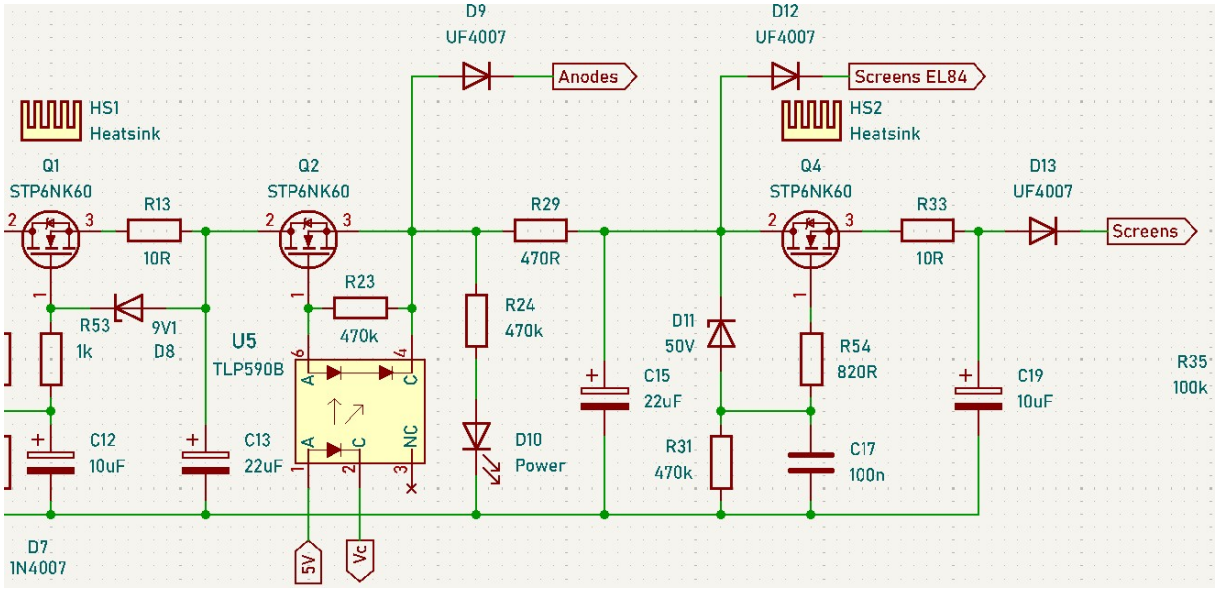


Figure 2: ripple voltage at Q1 drain at 200mA DC load.

I measured the ripple voltage at the B+ node (drain of Q1) at 200mA of load current. See above. 16.6V is quite a lot and some of this ripple is present on the source of Q1, because the regulator drops out of regulation at the valleys of the ripple. It's a matter of making a compromise and weigh up the pros and cons of either choice: clean but more voltage drop, less voltage drop but some ripple. Since we're not building an actual amp here, the latter is acceptable.

3. *Anode and screen voltages*

The stable DC voltage can now be applied to the anodes and screens of the valves under test. I use another mosfet (STP6NK60, rated at 600V_{dsmax} and 6A continuous current) to act as a switch. I want to be able to cut off the supply to the anodes and screens in case of



©Jygsawaudio 2026

a short between electrodes in a DUT. This is achieved by the optical mosfet gate driver TLP590B. The brilliant thing about these devices is that they work very efficiently without any extra circuitry needed to handle the high voltage that would otherwise be needed for driving mosfet gates in these high voltage circuits. The control circuitry is completely isolated from the HT circuitry. Alternatively, you could insert that mosfet on the low-voltage side (the circuit ground return line) to make it easier to connect control circuitry around the gate and source, but that would keep the high voltage on the anodes and screens, even though the return line is open. With the tube sockets exposed as they are, it seems safer to open the HT line itself in case of a fault.

When current flows through the LED side of the photocoupler, a voltage is generated across the detector (photodiodes), biasing the mosfet gate positive with respect to the source and turning it on. The detector diodes can only deliver microamps of current, but virtually no current is needed to bias the gate (except to charge the parasitic gate-source capacitance). The datasheet recommends 20mA of forward current through the LED. This results in gate-source bias voltage of +10V, which is enough to turn Q2 fully on. Drain-source resistance R_{dson} with $V_{gs} = 10V$ is minimal (typically 1Ω). R23 is necessary to decrease the turn-off time. Without R23 it takes Q2 too long to turn off. According to the datasheet turn-on time is 0.2ms (from 1 to 5V), turn-off time with a 510kΩ shunt resistor is 1ms (5V to 1V), with a load capacitance (i.e. the gate-source capacitance of the mosfet) of 1nF. There are faster devices with integrated turn-off circuitry such as VOM127. Turn-off time according to the datasheet is 65us, but that is for a load capacitance of only 200pf, so turning off a power mosfet will probably take longer due to the larger capacitance. Anyway, TLP590B with shunt resistance works fine. I used 470kΩ here as a shunt resistor.

Q2 will be normally on, applying voltage to the anodes and screens. When a fault condition is detected, the current through the gate driver's LED is cut off and Q2 goes into cut-off. See later.

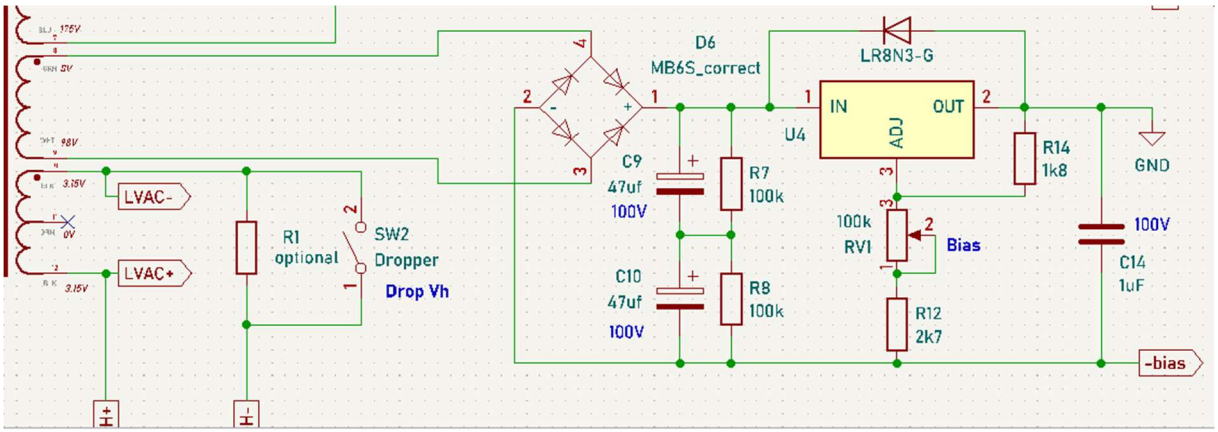
At this point I inserted a 'power on' LED to let me know that high voltage is present at the anodes.

Next up, the screen supply. Again, a STP6NK60 is used as a pass device. I wanted to keep the screen voltage around 50V lower than the anode voltage, because some of the valves I intended this device for have a screen voltage maximum rating of 400V. Just to stay on the safe side of things, I fixed the screen voltage for all octal valves to $V_a - 50V$. Since the anode supply in LO mode is only 250V, I tapped off the screen voltage for EL84's from the anode supply with a simple resistor and extra capacitor, mimicking the usual guitar amp configuration. Q3 is best heat-sunk as well, but a simple clip-on heatsink should suffice. The 1N4007 diodes prevent current to flow into the screen nodes in case of a short between anode and screen in the DUT.

©Jygsawaudio 2026

Negative bias power supply

The 290HX has a separate winding for low-current circuitry. Ideal for a negative bias supply to the control grids. Since this is an isolated winding, all circuitry can be treated as if it were a positive voltage supply. Referencing the 'positive' output to ground creates a negative supply from the viewpoint of the HT circuit.



Nothing really special here. The bridge rectifier is a MB6S surface mount device. The raw DC peaks at around 145V, so to avoid a larger high-voltage capacitor I connected two 47uF 100V-rated caps in series, each with a balancing resistor. Extensive filtering there is not needed, since I use a high-voltage linear regulator LR8 for the voltage regulation. This device is only intended for very low-current applications. It comes in a TO-92 package and cannot dissipate much power. It has a maximum differential voltage rating of 450V and a minimum of 12V (the input must always be >12V higher than the output). At this minimum V_{diff} the output current is limited to about 20mA. Obviously, the higher V_{diff}, the lower the amount of current the LR8 can deliver. It has internal current-limiting circuitry to prevent overheating.

The output voltage is adjustable and is determined by the combination of R14, RV1 and R12. The voltage difference between the output and adjust pins is a stable 1.2V, so R14 has a constant 1.2V across itself, resulting in a constant current of $\frac{1.2V}{1800\Omega} = 0.67mA$. For good regulation the LR8 requires at least 0.5mA of output current. This then flows through RV1 and R12, together with a 10uA adjust pin bias current. RV1 = 100k (full CCW) and R12 = 2k7, so the output voltage at full CCW will be $1.2V + (0.68mA \cdot 102700\Omega) = 71V$. This is ample voltage to drive any common power valve into cut-off. The minimum voltage is 1.8V.

The minimum output capacitance according to the datasheet is 1uF, and that is what I used. The output remains nice and stable. I like to keep the capacitance as low as I can, so when I turn the bias pot the voltage tracks well. Too large a capacitance can cause the voltage change to lag behind the turning action and that is annoying.

As said before, reference the positive output to ground and take the output from where you would normally reference to ground and you have a negative bias supply. Thanks to the galvanic isolation of separate transformer windings!

©Jygsawaudio 2026

Heater supply

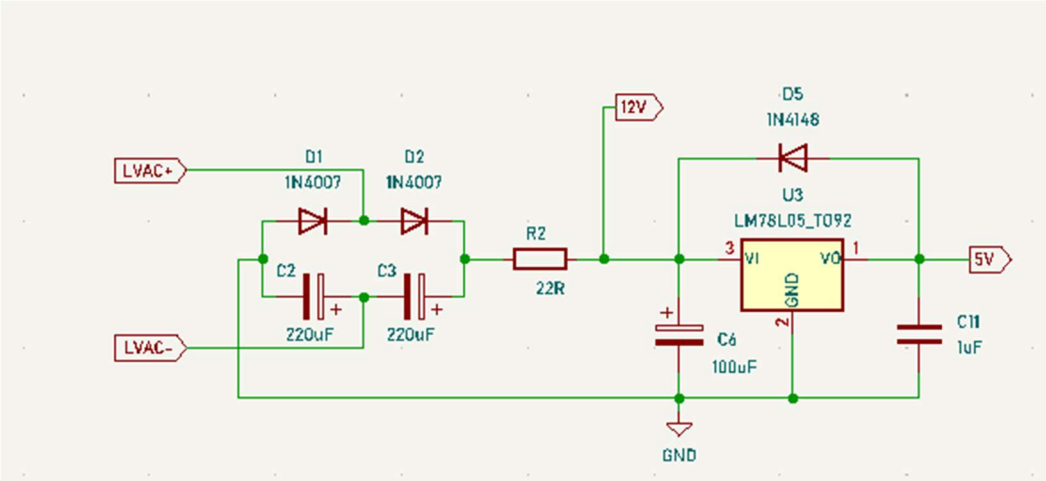
I kept things very simple here. AC current will do fine. I initially thought I might have to drop some voltage for lower power valves such as 6V6's (0.5A rated heater current), but since the mains voltage here is mostly 230V and the 290HX is wired for 240V operation, the heater voltage is slightly lower than nominal. It stays within the limits of ±10%. So I dropped the dropper and the switch.

You may notice there is no ground reference (centre tap left unconnected). That is because of what follows: I needed another supply for opamps and microcontroller.

Low voltage DC power supply

To power the opamps and microcontroller in the circuit I tapped off of the heater winding. $6.3V_{rms}$ isn't much to work with though, so I used a voltage doubler to get around $12V_{dc}$ to power the fan and semiconductors. R2 and C6 give some extra smoothing. That is followed by a 5V regulator to power the microcontroller, temperature sensor and opamps.

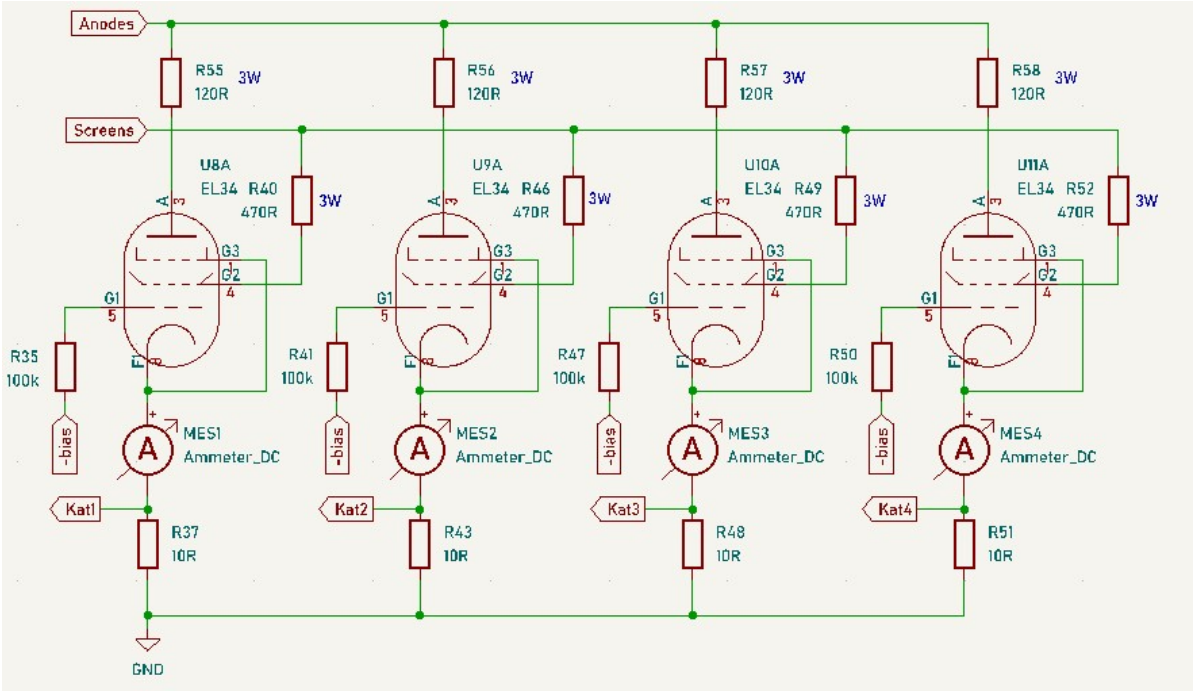
The 6.3V winding on the transformer sits at about half the raw DC voltage because of this configuration, so ground-referencing the centre tap is not a good idea in this case. D1 would probably not survive for long, and C2 would be subjected to substantially more AC current as well.



The valves under test

Anode voltage is applied straight to the anodes. Each of the screen grids, however, must have some resistance isolating them from one another. I chose the common guitar amp value of 470Ω (which may be on the low side for some true pentodes like EL34, in which case a value over 1kΩ would be preferred). Without this resistive 'decoupling' between screen grids, there is a very good chance that the system will become unstable and start oscillating. The figure below shows the octal valves alone. The same configuration is used for the noval sockets. Screen grid resistors are mounted on the sockets.

©Jygsawaudio 2026



The control grids each have 100kΩ stoppers. That is the maximum value for several octals such as 6L6's in a fixed bias configuration, so I chose that as the generic value for all valve types. If one of the valves gives excessive grid current as it heats up, the voltage developed across the stopper will make the valve conduct more and that will show clearly on the meters.

Update 09/2026: with four 6L6's under test the circuit was still prone to VHF oscillations at higher anode currents. I would notice that, when increasing the bias voltage – i.e. making it less negative – to a certain point, there was a significant drop in cathode current. After some measuring and head-scratching, the scope revealed parasitic oscillations and a drop in anode voltage due to current-limiting. The cause: the anodes were daisy-chained and a transmission line effect occurred where the stray inductances and capacitances due to the poor wiring scheme conspired to cause large oscillations. The solution: I removed the daisy chain wiring, soldered in 120ohm 3W resistors to the anode lugs on each socket pair (octal/noval) and in so doing converted to a star wiring configuration as well. The oscillations are gone now. I also kept the anode wires well away from the grid wires and tucked close to the grounded aluminium top plate.

A DC ammeter is put in series with the cathode of each valve. The scale of these meters goes to 100mA with a 4mA subdivision, which is exactly right for this application. I don't expect to have to push any valve beyond say 75mA when testing. The ammeters are analogue types to avoid the need for yet more power supply trickery. I calibrated each meter before installing them in circuit. Valves matched to within 2mA are ok, and that can easily be read off these meters.

Yes, I measure the cathode current and not the anode current without screen grid current. This test equipment is not intended for very precise measurement. It is intended for monitoring the DUTs' performance under prolonged stress. I could have connected

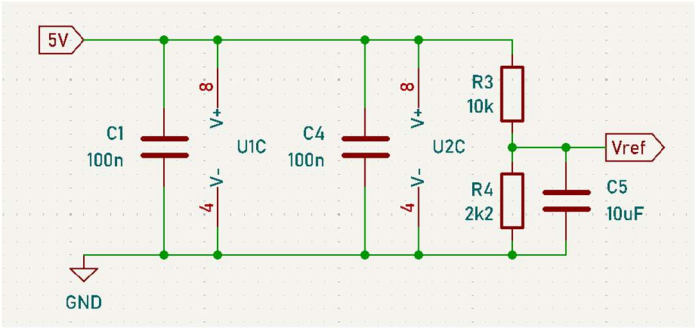
the meters in series with the anodes, but in terms of safety that is a poor choice and unnecessary in this case. It's safer to have them on the low voltage side of things. In another tester I built, very different from this one, the aim was to make precise measurements of anode and screen current, and the meters are connected in series with anode and screen respectively.

I connected a 10Ω resistor in series with each cathode to ground to monitor the cathode current. A microcontroller keeps an eye on the voltage across each resistor and as soon as the voltage exceeds .9V (i.e. 90mA of cathode current), the micro cuts off power to the photocoupled mosfet gate driver, cutting off power to the anodes and screens.

The microcontroller

1. *Current monitoring*

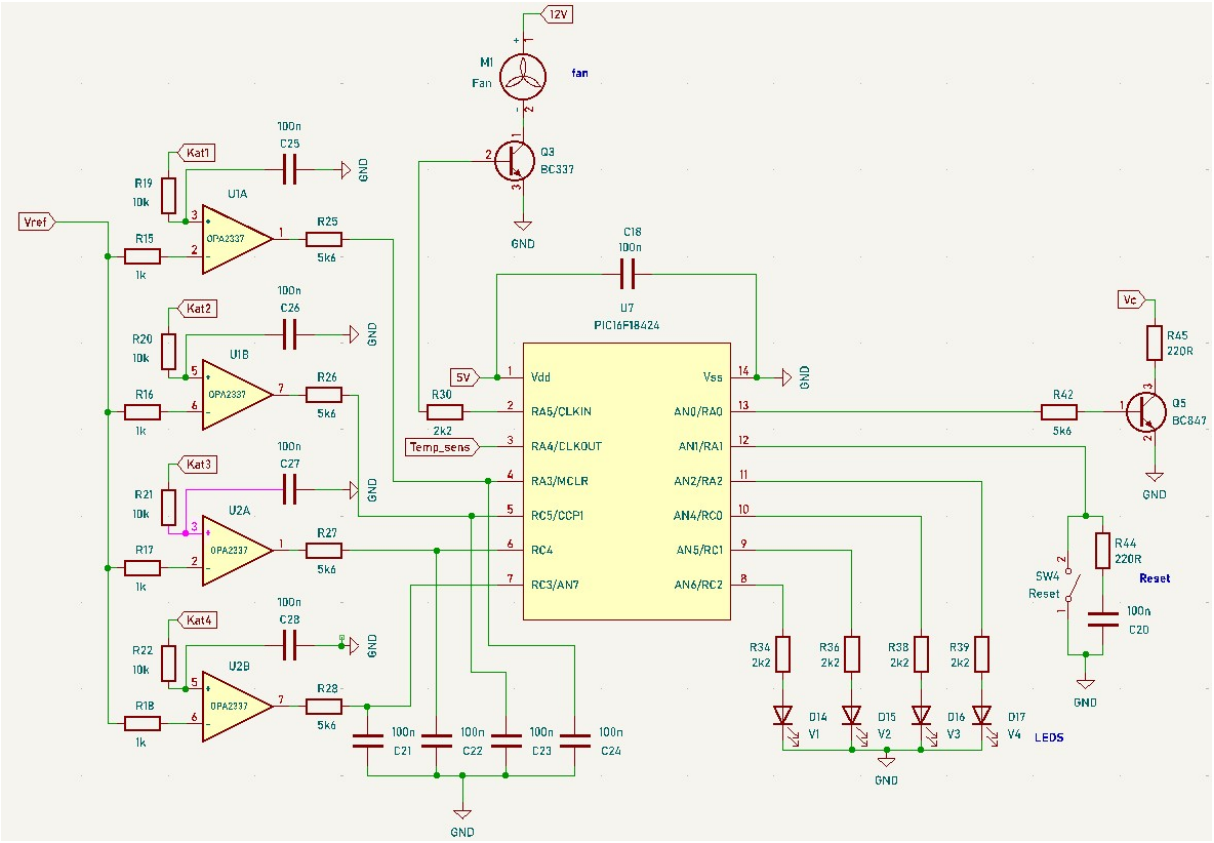
As said before, a microcontroller is used to monitor the current through each valve. I couldn't connect the cathodes directly to the microcontroller since the 'tripping' voltage is only .9V. I used a pair of OPA2337 opamps as comparators to sense the voltage at the sense resistors. The OPA2337 is a rail-to-rail opamp, CMOS inputs, low-supply-voltage (5.5Vdiff max). So the output of the comparators will swing between 0V and 5V in this circuit. The PIC16F18424 senses the output of the comparators. The inputs are configured as Schmitt trigger CMOS inputs and respond to CMOS logic levels: 0 is anything below about 1V, 1 is anything above about 4V. It is possible to configure the inputs to respond to TTL logic levels (<0.8V is LO, >2V is HI), but that increases the chance of false triggering somewhat.



R3 and R4 form a voltage divider to set the reference voltage for the comparators at 0.9V. A soon as one of the cathode voltages exceeds 0.9V, the opamp outputs 5V and the microcontroller brings its pin 13 to ground, shutting down HT power to the valves.

©Jygsawaudio 2026

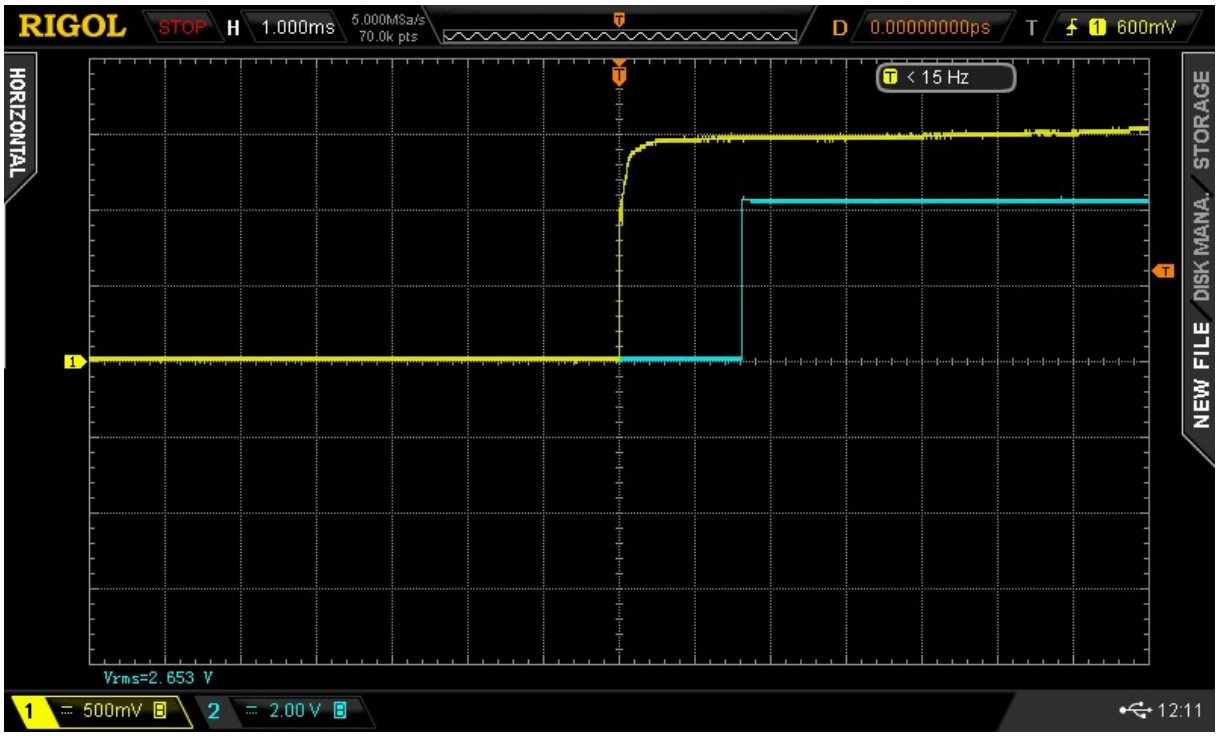
As can be seen on the schematic below, pin 13 connects to the base of a BC847 transistor through a 5k6 resistor. In normal operation pin 13 is high and Q5 passes current through



R45, a 220Ω resistor, to activate the photocoupler LED. This LED is powered from the 5V supply. It has a forward voltage drop of 1.5V, so 3.5V is dropped across R45, which means $\frac{3.5V}{22} = 16mA$ is the forward current through the LED. That is enough to generate around 10V of bias voltage in the photodiodes.

©Jygsawaudio 2026

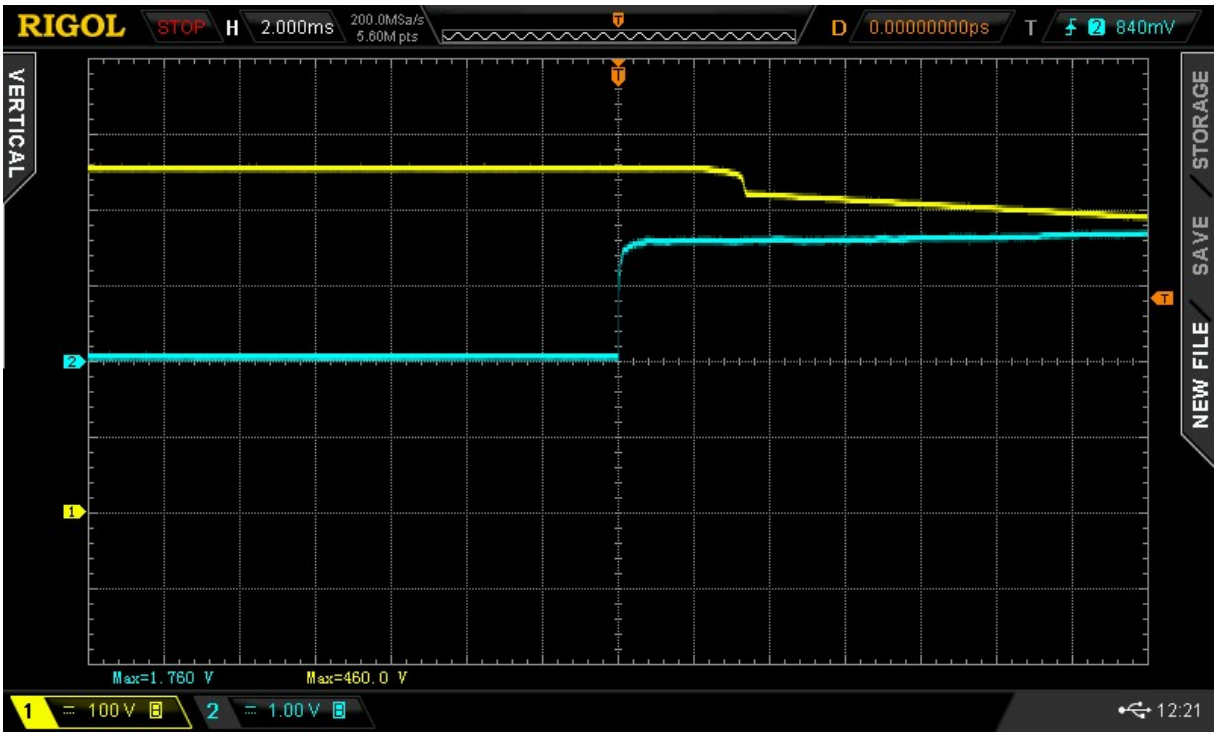
When a fault condition occurs, a LED is lit to indicate which valve failed. Power stays off until the reset button is pressed.

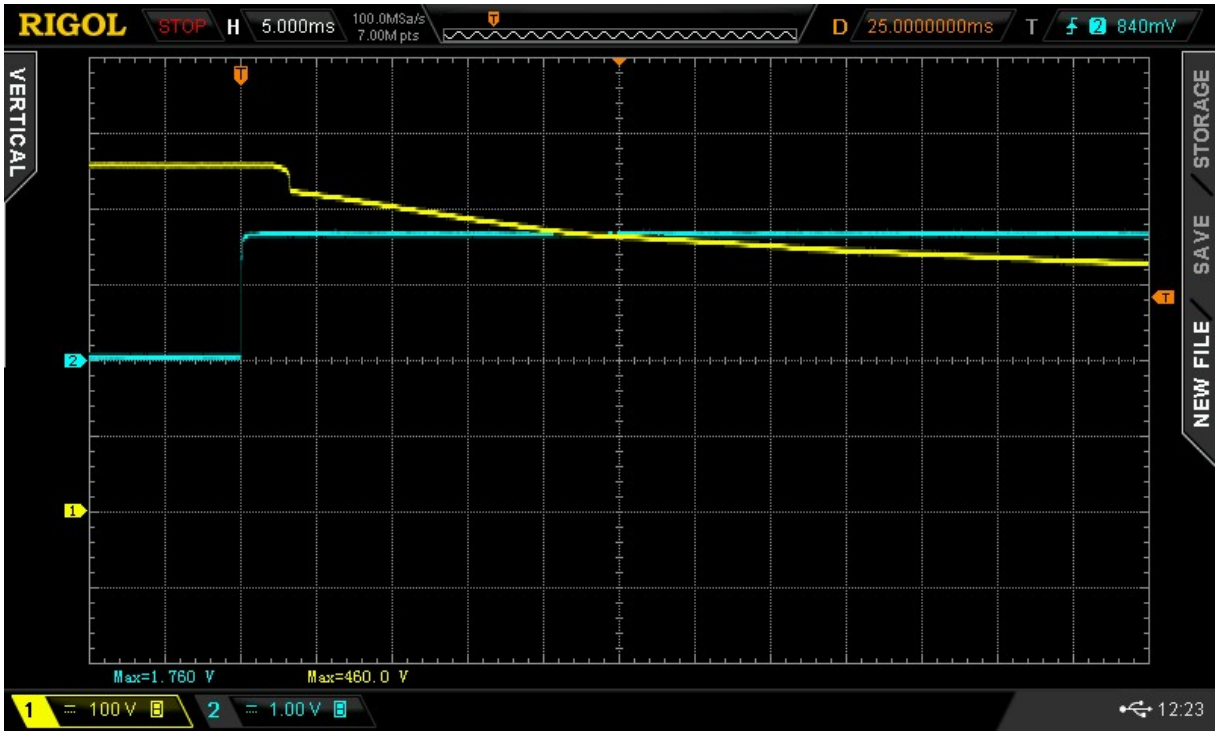


The picture above shows the response time. The yellow trace is a 1.5V_{dc} voltage applied to the cathode pin of one of the sockets. The blue trace is the voltage at the collector of the BC847 that controls the current through the optoisolator LED. When collector current is cut off, its voltage goes high. You can see that it takes about 1.5ms to respond to a fault condition. Most of that can be accounted for by the capacitors at the comparators' non-inverting input and output. They serve the purpose of removing any noise causing false triggering of the comparator and microcontroller, but they do slow down the response time, obviously. It's a compromise. I have measured the response time of the microcontroller itself to a pulse at the sensing pins. There is only 50 nanoseconds time difference between the rising edge of the pulse and the rising edge of the response signal (to the base of the BC847).

The pictures below show the response time again, but now the blue trace is the cathode voltage and the yellow trace the anode voltage. You can see that it takes about 3ms for the anode voltage to be effectively cut off. I had two 6L6's in circuit for this measurement. The anode voltage drops off gradually because there was no real short. It can be deduced that the optoisolator takes 1.5ms to bring the pass transistor in cut-off.

©Jygsawaudio 2026





In case of a real short in a valve, the voltage across the sense resistor may wel be above 5V. Although the current through the main pass transistor is limited to 400mA, the caps will discharge through the shorted valve and a large pulse of current is certainly possible. That will increase the rate at which the capacitor at the non-inverting input charges, making it reach the .9V threshold sooner and trip the comparator. The OPA2337

©Jygsawaudio 2026

has internal clamping diodes to protect the CMOS inputs. A current-limiting resistor remains necessary, however, to avoid damage to the opamps. That is 10kΩ in my circuit.

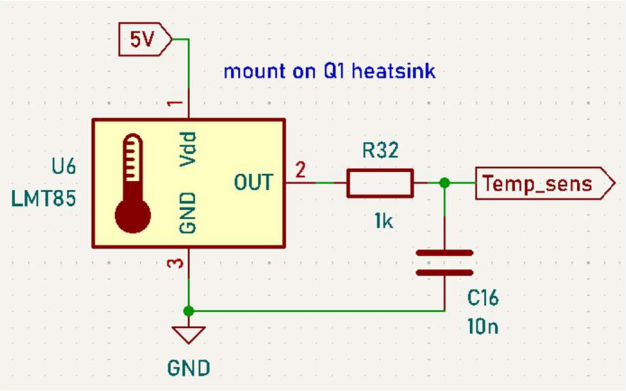
The microcontroller programme. An important consideration when configuring the microcontroller was reliability and speed. I wanted to make sure I wouldn't lose any more fault response time. The PIC16F18424, as many of the 8-bit PICs, contains a Configurable Logic Cell. I used this peripheral to handle the sensing and triggering. It is part of the hardware inside the PIC and once configured and initialised works independently of the micro's CPU. That ensures that there is very little time lost and that its operation remains consistent, independent of any interrupts or other software stuff.

The four sensing inputs are sent to a gate, configured as an OR gate. These inputs are normally LO or at logic 0. When a fault condition occurs and one of these inputs goes HI, the OR gate output goes HI. This output then goes to the RESET input of a S-R latch. The latch output then drives the base of the control transistor. When a logic 1 appears at the RESET the output latches LO, no matter what the condition at the SET pin is at that instant. Since power is cut off to the valves, current will drop fast and the comparator output will go LO again. So when anything goes wrong, it is sure that power to the anodes will be cut off and stay off. Only when RESET is LO and a HI pulse is applied to the SET input is power restored. That is done by means of a button switch. Pressing the button sends a logic 1 to the SET input.

I installed four LEDs to indicate which of the valves failed. This is done with an interrupt-on-change. Response speed is not important here, so this is done in software. Pressing the button resets all the LEDs.

When the temperature of the main pass transistor heat sink exceeds 100°C, power to the valves is cut off as well, and can only be reapplied after switching off the entire tester. All the LEDs are lit in this condition.

2. *Temperature sensing*



I wanted to keep an eye on the temperature of the main pass transistor Q1. When a quad of high-power valves are under test, this one can get pretty hot after a while. Since I had a spare output and input on the microcontroller, I decided to clip an LMT85 temp sensor on Q1's heatsink. This sensor in a TO-92 package outputs a decreasing voltage as temperature increases (negative temperature coefficient). At 60°C the output is about 1.07V, at 100°C 0.74V. The microcontroller monitors this voltage by means of an internal ADC. As soon as

©Jygsawaudio 2026

the voltage drops below 1.07V, the fan turns on. Should it drop even further, beyond 0.74V, the micro will shut down power to the valves. It's a failsafe. I really don't expect that to happen, except maybe when the fan fails miserably.

The sensor is clipped onto the heatsink right next to Q1. I used some heatsink compound for good thermal coupling.

Test data

No load anode voltage HI	463V
200mA DC load anode voltage HI	424V
No load anode voltage LO	235V
120mA load anode voltage LO	220V
No load voltage at reservoir cap	490V / 245V
Full load voltage at reservoir cap	460V / 225V
Fan voltage	13,4V
Heater voltage full load	6,2Vac

All measurements taken at 230V mains.

©Jygsawaudio 2026

Pictures



Figure 3: the control panel with (from left to right) the HI/LO switch, a needlessly drilled and capped hole, the bias control pot, ON led and OL leds, reset button. Four DC ammeters corresponding to each valve socket from left to right.



©Jygsawaudio 2026



Figure 5: a hole to pull air from for the fan on the other side of the box. The handle is placed on the side, because I intend to store it on its side.



Figure 4: the fan sits behind the grille. Rubber feet mounted on this side for storage. Reason being that the heavy transformer is mounted on this side.

©Jygsawaudio 2026



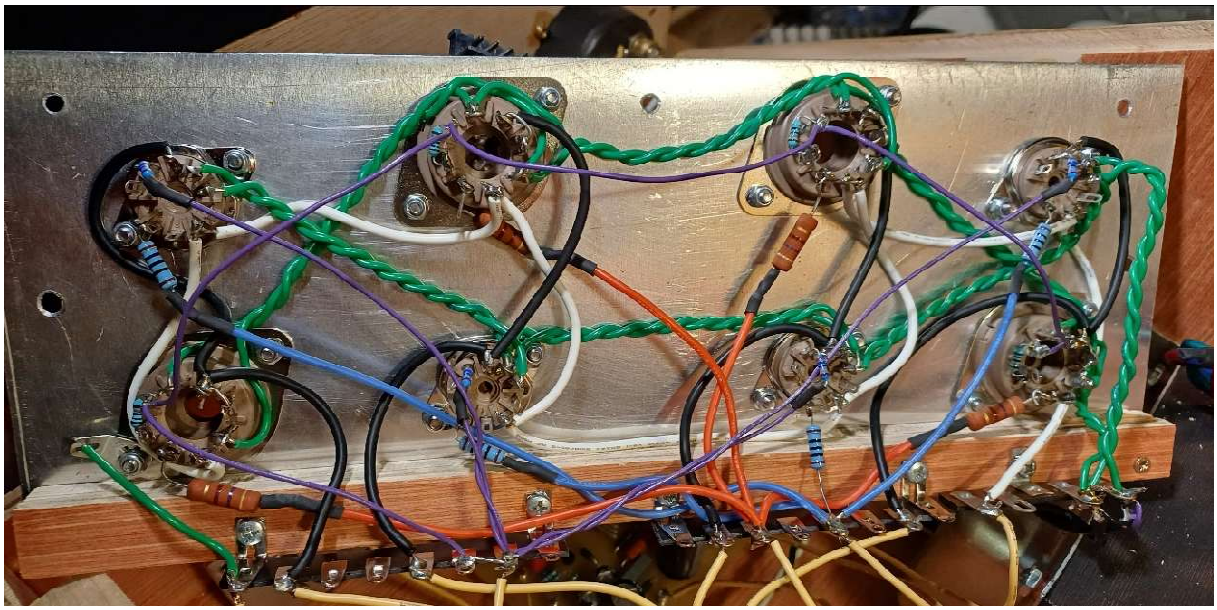
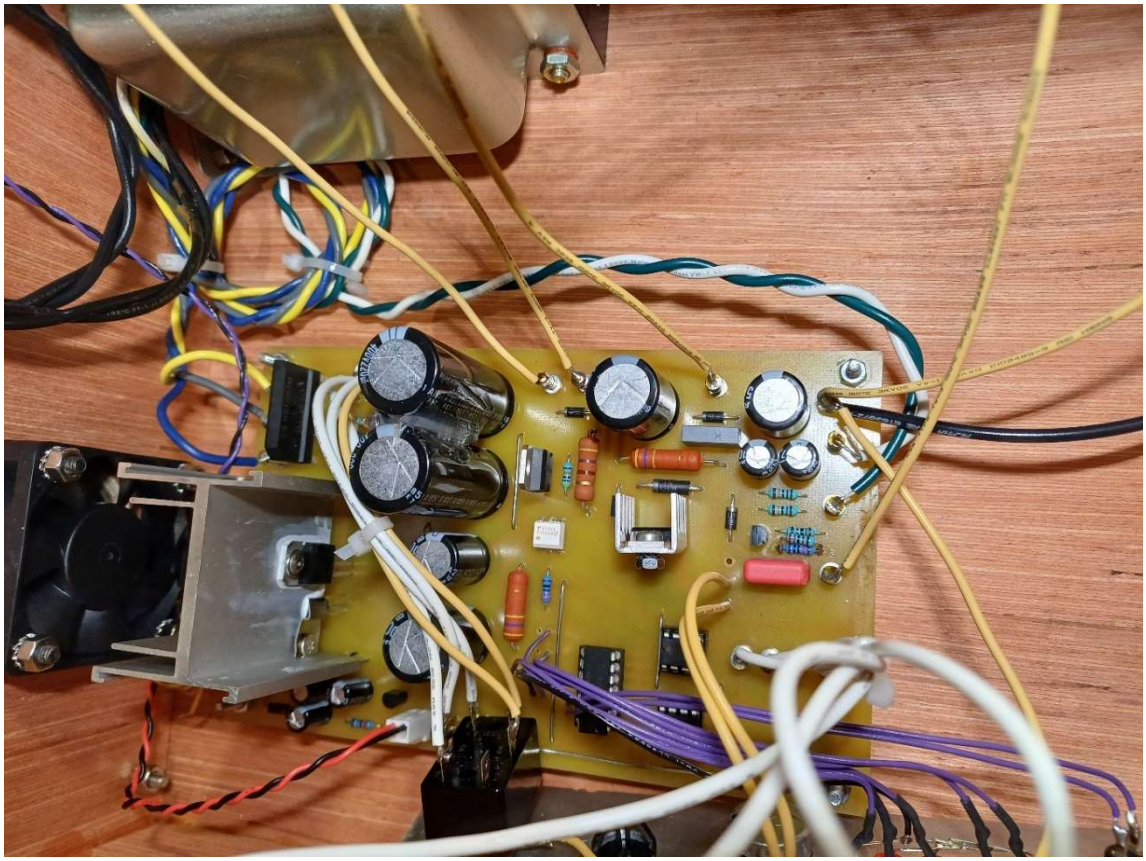


Figure 6: the wiring on the sockets. That took me a while to complete. Note: this picture was taken before the anode resistors were installed.

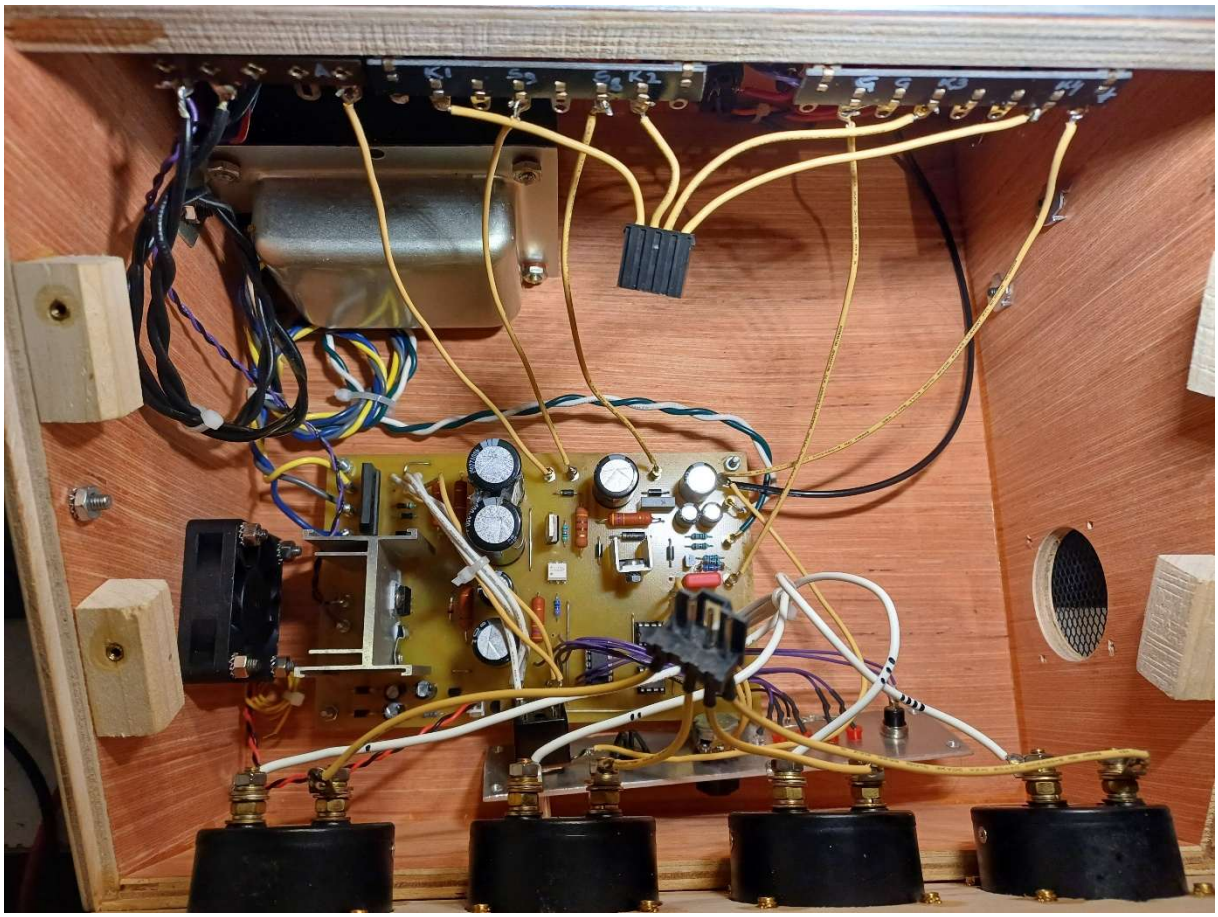
©Jygsawaudio 2026



Figure 7: all valve socket connections are made accessible from the top side, by use of tag strips. A heavy-duty Molex connector is used to connect the cathodes to the meters, to facilitate (dis)assembly.



©Jygsawaudio 2026



©Jygsawaudio 2026