A regulated 13.8V / 7A switch-mode power supply with over-voltage protection.

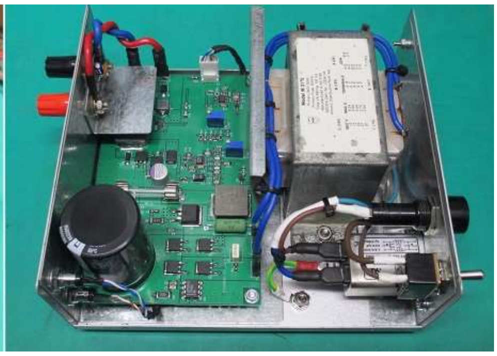

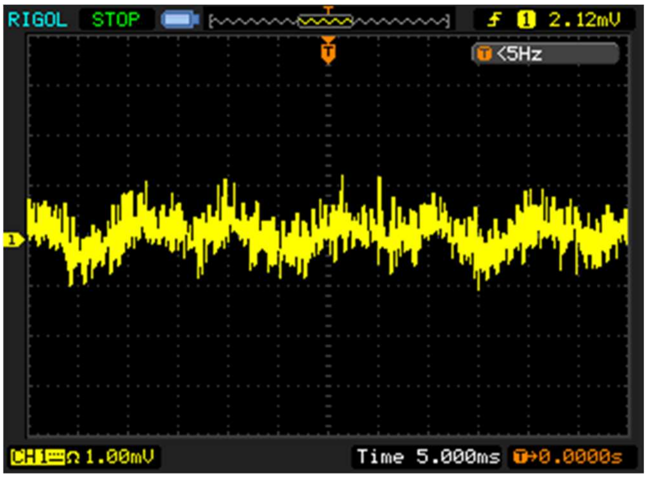

Display = AC ripple and noise on DC output voltage.
Load current = 6A.
Measurement bandwidth = 28Hz to 200MHz.
Probing = 50Ω back-terminated coax; actual vertical sensitivity = 2mV / div.

Glen Kleinschmidt

www.glensstuff.com

September 2023

Contents

Introduction

This project came about out of a sudden need for a number of 13.8V and 27.6V power sources for powering a host of newly acquired transceiver equipment.

A career change to a full-time role as a radio frequency person, with new colleagues having extra-curricular interests in HAM radio, has thrown me into a world that I had not blown much disposable income on, until now. Suffice to say, an introduction to amateur radio swap-meets and new friends looking for a sucker to become the custodian of their ~~old junk~~, err, fine vintage radio equipment, means that I suddenly have a number of new toys to play with.

On the low-power end of the spectrum, is a shelf in my shack newly dedicated to my assortment of radio scanners and mobile transceivers for both citizen and amateur bands, pressed into base-station usage. Except for my Kraco *SSB De Luxe*, which is the oddity on the shelf with its inbuilt mains (AC) power supply, these devices all require a 13.8V (DC) power source. The unit described herein was built for this purpose.

Being an electronics hobbyist with a hoard of parts to draw from, "home brewing" a solution was obviously a no-brainer. I rummaged through my junk and found a suitable power transformer. I also had a reasonably compact equipment enclosure, made of steel and aluminium, which I decided was just the perfect fit.

My initial thoughts were to quickly design and build a basic linear regulator. Inspiration, at the back of my mind, was the commercial units distributed by the Australian CB radio manufacturer GME. These things are ubiquitous down here and the same basic linear designs have been readily available for what must be at least two decades now. I have handled dozens of these things over the years.

These units are popular with HAMs because they do not generate radio frequency interference like many cheap switching supplies do. I have always half hated the darn things though, because you could almost fry an egg on the lid of one after running it at anywhere near its maximum rated output current for any significant amount of time - they do not have much transformer or pass-element heatsinking for their peak current ratings. For example, the "7-amp peak" model, at the time of writing, is the PSA 126, with fixed 13.8V regulated DC output rated to "6A max. @ 33% Duty Cycle". That is only about 2A and 27W on a continuous basis.

With the 100VA power transformer and the somewhat larger instrument case that I had at hand, I was set to better that specification significantly, just by being able to use much heftier heatsinking for the linear regulator. My 25W FM VHF rig running long-duration FSK will not cause a meltdown.

So, I sat down and started working on a schematic and although the path looked set, the intractable inefficiencies and attendant power output limitations of the linear approach started to gnaw away at me. To cut a long story short, I ditched the linear approach entirely in the end and decided to do something a little more ambitious. I elected to make an efficient synchronous power rectifier feeding a buck step-down switching regulator, with a design objective of generating unusually low levels of high-frequency switching ripple and noise.

The completed, wholly heatsink-free unit, now has pride of placed on the aforementioned radio shelf and I have yet to pick up any undue interference on any of my radio receivers – and that is including the scanner that is powered by this device and sits directly on top of it.

Specifications

The complete active bridge controller/rectifier, switching regulator and over-voltage protection circuitry is accommodated on a single double-sided PCB measuring approximately 140mm X 90mm. It requires a 15V (AC) input from a secondary winding of a mains power transformer.

Parameter	Value	Notes
Input voltage (AC)	15V	Recommended transformer 120VA or greater for continuous rated output.
Output voltage (DC)	13.8V	Trimable.
Continuous output current	7A	If paired with an appropriately rated power transformer.
Line regulation*	0.0075%	At a load current of 6A, 40V input delta.
Load regulation*	0.26%	6A load current delta.
Ripple and noise*	~6mVp-p	At a load current of 6A.

* Measured performance of prototype.

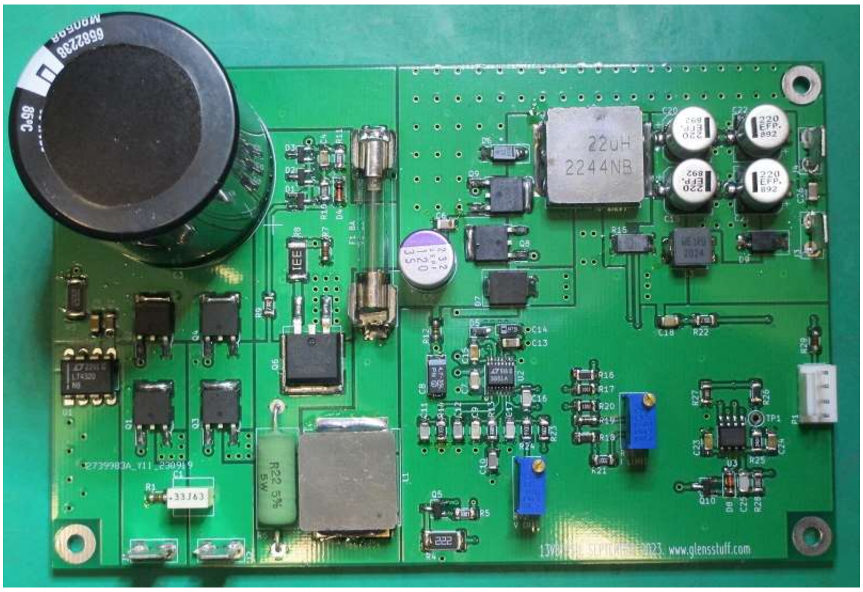

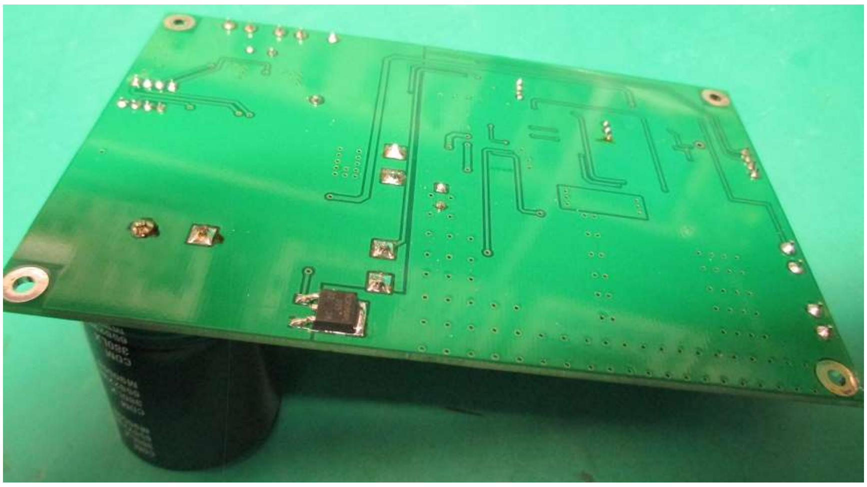

Choosing a buck controller

The choice here essentially boiled down to something readily stocked that is available in a package amenable to hand soldering. My preference was for a thermal-pad-free leaded package with a standard lead pitch of 0.635mm or greater. The shrink small-outline package (SSOP) variant of the LTC3851 fit the bill.

The LTC3851 is a +250kHz, fixed-frequency, current-mode synchronous buck controller with a V_{in} range to 38V. This IC has a number of desirable features, one of which is fold-back output current limiting.

In a basic circuit configuration, the part, however, is limited to regulator designs having an output voltage in the range of 0.8V to 5.5V. The lower limit is determined by the internal voltage-feedback reference potential of 800mV, while the upper limit is due to the current-sense comparator input pins (SENSE+ and SENSE-) having a specified *absolute maximum* rating of 6V.

Obviously, a work-around needs to be found here to apply the part to a regulator design supplying an output potential of 13.8V. The fix is a simple voltage divider at the current comparator sense pins (R16-R18 & R20). Note that the common-mode voltage now being divided down is large in comparison to the actual voltage being sensed across resistor R15, so the principal divider resistors need to be precision components (0.1%).

The addition of the common-mode voltage divider does, however, introduce a convenient means of calibrating the current-limiting threshold. R19, R21 and trimmer potentiometer RV1 provide a range of +/- pulling to the potential at the SENSE- pin such that the current-limiting threshold can be accurately set to 7A. There would, alternatively, be few practical options for calibrating the current-limiting threshold by tweaking the value of R15.

Over-voltage protection

A significant proportion of PCB area is devoted to circuitry that will, ideally, never be triggered into action after the power supply is verified to be fully functional. I do not like the thought of my vintage radio equipment, full of obsolete and "unobtanium" semiconductor devices, being fried by a faulty power supply.

An over-voltage condition at the output of the supply is indicative of a serious fault and the supply needs to be designed to disable itself to protect the load should this occur. A sure-fire way to make the supply disable itself is to apply an SCR crowbar to blow a fuse at the DC input to the regulator.

A triggering circuit based on integrated circuit U3 will fire SCR Q6, should the output voltage exceed 14.5V. That seals the fate of fuse F1, but F1 will not blow instantaneously. Furthermore, the potential across reservoir capacitor C3 will not drop and plateau below 13.8V instantaneously either. These limitations give the potential for an over-voltage condition to exist on the output of the supply for a duration of time that is long enough to cause damage to the load.

MOSFET switch Q7 serves to swiftly break current from the rectifier section to the regulator and the load upon the firing of Q6. Zener diode D9, connected across the output terminals, is included as a final measure to clamp the output potential during the microseconds of delay between U3 triggering and Q7 switching off.

Any potential for false triggering of the over-voltage protection circuit at power-on is prevented by Zener diode D8. This diode locks out the triggering circuit until the regulator output has ramped up to a potential sufficient to power U3.

Test point TP1 provides a means of triggering the over-voltage protection for testing purposes, as noted on the schematic. Although SCR Q6 has a brutal job in the clearing of fuse F1, the selected device is rather under-stressed and operated well within its ratings. I think the brief notes that I have included on the schematic for the associated stress-relieving components are self-explanatory.

Pictured below is some fuse torture porn and some of the units that expired during extensive testing of the effectiveness, durability and reliability of the over-voltage protection circuitry. I am happy to report that the fuses were the only casualties.

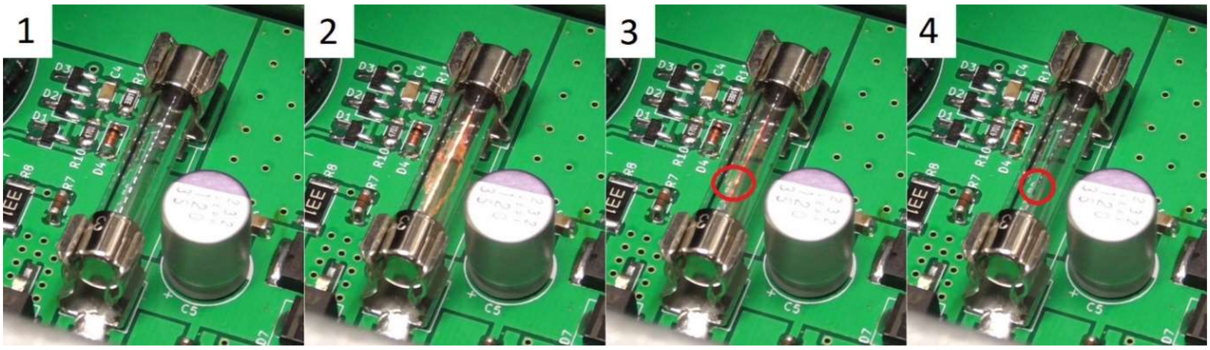

Construction

The enclosure that I used for this project measures 160mm (W), 184mm (D) and 70mm (H). The lid is nice and solid, being made of heavy-gauge steel, but the base is made of soft, thin-gauge aluminium that is not really up to the task of supporting anything as substantial as a large power transformer. In stock form, even pushing a tight banana plug into one of the binding posts mounted on the front panel would deliver sufficient force to bend the panel inwards.

I reinforced the front and rear panels internally with a pair of L-sections fashioned from 0.75mm-thick galvanised sheet steel. The base was also reinforced with the same material. This resulted in a really sturdy, solid enclosure. As a bonus, steel affords much better electromagnetic screening than aluminium too.

Everything fits snuggly inside and the form factor of this case is just right for my equipment shelf, but electrically there are a couple of shortcomings that I had to take mitigating measures against. An ideal enclosure, electrically, would have been a long and narrow one such that the circuit board as it currently sits, wedged between the front panel and the power transformer, could be rotated 90 degrees counterclockwise.

That would keep the magnetic flux of the power transformer at the rectifier end of the PCB, away from the buck controller, and the output binding posts could align directly with the spade terminals on the PCB and connect to them with minimal lead-lengths.

I had to shield the PCB from the magnetic flux of the power transformer, which was inducing current in the voltage feedback circuit of the LTC3851. Without the steel shield in place, under load I get distinct, low duty-cycle pulses of several mV peak amplitude on the regulated output that correspond to the bridge rectifier current pulses that top up the charge on reservoir capacitor C3 every 10mS (the mains supply frequency is 50Hz down here in this part of the world). I think trimmer potentiometer RV2, that serves as the output voltage adjuster in the voltage-feedback divider, acts as a good antenna when situated within the magnetic flux field of the power transformer.

Another problem with the chassis layout is that the output binding posts aren't ideally situated and the leads to them pick up switching spikes from the regulator. These high-frequency spikes are not seen at the PCB output terminals when properly probed using the ground-collar and tip method. I addressed this issue with the small metal shield and an additional pair of filter capacitors (a 10uF electrolytic in parallel with a 100nF ceramic), soldered to tag terminals mounted directly behind the binding posts.

Line regulation

This is a measure of the change in the output voltage of the power supply expressed as a percentage of a change in the input voltage to the power supply:

$$\text{Line regulation} = \frac{\Delta\text{Vo}}{\Delta\text{Vi}} \cdot 100\%$$

ΔVo is the change in output voltage and ΔVi is the change in input voltage. The picture below shows my bench setup for measuring line regulation.

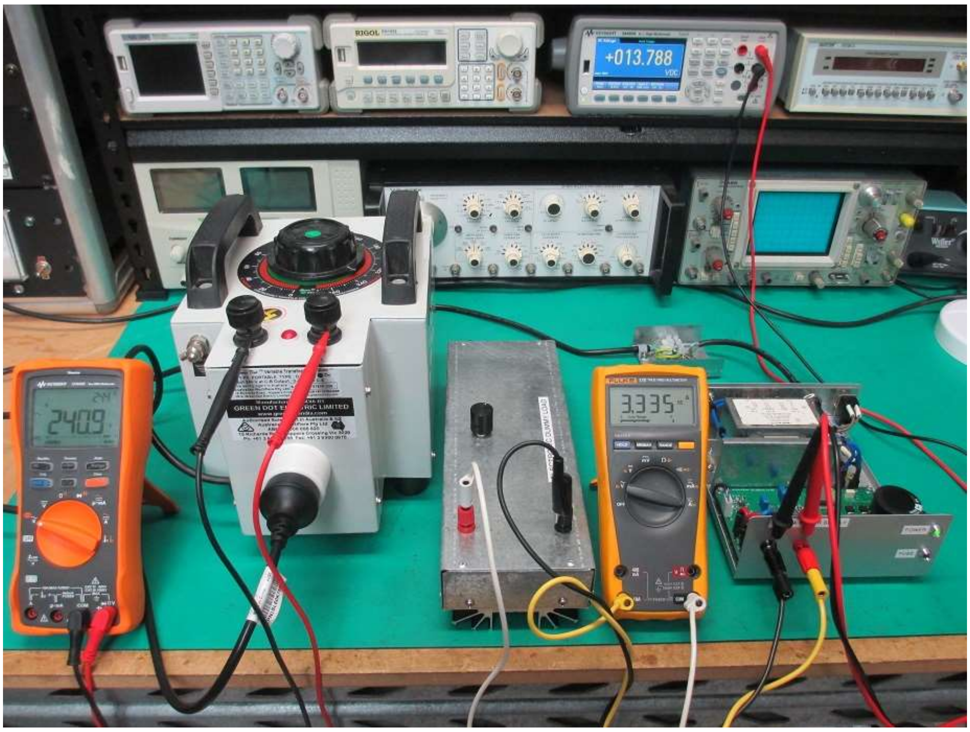

A variable autotransformer is used to vary the AC input voltage to the power supply. The rectangular metal box to the right of the autotransformer is an adjustable DC electronic load, being used to load the power supply. The Fluke 175 DMM to its right is monitoring the load current. The AC input voltage is being measured by the Keysight U1242C DMM at the far left of the photo, while my Keysight 34460A DMM on the shelf was used to monitor the DC output voltage.

The dummy load was adjusted for a load current of 6A. I varied the (nominally 240V) AC input voltage back-and-forth between 220V and 260V and subsequently got a consistent delta of 3mV at the regulated DC output.

Applying the formula shown above yields a line regulation value of 0.0075%.

I think that it is fair to say that the output voltage stability of this power supply is not unduly influenced by perturbations on the AC mains supply side.

Load regulation

This is a measure of a devices ability to deliver a constant voltage to the load over a specified change in load current. There are two common ways to express load regulation:

- The % rise in output voltage when a specified load is removed.
- The % fall in output voltage when a specified load is applied.

The first is commonly used to specify power transformers:

$$\text{Load regulation} = \frac{\text{Vul} - \text{Vl}}{\text{Vl}} \cdot 100\%$$

Vul is the unloaded output voltage and Vl is the loaded output voltage. For example, the 100VA, M2170 model power transformer used in this project has a secondary rating of 15V @ 6.6A. The secondary output voltage measures 15V at full load and 17.1V unloaded.

This yields a measured load regulation of 14% - exactly what the manufacturer specifies!

For a regulated power supply, however, it is more typical to specify the percentage fall in output voltage at the maximum specified load current:

$$\text{Load regulation} = \frac{\text{Vul} - \text{Vl}}{\text{Vul}} \cdot 100\%$$

The equation only differs in that the denominator has been changed from Vl to Vul.

To measure the load regulation, I reset the autotransformer to provide the nominal specified AC input voltage of 240V. I then adjusted the load current from 0A to 6.9A, whilst monitoring the output voltage. I only measured to 6.9A as the output current limiter is set to kick in at 7A. Here are the measured values, tabulated and charted:

I load	Vout
0	13.801
0.5	13.799
1	13.797
1.5	13.794
2	13.792
2.5	13.79
3	13.787
3.5	13.785
4	13.783
4.5	13.781
5	13.778
5.5	13.775
6	13.772
6.5	13.769
6.6	13.768
6.7	13.767
6.8	13.766
6.9	13.765

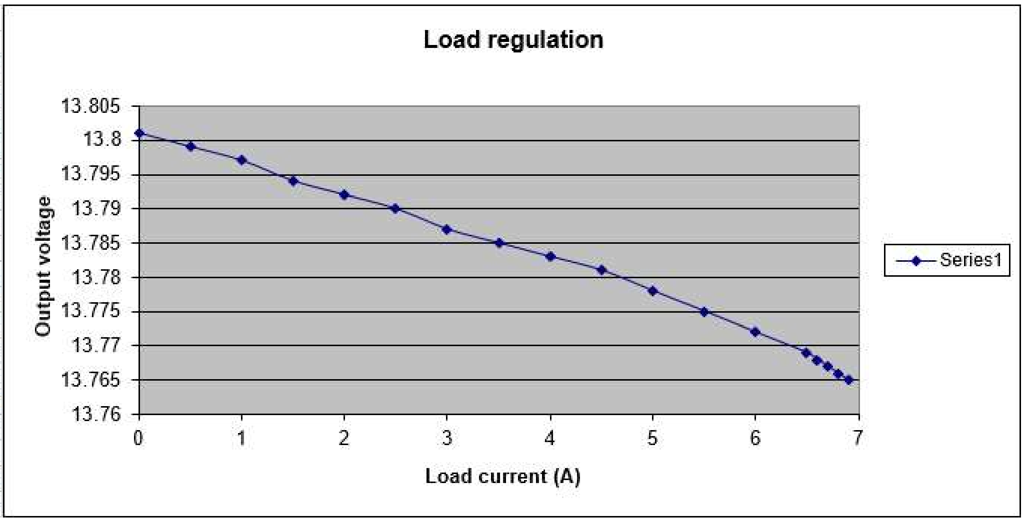

Over the full range of output current, from no load to a load current of 6.9A, the output voltage dropped by 36mV.

This yields a measured load regulation of 0.26%

The effective DC output resistance of the power supply can be found by applying Ohms Law and dividing the change in output voltage by the corresponding change in load current.

Between 0A and 6A, there is a change in output voltage of 13.801 – 13.772 = 29mV. Dividing by 6A gives an output resistance value of 4.8 milliohms. As a wild guess, the binding posts and connecting wires to them might contribute up to 10% of that.

Note that the transformer that I used for this project is actually a little undersized for the output current capacity of the regulator. I recommend that a transformer of at least 120 VA be used. That said, I am not particularly worried about keeping the current limit set to 7A in my unit.

As far I have managed to determine by measurement, my 100 VA transformer becomes maximally loaded at an output current of about 6.2A. You can never get the full VA rating, in watts, from the DC-side after rectification. I have contemplated winding the current limit threshold down to 6A, but there is no problem delivering the full 7A on a transient basis and even on a continuous basis things are not particularly worrisome.

I have run this thing for a couple of hours flat out at 7A and whilst the transformer did get quite hot, it never got to the point where anything bad would happen. Furthermore, under abnormal operating conditions, such as a short-circuited load, the fold-back output current limiting dramatically throttles back the power that the regulator actually draws from the power transformer.

Soft start

Capacitor C12, tied to the soft start pin of U2, sets the rate at which the regulated output voltage linearly ramps up to 13.8V upon the supply being turned on. The oscilloscope screen capture below shows this power-on characteristic. The output of the supply was loaded with a 3-ohm power resistor. The yellow trace is the unregulated potential across reservoir capacitor C3. The blue trace is the regulated output voltage.

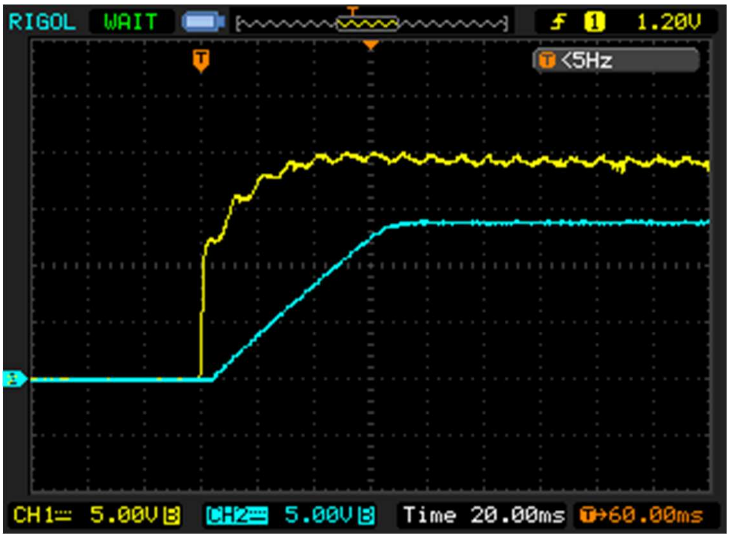

Ripple and noise

Of all the performance parameters that can be used to judge the merits of both the design and the implementation of a switch-mode power supply, this one is near the top of the checklist.

The ripple component is the rectified mains-frequency element that makes it through to the output terminals. “Noise” is a bit of a misnomer and is generally applied to refer to all of the high-frequency, switching-associated stuff that makes it to the output of the power supply. These high-frequency components are completely deterministic and do not exactly qualify as noise in the technical sense.

The “noise” typically consists of two main components:

- A ripple waveform at the switching frequency of the regulator.
- Brief spikes that occur during the power MOSFET switching transitions.

The latter can have high-frequency energy components approaching 100 MHz or beyond. This is one of the reasons why manufacturers of switch-mode power supplies often like to specify “ripple and noise” in a 20MHz bandwidth – it makes their designs measure better!

There is an opportunity here to segue into the power supply design details a bit more, before getting back to the measurements.

The LTC3851 is capable of operating at switching frequencies as high as 750kHz. What benefits a switching frequency that high might bring to a regulator of this power class is an interesting question, but I do not, personally, have the R&D time budget for that kind of esoterica right now. Furthermore, the thermal-pad-free SSOP variant of the LTC3851 that I went with, for the practical reasons outlined in *Choosing a buck controller*, does not have the thermal capacity to drive the gate capacitances of the chosen power MOSFETs that fast.

A single resistor from the FREQ/PLLFLTR pin of the LTC3851 to ground sets the operating frequency. I poked in a value of 110k, to give and operating frequency of approximately 324kHz. The current consumption of U2 can be monitored by measuring the voltage drop across resistor R12.

For a general-purpose, commercial switch-mode power supply, a ripple and noise performance specification that can be considered good, is where the peak-to-peak ripple and noise component measures not more than 1% of the DC output voltage. The measurement bandwidth should be at least 20MHz.

My design goal was to better that by at least an order of magnitude, which I achieved with a moderate margin. The elaborations upon the basic buck regulator circuit design that were required to achieve this were not particularly onerous.

The low switching noise at the output of my power supply it attributable to the 4-pole output filter design and low inductor ripple current. A buck regulator usually only has a 2-pole output filter. All else being equal, a buck regulator with less inductor ripple current will always generate less EMI and have less noise at the output, but this is a trade-off with transient response performance.
The addition of two extra poles of filtering significantly lightened the compromise, in this regard, that I would have otherwise had to make in order to achieve anything like the same noise performance. The transient response of this supply is still respectable and abundantly adequate for the intended application.

The two additional poles of filtering do have further implications for the transient response of the power supply that are not related to the ripple current of the main filter inductor. Because of this, the component values were chosen wisely, but more about this in the next section on transient load testing.

The sub-5-milliohm DC output resistance of this power supply was achieved in part by sensing the output voltage directly at the output terminal on the PCB. A conventional buck regulator, ordinarily, would not be stable with two additional poles of filtering introduced into the voltage-sensing feedback loop. I solved this problem with the feedback network consisting of C18, R22 and R23. DC feedback directly from the output terminal is provided via R22, but C18 effectively bypasses R22 as frequency increases, such that at high frequencies the two additional poles of filtering are removed from the voltage-sensing feedback loop entirely.

It would be nice if there was a practical way to star-earth the ground net of the switching regulator controller (like you can with a linear regulator controller) directly off the negative output terminal on the PCB. That would be great as far as the accuracy of the voltage-sensing feedback loop goes, but a big inductive loop would likely spell disaster for the gate drive to the power MOSFETs.

When it comes to PCB grounding around the controller IC, the MOSFET gate-drive loops have absolute priority over any ground feedback paths used for sensing the output voltage. A combination of brute-force ground fills and attention to where ground currents are actually flowing across the board is how I achieved the measured performance detailed herein, with a plain, double-sided PCB.

Now back to ripple and noise - the first thing to discuss is applicable measurement methods.

In the *Construction* section I mentioned how the instrument case/packaging constraints meant that I had to live with putting the output binding posts in a less than optimal position. This was problematic in regards to the leads to the binding posts picking up switching spikes and I had to take mitigating measures against this.

Manufacturers of switch-mode power supplies seem to be aware of this kind of problem, which is why there are various formal and non-formal standards and approaches to measuring ripple and noise that incorporate filtering components at some specified distance from the supply.

MEAN WELL currently has an instruction video on YouTube detailing such a setup, re-drawn here:

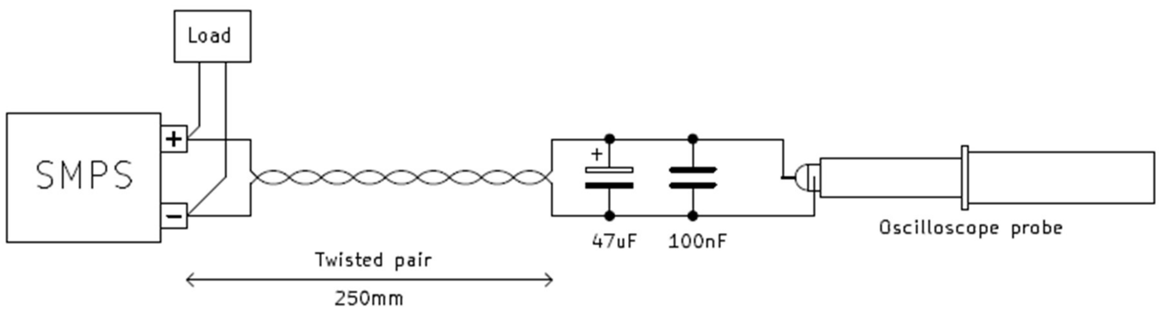

The oscilloscope probe connects to a DSO with the 20MHz bandwidth limiter switch on, as MEAN WELL specifies ripple and noise in a 20MHz bandwidth.

In my opinion, as a full-time radio communications technician and radio HAM, limiting the bandwidth to 20MHz is technically inadequate, but a good way to make your product measure better!

Not all manufacturers are this sneaky. Omron, for example, go to 100MHz. Here is a snippet from one of their switch-mode power supply data books that even quotes a formal standard:

Ripple Noise Voltage

The specified standard for the ripple voltage noise was measured with a measurement circuit that is based on JEITA standard RC-9131A.

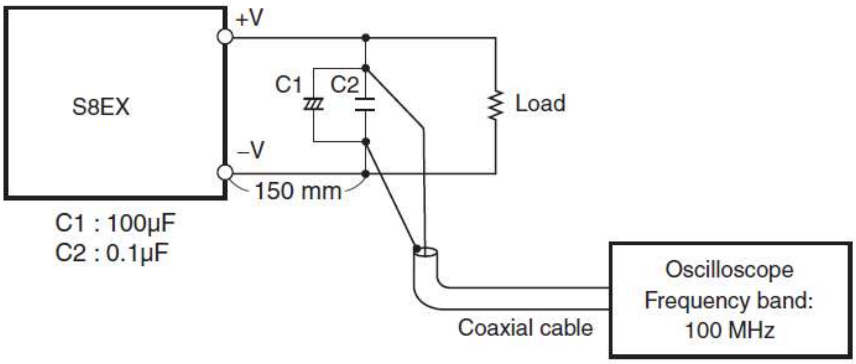

Some eyebrows might be raised about the addition of the filter capacitors. The 0.1uF directly at the probe input, an appreciable lead-length away from the power supply, does appreciably attenuate the conducted switching spikes, in my experience. It could be argued that this is cheating to some degree, but it could also be argued that any electronics item energised by the supply will also have some appreciable amount of capacitance bypassing its power supply input.

The 100nF (presumably ceramic) capacitor could form a high-Q resonant circuit with the inductance of some appreciable lead lengths (that might be threaded through ferrites inside the supply), so I presume that might be the reasoning for adding the electrolytic capacitor in parallel. A plain 47uF or 100uF aluminium electrolytic capacitor with some appreciable ESR would have a reliable dampening effect here.

Here is the setup that I used – a hybrid of the two methods detailed above:

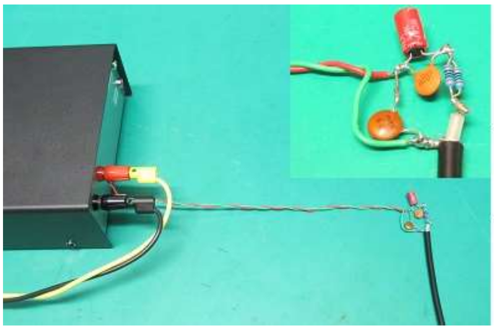

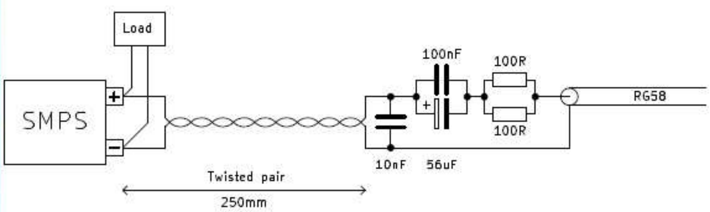

I found a single 10nF ceramic capacitor to be more than large enough to practically eliminate any residual conducted switching spikes.

The coax connects directly to my Rigol DS1202CA oscilloscope. The Rigol’s 50-ohm input termination is switched on, so the coax is properly back-terminated. The AC coupling capacitors were required because this DSO does not do AC-coupling when the 50-ohm input termination is turned on.

The ripple and noise measurements that follow were all taken in the full 200MHz bandwidth of my oscilloscope. Here is a baseband measurement taken with the power supply turned off.

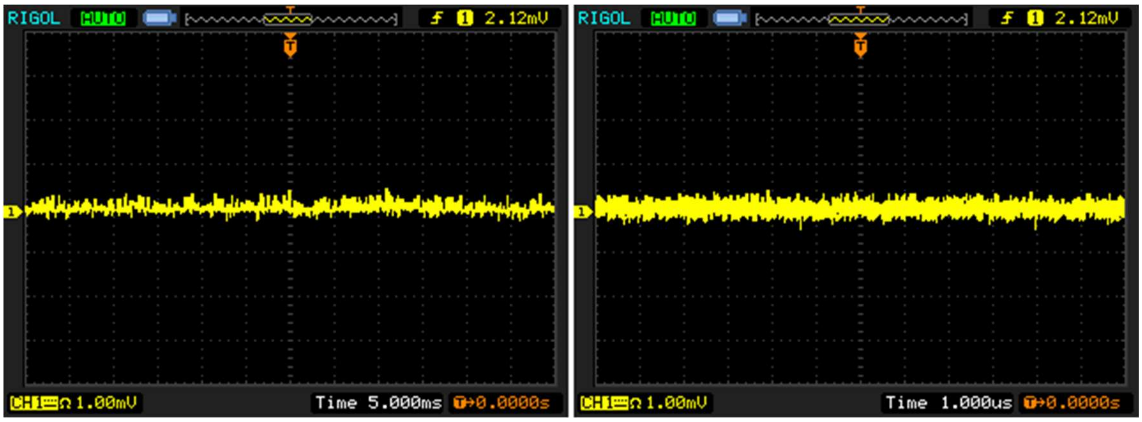

The difference between the two images is the setting of the horizontal timebase.

There is just under 1mV, peak-to-peak, of broadband noise here. This is the noise floor of the oscilloscope's vertical amplifier.

The following oscilloscope screen shots, with and without infinite persistence enabled, are with the power supply turned on and delivering 6 amperes of current into my electronic dummy load:

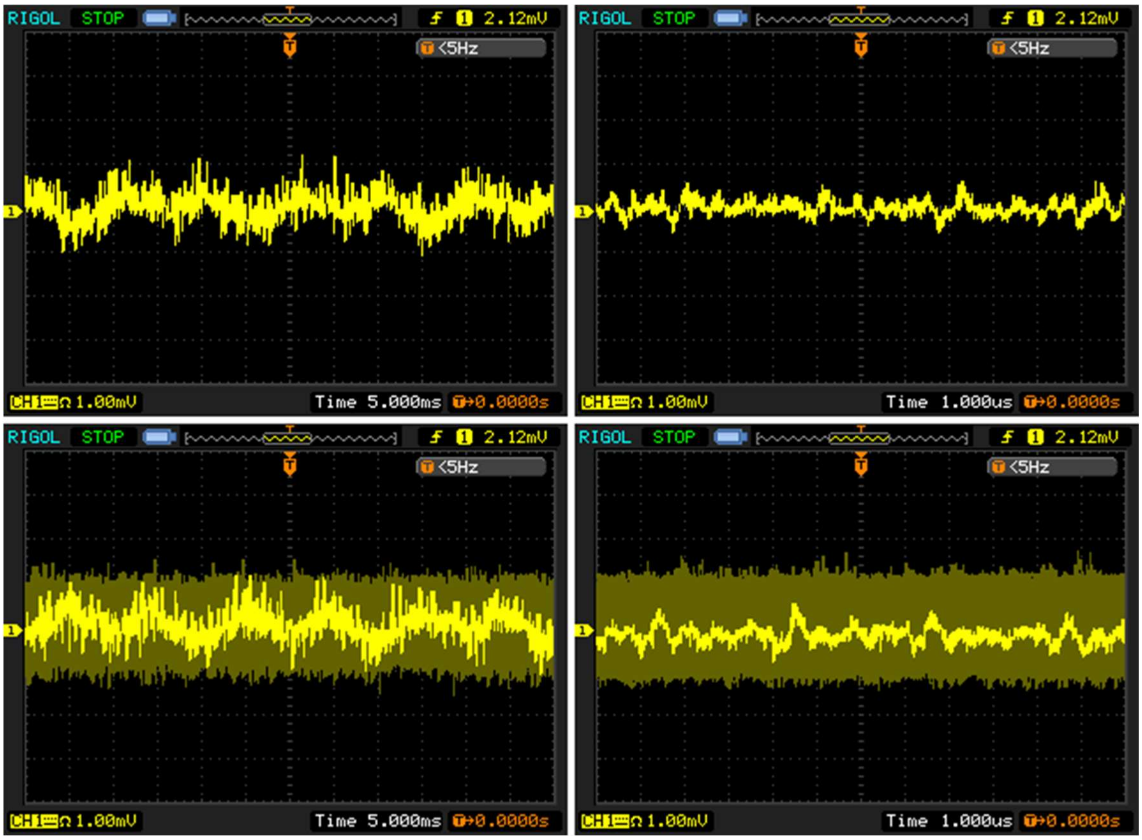

Note that the actual vertical sensitivity for this external signal, at 2mV per division, is double that indicated, due to the 6dB attenuation of the back-terminated coax. The omega symbol to the left of the sensitivity readout indicates that the 50-ohm input termination is enabled.

A vertical amplifier menu item permits you to select the division ratio of your probe, and the DSO will correct the vertical sensitivity readout accordingly. Somewhat annoyingly though, the firmware version, in my oscilloscope at least, does not have an X2 selection. It jumps straight to X5 from X1.

There are 3.2 divisions of peak-to-peak noise shown in the 1uS/div. display with infinite persistence turned on. That is 6.4mV of peak-to-peak noise. I think that I can round that down to 6mV, given that the displayed signal is only about 12dB above the oscilloscope's displayed noise floor.

At approximately 6mVp-p, the measured ripple and noise at the output of this power supply is 0.043% of the DC output voltage.

The rectified mains ripple component at 100Hz is quite obvious in the 5mS/div. oscilloscope screen captures. There must be some amount of 50Hz in there, but it is obviously well buried. The transformer flux shielding plate that I found necessary to include (discussed earlier, under *Construction*), is obviously quite effective.

It is worth noting that this mains ripple performance is considerably better than what a lot of commercial *linear* power supplies are capable of at the same load current.

I will close this section with a construction-related note that is relevant to the topic of ripple and noise. For EMC and shielding purposes, the regulator must be electrically grounded to the steel chassis. A single connection point is provided via the PCB mounting hole next to the negative output terminal, specifically for this purpose.

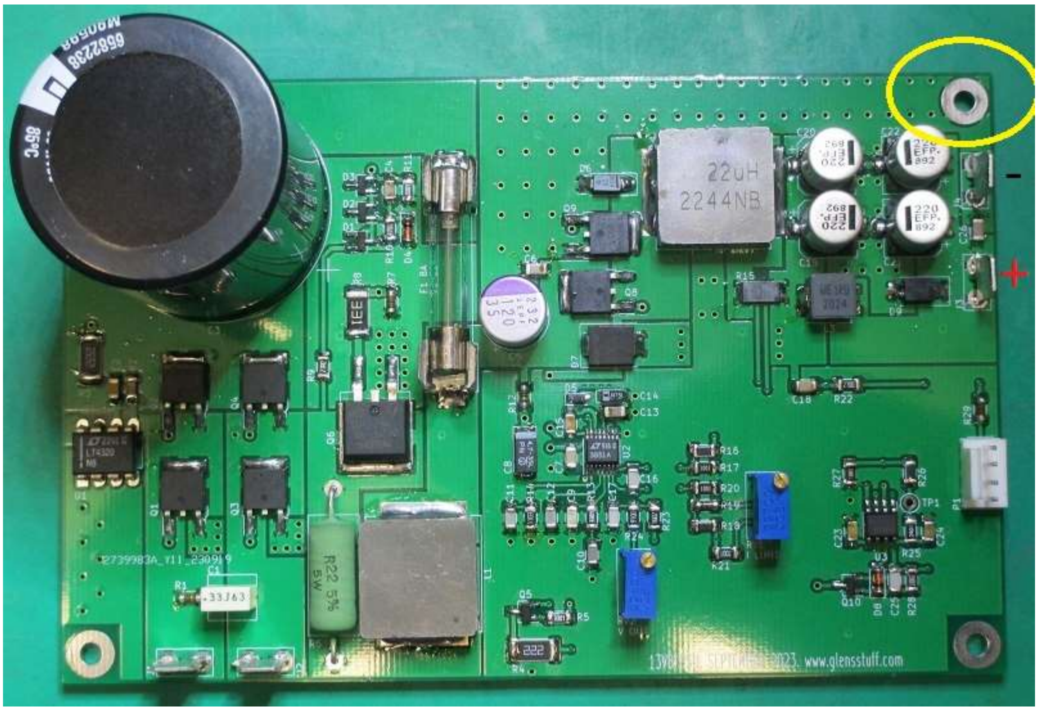

Transient response

As mentioned in *Ripple and noise*, the unusually low switching "noise" at the output of this power supply was achieved in part by adding an additional two poles to the conventional two-pole output filter.

For stability reasons, a feedback network was designed so that these two additional poles are bypassed out of the voltage-sensing negative feedback loop at high frequencies. Consequently, the impact that this additional filtering has on the transient response of the power supply is not actively compensated for by the control loop. This filter section is "freewheeling" and the component values need to be chosen wisely.

If there was no need to have a regard for transient response performance, one could just pick an arbitrarily large amount of inductance and a large-value capacitor with negligible equivalent series resistance (ESR). That would work great as a filter for high-frequency switching noise, but with a low-frequency, high-Q resonance, the output voltage of the power supply will over-shoot and ring like crazy in response to any transient shift in load current that does not have a very pedestrian di/dt rate of change.

The image below shows a frequency response simulation for a 2-pole LC low-pass filter. The capacitor has a value of 440uF and the inductor value is stepped from 1uH to 100uH.

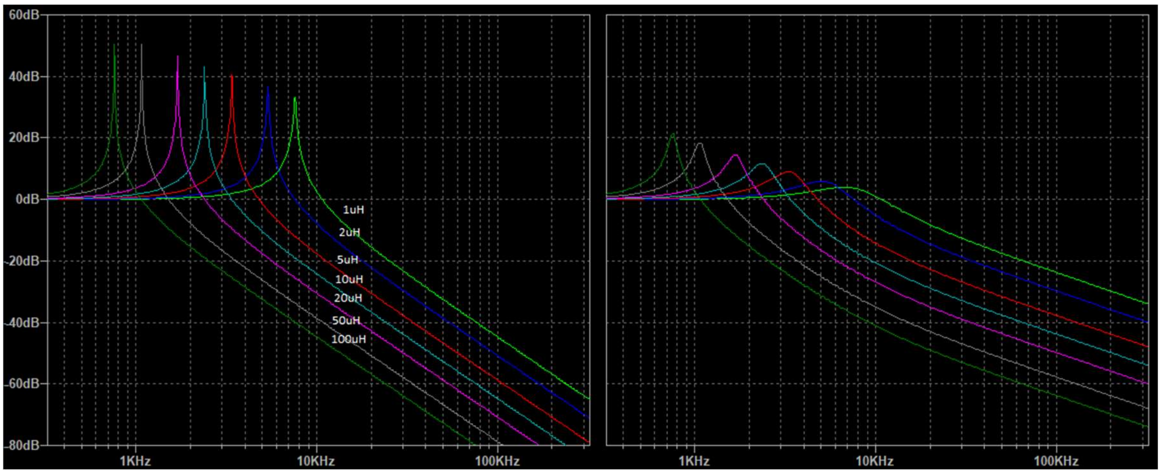

The frequency scale ends at 324kHz, which is the switching frequency of this power supply.
The graph on the left is using ideal components and the graph on the right shows what importantly happens when 40mΩ of ESR is added to the capacitor.

While these plots paint an idealised reality, the purpose is to show the fundamental relationships between filter Q, cut-off frequency and the ultimate attenuation achieved at the switching frequency of the power supply.

For a given value of capacitance, the two things that you can do to increase the amount of attenuation at the switching frequency (and beyond) will proportionally cost you in terms of worsening the transient response performance. These two things are increasing the value of the inductance and minimising the capacitors ESR.

Increasing the value of inductance will lower the cut-off frequency and increase the amount of peaking in the passband response by increasing the Q at resonance. That means more transient overshoot and ringing along with a longer settling/decay time.

The ESR of the output capacitor is a critical design parameter when it comes to taming the transient response of the filter. Provided that the cut-off frequency is not excessively low and accompanied by an excessively high Q at resonance, a modest value of ESR will generally be sufficient to adequately dampen and reduce the amount of peaking in the filter's passband response.

In theory, enough ESR could critically dampen the filter, but that amount of resistance might be too detrimental to the passband characteristics of the filter in other ways. For example, The ESR of the output capacitor is what predominately determines the mid-frequency output resistance of the power supply and therefore the amplitude of the initial output voltage shift induced by a sudden, large, high di/dt change in load current – a bit more on this a little later.

Note that while the DC resistance of the inductor must obviously have some influence in dampening the filter response, for any practical value of inductance it should be negligible in comparison to the influence of the capacitors ESR. You certainly should not wish to use the DC resistance of the inductor as a parameter to dampen the filter, due to, amongst other things, the I^2R power losses involved.

My improvised test bench setup for measuring the transient load response of the power supply is shown here:

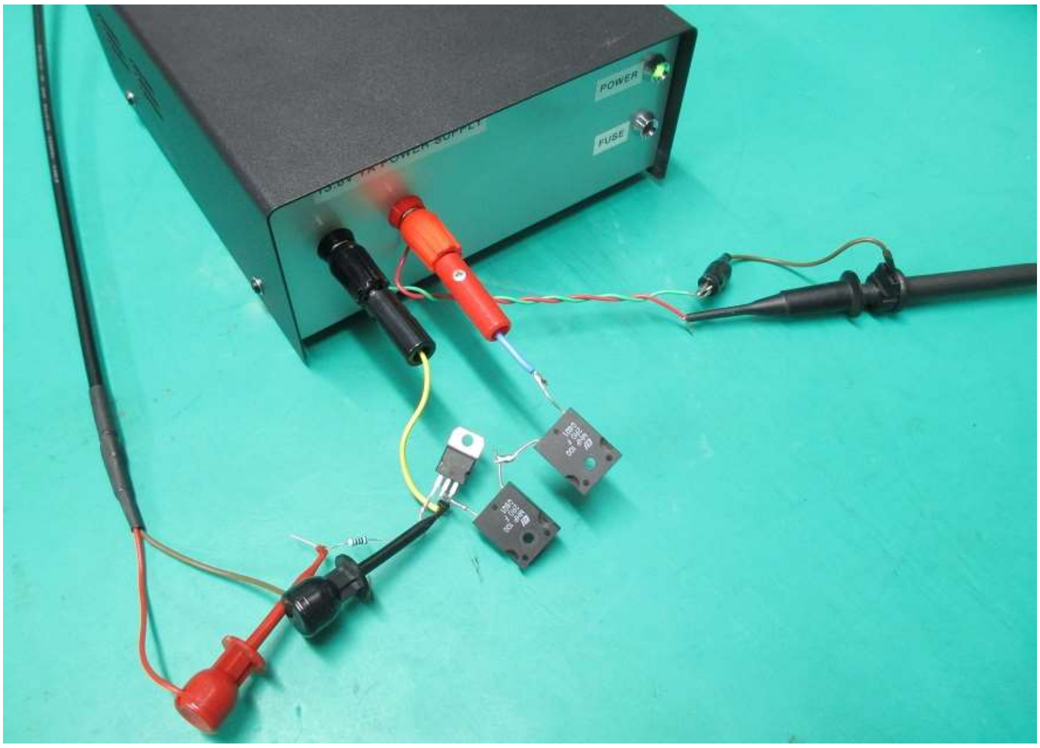

The setup is so simple I do not think that I need to draw a schematic diagram. A 4Ω load, consisting of a series-connected pair of non-inductive, high-power 2Ω resistors, is switched in and out of circuit with a generic TO-220-packaged power MOSFET, which I conveniently had lying about.

The gate-drive pulse signal for the MOSFET is provided by one of my test bench function generators. The resistor connecting to the red hook test clip is just a small-value gate stopper to kill any ringing. None of the power components needed to be mounted onto a heatsink due to the low duty-cycle on time (1ms in a 1s period) that made the average power dissipation negligible.

To avoid the voltage drops across the connecting wires carrying current to the power resistors, the output voltage of the supply is probed directly at the base of the binding posts courtesy of a short pair of twisted wires connecting the oscilloscope probe.

Here is the test result:

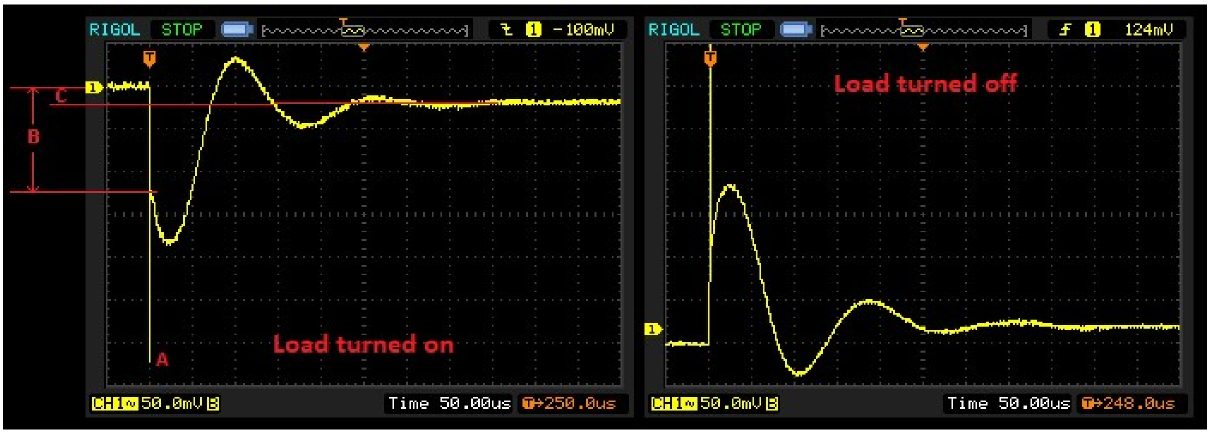

The transient response test exercises the supply with a large and abrupt change in load current that occurs much too fast for the control loop to respond to dynamically. A regulator circuit has a finite bandwidth and it will therefore take an amount of time for the regulator to correct for the initial change in output voltage that occurs due to the open-loop output impedance characteristics of the regulator.

The spike labelled (A) occurs at the instant that the load current is switched on and is due to the output inductance of the power supply. This inductance is mostly that of the wires between the binding posts and the regulator PCB.

Keep in mind that with a sufficiently fast turn-on di/dt, a large-amplitude voltage spike can be induced across even the tiniest amount of inductance, so the amplitude and presence of this brief spike is not much of a performance indicator for the power supply. Almost any practical length of wire lead, used to connect the power supply to an item of equipment requiring power, will contribute at least an order of magnitude more inductance to the effective output impedance of the power supply. If that lead is threaded through or has any clamped-on ferrites for noise suppression, then there might be another order of magnitude or more inductance added yet again.

The amplitude step labelled (B) is caused by the mid-frequency, open-loop output resistance of the regulator. This resistance is basically the ESR of the bulk-value output capacitors.

The amplitude step labelled (C) is the reduction in output voltage that occurs due to the DC output resistance of the power supply that will always be some value greater than zero. A measure of the DC output resistance was provided in the *Load regulation* section.

The 8-bit analogue-to-digital sampling converter and the chunky screen resolution of my DSO precludes any precision voltage measurements being made here, but we can see that it takes approximately 250uS for the nominally 13.8V output potential to settle to something like +/- 0.1% of the newly loaded value. I think that this is quite respectable performance, given the amount of filtering employed and the degree to which switching noise has been eliminated from the output.

Potpourri

Here are just some random things that I did not fit into the previous sections.

Capacitor C5

This is the bulk-value bypass capacitor at the power input of the switching convertor. It is a conductive polymer aluminium solid electrolytic capacitor; the only one of this type on the PCB. These capacitors are distinct from ordinary aluminium electrolytic capacitors in terms of their significantly lower ESR, which relates to their ability to handle much higher ripple current.

It is not an uncommon mistake for the inexperienced switch-mode designer to underestimate the magnitude of the RMS ripple current that a bulk capacitor in the position of C5 is actually subjected to.

I once worked for a company that made a speed controller for a three-phase BLDC motor. This motor could draw as much as 30A or so. The power input to this controller, connected via long inductive leads to a battery, was bypassed with a solitary 4700uF, plain-Jane electrolytic capacitor. The engineers were baffled when these controllers started coming back in large numbers under warranty with the connecting leads of this capacitor melted through.

This was a through-hole radial capacitor mounted a distance above the board axially. The "fix" that was decided upon was to start mounting this capacitor with the leads bent closer to the body and cut shorter for less resistance. I do not know if anyone actually twigged to the fact that the melted leads might be indicative of an amperage of ripple current a tad unhealthy for the poor capacitor, but I had already thrown my hands into the air before then anyway. I was just a lowly technician and any advice that I could have given would have only been brushed aside with contempt.

But getting back on topic now – under no circumstances should C5 be substituted with an ordinary electrolytic capacitor of the same capacitance value. Such a substitute will be stressed beyond its ratings and it will probably not have an impressively long lifespan if the supply is operated at or near its maximum output current.

MOSFET selection

The MOSFET gate drive signals generated by the LTC3851 are derived from its internal +5V regulator. This regulated potential is made available for the high-side bootstrapped-supply driver circuit via the $INTV_{cc}$ pin. With only 5V of gate drive available, it is quite important that a MOSFET with an adequately low gate threshold voltage is chosen. Note that the gate drive potential available is actually a little less for the high-side MOSFET, due to the losses associated with the bootstrap circuit - principally the forward voltage drop of the commutation diode, D5.

The TK3R1P04PL power MOSFET manufactured by Toshiba, that I specified, performs adequately in this regard. With a gate potential of 4.5V this part has a typical $R_{DS(ON)}$ of only 3mΩ.

Active bridge rectifier

I have not said much about this part of the design so far. This is not because the LT4320 "ideal" diode bridge controller is not an interesting or nifty part, but rather because applying it is not particularly involved. A conventional silicon diode bridge rectifier passing an average current in the region of 4A or so will get stinking hot if it is not attached to a reasonably substantial heatsink. Contrastingly, the surface-mounted power MOSFETs of the "ideal" rectifier as implemented here suffer a rise in package temperature above ambient that is barely perceptible to my fingers. Quite an improvement, I say!

PCB assembly

The highest wattage soldering iron that I own that is equipped with tips suitable for SMD work has a power rating of 80W. I have assembled two of these PCBs now. The first one was done before I purchased a PCB pre-heater. The 80W iron was barely adequate for the job and it was a struggle to get enough heat to make a decent joint into some of the pads which intentionally do not have thermal reliefs.

The second board soldered up an absolute doddle with the same iron, thanks to the assistance of the pre-heater. I purchased one of the better "quartz infrared" pre-heater units in what could be considered the budget category. For all intents and purposes, it appears to be a solidly constructed unit which does the basic job of pre-heating adequately.

I have used hot-air units made by Hakko in a production repair environment for power electronics PCBs, which were nice, but I could not bring myself to spend in the vicinity of $900 for something that looks less involved than the $30 toaster that I bought at Kmart.

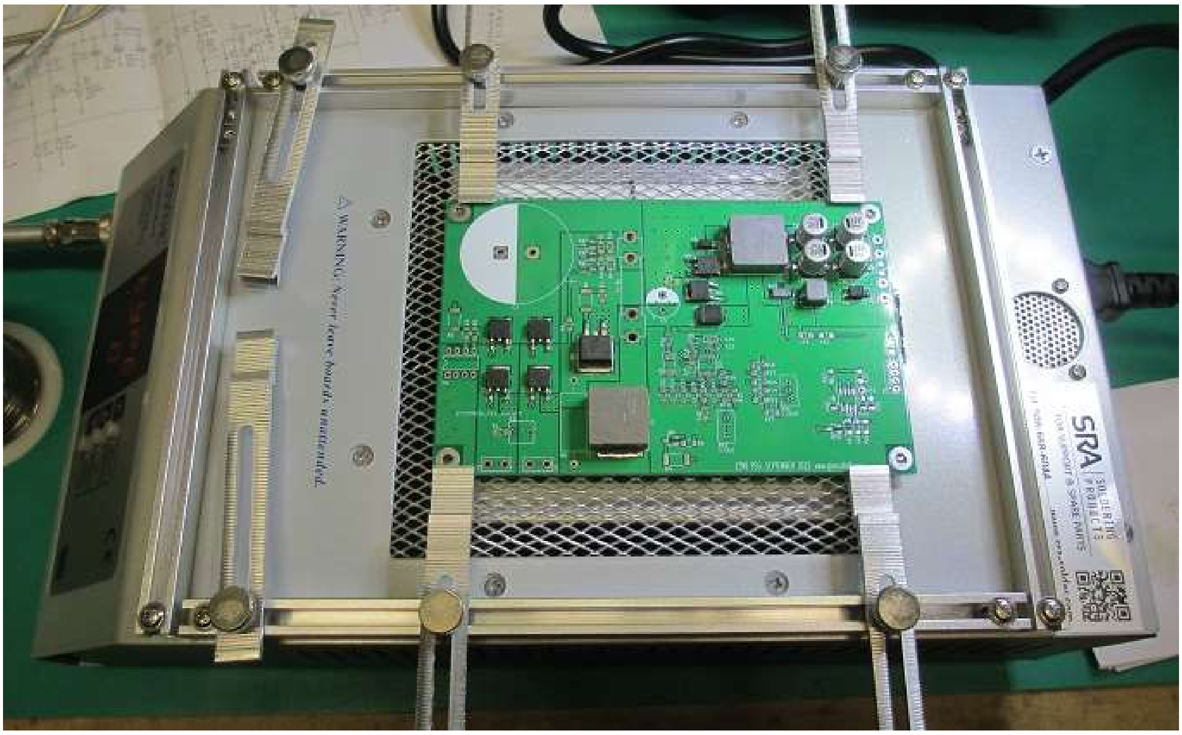

13.8V 7A POWER SUPPLY WITH OVER-VOLTAGE PROTECTION

SEPTEMBER 2023
www.glensstuff.com

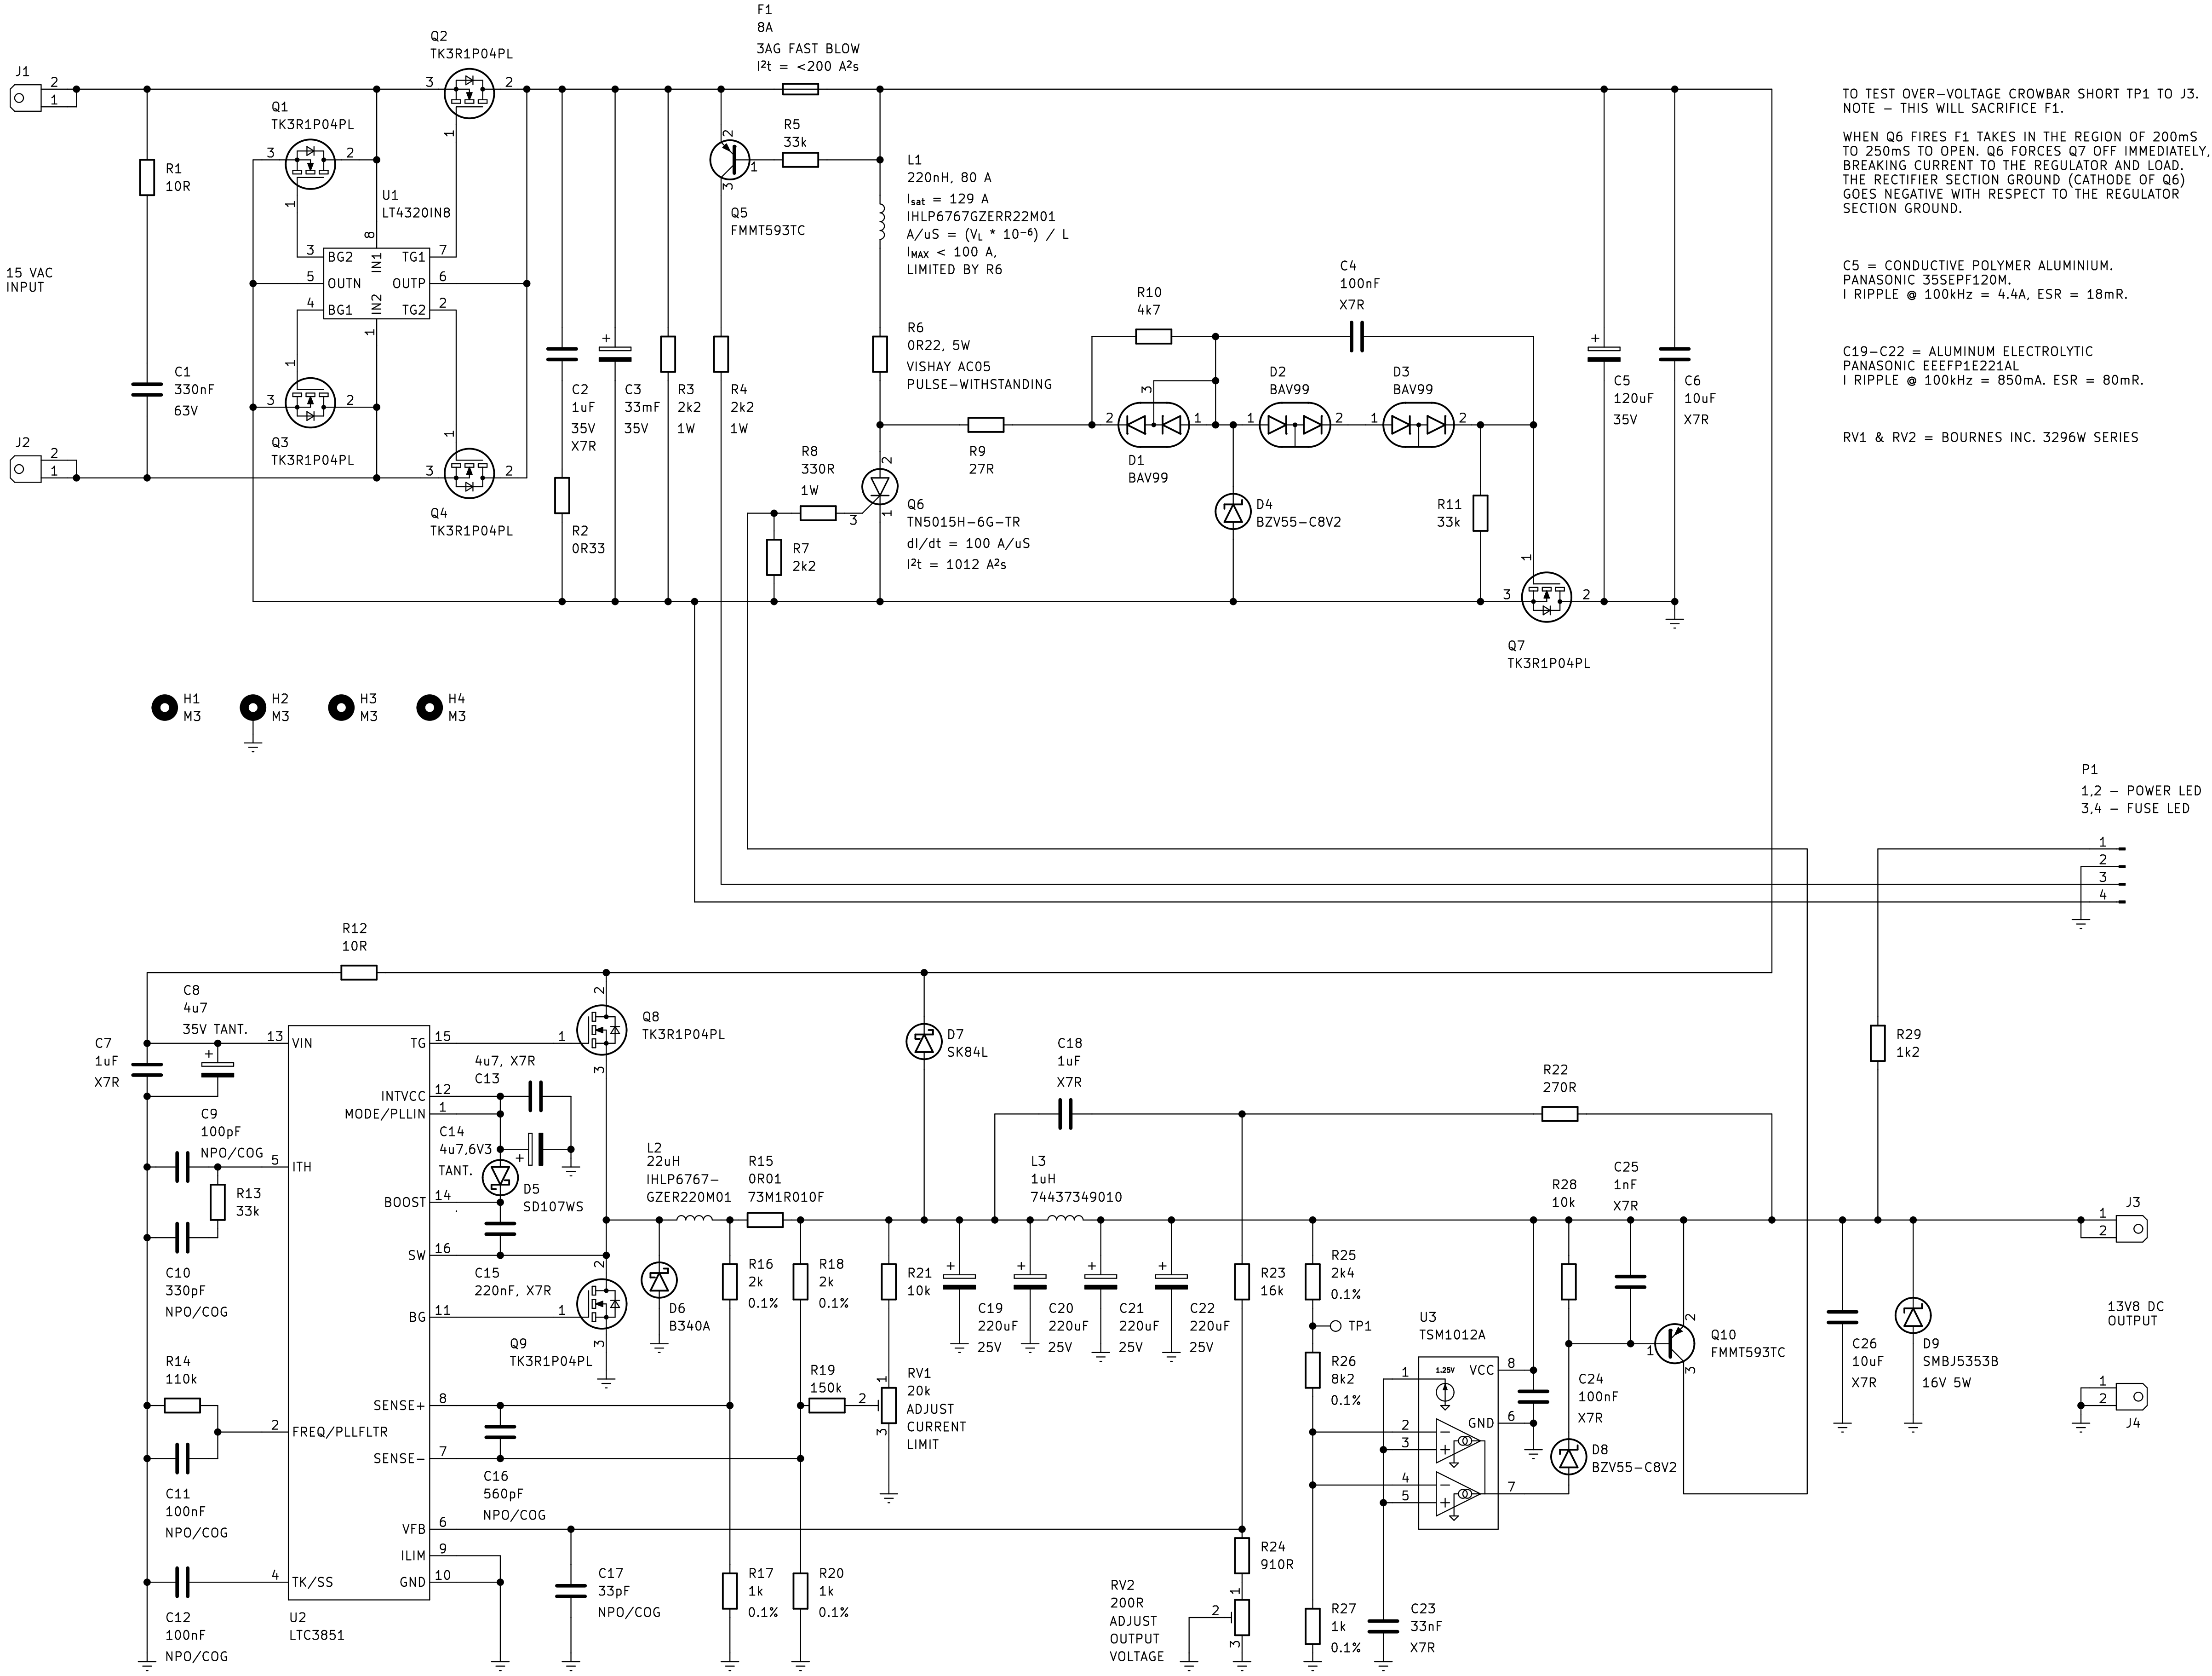

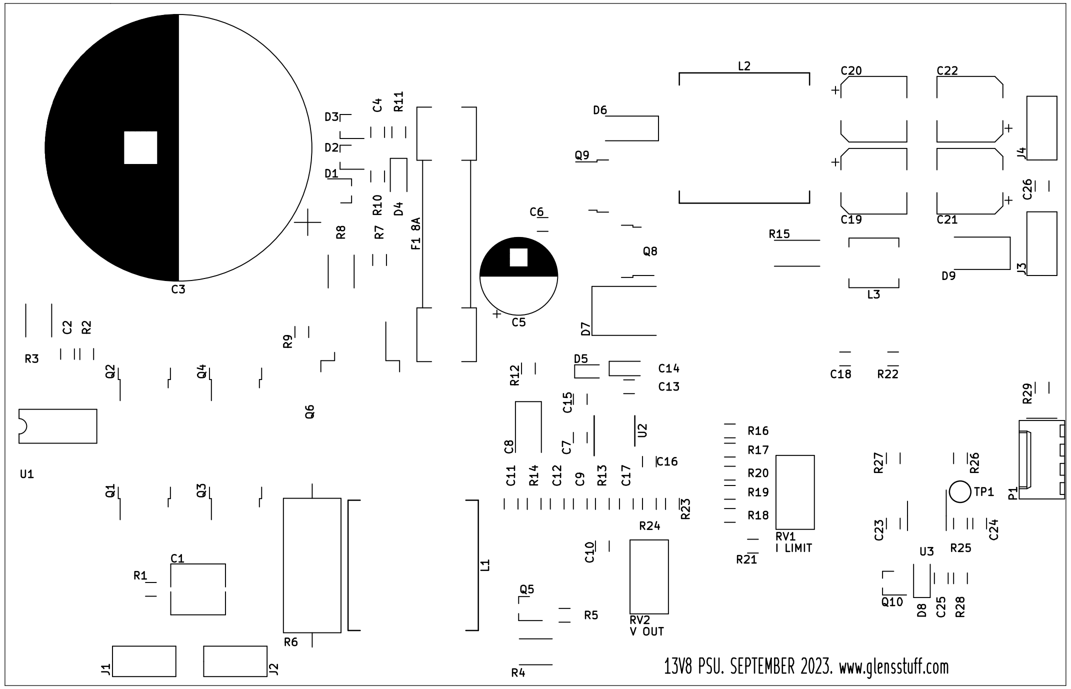

C3
D3
D2
D1
C4
R11
R10
D4
R8
R7
F1 8A
C6
C5
Q8
D6
Q9
D7
L2
C20
C22
C19
C21
J4
J3
C26
R15
L3
D9
C18
R22
R3
C2
R2
Q2
Q4
R9
Q6
U1
Q1
Q3
R12
D5
C14
C13
C15
C8
C7
U2
C16
C11
R14
C12
C9
R13
C17
R23
R16
R17
R20
R19
R18
RV1
I LIMIT
R21
R24
C10
RV2
V OUT
R29
P1
R27
R26
TP1
C23
C24
R25
U3
Q10
D8
C25
R28
C1
R1
R6
L1
Q5
R5
R4
J1
J2
13V8 PSU. SEPTEMBER 2023. www.glensstuff.com

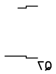

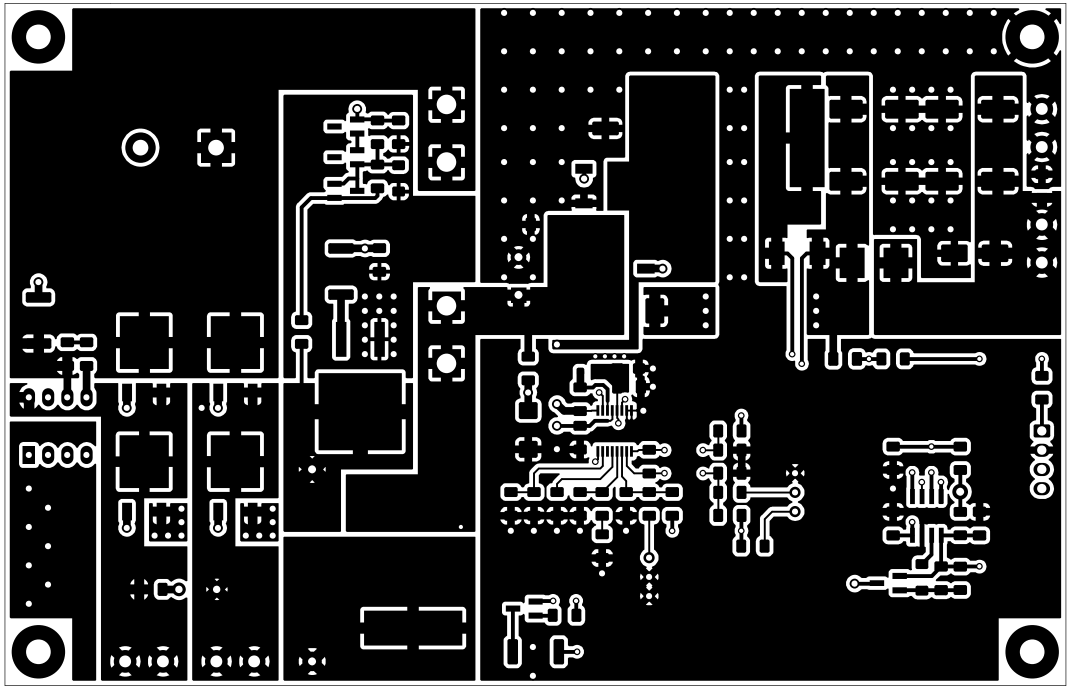

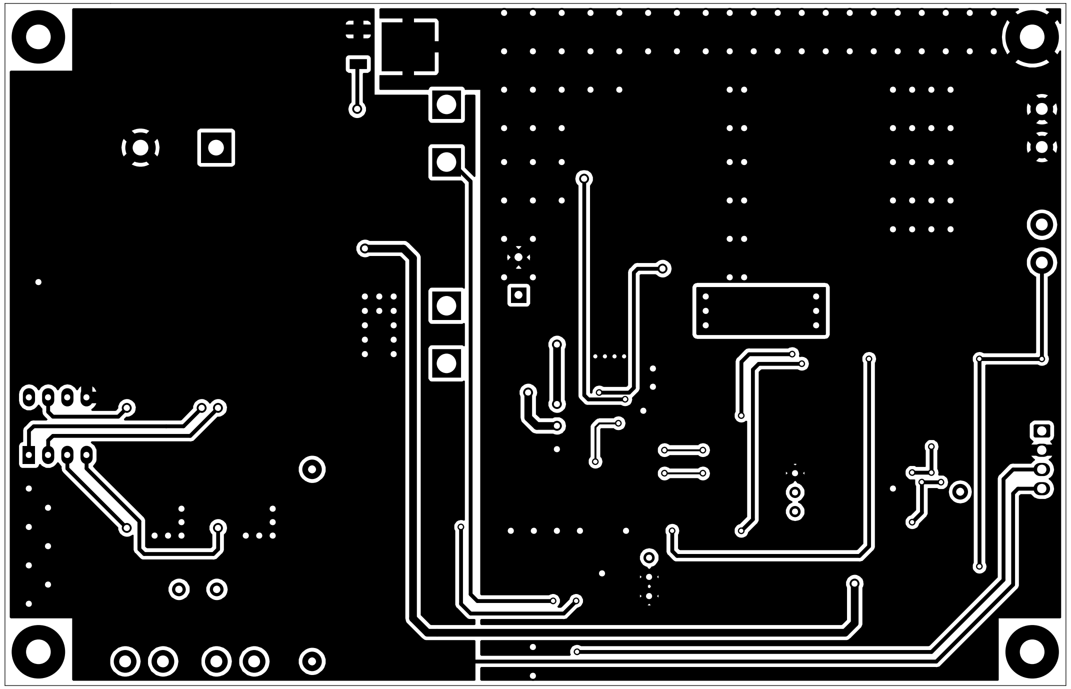

Parts for PCB SMPS13V8

Designator	Component type	Tolerance +/-%	Value	Part #	Manufacturer	Package	Description	Quantity
R1, R12	Resistor	1	10R			1206	Chip	2
R2	Resistor	1	0R33			1206	Chip	1
R3, R4	Resistor	5	2k2, 1W			2512	Chip	2
R5, R11, R13	Resistor	1	33k			1206	Chip	3
R6	Resistor	5	0R22, 5W	AC05000002207JAC00	Vishay	Through hole	Axial, pulse withstanding	1
R7	Resistor	1	2k2			1206	Chip	1
R8	Resistor	5	330R, 1W			2512	Chip	1
R9	Resistor	1	27R			1206	Chip	1
R10	Resistor	1	4k7			1206	Chip	1
R14	Resistor	1	110k			1206	Chip	1
R15	Resistor	1	0R01	73M1R010F	CTS Resistor Products	2512	Chip, current sense	1
R16, R18	Resistor	0.1	2k			1206	Chip	2
R17, R20, R27	Resistor	0.1	1k			1206	Chip	3
R19	Resistor	1	150k			1206	Chip	1
R21, R28	Resistor	1	10k			1206	Chip	2
R22	Resistor	1	270R			1206	Chip	1
R23	Resistor	1	16k			1206	Chip	1
R24	Resistor	1	910R			1206	Chip	1
R25	Resistor	0.1	2k4			1206	Chip	1
R26	Resistor	0.1	8k2			1206	Chip	1
R29	Resistor	1	1k2			1206	Chip	1
RV1	Trimpot	10	20k	3296W-1-203LF	Bourns Inc.	Through hole	Cermet, top adjust, 25 turn	1
RV2	Trimpot	10	200R	3296W-1-201RLF	Bourns Inc.	Through hole	Cermet, top adjust, 25 turn	1
C1	Capacitor	5	330nF, 40VAC, 63VDC	R82DC3330AA60J	Kemet	Through hole	Polyester, 5mm lead pitch	1
C2, C7, C18	Capacitor	10	1uF, 50V			1206	Chip ceramic, X7R	3
C3	Capacitor	20	33,000uF, 35V	380LX333M035A052	Cornell Dubilier	Through hole	35mm dia., 10mm lead pitch	1
C4, C24	Capacitor	10	100nF, 50V			1206	Chip ceramic, X7R	2
C5	Capacitor	20	120uF, 35V	35SEPF120M	Panasonic	Through hole	Aluminium polymer	1
C6, C26	Capacitor	10	10uF, 35V			1206	Chip ceramic, X7R	2
C8	Capacitor	10	4u7, 35V			2312	Moulded tantalum	1
C9	Capacitor	5	100pF, 50V			1206	Chip ceramic, COG/NPO	1
C10	Capacitor	5	330pF, 50V			1206	Chip ceramic, COG/NPO	1
C11, C12	Capacitor	5	100nF, 50V			1206	Chip ceramic, COG/NPO	2
C13	Capacitor	10	4u7, 25V			1206	Chip ceramic, X7R	1
C14	Capacitor	10	4u7, 6V3			1206	Moulded tantalum	1
C15	Capacitor	10	220nF, 50V			1206	Chip ceramic, X7R	1
C16	Capacitor	5	560pF, 50V			1206	Chip ceramic, COG/NPO	1
C17	Capacitor	5	33pF, 50V			1206	Chip ceramic, COG/NPO	1
C19, C20, C21, C22	Capacitor	20	220uF, 25V	EEE-FP1E221AL	Panasonic	Radial can, SMD	Aluminium electrolytic	4
C23	Capacitor	10	33nF, 50V			1206	Chip ceramic, X7R	1
C25	Capacitor	10	1nF, 50V			1206	Chip ceramic, X7R	1
U1	Integrated circuit			LT4320IN8	Analog Devices	DIP 8-pin	Bridge rectifier controller	1
U2	Integrated circuit			LTC3851AEGN	Analog Devices	SSOP 16-pin	Buck regulator controller	1
U3	Integrated circuit			TSM1012AIDT	STMicroelectronics	SOIC 8-pin	Voltage and current controller	1
Q1, Q2, Q3, Q4, Q7, Q8, Q9	MOSFET			TK3R1P04PL	Toshiba	DPAK	N-channel	7
Q5, Q10	BJT			FMMT593TC	Diodes Incorporated	SOT-23	PNP	2
Q6	SCR			TN5015H-6G-TR	STMicroelectronics	D2PAK	50A standard recovery	1
D1, D2, D3	Dual diode			BAV99		SOT-23	Silicon switching diode	3
D4, D8	Zener diode			BZV55-C8V2,115	Nexperia	MiniMelf	8V2, 500mW	2
D5	Schottky diode			SD107WS		SOD-323	30V, 100mA	1
D6	Schottky diode			B340A		DO-214AC (SMA)	40V, 3A	1
D7	Schottky diode			SK84L		DO-214AB (SMC)	40V, 8A	1
D9	Zener diode			SMBJ5353B		DO-214AA (SMB)	16V, 5W	1
L1	Inductor	20	220nH, 80A	IHLP6767GZERR22M01	Vishay Dale	Nonstandard	Shielded	1
L2	Inductor	20	22uH, 11A	IHLP6767GZER220M01	Vishay Dale	Nonstandard	Shielded	1
L3	Inductor	20	1uH, 10A	74437349010	Wurth Elektronik	Nonstandard	Shielded	1
F1	Fuse		8A, fast blow		Littlefuse		3AG, 1/4" x 1-1/4"	1
F1	Fuse clip			01020076Z	Littlefuse	Through hole	PCB-mount, straight leads	2
J1, J2, J3, J4	Spade terminal			1217861-1	TE Connectivity	Through hole	PCB-mount, 0.25"	4
P1	Connector			22272041	Molex	Through hole	KK-254 vertical header, 4-way	1