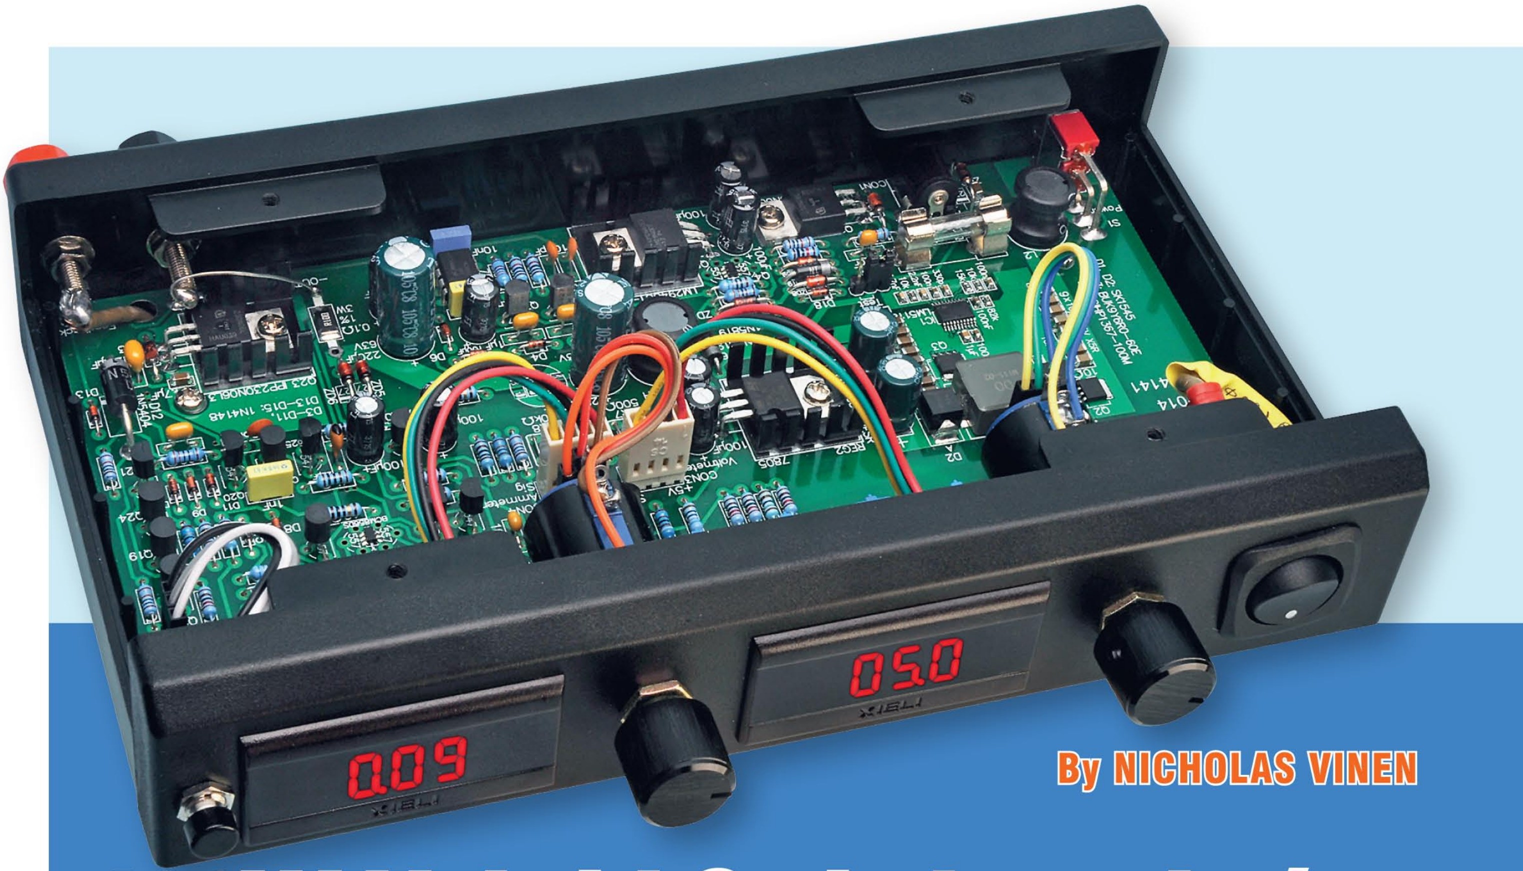

By NICHOLAS VINEN

100W Hybrid Switchmode / Linear Bench Supply – Part 2

Last month, we introduced our new 40V switchmode bench supply which can deliver up to 5A and fits into a compact rack-mount case with dual metering. This month, we describe the operation of the linear regulator circuit, discuss the PCB layout design and start the assembly by installing the parts on the PCB.

SINCE WE didn't finish describing the circuit last month, here's a quick refresher. The low-dropout (LDO) linear regulator determines the final output voltage and current, with the preceding switchmode section 'tracking' its output so that the input to the LDO is slightly higher than its output (by about 0.75V).

This gives the linear regulator sufficient 'headroom' to operate while keeping dissipation low – at 5A, the dissipation is around 5A × 0.75V = 3.75W, which is manageable with a small heatsink. This arrangement also means that if the current limit is set to a high level and the output is shorted, the switchmode regulator's output quickly drops to a low level and so the overall dissipation is kept reasonable.

A couple of different approaches can be taken when designing a low-dropout linear regulator. One is to use a PNP transistor or P-channel MOSFET as the pass element, and thus its base or gate is pulled towards ground to turn it on. The minimum drop-out voltage is simply the pass transistor's saturation voltage at full load current.

However, the fact that the base/gate voltage has to be pulled lower to increase the output current and higher to reduce it complicates the feedback system, since this type of arrangement can't have any 'local feedback'. This is why monolithic LDO regulators often have quite specific requirements for the output capacitor value and ESR to ensure stable operation; they rely solely on global feedback and the phase shift from the output capacitor forms a critical parameter for correct operation.

The other option is to use an NPN transistor or N-channel MOSFET with

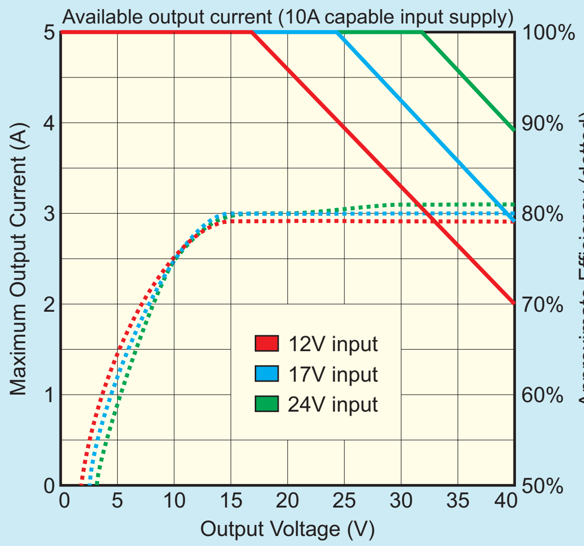

Fig.4: the circuit is capable of delivering 5A but this is limited at higher voltages by the power delivery capabilities of the DC supply and the 10A input limit. This graph shows how much current is available over the full output range for three common supply voltages. Note that the efficiency is best at lower output voltages. Note also that while the unit is capable of the indicated current, the switchmode section will get warm if operated at these limits for extended periods.

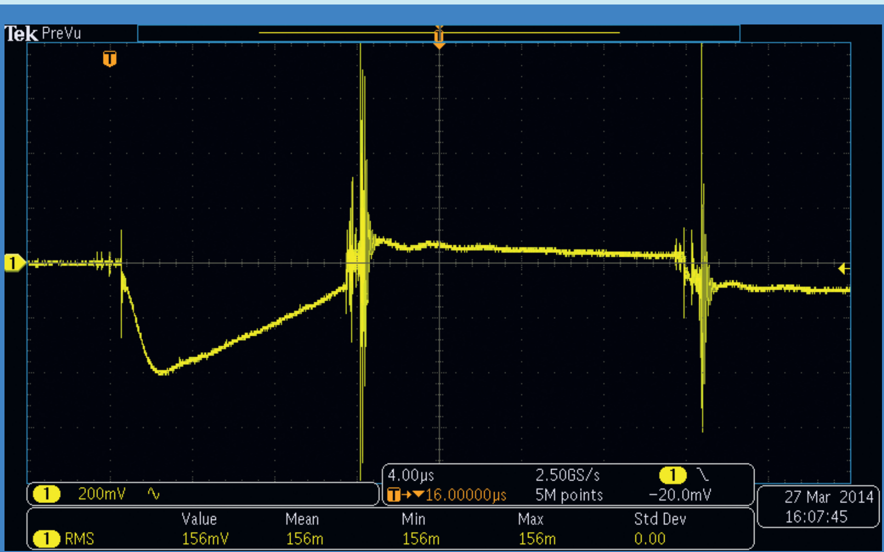

Fig.5: the output response with a 12V input, a 15V output and with a 1A load being rapidly connected and disconnected with no external output capacitor. The vertical scale is 200mV/div and the timebase is 4µs/div. As you can see, when the load current suddenly increases, the output drops but quickly recovers. There is a small amount of overshoot when the load is removed but it is well-controlled. The undershoot when the load is re-applied soon after is smaller than the first time as the switchmode section has not yet returned to idle operation.

its base/gate voltage driven from a 'boosted' supply rail somewhat above the main supply rail. In this circuit, we're using a MOSFET with a boosted supply that's around 10V above the output voltage. This allows us to vary the MOSFET's on-resistance from a very high value of many megohms when the output voltage is low and the load is light, to a very low resistance of around 15mΩ when it's delivering full-load current.

This arrangement gives superior regulation and filtering since it will inherently self-regulate to a certain extent. If the MOSFET's gate voltage is held constant, its source voltage (ie, the output) will be a certain amount lower than this and it will only vary over a small range (~1V), regardless of the drain voltage (ie, upstream supply).

Consider what happens if the output voltage (source terminal) drops and the gate voltage is constant. In this case, the MOSFET's gate-source potential increases and that turns the MOSFET on harder so that it conducts more current and thus pulls the output voltage up. Conversely, if the output voltage increases, the gate-source voltage drops and so the MOSFET conducts less current, thus reducing the output voltage.

While this mechanism is 'local' and therefore very fast, it isn't very accurate, as the gate-source voltage varies somewhat with temperature and channel current. So there still needs to be a global negative feedback mechanism to give an accurately regulated DC output voltage. However, this feedback system is less critical to performance thanks to that inherent local feedback.

Also, because there is less phase shift in this arrangement, the feedback loop doesn't have to be as heavily compensated and this allows the global feedback to act more rapidly, responding more quickly to sudden changes in load impedance.

Design details

There are a few regulator ICs which operate in this manner, but all the ones we could find have a fixed current-limit threshold, set by a low-value resistor in the main current path. That makes it awkward to implement a wide-range adjustable current limit.

As a result, we built our own regulator circuit. This is obviously more complex than using an IC, but the parts are cheap and commonly available, whereas ultra-LDO regulator controller ICs are somewhat expensive and hard to get.

Fig.6 shows the circuit of this linear regulator section. The labels at the edges match up to the labels on Fig.3, published in Part 1 last month, to show the connections between the two circuit sections. Taken together, these form the complete circuit of the bench supply.

The incoming supply rail (VIN) comes from the output of the switchmode regulator and its ripple filter, described last month. MOSFET Q23 controls current flow from this supply to the output (VOUT+), as described above, with its gate voltage typically 2V above the output voltage.

The regulator circuit is somewhat similar to that of an audio amplifier due to the need for accurate and fast-acting negative feedback. If you compare the two, you will find that there are broad similarities but subtle differences. The most obvious difference is that there are two differential input pairs, one to control the output voltage and the other to enforce the current limit. These are based around PNP transistor pairs Q14/Q15 and Q9/Q10 respectively.

Features and Specifications

Size and weight: 209 × 43 × 162mm, 400g

Input supply: 12-24V at up to 10A

Input under-voltage lockout: 11.3V

Output range: 0-40V at up to 5A (see Fig.4)

Output power: 100W+, depending on input supply voltage and current

Output ripple and noise: typically <5mV RMS @ 1A, ~1mV RMS at light load

Output capacitance: 2.2µF internal, handles any external capacitive load

Load regulation: <10mV for a 1A load step (measured at PCB terminals)

Line regulation: <5mV, 12-24V

Transient response: <500mV undershoot/overshoot for a 1A load step, recovery in ~10µs (no external capacitor) (see Fig.5)

Current limit response time: <150µs (short circuit @ 40V); <2ms to resume voltage regulation (depending on current limit)

Efficiency: ~70-80% (see Fig.4)

Voltmeter: resolution 0.1V, accuracy ±0.1V

Ammeter: resolution 10mA, accuracy ±10mA

Protection: fuse, cycle-by-cycle input current limiting, output current and voltage limiting

Current limit: continuously adjustable 0-5A, typically stable within ±1mA

Other features: view current limit, load switch

For accurate DC control, we need the base-emitter voltages of these transistor pairs to be fairly accurately matched (or rather, for the difference between them to remain constant) but this depends on temperature. So differential heating or cooling of these transistors due to air currents and so on, even of a fraction of a degree, can affect operation. As such, these transistor pairs are thermally bonded so that they remain at the same temperature.

This can be achieved in one of two ways: either by using two transistors in a single package or by bonding two separate transistors with thermally conductive paste. The PCB is designed for either approach, but thermal tracking is better when the two transistors are in a single package, so we have used the BCM856DS dual-matched transistors for our prototypes. We suggest you do the same.

These are rather neat devices, being equivalent to two BC556s in a 6-pin surface-mount package. The current gain (h_{FE}/beta) and base-emitter voltages are matched to within 10% and 2mV respectively. They are quite affordable and available in three different packages; we are using the largest one since it is easier to solder.

To understand the operation of the regulator as a whole, start by considering voltage-monitoring transistor pair Q14 and Q15. The two emitters are fed with a constant current of about 1mA by PNP transistor Q13, with this current set by the 680Ω resistor at its emitter. A bias voltage fed to Q13's base via a 2.2kΩ resistor from Q18 acts to keep around 0.6V across the 680Ω resistor.

This differential pair has 47Ω 'emitter degeneration' resistors to reduce the overall gain somewhat and improve linearity, which aids stability. The collector currents are kept more or less equal by a current mirror consisting of NPN transistors Q16 and Q17. This keeps the circuit operating consistently despite large variations in supply voltage as the output voltage varies.

Q14's base is tied to the negative output terminal (ie, effectively ground) via a 22Ω resistor, while Q15's base goes to the output feedback divider (shown in Fig.3 last month) via 'OV-feedback'. This is a divided-down version of the output voltage, as set by VR1, the voltage adjustment pot.

When the regulator's output voltage increases, the voltage at Q15's base also increases, reducing the current through Q15. Since the collector currents are mirrored, this means that more of the 1mA total emitter current must flow from Q14 to D7 to maintain an equal current through Q16 and Q17.

D7 feeds the base of Q25, so as this current flow increases, its collector voltage drops, reducing the drive voltage to the NPN/PNP emitter-follower push-pull pair of Q21 and Q22. This in turn pulls down Q23's gate, reducing the output voltage until it is back

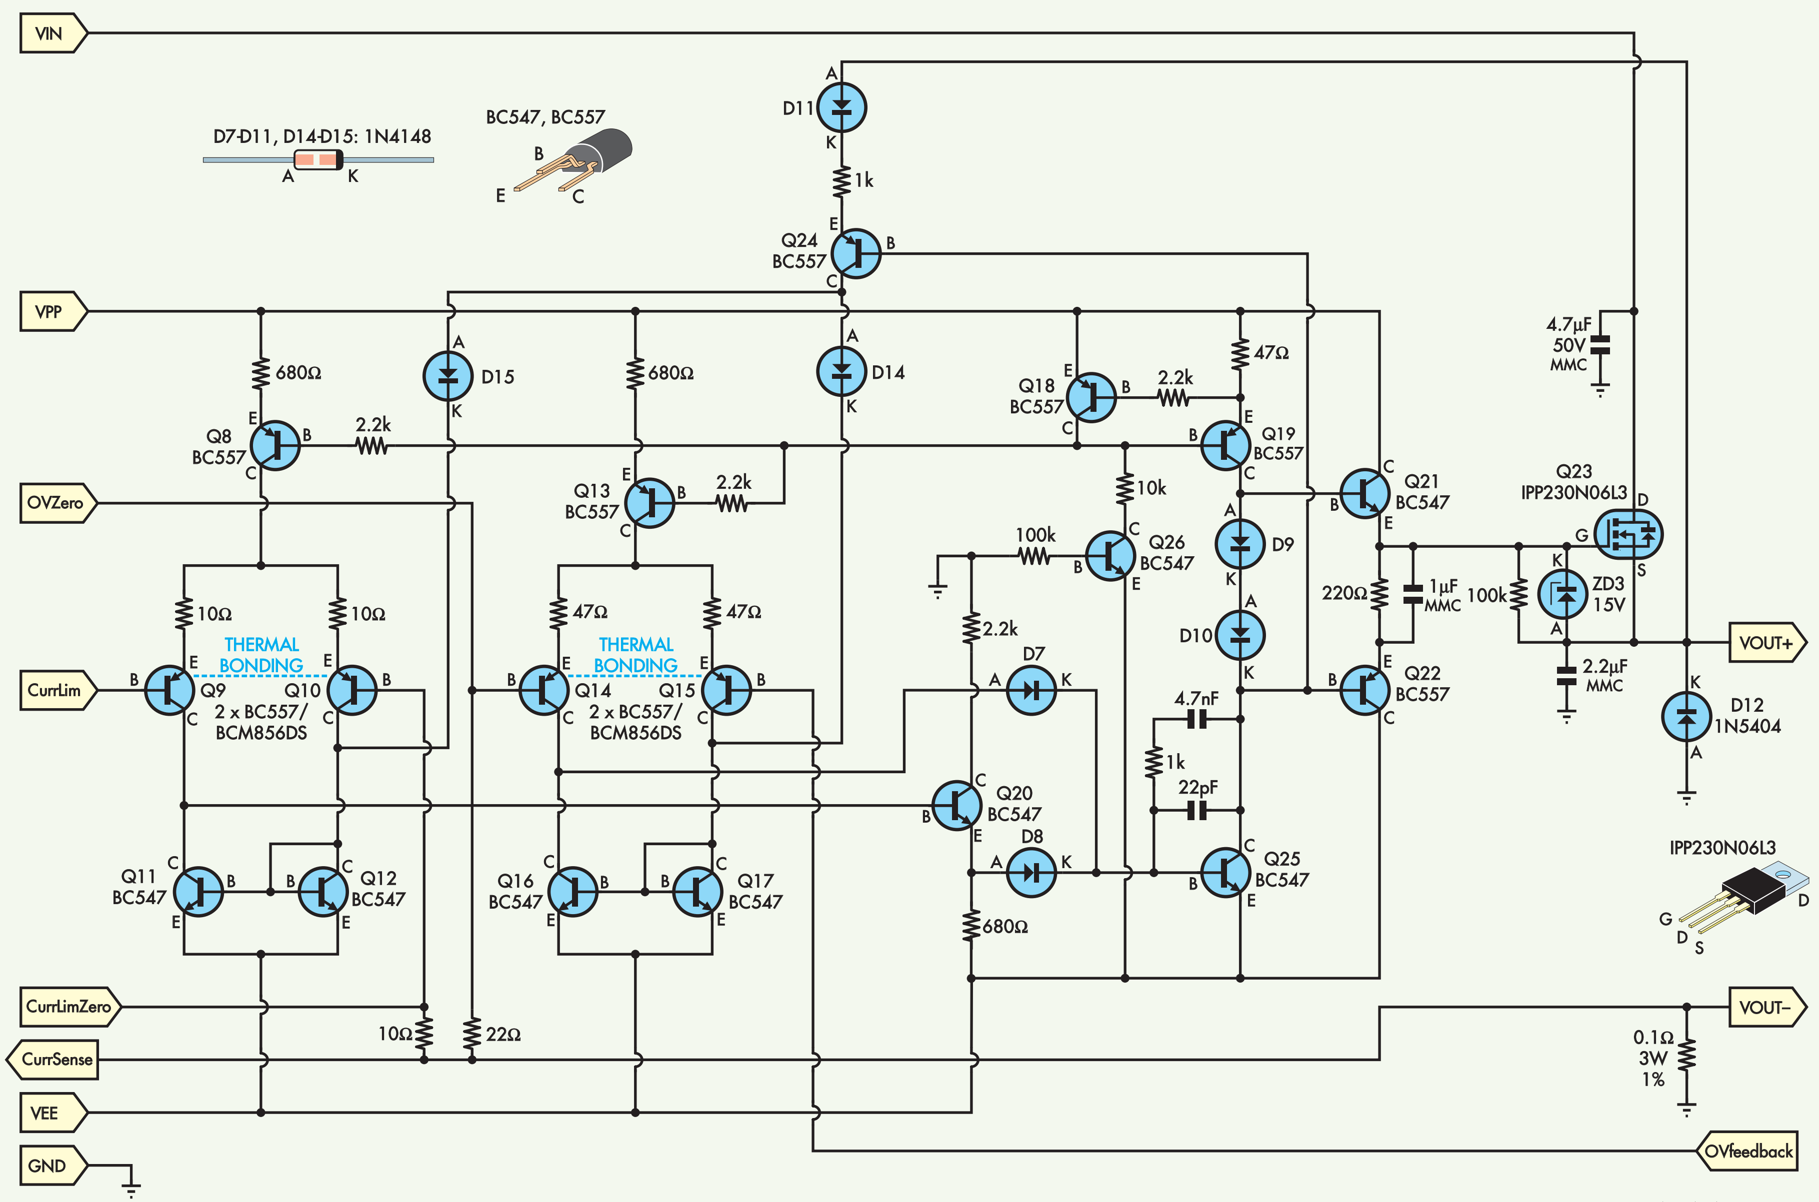

SWITCHMODE-LINEAR HYBRID BENCH POWER SUPPLY LINEAR REGULATION

Fig.6: the linear section of the circuit diagram. This matches up with the switchmode section published last month (Fig.3) via the labelled inputs and outputs. This is essentially a self-contained ultra-low-dropout linear regulator with adjustable current limiting. Q14 and Q15 compare the feedback voltage to a ground reference and regulate the output voltage while Q9 and Q10 enforce the current limit by comparing it against the voltage across the 0.1Ω shunt. There is provision to zero (trim) the output voltage and current limit via OVZero and CurrLimZero respectively.

Parts List

1 double-sided PCB, available from the *EPE PCB Service*, coded 18104141, 198 × 95mm
1 half-rack plastic instrument case with two integrated LED panel meters and SPST rocker switch (available from Altronics)
2 M205 fuseholder clips
1 10A M205 fast-blow fuse (F1)
1 10µH 15.5A 5MHz shielded SMD inductor, 14 × 14mm (L1) (SCIHP1367-100M; Digi-Key Cat 595-1400-1-ND)
2 3.3µH 5.6A bobbin inductors (L2,L3) (RLB1314-3R3ML; element14 Cat 2333682, Digi-Key Cat RLB1314-3R3ML-ND)
1 PCB-mount DC socket (CON1)
1 pair red and black chassis-mount binding posts (CON2)
2 4-way polarised headers and matching header plugs with crimp pins (CON3, CON4)
2 3-way polarised headers and matching header plugs with crimp pins (CON5, CON6)
1 SPDT PCB-mount right-angle toggle switch (S1)
1 small chassis-mount SPDT momentary pushbutton switch (S2)
2 3-pin headers (LK1,S2)
1 2-pin header (LK2)
2 jumper shunts (LK1,LK2)
1 3-pin female header plug or cable with suitable plug (for S2)
2 10kΩ linear 10-turn panel-mount potentiometers (VR1 - VR2) (eg, Rockby 41645***, Element14 1144798/1612609/1386483 etc)
OR
2 10kΩ linear chassis-mount standard potentiometers
4 500Ω mini horizontal trimpots (VR3-VR6; VR7 optional)
1 20kΩ mini sealed horizontal trimpot (VR8, optional)
2 knobs to suit VR1 and VR2
1 ferrite bead (FB3)
3 6073B-type TO-220 'Micro-U' heatsinks
4 M3 × 6mm machine screws/nuts
1 150mm-length rainbow cable or assorted light-duty hookup wires
1 100mm-length tinned copper wire
1 500mm-length extra-heavy duty hookup wire
2 6.8mm female spade crimp connectors to suit extra-heavy duty wire
4 No.4 × 6mm self-tapping screws
2 M3 × 5mm black machine screws
*** Limited stock available

Semiconductors
1 LM5118MH(X) buck/boost switchmode regulator IC (IC1) (Element14 Cat 1606457, Digi-Key Cat LM5118MHX/NOPBCT-ND)
1 7555 CMOS timer IC (IC2)
1 LM2940CT-12 12V 1A low-dropout regulator (REG1)
1 7805 +5V 1A regulator (REG2)
1 79L05 -5V 100mA regulator (REG3)
1 LM285Z-2.5 voltage reference (REG4)
1 IRF1405 or IPP230N06L3 N-channel MOSFET (Q1)
2 60V 100A N-channel SMD logic-level MOSFETs (Q2,Q3) (Digi-Key Cat. 568-10984-1-ND)
3 BCM856DS dual PNP SMD transistors (Element14 Cat 1829188, Digi-Key Cat 568-6834-1-ND) OR
6 BC557 100mA PNP transistors (Q4,Q5,Q9,Q10,Q14,Q15)
1 BC327 NPN transistor (Q6)
1 BC337 PNP transistor (Q7)
6 BC557 100mA PNP transistors (Q8,Q13,Q18,Q19,Q22,Q24)
8 BC547 100mA NPN transistors (Q11,Q12,Q16,Q17,Q20,Q21,Q25,Q26)
1 IPP230N06L3 N-channel MOSFET (Q23)
2 SK1545 45V 15A SMD Schottky diodes (D1,D2) (Digi-Key Cat. SK1545-TPCT-ND)
13 1N4148 signal diodes (D3-D11,D13-D15,D18)
3 1N5819 1A Schottky diodes (D16,D17,D19)
1 1N5404 3A diode (D12)
3 15V 1W Zener diodes (ZD1,ZD3,ZD8)
4 27V 1W Zener diodes (ZD2,ZD5-ZD7)
1 4.7V 0.4W or 1W Zener diode (ZD9)

Capacitors
2 220µF 50V/63V low-ESR electrolytics
8 100µF 25V electrolytics
2 47µF 50V/63V low-ESR electrolytics
9 10µF 25V X5R SMD ceramic, 3216 (imperial 1206) or 2012 (imperial 0805) package
10 4.7µF 50V X5R SMD ceramic, 3216 (imperial 1206) or 2012 (imperial 0805) package
1 2.2µF 50V MMC*
3 1µF 50V MMC*
1 1µF 50V X5R SMD ceramic**
2 100nF 50V ceramic disc or MMC*
4 100nF 50V X7R SMD ceramic**
1 10nF 50V MKT or MMC*
1 10nF 50V X7R SMD ceramic**
1 4.7nF 63V MKT
1 4.7nF 50V X7R SMD ceramic**
1 2.2nF 63V MKT
1 2.2nF 50V X7R SMD ceramic**
1 330pF 50V C0G/NP0 SMD ceramic**
3 100pF 50V ceramic disc
1 22pF 50V ceramic disc

Resistors (0.25W, 1% unless stated)

2 10MΩ	7 2.2kΩ
1 1MΩ	1 1.8kΩ
1 910kΩ	2 1kΩ
5 100kΩ	2 820Ω
1 82kΩ 1% SMD**	5 680Ω
1 15kΩ 1% SMD**	2 470Ω
4 10kΩ	1 220Ω
2 10kΩ 1% SMD**	3 47Ω
1 9.1kΩ	1 22Ω
2 3.3kΩ	5 10Ω

1 10Ω 1% SMD**
1 0.1Ω 1% 3W SMD 6432 (2512 imperial) (element14 Cat 1435952, Digi-Key Cat CRA2512-FZ-R100ELFCT-ND)
OR
1 0.1Ω 1% 3W through-hole resistor (Welwyn OAR-310F or similar)
1 15mΩ 0.75W or 1W SMD 3216 (1206 imperial) (element14 Cat. 1887165, Digi-Key Cat. MCS1632R015FERCT-ND)

** Monolithic Multi-layer Ceramic*
*** These SMD passive components can be in either 1608 (imperial 0603) or 2012 (imperial 0805) packages*

where it should be and the bases of Q14 and Q15 are then at the same voltage.

Conversely, if the output voltage drops, the reverse occurs and the voltage at Q23's base increases to compensate.

Q25 has a 22pF Miller capacitor to reduce the AC open-loop gain at high frequencies, to keep the feedback loop stable, otherwise Q23's gate voltage would not settle down. In parallel with this Miller capacitor is a 4.7nF capacitor with a series 1kΩ resistor. This improves stability when a high-value capacitor is connected to the output of the supply, avoiding excessive voltage overshoot. It also helps stabilise the voltage during current limiting.

The 2.2µF output capacitor is also important for stability and this and the Miller capacitor component values have been chosen as a compromise between stability, fast transient response and a low output capacitance so that the supply can more closely approximate an ideal current source.

Q25's collector load is another constant current source, this time providing around 12.5mA. It's controlled by Q19 which (like Q8 and Q13) is biased by Q18, in turn biased by a 10kΩ resistor to the negative supply via Q26. The 47Ω resistor sets the current through Q19 to 0.6V / 47Ω = ~12.5mA. This value was chosen based on the 625mW dissipation limit for Q25, given the maximum possible voltage across it of around 46V, when the output is at 40V.

Diodes D9 and D10, in series with Q25's collector, bias the base-emitter junctions of buffer transistors Q21 and Q22 so that they are both slightly conducting all the time. This speeds up the 'hand-off' between them as the output switches from slewing positive to negative and vice versa.

The quiescent current through this pair is limited by a 220Ω resistor which is bypassed with a 1µF capacitor so that Q23's gate can be quickly discharged by Q22. Note that the relatively high 12.5mA through this stage is required so that Q23's gate can be quickly charged when the output is slewing in the positive direction, eg, when recovering from a brief short circuit.

Zener diode ZD3 clamps Q23's gate voltage to no more than 15V above its source. Normally, the supply rails guarantee this, but under some conditions it could be exceeded, hence the zener clamp.

PCB design and layout

The main part of the article describes how the circuit works, but that isn't the end of the story. Let's look at a couple of the trickier aspects of the PCB design.

The most obvious place where layout is critical is around the switchmode regulator, ie, IC1, D1, D2, Q2, Q3 and L1. We've purposefully chosen small components here, while also taking into account ease of soldering. The rationale behind this is that by keeping the components small, the distance through which high switching currents must flow is kept to a minimum and thus the resistance and parasitic inductance of the short, wide PCB tracks used is kept to a minimum. The current flowing through these components thus also stays close to the PCB's ground plane.

This ground plane acts as a shorted turn for the various parasitic inductors (transformers) formed by loops in the circuit, but this isn't perfect – it is on the other side of the PCB (~1.5mm away) and does not have zero resistance. So it's good practice to keep those loops as small as possible.

The layout of this switchmode section is based on the demonstration board for the LM5118 IC*. This is a rather clever scheme whereby the current runs around the edge, from the input at lower left up to the top, then across to the right and then down to lower-right. The central area is a large power groundplane to which the IC is connected and there are dozens of vias connecting this to the underside of the PCB where the groundplane covers nearly 100% of the area under the current-carrying components.

The idea behind this is that while current flows in a clockwise direction around this section, the return current flowing through ground goes in an anti-clockwise direction around the groundplane. This is because current follows the path of least impedance (not just resistance) and this is true when the parasitic inductance is minimised, ie, when the return current flows directly underneath the main current path.**

There is a separate analogue ground plane below pins 1-10 of the IC (ie on the underside), above which the analogue components are mounted (eg, compensation and feedback networks). The two groundplanes are joined under the IC. This keeps the switching noise out of the analogue components.

The LM5118 has a large pad on its underside which is soldered to the PCB to provide heatsinking for the IC. The large copper area it's connected to helps draw heat away, too. Since constructors won't necessarily have a hot air or reflow station to solder the IC, we have placed three large vias through this pad and onto the bottom side, so that solder can be flowed through to this pad from underneath.

Finally, a note on the layout of the analogue section. The ground return paths for the two panel meters have been brought back to the earth plane separately from other tracks so that the relatively high current flow (250-300mA) does not create ground voltage shifts to upset the meter readings or other circuitry. After all, the lowest digit on each meter has a resolution of just 0.1mV.

** See Texas Instruments application note AN-1819 at:* ***www.ti.com/lit/ug/snva334b/snva334b.pdf***

*** Some good information on current flows in double-sided PCBs can be found here:* ***www.analog.com/library/analogdialogue/archives/41-06/ground_bounce.html***

Q23's gate voltage is prevented from going below the output voltage by Q24. As the gate goes negative, Q24's base is pulled below its emitter and thus delivers current to the collectors of Q17 and Q12 via diodes D14 and D15. This forces the drive to Q25's base to be reduced, thus preventing Q23's gate from dropping any further. The zener will also do this job, however the advantage of including Q24 is that it also acts to clamp Q25's collector voltage, giving much faster recovery from current limiting.

Q26 ensures a clean start-up when power is first applied. It prevents the bias current for Q8, Q13 and Q19 from flowing to VEE via the 10kΩ resistor until the VEE negative rail has dropped below about –1V. Without this bias, Q19 remains off and so Q25's collector voltage remains low until the supplies stabilise. The 100kΩ resistor across ZD3 keeps Q23's gate discharged

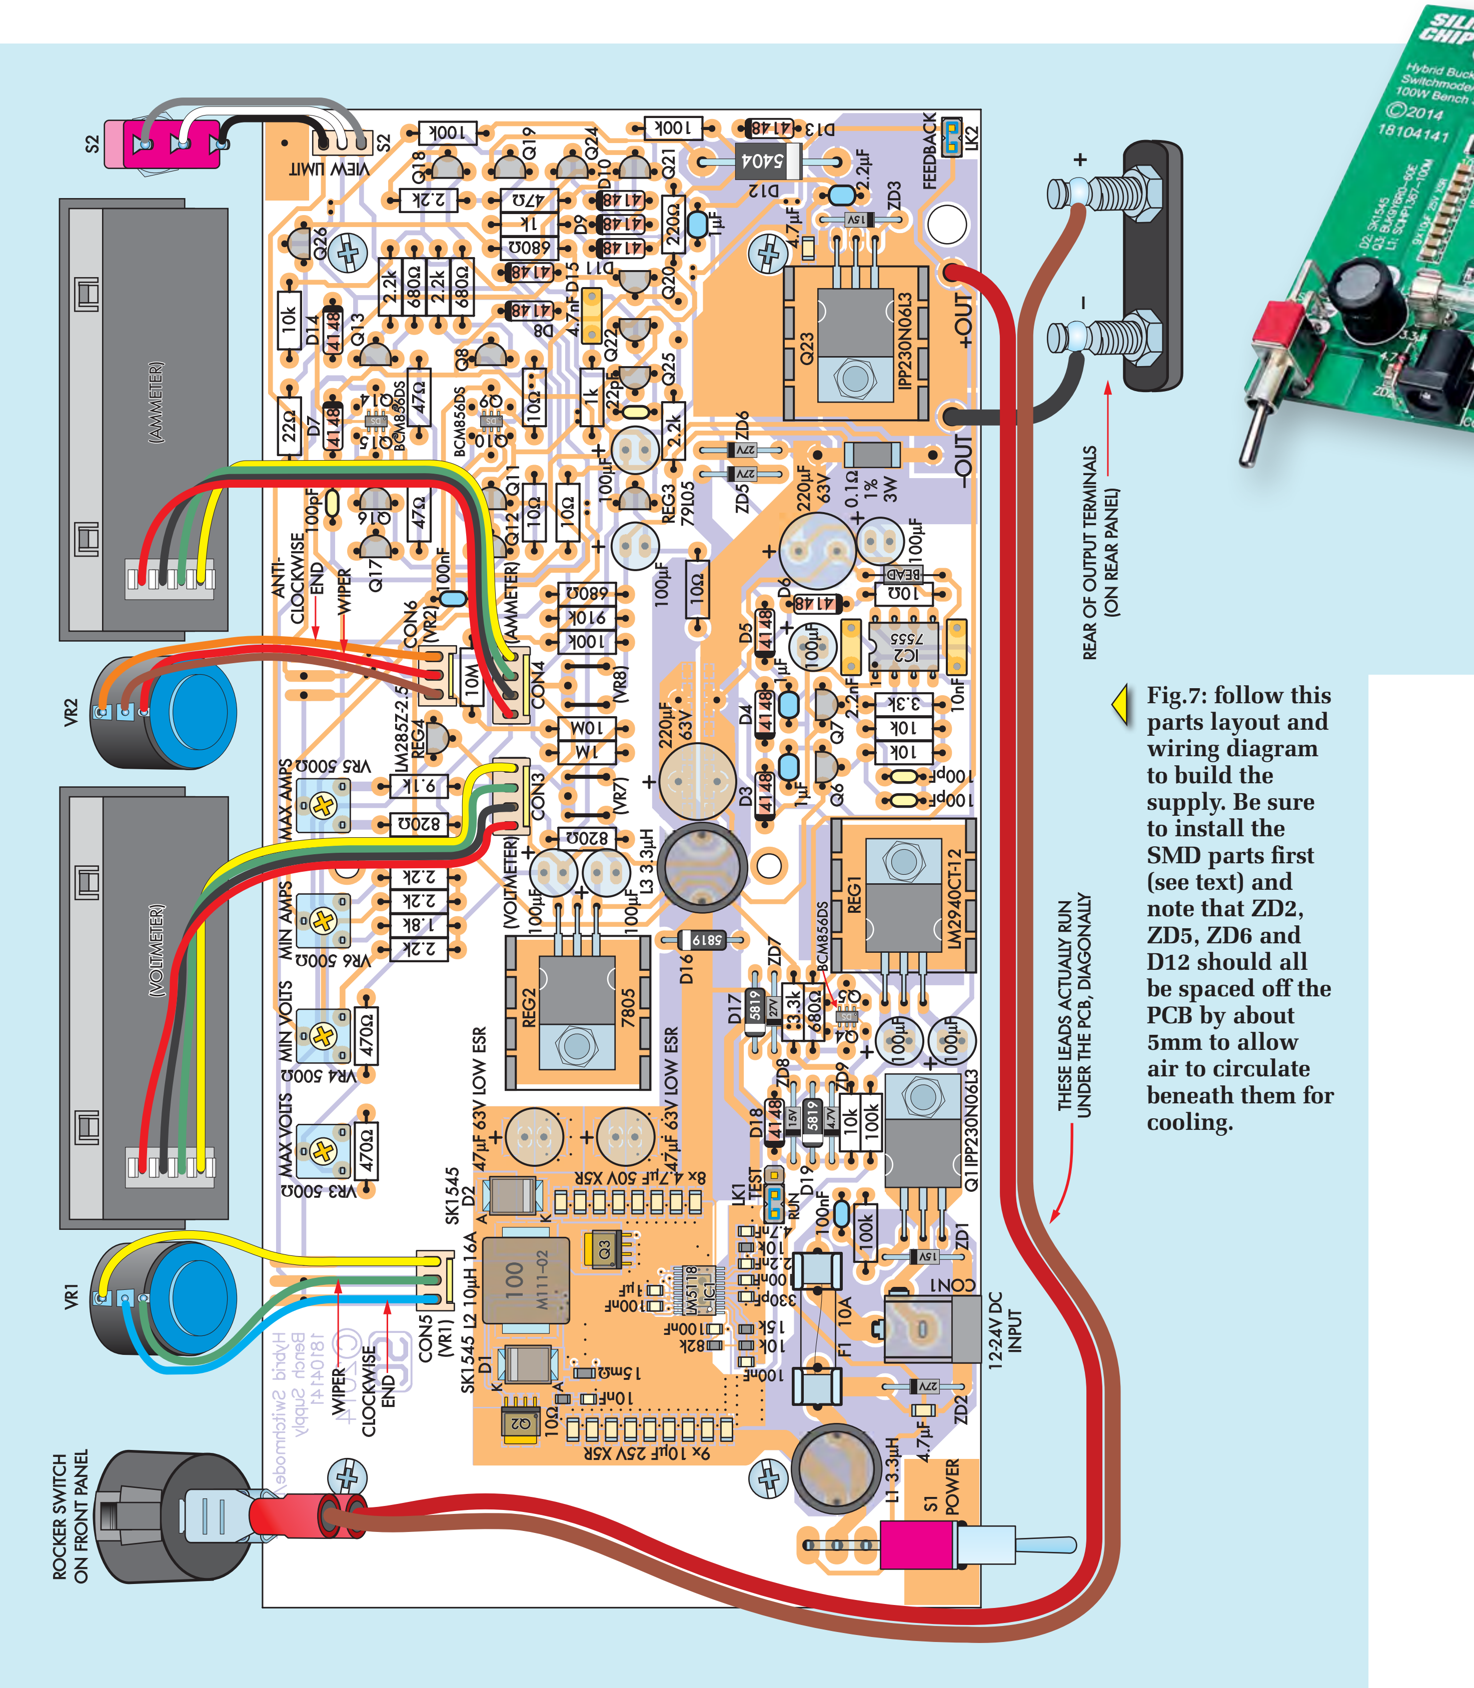

Fig.7: follow this parts layout and wiring diagram to build the supply. Be sure to install the SMD parts first (see text) and note that ZD2, ZD5, ZD6 and D12 should all be spaced off the PCB by about 5mm to allow air to circulate beneath them for cooling.

during this time, so it remains off and the output voltage stays low.

Note that we want to be sure that when OVFeedback is connected to VOUT+ via a low resistance (ie, voltage adjustment pot VR1 is at minimum), the output voltage is zero. This depends on how accurately Q14 and Q15 are matched and also on VR1's tolerances.

These error sources are trimmed out by injecting a small current from 'OVZero' into the base of Q14 (from trimpot VR4, in Fig.3 last month). This develops a maximum of a few

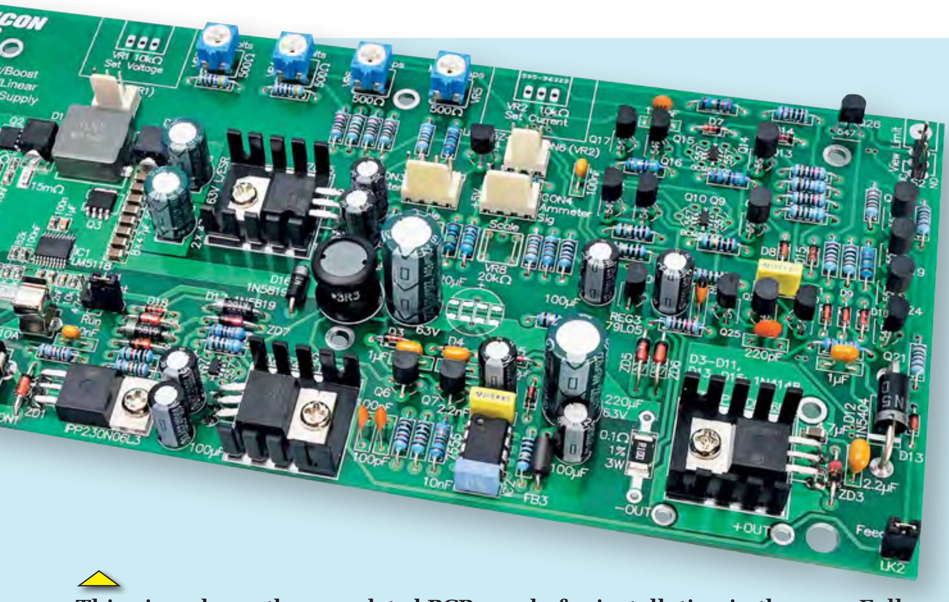

This view shows the completed PCB, ready for installation in the case. Follow the procedure described in the text to solder in the SMD parts.

millivolts across the 22Ω resistor and allows the output to be nudged one way or the other.

Current limiting and regulation

The current-regulation mechanism operates similarly to the voltage feedback described above. Differential input pair Q9 and Q10 is configured identically to Q14 and Q15, except that 10Ω emitter-degeneration resistors are used rather than the 47Ω. That's because the voltage levels applied to this differential pair are much lower (500mV max).

Also, rather than the output current from this differential front-end driving Q25 directly, it's amplified by Q20. This is required because if the output current exceeds the set level by 1mA (for example), the voltage difference between the bases of Q9 and Q10 is just 0.1mV. However, we want the output voltage to drop rapidly as soon as the current limit is exceeded by even a small amount, and so Q20 provides additional current gain of around 250.

Q9's base is the non-inverting input of the pair and this is connected to VR2's wiper. This is the current limit adjustment potentiometer ('CurrLim') and it supplies a 0-500mV signal to set the current limit to between 0 and 5A. Q10's base is connected to the top of the 0.1Ω current-sense resistor via an additional 10Ω resistor. This allows a small current injection from the 'CurrLimZero' input to pull this a few millivolts one way or the other, to cancel out the differences in the base-emitter voltages of Q9 and Q10.

Thus, if the voltage across the sense resistor exceeds the voltage from VR2's wiper, current flow from Q9 to Q20's base increases and thus D8 becomes forward-biased. This pulls Q25's collector down and thus reduces the output voltage until the current flow stabilises at the set level.

Q20 is linearised by a 680Ω emitter resistor for a more progressive action (no additional Miller capacitor is needed). A 2.2kΩ collector resistor (from ground) limits the current delivered to Q25 under a hard short-circuit condition and thus limits the dissipation in Q25. Diode D12 protects the circuit in case an external load pulls the output terminal negative.

Remaining circuitry

The regulator's positive rail is labelled VPP and comes from the charge pump described last month (see Fig.3). This tracks VIN and is generally boosted to be 10V higher. VEE is a regulated –5V rail, derived from the same charge pump. These supply rails are 'wider' than the input supply (VIN) and ensure that Q23's gate can vary over a wide enough range (approximately 0-43V) to control the output over the full range of 0-40V.

There are a couple of aspects of this circuit which are not obvious at first glance. The top end of the 0.1Ω current-sense resistor is connected both to 'VOUT-' and 'CurrSense'. The former is for return current to flow from the load while the latter goes to the ammeter divider circuit. This sense resistor can dissipate up to 2.5W at 5A; its value was chosen as a compromise between keeping the dissipation reasonable and giving enough of a voltage swing for the current-limiting circuitry to operate quickly and accurately.

In addition, the non-inverting input for voltage regulation (Q14's base) is connected to the top of this resistor (ie, CurrSense) rather than to ground as you might expect. This is because we don't want the output voltage to drop as the load current increases due to the increase in voltage across the current sense resistor.

This connection for Q14, in combination with the fact that the output divider's –2.5V reference is also connected to the top of this sense resistor (via Vout–; see Fig.3 last month), means that the voltage as set by VR1 is actually VOUT+ – VOUT-. As a result, when VOUT- increases, so must VOUT+ due to the negative feedback action.

Minor changes

Since publishing the main circuit diagram last month, we have made a few minor changes to the circuit. First, we added a 100µF bypass capacitor at the input of REG2, as the latter was moved away from REG1 and its 100µF output filter capacitor to aid heat dissipation.

We also increased the value of the 2.2MΩ resistor (near VR8) to 10MΩ and reduced the associated 1kΩ resistor to 680Ω, so that the ammeter reads zero when there is no current flow.

In addition, the 680Ω resistor connected to VR5 has been changed to 820Ω to ensure that the maximum current can be set to 5A.

Finally, the 1kΩ resistor at the ground end of VR7 has been reduced to 820Ω, although if you link out VR7 as suggested, either value will work.

Building it

Despite the circuit complexity, the assembly is quite straightforward. All the parts are mounted on a PCB coded 18104141 and measuring 198 × 95mm. This PCB is available from the *EPE PCB Service*. Fig.7 shows the parts layout and external wiring details.

It's easiest to fit all the surface-mount (SMD) parts on the PCB first, starting with IC1. Begin by removing it from its packaging and locating the pin 1 dot (which isn't all that obvious).

If using a hot-air station or reflow oven, it's simply a matter of sparingly applying fresh solder paste to the pins

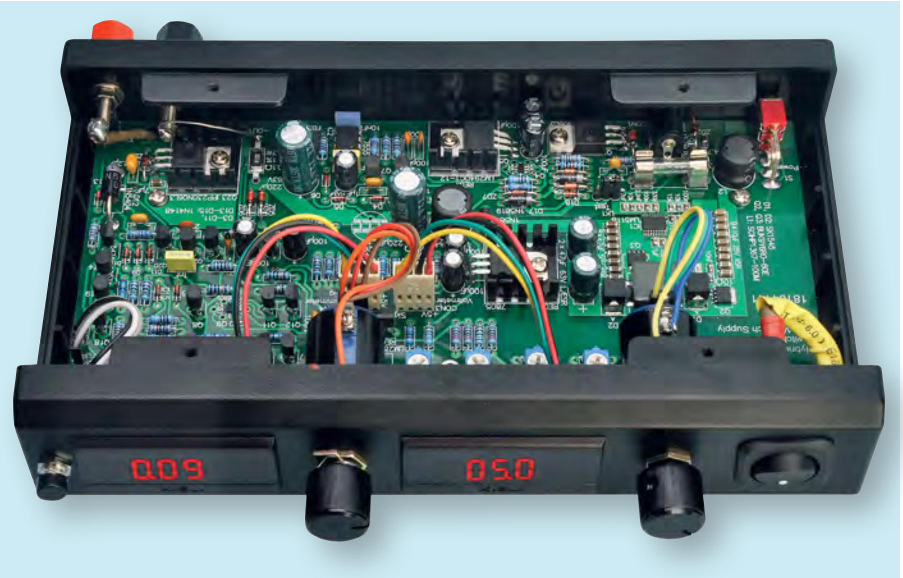

The PCB fits neatly inside this standard instrument case which comes complete with 3.5-digit LED readouts for simultaneous voltage and current display. The final assembly details are in Part 3 next month (prototype PCB shown).

and central pad, placing the IC with the dot at lower-left and then heating the device until the joints are formed. Or, provided you have a temperature-controlled soldering iron with a fine bit, you can certainly do it by hand.

The procedure for hand-soldering IC1 is as follows. First, it's a good idea to place a blob of flux gel on the central pad. This helps hold the IC in place during soldering and will also come in handy when it's time to solder the thermal pad. That done, put a very small amount of solder on one of the pads, check that its orientation is correct, then heat this solder and slide the IC into place.

Now check that all the pins are correctly aligned with their pads using a magnifying glass. Each pin should be on top of its associated pad and not protruding over the side. Unless you're very lucky, it won't be right the first time, in which case you need to reheat the single soldered pad and give the IC a gentle nudge in the right direction. Repeat this procedure until you're happy with the placement, then solder the diagonally opposite pin and re-check the orientation.

Having tacked it down, you now have several possibilities for soldering the remaining pins. You can solder two or more pins at a time by placing a standard chisel/conical tip between a pair of pins and feeding a small amount of solder in (pre-fluxing the pins helps) or you can run flux down both sides of the IC and then drag solder all the pins in one go using a hoof or mini-wave tip (this can also be done with standard tips but not as easily). In either case, using a high-quality flux paste makes the process a lot easier.

Don't worry if any or even all the pins are bridged after soldering; it's just a matter of applying some more flux paste, placing some solder wick over the bridge and then heating it until it 'sucks up' the excess solder. Just be careful not to pull on the wick or otherwise apply force until the solder has melted or you could damage the IC pins or the board.

Once the chip is in place, flip the board over and apply some liquid flux or flux gel to the three vias under IC1. Now melt some solder into these holes; there is an exposed pad surrounding them to make it flow better. The combination of capillary effect and flux on both sides of the board should cause the solder to flow through and form a joint between the IC and the pad on the other side. You can confirm this is the case by touching the IC after doing this; it should be quite hot due to the solder that's adhered to its thermal pad.

Finally, clean off any flux residue or other contaminants using a good solvent and then carefully examine the joints using a magnifying glass and lamp to ensure that every pin has been soldered to its pad and no bridges remain.

(Note: if you have to remove this IC, it's easy to do using a hot-air gun (which can be bought quite cheaply). It's virtually impossible to do using any other method, without damaging the board.)

SMD MOSFETs

These are next on the list. Again, you can use solder paste and hot air/reflow; however, these can also be soldered using a regular iron. Put some flux on the large PCB pad for the tab and then flow a small amount of solder onto it. Ensure this forms a thin, even layer on the pad; if not, clean off the excess using solder wick.

Now do the same to the underside of the MOSFET tab itself; you will probably have to place it in a mini vice or some other clamp while doing so.

Next, spread some flux paste onto the tinned PCB pad, then put a small amount of solder onto one of the four smaller mounting pads on the PCB. It's easiest to start with the topmost (Q2) or right-most (Q3) pad (ie, the MOSFET gates) as these have the least thermal mass.

As with IC1, heat this solder and slide the MOSFET into place, then check that its tab and remaining pins are properly centred on their pads. You can now apply a little flux along the edge of the large tab and heat it with your soldering iron until the layers of solder on both melt and merge.

That done, it's just a matter of letting it cool a bit, soldering the three remaining pins, then touching up the gate joint (ie, the first small pin you soldered) using a bit of extra flux.

You can now use a similar procedure to fit inductor L1, ie, tin the underside of its leads, tin the PCB pads, add flux to both, then flow them together. Press down gently on the inductor while soldering the second pad and then re-flow the first one to ensure it's right down on the board. Note that it might be necessary to add more solder to the initial pad, to form a good joint. That's the easiest method we've found to solder such a large SMD component.

While doing this, you will need to hold it with tweezers or similar as it gets hot. Use this same technique to fit diodes D1 and D2, but be careful with their orientation; their cathode stripes face in opposite directions.

Now fit the remaining SMD capacitors and resistors. For each, it's just a matter of placing solder on one of the pads, sliding the part in, waiting for the joint to cool and then soldering the other side. Be careful – it's quite easy to get solder to flow onto the part but not the PCB pad; use plenty of flux, clean off the residue and inspect the joints carefully under magnification.

Note that the resistors will be labelled with their value (eg, 15kΩ = 153, 10Ω = 100). On the other hand, the ceramic capacitors are not labelled and you will have to check the value on the packaging before removing them. Note also that there are some SMD components away from the switch-mode section. These are the two additional 4.7µF ceramic capacitors near lower-left and lower-right and the 0.1Ω shunt near the negative output terminal. You can use a through-hole shunt instead of an SMD type, if you can get one with that will fit and has the correct rating.

Now is also a good time to solder in the three BCM856DS dual transistors (assuming you are using these, as recommended). Their pin layout is symmetrical so orientation doesn't matter. You may be able to solder the pins individually, then clean up any bridges with flux paste and solder wick. In fact, we like to add some flux and apply solder wick anyway as reflowing the joints in this manner gives a more consistent and reliable result.

Through-hole parts

With the SMDs out of the way, the next job is to fit the passive through-hole components, starting with the 1N4148 diodes. Note that these do not all have the same orientation, so check the layout diagram (Fig.6) carefully. Follow with the resistors (check the values with a DMM if unsure) and then the medium-sized diodes such as the zeners and 1N5819s, again taking care with the orientation.

Leave 27V zener diodes ZD2, ZD5 and ZD6 out for now. Diode D12 (1N5404) should also be left out at this stage.

Now solder IC2 in place (don't use a socket), making sure that its notch or dot faces towards the top. Follow with ferrite bead FB3 (FB1 and FB2 were removed from the design). If you have a plain bead without leads, run a component lead off-cut through it.

The next job is to fit Q1. First, bend its leads down through 90° about 6mm down from the tab, then feed them through and fasten its tab down using an M3 × 6mm machine screw and nut. Solder and trim the leads, then install

Another view inside the case, this time from the rear. The rear panel carries the power switch (S1), a hole to access the DC socket and the two output terminals (note: prototype PCB shown).

all the ceramic capacitors, including the multi-layer types, followed by the MKT capacitors.

That done, fit the zener diodes you left out earlier (ZD2, ZD5 and ZD6) but space these off the board by about 5mm to allow the air to circulate beneath them for cooling (they get hot if they conduct a significant amount of current). Next come the trimpots; remember that you probably don't need to fit VR7 and VR8 but if you do, VR8 should be the 20kΩ pot. If you aren't fitting them, each should have two wire links soldered in its place where shown.

You can now install the pin headers for LK1, LK2 and S2, followed by the two fuse clips. Check that these have the retaining lugs on the outside or the fuse will not fit and make sure they are fully inserted before soldering.

Next on the list are all the TO-92 devices, ie, the small signal transistors plus REG3 and REG4. Check their markings to ensure each one goes in the right place and bend the leads with small pliers if necessary, so that they fit the PCB pad layout. Note that if you are not using the BCM856DS chips, you will also need to install BC556s for Q4, Q5, Q9, Q10, Q14 and Q15. Be sure to smear thermal paste on the faces of these transistors and push each pair together so that they are in close contact.

The remaining diode (D12) can go in. It too should be spaced off the board by about 5mm. Follow this with the smaller electrolytic capacitors (47µF and 100µF), all of which are inserted with the positive (longer) lead towards the top of the board. The two bobbin inductors (L2 and L3) and the DC socket (CON1) can then go in. Be sure to push the latter all the way down onto the PCB before soldering.

Now for the heatsinked TO-220 devices, ie, REG1, REG2 and Q23. Don't get these mixed up; they are installed in the same manner as Q1 except that you will need to slide the heatsink under the device package before fastening each assembly down using an M3 x 6mm machine screw and nut. If you like, you can smear some heatsink paste on each device tab before its fastened down, although this isn't strictly required.

Finally, complete the PCB assembly by fitting the power switch (S1), polarised connectors (CON3-CON6) and the two 220µF electrolytic capacitors. If you're using a through-hole shunt, don't forget to install that too.

Next month

That's all we have space for this month. In the final article next month, we'll run through the test procedure and the trim adjustments, describe how to build it into the case and give some tips on using it.

Reproduced by arrangement
with SILICON CHIP
magazine 2016.
www.siliconchip.com.au