Construction project:

Versatile 40V/3A lab power supply - 1

This new supply should handle all but the most specialised tasks on a home or professional workbench. Based on readily available components, it has full electrical and thermal overload protection, adjustable current limiting, dual meters, and can be used as either a single or dual-tracking supply.

by ROB EVANS

In the August 1991 issue of *Electronics Australia*, we presented a 'mean, lean, recession-compatible' 18V/1A benchtop supply, which included what we felt were the essential features for a useful but low-cost unit. This turned out to be a very popular project, and we subsequently described two beefed-up versions rated at 18/2A and 30V/1A in the following October issue.

It's a safe bet however, that there are a reasonable number of us that need a rather more elaborate supply, and are prepared to pay the associated premium — provided of course, that it still represents good value for money. Such a supply would need to have a higher output voltage and current capability and a wider range of features, while hopefully maintaining the simple construction techniques and cost-effective design of our original 18V/1A unit.

To achieve this end, we once again took a careful look at just how such a supply is mostly used, and included features into the new design based on this knowledge — while keeping a close eye on the cost of the individual components.

Our final design bears little resemblance to the previous 18V/1A supply as it happens, and uses a larger case and power transformer, has meters for both voltage and current readings, sports two heatsinks on the rear panel, and uses a much larger circuit board and number of components.

Until you notice the additional '0V' binding post on the front panel, the unit looks like a reasonably conventional 40V power supply with variable current limiting. It's the combination of this connector and the current limiting control that sets the design apart from all previous efforts, since now we effectively have a dual-tracking supply *with* fully variable current limiting.

These two features have been quite mutually exclusive in the past, where single-ended supplies tended to offer variable current limiting, and dual-tracking units would have a fixed maximum current limit to protect the supply itself. One reason for this shortcoming in dual-tracking supplies is the difficulty in deriving a simple circuit which is able to both monitor and shut-down the current in each half of the supply, yet can be adjusted by a single 'dual tracking' current limit control.

FEATURES

- Single or dual tracking operation
- Variable current limiting in two ranges
- Individual Volts and Amps meters
- Over-temperature shutdown
- Output load switching
- Current setting switch
- Fully floating circuit with separate earth terminal
- Dropout (output ripple) indicator

In our new circuit however, we've come up with a relatively straightforward arrangement using just three op-amps, which monitors the current in both sections of the supply, produces an output proportional to the *greater* of the two, then limits the supply's output in response to this level when compared to a preset current limit control. And as a bonus, its output can also drive the supply's current meter, which will then display whichever current is greater.

In practice though, both the meter readings and current limiting action are just as you would expect from a conventional supply, and you're not aware that the circuit is picking the higher figure. Note that the two currents will be almost identical when the unit is used as a single-ended or dual-tracking supply, and one will be at zero when only one half of the supply is used — so in all cases, both the meter and the limiting circuit are responding to a realistic level.

The end result is a very flexible arrangement, allowing the supply to perform a wide range of tasks with a minimum number of controls.

As you can see from the summary panel it offers both a host of useful features and impressive performance, and by our reckoning, represents excellent value for money. To be optimistic about Australia's economy, we might present this project as a mean, not so lean, *recovery*-compatible power supply...

Circuit description

As you can see from the schematic, we've used a number of op-amps to perform the supply's various functions, rather than calling on one or more dedicated regulator ICs (such as the LM723) as used in the previous circuits. This allows for a more flexible design approach and poses few restrictions on the unit's performance or features.

Despite the reasonably high component count though, the circuit itself is quite straightforward and can be divided into several sections, with each op-amp and its associated components performing quite a distinct function.

In general terms, IC1A (with Q1 and Q2) forms the positive regulator stage, and IC1B 'measures' its output current across R4 or R5, while IC2A (with Q3 and Q4) and IC2B (with R14 and R15) perform the same functions for the negative supply.

Other than that, IC3 acts as a temperature detector, IC4A senses the

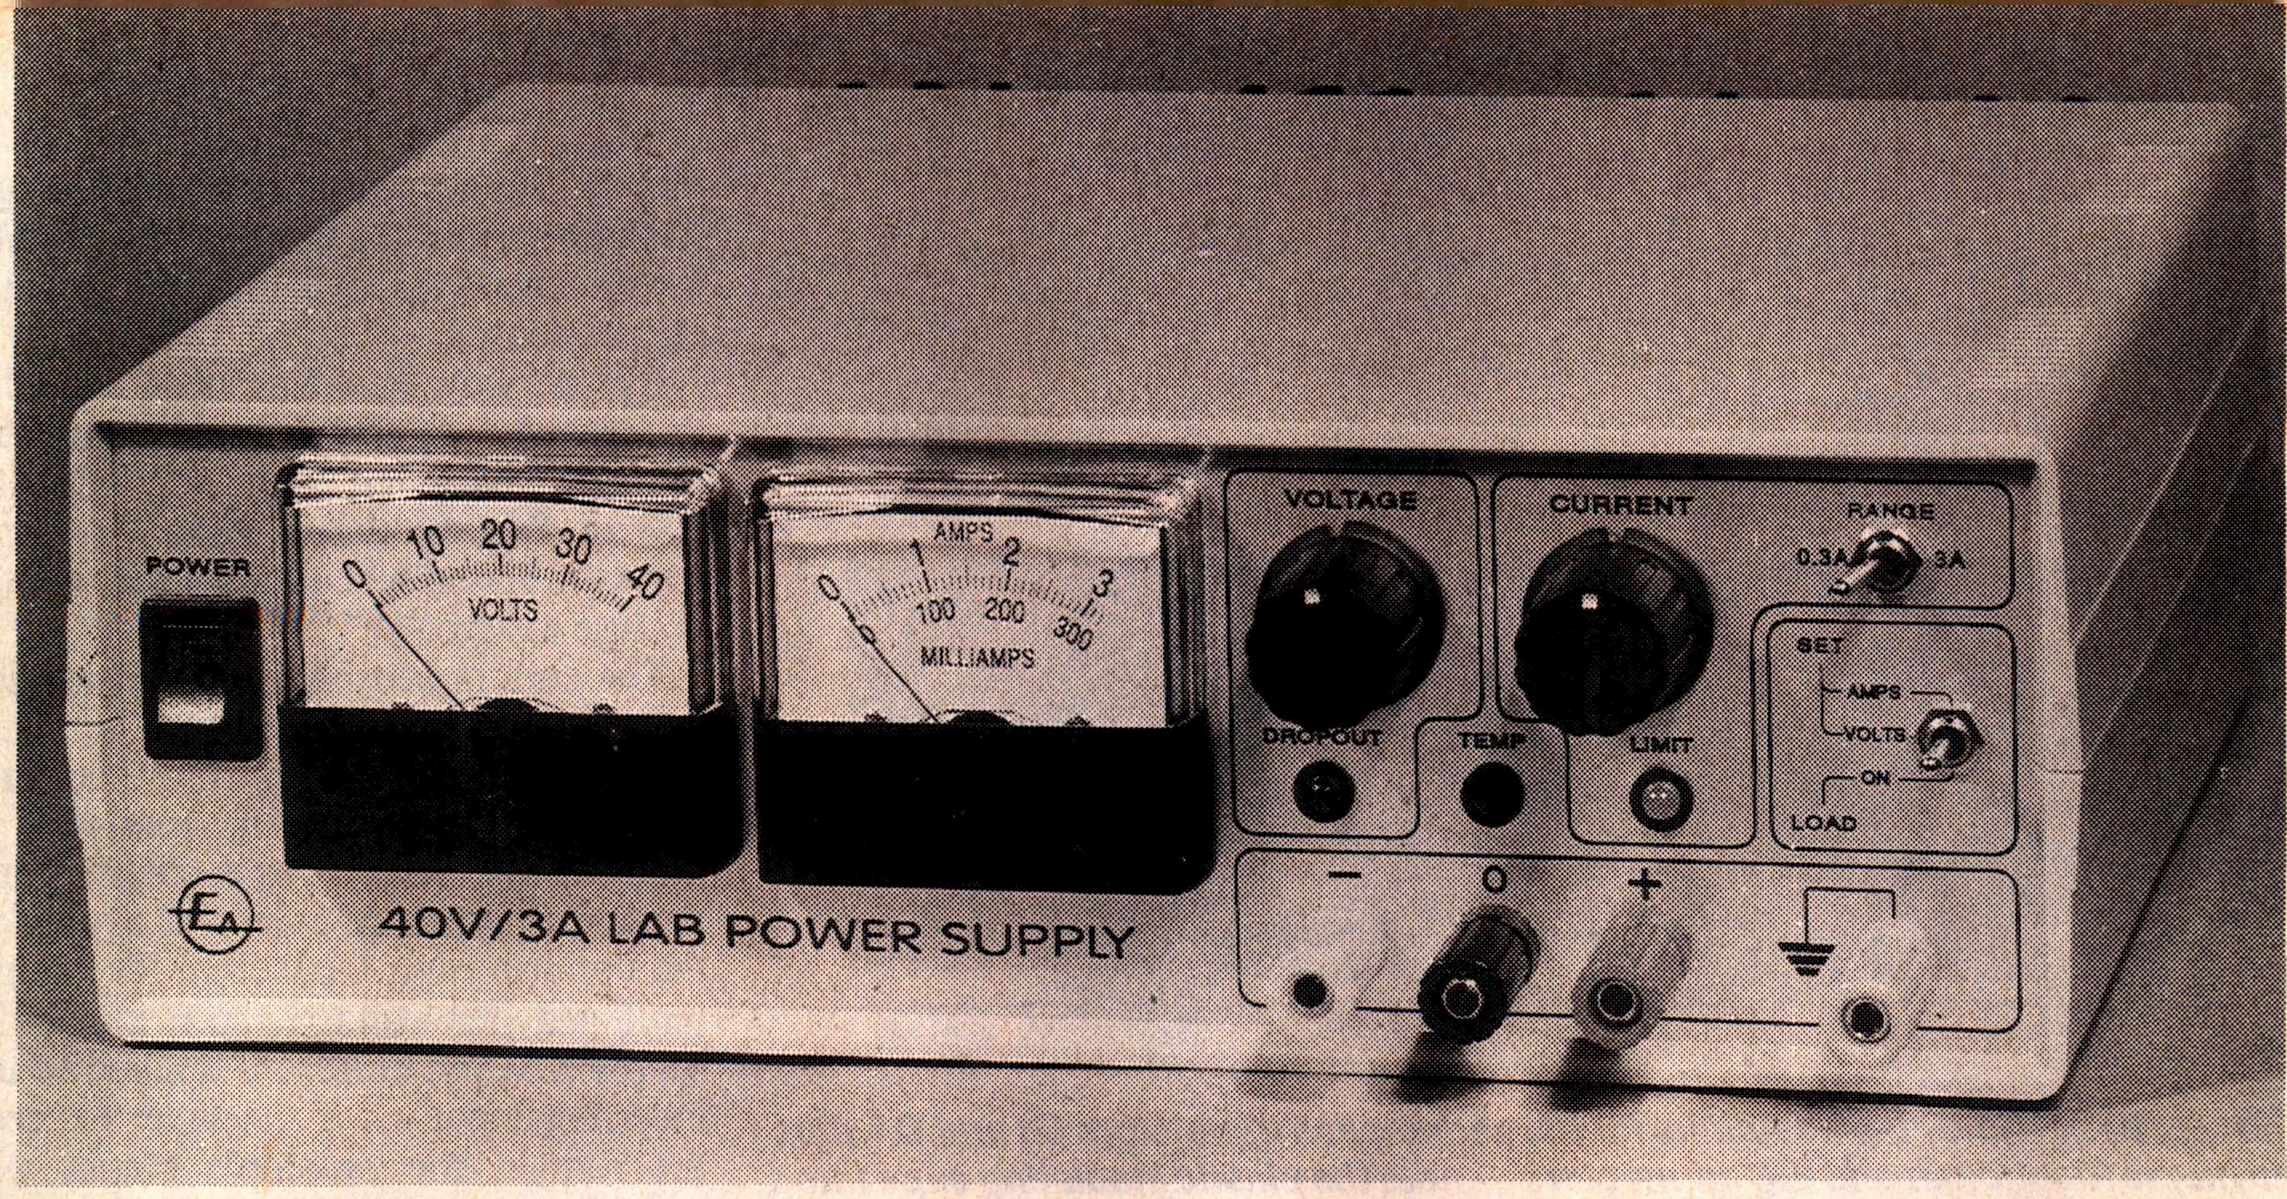

overall output current, and IC4b monitors the supply's output terminals for excessive ripple.

In more detail, power for the supply is derived from the 15-0-15 volt secondary winding of TX1, and full-wave rectified by the encapsulated diode bridge BR1 producing a raw supply of about +/- 21V (referenced to 0V or ground). This in turn is filtered by reservoir capacitors C1 (positive supply) and C2 (negative), then applied to both the main circuitry and the 12V regulators IC5 and IC6 — these provide +/-12V rails for the low-powered circuitry.

The main positive regulator circuit is based around IC1A, and can be considered as a standard non-inverting amplifier, with its output current capability boosted by the darlington pair formed by Q1 and Q2. This stage has a gain of two, as set by the combination of R1 and R2, includes both the darlington pair and resistors R4 to R5 in its feedback loop, and has its input connected to the wiper of the Voltage pot, RV1.

Note that to assist in understanding and faultfinding the circuit, we've included a range of typical voltage readings on the schematic itself. These are for an output voltage setting of +/-10V (20V overall), and correspond to a 1A output current — which would be flowing into a load of 10 ohms per side (or 20 ohms overall) connected via the output switch SW2.

If we also assume for the moment that the voltage across RV1 is constant (at around 10.5V) and the wiper is set to a level of 5V as shown, then the positive regulator's output will be at 10V since the op-amp will drive its output so as to make the voltage at its two inputs (pins 5 and 6) at the same level (in this case, 5V). And of course if the output voltage is say tending to fall, the reduced level at the op-amp's inverting input (pin 6) will cause its output to rise, thereby increasing the drive to Q1 and Q2 and correcting the fall.

The remaining parts of this stage include C4 to restrict the overall bandwidth, C3 to remove any stray interference induced in RV1 (as it's remotely connected to the PCB), and power supply bypass components R3, C8 and C9.

R3 has been included to reduce the overall voltage applied to IC1 to a safe level of around 32V when the supply is unloaded, where the raw positive rail rises to about 22V. Without the small drop induced by the 680 ohm resistor, the voltage across the TL072 would be around 34V, which is uncomfortably close to the IC's maximum supply rating of 36V.

Next in line is the current sensing stage, based around the differential amplifier IC1B, which monitors the voltage across the regulator's current sensing resistors R4 or R5/R6 as selected by the current range switch SW1. In this circuit the two input networks (R16 to R19 and RV2) set the *differential* gain to 4.5 and the *common mode* gain to near zero, so that voltage levels common to both inputs (the supply's output voltage) are ignored and only the voltage drop across the sensing resistors (proportional to the supply's output current) will be amplified.

In the situation depicted in the schematic, R5 and R6 are bypassed by SW1A (shown in the 3A position), which leaves R4 as the sensing resistor. As is indicated by the voltage levels, the 1A current flowing into the load induces a drop of 0.22V across R4, which is then amplified by IC1B to produce an output level of 1V at the cathode of D3 — more about this diode in a moment.

As you would also expect, when the load is disconnected and the voltage drop *across* R4 falls to zero, the differential amp's output will drive to 0V — in short, we have a current-to-voltage converter with floating inputs, where in this case, the output rises at a rate of 1V per amp. This then drives the current reading meter M1 (via RV6 and R32), and the following current limiting stage based around IC4A.

When we wish the supply to have a current reading and limiting range of 0.3A rather than 3A as above, SW1 is moved to the 0.3A position where R4 is bypassed and R5 and R6 are now in circuit. If the parallel combination of these resistors had the value of 2.2 ohms, we would then have a theoretical output rate of 1V per 100mA from the differential amp — the required ten-fold increase in sensitivity.

In practice however, we have made up the required value from the parallel combination of a 2.7 ohm resistor and a 12 ohm 'padding' resistor, where the latter is used to adjust the overall value so that we have *exactly* ten times the resis-

SPECIFICATIONS

Output voltage: 0 to 40V (0 to +/-20V)
Output current: Up to 3.5A (see Fig.1)
Current limit: Approx 30mA to 3.5A in two ranges
Load regulation: Better than 0.2% at 3A output current
Output ripple: Less than 5mV at 3A output current
Overload duration: Indefinite

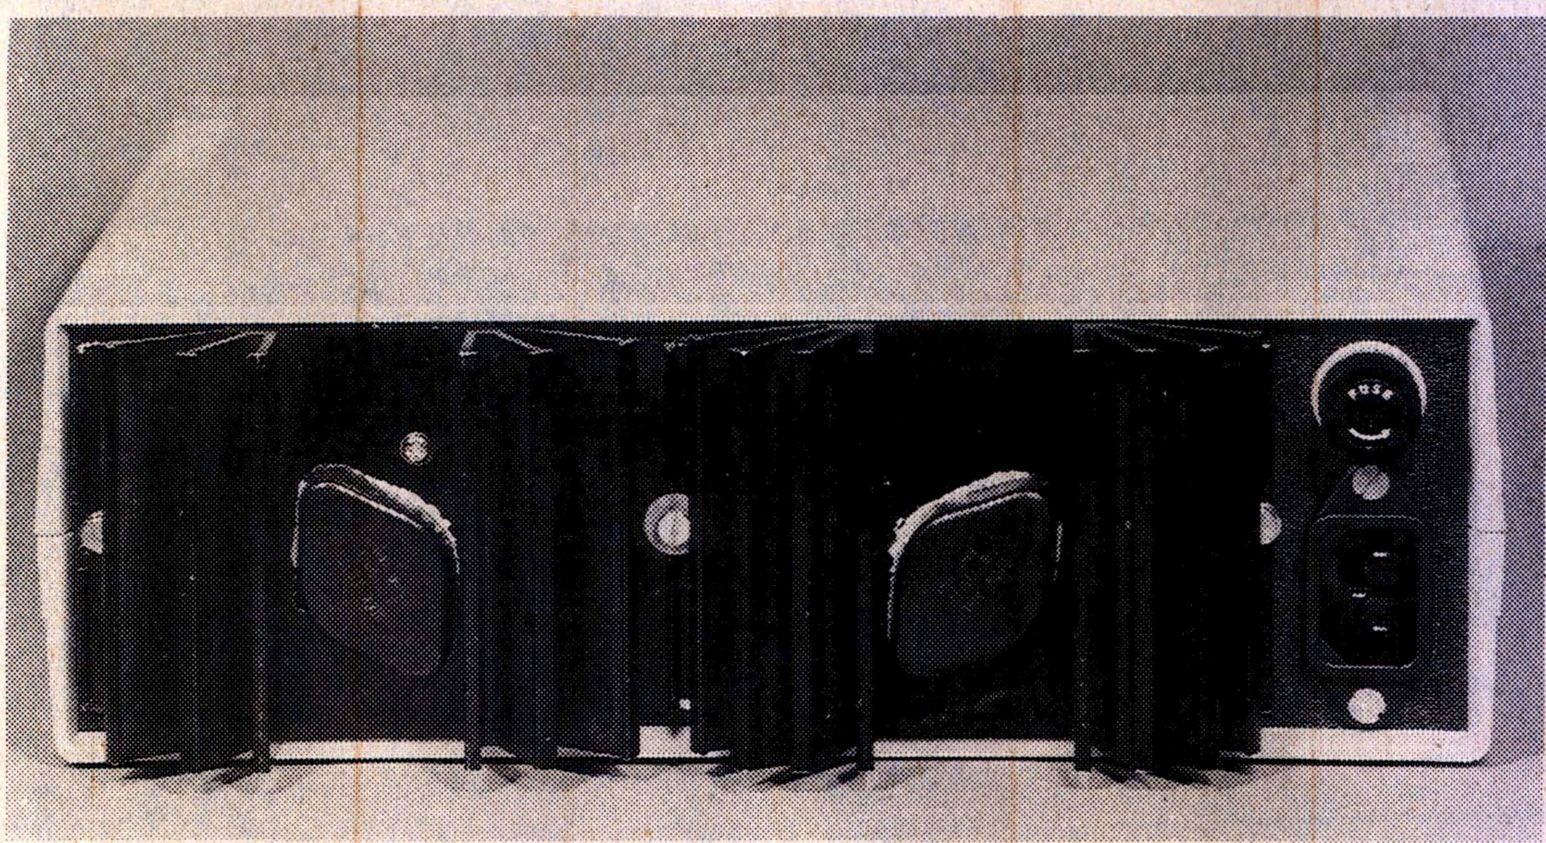

The rear panel holds the heatsinks for the series-pass transistors Q1 and Q4, plus the mains fuseholder and IEC connectors — note the plastic insulating caps fitted to the transistors.

tance that was present when SW1 is in the 3A position.

Since the resistors involved won't precisely match their rated value, and SW1 and its associated wiring will effectively add resistance to R4, we need to adjust one of the sensing resistors so that both the current readings and limiting settings make sense between ranges.

While our prototype supply showed only a small error between current ranges with just a 2.2 ohm resistor installed (rather than the parallel pair), we felt that the variables involved would surely work against some constructors, and the trimming scheme was required. Fortunately, the adjustment process itself is quite straightforward (see 'Setting up, testing' in part 2).

The remaining components in the differential current amp are the input balancing trimpot RV2, and the output isolating diode D3 — which allows two of these stages to control a single output line in an OR fashion — as you've probably noticed, the negative supply based around IC2A drives an identical current sensing amp, formed around IC2B, which despite having negative rather than positive voltages at its inputs, operates in the same manner.

While a first sight this ORing scheme appears quite simple — since two diodes alone form an OR gate — its actual operation is quite subtle. Considering the action of IC1B with the voltages as shown on the circuit, you can see that any fall in the final output voltage (at the cathode of D3) will be automatically corrected by the negative feedback action via R19 — note that D3 is always forward biased in this scenario.

However, if the final output is tending to rise (say, pulled high from some other drive source), the negative feedback action will force the op-amp's output (pin 1) to drive low, in an effort to correct the situation. Diode D3 then becomes reverse biased, and the op-amp is no longer driving the output, which continues to rise. As you've no doubt gathered, this 'other' drive source is in fact the output from the negative supply's current sensing amp (IC2B), which is (say) responding to an increase in load current.

Thus we have a situation where the section of the supply (that is, positive or negative) which is handling the most output current will be driving the current meter and limiting circuits. With the voltage levels shown on the schematic, the negative supply is handing slightly less current than the positive, and consequently, the output of IC2B (pin 7) has driven low. Of course, if the output current was flowing *directly* from the positive to negative output terminals, this would mean that the negative current sensing circuit has a slightly lower sensitivity to that of the positive current sensor.

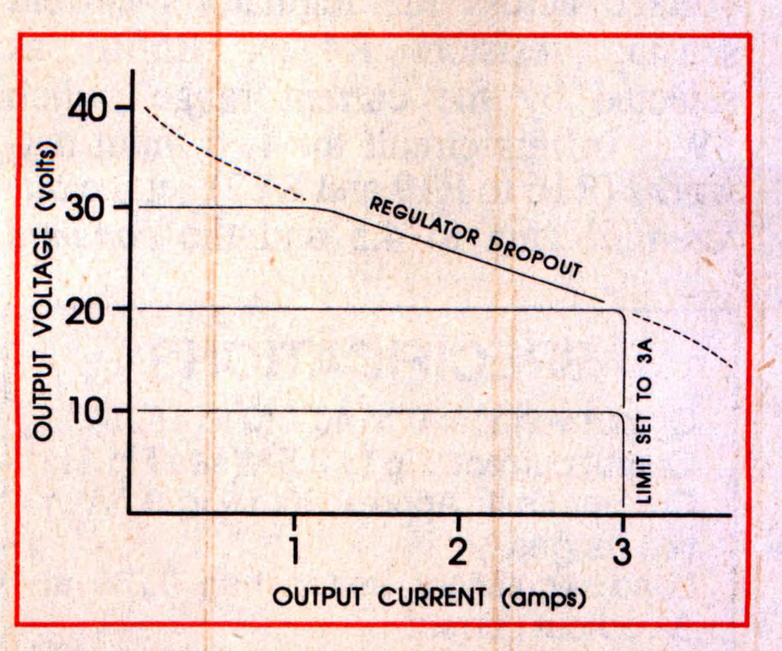

Fig.1: The supply's output current capability at voltage settings of 10, 20 and 30 volts, with the current limit set to 3A.

By the way, R24 has been included to provide a standing bias current through whichever diode is conducting (D3 or D4), so that it cannot become reversed biased when the output voltage is close to zero.

The current sensing output is applied to the metering circuit (M1) as mentioned, and the current limit detector based around IC4A. The best way to think of this stage is as a simple comparator, where a preset level is applied to the non-inverting input from RV7 (Current) and our voltage representing the supply's output current is applied to the inverting input — if the latter exceeds the former, the op-amp's output will swing low.

With the voltages as shown on the circuit, the Current pot (RV7) is set to detect a level corresponding to 2A (2V) and the current detector is generating a 1V level in response to the supply's output current of 1A. In this situation, the output of IC4A is at a high level (around 11V), Q8 is biased off via R33, and the limiting indicator LED2 is off. And most importantly, the high level at pin 7 of IC4A is passed to the Voltage pot RV1, via isolating diode D1 — providing the supply's overall reference voltage (note the connection at point 'A' on the schematic).

If the supply's output current rises above 2A however, the level presented at IC4A's inverting input will exceed that at its non- inverting input (as set by RV7), and its output will swing low. This then biases Q8 on via R33, which in turn activates LED2 via limiting resistor R34, indicating that current limiting is taking place. As you expect, the fall at the output of IC4A will also cause a similar drop in voltage across RV1, which reduces the regulator circuit's source voltage and causes the supply's output level to fall.

This in turn reduces the current flowing into the load and the voltage level applied to IC4A's inverting input, which halts the downward trend at the op-amp's output, and so on. Thus the whole feedback process causes the output voltage to balance at a point where our preset current (set by RV7) is flowing into the load.

Despite this seemingly convoluted process, the comparator stabilises quickly and smoothly, as IC4A has its gain set to 100 by R37 and R35, and the transient response is tailored by feedback components C16 and R36.

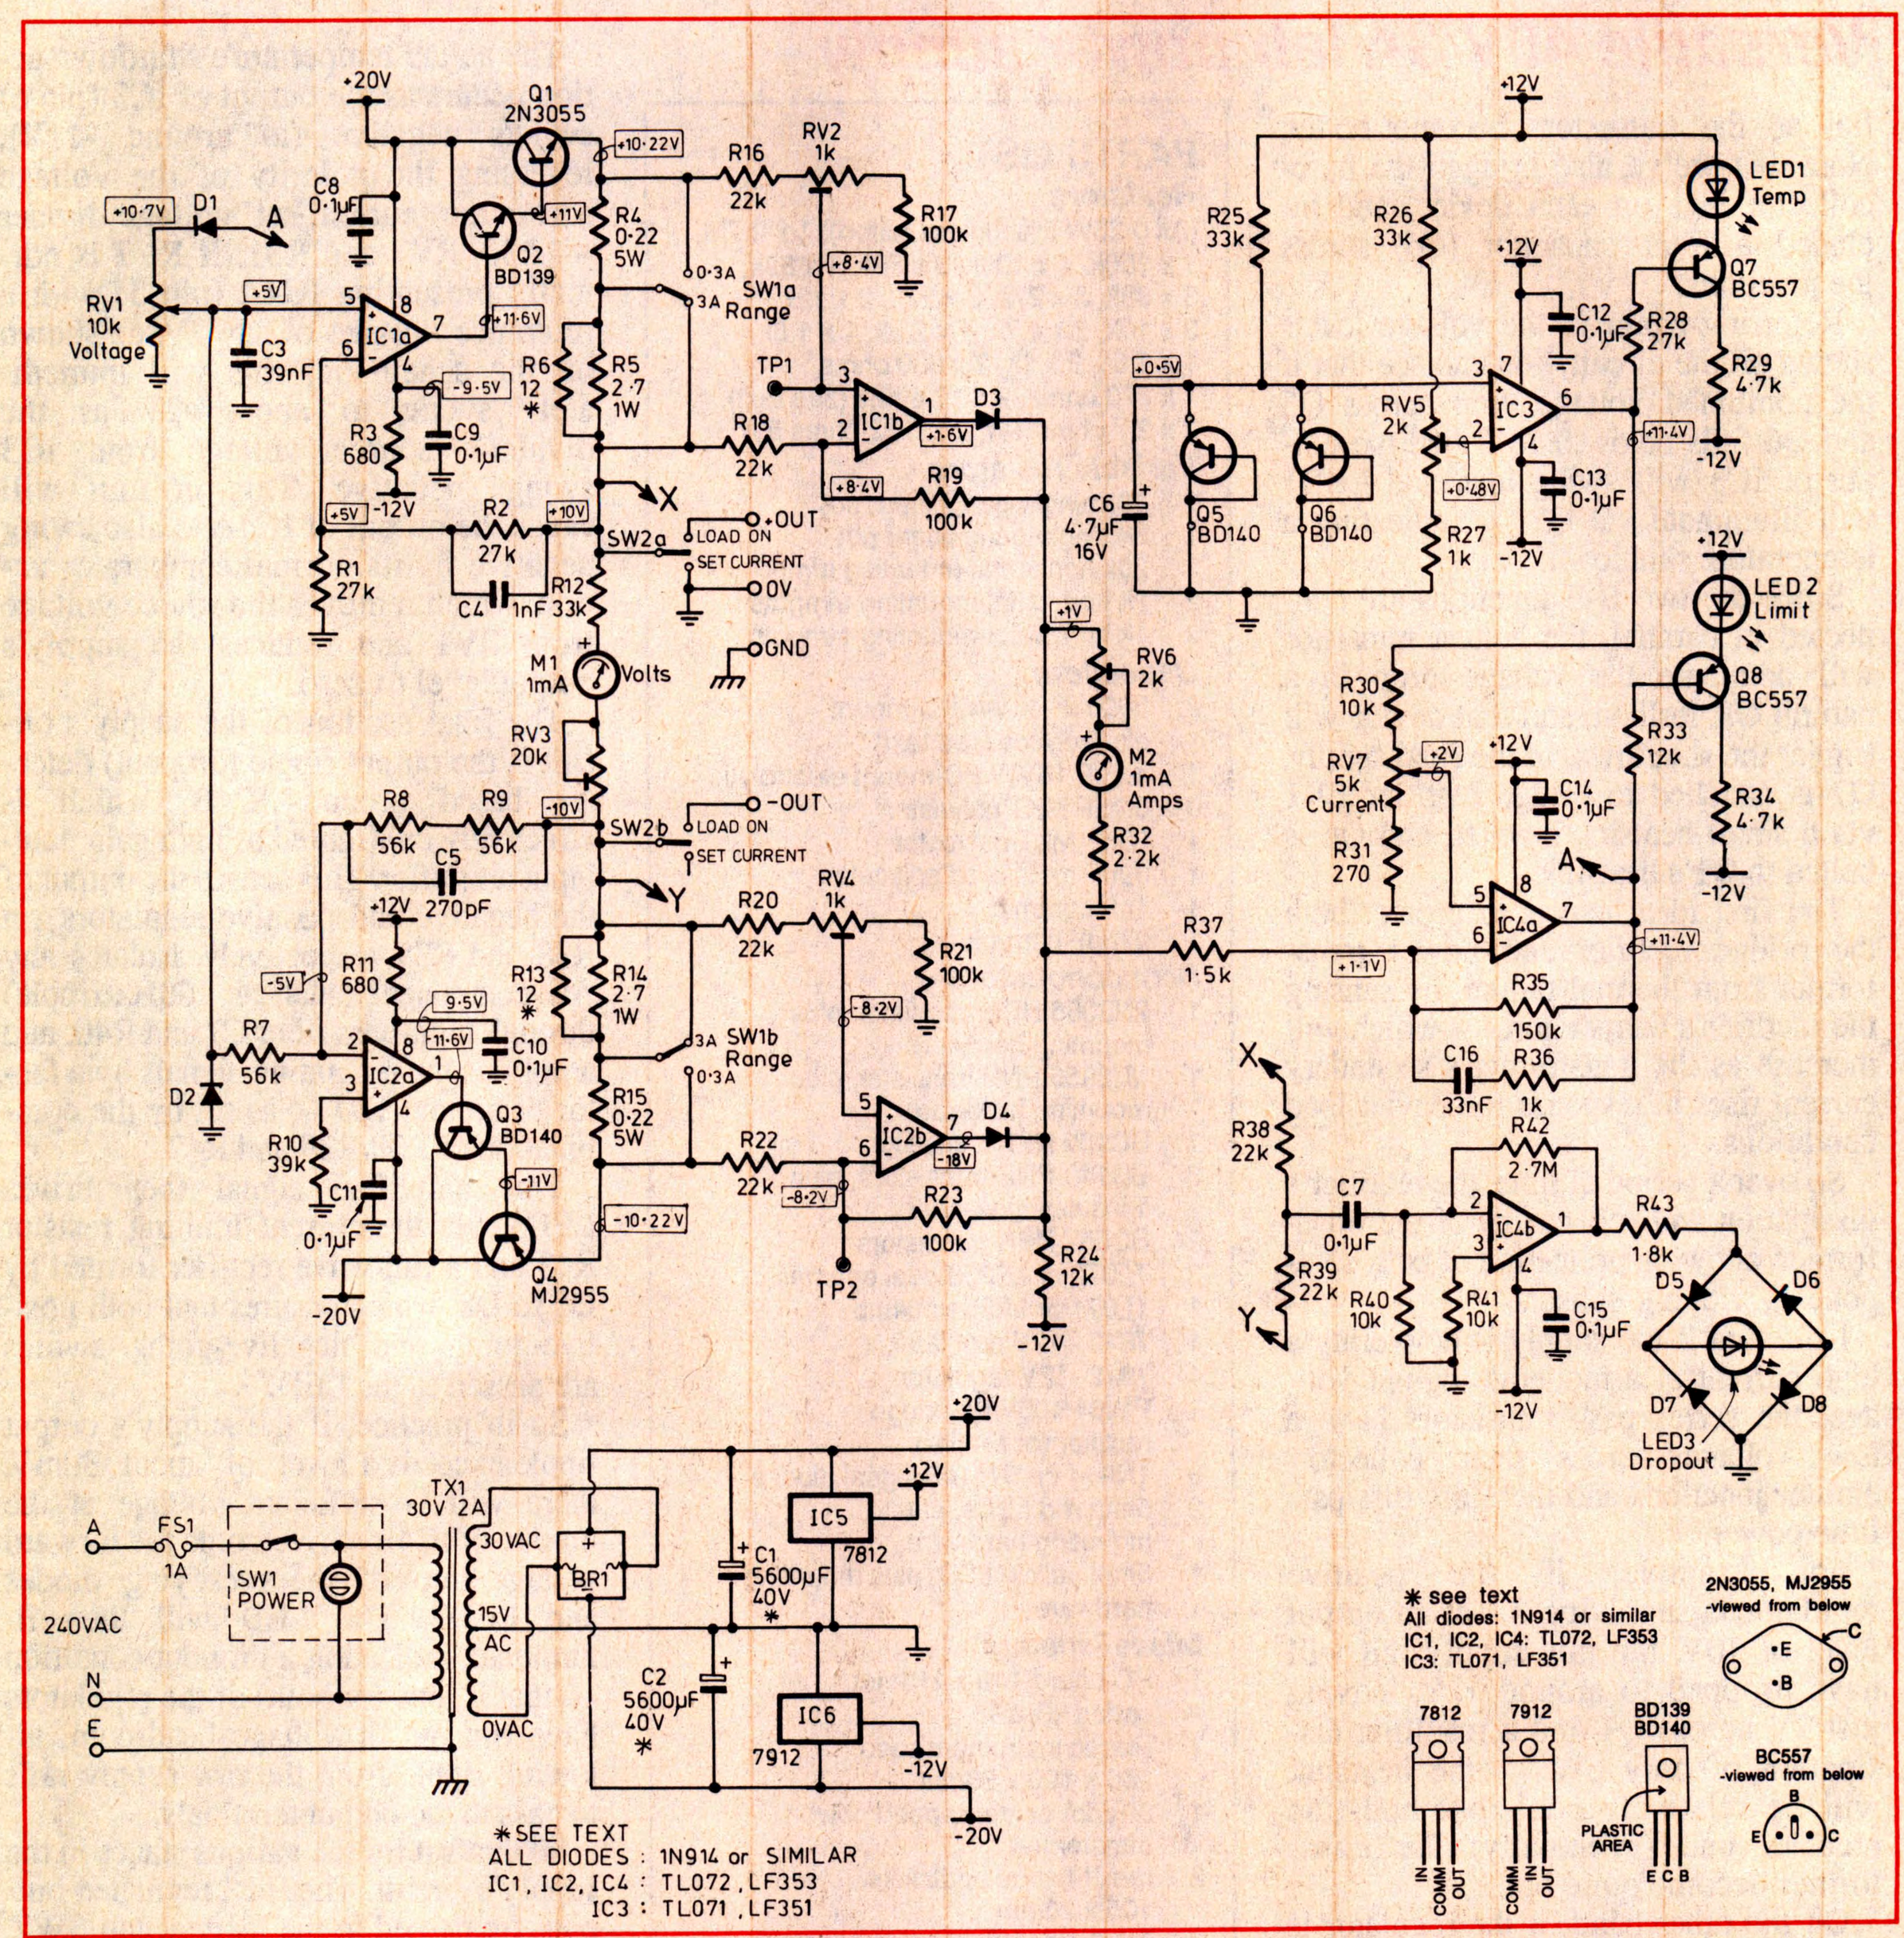

The supply's overall schematic diagram. The voltage readings shown on the circuit correspond to an output voltage of +/-10V (20V overall), and a load current of 1A.

Any voltage changes at the wiper of RV1 also effect the *negative* regulator circuit, as you'd expect. This circuit is based around IC2A, and acts as an *inverting* amplifier with a gain of two (rather than a non-inverting stage as for the positive regulator circuit). So for a source voltage of 5V, the negative regulator will produce an output of -10V.

In this circuit the gain is set by the combination of resistors R7 to R9, the bandwidth restricted by capacitor C5, and the output current capability boosted by PNP transistors Q3 and Q4. Other than that, R11 restricts the supply voltage as before, C10 and C11 bypass the IC2A's supply pins, and R10 locks the non-inverting input (pin 3) to ground. Finally, D2 has been included to prevent the source voltage swinging below -0.6V — which won't occur in normal circumstances, by the way.

The final parts of the supply's circuit concern the output ripple detector based on IC4B, and the temperature cutout stage formed around IC3. The temperature sensing circuit uses IC3 as a conventional comparator, where a reference voltage is applied to the inverting input (pin 2) via a voltage divider formed by R26, R27 and RV5, and the inverting input (pin 3) monitors the voltage across transistors Q5 and Q6.

These PNP flat-pack power transistors (BD140's) are effectively connected as diodes, where the current supplied by R25 flows through their base-emitter junctions to ground.

And as you would expect from the natural temperature response of the silicon junctions, the voltage drop across the base-emitter connections (Vbe) will

fall as the transistors become hotter. Note that we've also bridged the base-collector legs on each device, and included a filter capacitor (C6) across the junctions.

Referring again to the voltage levels shown on the circuit, you can see that if the combined voltage drop across Q5 and Q6 falls below 0.48V then the output of IC3 will swing low, and LED1 will illuminate via Q7 — indicating a temperature shutdown.

Since the two B-E junctions are connected in parallel, the hottest transistor will dominate the voltage drop in a natural OR fashion, and that device will trigger the shutdown. In the final supply, Q7 is attached to bridge rectifier BR1 via a small heatsink bracket, and Q6 is bolted to Q1's heatsink.

The first idea here is to protect both the bridge rectifier and power transformer from thermal failure, by sensing the rectifier's temperature — which will increase as the transformer's secondary current rises in response to heavier load conditions.

So over a period of time, the rectifier's dissipation will reflect that of the transformer, allowing protection of both components with one sensor.

Note that if the supply is delivering a high current at a moderate output voltage, the series pass transistors have a *low* voltage across their collector-emitter junctions and therefore dissipate little power.

For example, if the positive regulator is supplying 3A at an output level of 10V, the raw supply rail will have dropped to around 12V, leaving just 2V across Q1 and a resulting dissipation of only 6W — the heatsink will be relatively cool, but both the rectifier and (eventually) the transformer become quite hot.

On the other hand, if the regulator is delivering only half of that current (1.5A) into say a short circuit, the full supply rail (around 20V at this current) is impressed across Q1, forcing it to dissipate a much higher power level of about 30W — in this case, the rectifier and transformer remain relatively cool, but the heatsink becomes hot.

So as you've no doubt gathered, the second temperature sensor attached to Q1's heatsink is intended to protect the series-pass transistors against thermal failure. While the temperature of the negative regulator's pass transistor (Q4) is not sensed directly, its dissipation will be very close to that of Q1 in virtually all circumstances.

PARTS LIST

Resistors

(All 0.25W 5%, unless noted): 1 x 2.7M, 1 x 150k, 4 x 100k, 3 x 56k, 1 x 39k, 3 x 33k, 3 x 27k, 6 x 22k, 2 x 12k, 3 x 10k, 2 x 4.7k, 1 x 2.2k, 1 x 1.8k, 1 x 1.5k, 2 x 1k, 2 x 680 ohms, 1 x 270 ohms, 2 x 12 ohms (see text), 2 x 2.7 ohms 1W, 2 x 0.22 ohms 5W

Variable resistors

1 10k single-gang linear pot
1 5k single-gang linear pot
1 20k horizontal-mounting trimpot
2 2k horizontal-mounting trimpots
2 1k horizontal-mounting trimpots

Capacitors

2 5600uF 40VW PC-mount electrolytics (see text)
1 4.7uF 16VW PC-mount electrolytic
9 0.1uF MKT polyester
1 33nF MKT polyester
1 39nF metallised polyester
1 1nF ceramic
1 270pF ceramic

Semiconductors

1 2N3055 NPN transistor, plus mounting hardware
1 MJ2955 PNP transistor, plus mounting hardware
1 BC139 NPN transistor
3 BD140 PNP transistors, plus 1 insulating washer
2 BC557 NPN transistors
3 TL072 or LF353 dual opamps
1 TL071 or LF351 opamp
1 7812 +12V regulator
1 7912 -12V regulator
1 PW04 400V/6A bridge rectifier (or similar)
8 1N914 or 1N4148 signal diodes
2 5mm red LEDs, plus mounting hardware
1 5mm yellow LED, plus mounting hardware

Miscellaneous

1 PC board 149 x 113mm, coded 93ps9
1 Plastic instrument case, 260 x 190 x 80mm
1 30V/2A center-tapped power transformer
2 radial-finned heatsinks, 105 x 70mm
2 TO-3 plastic insulating caps
2 MU-45 type meters, 1mA FSD
1 DPDT miniature toggle switch
1 DPDT miniature toggle switch, centre off, momentary action on one side.
2 Plastic knobs
4 Banana-type binding posts (yellow, black, red & green)
1 miniature rocker power switch, with internal neon
1 Panel mounting IEC-type mains plug (male)
1 3AG panel mounting fuse holder, with 1A fuse

Solder lugs, heavy-duty hookup wire, mains-rated hookup wire, rainbow cable, PCB pins, aluminium plate, heatshrink or cambric tubing, nuts, bolts, lockwashers, etc.

The actual temperature shutdown action occurs as the output of IC3 (pin 6) swings negative (to around -11V), reversing the polarity of the voltage across the 'current set' voltage divider — R30, RV7 and R31. If RV7 is normally presenting IC4A (pin 5) with a reference voltage of say 2V, as shown on the schematic, this will immediately swing to about -2V as the divider's source voltage from IC3 swings negative. This in turn will cause the output of IC4A to also swing negative by its normal comparator action, which removes the source voltage from RV1 and reduces the supply's output level to zero.

The final section of the supply's circuit is the output ripple (dropout) detector based around IC4B, which is effectively configured as high gain dual-input amplifier. This senses the output of the negative and positive regulators via R38 and R39 respectively, isolates any AC component (such as 100Hz ripple) through high-pass filter C7 and R40, and amplifies the remaining signals by a factor of around 120 — as set by the combination of R38/R39 and R42.

This amplified signal then drives LED3, via the current limiting resistor R43 and a full-wave rectifier formed by D5 to D8, which ensures that both positive-going and negative-going swings are passed to the LED.

So in practice, if the supply's output ripple exceeds a level of about 30mV, there will be sufficient voltage at the output of IC4B to overcome the forward voltage drop of the rectifying diodes and LED3. The LED will then illuminate, indicating a dropout condition — that is, one (or both) of the regulators have run out of voltage headroom, allowing ripple from the raw supply rails to pass to the output terminals.

So that's it for the various stages in the supply's circuit. The final regulated outputs are passed to the load switch SW2, which isolates the load in the 'set volts' position, connects the load when 'load on' is selected, and shorts the positive output to 0V in the 'set current' position so the maximum output current can be set in advance.

The voltage meter (M1) and its associated resistors (R12 and RV3) are tied between the supply's positive and negative outputs, and will therefore respond to changes in the unit's overall output voltage.

Unfortunately, that's all the space we have available in this issue. In the next and final part, we'll look at constructing and setting up the lab supply.

(To be continued) ❖