ELECTRONIC JUMPER LEADS

What a great idea, eh? . . . Don't get your hands dirty next time your car has a flat battery. Charge it electronically through the cigarette lighter!

Ian Thomas

EVER HAD THAT fun experience of coming out to your car after a great evening to find you left the headlights on and the battery's dead flat? Usually you're all done up in your tux (or dinner stubbies according to taste) and you have to rummage around in the boot among all that unspeakably grotty junk to find your jumper leads. Then you have to fumble around under the bonnet of both your car and your friend's to make the connections to the batteries. Many blinding sparks fly and you nearly melt the jumper leads into a molten puddle when you inadvertently get the polarity wrong while your companions nearly fracture themselves laughing. Great fun, isn't it.

The easier way is to recharge the battery from a good one electronically. All cars have a cigarette lighter which has become the de facto power outlet in recent, more health conscious times. These sockets are good for about 10 amps only. If a simple connection were to be made between the lighter outlets of two cars, spectacular but not useful events would occur. If the battery of one car were dead flat and the other OK, all that would happen would be a blown fuse(s). However, if a little electronics were to be interposed between the two outlets much more productive results could be created.

A regulator

What is necessary is some sort of voltage regulator in the lead connecting the sockets. If the car to be used as a source has a good battery and, even better, its engine is running, then the voltage available is about 13 to 14 volts. If the flat battery is dead flat then the load to be charged into is close to zero volts. As the flat battery charges, its voltage will rise to 15 or 16 volts or a higher voltage than the good battery. Resistive losses in the leads require a regulator to be able to step up as well as down. So, what is needed is a regulator that is able to draw power from the good battery and deliver a constant current to the flat one regardless of the output voltage.

Another problem that has to be borne in mind is the fact that quite considerable amounts of power have to be handled. In the case of our dead flat battery, let's consider what would happen if a simple series pass regulator were used. If a charging current of 10 amps were used then the power that the regulator would have to dissipate at the start of the charge would be 13 x 10 or 130 watts. This sort of power makes heatsinks *very* hot and also wastes an awful lot of power.

The only answer to the power dissipation problem would be some sort of switching regulator. A conventional voltage regulator works by acting as an electronic variable resistance between the source and load to keep the voltage of the load constant. A switching regulator works on an entirely different principle. Voltage control is achieved by switching the source on and off very rapidly to control the output voltage and filtering the output square wave to obtain a (reasonably) smooth output voltage. As the main power handling device is either on like a switch, and hence has no voltage drop across it, or off, and has no current flowing through it, then it doesn't have to dissipate

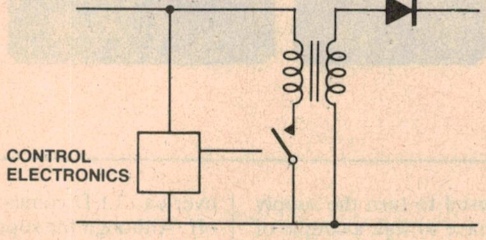

Figure 1. Simplified circuit.

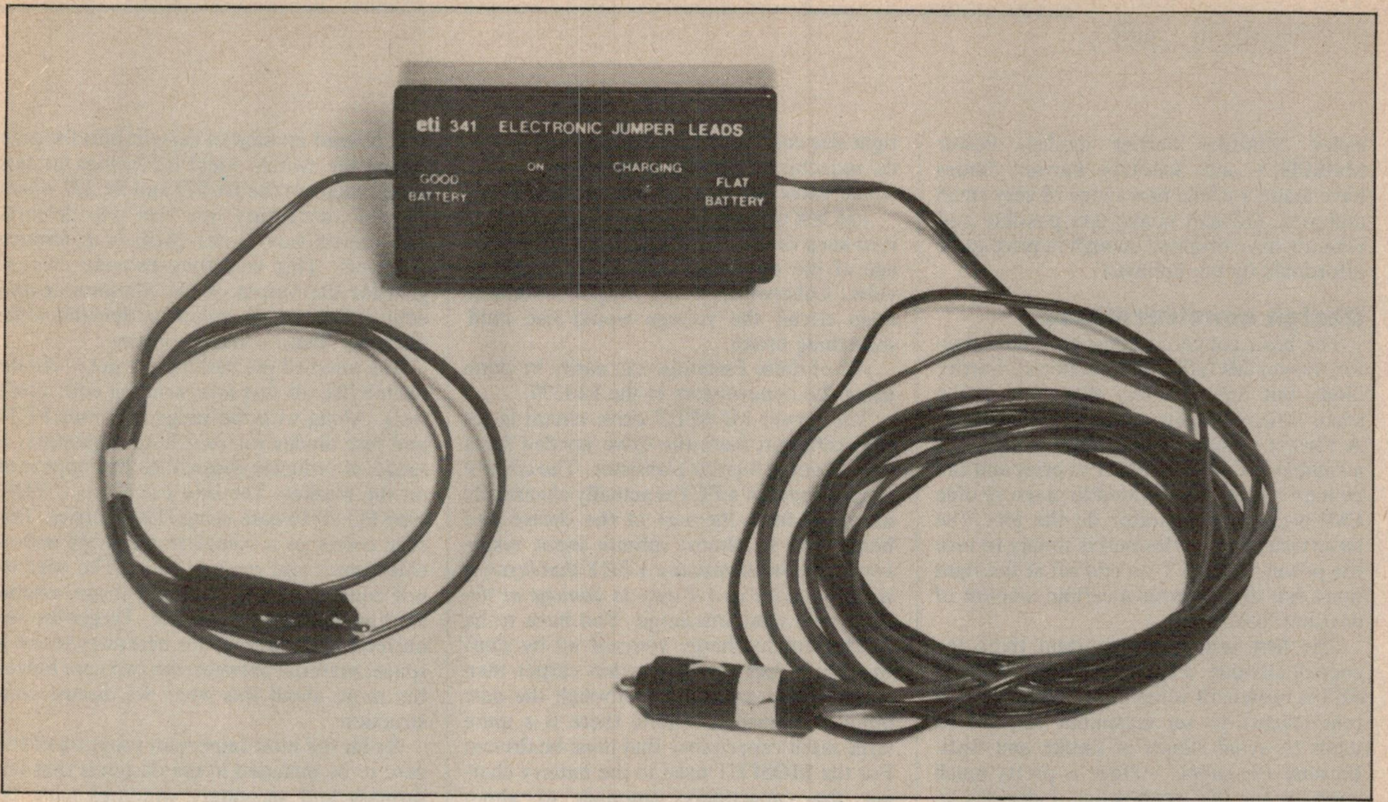

much power.

This is all fine if the output voltage is less than the input voltage, but in our case the output must be either greater or less than the input. Fortunately, switching regulators come in many flavours and by using the filter inductor in various ways almost any desired characteristic can be obtained.

One of the most versatile forms of switching regulator is the flyback or ringing choke converter, which has no direct connection between the input source power and the output circuit apart from the main transformer/inductor. The converter operates by applying the source voltage to the inductor through the main switching transistor (see "How it works" section). This allows current to build up in the coil primary and hence stores energy in the inductor equal to $\frac{1}{2}LI^2$. When sufficient current has built up in the primary winding, the main transistor is switched off. One of the properties of an inductor is that it will not allow the current through it to change instantaneously. To do so requires infinitely large voltages. This is the exact dual of a capacitor where the voltage across a capacitor cannot be changed instantly — to do so requires infinite current (try shorting out a 10,000 μF electrolytic and you'll see what I mean!). If the coil primary is switched off then the current flowing through it is immediately transferred to the secondary and, providing a current path is available to it, will flow there.

The handy part about this is that we're only interested in currents — the voltage across the primary and secondary doesn't matter. If the voltage across the secondary is high then the current will decay very quickly and if the secondary voltage is low then the current takes a long time to decay, but power is transferred to the output circuit regardless of the output voltage. This type of circuit would seem to be ideal for the battery charger and is the one I used. In practice, the flyback converter is more critical in the design of the inductor than the more conventional down converter but its versatility often makes the design worthwhile.

Controlling the current

Some quick sums showed that if reasonable charging currents were to be achieved a few tricks were going to be needed if the charger was to be small. The basic problem with the flyback converter is that all the power must first be stored in the inductor before it can be transferred to the output. This makes for a big lumpy inductor. To minimize this, the first trick is to never allow the coil current to decay to zero but turn the transistor on again after only part of the coil energy has been removed. This helps because the energy stored in the coil is proportional to the current squared, but the current rises in the coil linearly with time. Suppose we had a monster coil with a one henry inductance and applied one volt across it for one second. After the second, half a joule of energy would be stored in it. If the current were allowed to run up for another second then 2 amps would be flowing *and the energy stored would be 2 joules*. In other words, three times as much energy is stored in the inductor during the second second as was stored during the first. Clearly it's smarter to delete the first second altogether. This is quite easy to do as all that's needed is to sense the secondary current and turn the main switching transistor on again after it drops to a preset level.

The next trick is to use an inductance of only a very few microhenries. This is fine (and obvious) and makes for a small inductor physically, but it means that when the input voltage is switched across the coil its current rises very fast and hence the main switching transistor on time becomes very short if the primary current isn't to become excessive. It has been implicit in all that I've said so far that the transistor can be switched on and off instantly. Sad to say but it can't. An average power transistor takes about a microsecond (less for fancy ones — maybe 0.2 to 0.3 microseconds) and when the switching transistor is neither on nor off it has volts across it and current flowing through. This means that when the transistor is switching it has to dissipate power. If you try to switch it on and off too fast then it spends most of its life halfway and gets very hot!

Fortunately, in recent years a new breed of transistor has appeared on the market called the power MOSFET. These devices have the wonderful property that they switch very, very fast. Instead of taking a microsecond to switch they can, with correct drive, switch in about 50 nanoseconds or so, allowing much smaller inductors to be used while still getting good converter efficiency.

If the main switching transistor is blindingly fast then the output diode in the coil secondary must be fast to match. Once again devices are available, in this case ▶

called Schottky barrier diodes. Power MOSFETS and Schottky barrier diodes have made a whole new range of very small and very efficient converters possible and recently have dropped enough in price to be affordable by the hobbyist.

Design considerations

The basic converter consists of an inductor, power MOSFET and Schottky barrier diode but unfortunately there has to be some control circuitry to make them work. A variety of special integrated circuits are available to do this but they're pricy and can be hard to get, so I decided to make ye olde LM339 quad comparator do the job. The basic task to be performed is simply to turn the power MOSFET on and off at the right times but this requires a certain amount of decision making.

The first and most important thing the control circuitry must do is ensure that the FET is turned off when the coil primary current reaches its set maximum value. This tends to avoid clouds of smoke and such. Because the supply voltage is pretty much constant for this application, I decided to sense the actual voltage across the FET to determine the right turn off time. Power FET characteristics are fairly tightly controlled (that's part of what you pay the dollars for) so for a given control voltage it's fairly certain that for a given drain-source voltage the required current is flowing. Normally a sense resistor would be inserted between the source of the FET and ground to sense FET current, but this has undesirable side effects such as unwanted inductance which degrades switching time.

Another decision is when to turn the FET on again. In this case a sense resistor is necessary. It must be in the output winding circuit and must not cause excessive power loss. As peak current in the secondary (and primary) is about 15 amps, the maximum resistance is of the order of tens of milliohms before the resistor starts to get hot. It is not easy to make a resistor this small and you can't buy one off the shelf. I simply used a short length of copper wire of the right size carefully twisted so as to minimize its inductance (see "Construction"). The actual value finally used was 13 milliohms.

The last control function to be performed is to keep the FET off when the output voltage of the converter rises above about 17 volts. Otherwise, if the output were left open circuit the voltage would rise until something broke.

All of these functions can easily be done using the comparators in the LM339.

The power MOSFET drive circuit is totally different from the drive needed for a conventional bipolar transistor. The control or gate input of a FET is actually electrically insulated from the rest of the device and hence has an almost infinite input resistance. *This insulation is a VERY thin layer of silicon dioxide and is easy to damage or destroy with static discharge.* You have to be careful not to charge yourself up by shuffling your feet across a nylon carpet then touching the gate lead. Although the gate input resistance is infinite there is a quite large input capacitance that must be driven. For the MOSFET used in the battery charger the capacitance amounts to about 1.2 nF but this is magnified during switching by miller capacitance from the drain to gate. To turn the FET on, the gate must be taken positive with respect to the source and the basic principle is the more positive the better (up to the maximum V_{gs} of 20 volts). Above about 12 volts the on resistance of the FET is 0.12 ohms and if the gate is taken to +20 volts this can be reduced to about 0.1 ohms. Since the dc supply voltage is 13.6 volts it doesn't seem worthwhile going to the trouble of generating an extra high voltage just to gain the extra little bit of efficiency.

The actual drive that needs to be applied to the gate is a square wave generated by the control circuitry, which needs to deliver almost no steady bias current but *must* be able to give very high current spikes during switching to charge and discharge the gate capacitance. The total gate power needed is very low so small T092 plastic transistors can be used but they must be able to provide at least an amp of current into the gate during the switch. Ideal transistors for this application are the BC327 and BC337 which are complementary npn and pnp devices. The easiest way to use them is as emitter followers. This way they provide current gain for the output of the comparator and deliver the current spikes on demand without any transistor bias problems.

The heart of the battery charger is the inductor (we all just *love* winding coils don't we!). Along with the range of power FETs and fast diodes for switching regulators, a range of inductor assemblies has appeared on the market. The one I chose is a TDK type PQ 20/16 core using H7C1 ferrite. The core comes as a complete assembly with a coil former and retaining clip. The core itself consists of two halves that are assembled around the coil former. To get the required coil inductance it's necessary to use a spacer material between the two core halves but more about this when we discuss construction.

By far the most important property of the core to be included in the design is that the primary and secondary windings must be tightly coupled together. If the two windings are simply wound one on top of the other on the former, the coil will show an excessive leakage inductance. This means that part of the energy stored in the coil primary will not be transferred to the secondary and must be dissipated by the power FET. To ensure the tightest possible coupling the primary and secondary windings are actually twisted together and then both wound on the core as one.

As the converter is switching at quite high frequencies (about 80 kHz) yet another parasitic effect must be allowed for in the coil. It's been known for a long time that if ac current is required to flow along a piece of wire then it tends to flow only along the outer surface of the wire. This is called the skin effect and has the effect of increasing the resistance of a piece of wire (ie a coil winding) and the higher the frequency of the current the higher the resistance. For copper it works out that for 100 kHz a 0.5 mm diameter wire has its resistance increased by about 12% so this seems to be the thickest wire that can be used if all the copper is to do work. In order to make up the total winding, many strands of 0.5 mm wire are twisted together then wound into the coil. To make the twisted wire strands lie neatly together I used seven lengths of wire in a neat hexagonal arrangement. This wasn't as hard as it sounds and resulted in a very neat inductor (it worked well too!).

One last remark must be made about the electrolytic filter capacitors on both the input and output sides. The converter is switching high currents and, because of the long leads into and out of the converter, most of the ac component of this current must be carried by the filter electrolytics.

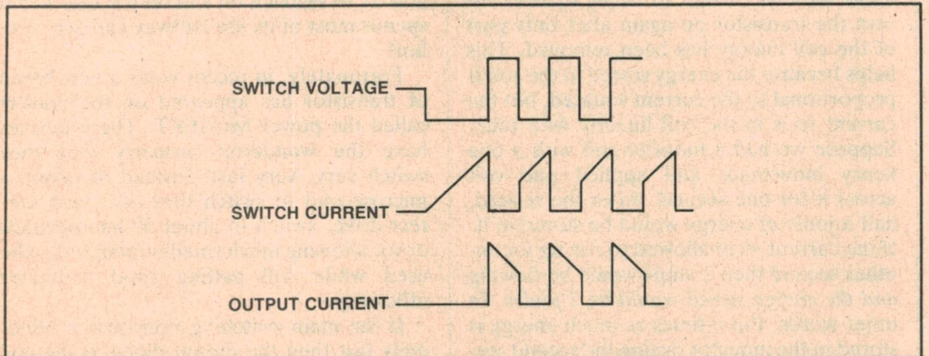

Figure 2. Circuit waveforms.

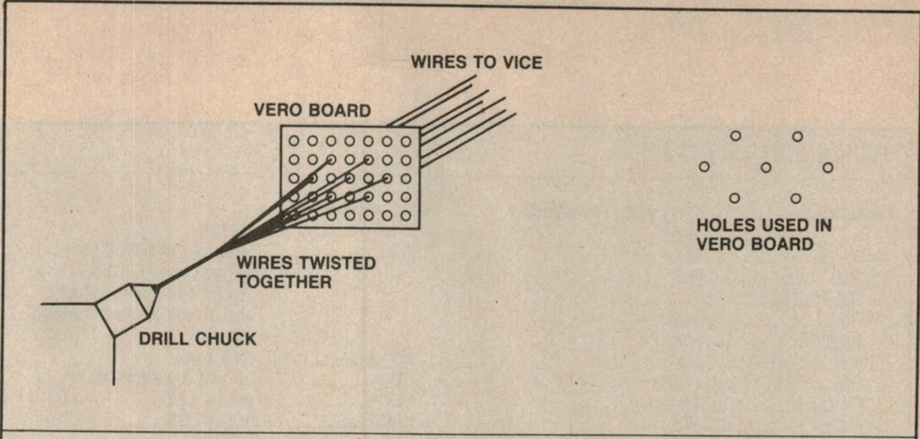

Figure 3. Wires are fed through the Veroboard, twisted tightly and placed in the drill chuck.

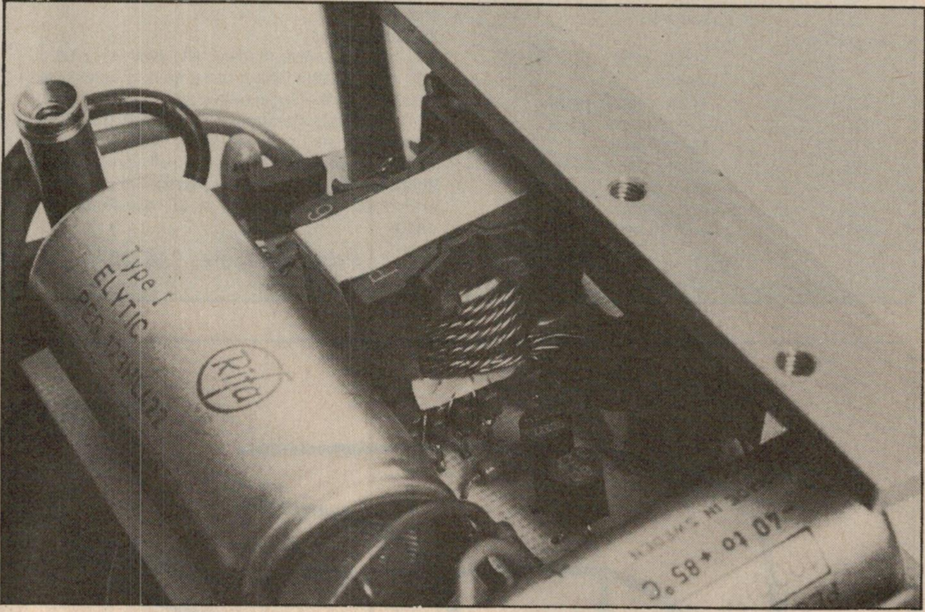

Close up view of mounted coil.

Fortunately the converter doesn't have to run reliably for 10,000 hours (you don't flatten your battery *that* often do you?), but if the electrolytics don't have a sufficient low impedance it may upset the operation of the converter. A good range of electrolytics is San-hwa. In the prototype I used Rifa capacitors. As a test in the prototype I purchased some samples from a certain aeronautically oriented gentleman's establishment which were brand named REC. The capacitors were beaut and small but their series resistance was far too high for switching regulators and they didn't work!

Construction

By far the trickiest part of the project is the winding and testing of the coil so let's get that one over with first. The winding consists of 10½ turns of seven strands of 0.5 mm wire twisted together. If you look closely at the photo (right) you'll see that the strands are twisted together to make a neat cable. The total length of this cable needs to be about 600 mm long so to be sure you need to cut seven lengths of wire about 700 mm long. To twist them together you'll need an eggbeater drill or variable speed electric drill and a small piece of Vero board with holes drilled in it on 0.1 inch centres. Feed all seven wires through holes in the piece of Vero board as shown in Figure 3. Notice that they form a neat hexagonal pattern with the seventh wire in the centre. Twist all seven wires together tightly at one end for about 10 mm and place the twisted end in the chuck of the drill. Next make *absolutely sure* that all the wires are free and untangled and have no kinks in them. Stretch all the wires out straight and even, and clamp their free end in a bench vice.

Now comes the fun bit. Pull all the wires tight with the drill chuck and rotate the chuck slowly so the wires twist together. Have a helper hold the piece of Vero board so it's about 15 mm from where all the wires are lying together and move the Vero board along the wires as they twist up to maintain the constant 15 mm spacing. This should result in a beautiful neat cable being formed. The centre wire will become loose as the other wires are wrapped around it, though this didn't give me any trouble. The end result of this exercise should be about 650 mm of wire ready to be wound onto the former.

The first step is to trim about 20 mm off the end of the wire, which is probably a bit ratty, and then spread out 15 mm of the end of the cable. Tin about 5 mm of the end of *four* of the seven wires. You'll see that the coil former has two groups of four pins on one side and two groups of three on the other. Stand the former on its pins with the two groups of four facing you. The four tinned leads go onto the pins on the *right.* This is the start of the primary and is connected to the drain of the MOSFET on the circuit board. Check what you're doing against the board layout to make sure you've got it right. The last three wires are tinned about 15 mm in from the first four and are the start of the secondary. If you look at the top of the former with the pins facing away from you the coil is wound *anti-clockwise* so you'll wind on about a quarter turn of the four strand primary before the three strand secondary joins the cable. The start of the secondary is connected to the power diode so once again you can check that you're doing it right by comparing the coil former to the board layout.

The next bit is easy. Wind the cable around the former until you've got one full layer. This should take about five to five and a half turns. Then wind on a second layer to give the full 10 turns. Finish the winding next to the unused three pin group *but don't cut off the wires.* Untwist the cable for about 40 mm from where the winding ends next to the three unused pins and tin the ends of all seven wires. This is so you can check for continuity and find which wire should connect where. Using a multimeter on the ohms setting, test for continuity between the group of three pins and the end of the wires. Terminate the end of the secondary three wires on the unused three pins then wind the other four wires another quarter turn and terminate them on the unused group of four pins. You should now have a complete wound coil with all wires terminated.

As a last check take four pieces of tinned wire (resistor leads that have been cut off are fine) and tack together the four wire start of the primary, the four wire end of the primary, the three wire start of the secondary and the three wire end of the secondary. Now check that there is no continuity between the primary and secondary windings. If there is, you've made a mistake and will have to disconnect the end of both the primary and secondary and try again.

Next you must assemble the ferrite cores around the coil. As previously mentioned, the cores must be spaced slightly apart to ▶

PARTS LIST — ETI-341

Resistors	all ¼ watt, 5% unless noted
R1, 6, 7, 11	2k7
R2	4k3
R3, 15, 18	1k
R4	330k
R5, 20	10k
R8	1M
R9	470R
R10	47k
R12	2k2
R13	680R
R14, 17	120k
R16	56k
R19	15k
R21	100R
R22	see text
Capacitors	
C1, 2	1000µ (1mF) 16 V A1 electro
C3, 6	100p ceramic plate
C4, 5	470p ceramic plate
C7, 11	1n met poly
C8	2200µ (2m2) 25 V A1 electro
C9	100n ceramic monolythic
C10	4µ7 35 volt tag tant

Semiconductors	
D1, 3, 4	1N914
D2	Motorola MBR735 (any Schottky barrier diode in a T0220 pack with a 7 amp and 35 volt or greater rating will do)
ZD1	6V8 zener
ZD2	75 volt 1 watt zener
LED1	yellow LED
LED2	green LED
IC1	LM339N
Q1	Motorola MTP12N10 or MTP12N08
Q2	BC327
Q3	BC337
Inductors	see text

Miscellaneous

ETI-341 pc board; medium jiffy box (41 x 68 x 130 mm); 35 mm length radial fin heatsink; 2 car cigarette lighter plugs; heavy duty speaker cable; LED mounting grommets; 2 x T220 mounting washers, nuts and insulating washers; 4 x 25 mm tapped pc board stand-off spacers and screws; 2 small self-tapping screws or 6BA nuts and bolts; 75 mm length of 25 mm aluminium angle.

Price estimate: $44.00

get an air gap. The gap needs to be about 0.6 mm to give the right inductance. Rummage around and find some suitable material, remembering that *you can't use metal.* As an example, the cardboard used in manila folders is about 0.25 mm thick so a couple of thicknesses of this would be OK or you may choose to find something else. A vernier caliper or micrometer is a great help here! Cut the material out to the same shape as the two mating faces of the cores and assemble them with the spacer material in place. Clip the cores together using the mounting clip provided and the core's ready for testing.

To test the coil leave the four wires you've tacked on to test continuity in place and add a 1k resistor in series with one side of the primary. Tack a 1 microfarad capacitor in parallel with the winding. Beg, borrow or whatever an audio oscillator and connect the output to the oscillator to the 1k resistor. The earth end of the oscillator goes to the other end of the primary. Now monitor the voltage across the coil with an oscilloscope or wide band ac voltmeter. You should find a nice sharp resonance at

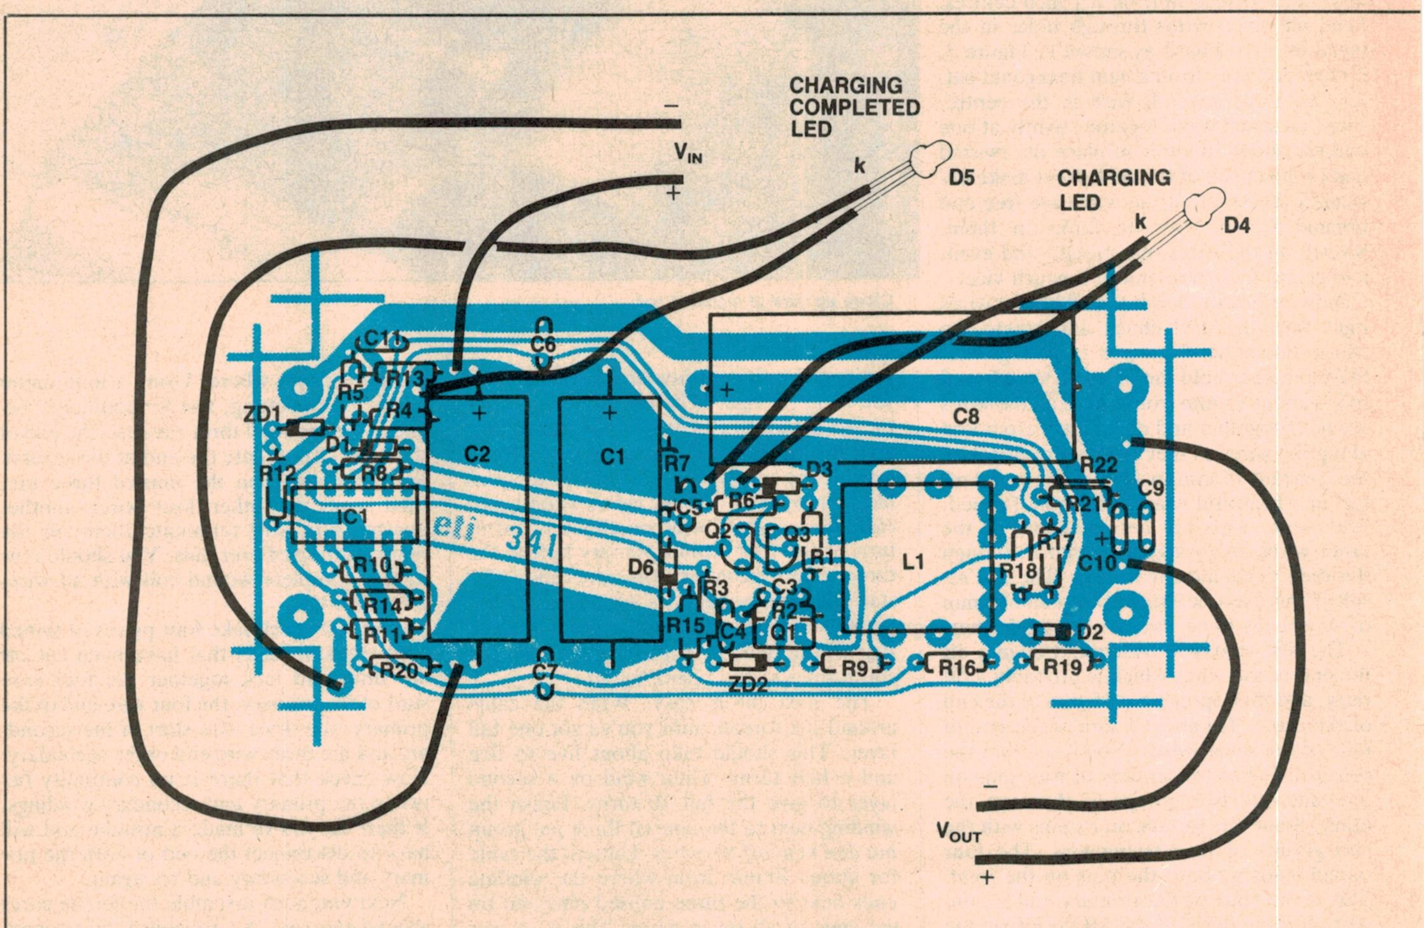

around 50 kHz. Anywhere between 30 kHz and 70 kHz is fine — the current sensing in the electronics will look after any discrepancies. To a large extent this test is optional provided the spacer is the right thickness; the electronics will look after any minor differences. That's the hard bit done.

The charger was built into a plastic box with a metal base bought from Dick Smith Electronics. I wanted it to be as small as possible so, after working out the component area needed, I opted for the 130 x 67 x 41 mm size (they seem to have about 800 different sizes!). If you're making your own artwork watch out for the board dimensions as the components use all the available space. There are also a few tricks used in the board layout to avoid unwanted effects so, all in all, it's probably better to use the

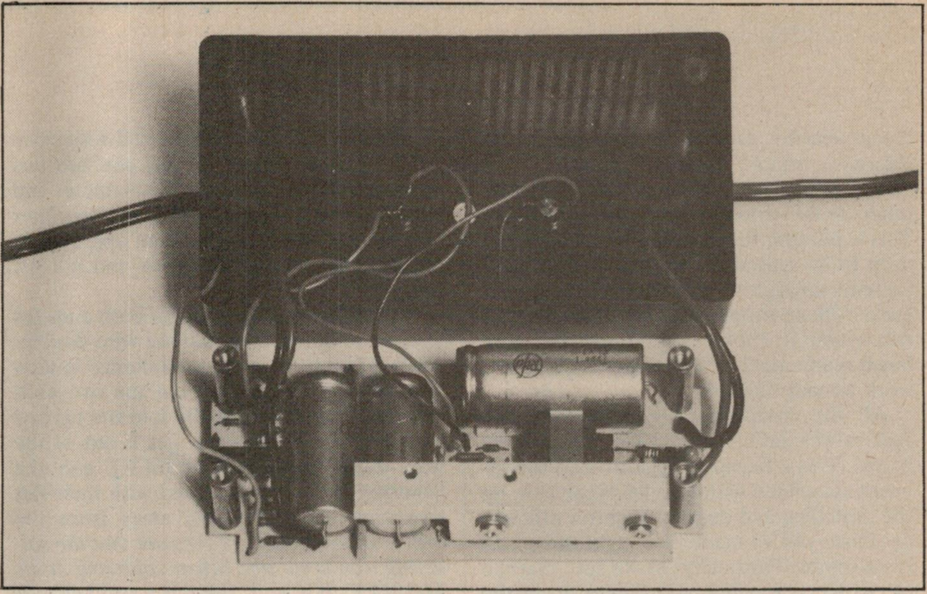

The assembled board is positioned exactly in the middle of the box lid.

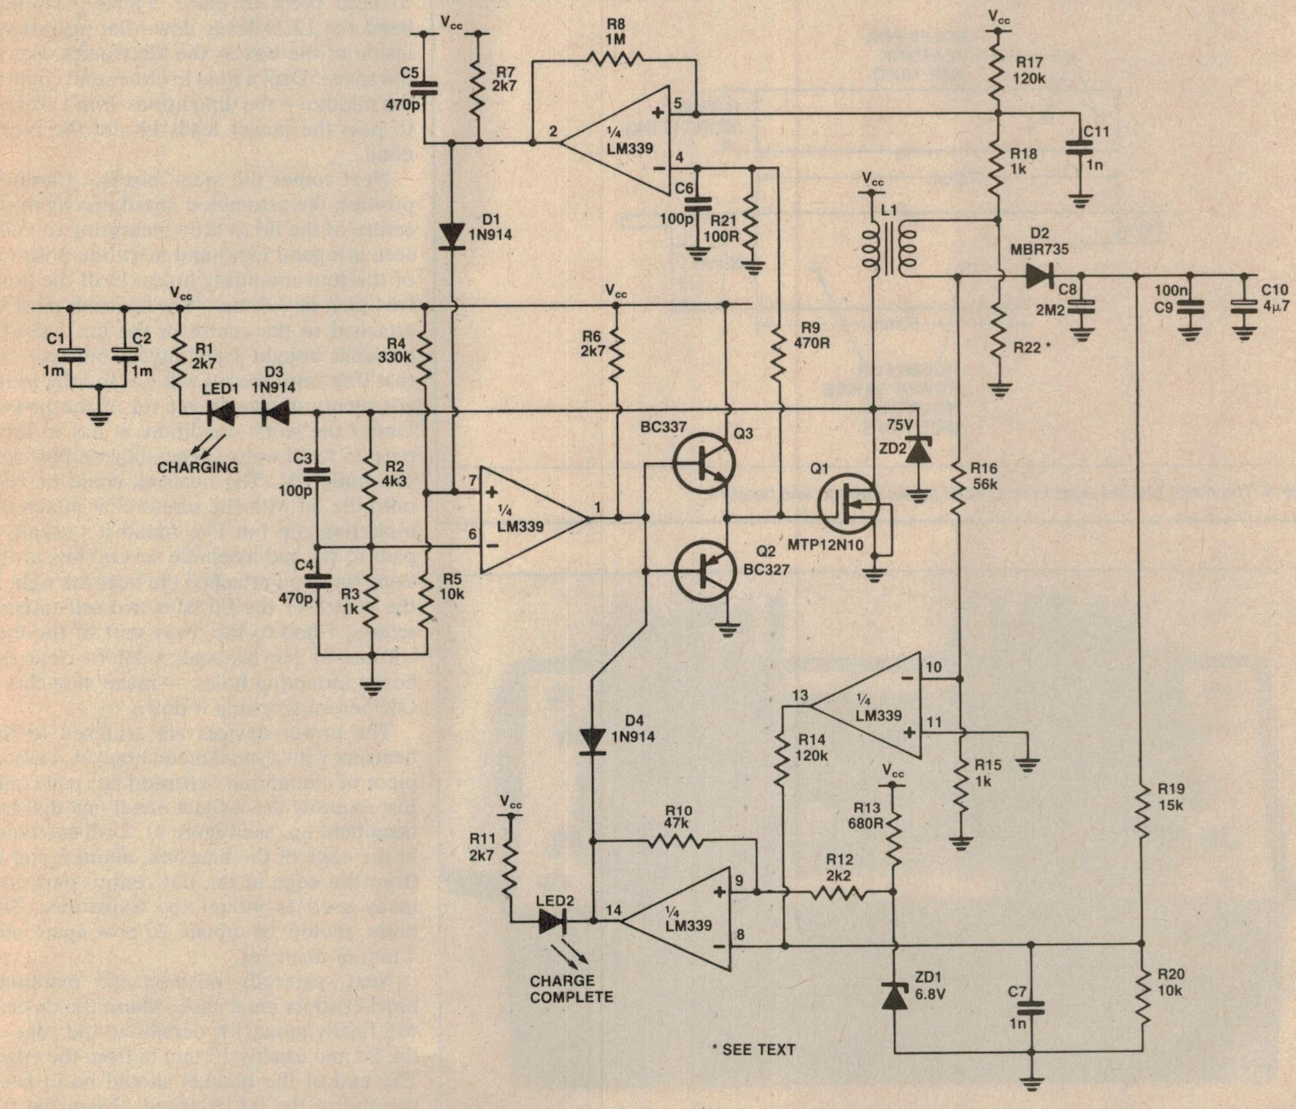

layout exactly as it is. For example, the reference 'ground' for both inputs for pins 4 and 5 of the comparator are next to each other (ie the ground ends of R21 and R22). This is because there is only a small voltage drop to be sensed across R22 and the high currents through the inverter can easily generate voltage drops along earth tracks that will swamp it. Probably the safest way is to buy a ready-made board or get the ETI artwork already checked.

All the components including the two high power devices mount directly onto the board. This is because the stray lead inductances associated with long leads can slow up the switching and degrade inverter efficiency. Once you've made or bought the board it's assembled normally. Make sure that the electrolytic capacitor and diodes go in the right way. R22 is not actually a resistor but a piece of copper wire 0.36 mm in diameter (27 AWG). To get the required resistance you'll need 80 mm of wire if you use this diameter but you can vary the diameter and vary the length to suit. Don't use too finer and shorter wire though, as an appreciable amount of power is dissipated and it'll get hot.

To mount the wire, first tin both ends for about 2 or 3 mm then fold the wire double. Twist it together tightly to minimize inductance effects then spread out the two ends and solder them in. Finally fold the twisted wire to make it neat and get it out of the way. Both the power MOSFET and the Schottky diode are mounted with their flat mounting surfaces facing away from the centre of the board. *Make sure that the soldering iron is earthed before soldering in the MOSFET.* If it isn't, there is a good chance you'll damage the gate insulation when you solder the gate lead — and the MOSFET is the most expensive device on the board! When both power devices are correctly placed their mounting surfaces should be in the same plane ready for attachment to the heatsink mounting bracket.

To prepare the box ready for the board drill two ¼-inch holes for the indicator LEDs and insert the plastic mounting clips. Cut four pieces of lead wire about 100 mm long (preferably of different colours so you don't have to peer into the box all the time to figure out which is which) and bare and twist the ends for about 5 mm. For the two LEDs the longer lead *must* be connected to the positive. Carefully clip off the *shorter* (cathode) leads about 7 mm from the base and solder your chosen coloured leads to them. Make life easy for yourself by writing down on a piece of paper what connects to what then repeat the process for the two anode leads. Insert the LEDs into the clips and push the mounting rings over the clips to hold them in place. Finally, carefully bend the LED leads down flat against the inside of the box as the electronics uses all the space. Drill a hole in either end (right in the middle — the dimensions aren't critical) to pass the power leads in and the box is done.

Next comes the metal box lid. Carefully position the assembled board *exactly* in the centre of the lid (a little measuring to make sure is a good idea) and mark the positions of the four mounting holes. Drill the holes 9/64 inch or 3.6 mm. The heatsink must be attached to the centre of the lid. I used a heatsink bought from Jaycar but any one that'll fit onto the lid will do, so long as it's got plenty of size to get rid of the power. Under the worst conditions it has to dissipate 15 to 20 watts (depending on how well you build it!). The heatsink could be held onto the lid with the screws that attach the mounting clip but I've found it's usually a pain to try and assemble several bits at the same time so I attached the heatsink right in the centre of the lid with two self tapping screws. I had to file away part of the four corners of the heatsink a bit to clear the board mounting holes — make sure this is OK before screwing it down.

The power devices are attached to the heatsink with a small metal bracket. I used a piece of aluminium extrusion but you could just as easily use a sheet metal one that has been bent up (see Figure 4). Drill two holes in the edge of the heatsink, about 5 mm in from the edge in the flat central part normally used to mount the transistors. The holes should be about 20 mm apart and 4 mm in diameter.

Next, carefully position the mounting bracket so its inner face (where the devices will finally mount) is parallel to the edge of the lid and exactly 10 mm in from the edge. The end of the bracket should be in from the end of the lid by about 13 mm but it's

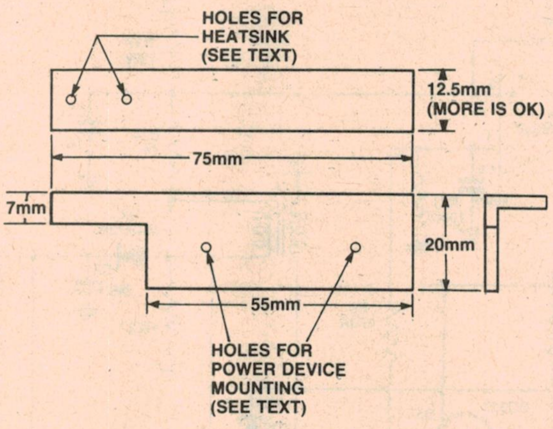

Figure 4. The metal bracket used to mount the power devices and heatsink.

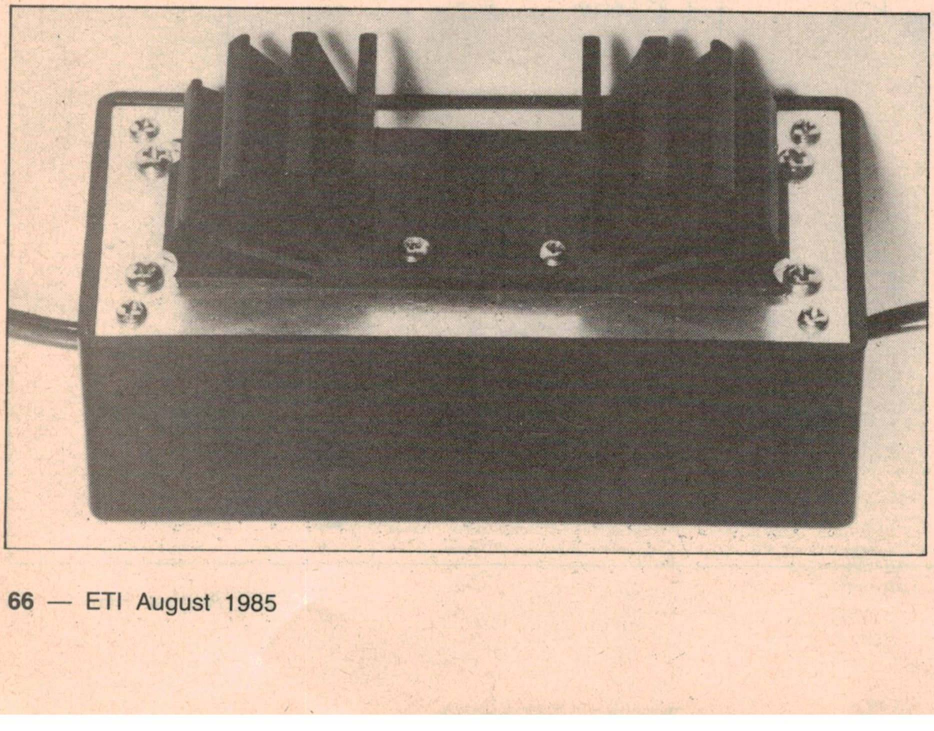

HOW IT WORKS — ETI-341

The battery charger is a basic flyback or ringing choke converter using voltage sensing across the main switching transistor to determine peak coil and current sensing in the coil secondary to set coil minimum current. Switching is done using a power MOSFET so the converter can operate at high frequency and an LM339 quad comparator is used to perform all control functions.

The converter operates by switching on the main transistor Q1 and allowing current to flow through the primary of L1. Since the voltage across the coil is constant, the current through it will rise linearly with time. As a power MOSFET which is turned hard on behaves very much like a resistor, then with the rising current in the coil the voltage across Q1 also rises. As Q1 is on, its gate is held at about 13 volts by the output of the comparator IC1 pins 1, 6 and 7 through the driver transistors Q2 and Q3.

The positive voltage on the MOSFET gate means that the negative input of the comparator IC1 pins 2, 4 and 5 is also held positive through R9. This makes the output of this comparator, which is used to sense coil secondary current, negative and reverse bias (or at least not forward bias) D1. This leaves the negative input of the MOSFET drive comparator pins 7, 6 and 5 to be set by the potential divider R4 and R5. This situation remains stable as long as the coil current doesn't exceed its peak value.

As soon as the coil primary current reaches its maximum value, the voltage on the comparator pin 6 rises above that set on pin 5 and the comparator output changes state. This turns off Q1 and positive feedback through C3 and C4 ensures a fast and clean switch. The 12 to 15 amps that was flowing in the primary winding is immediately transferred to the secondary circuit and flows through D2 and R22. R22 is the secondary current sense resistor and the current flows in such a direction as to take the end of the coil and the end of R22 negative. As the power FET is off, its gate is at ground potential and hence pin 4 of the current sense comparator is also at ground. R17 and R18 couple the negative voltage from the current sense resistor to the positive input pin 5. As current runs down in the coil secondary the voltage across R22 decreases linearly until pin 5 of the current sense comparator starts to go positive.

At this time the output of the current sense comparator goes rapidly positive and forward biases D1. This takes the positive input of the FET drive comparator positive and hence the output of the comparator goes positive and turns the FET on again to repeat the cycle. Thus the current through the coil is switched between a maximum value set by the on state voltage drop across the power MOSFET and a minimum value (nominally half of the peak value) set by a sense resistor in the secondary circuit.

To monitor the converter output voltage a zener diode is used to generate a reference voltage which is applied to a control comparator, IC1 pins 8, 9 and 14. As the reference is applied to the positive input of the comparator and some hysteresis is needed in the voltage control, R12 is in series with the input and R10 provides positive feedback around the comparator. A simple potential divider R19 and R20 transfers the output voltage and applies it to the negative input of the comparator. Thus if the output voltage of the converter rises above about 17 volts, pin 14 of IC1 goes negative and holds the gate drive to the MOSFET off.

One last problem remains to be solved which is that when current is flowing into the secondary circuit it momentarily raises the output because of the unwanted series resistance of the filter capacitor C8. This is a rather unknown property as the quality of the electrolytics will vary with the brand. Its effect is to raise the voltage on pin 8 during the current rundown and prevent the current control comparator from turning the FET on again until all the current had been transferred. This reduces the charge current out of the converter.

To avoid this, the fourth comparator in the IC is used to sense when current is flowing in the secondary circuit and disable the output voltage comparator pins 8, 9 and 14. When current is flowing in the secondary circuit the anode of the main diode D2 is positive and when the power FET is on it goes negative. Thus it's an easy business to translate this large voltage swing to suitable levels for the comparator input.

R15 and R16 ensure that the large diode voltage swings don't exceed the maximum input voltages. When current is flowing in the secondary, the output pin 13 goes to ground and pulls pin 8 of the output voltage comparator negative. This ensures that pin 14 remains positive during the secondary current cycle of the inverter and allows the power FET to be turned on again at the correct time. When the output in 14 is negative and holding the converter off it also passes current through the green LED2 to indicate that the output voltage has reach 17 volts and charging is completed.

Also, when the converter is running and current is flowing in the secondary circuit the drain of the MOSFET is above the positive rail. This turns on D3 and passes current through the yellow LED1 and R1 to indicate that the converter is running.

Zener diode ZD2 (75 volt) is placed between the drain of the FET and ground as there is always a small amount of energy that isn't coupled to the coil secondary winding. This energy is stored in what is called the transformer primary leakage inductance. Since the power MOSFET switches in a few tens of nanoseconds there is no time for the transistor to safely dissipate this energy and it will produce excessive voltage spikes on the primary and drain. ZD2 absorbs these spikes and prevents the FET from being damaged.

best to hold the board up to the assembly you're marking to see what you're doing.

Finally, drill the two holes you've marked out. In the prototype I actually tapped threads in these two holes but self-tapping screws would work as well so long as they are big. As this bracket must conduct waste heat to the heatsink it *must* be clamped good and tight. Also, nuts and bolts can't be used as the nuts would be a bit hard to get at.

The board is mounted on the lid with four 1-inch tapped metal spacers. Attach the spacers to the board and the board to the lid. The heatsink bracket should lie nicely against the power device mounting surfaces. Mark where the holes in the power devices lie on the bracket (it's a bit awkward but I did it OK) then take the board off the lid and the bracket off the heatsink. Drill the last two holes in the bracket to clear the insulating hardware that comes with the power devices and mount the bracket onto the power devices using all the insulating hardware. Connect up the power leads to the board and the cigarette lighter plugs to the other ends of the leads *making absolutely sure that the positive connects to the centre pin of the plugs.* The input lead should only be a metre or so long but the output can be as long as you like. Finally, using the piece of paper you've kept, connect up the four leads from the LEDs in the lid to the board and it's ready for testing.

Testing

Before turning anything on, double-check that the zener diode ZD2 is in and the right way round. If it's the wrong way round it will burst into flames and if it isn't in at all the MOSFET will have a hard time. Connect a voltmeter across the output and set it to 20 or 30 volts range. Finally, apply power to the input (if you're using a power supply watch the polarity!). There should be a current in-rush as the power is applied then the current should drop to about 20 mA. The output voltage should jump to about 17 volts. The 'charging' LED1 will not come on as the converter has only to maintain the charge on the output capacitor, but the 'charge completed' LED2 should light. Anything else means you've got trouble.

If the output voltage goes off scale then the regulation circuit isn't operating or D4 is in the wrong way. As a final check disconnect the voltmeter from the output and, if it has a 10 A range, set it to this then *briefly* connect it to the output. The meter should read about 6 amps and the 'charging' LED1 should come on. Also the 'charge completed' LED2 should go off. Don't leave the ammeter across the output for too long though, as the power devices aren't connected to the heatsink. If all is well you can screw the board onto the lid and the heatsink mounting bracket onto the heatsink.

Finally, ease the whole works into the box and screw down the lid. The whole unit is then ready for use.

A final word of warning: if the charger is connected the wrong way round and attempts to draw charge off the flat battery, then that is exactly what it will do. Unless the flat battery is absolutely dead flat then everything will appear to be just fine except that the flat battery will be sucked dry to charge the good one — a waste of time. Also, when you're using the unit, the higher the input voltage the quicker the battery will be charged so leave the engine idling while it's running unless you're not in a hurry. ●