Campsite AC Monitor

"I packed my bag and in it I put…"

Ton Giesberts (Elektor Labs)

Those of you who go camping regularly will have experienced this before: you switch on a powerful electric hob just when the fridge is on and the campsite fuse blows. It's a pain because the campsite manager usually has to be called to replace or reset the fuse and the exercise can cost a lot of money. This campingsite AC monitor makes these occurrences a thing of the past.

Camping hook-ups for power line supplies usually have a limit on the amount of current that can be drawn. When a larger current is drawn it trips a fuse, which most likely has to be reset by the campsite manager and which probably results in a fine (or 'service charge') to be paid.

To prevent such inconveniences we have designed this controller that can quickly limit the maximum current drawn.

Primary operation

This circuit ensures that the current cannot increase any further once a presettable value has been exceeded. One consequence of this is that electronic devices (such as TV sets, radios, energy saving bulbs) must not be connected to this circuit. The circuit is primarily intended to be used with energy-hungry appliances such as electric ovens, hobs (without electronic controls!) and pressure cookers. These can sometimes consume as much as 3 kW. The turning on of such an appliance can result in the immediate dropping out of the power line voltage.
To avoid this it is best to connect the controller between these devices and the AC power line. The smaller appliances can then be connected directly to the power line. You should still bear in mind what the total current consumption of these devices is, however. For example, if this is just under 1 A then the value set in the controller should be one Amp less than the maximum current that can be drawn from the camping hook-up.

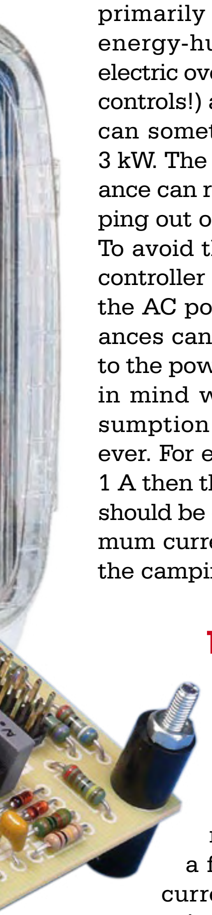

The circuit

At the hart of the circuit is IC1, a U2008B made by Atmel (see **Figure 1**). This 8-pin phase controller requires only a few external components. The IC has a facility to measure the load current, which is ideal for the prevention of an overload. The inclusion of automatic retriggering means that inductive loads won't be a problem. The IC also offers a soft-start function (connect a capacitor between pin 1 and ground) or detection of the load current via a shunt resistor in series with the triac (also between pin 1 and ground). We've chosen to use the soft-start function here.
The current through the load is measured using a shunt resistor in series with the triac. A separate detector circuit drives the control input (pin 3) of the IC.
The ('negative'!) supply of the circuit is internally regulated by the U2008B. In our prototype the voltage was found to be just below 16 V. The IC requires at least 3 mA, the applied rail-to-rail opamp (TS922IN) needs at most 3 mA (unloaded), the LED 4 mA (pulsed) and the reference, 1 mA. This is the reason for increasing the current of the supply for the U2008B to 10 mA. For this we've used two 5-watt resistors of 4kΩ7 and D1 in series with the power line AC. The voltage across C4 is actually determined by the average current flowing through R1 and R2. This current can be calculated using the following formula:

$(U_{mains} - U_{supply}) \times \sqrt{2} / [\pi (R1+R2)]$

We've ignored any voltage drop across D1. This formula is very similar to the one for the resistor found in the datasheet for the U2008B.
The effective value of a half-wave rectified voltage is $U_{peak}/2$. The average value, however, is U_{peak}/π. Because of the half-wave rectification, the heat generated in resistors R1 and R2 (U^2/R) is about 2.5 times greater than it would be for a DC current with a value equal to the average value of an AC current. From a safety aspect we

Technical specification

- For 230 VAC or 110 VAC campsite power lines
- Limits at 3 A, 4 A, 5 A, 6 A, 7 A, 8 A or 10 A
- Indicator LED
- Activation level configurable by jumper or rotary switch

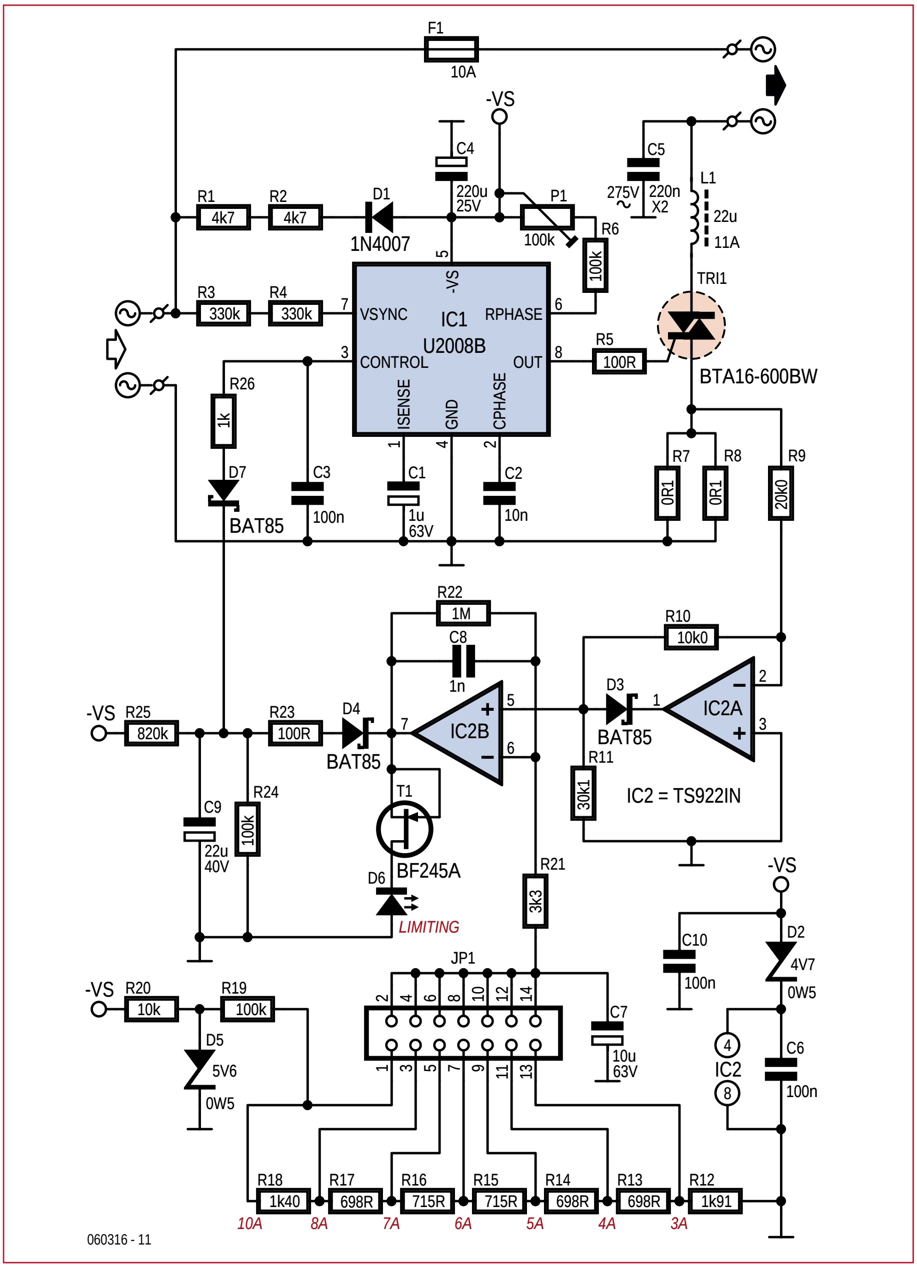

Figure 1. The circuit diagram of the campsite AC monitor illustrates that it isn't an easy task to limit the mains supply.

decided to use two 5 watt resistors for R1 and R2, even though the total dissipation is only 2.5 W.
The supply for the opamp should not be greater than 12 V. A zener diode (D2) has been connected in series with the supply to keep the voltage within safe limits.

Limiting

The circuit was designed for 230 VAC power lines. The above formula and what follows on the circuit operation should allow adaptation to 110-117 VAC networks by changing relevant component values.
We decided not to make the threshold for the limiting adjustable, but instead have a jumper on the board that selects one of seven limits via a 14-way header. This also gives the opportunity to use a (rotary) switch to change the setting after completion of the installation. The shunt consists of two 5 W resistors of 0.1 Ω connected in parallel. At 10 A the power loss is only 5 W and at the lower end of the scale the voltage can still be measured without any extra amplification.
The measured signal is first rectified. A simple full-wave rectifier is built around IC2a. When the signal is positive IC2a functions as an inverting amplifier and the voltage across R11 via D3 is determined by the relative values of R10 and R9. The signal is therefore attenuated by a factor of 2. The attenuation is necessary because of the simplicity of the full-wave rectifier; with this particular design only one diode is required.
When the signal is negative the output of IC2a stays at ground level and D3 is reverse-biased. R9–R11 now form a potential divider. The voltage across R11 becomes half of the input signal. The rectified signal is then fed to comparator/amplifier IC2b. This compares the peak value of the current with the reference set up.

The reference can be chosen from seven values. These have been configured such that the current limits are 3, 4, 5, 6, 7, 8 or 10 A. Jumper JP1 can be used to select the required limit. The speed at which IC2b switches is limited somewhat by C8, R22 and R21 in order to obtain a stable output. Any spikes, glitches and other high-frequency interference are filtered out by these components. The pulse at the output of IC2b that is produced when the current exceeds the limit is used to create the control voltage for IC1.
Junction FET (JFET) T1 has been configured as a current source and drives an LED that lights when the limiting is active. The current source is necessary because when a resistor is used a small pulse will hardly light up the LED (even when a low-current device is used). C9 is charged via D4 and R23. The voltage across C9 then provides the control signal for IC1 via D7 and R26. C3 provides extra decoupling.
Due to the adjustment range of the control input and the amplification of the measured signal an increase of the load current is required before the voltage can be reduced to its minimum value. This increase is about 1 A. Because of this the levels of the reference voltages have been adjusted to a slightly lower value than is indicated on the board. The limiting is therefore applied a little bit 'before' the current limit is reached.
As the overload increases, the limiter will reduce the load current a little further below the selected value. The lower load resistance means that the current pulse becomes shorter and greater in amplitude. The control voltage across C9 then increases. The advantage of this is that the peak currents aren't as large.

Reference

The reference voltage is derived from a standard 5.6 V zener diode (D5), which has a current flowing through it of about 1 mA, set by R20. The reference voltages for IC2b are relatively small and are between 90 mV (3 A setting) and 340 mV (10 A setting). These voltages are generated by the potential divider formed by R12–R19.

For the calculations we assumed that about 50 µA flows through the resistor network. The voltage across zener diode D5 will be less than 5.6 V. For this reason we've chosen a value of 100 kΩ for R19. The total resistance of R12–R18 can then be neglected and the current remains fairly constant. The resistor values are then easy to calculate.
C7 decouples the reference voltage selected via JP1. There is some non-linearity present in the feedback loop, because the initial increase of the phase angle has much less of an effect on the effective voltage than when the phase angle is changed at 90 degrees.
The control also has a dead area. Because the control input only becomes

Meters don't lie, or do they?

The supply voltage for the U2008B is internally regulated. Our circuit hardly differs from the standard application. The only difference is that our circuit requires a higher current. We have designed the supply for a current of about 10 mA, as already mentioned in the article. However, during the test and measurement phase we found out a few things about RMS multimeters.

C4 is charged up by a single-phase rectified current. Theoretically, the average value of a pure sine wave is 2Upeak/π, which also applies to a full-wave rectified sine wave. The effective value is also the same in both cases: Upeak/√2. This is where the well-known crest factor of 1.11 (π/2√2) comes from. But it's all very different when it concerns a half-wave rectified sine wave. The average value here is Upeak/π. The effective value is a lot larger though, Upeak/2. If we now use a standard AC voltmeter to measure the voltage across R1 and R2 (and leave the blade terminal disconnected), we will measure the average value. In our case, at the castle (Elektor House) with a somewhat low power line voltage, it was only 220 VAC instead of 230 VAC. The peak voltage across R1 and R2 is then √2×220 – 16.6 V. The average value should be about 94 V, which is also what we measured. If we now use a true-RMS meter we expect to see the peak voltage divided by two, which is 147 V. The strange thing was that the true-RMS meter displayed 115 V. Our first thought was: 'Perhaps something is wrong with the meter, after all, it's getting quite old'? But another (newer) meter, although from the same manufacturer, gave exactly the same result. A more advanced meter provided the solution. This was able to measure AC and DC and displayed the expected 147 V. It appears that many true-RMS meters can't cope with a DC component. Those that do usually have it specifically stated on the meter, so bear that in mind.

This also applies to digital, square-wave shaped signals and the like, not just for half-wave rectified sinusoidal signals. This just goes to show that you have to know what you're measuring and select an appropriate meter accordingly as the wrong choice can lead to an inaccurate result.

active when the voltage falls below –2 V, the load current has to increase by about 0.2 A compared to the selected current before the limiting starts. This is another reason for making the thresholds a bit lower than indicated. The resistor values for the divider have been rounded to values from the E96 range, otherwise the total error would become too much.

P1 is used to set the maximum phase angle. At the lowest setting it is still possible to limit the current to 3 A during a large overload, otherwise it stays a bit lower.

Heatsink

For the triac we have selected a snubberless type made by STMicroelectronics, the BTA16-600BW. The triac is available in an isolated version (still with a metal tab) and can handle 16 A. The BW version needs a trigger current of at least 50 mA. The value of 100 Ω chosen for R5 provides a gate current slightly higher than this.

The disadvantage of the isolated version is a higher internal thermal resistance: 2.1 K/W instead of 1.2 K/W. On top of this, the maximum permitted junction temperature is only 125 °C. The voltage drop across the triac is partially dependent on the junction temperature. At 10 A the voltage drop at a junction temperature of 125 °C is about 0.25 V lower than at 25 °C.

For calculating the size of the heatsink we assumed that the ambient temperature could be 50 °C. This may seem high, but when the circuit is housed in a case and placed in a caravan in the summer it could possibly be higher. The dissipation at 10 A is about 11 W. The maximum total thermal resistance is then

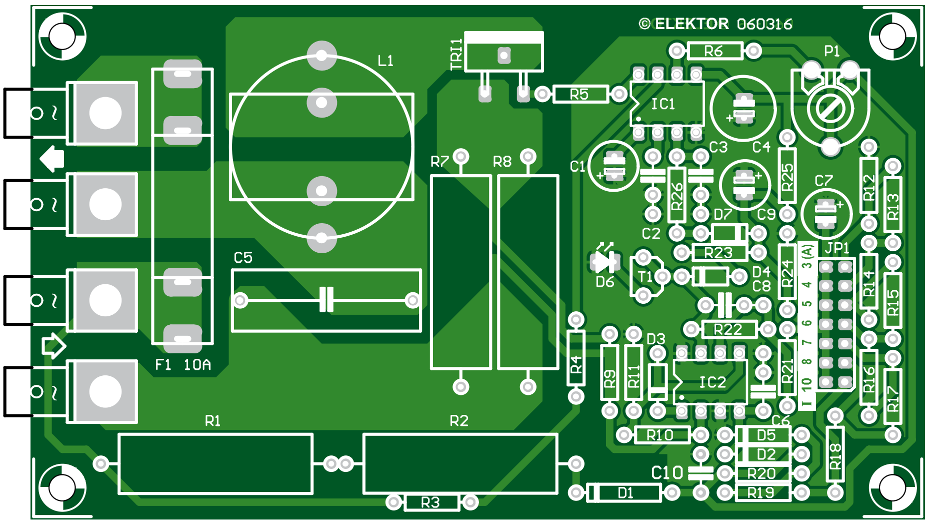

Figure 2. You can see from the component layout and the lighter coloured track layout that the AC power section is more spacious in order to comply with electrical safety requirements.

COMPONENT LIST

Caution.
Circuit designed for 230 VAC power lines. Please refer to circuit description for component adaptation to suit 110-117 VAC networks.

Resistors

R1,R2 = 4kΩ7 5W, lead pitch 30mm max.
R3,R4 = 330kΩ
R5,R23 = 100Ω
R6,R24 = 100kΩ
R7,R8 = 0Ω1 5W, lead pitch 30mm max.
R9 = 20kΩ0
R10 = 10kΩ0
R11 = 30kΩ1
R12 = 1kΩ91
R13,R14,R17 = 698Ω
R15,R16 = 715Ω
R18 = 1kΩ40
R19 = 100kΩ
R20 = 10kΩ
R21 = 3kΩ3
R22 = 1MΩ
R25 = 820kΩ
R26 = 1kΩ
P1 = 100kΩ preset

Capacitors

C1 = 1µF 63V, radial, lead pitch 2.5mm, diam. 6.3mm max.
C2 = 10nF lead pitch 5mm or 7.5 mm
C3,C6,C10 = 100nF, lead pitch 5mm or 7.5mm
C4 = 220µF 25V, radial, lead pitch 2.5mm, diameter 8.5mm max.
C5 = 220nF 275VAC X2, lead pitch 22.5mm
C7 = 10µF 63V, radial, lead pitch 2.5mm, diam. 6.3mm max.
C8 = 1nF, lead pitch 5mm or 7.5mm
C9 = 22µF 40V, radial, lead pitch 2.5mm, diam. 6.3mm max.

Inductors

L1 = 22µH 11A e.g. 1422311C Murata Power Solutions (Farnell order code 1077056)
or
22µH 10.3A e.g. 2205-V-RC (J.W.Miller Magnetics), Digi-Key # M8868-ND

Semiconductors

D1 = 1N4007
D2 = 4.7V 0.5W zener diode
D3,D4,D7 = BAT85
D5 = 2.7V 0.5W zener diode
D6 = red low current LED
T1 = BF245A
TRI1 = BTA16-600BWRG (TO220AB insulated) (Farnell # 1175636)
IC1 = U2008B (Atmel), 8-pin DIP
IC2 = TS922IN (ST), 8-pin DIP

Miscellaneous

K1 = 14-way (2x7) pinheader + 1 jumper
F1 = 10A 1¼" x ¼" e.g. Farnell # 1175149 + 2 fuse clips 15A rated, Farnell # 1175125
4 AMP connectors , M4 screw mounting, + 4x 10 mm M4 bolt + nut + washer + locking ring
Ceramic isolation 4.5mm e.g. type AOS220SL (Fischer Elektronik)
Heatsink for 10A: R_{th} < 3.7 K/W (< 9.4 K/W voor max. 6A)
PCB # 060316-1 from www.thepcbshop.com

(125–50 °C) / 11 W = 6.8 K/W.

This value is reduced by 2.1 and 1 K/W for the triac and the isolation. The heatsink used should therefore have a thermal resistance better than 3.7 K/W to provide sufficient cooling for the triac at the maximum current of 10 A.
When the maximum available current on your favourite campsite is only 6 A you could use a smaller heatsink:

(125–50 °C) / 6 W –2.1–1 = 9.4 K/W.

This calculation goes to show that a few amps difference result in a big change in the size of the heatsink.
We should make one thing clear though: these calculations are for the maximum operating temperature of the triac. This isn't beneficial to the lifespan of the semiconductor. If we want to give the triac a longer life then the following applies: the more cooling the better.

Printed Circuit Board

Since the circuit board is single-sided and the 10 A tracks require a certain minimum amount of copper and the triac is at the edge of the board for ease of construction, we decided to use a thick ceramic isolator (see parts list). This isolation wasn't chosen to comply with electrical safety regulations, as that is already covered by the internal isolation of the triac. It provides more space on the board for the A2 connection of the triac (on the non-isolated version this is also the tab). The connections for the gate and A1 have been placed a bit further away from the triac in order to make more room for the copper to the A1 connection.

Due to the large currents involved, we haven't held back on the copper used to make the connections to the other components (F1, L1, R7, R8, the mains and load connectors). For the 10 A fuse there are two separate fuse clips for a 32 mm fuse mounted on the board. These clips are rated at 15 A. The connections for L1, R7 and R8 have been made without the usual thermal reliefs. This does mean that you will need a higher power soldering iron to solder these components compared to the rest of the components.
For the load and power line connectors we've used automotive blade (AMP) connectors that have been mounted on the board using 4 mm screws. The distance between these connectors is slightly more than the minimum required 3 mm. When you screw the connectors in place you need to make sure that they are mounted perfectly straight. On the prototype they were mounted on the component side, but depending on the way the board is housed, you may consider mounting them on the solder side. This has the advantage that losses caused by the resistance of the screw connection are avoided. You should still make sure that there is a minimum separation of 3 mm between the connectors (and PCB tracks) that carry mains voltages.
The current limits have been clearly printed on the board next to JP1. To reduce the stress on the solder joints it's advisable to provide the power resistors with a small kink in their legs before soldering them onto the board (this also applies to the central connection on the triac, A2).

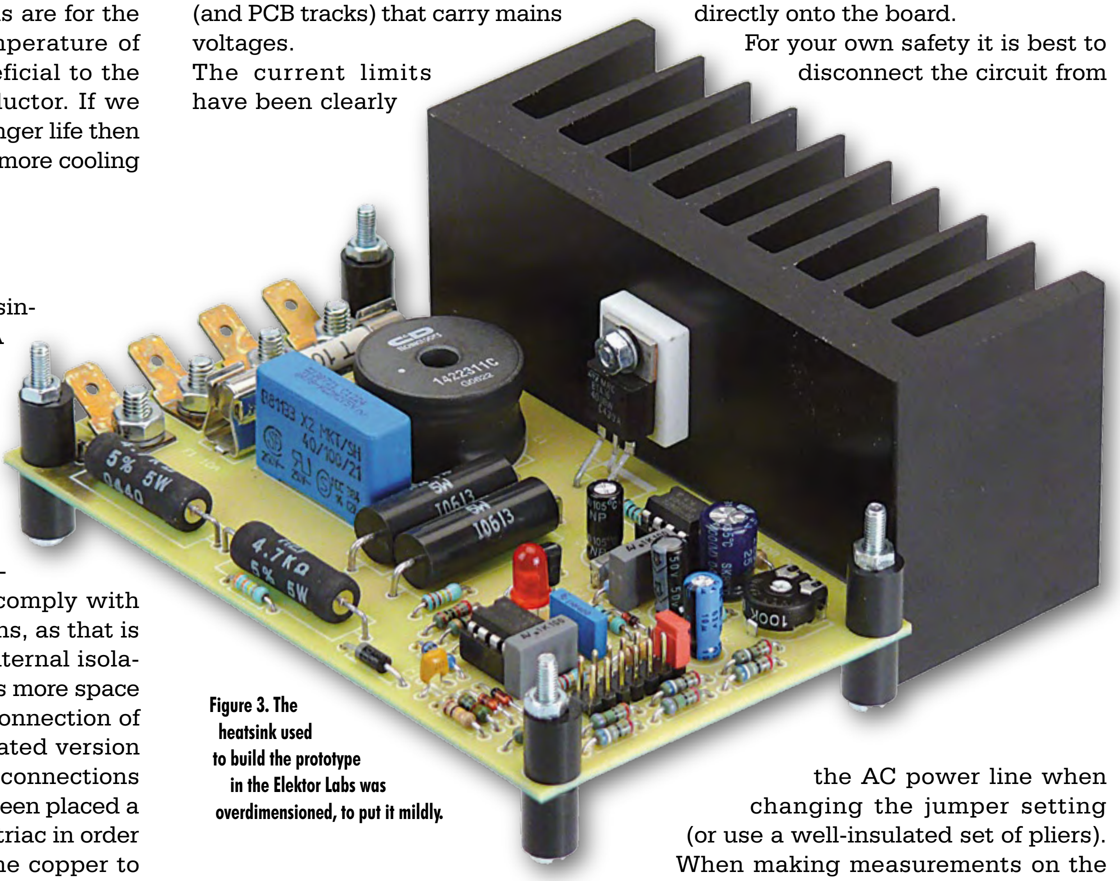

Figure 3. The heatsink used to build the prototype in the Elektor Labs was overdimensioned, to put it mildly.

Safety and earthing

Since the triac has been positioned at the edge of the board, it follows that it's easiest to mount the heatsink on the edge of the board as well. To comply with isolation requirements you should stick some insulating tape between the heatsink and the board (because of the required 3 mm separation). For safety reasons the heatsink has to be earthed.
The controller is intended for use as a Class-I electrical device. This means that if it's part of a distribution box it should include a reliable protective earth (PE) line. An LED is connected to the output of IC2b to indicate when the circuit is limiting the current. The whole circuit is connected directly to the mains, including the LED. For this reason it is not permitted for it to be mounted such that it protrudes through the case and can be touched. It is therefore best to mount the LED directly onto the board.
For your own safety it is best to disconnect the circuit from the AC power line when changing the jumper setting (or use a well-insulated set of pliers). When making measurements on the circuit you also have to take great care, especially with points connected to AC power: between R1 and the blade terminal on the corner, for example. A short between these two can cause the track next to the blade terminal to disappear in a puff of smoke, as we found out the hard way during testing.
The way in which this circuit is housed depends whether you want to include it in an existing installation or if you want to create a type of enhanced distribution box. You should also bear in mind the size and temperature of the heatsink and the amount of heat released.

(060316-I)