Arduino-based Digital Inductance/Capacitance Meter

Do you ever need to check or confirm the values of inductors or capacitors? This Arduino-based LC meter will give you a digital readout and can even measure parasitic inductance or capacitance present in a circuit. It's much more accurate than most DMM-based LC meters.

Many digital multimeters (DMMs) have capacitance ranges but they are not normally accurate for values below about 50pF. And those few DMMs that can measure inductance are often not very good at measuring inductance in the range of 1-100µH – those that are typically used in audio and RF circuits.

An inductance meter with a 10µH resolution (typical for DMMs) isn't very helpful if you want to wind a choke of say 6.8µH, for an amplifier output filter.

Professionals tend to rely on digital LCR meters for these types of measurements. They allow you to measure almost any passive component quickly and automatically, often measuring not just their primary parameter (like inductance or capacitance) but one or more secondary parameters as well. However, many of these you-beaut instruments also carry a hefty price tag, keeping them well out of reach for many of us.

Fortunately, thanks to microcontroller technology, much more affordable digital instruments are becoming available. These include both commercial and DIY instruments like the low-cost unit described here.

By JIM ROWE

Essentially it's an improved version of the PIC-based Digital LC Meter we described in the May 2008 issue of Silicon Chip. This time, we're basing it around an Arduino Uno or equivalent module.

Main features

Our new Digital LC Meter is compact and easy to build, since the Arduino board comes pre-assembled. It also has a better LCD readout than the previous version. It fits snugly inside a UB3 utility box and you should be able to build it for under $100.

It offers automatic digital measurement of both inductance (L) and capacitance (C) over a wide range and with 5-digit resolution. Measurement accuracy is better than ±1% of reading over most of the ranges.

It operates from 5V DC, drawing an average current of about 62mA, so it can run from a 5V USB supply (either mains or battery) or from a spare USB port on your PC.

How it works

The meter's impressive performance relies on an ingenious measurement technique developed almost 20 years ago by the late Neil Heckt in the USA.

It uses a wide-range test oscillator and its frequency is varied by connecting the unknown inductance or capacitance you're measuring. The resulting change in frequency is measured by the microcontroller and used to calculate the component's value, which is displayed directly on a small LCD panel.

To achieve reliable oscillation over a wide frequency range, the test oscillator is based on an analog comparator with positive feedback around it, as shown in Fig.1. This configuration has a natural inclination to oscillate, because of the very high gain between the comparator's input and output.

When power (+5V) is first applied, the comparator's positive input is held at +3.3V by the divider formed by the two 100kΩ resistors and the 100kΩ and 4.7kΩ resistors. Initially, the voltage at the negative input is zero because the 10µF capacitor at this input needs time to charge via the 47kΩ resistor.

So with its positive input much more positive than the negative input, the comparator initially switches its output high, to near +5V.

Once it does so, the 10µF capacitor connected to the negative input begins charging up via the 47kΩ resistor and the voltage at this input rises. As soon as it goes above +3.3V, the comparator output switches low and the positive input is brought to 1.67V due to the 100kΩ feedback resistor pulling the 100kΩ divider low.

The low comparator output voltage is also coupled through the 10µF input capacitor to the tuned circuit formed by inductor L1 and capacitor C1. This makes the tuned circuit "ring" at its resonant frequency.

As a result, the comparator and the tuned circuit now function as an oscillator at that resonant frequency. In effect, the comparator functions as a negative resistance across the tuned circuit, to cancel its losses and maintain oscillation. Once this oscillation is established, a square wave of the same frequency is present at the comparator's output and it is this frequency (F_{OUT}) that is measured by the microcontroller.

In practice, before anything else is connected to the circuit, F_{OUT}, will simply correspond to the resonant frequency of the tuned circuit comprising L1, C1 and any stray inductance and capacitance that may be associated with them.

When power is first applied to the circuit, the microcontroller measures this frequency (F1) and stores it in memory. It then energises reed relay RLY1, which switches capacitor C2 in parallel with C1 and thus lowers the oscillator frequency. The micro then measures and stores this new frequency (F2).

Next, the micro uses these two frequencies plus the known value of C2 to accurately calculate the values of both C1 and L1. The equations it uses to do this are shown in Fig.1. Following these calculations, the micro turns

Fig.1: operating principle of the LC Meter. L1 and C1 form a tuned circuit in combination with an external capacitance/inductance connected via S1. Feedback for oscillation is provided by a comparator and the frequency of oscillation depends on the known values of L1/C1/C2 and the unknown external component. The unknown value can be computed from the frequency of F_{OUT}, as described by the accompanying equations.

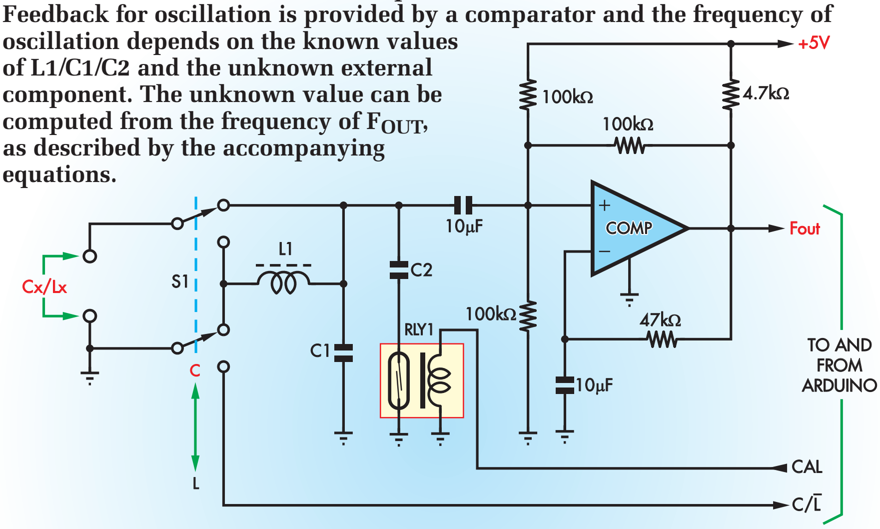

HOW IT WORKS: THE EQUATIONS

(A) In calibration mode

(1) With just L1 and C1: $$F1 = \frac{1}{2.\pi.\sqrt{L1.C1}}$$

(2) With C2 added to C1: $$F2 = \frac{1}{2.\pi.\sqrt{L1.(C1+C2)}}$$

(3) From (1) and (2), we can find C1:

$$C1 = C2 \cdot \frac{F2^2}{(F1^2 - F2^2)}$$

(4) Also from (1) and (2), we can find L1:

$$L1 = \frac{1}{4.\pi^2.F1^2.C1}$$

(B) In measurement mode

(5) When Cx is connected: $$F3 = \frac{1}{2.\pi.\sqrt{L1.(C1+Cx)}}$$

so $$Cx = C1 . \left(\frac{F1^2}{F3^2} - 1\right)$$

(6) Or when Lx is connected:

$$F3 = \frac{1}{2.\pi.\sqrt{(L1+Lx).C1}}$$

so $$Lx = L1 . \left(\frac{F1^2}{F3^2} - 1\right)$$

NOTE: F2 & F3 should always be lower than F1

Features & specifications

Inductance range:	10nH to 100mH+
Capacitance range:	0.1pF to 2.7µF+ (non-polarised only)
Measurement resolution:	five digits in either mode
Range selection:	automatic
Sampling rate:	approximately one measurement per second
Accuracy (when calibrated):	±1% of reading, ±0.1pF or ±10nH
Supply voltage:	5V DC @ <65mA (including backlit LCD)
Supply type:	USB charger or the USB port on a PC

RLY1 off again to disconnect C2, allowing the oscillator frequency to return to F1. The unit is now ready to measure the unknown inductor or capacitor (L_X or C_X).

As shown in Fig.1, the unknown component is wired to the test terminals at far left. It is then connected to the oscillator's tuned circuit via switch S1.

When measuring an unknown capacitor, S1 is switched to the "C" position so that the capacitor is connected in parallel with C1. Alternatively, for an unknown inductor, S1 is switched to the "L" position so that the inductor is connected in series with L1.

In both cases, the added L_X or C_X again causes the oscillator frequency to change to a new frequency (F3). As with F2, this will always be lower than F1.

So by measuring F3 as before and monitoring the position of S1 (which is done via the C/$\overline{\text{L}}$ line), the micro can calculate the value of L_X or C_X using one of the equations shown in the right section of the equations box in Fig.1.

From these equations, you can see that the micro has some fairly solid number-crunching to do, both in the calibration mode when it works out the values of L1 and C1 and in the measurement mode when it must work out the value of L_X or C_X.

Each of these values needs to be calculated to a high degree of resolution and accuracy, using floating-point maths.

Circuit details

The full circuit diagram is shown in Fig.2. It mainly consists of the Arduino microcontroller module and the serial I^{2}C LCD module together with the oscillator circuit we've already introduced, built using an LM311 high-speed comparator (IC1).

The Arduino controls RLY1 to switch calibrating capacitor C2 (1nF) in and out of circuit, via its IO3 pin. Diode D1 is connected across the relay coil to prevent the Arduino's internal circuitry from being damaged by inductive spikes.

The Arduino senses which position L-C switch S1 is in using its IO2 pin, which is pulled high internally when it's not pulled low by S1b (in the L position).

The output of the oscillator at pin 7 of IC1 is taken to pin IO5 of the Arduino via a series 6.8kΩ resistor. It needs to be taken to this pin because this is also the external input pin for the 16-bit timer/counter inside the ATmega328P micro which forms the heart of the Arduino Uno.

As a result, we are able to use the Arduino to easily measure the oscillator's frequency.

The results of the Arduino's measurements and calculations are displayed on a blue back-lit 16x2 alphanumeric LCD module.

This has a serial I^{2}C module fitted, so it can be controlled from the Arduino via its I^{2}C port lines (SCL and SDA).

Its features were fully described in Silicon Chip March 2017 issue.

Calibration functions

The firmware sketch running in the Arduino is designed to perform its "zero calibration" adjustment just after initial startup.

But pushbutton switch S2 is also provided to allow zero calibration to be performed at any other desired time as well (to allow for temperature drift, for example).

S2 pulls the Arduino's $\overline{\text{RESET}}$ pin (pin 4) down, so that it is forced to reset and start up again, readjusting its zero setting in the process.

LK1 and switch S3 can be used to nudge or tweak the calibration in small increments or decrements, if you have access to an accurate reference capacitor. When LK1 is fitted, pulling input pin IO7 low, the micro will increase

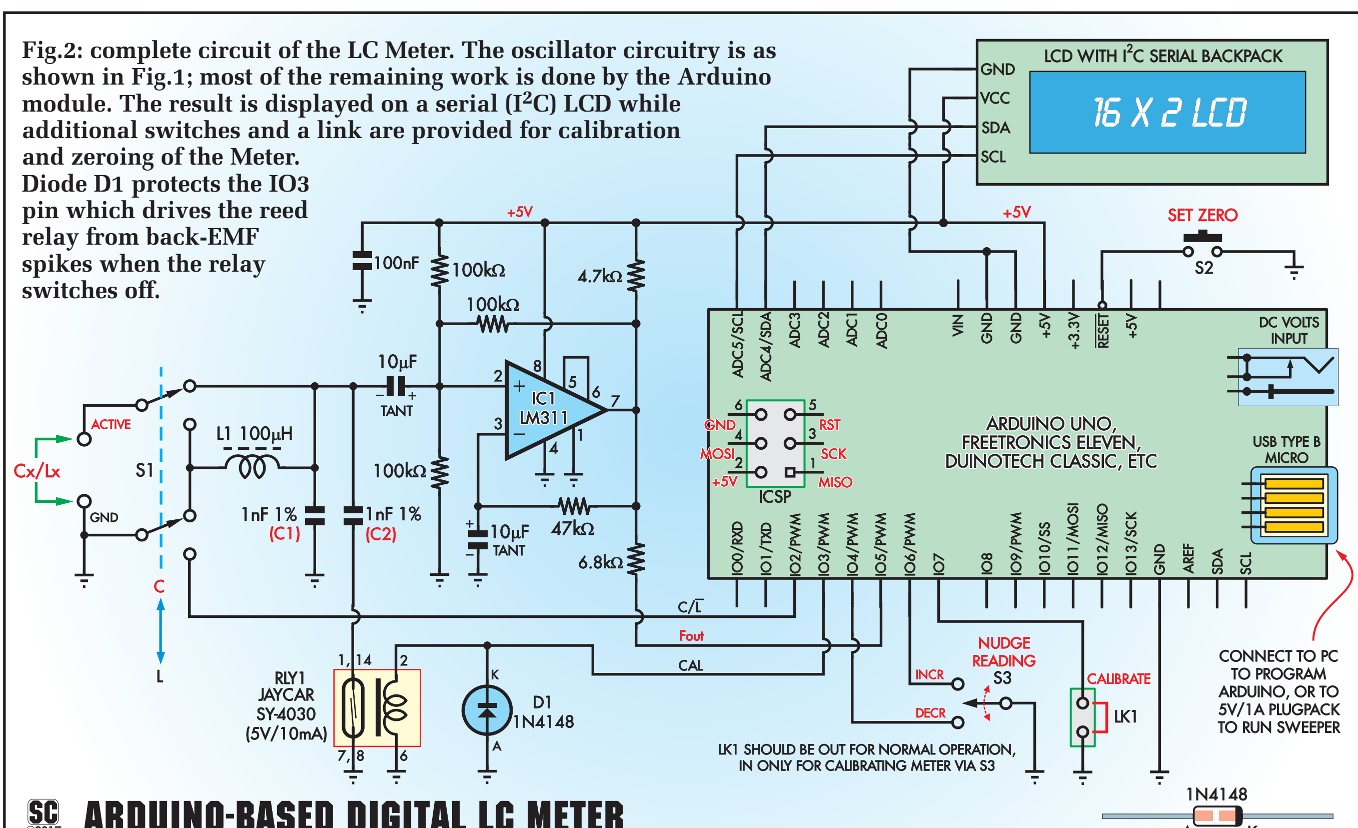

Fig.2: complete circuit of the LC Meter. The oscillator circuitry is as shown in Fig.1; most of the remaining work is done by the Arduino module. The result is displayed on a serial (I^{2}C) LCD while additional switches and a link are provided for calibration and zeroing of the Meter. Diode D1 protects the IO3 pin which drives the reed relay from back-EMF spikes when the relay switches off.

the capacitance reading by about 0.5% each time S3 (a centre-off rocker switch) is pushed to the upper "INCR" position, or alternatively decrease the reading by the same amount if S3 is pushed to the lower "DECR" position.

So the idea is to push S3 in one direction or the other until the reading is correct.

Each time a change is made, the adjustment factor is stored in the Arduino's EEPROM memory, so it's remembered for future sessions. When link LK1 is not fitted, pressing S3 in either direction has no effect at all.

This is a safety feature, to prevent unintended changes to the meter's calibration during normal use. Although this calibration is normally done using a reference capacitor, it also improves the accuracy of inductance measurements.

Construction

There is no custom PCB used for the LC meter's circuitry; instead, most of the added circuitry is fitted on a prototype shield board which simply plugs into the top of the Arduino PCB.

There aren't that many components involved, so it's a straightforward job to wire it up as shown in the wiring diagram, Fig.3.

The only components which are not mounted on the ProtoShield are the serial LCD module, switches S1-S3, the test terminal binding posts and reference components L1 and C1.

As shown in Fig.3 and the photos, these are all mounted on the lid of the UB3 box, which forms the meter's front panel. These off-board components are all linked to the ProtoShield board via short multi-wire interconnection leads and SIL connector plugs and sockets, which are also shown in Fig.3.

You can get an idea of how everything fits together from the internal cutaway diagram of Fig.4, along with the internal photos.

The Arduino module mounts in the bottom of the box via four 9mm long M2.5 machine screws and four M2.5 nuts, with another four M3 or M2.5 Nylon nuts used as spacers.

The ProtoShield is plugged into the top of it. The rest of the meter circuitry connects via the 90° pin headers on the ProtoShield.

Follow the wiring diagram (Fig.3) and internal photos to build the ProtoShield. Start by soldering the components into place where shown in

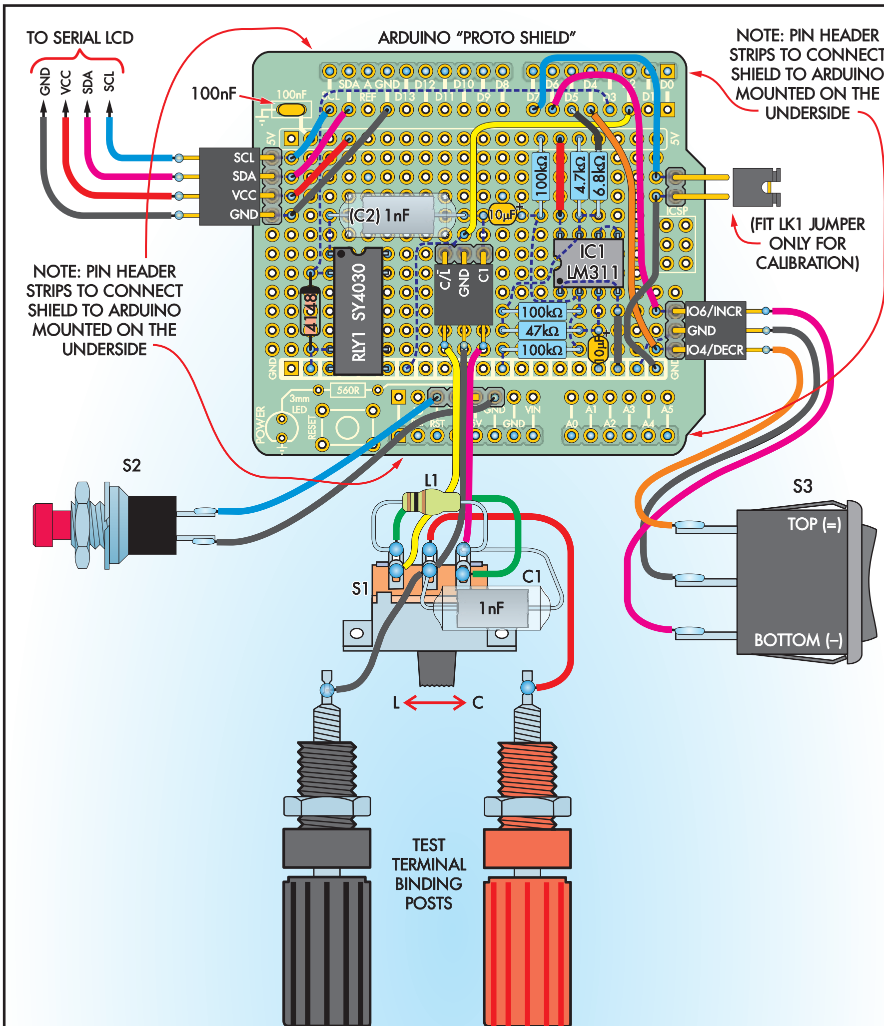

Fig.3: follow these diagrams to fit the components to the ProtoShield and also to wire up all the external connections. Connections made between component pads on the underside are shown below, significantly larger than the 1:1 diagram above, for clarity. These should be made with insulated wire to avoid short circuits.

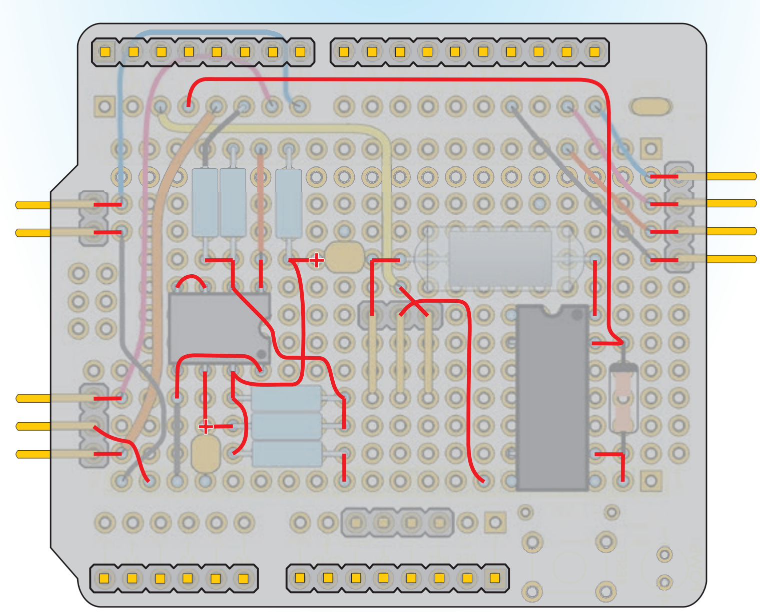

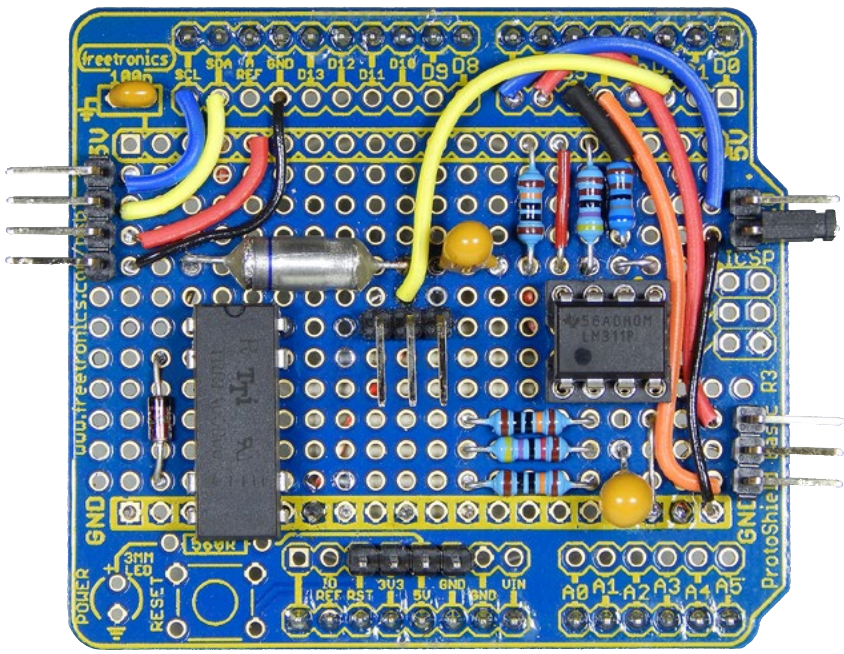

"Larger than life" photo of the wiring on the top side of the Freetronics Arduino ProtoShield board (actual size is shown in Fig.3, below). This board "plugs in" to the Arduino Uno (etc) board via the rows of pin headers on the underside; the I²C LCD board plugs into the ProtoShield board.

Fig.3, ensuring you use the correct orientation for polarised components: IC1, diode D1, RLY1 and the two 10µF tantalum capacitors.

Next, add the wiring on the underside, as shown in the underside wiring diagram of Fig.3. Use insulated wire because several of these wires cross over each other.

In cases where adjacent pads are connected, you can simply place a solder bridge between the two pads or alternatively, bend the component leads while fitting them and trim them so that they reach the adjacent pads. For longer connections, use component lead off-cuts, routed carefully to avoid the possibility of shorting anything else, or short lengths of light-duty hookup wire (eg, stripped from a piece of ribbon cable) or bell wire.

Here's our suggested order of fitting the components and wiring the ProtoShield board; check Fig.3 for the exact placement in each case:

1. Fit the four 90° SIL headers.
2. Fit a four-pin vertical header for switch S2.
3. Fit the four SIL pin headers to the underside, along the upper and lower edges of the ProtoShield, which connect it to the Arduino. These comprise a 10-pin header at upper left, two 8-pin headers (one at upper right and the other at lower centre) and a 6-pin header at lower right. Do not fit a 3x2 DIL pin header in the ICSP position at centre right on the ProtoShield board.
4. Fit the 8-pin DIL socket for IC1, with its notched end to the left, then relay RLY1, with its notched end towards the top.
5. Mount the six resistors, the 100nF capacitor and the two 10µF tantalum caps. Note that the last two are polarised, so make sure you fit them with the orientation shown.
6. Fit the 12 insulated wires on the top of the board and any insulated wires required to complete the wiring on the underside. This will require you to strip the insulation from each end by about 5mm or so.
7. Fit diode D1, making sure its end with the cathode band is uppermost and adjacent to pin 2 of RLY1, then plug IC1 into its socket.

Box and lid preparation

There are four holes to drill in the bottom of the box for mounting the Arduino module and two larger holes to cut in the left-hand end for the USB plug and alternative DC power plug.

The locations and dimensions of

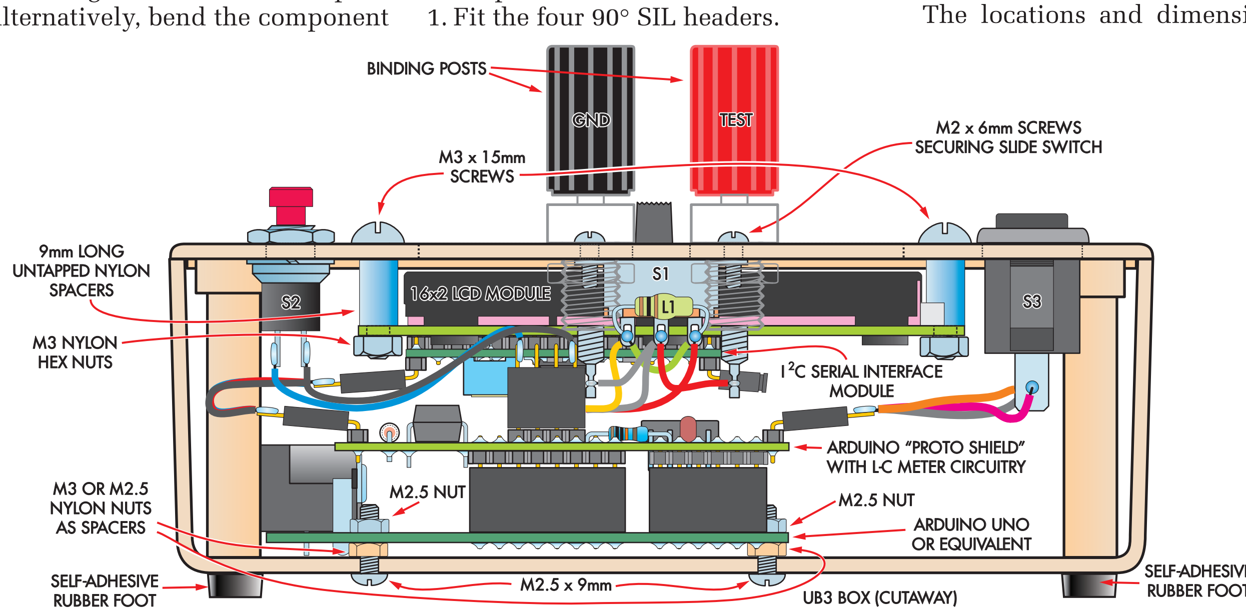

Fig.4: this shows how everything fits together inside the UB3 "Jiffy" box. The Arduino module is attached to the bottom of the case with the proto-board hosting most of the remaining circuitry plugged on top. The three switches, two binding posts and the I²C LCD module are mounted on the lid and connected to the ProtoShield via flying leads.

The Freetronics Eleven (Uno equivalent) board, mounted in the bottom of the case (see drilling template on pages 35 and 36).

all of these holes are shown in Fig.5, the drilling template, while the corresponding information for the holes to be drilled and cut in the lid/front panel are shown in Fig.6.

For best results, start the larger holes with a smaller pilot drill and enlarge with a stepped drill bit, series of larger drill bits or a tapered reamer. Rectangular or other non-round holes can be made by drilling a series of holes, knocking out the centre section and then filing the hole to shape.

We fixed four self-adhesive rubber feet to the underside of the box to protect any surface it's placed on.

Making all the required holes in the lid is rather tedious as there are twelve, including three rectangular cut-outs and two holes with flat edges.

To save time and guarantee a neat result, you can purchase a laser-cut clear acrylic lid (which replaces the lid supplied with the box) from the SILICON CHIP online shop (see parts list).

As the acrylic panel is transparent the lid doesn't need a cut-out to view the LCD. Note that since the 3mm acrylic is slightly thicker than the lid supplied with the UB3 box, depending on the length of the screws that came with it, you may need to use slightly longer self-tapping screws to attach it.

We have also prepared artwork for the front panel, to give it a professional look. You can download this as a PDF file from the SILICON CHIP website.

There are two ways to go here: after you print it, it can be hot laminated, then attach it to the box lid using double-sided adhesive tape or spray glue. After that, you can cut out the holes in the front panel to match those in the box lid using a sharp hobby knife.

Or, for longest life and an even more professional finish, consider fitting the label to the *underside* of the lid – it's more fiddly to fit but doesn't require laminating, nor double-sided tape to hold it in place (the switches and terminals hold it in position; a very light mist of clear spray adhesive will also ensure it stays tight against the lid).

Perhaps it's gilding the lily somewhat but if you can print the label onto clear film, you can see the “works” through the label, as we did with the photo on page 28.

Just make sure you get the right film to suit your type of printer (eg, laser printer or inkjet printer, etc).

Once the lid/front panel is finished, fit switches S1-S3 to it, along with the two test terminal binding posts and the serial LCD module.

Slide switch S1 attaches to the front centre of the lid via two 6mm long M2 machine screws, while switch S2 mounts using the spring washer and nut supplied with it and S3 simply pushes into its rectangular mounting hole until its two barbs spring outwards to hold it in place.

Just make sure that you fit it with the “=” sign on its rocker actuator uppermost (see photos).

The two binding posts are mounted using the mounting nuts and lock washers provided.

Take care doing so, however, as the upper and lower mounting bushes have D-shaped sections which should mate with the matching holes in the lid/front panel.

The serial LCD module mounts under the lid in the top centre position,

Parts list – Arduino-based LC Meter

1 Arduino Uno R3, Duinotech Classic, Freetronics Eleven or equivalent microcontroller module
1 Serial I^2C 16x2 LCD module with back-lighting (SILICON CHIP online shop Cat SC4198)
1 Arduino Uno Prototyping Shield (eg, Freetronics SH-Proto-Basic)
1 UB3 “Jiffy” box, 130 x 68 x 44mm
1 laser-cut clear acrylic lid for UB3 box [optional but recommended] (SILICON CHIP online shop Cat SC4274)
4 self-adhesive rubber feet
1 5V/10mA DIL reed relay (RLY1; Jaycar SY4030)
1 100μH axial RF inductor (L1; Jaycar LF1534)
1 DPDT subminiature slide switch (S1; Jaycar SS0821)
1 panel-mount SPST NO momentary pushbutton switch (S2; SP0710)
1 panel-mount SPDT on-off-on momentary rocker switch (S3; Jaycar SK0987)
1 8-pin DIL IC socket
1 40-pin header, 2.54mm pitch
1 40-pin right-angle header, 2.54mm pitch
1 150mm socket-to-socket jumper ribbon cable (Jaycar WC6026)
1 jumper shunt
2 binding posts with integral banana socket (1 red, 1 black)
4 9mm Nylon untapped spacers, 3mm inner diameter
4 15mm M3 machine screws
8 M3 Nylon hex nuts
4 9mm pan head M2.5 machine screws
4 M2.5 hex nuts
2 6mm M2 machine screws (for S1)

Semiconductors
1 LM311 DIP high-speed comparator (IC1; Jaycar ZL3311)
1 1N4148 small signal diode (D1)

Capacitors
2 10μF 16V through-hole tantalum
1 100nF multilayer ceramic
2 1nF 1% NP0 ceramic, mica, MKT, polypropylene or polystyrene (SILICON CHIP online shop Cat SC4273)

Resistors (all 0.25W, 1%, through-hole mounting)
3 100kΩ 1 47kΩ 1 6.8kΩ 1 4.7kΩ

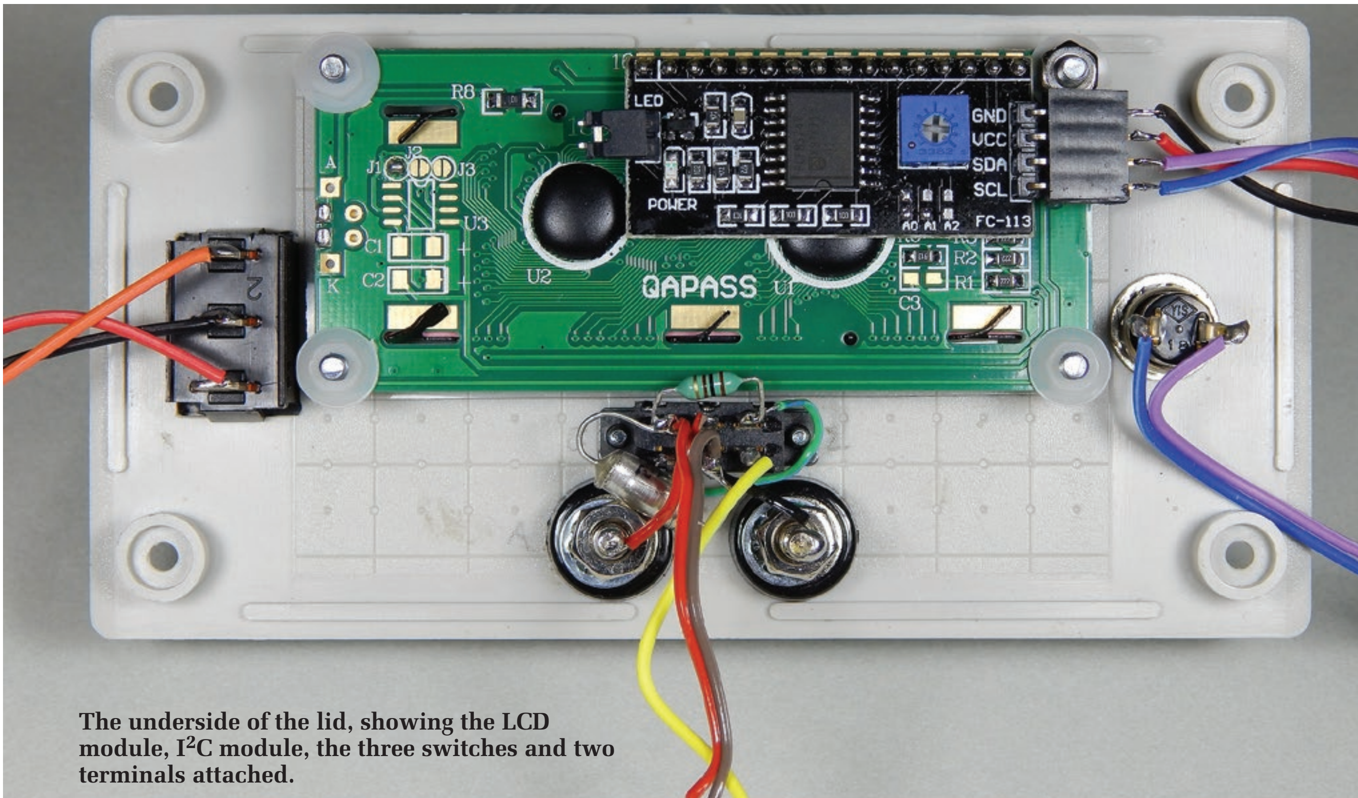

The underside of the lid, showing the LCD module, I²C module, the three switches and two terminals attached.

using four 15mm long M3 machine screws passing down through four 9mm long untapped Nylon spacers and fastened using four Nylon M3 nuts (under the module PCB).

With the LCD module in position, your front panel assembly is ready to be wired up and provided with its various leads to connect to the ProtoShield board.

Refer back to Fig.3 and the internal photo, following them carefully to make the correct connections between S1, the test terminal binding posts and L1 and C1 in particular.

Note that the leads of L1 and C1 should be kept as short as possible, to keep stray capacitance low (and stable). You can then make up the various short leads which will connect the front panel components to the ProtoShield board.

Note that the lead which connects S1, L1, C1 and the test terminals to the ProtoShield ends in a three-way SIL header socket, as does the lead from S3.

In contrast, the lead which connects to the serial LCD module has a four-way SIL header socket at each end, while the lead to connect zero/reset switch S2 (although of only two wires) ends in a four-way SIL header socket, with the wires connecting only to the pins on each end.

The two pins in the centre of the socket can be either cut short or pulled out, since they are not used.

Rather than using SIL sockets like we did on the prototype, we suggest you simply split a 40-way ribbon jumper cable with individual "DuPont" sockets on each wire.

This makes the job really easy; you simply pull off the required number of wires and then cut the cable to length and strip the free end, to solder to the switch or connector.

You don't even need to cut the cable for the LCD, you can just plug it in at both ends.

In each case, make sure each wire goes to the correct pin as with individual sockets, it's easy to get them out of order.

Having made up all the required leads, complete the LC Meter assembly with the following steps:

1. Mount the Arduino module inside the bottom of the box using four 9mm M2.5 screws and nuts, using four Nylon M3 nuts as spacers.
2. Plug the LC Meter ProtoShield into the Arduino, making sure you have all four SIL pin headers lined up correctly.
3. Holding the front panel assembly close to the top of the box and orientated correctly, plug the various connection leads into their matching pin headers on the ProtoShield. Be especially careful to get the correct connections between the ProtoShield and the LCD module, as shown in Fig.3.

Resistor Colour Codes

	No.	Value	4-Band Code (1%)	5-Band Code (1%)
❑	3	100kΩ	brown black yellow brown	brown black black orange brown
❑	1	47kΩ	yellow purple orange brown	yellow purple black red brown
❑	1	6.8kΩ	blue grey red brown	blue grey black brown brown
❑	1	4.7kΩ	yellow purple red brown	yellow purple black brown brown

Here's the alternative finish using a paper-printed label fixed to the outside of the UB3 Jiffy box lid, after it has been drilled and cut to suit. (You can, of course, glue a paper label to the laser-cut lid purchased from the SILICON CHIP online store). In this case the meter is measuring a nominal 100μH inductor and showing it is slightly high at 103μH.

4. Lower the lid assembly down into the box and fix it into place.
5. Program the Arduino, as described below.

Uploading the firmware

In order to do this, you need to have the Arduino IDE installed on your PC.

The latest version of the IDE can always be downloaded from the Arduino website (www.arduino.cc/en/Main/Software).

At the time of writing, the latest version is V1.8.2, dated 22/03/2017. There are various versions available to suit different operating systems: Windows (32-bit or 64-bit), macOS and Linux (32-bit, 64-bit and ARM).

After the IDE has been installed, download our firmware sketch for the LC Meter from the SILICON CHIP website (www.siliconchip.com.au). It's called "Arduino_LC_meter_sketch.ino".

Now plug your LC Meter into one of your PC's USB ports, using a suitable USB cable (usually with a Type A plug on one end, and a micro Type B plug on the other). You may need to install the correct USB VCP driver for it if this is not already installed.

If you're using a Freetronics Eleven module, you can download the appropriate driver from their website (www.freetronics.com.au). All of their drivers are zipped up in a file called "FreetronicsUSBDrivers_V2.2.zip", and there's also a document which explains how to install it.

Once the USB driver has been installed and your operating system confirms that it can communicate with the Arduino in your LC Meter, use Control Panel to find out which COM port the Meter's Arduino has been allocated (eg, COM5, COM7, or whatever).

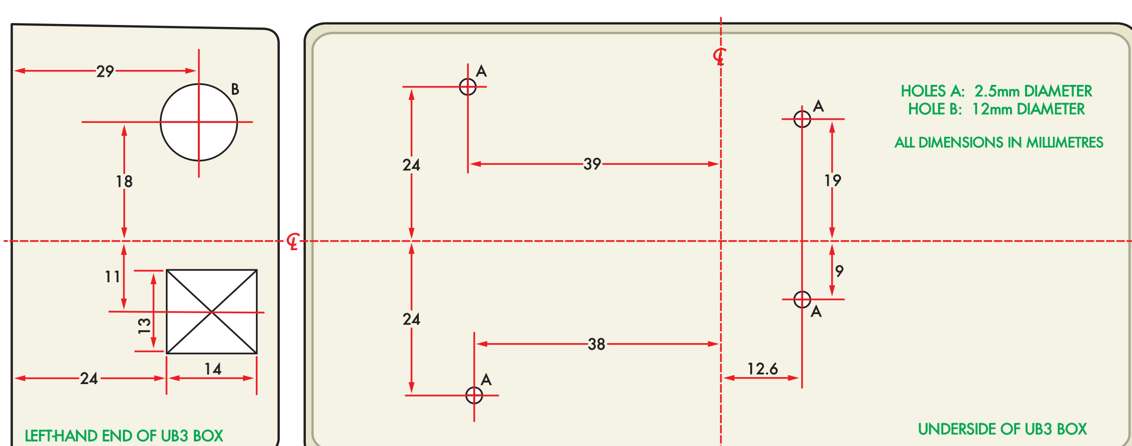

Fig.5: the drilling templates for the four Arduino mounting holes in the bottom of the box along with the USB and DC power access holes in the left-hand end.

Set the port for communication at 115,200 baud with the 8N1 “no handshaking” protocol. The COM port number should be entered into the Arduino IDE’s Tools->Port pull-down menu after you start it up.

Now open the LC Meter firmware sketch in the Arduino IDE, verify and compile it, and then upload it into the LC Meter’s Arduino flash memory. Soon after it has been uploaded, your meter should spring into life, flashing this message on the LCD screen:

```
Silicon Chip
Digital LC Meter
```

This should remain visible for two seconds, after which the screen should go black, before the Meter begins its initial zero calibration.

If you don’t see this initial message, this may be because the contrast trimpot on your LCD display module’s serial interface PCB is not set to the correct position.

The remedy is to swing open the lid of the box just enough to fit a very small screwdriver or alignment tool into the trimpot’s adjustment slot, turn it and then press switch S2 to force the Arduino to reset and start again.

Try changing the pot setting in one direction or the other until the message becomes clearly visible, pressing S2 after each adjustment.

Startup & calibration

At start-up, the Meter normally expects slider switch S1 to be set in the Capacitance (C) position, and no external capacitor to be connected to the test terminals. If you have done this it will now display the message:

```
S1 set for C: OK
Now calibrating
```

But if you have set S1 in the Inductance (L) position instead, you’ll see a different message:

```
Fit shorting bar
Now calibrating
```

This means that the Meter has detected that S1 is set to the L position, and is assuming that you want to do the zero calibration in this mode. As a result, it’s advising you to fit a very low inductance shorting bar between the test terminals. This can be in the form of a 40mm long piece of 1.66mm diameter copper or brass rod between the terminals, or (better still) a 40 x 30mm rectangular piece of 1mm thick copper or brass sheet between them. In either case, the rod or sheet must be shiny rather than oxidised.

Note that if you have set S1 to the L position accidentally and don’t have a shorting bar available, there’s no harm done. Simply flick S1 to the C position and then press switch S2 to get the Meter to reset and begin over again.

Or if you do want to calibrate in inductance mode, simply fit the shorting bar between the terminals (if you haven’t already done so) and then press S2 to reset and begin over again.

In either case, there will be a brief pause after which the meter will show the values for C1 and L1 it has found from the initial calibration. This will be something like:

```
C1 = 1084.2 pF
L1 = 91.24 uH
```

The actual values displayed will depend on the components in your unit, as well as the stray capacitance and inductance. They’re shown at this stage mainly as reassurance that the Meter is working correctly. The measured values of C1 and L1 will be displayed for three seconds, after which this message will appear:

```
Calibration done
Ready to measure
```

Credit where it’s due

As mentioned earlier, this Digital LC Meter, like our earlier May 2008 design, is based on a 1998 design by the late Neil Heckt, of Washington, USA.

Since then, various people have produced modified versions of the design, including Australian radio amateur Phil Rice VK3BHR, of Bendigo in Victoria. Mr Rice and others also modified the firmware and adapted it to use the PIC16F628 micro with its internal comparator. They also added a firmware calibration facility.

So a significant amount of credit for this latest version of the design must go to these earlier designers. The author is happy to acknowledge their earlier work.

This will display for one second, after which the Meter will begin making measurements. If you have done the initial calibration in C mode and S1 is still in this position but no unknown capacitor is as yet connected to the test terminals, you should now get a display like this:

```
Cx = 0.004 pF
(F3 = 515838 Hz)
```

where the value shown for C_x is very close to zero, while the frequency F3 shown on the second line is for the current oscillator frequency; essentially the same as F1 at the current ambient temperature.

The Meter’s oscillator frequency does drift a little with temperature. This means that after a while, the value shown for Cx with no external capacitor connected may creep up from the almost-zero reading you get initially. At the same time, the reading for F3 would slowly decrease.

If you find the value shown for Cx

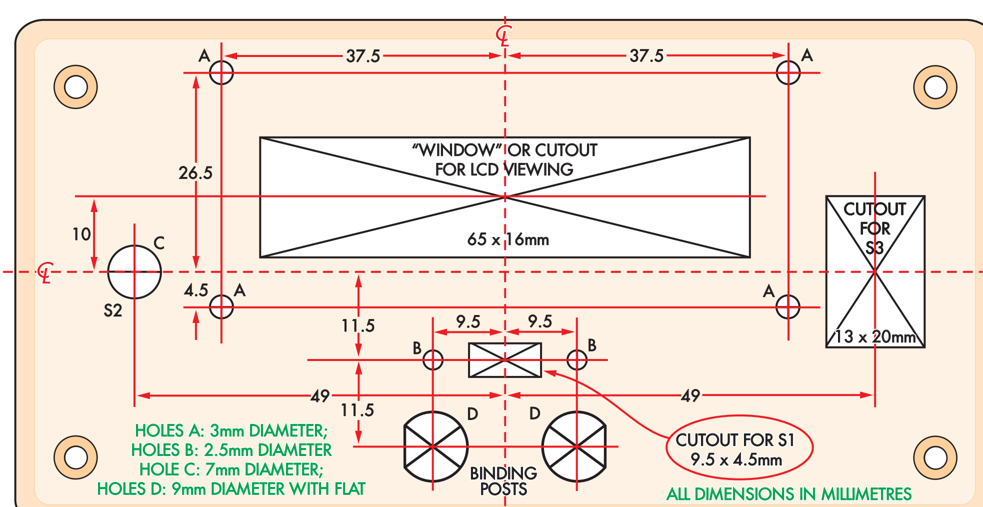

Fig.6: you can either drill and cut the twelve cut-outs required in the lid supplied with the UB3 “Jiffy” box, as shown in this diagram, or (much easier!) purchase a laser-cut acrylic lid from the SILICON CHIP online store and use that instead of the lid that came with the box.

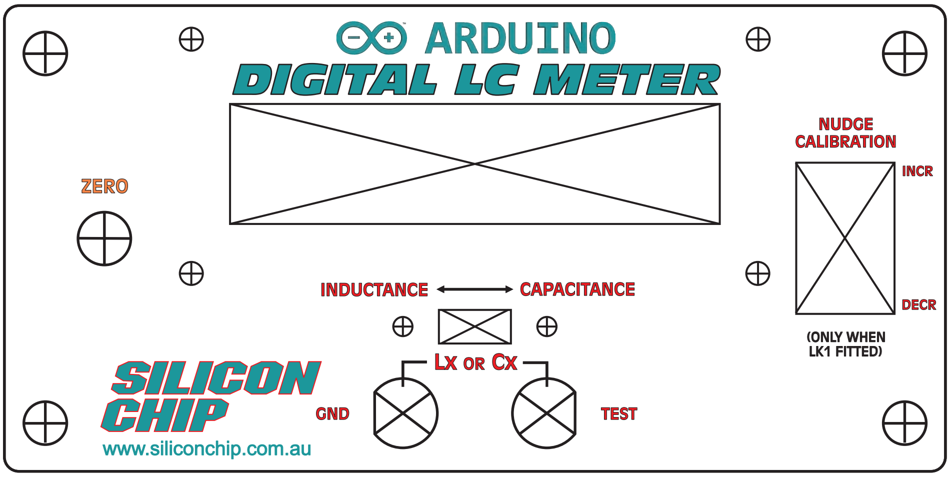

Fig.7: same-size front panel artwork designed to fit a UB3 Jiffy Box. It will also fit the laser-cut acrylic front panel from the SILICON CHIP online store. This, along with the two cutting/drilling diagrams, can also be downloaded (as a PDF) from www.siliconchip.com.au

with no external capacitor has crept up to 0.1pF or more, simply press S2 again to get the Arduino to perform a new zero calibration.

On the other hand, if you've done the initial calibration in L mode and S1 is still in this position but the shorting bar is still connected across the terminals, you should get a display like this:

```
Lx = 0.002 uH
(F3 = 516615 Hz)
```

The value shown for Lx is again very close to zero, and the frequency F3 shows the current oscillator frequency, again very close to F1 at the current ambient temperature. Now if you remove the shorting bar in this mode, you'll find the display will change to something like this:

```
Over Range!
(F3 = 2 Hz)
```

This simply shows that in this mode, an open circuit between the terminals is equivalent to a very high inductance, because it causes the oscillator frequency to drop to near zero.

When you connect a real inductance between the test terminals, the Meter will measure its inductance and display it (assuming its value is within the Meter's range, which is from 10nH to 150mH).

As before, drift in the Meter's oscillator may cause the Lx reading for the shorting bar to creep up gradually. So before making a particularly critical measurement, it's a good idea to fit the shorting bar between the test terminals and press S2 again to force the Arduino to reset and perform a new zero calibration.

Optimising accuracy

If all is well so far, your Digital LC Meter should be operating correctly and ready for use. If you have been able to procure a couple of 1% tolerance (or better) capacitors for C1 and C2, it should also be able to deliver that order of accuracy without any extra calibration.

But as mentioned earlier, it is possible to achieve even better accuracy with the meter providing you have access to a reference capacitor whose value is accurately known (because you've been able to measure it with a high-accuracy LCR meter).

Ideally, this reference capacitor should have a value of between 10nF and 100nF, but even one with a value between 1nF and 10nF would be OK.

This is achieved by tweaking or "nudging" the Meter's reading for the reference capacitor using switch S3. Here's how you do it:

1. Remove the 5V supply from the Meter
2. Lift the lid/front panel up from the box and carefully fit the jumper shunt over the pins for LK1, down on the ProtoShield.
3. Close the box and slide S1 to the C position but don't connect your reference capacitor to the test terminals.
4. Re-apply the 5V power and let the Meter go through its initial zero calibration.
5. Wait a couple of minutes, watching the reading for Cx to see if it drifts up appreciably from the initial near-zero figure. If it does, press switch S2 to force a reset and bring the reading back to less than 0.01pF.
6. Connect your known-value capacitor to the test terminals and note the Meter's measurement reading. It should be fairly close to the capacitor's known value, but may be a little higher or lower.
7. If the reading is too low, press the rocker of switch S3 at the upper ("=") end for a second or so; if it's too high, press the lower end ("-") instead. The reading should change by about 0.5%. Continue until the reading is as close as possible.
8. Remove power, open the lid and remove the jumper from LK1.
9. Re-attach the lid.

Note that since the Arduino always saves the revised calibration factor in its EEPROM after every measurement during this nudging procedure, so you only have to do the calibration once.

Also, when you calibrate the meter in this way using a known value capacitor, it's also calibrated for inductance measurements too. SC

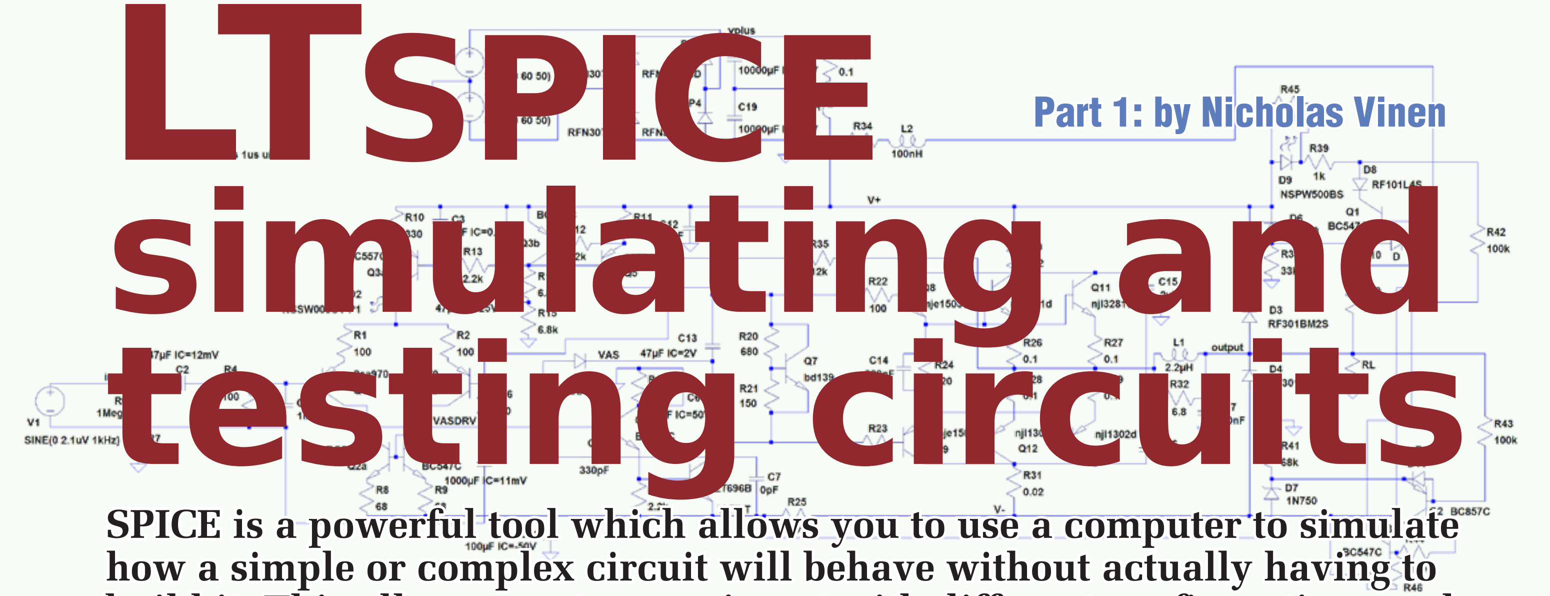

LTSPICE simulating and testing circuits

Part 1: by Nicholas Vinen

SPICE is a powerful tool which allows you to use a computer to simulate how a simple or complex circuit will behave without actually having to build it. This allows you to experiment with different configurations and examine the internal operation of a circuit before building it, saving you time and effort.

In this series of articles, we'll take you through installing and using LTspice, a free, easy-to-use and yet very powerful circuit simulation package.

Once you're familiar with LTspice, you can draw up a circuit and start simulating it. Testing circuits in LTspice is a lot cheaper and safer than building them – if you blow up components in LTspice you don't have to buy new ones! Just modify the circuit and try again.

Besides just figuring out whether a given circuit will do what you expect, you can also use SPICE (which stands for **S**imulation **P**rogram with **I**ntegrated **C**ircuit **E**mphasis) to determine certain performance parameters such as stability, efficiency, distortion, noise, reaction time, overshoot, frequency response, power consumption and dissipation, and so on. Throughout this series we'll show you examples of how to calculate all of these parameters.

While SPICE isn't perfect and may sometimes fail to simulate some complex analog circuits reliably, it is quite surprising how close the results of simulations can match the real-world behaviour of a circuit. Note that accurate simulation does rely upon accurate component models and these are not always available.

Simulating a circuit starts with drawing it. During this process you will place component symbols on a sheet and "wire them up".

You will then need to tell the simulator the type code of each component so that it can select an appropriate model. In many cases, for components like resistors, capacitors and inductors, totally realistic behaviour is not terribly important and you can simply use a default "ideal" component.

To get accurate results with devices like transistors and diodes, you would be better off picking one of the available component models which exactly matches the part you intend to use, or at least has similar characteristics. We'll discuss this aspect in more detail later.

Installing LTspice

We're going to use LTspice for Windows in this tutorial series because it's free, easy to install and use and most importantly, is supplied with a fairly large and mostly complete library of component models so that you can get up and running right away.

A component model defines its characteristics. For example, each type of transistor has a different curves for V_{BE}, V_{CE}, h_{fe}, maximum voltage and current and so on. The model provides coefficients so that the simulated component behaves similarly in these respects to an average, real component.

To start off, download the latest version of LTspice from www.linear.com/designtools/software/ It's available for 32-bit or 64-bit Windows 7, 8 or 10; there is also an older version available for Mac OS X 10.7+. Simply download the executable file, run it and follow the prompts to install it. It's a straightforward process. Once installed, run the program and you will see a blank window like in Fig.1.

Now select the "New Schematic" option from the "File" menu. Not much will appear to have changed

Toolbar Icons

Concentrate your gaze on the right-most section of the toolbar, blown up in Fig.1. From left-to-right, the buttons are:
Wire – connects two or more component pins
Ground – place a ground (0V) symbol on the circuit
Label Net – assign a name to a "net" (more on that later)
Resistor – place a resistor in the circuit
Capacitor – place a capacitor in the circuit
Inductor – place an inductor in the circuit
Diode – place a diode in the circuit
Component – place something else in the circuit, such as a transistor, IC, regulator, voltage or current source, etc
Move – move something around in the circuit diagram
Drag – same as Move, but keeps any wire connections to the selected component(s) intact
Undo/redo – revert the last change to the circuit, or reinstate it
Rotate component – rotates the selected component/components by 90°
Mirror component – flips the selected component horizontally
Text – add text to the circuit diagram
SPICE Directive – add an instruction to the circuit diagram which tells the simulator how to behave

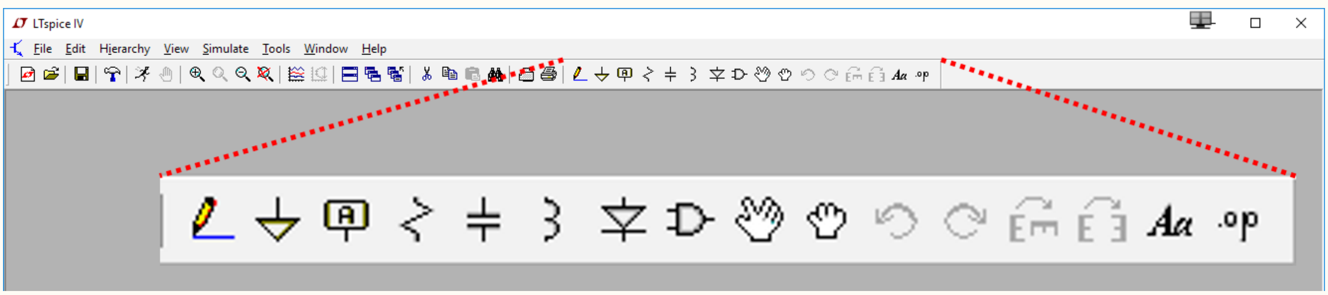

Fig.1: how the LTspice window looks just after creating a blank simulation. The toolbar at top has been blown up to show the important buttons, which are (from left-to-right) Wire, Ground, Label Net, Resistor, Capacitor, Inductor, Diode, Component, Move, Drag, Undo, Redo, Rotate, Mirror, Text and Spice Directive.

but you are now ready to start drawing your circuit. First though, it's best to give it a name. Select "Save As" under the "File" menu, type in "tutorial1" and press Enter.

Chances are that it will say that you don't have permission to save the file into the "C:\Program Files" directory and it will ask if you want to save it in the User folder instead. That's a good idea, so say Yes and then press Enter again to save your file.

We'll now draw up a simple mains power supply circuit. But first, let's look at the toolbar at the top of the window. This is important since you will be using these buttons a lot. The description of each icon in the toolbar is under "Toolbar Icons" at the top right of the next page.

You'll find that you will need to use nearly all these icons when drawing up the circuit you want to simulate. We're going to start by creating a source of 230VAC.

Click on the Component button (which looks like a logic gate). You will then be presented with a list of components and folders (which are in square brackets). Scroll to the right and click on "voltage", then "OK" (or just double-click "voltage").

Click somewhere in the blank circuit to place your first voltage source. This will be simply shown as a circle with positive and negative symbols inside, corresponding to the two output terminals. Note that a voltage source will always take the same form, whether it is intended to produce AC or DC.

Now right click your mouse or press escape, since we only want one voltage source for now.

This is one of the most fundamental parts of a circuit to simulate; the voltage source can generate AC, DC, both AC and DC or a function such as a sinewave or pulse train and is used to feed other components in the circuit. Voltage sources can be combined in various ways.

Voltage source mode setting

There are three different kinds of voltage sources and we need to use the right one to simulate 230VAC mains; refer to the panel titled "Simulation Types" for an explanation.

Having read that, right-click on the V1 element you have placed and then click Advanced. You can now select SINE from the list on the left, and enter 0V for DC offset, 325V for Amplitude (this is the peak value; not RMS), 50Hz for the frequency and leave the rest blank.

Units in these values are optional, however, for clarity it's usually best to include them.

Click OK and the circuit updates to include these parameters. Now click on the Ground button in the toolbar and place a ground symbol directly below the "negative" end of your voltage source.

You need to define 0V somewhere in the circuit if you want to simulate it and this (effectively, the incoming Neutral line), is as good as anywhere. As before, right-click your mouse or press escape to stop placing components.

Now use the Wire tool (leftmost on the section of the toolbar described above) to draw a wire between the negative end of the voltage source and the ground symbol. Click at one end, then the other, then right-

Simulation Types

There are two common types of simulation you can perform, plus several other less common types. The two most common types are "transient" and "AC".

A transient simulation is essentially equivalent to hooking an oscilloscope up to various points in the circuit and then freezing its display to examine how the voltages and currents vary over time. An AC analysis is more like connecting a spectrum analyser with tracking generator up to a circuit.

AC voltage sources in SPICE are primarily useful for AC analysis. For transient analysis, you need a combination of DC voltages or "function"-based voltage sources which are generally one of the following: PULSE, SINE, EXP (exponential), SFFM (single frequency FM) or PWL (piecewise linear).

Basically, if you want an AC voltage source in a transient analysis, you use the SINE function. If you try to use an AC voltage source in this situation, you'll find it won't do anything useful.

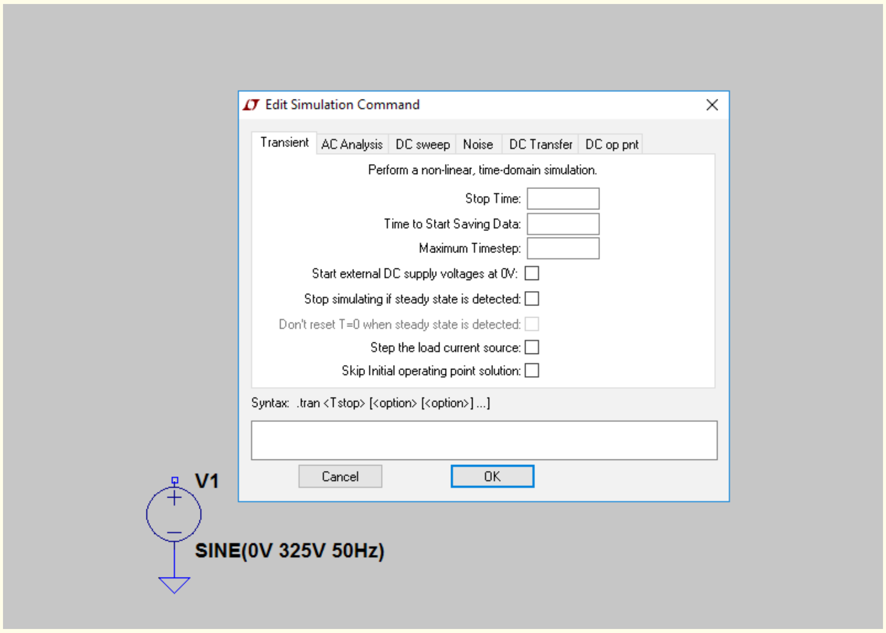

Fig.2: the Edit Simulation Command dialog comes up the first time you select the Run option from the Simulate menu. Select the simulation type from the tabs at the top and then fill in the details below. For a Transient analysis, the most important ones are: Stop Time; Time to Start Saving Data; and Skip Initial operating point solution.

click or press escape to stop drawing wires. Note that if at any point you make a mistake, you can press F9 or click the Undo button on the toolbar to revert to the previous state.

Now we can run the simulation for the first time. Select the "Run" option under the "Simulate" menu. As this is the first time, you will need to set up the simulation conditions, using the dialog which appears (see Fig.2).

"Transient" is the default simulation mode (tab) selected so all you need to do is enter a Stop Time (let's use 100ms) and then click OK. A SPICE Directive automatically appears on the circuit, which reads ".tran 100ms", and you will find a black box appears at the top half of the screen, with the circuit shrinking below.

This is our virtual scope display. Move your mouse cursor down to hover just over the little square box at the positive end of the voltage source in the circuit diagram below and the mouse cursor should change to look like a probe. Click there and you should get a display like Fig.3.

This shows our simulated mains voltage. Of course, the real mains sinewave is nowhere near as clean as this but it's a good start!

Note the text reading "V(nc_01)" at the top. This indicates that the green trace is showing the voltage at the node labelled "nc_01" which is a name automatically generated for this part of the circuit, as we have not provided our own name yet.

Hold down the CTRL key on your keyboard and click on this label. You will get a dialog box showing information about the "trace" including the start and end times, the average (which is very close to zero, as it should be) and the RMS value which is just under 230VAC; exactly what we wanted. You can now dismiss this dialog.

By the way, if you want to change the parameters later, you can right-click on the ".tran" directive to re-open the simulation dialog.

Building the circuitry

Note that if you already know how to build a circuit in LTspice, you can download the tutorial1.asc file from the SILICON CHIP website and skip to the next cross-heading.

If you find yourself confused by the following instructions, refer to Fig.4 to see how the finished circuit looks.

Let's start by adding a capacitor connected to the 230VAC "positive" terminal (effectively mains Active). Click somewhere inside the circuit diagram, then click the Capacitor button in the toolbar and place the capacitor above the voltage source. Right-click the capacitor to set its Capacitance value to 470nF. Set the voltage rating to 400V (peak) at the same time and the RMS Current Rating to 250mA.

Use a similar process to add a resistor to the right of that capacitor and set its value to "10Meg". Note that one of the traps when using SPICE is that "10M" would be interpreted as "10m" (ie, 10 milliohms) so you need to write it with "Meg" on the end. You can set the tolerance to 5% and power rating to 1W at the same time.

Now use the Wire tool to wire the two components up in parallel and connect the common bottom end to the voltage source. Add a second resistor, in series with the capacitor/resistor combination, and set its value to 470 (ohms), tolerance to 5% and power rating to 1W.

The next step is to add two diodes to form a half-wave rectifier. Click on the Diode tool in the toolbar, then move the mouse down into the circuit. You will notice that if you place it, its cathode will face towards the bottom of the circuit but we want it at the top.

So before placing it, move the mouse back up to the toolbar and click the rotate button twice (note that this button will be disabled before moving the mouse down into the circuit area, so after clicking the diode button, you need to move it down and then back up).

Now place the diode above and to the right of the existing components. Right-click the diode symbol, which is currently configured as a generic (ideal) diode, and click the "Pick New Diode" button. You will now get a list of the diode models built into LTSpice, which includes silicon/switching/Rectifier (standard) diodes, fast recovery diodes, schottky diodes, zener diodes, LEDs and transient voltage suppressors (TVS/varactor).

Scroll down to where the "silicon" type diodes are listed and click on the MURS120 which is roughly equivalent to the 1N4002, then click OK. If the placement of the diode is not ideal, click the "Move" button in the toolbar (or press F7 on the keyboard) and click on D1 to move it to a better spot.

Now we need a second, identical diode so the easiest solution is to clone the one we have. Press F6 on the keyboard, then click on D1 and place the new diode (D2) directly above it. Join the adjacent anode and cathode pins, then connect the free end of the 470Ω resistor to this junction, all using the Wire tool. Connect the free anode at the bottom to ground, as we did with voltage source V1.

Now we need a zener diode. You can clone one of the two existing diodes, placing it immediately to the right of voltage source V1, then right-click on and select "Pick New Diode" to change its type. Scroll down to the zeners and you will find multiple 15V zener diodes in the list (look for 15 in the Vbrkdn(V) column). Pick the KDZ15B as this is a 1W type, then click OK. Move D3 if necessary, to avoid labels from overlapping.

Now connect the zener's anode (bottom end) to ground and the cathode (top end) to the free cathode of the rectifier diode above. Having done that, add a 220µF 25V capacitor in parallel with D3, with a 500mA ripple current rating and ESR of 0.1 (ohms). Also add a 1.5kΩ 10% 5W resistor, simulating a power supply load, in parallel with both. When finished, your circuit should look similar to that shown in Fig.4.

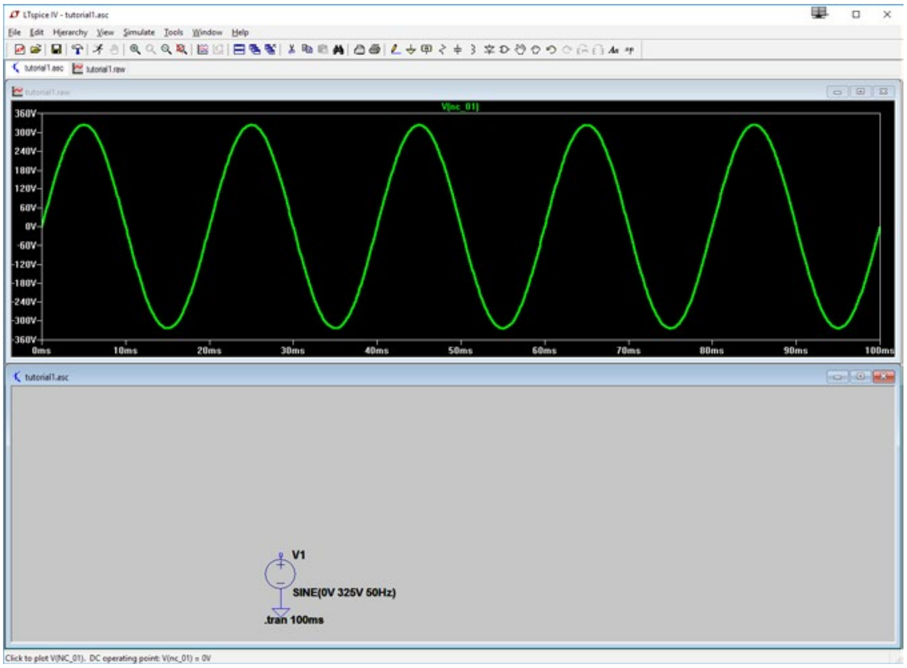

Fig.3: the result of our first Transient simulation, showing the voltage at the top of voltage source V1 over a 100ms period. Note that the 325V figure selected defines the peak voltage, not RMS and that several parameters have been left blank and so default to zero, including the DC offset and phase values.

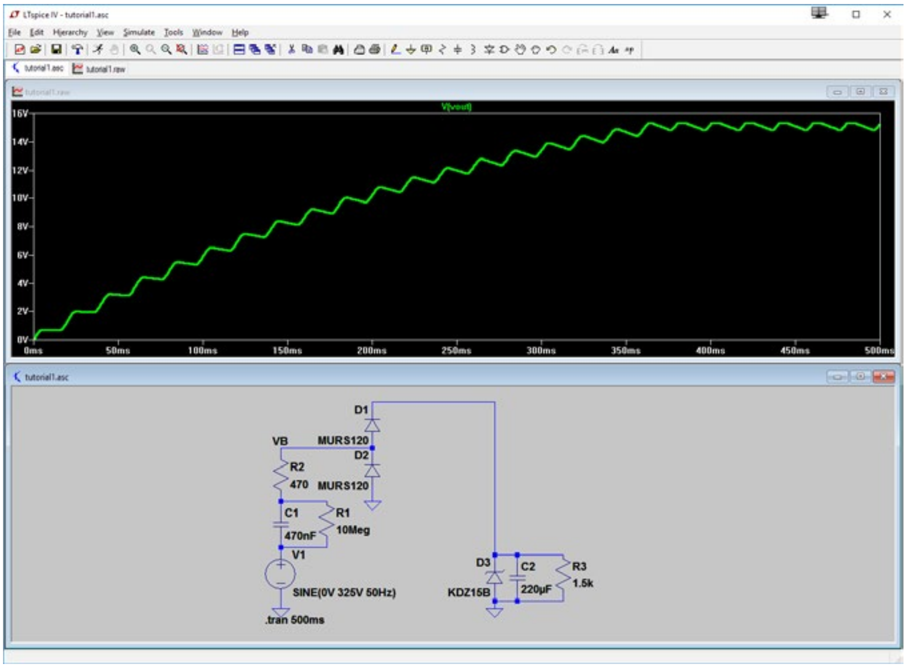

Fig.4: now we've built up a basic mains power supply with a simple resistive load and can observe how the main 220µF filter capacitor charge increases every 20ms during the peak of each mains cycle. We can see that D3 (a 15V zener) begins to conduct after around 350ms, but some ripple remains.

Making some measurements

Right-click on the ".tran 100ms" directive and change the Stop Time to 500ms, then re-run the simulation ("Run" option under the "Simulate" menu). Click on the "wire" at the cathode of D3 to view the resulting voltage. Your result should be the same as shown in Fig.4.

As you can see, it takes around 370ms from the application of mains

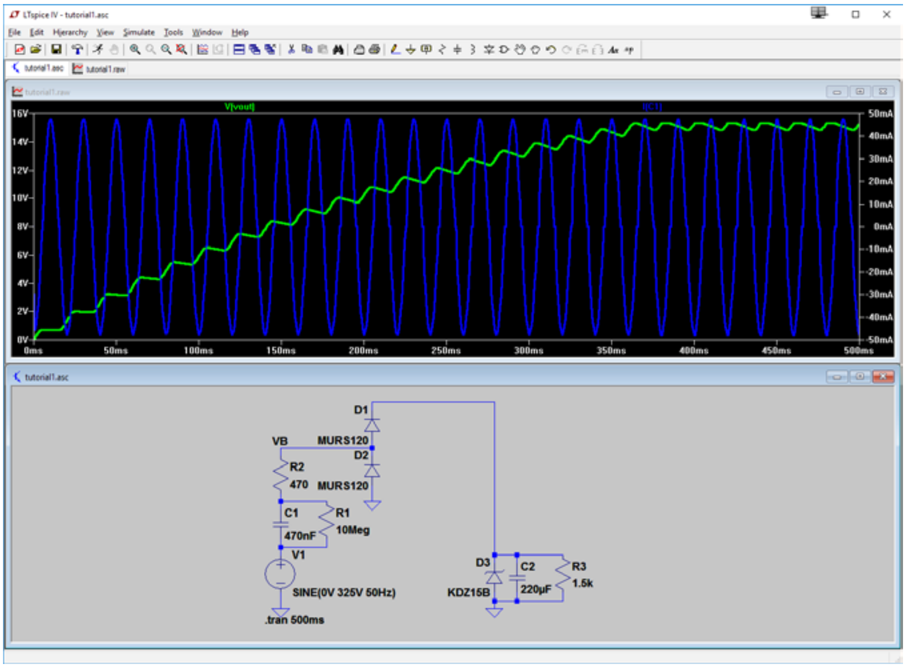

Fig.5: not only can we see the voltage across C2 but now we can also observe the current drawn from the mains as it charges – all without having to wire up a single component and without any test equipment! One of the benefits of using SPICE is how easy it is to make multiple voltage and current measurements.

power before the 220µF capacitor is fully charged to 15V. You can drag a box around the waveform at the top of the screen to zoom in and examine it in more detail (right-click and select "Zoom to Fit" or press CTRL+E to go back to the normal view).

Once zoomed in, you can see that the peak voltage across the capacitor is clamped to around 15.35V and with the 1.5kΩ load, the minimum voltage is around 14.85V, giving a ripple of around 0.5V.

You can make reasonably accurate measurements by placing the mouse cursor over the trace and then reading the time and voltage values shown in the bottom-left corner of the LTspice window. Also, once you've zoomed in, if you CTRL-click the V(n001) text at the top of the screen, it will calculate average and RMS values for the time period displayed, in this case, both around 15.124V.

Now click the mouse in the circuit window at bottom and move the cursor over capacitor C1. You will note that the cursor changes to what looks like a clamp meter. Click here and the current through this capacitor will also be shown in the top window. Note that it is essentially symmetrical and looks like a sinewave with zero-crossing artefacts.

Note also that a new y-axis appears on the right-hand side of the plot, allowing you to see that the peak current through C1 is just below 50mA. You can CTRL-click the label at the top of the display to read off the RMS current which is 33.5mA (see Fig.5).

Efficiency calculations

The efficiency of this circuit is the power delivered to the load (R3) divided by the power drawn from the mains (V1). In both cases, we can compute power as V × I. We could use $V^2 \div R$ for R3 but then we could need to change the calculation if we changed the value of R3, and it would also make it harder to change the circuit to a more realistic load.

To make it easier to calculate both power figures, let's label the two voltages. Click the "Label Net" button in the toolbar and type in "VIN", then press OK. Place the label at the junction of V1, C1 and R1.

Similarly, label the junction of D1, D3, C2 and R3 as "VOUT". Press the DEL key on your keyboard and click on the labels at the top of the simulation output to delete the traces, then re-run the simulation.

Now right-click on the (now blank) top half of the window and select "Add Trace" (or, having clicked in this sub-window, press CTRL-A). It will prompt you for "Expression(s) to add". Type in "V(VIN) * -I(V1)" and click OK. A new trace will appear showing the instantaneous power being drawn from V1. V(VIN) refers to the voltage at the node labelled VIN and I(V1) refers to the current through voltage source V1. "*" is the multiplication operator so giving us the product of the two.

The minus sign before I(V1) just sets the polarity of the result and is something you'd normally need to determine experimentally. You will see that the instantaneous power goes positive and negative at different times in the mains cycle.

This is because sometimes, current flow into C1 is in-phase with the mains voltage and sometimes it is out-of-phase. In other words, there are times when power is flowing from the mains into C1, and times when it is flowing out of C1 and back into the mains.

If you CTRL-click the expression at the top of the window, you will see that the average is 712.21mW and its integral (ie, total energy consumed in the 500ms window) is 356.11mJ.

But note that this includes the time that C2 is charging. So to get an accurate result, right-click on the ".tran 500ms" directive and change the "Time to Start Saving Data" to 400ms, then re-run the simulation. The average is now 783.93mW, which represents a steady-state value, and you will notice that the waveform is consistent across the five mains cycles (100ms) shown.

By the way, if you want to change the expression used to plot the power, you can do this by right-clicking where it's shown at the top of the window.

Now, to compute the power consumed by R3, right-click in the top window (or press CTRL+A) and enter the similar expression "V(VOUT) * I(R3)". If you CTRL+click the new expression which appears at the top of the window, you will see that the average power is 152.69mW (see Fig.6). This is in line with what you'd expect from 15V across a 1.5kΩ resistor ($V^2 \div R$ = 15 x 15 ÷ 1500 = 150mW).

So we can calculate the efficiency as 152.69mW ÷ 783.93mW = 19.5%. That's pretty lousy! That means that 80.5% of the energy drawn from the mains (630mW or so) is being dissipated elsewhere in the circuit, just turned into useless heat. Luckily, we can use LTspice to figure out where and improve the situation.

First, let's see how much power is

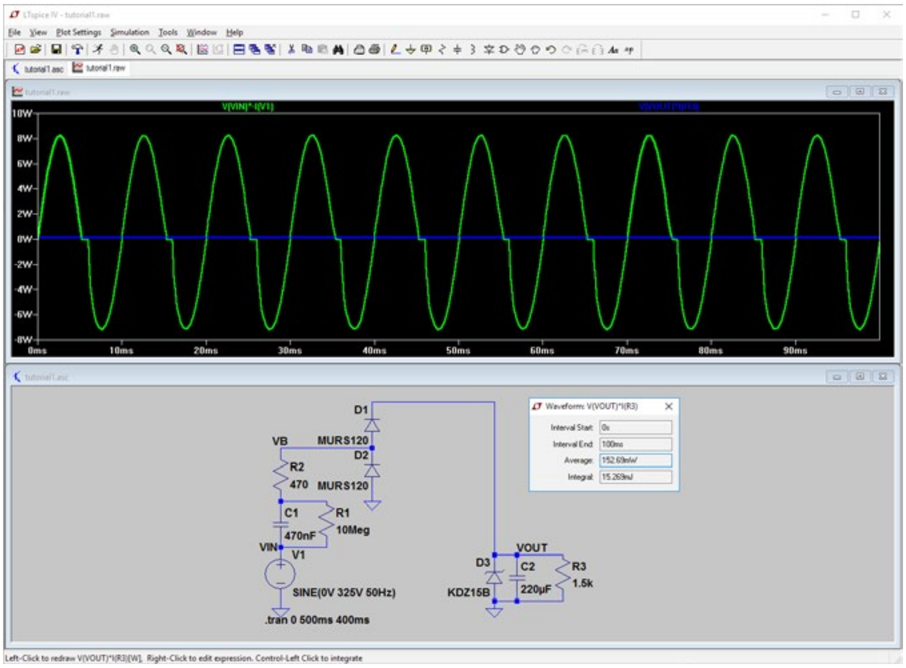

Fig.6: plot of the product of the input voltage and current; LTspice automatically shows the result in watts and changes the Y-axis to suit. The area enclosed by the power curve below the horizontal axis is smaller than that above, with the net power consumption shown in the average (in the box to the right of the circuit).

dissipated in D3, the zener clamp diode. We can simply plot the expression "V(VOUT) * I(D3)" and integrate it as before, to yield a figure of 73.282mW. Well, that's barely more than 10% of the energy being wasted, so that isn't the culprit; we may still be able to make some tweaks to reduce this figure and improve efficiency but let's figure that out later.

What about R2? To calculate the voltage across that, we need to label the wires (nets) at both ends. Let's label the one junction of C1/R1/R2 as "VA" and the junction of R2/D1/D2 as "VB".

We can then plot the expression "(V(VB) – V(VA)) * I(R2)", in other words, the difference between the voltage at points VB and VA (ie, the voltage across R2) times the current through R2.

Integrating this gives us a figure of 529.33mW. Adding this to the power dissipated in D3 gives a result of 602.6mW, explaining over 95% of the power lost in the circuit (the other ~5% is probably in R1). So to improve the efficiency we need to do something about R2.

Improving the efficiency

R2's purpose is to reduce the inrush current into C1 when the circuit is first connected to the mains, especially if that happens to be in the middle of a cycle. If we reduce R2's value, that will reduce its dissipation and improve the overall efficiency but we need to check that this won't cause any problems and also quantify just how much of an improvement we can achieve.

So let's simulate the (almost) worst case, where the circuit is connected to the mains at the peak of 325V and C1 is discharged, and see how low we can make the value of R2 before we risk damaging something.

To do this, rightclick on the body of V1 and enter 90 for "Phi(deg)". We also need to make two changes to the simulation directive, which we can access by right-clicking on the ".tran 400ms 500ms" text.

First, change the "Time to Start Saving Data" back to 0ms so that we can see the initial conditions, then also tick the "Skip Initial operating point solution" box towards the bottom. This tells the simulator to start with all capacitors and inductors fully discharged (although you can specify an initial charge on a case-by-case basis if necessary; we'll explain how to do this in a future instalment).

Re-run the simulation, clear all the traces and plot the current through C1; you can achieve the latter two simply by moving the mouse cursor over C1 until it turns into the clamp symbol and then clicking twice. The first time it will show the current plot for

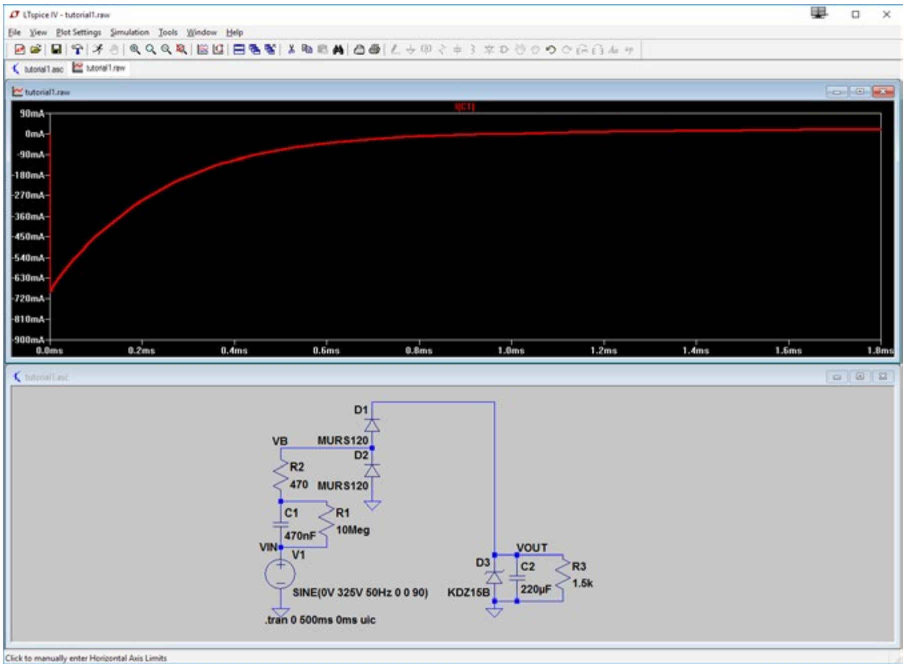

Fig.7: by zooming into the early part of the current trace for C1, we see the inrush current is around 700mA for a fraction of a millisecond. The "uic" on the end of the ".tran" directive is critical; it stands for "use initial conditions" and without it, capacitors and inductors start in a "steady state" condition.

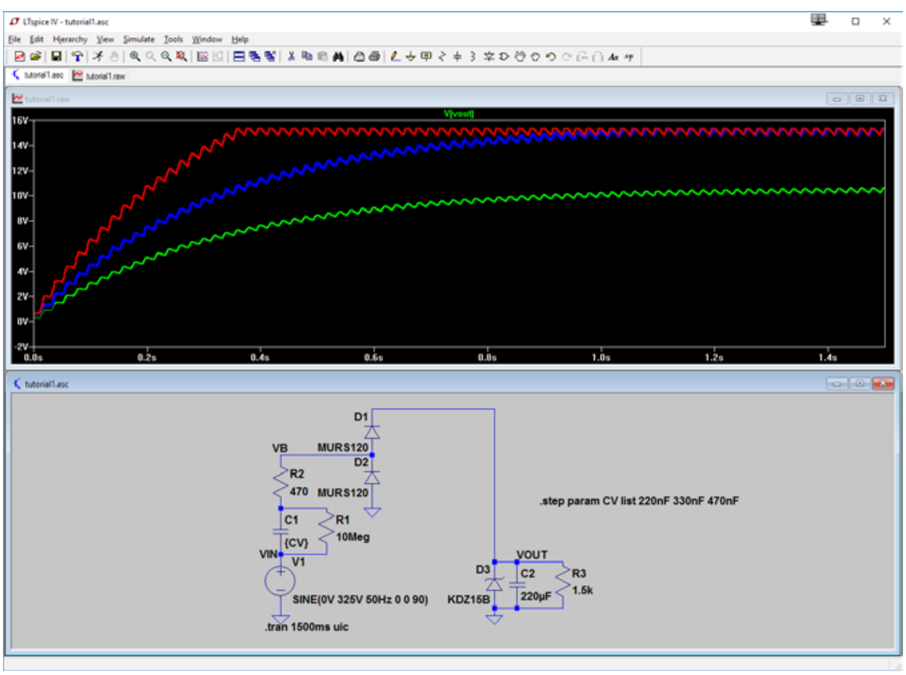

Fig.8: parameter stepping is a valuable method for optimising component values. Here we can see how varying the value of C1 between 220nF and 470nF affects circuit operation. You can also use this method to vary the simulated ambient temperature or to see how component tolerance affects circuit operation.

C1 and the second time, it will erase all the other traces except for that plot.

If you zoom into the first few milliseconds you can see that the peak current is around 700mA but this drops very rapidly, to just a few milliamps after 1ms or so (see Fig.7). In retrospect, we could have calculated the 700mA figure simply by assuming that C1 is initially a short circuit and doing the calculation 325V ÷ 470Ω = 0.7A. This suggests that whatever we do to reduce the value of R2 is inevitably going to increase the inrush current but the simulation shows that this is really very brief as C1 rapidly charges up.

Consider that in a real circuit, this would be an X2 capacitor which is designed for direct connection across the mains supply with no real current limiting whatsoever so it should be able to tolerate a high inrush current. So on that basis, let's reduce R2 to 68Ω, giving an inrush current of just under 5A.

The only other components which need to handle this current are R2 (which should be OK given how brief the spike is) and D1/D2 (which will handle much larger spikes as long as they're short or non-repetitive).

At the load end, how much of the initial spike will be borne by D3 and C2 depends on the polarity of the applied mains voltage (ie, whether D1 or D2 conducts) and C2's ESL (equivalent series inductance). Typical ESL of a moderately-sized electrolytic capacitor appears to be pretty low at around 1nH so C2 should safely absorb the brief initial spike, but even if it doesn't, it should not pose much difficulty for D3.

We can now re-run the simulation, adjusting the time to start saving data back to 400ms and calculate the steady-state figures as input power: 327mW, output power: 152.7mW, efficiency: 46.7%. That's a lot better but still not great.

Let's look again at the power consumed in D3, the zener diode. It's virtually identical to before at 73.75mW but now this is around 50% of the power loss. We can reduce this by lowering the value of C1, so that it doesn't deliver more current than the load requires and D3 will then only conduct rarely (eg, if the mains voltage is higher than nominal or the load is lighter than expected).

Parameter stepping

Now we consider whether changing the value of C1 will affect efficiency. It will because if the value is too high, D3 will shunt more of the current coupled through it, effectively wasting power whereas if the value of C1 is too low, the voltage across D3 will not rise to the desired value of ~15V.

What we really want to do to figure out the ideal value is look at the effect of changing the value of C1 with everything else the same. We can do this by stepping its capacitance through different values. To do this, click on the SPICE Directive ("op") button in the toolbar and then type in ".step param CV list 220nF 330nF

470nF". This creates a parameter called "CV" which steps through three different capacitance values.

Now change the value of C1 from 470nF to {CV}. Re-run the simulation, with a start time of 0ms and finish time of 1500ms and plot VOUT. The result is shown in Fig.8.

Unfortunately, LTspice doesn't provide a colour-coding legend but it's fairly obvious that the green curve is for C1=220nF, blue for C1=330nF and red for C1=470nF. 220nF is too low as VOUT doesn't even reach 10V, while with both 330nF and 470nF it reaches the same final voltage, albeit after a different time delay.

So it seems that 330nF is probably close to the ideal value. Let's set the capacitance value of C1 back to 330nF, delete the step directive (press DEL on the keyboard, then click on the directive) and then re-run the efficiency calculations.

Final results

After changing the "Time to Start Saving Data" back to 1400ms and using the same steps as before, we can now compute the input power as 219mW and the power consumed by the load at 151.94mW, only a tiny bit lower than before, giving an efficiency figure of 69.4%. That's pretty reasonable for such a simple circuit, and with a virtually identical load voltage.

So we've barely sacrificed any performance for what is a pretty large improvement in efficiency, all thanks to the ease of simulating such a circuit. Compare this to the difficulty of measuring it, especially when you consider it would be directly connected to the mains!

Apparent power consumption

There are a couple of final issues to discuss regarding simulating this circuit. Firstly, our method of integrating the instantaneous power gives us the real power consumption of this circuit, as would be measured by your power meter (and which would be used to charge you for electricity).

But note that the RMS current drawn from the "mains" (V1) is now 23.65mA with an RMS voltage of 230VAC. That gives an apparent power consumption of 0.02365A x 230VAC = 5.44W.

That tells us that this circuit has a very low power factor. In fact, we can calculate it, it's simply the real input power of 219mW divided by the apparent input power of 5.44W, giving a power factor of 0.04 or 4%. Note that because this is so low, many domestic power meters would have trouble giving any kind of reading at all and the power reading could range from zero all the way up to several watts.

The low power factor is due to the fact that so much of the energy drawn from the mains goes into simply charging up C1 and this is returned later in the cycle, so the power moving into and out of the unit via the mains socket is much higher than the actual net consumption.

Secondly, there's nothing to stop you from taking the simulation further and actually drawing up a real load instead of using resistor R3. This would give a more realistic depiction of the voltage regulation of this power supply circuit in the face of changing load demands.

For example, this sort of circuit is commonly used to power a relay, either to act as a mains timer or some sort of load-detecting switch.

Actually, if you look at our SoftStarter in the April 2012 issue, Soft Starter for Power Tools in the July 2012 issue and Mains Timer for Fans and Lights project in the August 2012 issue, you will see just this type of circuit.

In those cases, the load current depends heavily on whether the relay is energised and it's acceptable for the supply voltage to drop once the relay has latched, as a lower voltage is required to hold the relay than to switch it initially. So further simulation would definitely help optimise such a circuit. SC

Next Month

Modelling relays in SPICE is a little tricky but it can be done, as we will demonstrate by building a fairly realistic relay model next month. We'll also get into some more advanced techniques that are possible with LTspice.

SERVICEMAN'S LOG

Fixing the food processor that wouldn't

Dave Thompson*

I'm not sure whether it is Sunspots, Murphy's Law or just plain bad luck, but it appears there is a conspiracy among our household appliances to all fail around the same time. Sitting among those appliances that fail is one of our food processors which is only a few years old.

Items Covered This Month

- A Serviceman's kitchen
- Brownout protection for a TV
- TEAC HDR PVR

**Dave Thompson runs PC Anytime in Christchurch, NZ.*
Website: www.pcanytime.co.nz
Email: dave@pcanytime.co.nz

Last month it was the vacuum cleaner. Then a fan heater decided to stop fanning and heating. Even my own computer has been increasingly throwing up those cloying "your PC has encountered a problem so we have shut it down" screens.

Between crashes, it often displays weird on-screen artefacts, such as coloured blocks of pixels and very thick black lines appearing randomly. As I write this article, it is as if someone is trying to redact what I am writing in real time. Maybe the CIA really can hack into our homes via our smart fridges (or is it ASIO?)! Hmm, we don't have a smart fridge...

Obviously, the computer has a problem and I suspect my graphics card's VRAM is faulty, though this is a wasted diagnosis because I can't actually do anything about it except replace the card anyway. I'm not about to start stripping memory chips off the thing and replacing them.

Considering that it cost a small fortune and was the single most expensive component in this system, I'm not that happy it has decided to fail just outside of warranty. I would expect more than 16 months out of a high-end graphics card, but that's how it often goes with high-performance hardware.

The irony is that I haven't had time to play the games I originally obtained it for as I am too busy doing far more mundane things, such as renovating workshops, which is proving surprisingly difficult to do whilst actually working out of them.

It's a bit like changing the tyre on your car while driving it down the road; not impossible, as demonstrated by many car-crazy middle-eastern YouTubers, but for the rest of us it is definitely pretty tough.

Now our brand-new dishwasher, purchased when we renovated the kitchen nine months ago, has started making odd noises and during the last few nights, the LED display has been randomly flickering between showing all 8's or nothing, to the time remaining and back again.

I can hurry things up with a light tap on the door, so it looks as if something is not quite making contact somewhere in the electronics behind the door panel.

That'll mean a trip to the repair agent; I have to resist the temptation to go searching for the fault myself. For starters, I haven't the room to pull it to bits in my workshop and besides, Mrs Serviceman wouldn't be too keen on me voiding the warranty!

The curse takes another victim

Now to top it off, one of our food processors has decided it wants to stop processing. This appliance is one of the better and most-used of our kitchen tools so having it give up is a bit of a curse, as it is a few years old now and I'm reasonably sure they don't make them any longer, so we can't just go out and buy another one (which is I'm sure what the manufacturer intended). Then I remembered; I'm a serviceman! This shouldn't be a problem for the likes of me!

This particular mixer has a solid Pyrex glass mixing bowl forming the bottom half of the appliance. A tough plastic lid then clamps securely onto the glass bowl and sitting on that lid is a chrome-and-black plastic housing which contains the motor and drive assembly.

This spins a centrally-mounted twin blade system down through a hole in the lid via a splined drive shaft. This machine is so powerful, you could stuff the bowl full of ice cubes and with a few short bursts of the motor, turn them into slush.

The four curved and razor-sharp blades make small work of anything in the bowl, and while Mrs Serviceman and I aren't great foodies (or cooks for that matter), when the urge takes us, we like to have decent tools with which to do the job.

The problem with this thing was that it no longer powered up. There is no obvious on/off switch; instead, it has an interesting push-switch arrangement. On the very top of the motor housing is a black plastic cap that activates a switch when pushed with the palm of one's hand.

However, for (I assume) safety reasons, there is another switch activated when the whole assembly is pushed against the lid of the bowl. This way, the motor will only run when significant pressure is put on the whole motor assembly and both switches are pushed.

It sounds more awkward than it really is; in use, it is quite natural to push down to start operation. It also makes sense because the motor part of the appliance isn't held down by anything other than the operator's hand pressure.

The splined output shaft of the motor assembly locates into a circular moulding in the thick plastic lid, which is well-fixed to the glass bowl underneath, but other than that hand pressure, the motor assembly is free to move, so it makes sense to hold it tightly when it is running and makes even more sense to require downwards pressure before power is switched to the motor.

Except for the fact that no matter how hard I pushed on it, nothing happened.

So I had to assume that either something was wrong with the switch arrangement or the motor had burned out. I was hoping it wasn't a dead motor, because then we'd likely have to bin the whole thing (although I suppose we could use the bowl for something else). However, this is a classic example of the way everything is going these days. We have become a consumable, throw-away society, and this is very apparent in the kitchen.

Start of Serviceman's rant

By comparison, my mother still uses a mixer she bought in the seventies. Sure, it's had the motor reconditioned a few times and there are a few minor cracked mouldings here and there, but the point is that it was designed to be repaired and there are still parts available for it.

Most of the high-end appliances you buy today don't have anything like the parts backup these older brands have and in 40 years they'll be part of a landfill somewhere while those like Mum's will probably be still going. That's not only hugely wasteful, it's bordering on criminal.

People these days buy new printers rather than buy hyper-expensive cartridges for their old one. Tablet owners chuck a tablet with a broken digitiser rather than repair it, even though in most cases it costs far less than a new one. And people chuck food processors because a part that would cost just a few bucks to manufacture gives out and because the parts aren't there to repair them, and who can blame any of them?

Most manufacturers today are only concerned with moving as many units as they can and don't give a toss what happens to their products once they break down; that becomes someone else's problem.

Increasingly, I cannot source parts for even newer models of computers,

I USED ONE OF MY FAVOURITE TOOLS, AN OLD DENTIST'S PICK

forcing me to look to the second-hand or refurbished market. Most punters these days accept that having to buy a whole new anything is inevitable, the collateral damage of technological progression. And maybe you can't really blame manufacturers for not wanting to have capital tied up in parts sitting around gathering dust on a warehouse shelf somewhere. It's easier to just sell more computers and let someone else create a second-hand or refurbished market.

And that's exactly what has happened; these days many companies are buying up old appliances – whether computers, food processors or washing machines – just to strip them down for spare parts. They know what car wreckers have known for decades; that they can make good money from selling parts rather than selling complete appliances.

This is exactly why I am so keen to get stuck into trying to repair anything that breaks down. I consider it a challenge to buck that wasteful philosophy and try to keep things going for at least a reasonable lifespan. If something wears out to the extent it cannot function any longer, then that's fair enough, but when an entire printer is junked simply because a 20-cent proprietary part is not made available, that is not OK in my book.

Normal service(man) resumes

Getting into this blender motor looked like it could be a challenge. From the outside, there wasn't anything much to suggest how it was held together. There was a hole where the power cable entered the case but that was about it.

After having a closer look at the bottom of the motor assembly, I could see a black plug about 5mm in diameter that could be hiding something. I used one of my favourite tools, an old dentist's pick, to pry the bung out and sure enough, it was covering one of those annoying safety-type screws. I found the right bit for my driver and removed the screw.

After a few attempts to pull the cover off, I concluded that something else must be holding it on. However, I couldn't find anything, and so after sitting and contemplating it for a few minutes, I decided to give it a gentle twist.

Sure enough, the cover slowly gave way to the sideways pressure and after turning ninety degrees, the whole motor mechanism began to pull out from the plastic case. It only came out part way as the power cable held everything in, so I forced the cable through the stress-relieving sleeve bit by bit and pulled the motor far enough out until I could see the power switching arrangement.

The series-wound universal (brushed) motor was switched by a large two-stage microswitch that was actuated by two formed, plastic-capped copper-coloured springs. One of these springs was actuated by pressing down on the plastic cap by way of a captive plastic rod. When the cap was pressed, the rod was forced down onto the spring and that flexed and pressed onto one toggle of the power switch.

The bottom spring actuated the second pole of the switch via a clever little plunger arrangement; when the motor body was pressed down onto the lid of the bowl, a corresponding solid piece of plastic in the lid pressed onto the plunger, forcing it upwards into the second spring and causing that to push the switch's second toggle. Only when the two toggles were actuated would power flow to the motor.

This was actually quite a clever "manual" safety system, as it meant the motor could neither be run without the load of the blades to govern its speed nor without holding it down quite tightly. Pressing down also engaged the blade assembly down in the bowl and unless this was all connected properly, the bottom plastic plunger would not be actuated. Simple and effective; and broken.

The motor wouldn't power on because the bottom spring had somehow broken away from its mount, meaning that it simply flexed out of the way when the plunger touched it, instead of transferring that pressure to the switch.

I hadn't noticed any plastic shrapnel floating around in the case but then again, any fragments might have been small enough to work their way through the small gap for the blade shaft in the lid and fall down into the bowl below. I suppose we'll find them when we defrost and eat the Ćevapi my wife made when she last used the blender!

Now though, I would have to rebuild the support for the bottom spring so I made a rough-and-ready former using strips of gaffer tape stuck to the remainder of the plastic mount. It was originally a rectangular shape and thus easy to recreate. To build it up, I utilised a two-part compound that comes in a tube and one simply slices off a small amount and rolls it to mix it.

It sets rock-hard in about five to 10 minutes, depending on the ambient temperature, so you can't faff about once mixed. Another smaller piece of tape held the contact in place while the pressed-in compound set and I left it overnight to be certain it was properly cured.

The next afternoon, I removed the tape and tested the contact's travel, making sure it reached and actuated the microswitch. It did, so I reassembled the motor back into the case and with power applied, tried it a few times to ensure it started consistently with the usual amount of pressure. Hopefully, that fix will last and we'll get a bit more use before we have to consider chucking it away.

DIY circular saw upgrade goes awry

And in one of the funnier service moments this month, a customer (and neighbour) brought in one of those bright green, plastic-bodied benchtop table saws that he'd stripped down in order to mount the guts into another more solid benchtop. He'd made a nice job of swapping everything over but had needed to remove the power switch connections and associated wiring in order to get it out of the original housing.

After throwing away all the old parts, he'd fixed the saw and motor assembly into a recessed pocket in a sheet of 32mm custom-wood. That would certainly make a far sturdier base for it than the original flimsy plastic body and thin, cast aluminium table top.

The problem was that he'd chucked away the diagram he'd made of the wiring and was too smart (or afraid) to mess about with mains wiring. Fair enough; sometimes power tools can be wired up in weird ways and just randomly re-connecting spade connectors and plugging it in isn't the best way forward.

The first thing I did was look for a circuit diagram for the saw. I found plenty, all American and not much use for our configuration. I Googled the paddle-style safety switch's part number and found a suitable usage diagram in the datasheet, and with it all wired up correctly, I plugged it into my Variac and wound on 50VAC before pushing the green button. The motor spooled up but there wasn't enough juice to hold the button in. I added another 100 volts and away it went.

Pushing the red "Stop" button switched it off as expected so I called my neighbour to tell him it was ready. The next day, he brought it back, claiming the motor was running backwards. Puzzled, I drew a quick wiring diagram and looked at how changing the connections on the switch would make the motor run in reverse.

There wasn't any configuration that would do that; to swap rotation, I'd have to change either the armature polarity or the field polarity with respect to each other and as it had run properly before, I couldn't see how that was all necessary.

After pondering it for a while, it hit me; he'd put the blade on backwards, which of course made it look as if the motor was running the wrong way. After flipping it, all was well. At least it was an easy fix!

Browned off by brownout protection

G. M., from Pukekohe in New Zealand, recently had a call-out to a remote location to a malfunction in a new TV installation which proved to be a case of "too much protection is too much".

I am the local service agent for a well-known and popular range of home entertainment products. I normally restrict my call-out radius to about twenty-five kilometers from base but this was one of those occasions when Murphy's Law kicked in, and a faulty TV as far away from base as I could be enticed to go became a real head scratcher. This law is closely related to the one which decrees that when a screw is dropped, it will always roll to the darkest, dustiest, most inaccessible corner of the workshop!

A local retailer phoned me with an appeal to talk with one of their customers who had purchased a 55-inch TV three months earlier and he was now convinced the TV was faulty. The farming owner lived near the end of a long peninsula, a 60km one-hour drive over reasonable but quite twisty roads.

I was very busy with other work so I was not keen to waste half a day unproductively driving back and forth all that way, despite the very scenic vistas I would enjoy on the journey. So I was eager to solve the problem by any other means, hoping for an installation or operation issue which I could talk the owner through.

I phoned the owner and he explained his problem. The TV would completely cut out for two to three minutes and then come on again with no sound or picture. It would repeat this several times an evening and it was happening virtually every night. Soon after cutting out, the screen would show the word "SAT" in the top left corner. When the picture and sound eventually came back on everything worked normally until the next time. He had only experienced these problems since purchasing the new TV. The old plasma TV had worked just fine.

I asked some obvious questions such as whether the problem was occurring on all sources which was responded to with a considerable pause. So I clarified my question by asking if there was anything else connected to the TV such as a set-top box or DVD player. Yes, there was a Sky box (pause again); not sure about a DVD player. Now we're getting somewhere – or are we? I gained the impression that he was not sure what a DVD player was. He only watched Sky broadcasts.

The owner explained that he had had the Sky people out twice attempting to solve the issue and they had eventually replaced the decoder, to no avail. Having eliminated that as a cause, he surmised that the problem now had to be with the TV. He even took a photo of the malfunctioning TV with his phone and sent it to me so he was obviously not completely technophobic! I talked the owner through the procedure to reset the TV to factory condition and left him to test it that night.

The next morning, I phoned a couple of satellite installers whom I knew and discussed the possibility of there still being a decoder or satellite signal issue since this was not my area of expertise, but neither of them were convinced that the problems were Sky related. However, since they had both been in the trade nearly as long as me, neither was prepared to lay odds that it was not a Sky reception fault.

Meanwhile, I forwarded the photo to the TV brand's technical manager and had several discussions with him by email and phone regarding the fault. He seemed to be quite certain that the "SAT" logo appearing on the screen was not something being generated by the TV so the problem had to be elsewhere. We were running out of elsewheres.

The last straw caused the frustrated owner to send a curt text to my phone a couple of days later that 'Now everything had cut out, nothing was working – I'll pay for you to call'. I phoned him and waited while he fetched an electric drill to plug into the wall to test that the power was on. The drill whirred into life so I said I'd think on it and let him know. He was slowly convincing me that the TV was to blame and had somehow now caused everything to fail.

I talked with the brand's head technician again and he agreed to cover the considerable cost of a call out if it proved to be a TV fault. I phoned the owner and he was happy for me to call on the basis that if it was not the TV at fault then he had to bear the cost. He was quite convinced that the TV was faulty, so was confident that he would not be writing a cheque.

I called the next day armed with a small loan TV of the same brand with a twofold job for it – first, to pacify the owner by leaving him something to watch and second, to act as a test unit while I checked his new TV at my shop.

The first thing I noticed when I arrived at his home was that there was indeed no power. The TV, decoder and sound bar/sub-woofer were all dead. Up until this point I had not known there was a sound bar. It was now I also learned that the TV and sound bar had been delivered and installed by the retailer at the time of purchase.

I knew there was power to the wall so I peered over the back of the TV and saw a nice new multi-outlet power box there, no doubt sold by the retailer as an add-on to give protection to the new equipment in the event of a power surge. In my opinion this is a duplication of the sort of protection which is now already built into most equipment, but having been a retailer myself in a past life, I was not about to cast aspersions on the practice of up-selling.

The box had an on/off switch on the top so I reached for it and toggled it to the other position at which point I was greeted by an encouraging green glow from beneath the switch button and now everything was working.

I quizzed the elderly couple about how the switch could have been bumped off since there were obviously no children around at which point the lady admitted that one of her cats had chased a mouse behind the TV the previous day and must have jumped on the switch. This seemed quite plausible since the switch button was quite large and the switch took little effort to operate.

So, the immediate no-power problem was solved but I was fairly sure this had no connection (pun intended) with the original complaint so I carried on with swapping the TVs and after the loan TV was connected and working, I checked some of the menu settings on the decoder and sound bar. These were both connected with HDMI cables to the TV.

I've experienced some odd behaviour in the past by allowing such appliances to talk to each other. I went into the menu of each device and switched off this interactive feature to eliminate arguments between them as a possible cause of the original complaint. If the owner wants this feature, it can be readily restored, once our present issues were behind us.

Back at base I connected the new TV to an HDMI source and left it running. It didn't miss a beat. That evening I received a text from the exasperated owner to say the same thing was happening with the loan TV except that this TV displayed the "HDMI1" in the top left corner, not "SAT" which the new TV had displayed.

It was now that he offered a little more information which in hindsight, I probably should have asked about earlier. The problem only seemed to occur at around dinner time each evening and it did not happen so often later in the evening.

Now armed with another clue in the timing of the fault occurring, I wondered whether this could be a power supply problem after all, but the owner wasn't convinced of this since despite living at the end of the road at the tip of a long peninsula, it couldn't be a power voltage drop because the lights remained bright.

Now clutching for straws and wanting to resolve the problem before I headed away for a brief break in a few days, I connected the new TV in my shop to a variable transformer just to eliminate that from the mix of possibilities. I slowly reduced the voltage from the normal 230VAC and I was surprised that the TV worked perfectly down to 80VAC, at which point it cut out. As soon as the voltage came up a bit, the set burst back to life without so much as a hiccup. I reset the TV a second time and adjusted some of the menu settings.

The next day was a public holiday. Despite that and being so keen to get on top of the problem, I loaded the owner's TV into the van. I included the variable transformer and set off with some trepidation, knowing that if I did not solve the problem this time around, I would have to admit defeat. Who would want to pay me for the many hours I had already invested in this job if I couldn't provide a solution?

On arrival, I connected the variable transformer to the wall socket and the multi-outlet box in turn to the transformer and turned everything on. As I slowly reduced the voltage from 230VAC to just above 200VAC there was a click and everything died. Turning the voltage back up a little produced another click and everything except the sound bar powered up again.

The clicking noise was coming from the power multi-box and it was now obvious to me that this was a feature of the device to protect connected appliances against brown-out damage due to prolonged low voltage supply. A press of the power button on the sound bar remote control turned that on again.

When I repeated the exercise, I noticed the TV came back on almost immediately after the voltage was restored but the Sky decoder took several minutes to reboot itself. In the meantime, the TV patiently displayed a blank screen with "HDMI1" showing in the top left corner, until it once again received a signal from the decoder and then displayed a perfectly normal picture.

These were exactly the symptoms the owner had described. Eureka! After substituting the power multi-box with a cheap discount store version, without the brown-out protection feature, he has had no further problems – so far.

Epilogue

So I had been hoodwinked by a technological feature of the power outlet multi-box which I did not realise it had. In any other location, the protection would most likely not have tripped until a real and rare brown-out event happened and I would have remained blissfully unaware of the feature.

After the power was restored and the protection reset, the viewer would continue to enjoy uninterrupted television for many more months or even years, before such an event might happen again to trip the protection.

Since I was never present in the owner's home when the unprompted fault happened, I can only suspect his mains voltage is regularly dropping below 200VAC. I was not planning to spend an evening there monitoring the mains voltage to find out. I suggested to him that he purchase a $30 plug-in power monitor available from Jaycar to check for himself; it is not uncommon for there to be power fluctuations in rural locations.

The owner was hoodwinked into believing he did not have a power problem by a different technology – CFL lights which maintain their brightness despite the voltage drop. Unlike incandescent lamps which give a visual indicator of the voltage of your power supply; flickering or obviously dimming when the voltage dropped.

Yet another hidden and overlooked clue was that the Sky decoder was housed in a cabinet with a tinted glass door, so it was not obvious to the owner that the decoder was going through its reboot sequence after the power interruption. All he saw was the blank TV screen which came back to normal by itself after a few minutes.

I suggested the owner have an electrician check his power supply as such a severe voltage drop could suggest a high resistance in one of the three phases supplied to the property, which in turn could pose a fire risk. All of this happened in late summer so it was destined to be much more of an issue come mid-winter when heaters, hot water and ovens are all working harder.

The "SAT" word which side-tracked me briefly must have been programmed into the new TV when it was installed by the retailer. I was surprised he had gone to this much trouble. It is a feature within the menu of some TVs which allows you to enter your preferred label for the various inputs. Resetting the TV the first time erased this entry and it defaulted back to the HDMI1 label.

Another little twist was that despite purchasing a nice sound bar and subwoofer, again maybe an upsell by a keen salesman, the new owner was missing out completely on the much-improved sound these systems offer over the standard TV speakers.

From the very first time the power dropout happened, the speaker system had not automatically restarted. It was the only device not to do so and despite a scan through its menu, I could not find any reference to power settings which would allow me to select auto power on. It probably relies on the HDMI interaction feature which I mentioned earlier, switching itself on or off in sync with the TV.

The owner had been using the TV speakers only and I know this because he complained that he had to hold the TV remote control high in the air to adjust the volume. The sound bar was blocking the infrared signal to the TV, unless the remote was raised sufficiently above it.

When I explained to the owner that he needed to deliberately power the sound bar back on after a power interruption, I could tell he had not been doing so. In fact, he was completely unaware that it had not been working and that he had been missing out on much improved sound fidelity from his entertainment system.

I had another job recently with a near new TV which came in for service with both speakers blown; this despite the owner having a sound bar. I wonder if this was another case of the sound bar not being able to come back on automatically after a power interruption?

Faulty capacitors don't need to be bulging

L. W., of Logan, Qld, replaced bulging capacitors on the main board of his malfunctioning PVR and thought that would fix it. But it wasn't that simple and it took a lot more detective work to complete the repair.

About 12 months ago, my TEAC HDR 1600T Personal Video Recorder began to exhibit erratic operation until it finally failed and would not complete its start up boot process, hanging approximately half-way through.

I had experienced problems before with the unit which had turned out to be a bulging electro in the power supply circuit, so it was with some hope that I removed the cover to take a look inside.

After spending some time examining the power supply and finding no signs of distress, I turned my attention to the main board. Only part of this board was visible as the hard drive is mounted above it. There was nothing obvious here either so I decided to switch on and measure the outputs from the power supply.

I waited for it to complete as much of the boot process as it could and then took some measurements. There was no joy here as all appeared to be well within tolerance. At this point I started to lose interest as the machine was now getting a little long in the tooth and much better machines with larger hard drives are available on the market.

However, there were a couple of weekly shows that I had recorded and I really did want to view them. I decided to remove the hard drive for a better look of the main board. And it was here that I struck pay dirt or so I thought. Under where the hard drive had been sitting I could clearly see three electrolytic capacitors with their tops bulging.

I didn't have any of the correct values (all three were 220µF 16V types) so after a quick visit to the local electronics store, I wasted no time fitting the capacitors and replacing the main board. The HDD was reinstalled, power supply connected and switched on.

Well, you can imagine how I felt when this time there was even less response from the unit than before I had started. In disgust, I switched off, pushed it to the rear of the bench and that's where it sat for 12 months.

During that time it got in my way on several occasions but I just ignored it as best I could. I did make an attempt to boot the HDD from my PC in an effort to retrieve the information from it but even with the assistance of the internet I wasn't able to achieve an end result.

Having recently purchased a replacement machine, the time came to decide whether to have another go at fixing it or consign it to the wheelie bin. It was still in a dismantled state and picking up the main board, I noticed that there were several other similar capacitors in the same area as the ones I had previously changed.

All were black with blue labelling while most of the other electros on the board were black with white labelling. So for no other reason I decided to change those too, as I felt the symptoms dictated a faulty electro somewhere.

I didn't have the correct values but decided to use what I had at hand rather than waste any more money buying new ones. In all I ended up changing CE2, 3, 11, 15, 16, 24, 29 and CE223.

Some of the capacitors I used were physically bigger, so I had to leave their legs fairly long and bend them over into what space I could find so as to reinstall the HDD which, as stated, fitted over the whole lot. The photo shows some of the replacement capacitors.

I didn't hold much hope of success, so it was a surprise when the machine booted up and all appeared well. That was several months ago now and I am gradually catching up on episodes of the weekly shows I had recorded some 12 months ago.

However, it did bug me in not knowing which capacitor was actually the culprit. So one night I fired up the scope and oscillator and measured the ESR of the capacitors that I had replaced the second time around.

Again to my surprise, while it was a 220µF capacitor like the bulging ones I had originally replaced that was obviously faulty, all the others were at various stages of deterioration, with much higher ESR than they should have exhibited.

So there you go; another piece of electronic equipment saved from the tip for just the price of a few electrolytic capacitors. You can be lucky sometimes. SC

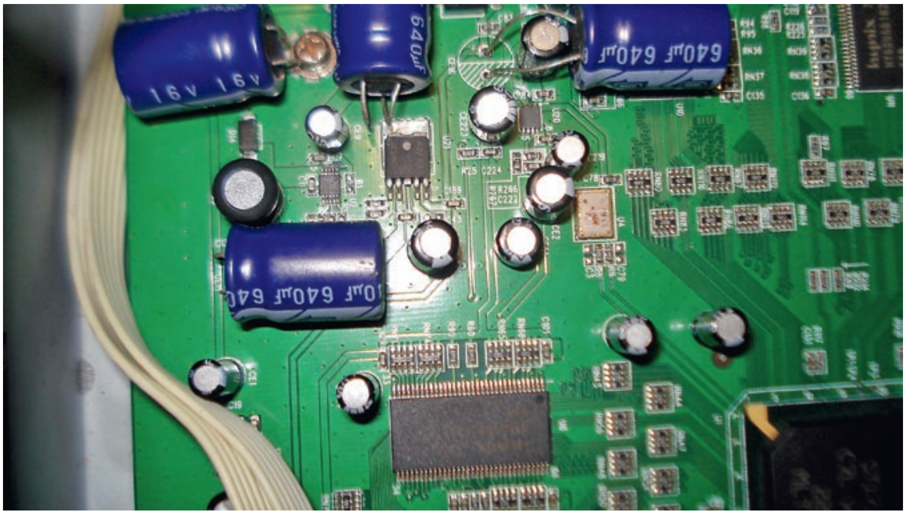

Bulging 220µF 16V electrolytics weren't the only issue in this video recorder as it wasn't until some of the larger capacitors were changed due to high ESR values that the set came back to life. This photo shows the new capacitors installed.

MID YEAR SALE

Jaycar Electronics

SAVE UP TO $150

60W ESD SAFE SOLDER/DESOLDER REWORK STATION

Complete solder/desolder station for production and service use. Microprocessor controlled for precise control of thermal performance. 60W Power.

- Backlit LCD displays
- Temperature range: 160 - 480°C
- 215(W) x 225(L) x 155(H)mm

TS-1574 **ORRP $349**

NOW **$199** SAVE $150

DUAL OUTPUT LABORATORY POWER SUPPLY

The two outputs can be operated independently, connected in parallel, or series for multiple output currents and voltages.

- Output voltage: 0-32VDC (x2)
- Output current: 0-3A (x2)
- 185(H) x 260(W) x 400(D)mm

MP-3087 **WAS $399**

NOW **$319** SAVE $80

2-IN-1 HANDHELD SCOPE & MULTIMETER

All the functions of a 4,000 count True RMS CAT III digital multimeter and a 10MHz oscilloscope. Includes USB interface and PC logging software.

- Autoranging, AC/DC voltage (1000V), current (20A)
- 50MSa/s sample rate

QM-1577 **WAS $579**

NOW **$499** SAVE $80

SAVE $150

16 CHANNEL 720P ANALOGUE HIGH DEFINITION (AHD) DVR

View live footage on a Smartphone. Supports monitors with VGA and HDMI input. Plug & Play easy network setup. Includes DVR, 1TB HDD, mains power adaptor and mouse.

- Dropbox cloud storage
- Motion trigger recording
- 300(W) x 227(D) x 53(H)mm

QV-3149 **WAS $749**

NOW **$599** SAVE $150

720P AHD BULLET CAMERAS TO SUIT

- Ideal for a wide range of surveillance applications
- 1/3" CMOS sensor
- IR LEDs for day/night vision
- IP66 rated
- Includes 18m cable and power supply

VARI FOCAL QC-8668 **ORRP $179**
PAN & TILT QC-8670 **ORRP $249**

$99 QC-8668 SAVE $80

$149 QC-8670 SAVE $100

SAVE 25%

BASIC DUINOTECH EXPERIMENTERS KIT

The entry point for learning and experimenting with Arduino®. Contains many parts to get you up and running including duinotech Nano board, breadboard, jumper wires and plenty of peripherals. XC-4285 **WAS $79.95**

NOW **$59.95** SAVE $20

SAVE $110

1080P REAR VIEW MIRROR CAR EVENT RECORDER

Clips over your existing mirror and features a cool blue anti-glare tint, IR night vision, motion detection recording, loop recording, emergency file lock, time and date display.

- 4.3" colour LCD
- MicroSD card required (up to 32GB)

QV-3860 **ORRP $199**

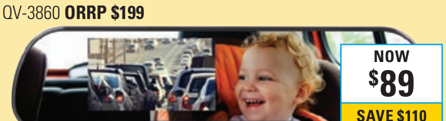

NOW **$89** SAVE $110

SAVE $60

NOW **$129** SAVE $60

NOW **$139** SAVE $60

2 X 50WRMS COMPACT STEREO PA AMPLIFIER

Delivers the quality of a Class AB amplifier with the efficiency of a Class D. Solid aluminium body, banana socket speaker terminals, 3.5mm stereo input & 6.5mm headphone socket.

- Includes power supply and audio cables
- 78(W) x 150(D) x 50(H)mm

AA-0488 **WAS $189**

2 X 75WRMS COMPACT STEREO AMPLIFIER

Powerful 2 channel (Stereo) in a compact size.

- Built-in digital signal processor
- Volume control and on/off switch on the front panel
- Includes power supply and cables
- 165(L) x 95(L) x 30(H)mm

AA-0505 **WAS $199**

2 CHANNEL WIRELESS UHF MICROPHONE SYSTEM

Ideal for schools, churches, weddings, gyms, karaoke, etc. Supplied with two microphones.

AM-4114 **ORRP $149**

NOW **$89** SAVE $60

900W DMX FOG MACHINE

Use remote control or hookup to a DMX512 controller for total customisation of stage or party effects.

AF-1213 **WAS $189**

NOW **$129** SAVE $60