Practical Electronics

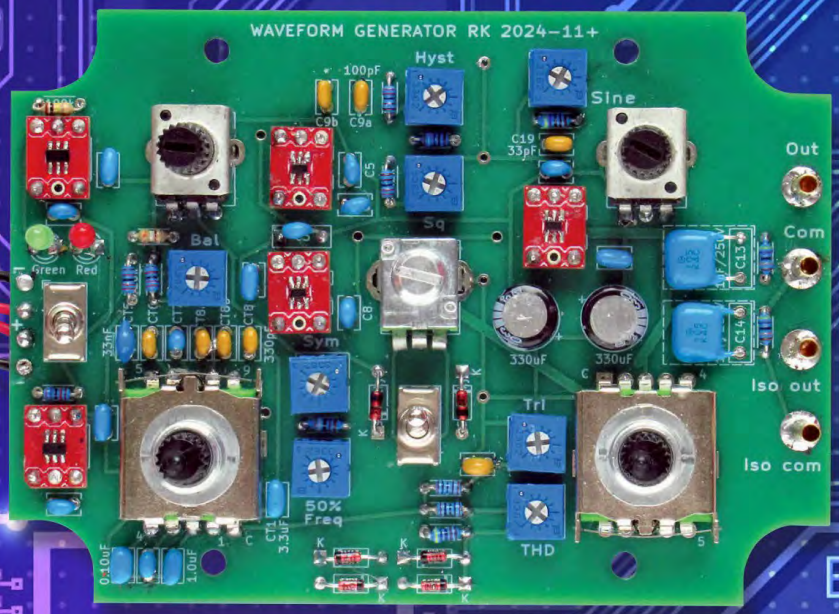
The UK's premier electronics, computing and maker magazine

Max's Cool Beans
Cartridges for the retro games console

Teach-In 12
Part 4: software-defined radio

Audio Out
A plethora of fuzz pedals

Circuit Surgery
Op amp gain, phase, stability & bandwidth

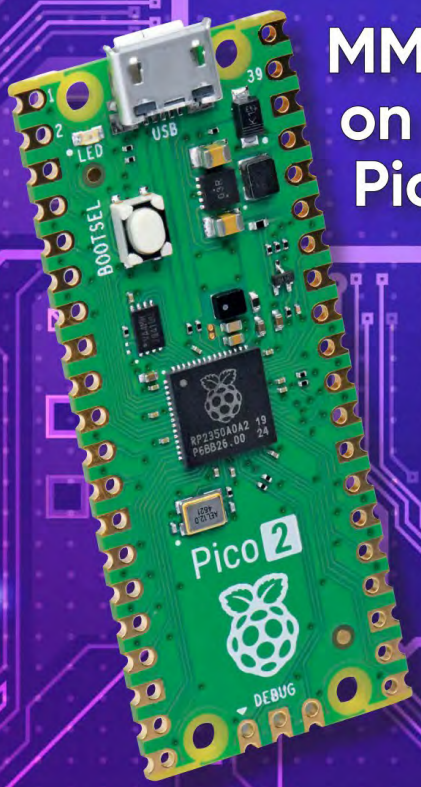


Multi-waveform Function Generator using op amps

High-current Inductor Tester



MMBasic on the Pico 2



PLUS!

Techno Talk – ultra-low-power comms

Raspberry Pi Pico 2 review

Mini "UPS" module review

Feb 2026 £5.99

9 772632 573030 02



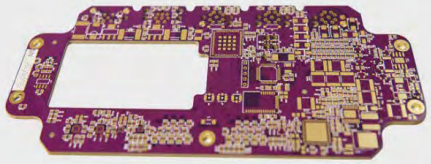
PCBWay

Beyond Standard —

Discover PCBWay's Advanced PCB Lineup

About PCBWay:

With dedicated workshops for advanced PCBs, PCBWay offers full-spec, high-precision, and large-scale production. We ensure customers are supplied with PCBs that meet stringent quality standards at the lowest cost.

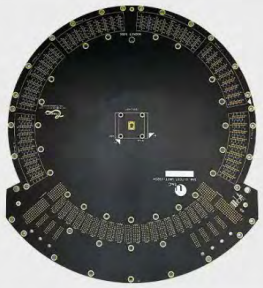


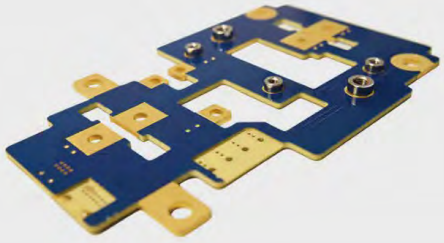
HDI PCB

- 4-60 layers
- FR-4, High Tg, Halogen-free, PTFE
- Finished Copper: 1/3OZ–8OZ
- Min Track/Spacing (inner): 2/2mil (H/HOZ base copper)
- Min BGA Pitch: 0.35mm

High Frequency PCB

- RO4003C, RO4350B, Ro3003, Ro3010, RT5880
- Min. Track/Spacing: 2mil
- Min. Hole Size: 0.15mm
- Finished Cooper: 1-2oz
- Board Thickness: 0.2–3.2mm





Other PCBs with Special Process

- Embedded Copper PCB
- Backplane PCB
- Cavity PCB
- White/Partially Hybrid PCB
- Black Core PCB (After-dark)



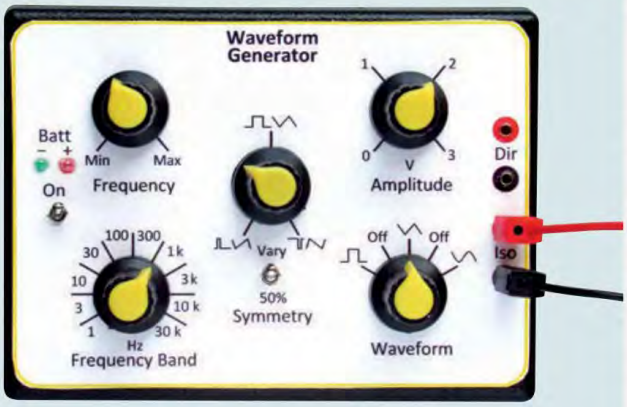
www.pcbway.com service@pcbway.com

Practical Electronics

Volume 55. No. 02
February 2026
ISSN 2632 573X

Contents







Projects and Circuits

Power LCR Meter, part one by Phil Prosser 10
This LCR meter can deliver up to 30A to inductors to determine their properties at higher power levels. That makes it particularly useful for determining when and how an inductor saturates. It can also measure very low resistances and high capacitances.

Versatile Waveform Generator by Randy Keenan 34
This versatile waveform/function generator is handy for a variety of uses, including audio equipment analysis and circuit development. It uses three op amps to deliver square, pulse, triangle, ramp and sine waves from 1Hz to 30kHz.

The PicoMite 2 by Geoff Graham 54
This MMBasic interpreter for the Raspberry Pi Pico 2 and Pico 2 W takes advantage of the new features of the Raspberry Pi Pico 2. It converts the Pico 2 into an easy-to-use and powerful platform for beginners and experts alike.

Series, Features and Columns

Max's Cool Beans by Max The Magnificent 4
Weird & Wonderful Arduino Projects, part 14: enough playing games, it's time to play games! This instalment focuses on wiring up the plug-in cartridge system that our games console will use to select and play games.

Teach-In 12.4: the World of Wireless by Mike Tooley 22
This month, our focus shifts to software-defined radio (SDR), which enables radio signal processing tasks traditionally handled by hardware to be managed via software. We examine different SDR solutions that will work from HF to UHF.

Circuit Surgery by Ian Bell 42
Measuring the frequency response of a circuit using a PC sound card, part 5: op amp frequency response & stability. We look into op amp open-loop vs closed-loop gain, gain bandwidth, phase shift, phase margin and stability criteria.

Raspberry Pi Pico 2 microcontroller module review by Tim Blythman 50
The Raspberry Pi Pico 2 microcontroller board is similar to the original Pico (released in 2021) except that it features the new RP2350 microcontroller chip.

Techno Talk by Max the Magnificent 60
Welcome to the sensorium. From breadcrumb-sized accelerometers to microwatt neural nets and backscattered wireless links, ultra-low-power technologies are turning everyday objects into intelligent, connected companions.

Audio Out by Jake Rothman 62
We've described some musical instrument fuzz boxes over the last few issues; this new PCB can be configured in many different ways to produce all kinds of fuzz pedals.

Review: Mini UPS module by Jim Rowe 70
This module uses a single Li-ion cell (or a few in parallel) to provide backup DC power. It's useful for powering devices like internet routers so they'll work during a blackout.

The Fox Report by Barry Fox 74
Pure's Classic C-D6i do-it-all DAB & FM radio/internet streaming audio player.

Regulars and Services

Editorial 2
Waiting for the AI bubble to burst

Reader services – Editorial and Advertising Departments 2

Electronics Teach-In Bundles 3

Subscribe to Practical Electronics and save money 73

Practical Electronics Back Issue DOWNLOADS – 2024 now available! 75

Practical Electronics PCB Service 76-78
PCBs for Practical Electronics projects

Classified ads and Advertiser index 79

Next month! – highlights of our next issue of *Practical Electronics* 80

Renewals paid by PayPal, Visa and Mastercard are now handled via our website: www.electronpublishing.com



Made in Australia & the UK.

Written in Britain, Australia, the USA and Ireland.

Read everywhere.

© Silicon Chip Publications Pty Ltd 2025

Copyright in all drawings, photographs, articles, technical designs, software and intellectual property published in *Practical Electronics* is fully protected, and reproduction or imitation in whole or in part are expressly forbidden.

The March 2026 issue of *Practical Electronics* will be published on Thursday, 5th of February, 2026 (see p80).

Practical Electronics

Volume 55. No. 02
February 2026
ISSN 2632 573X

Editorial

Editorial offices
Electron Publishing (Australia)
Tel +61 2 9939 3295
Email pe@pemag.au
Web www.electronpublishing.com

Address mail to:
Electron Publishing (Australia)
PO Box 194, Matraville NSW 2036
Australia

Advertising enquiries
+61 2 9939 3295 pe@pemag.au

Editor	Nicholas Vinen
Publisher	Nicholas Vinen
Digital subscriptions	Stewart Kearn **Tel 07918 614662**
Online Editor	Alan Winstanley
Web Systems	Kris Thain
Production	Bao Smith
Technical staff	Tim Blythman, John Clarke

Print subscriptions
Practical Electronics Subscriptions
PO Box 6337
Bournemouth BH1 9EH **Tel** 01202 087631
United Kingdom **Email** pesubs@selectps.com

Technical enquiries
We regret that technical enquiries cannot be answered over the telephone. We are unable to offer any advice on the use, purchase, repair or modification of commercial equipment or the incorporation or modification of designs published in the magazine.

Questions about articles or projects should be sent to the editor by email: pe@pemag.au

Projects and circuits
All reasonable precautions are taken to ensure that the advice and data given to readers is reliable. We cannot, however, guarantee it and we cannot accept legal responsibility for it.

Some projects and circuits published in *Practical Electronics* employ voltages that can be lethal. Do not build, test, modify or fix any mains-powered equipment unless you fully understand the safety aspects involved and you use an RCD (GFCI) adaptor.

Component supplies
Silicon Chip Publications may offer kits or other parts for making our projects, but not in all cases. When kits are not available, readers will need to find and source parts themselves. We advise readers to check that all parts are still available before commencing any project in a back-dated issue.

Advertisements
Although the proprietors and staff of *Practical Electronics* take reasonable precautions to protect the interests of readers by ensuring as far as practicable that advertisements are bona fide, the magazine and its publishers cannot give any undertakings in respect of statements or claims made by advertisers, whether these advertisements are printed as part of the magazine, or in inserts. The Publishers regret that under no circumstances will the magazine accept liability for non-receipt of goods ordered, or for late delivery, or for faults in manufacture.

Transmitters/bugs/telephone equipment
We advise readers that certain items of radio transmitting and telephone equipment which may be advertised in our pages cannot be legally used in the UK. Readers should check the law before buying any transmitting or telephone equipment, as a fine, confiscation of equipment and/or imprisonment can result from illegal use or ownership. The laws vary from country to country; readers should check local laws.

Waiting for the AI bubble to burst

I hope you aren't planning to upgrade your computer any time soon, given the explosion in memory prices over the last month or so.

It seems that the AI bubble (and it surely is a bubble) is doing to the memory market the same thing it did to the GPU (graphics processing unit) market over the last few years – gutting inventories and driving prices through the roof.

This is going to get worse before it gets better. RAM prices (and storage prices, to some extent) have always been cyclical. Someone builds a new factory, the price drops. The factory goes out of business, prices rise. But this is the worst it has been in a very long time.

We upgraded most of our computers around 2020-2022. I was starting to think about whether it's maybe time to get a new workstation soon, but that thought has evaporated now that a couple of sticks of DDR5 RAM cost more than a whole computer did a few months ago.

Since early 2023, I have been convinced that there will be situations where AI (or machine learning, if you prefer) will be incredibly useful. It has also become pretty clear that the applications of AI are, at least for now, quite limited due to its unreliability.

As a programming assistant, AI can be extremely useful. It is undoubtedly capable of reducing the time spent on repetitive tasks. But it can't replace a human in most roles yet, and possibly won't for a long time. So the current "stick AI in everything" push is getting a bit ridiculous.

Unfortunately, the bubble is still inflating. When it does pop, hopefully computer part prices will come back down to Earth. And maybe we can start to get a bit more realistic about what this technology can and can't do. I think we're more likely to put it to good use if we can be honest about its limitations as well as praise its capabilities.

So, for now, I guess I will be writing these editorials on the computer that I put together a few years ago. Hopefully you can stick with yours too, until prices go back to normal. It will have to happen eventually!

Nicholas Vinen, Electron Publishing (Australia)*
Publisher & Editor, Practical Electronics Magazine
* a division of Silicon Chip Publications Pty Ltd.

ETI BUNDLE (1) Teach-In 3, 4 and 5 – all on CD-ROM – only £18.95

ELECTRONICS TEACH-IN 3 – CD-ROM

Mike & Richard Tooley



The three sections of the *Teach-In 3* CD-ROM cover a huge range of subjects that will interest everyone involved in electronics – from newcomers to the hobby and students to experienced constructors and professionals.

The first section (80 pages) is dedicated to *Circuit Surgery*, *EPE/PE's* regular clinic dealing with readers' queries on circuit design problems – from voltage regulation to using SPICE circuit simulation software.

The second section – *Practically Speaking* – covers hands-on aspects of electronics construction. Again, a whole range of subjects, from soldering to avoiding problems with static electricity and identifying components is covered. Finally, our collection of *Ingenuity Unlimited* circuits provides over 40 circuit designs submitted by readers.

The CD-ROM also contains the complete *Electronics Teach-In 1* book, which provides a broad-based introduction to electronics in PDF form, plus interactive quizzes to test your knowledge and TINA circuit simulation software (a limited version – plus a specially written TINA Tutorial).

The *Teach-In 1* series covers everything from electric current through to microprocessors and microcontrollers, and each part includes demonstration circuits to build on breadboards or to simulate on your PC.

ELECTRONICS TEACH-IN 4 – CD-ROM A BROAD-BASED INTRODUCTION TO ELECTRONICS

Mike & Richard Tooley



The *Teach-In 4* CD-ROM covers three of the most important electronics units that are currently studied in many schools and colleges. These include, Edexcel BTEC level 2 awards and the electronics units of the Diploma in Engineering, Level 2.

The CD-ROM also contains the full *Modern Electronics Manual*, worth £29.95. **The Manual** contains over 800 pages of electronics theory, projects, data, assembly instructions and web links.

A package of exceptional value that will appeal to anyone interested in learning about electronics – hobbyists, students or professionals.

ELECTRONICS TEACH-IN 5 – CD-ROM JUMP START

Mike & Richard Tooley



15 design and build circuit projects for newcomers or those following courses in school and colleges.

The projects are: ■ Moisture Detector ■ Quiz Machine ■ Battery Voltage Checker ■ Solar-Powered Charger ■ Versatile Theft Alarm ■ Spooky Circuits ■ Frost Alarm ■ Mini Christmas Lights ■ iPod Speaker ■ Logic Probe ■ DC Motor Controller ■ Egg Timer ■ Signal Injector Probe ■ Simple Radio Receiver ■ Temperature Alarm.

PLUS

PIC'n' Mix – starting out with the popular range of PIC microcontrollers and *Practically Speaking* – tips and techniques for project construction.

The CD-ROM also contains:

■ Complete *Teach-In 2* book, a practical introduction to PIC microprocessors

■ MikroElektronika, Microchip and L-Tek PoScope software.

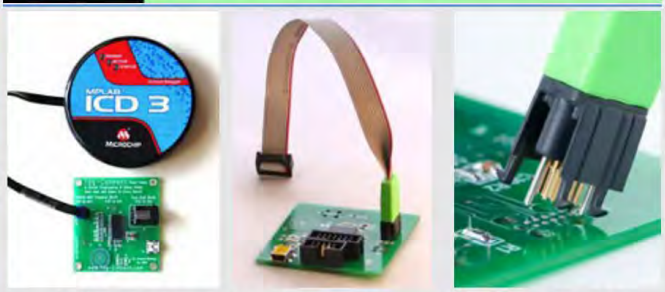
Three Teach-ins for the great price of £18.95

PLUS you also get the contents of the free CD-ROM from each issue...

... so that's another TWO Teach-Ins and The Full Modern Electronics Manual!

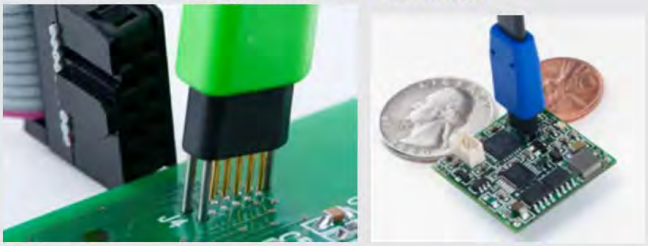
What a Bargain!!

Tag-Connect



JTAG Connector Plugs Directly into PCB!!

No Header! No Brainer!



Our patented range of Plug-of-Nails™ spring-pin cables plug directly into a tiny footprint of pads and locating holes in your PCB, eliminating the need for a mating header. Save Cost & Space on Every PCB!!

Solutions for: PIC . dsPIC . ARM . MSP430 . Atmel . Generic JTAG . Altera Xilinx . BDM . C2000 . SPY-BI-WIRE . SPI / IIC . Altium Mini-HDMI . & More

www.PlugOfNails.com

Tag-Connector footprints as small as 0.02 sq. inch (0.13 sq cm)

ENIGMA COMPONENT SHOP

www.enigma-shop.com

Electronic Components for the Hobbyist!



Cable & Wire • Capacitors • Connectors • Filters
Kits & Modules • Hardware & Cooling
and much more...

READY TO BE REWARDED? Join our loyalty scheme and **earn points** on every purchase you make which can be **used for discounts** in the future.

Read our reviews on Trustpilot

By Max the Magnificent

Weird & Wonderful Arduino Projects

Part 14: enough playing games, it's time to play games!

Hello there. But that's enough of the pleasantries. Hold on to your hat because I have exciting news. Believe it or not, we are finally ready to take all the disparate elements we've worked on in previous columns and bring them together to form our retro games console.

As a reminder, in addition to our Arduino Uno, we have the 14 × 10 array of tri-colour light-emitting diodes (LEDs) that we created in Part 3 (March 2025). We also have our six 7-segment displays (implemented as three two-digit prebuilt modules) that we already mounted on the front panel in Part 13 (January 2026).

Last, but not least, we have our switches, which we've already attached to the front panel, but have not yet wired up.

As a further reminder, in these discussions I'm assuming that (like me) you will be using a laser-cut front panel and a 3D-printed enclosure. As previously noted, you could simply attach everything to a timber board using hot glue. If anyone asks, you can tell them it's a prototype and then fix them with a steely glare (my dear old mother could give lessons in the art of the steely glare).

A rush of power

When running in standalone mode (not connected to a host computer via a USB cable), the Arduino Uno requires a 9V supply. This can be presented using its power jack or via its V_{IN} pin. The Arduino will use its onboard voltage regulator to drop the 9V supply to the 5V it requires for its internal operations.

Way back in the mists of time (April 2025), we decided to power our console using an external 9V supply. In my case, I'm using a very tasty NANKADF (gesundheit) 30V 10A bench power supply. On the off chance you are in the market for one of these little beauties, the same unit is available on Amazon UK (https://pemag.au/link/ac48).

Our 14 × 10 LED array and our 7-segment display modules both require 5V to power them. To address this, we decided, in the April 2025 issue, to use a cheap-and-cheerful DC-to-DC converter module (see Photo 1).

These devices, available from Amazon UK (https://pemag.au/link/ac2m), accept input voltages from 9V to 24V and claim to deliver up to 5A at 5V. In reality, the ones we tested struggled to deliver more than 2A. Happily, we require only around 1.5A for our games console, well within its capabilities.

A high-level view of our proposed power supply setup is shown in Fig.1.

Bits and pieces

Before we leap into action, let's take a moment to remind ourselves where all the bits and pieces go (Fig.2). When you look at the front of the console, as illustrated in Fig.2(a), you may note that I've transposed the Action and Direction control buttons compared to our original prototype.

This wasn't intentional. It was only after I'd attached the 7-segment display modules that I realised I'd mounted them on the wrong side of the panel. I could have fixed this, but then I thought, "What the heck!".

When we look at the back of the console, shown in Fig.2(b), we see that I decided to mount the DC-to-DC converter on the back of the 14 × 10 LED array panel (I used trusty hot glue for this purpose).

Tips and tricks (and other stuff)

Don't forget that when you're working with semiconductor devices, including LEDs (like our 7-segment display modules and our 14 × 10 array), you should always use an antistatic mat and an antistatic wrist band. Of course, these need to be properly grounded. We discussed this in our first *Arduino Bootcamp* column (January 2023).

I'm assuming that you know how to solder. If not, may I be so bold as to recommend *The Basic Soldering Guide Handbook* (pemag.au/link/ac9o), written by my friend and fellow PE contributor Alan Winstanley.

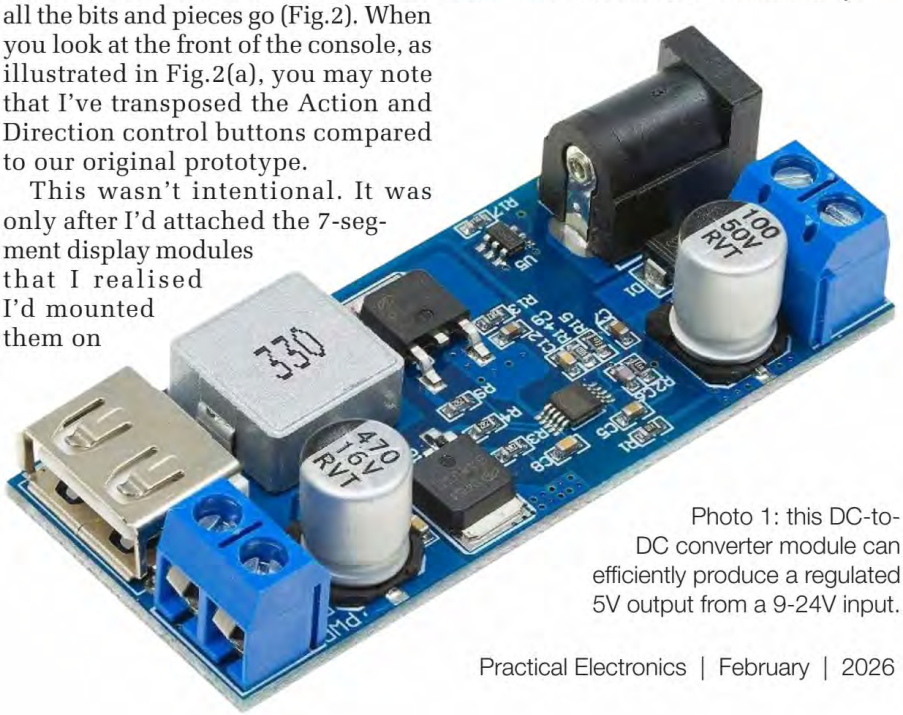


Photo 1: this DC-to-DC converter module can efficiently produce a regulated 5V output from a 9-24V input.

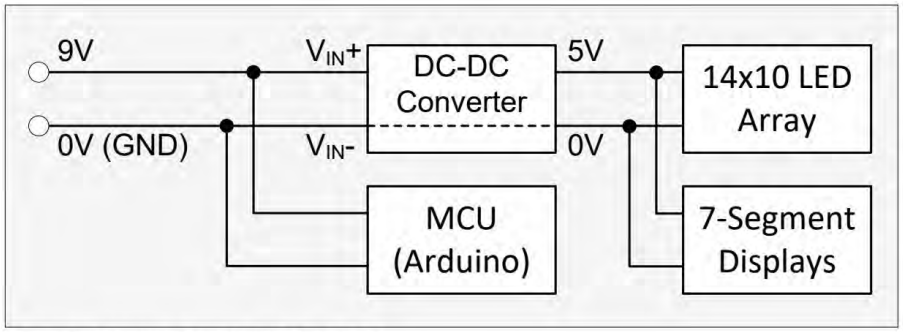


Fig.1: our power supply configuration.

Whenever possible, use heatshrink tubing to protect your soldered joints. For example, in a moment, we'll be soldering wires to a 25-way D-sub connector, and you should apply heatshrink tubing to each connection. This serves two purposes: it prevents unwanted short circuits, and it provides strain relief for the connections when the wires flex as you work with them.

When I first embarked on this project, I imagined a solid enclosure with a removable front panel. The more I thought about it, however, the more I realised what a pain that was going to be.

As you'll see in the photos of the finished unit later in this column, I eventually decided to permanently mount the front panel to the front of the enclosure, leaving the rear open, with just a band of material around the back edge. My plan is to add a proper back panel later, but for now, this open rear gives me full access to wire everything from the back.

Delightfully different

Back in the 1980s and 1990s, most home game consoles were built around the concept of plug-in game cartridges. This was long before the age of the internet, digital distribution or downloadable content. If you wanted a new game, you bought a physical cartridge, slid it into the slot, and—assuming you'd chanted the correct incantations—off you went.

In those days of yore, commercial game cartridges didn't include processors. The main processor was built into the console itself. The cartridges were essentially clever little boxes of read-only memory (ROM) that held the game code and data. When you inserted a game, the console's processor executed the code directly from the cartridge's ROM.

Having said that, some cartridges did get a bit fancy. To work around the limitations of the host console, manufacturers occasionally added extra hardware into the cartridges, such as bank-switching 'mapper' chips, additional RAM, custom audio chips, or even the occasional co-processor (the most famous being the Super FX chip used in Star Fox for the Super Nintendo).

Even so, the core idea remained the same: the console provided the brainpower, while the cartridge provided the instructions.

We dare to be different. In our case, each plug-in game cartridge will contain its own Arduino microcontroller with the game code stored in its flash memory, thereby making it an independent computing node. In this case, the console itself becomes more of a shared backplane that supplies power, buttons, and displays. The cartridge plugs in, commandeers the hardware, and away we go.

The bottom line is that, although

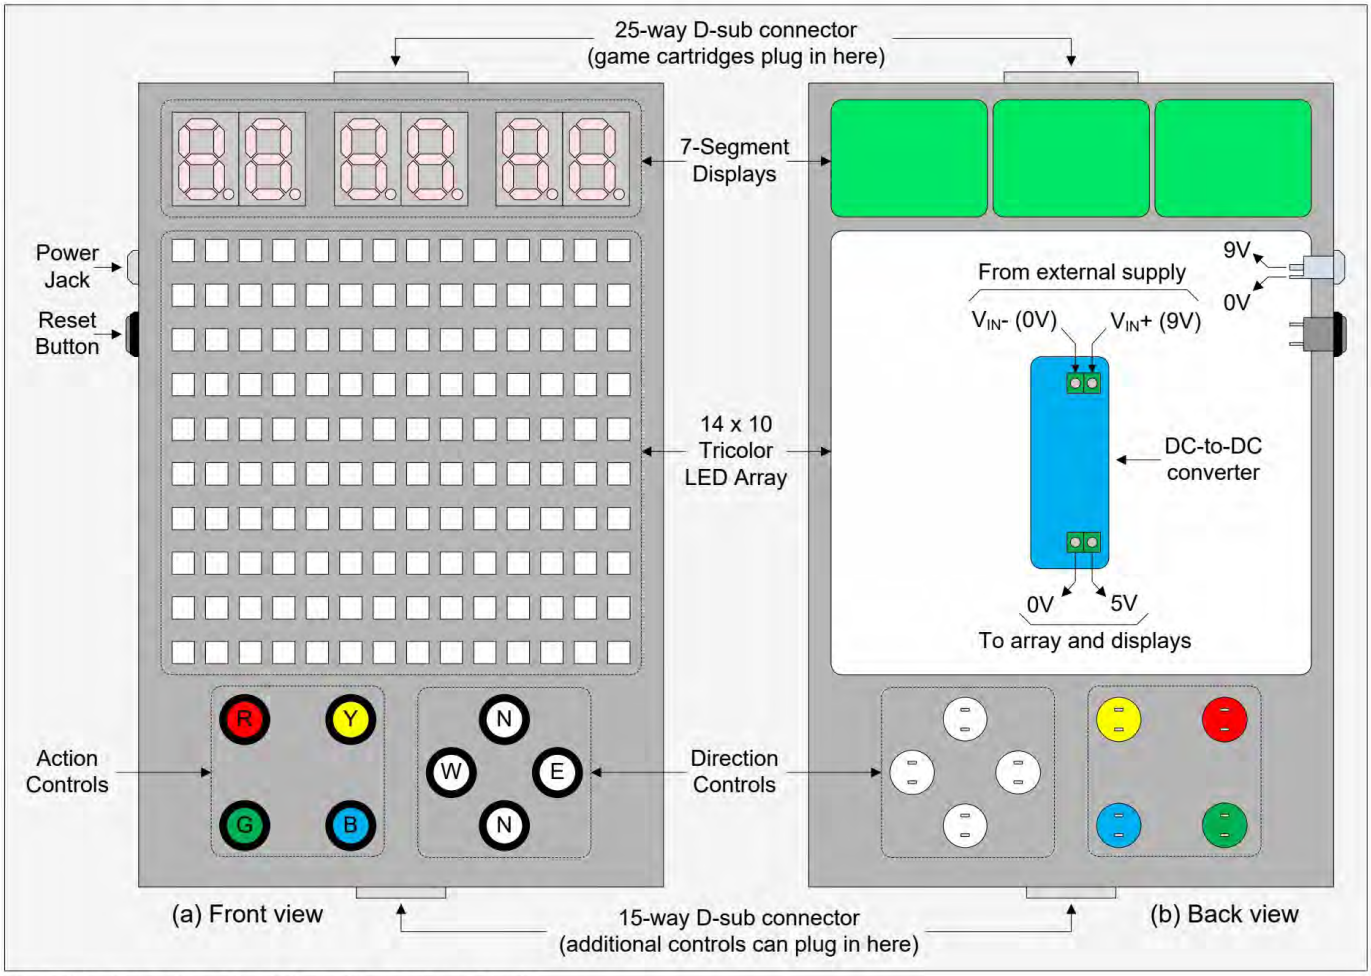


Fig.2: the main elements of the retro games console.

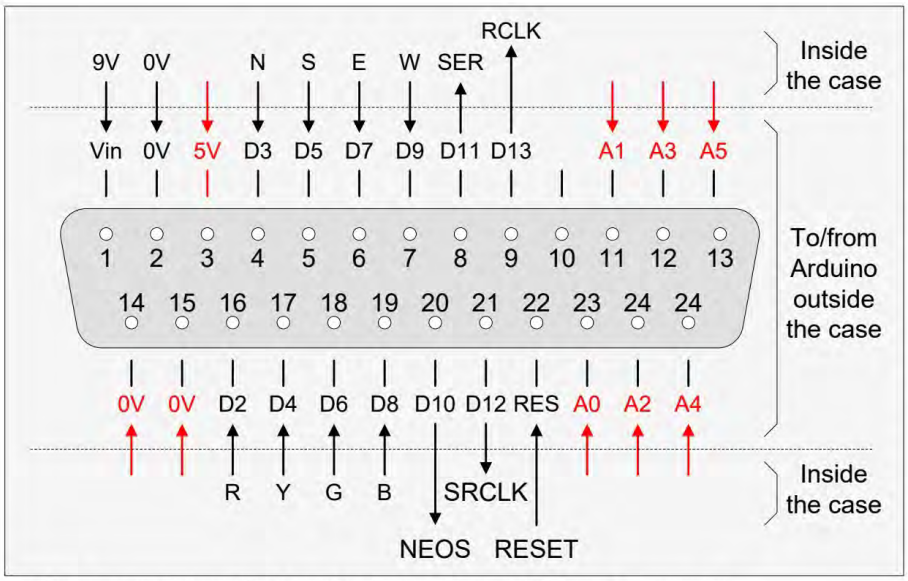


Fig.3: the connections that need to be made on the 25-way connector.

our design isn't a faithful recreation of how retro game cartridges worked internally, it does capture the spirit of slotting in a physical game module and instantly transforming the experience. And, honestly, that's half the fun.

Feeling connected

Let's start with the 25-way D-sub connector mounted at the top of the enclosure (Fig.2). If we imagine ourselves peering at it from the inside of the enclosure, the pinout we'd see is shown in Fig.3. In this first pass, I'm only going to wire (and use) the connections shown in black. I'll add the connections shown in red at a later date.

The 9V and 0V (GND) power supply signals inside the enclosure will be wired to pins 1 and 2 on the connector, respectively. In turn, these supplies will be fed to the Arduino's V_{IN} and 0V (GND) pins on the game cartridge.

The north (N), south (S), east (E) and west (W) direction control switches inside the enclosure will be wired to connector pins 4, 5, 6, and 7, respectively. In turn, these signals will be fed to the Arduino's digital D3, D5, D7 and D9 pins on the game cartridge.

The red (R), yellow (Y), green (G) and blue (B) action control switches inside the enclosure will be wired to connector pins 16, 17, 18, and 19, respectively. In turn, these signals will be fed to the Arduino's digital D2, D4, D6 and D8 pins on the game cartridge.

The black reset button on the side of the enclosure will be wired to pin 22 on the connector. This signal will be fed to the Arduino's **RESET** input on the game cartridge.

Meanwhile, the Arduino will control the 14 × 10 array of tri-colour LEDs (aka Neopixels) using its digital pin D10, which will be fed from the cartridge to pin 20 on the connector. In turn, we will wire this pin to the input of our array.

Similarly, the Arduino will control the 7-segment display modules using its digital pins D11, D12, and D13, which will be fed from the cartridge to pins 8, 21, and 9 on the connector. In turn, we will wire these pins to the **SER** (serial data in), **SRCLK** (shift register clock), and **RCLK** (output register clock) pins on our display modules, respectively.

I messed up here because I started by jotting down my connection plans on a piece of paper, which eventually became Fig.3. I then wired things up as per my jottings (Photo 2). It was only after I'd finished that I realised I'd transposed the **SRCLK** and **RCLK** signals on my original diagram.

That is, in our earlier programs, we used to drive each of these signals with the other's pin (if you see what I mean). Fortunately, the solution to this was easy-peasy lemon squeezy; all I had to do was swap these pins in my code, but we will return to that later.

For future reference

For this initial trial, I didn't bother wiring up the 15-way D-sub connector mounted at the bottom of the enclosure. As you may recall, this is intended to allow additional controls and sensors mounted on a mini console to be plugged in at some point in the future.

You can, of course, wire your system up as you wish. For the sake of interest, the way I'm currently planning to wire my 15-way connector is illustrated in Fig.4. Once again, I'm imagining that we are peering at this connector from inside the enclosure.

The way I'm envisaging things, we would feed 5V and 0V (GND) from the main console to power anything that needs powering on the mini console.

We could have a second set of action and direction buttons mounted on the mini console, perhaps for use in two-player games. Alternatively, we could mount two four-way momentary-contact joystick switches on

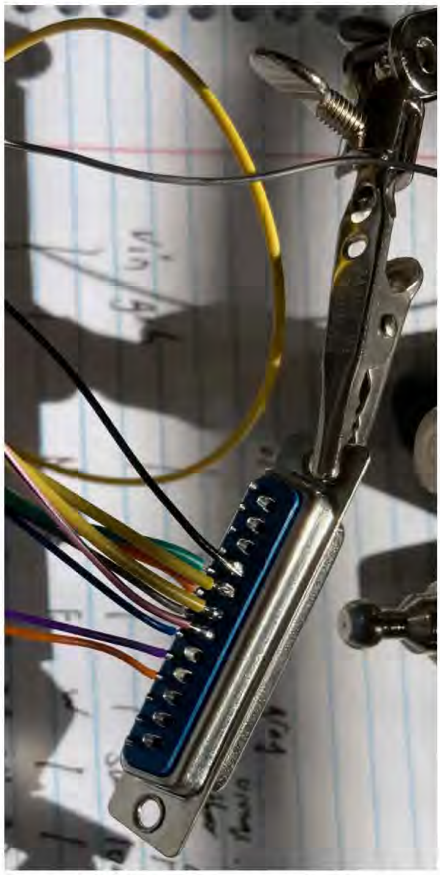


Photo 2: just call me a soldering fool.

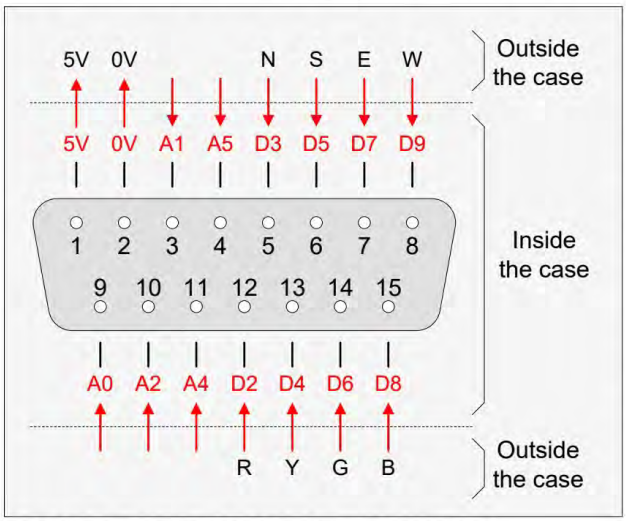


Fig.4: the connections to the 15-way connector.

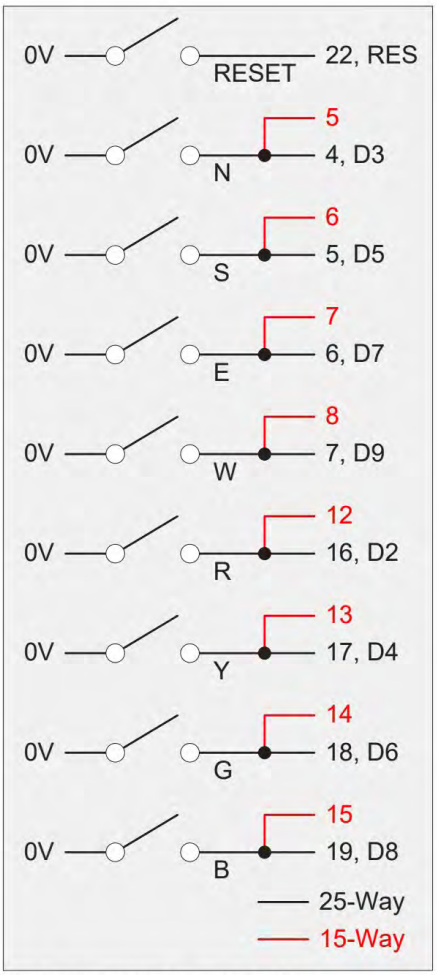


Fig.5: the console switch wiring.

the mini console. In both these cases, the signals from these switches/joysticks would be wired in parallel to the switches inside the main console.

Instead of using 4-way momentary-contact joystick switches, we could add two two-axis analog joysticks to the mini console. Each of these joysticks would have two analog outputs. We could feed these to the A0, A1, A2 and A4 pins on the connector.

Inside the main console, these pins would be wired to the corresponding A0, A1, A2, and A4 pins on the 25-way connector, which would then feed them into the processor on the game cartridge.

Did anything catch your eye in Fig.4 and in the preceding paragraph? You're right; there's no A3 pin. That's because I ran out of pins on the connector. We could move to using a bigger connector, but I'm happy with the 15-way one we currently have.

So, why do I show A0, A1, A2, A4, and A5 in Fig.4 (as opposed to A0, A1, A2, A3, and A4)? The reason is that on the Arduino Uno R3 (which we are using), the primary I²C pins are A4 (the data line, SDA) and A5 (the clock line, SCL).

Maintaining access to these lines means we have the ability to deploy all sorts of I²C sensors (such as a three-axis accelerometer and/or a three-axis gyroscope) via our mini-console.



Photo 3: wiring is fun… …but not as fun as finishing!

As one final point, just as we can create multiple plug-in game cartridges, we can also create multiple plug-in mini consoles. The world truly is our lobster (or crustacean of your choice).

The old switcheroo

Last, but certainly not least, we have to wire our switches to the 25-way connector (and the 15-connector, at some stage), as illustrated in Fig.5.

Switches are usually reliable. It's relatively rare in the scheme of things that you end up with a duff switch. On the other hand, it can happen, so it's a good idea to use the continuity tester on your multimeter to verify that your switches work before you attach them to the console (or, at least, before you start soldering wires to them). This takes only a minute, and it can save much hair-pulling later.

The result of my wiring frenzy is shown in Photo 3.

A plan of action

Once all the wiring was completed, the next thing I did was to create a plan of action ("If we don't have a plan, we're no better than the French", as the old saying goes). *[No doubt there's a French saying that translates as "If we don't have a plan, we're*

no better than the English" – Editor]
There are three main aspects of the console that we need to test before proceeding further:

- The 7-segment displays
- The 14 × 10 array
- The switches

I decided to start with the low hanging fruit in the form of the 7-segment displays. I took the final program from our previous column (January 2026)—the one that displayed a level from 1 to 9 coupled with a score counting up from 0000 to 9999—and I tweaked the pins to match my wiring 'gotcha' as discussed earlier in this column.

This means we now have the SER, SRCLK and RCLK signals connected to Arduino pins 11, 12, and 13, respectively.

I uploaded this new program into the Arduino before I connected it to the console, just to be sure I had a clue what it would do when I powered everything up. I've made a copy of this updated version of the code available in the file named **CB-feb26-code-01.txt** (as usual, all files mentioned in this column are available as part of the February 2026 download package from the PE website at https://pemag.au/link/ac9p).

A roll of drums

In our next column, we will create our Arduino-Uno-based game cartridge. In the meantime, I just happened to have a DB25 male-to-bare-wire connector in my treasure chest of parts (Photo 4).

These are available on Amazon (https://pemag.au/link/ac9q). The great thing is that all the wires are different colours, and each end has a little band with the pin number printed on it. They are quite expensive, so I wouldn't suggest buying one just for this project. On the other hand, if you happen to have one lying around...

I chopped the existing pins off the ends of the wires (they were too big to plug into my Arduino's headers), used my crimping tool to attach appropriate pins, and then plugged them into my Arduino. Next, I plugged the other end of the cable into my console. Remember that the Arduino will now be powered from the console, but we'll still need to use the USB cable to upload future programs into it.

This is the point in a project when one's heart is in one's mouth. When you apply power, will everything work as planned (hey, it could happen!), will your creation simply sit there sadly, or will it add a little excitement to your life by bursting into flames? When you've been at this game as long as I have, you'll have sampled the delights of all three scenarios.



Photo 4: a handy-dandy DB25 male-to-bare-wire connector.

The first thing I did was power up my bench power supply and set the voltage to 9V and the current to 2A. I connected the supply to the console, then I pressed the "Output" button. I was poised to power down if I heard or smelled anything untoward and... nothing happened. "Oh, bother", I said (or words to that effect).

I looked at the displays on the bench power supply. The voltage showed 9V, but the current showed 0A. Then I looked lower and realised that the positive wire had somehow become disconnected from the supply. Phew! So I powered it down, reconnected the errant wire, and repeated the power-up process.

It worked! The count commenced as expected. I was wearing my happy face. I pressed the reset button on the console. That worked too. I felt a roll of drums was in order, but there's never a drummer around when you need one, so I had to make do with an imaginary drum roll in my head (my ears are still ringing).

Testing the main 14 × 10 array

The next step was only a little more complicated. First, I took the latest and greatest program we created to drive our 14 × 10 array, from the May 2025 issue. This is the one that repeatedly draws randomly sized rectangles at random locations on the display. Furthermore, each rectangle has a random outline colour and a (different) random fill colour.

I changed the pin number driving the Neopixel string to 10 to match the wiring in the console. I also copied just enough of the 7-segment display test code to clear the 7-segment displays and set their control pins to their inactive states. If I hadn't done this, the control pins would have been left in high-impedance states, which would have resulted in the 7-segment

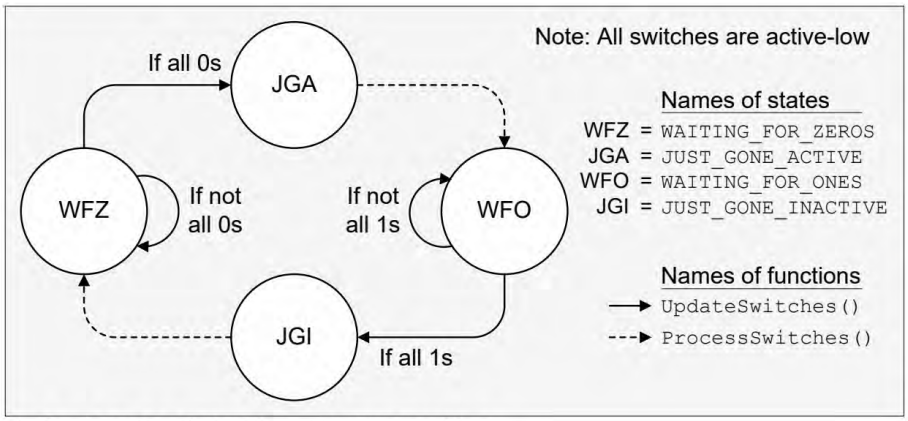


Fig.6: our new switch debounce state machine

displays flickering annoyingly with random values.

Oh Joy! This program also worked the first time. I was on a roll. I've made a copy of this code available in the file named **CB-feb26-code-02.txt**.

Testing the switches

My final test involved the switches. When we last looked at debouncing switches (August 2025), we were working with three breadboard-mounted momentary pushbuttons with a mix of active-high and active-low states. Since all the switches on our console are active low, I decided to update the state machine we are using. The new version is depicted in Fig.6.

Next, I took a copy of our latest switch debounce code (August 2025) and modified it to handle our eight switches using our new state machine.

VERY IMPORTANT: I don't know if you noticed in Fig.5 that we aren't using any pull-up resistors with our switches. In this case, we declare these pins as INPUT_PULLUP in our code, which activates the Arduino's internal pull-up resistors (yes, of course we could have done this before, but where would the fun be in that?).

Once again, I copied just enough of the 7-segment display test code to clear the 7-segment displays and set their control pins to their inactive states. Also, I copied just enough of the 14 × 10 array code to ensure that the array was initialised and deactivated. I've made a copy of this code available in the file named **CB-feb26-code-03.txt**.

As you may recall, in our switch handling code, we first call our UpdateSwitches() function, which cycles through all the switches, reads their states, and updates their state machines. Then we call our ProcessSwitches() function, which deals with any game-specific actions.

For example, the game-specific code associated with the north (N) switch in our current test program is shown in Listing 3(a). Remember, we're using the convention that the listing number (3 in this case) corresponds to the numerical part of the matching code file name ("03" here).

The way we've architected this means we can use the same code across all our future game programs. All we will need to do is modify the 'Optional (game-specific)' part of the code. For the moment, all we are doing is pressing the switches and ensuring that the appropriate messages appear in the serial monitor window. In my case, I'm seeing:

```
:
Switch North just went active
Switch North just went inactive
Switch South just went active
Switch South just went inactive
Switch East just went active
Switch East just went inactive
Switch West just went active
Switch West just went inactive
:
```

This is where we ensure the switches are wired to the correct pins on the connector. If we've made any mistakes, we can easily rectify them by tweaking the pin definitions in our code.

Next time

We've made great progress in getting the games console working this month!

Our next column will be the final one in this project (although we will still add new games and things in the future). We will start by building our first game cartridge and creating our first game. In the meantime, you could create a test program of your own.

How about a simple Etch-A-Sketch program? You could start by flashing a cell in the middle of the array with one of our four action button colours (red, yellow, green, or blue). Pressing one of the north, south, east, or west direction buttons could lock this colour into this cell and advance the cursor (the active flashing cell) in the specified direction.

Pressing one of the action buttons could change the selected colour.

If you have any thoughts on anything you've read here, please feel free to email me at max@clivemaxfield.com PE

```
147   // Switch North
148   if (Switches[SW_NORTH].swState == JUST_GONE_ACTIVE)
149   {
150     // Mandatory
151     Switches[SW_NORTH].swState = WAITING_FOR_ONES;
152
153     // Optional (game-specific)
154     Serial.println("Switch North just went active");
155   }
156   else if (Switches[SW_NORTH].swState == JUST_GONE_INACTIVE)
157   {
158     // Mandatory
159     Switches[SW_NORTH].swState = WAITING_FOR_ZEROS;
160
161     // Optional (game-specific)
162     Serial.println("Switch North just went inactive");
163   }
```
Listing 3(a): a portion of our new switch handling routine.

Devices and components used in this issue

Arduino Uno R3 microcontroller module	https://pemag.au/link/ac2g
Dual 7-segment display modules	https://pemag.au/link/ac8i
PCB spacer kit (black, M3)	https://pemag.au/link/ac4u
Pushbuttons (assorted colours)	https://pemag.au/link/ac2t
2.1mm panel-mount DC power socket	https://pemag.au/link/ac2v
2.1mm standard DC power plug	https://pemag.au/link/ac2w
25-way D-Sub panel-mount socket	https://pemag.au/link/ac2x
15-way D-Sub panel-mount socket	https://pemag.au/link/ac2z
22 AWG multicore wire kit	https://pemag.au/link/ac3m



Part 1 by Phil Prosser

POWER LCR METER



While we have published plenty of LC and LCR meters over the years, this one is quite different. It can deliver up to 30A to inductors to determine their properties at higher power levels. That makes it particularly useful for determining when and how an inductor saturates. It can also measure very low resistances and very high capacitances.

Inductors are the easiest of the basic components to make yourself. This is typically done by winding enamelled copper wire around a core or former. For air-cored inductors, you're generally only worried about the inductance (which can be measured with an LC meter) and DC resistance (measured with a low-ohms meter).

It's quite a bit more complicated for inductors with a core, though. Cores are typically made of ferrite, compressed powdered iron or mu metal, and they all behave quite differently at higher current levels. As the current through the inductor increases, eventually the core saturates and the inductance drops.

This meter will let you determine at which current the inductance starts to drop off and how fast it drops off. It isn't just handy for self-wound inductors; any that you recover from a piece of equipment will likely have unknown properties – this device can erase that mystery.

It can also measure very high capacitances and low DC resistances, which is handy for characterising the series resistance of any inductor, including air-cored types.

As an example of when this device might come in handy, if you use a ferrite or iron-cored inductor in a loudspeaker crossover, its inductance will fall as the current through it increases past its saturation point. The result of this is non-linear behaviour that will be heard as distortion.

High-quality speakers use air cored inductors because they do not suffer this problem. However, large inductors in speakers still often have ferrite or iron cores to manage cost, and the bulk and high resistance of a large coil of wire.

Another common application of power inductors is in switch-mode power supplies. At high currents, saturation in the core can reduce the inductance and degrade the performance of the power supply. While this tester will not characterise inductors at high frequencies, it is unusual in that it allows characterisation of inductors at very high currents, up to about 30A.

Our Power LCR Meter measures inductance, resistance and capacitance for larger power devices. For measuring low resistances, it will push up to 1A through the resistor, although only for a very short time.

This is not a general-purpose meter; it is for those chunky passives you are considering for your switch mode power supply, Class-D

Power LCR Meter Features & Specifications

- Measures capacitance from 50nF to more than 1F
- Measures resistance from 1mΩ to 300Ω
- Measures inductance from 50µH to 1H+
- Measures inductance saturation from 10µH to 1H at up to 30A (limited by internal resistance)
- Optional Kelvin probes for measuring low resistances
- Power supply: 12-20V DC at 1A

Power LCR Meter, part one

amplifier output or loudspeaker crossover. It will give you insight into your parts that you won't find in many other testers.

Operating principles

Rod Elliot describes a circuit that can be used to manually measure the saturation of power inductors on his website at https://sound-au.com/project250.htm (we've seen it elsewhere but his description of how it works is pretty thorough). The problem with that approach is that you need an oscilloscope to make the measurements.

This circuit uses an interesting technique to make that measurement, and in the deal we are forced to measure capacitance and resistance as part of the overall system, which makes for an unusual and capable device.

The concept is to monitor the transient behaviour after we apply a voltage step across the inductor, analysing the current through and voltage across the device under test (DUT) over time. Pretty much all the similar circuits on the internet that we found use a variation of the simple circuit shown in Fig.1.

We will avoid lots of maths here, but the following principles are used in this project. Most of us are familiar with the resistance equation (Ohm's law): V = IR.

The capacitor equivalent to this is C = q/v and its differential is C = (dq/dt) ÷ (dv/dt). Realising that dq/dt is simply current, we have C = I ÷ (dv/dt). For inductors, the equivalent formula is L = V ÷ (di/dt).

This circuit uses the property that if we turn that driving Mosfet on and apply 1V across a 1H inductor, we can expect to see the current increase at 1A per second. Similarly, if the inductor is 100µH and the rail voltage is a 10V or so, then we get the formula 10^{-4}H = 10V ÷ (di/dt), so we can expect di/dt to be 10^{5}A/s or 100mA/µs.

For a 10µH inductor with 10V applied, we can expect a rather lively 1A/µs rate of increase in current. This means that we will have to switch the Mosfet off pretty quickly after it turns on, or be ready to handle some very high currents after a few millionths of a second!

By measuring the rate of change of current through the inductor, we can measure its inductance. If we measure that repeatedly at a series of points on that current curve, we can monitor how the inductor behaves at different currents. Neat!

Rod Elliot makes a case that running this manually with a programmable pulse generator makes sense. We found ourselves in a situation somewhat akin to the serviceman looking at a broken widget that probably ought to be thrown in the bin. The temptation to 'have a go' at automating the measurements was irresistible.

Challenges this presents include:

1. Inductors have resistance, which really messes up the measurements if you ignore it.
2. The precise voltage across the DUT is important; if it droops during the measurement, you need to know by how much, or your results will be inaccurate.
3. It has to work over a wide range of inductor values and sizes.
4. We should ideally make sure that if someone connects an unexpected component, it won't explode.
5. How do we make measurements of a process that can be over in 50 millionths of a second, and even get multiple results in that short a time?
6. If we have an inductor being tested with 30A flowing through it, we have the potential for a massive back-EMF spike with significant energy once the test finishes.

Those first three challenges mean that this meter needs to be able to accurately measure the resistance of the DUT and also the test voltage applied, as this is essential to get a good measurement of the inductance. This meter needs to be able to tell whether the DUT is a capacitor or not. Only after determining these things can the meter then run the saturation current test.

To measure resistance and capacitance, we can use a current sink in a fairly conventional way, which involves driving a constant current through the device under test, not a constant voltage.

Dealing with the back-EMF that will be created when we switch off

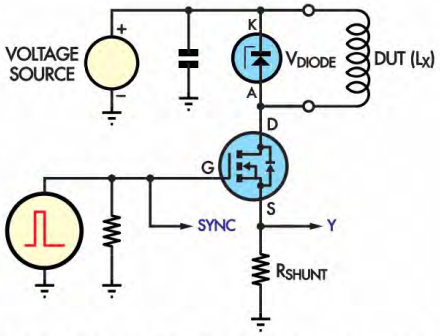
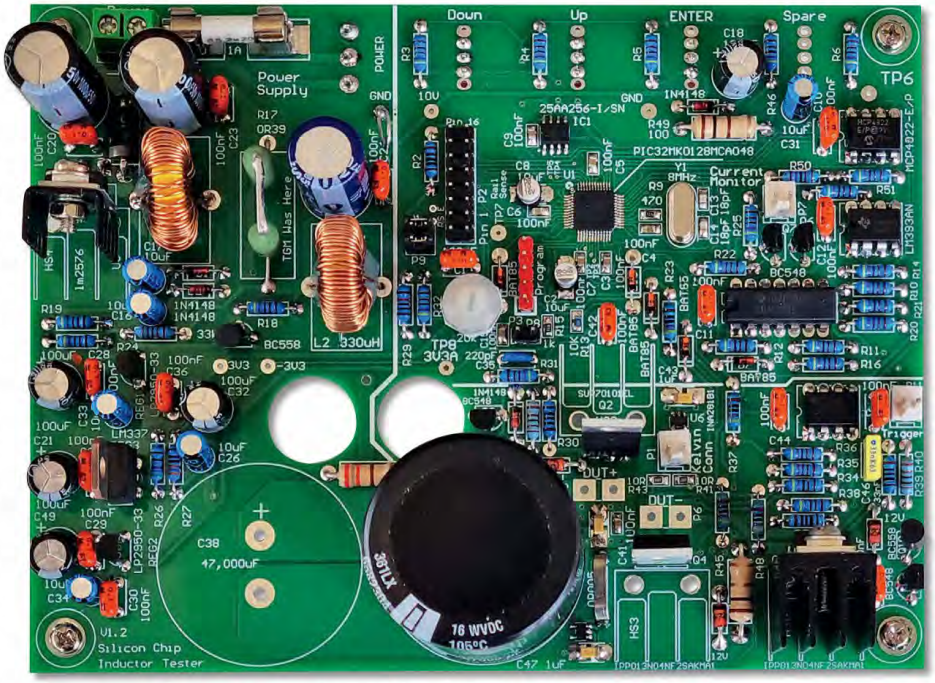


Fig.1: the basic principle of measuring inductance at various currents. A brief pulse applied to the Mosfet gate results in the voltage source being applied across the DUT, resulting in a current ramp through R_{SHUNT} that can be measured. The inductance can be calculated from its slope.



There are a few parts on the PCB but it isn't overly complex. The second large capacitor is optional and zener diode ZD13 is not required on the final board.

the Mosfet drive after an inductance test requires a reverse diode across the DUT. As it turns out, our circuit needs a Mosfet there to discharge capacitors while testing them, so we can use a P-channel Mosfet to deal with this.

With a bit of head scratching, we arrived at an arrangement that allows us to parallel the high-current saturation drive Mosfet with a constant current sink that's used for resistance and capacitance tests. Since we have a system to measure resistance and inductance, and also to sense whether the DUT is a capacitor or not, we might as well measure what the capacitor value actually is.

In Fig.1, the Mosfet basically shorts the DUT across the 10V supply rail. If we leave it inactive, we can dial up a current on a programmable current sink to make those other measurements. This leads to the arrangement shown in the block diagram, Fig.2.

The key system components are:

- A power supply capable of delivering 10V at up to 30A for brief periods
- The DUT
- A Mosfet that can connect the DUT between the power supply rail and the current sense circuitry. This includes a hardware-based current limiter.
- A discharging Mosfet that can apply a load across the DUT.
- A programmable current sink.
- Differential current and voltage sensing circuits for the DUT.
- The PIC32MK0128MCA048 microcontroller to manage all this (in a 48-pin TQFP package).

The PCB includes headers allowing the trigger timing and DUT current to be monitored on an oscilloscope so you can look at those waveforms, but the microcontroller samples all relevant signals and provides measured results. To see the waveforms on an oscilloscope, you need to set it on single shot and run a measurement, as we describe later.

The power supply operates from a DC plug pack. This is nominally 12-20V at 1A. The average current draw is not great, but tests will demand up to 1A from time to time.

The software implements four distinct algorithms, for measuring resistance, capacitance, inductance and the saturation of an inductor (plus low inductances).

Measuring resistance

In this mode, the Mosfets are left off and the constant current sink is switched on at ~10mA. We have two channels that can monitor the voltage across the DUT, one with a gain of 20 and one with unity gain.

We can measure up to 3.3V across the DUT, which is a maximum of 330Ω at 10mA. The software looks at the value determined by the voltage, and if the resistance is less than about 30Ω, it increases the current to 100mA. If the resistance is below about 15Ω, the high gain channel is used with a current of 10mA, and so on through to the meter driving 1A with a high gain.

At 1A drive with high gain, the maximum value is 0.165Ω (165mΩ) and with the 12 bits of analog-to-digital converter (ADC) resolution, we should be able to resolve under 1mΩ. You will need to be using the Kelvin probes to measure resistances down at this level.

The precision of these measurements is a result of the current sink and its calibration, the differential amplifier, the Kelvin probes and the ADC itself. If you are reasonably careful with your current calibrations and use 1% resistors, you will see accuracy in the region of a few percent.

Measuring capacitance

The software can use the same constant current sink along with the discharging Mosfet to determine if there is a capacitor on the DUT terminals. It does this by discharging the DUT by shorting the terminals, then feeding current to the DUT for a short period, then monitoring the voltage across the DUT after this current is removed.

If the DUT is a resistor or inductor, the voltage will rapidly fall to 0V. In fact, for an inductor, the back-EMF will generate a negative voltage across the DUT. If the DUT is a capacitor, it will hold charge and the software will see this positive voltage.

We can control the magnitude and duration of the current applied to the DUT, and we have a pretty decent ADC that can measure the voltage across it. So the software can also read the capacitance.

To achieve this, the positive terminal of the DUT is connected to the positive rail and the current sink draws 10mA from the negative terminal. At the same time, the software switches the Mosfet across the DUT switches on. This discharges the capacitor we are measuring and also provides a path for the 10mA to flow.

After the current sink has stabilised, the software clears its measurement buffer and starts sampling at the maximum sampling rate of 3.75Msa/s. The software then switches the DUT discharge FET Mosfet off, allowing the DUT to start charging. If the capacitor exceeds a predefined voltage, the software stops sampling, switches the current off and discharges the DUT.

The data in the measurement buffer is similar to an oscilloscope trace of the capacitor (DUT) voltage. As shown in

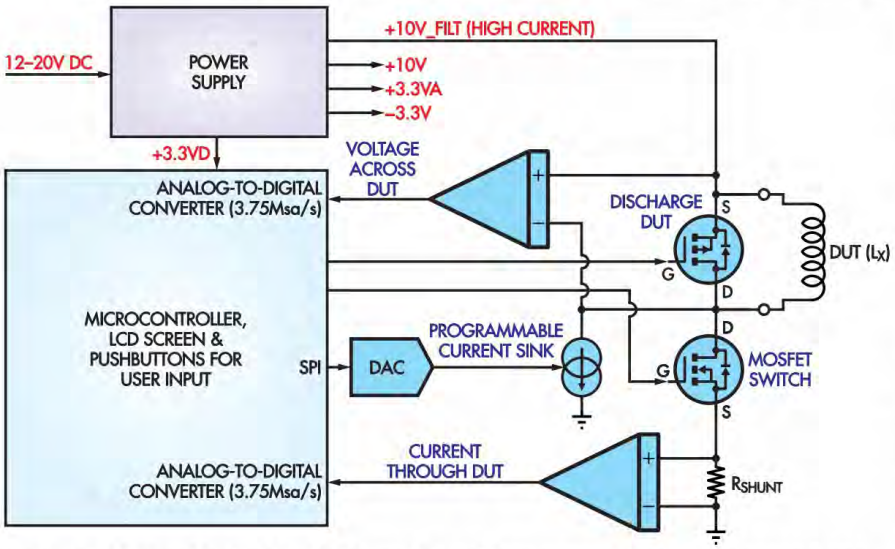


Fig.2: in addition to the Mosfet to switch voltage across the DUT, a second one can be used to discharge it (in case it is a capacitor). The programmable current sink allows for lower-current testing, with a DAC controlling the current level. Two differential amplifiers are used to monitor the voltage across the DUT and the current through it (via the voltage across the shunt).

Fig.3, the software looks for two points on this charging trace, V1 & V2. It also counts the number of samples between them (T1 & T2). From this, we can calculate dv/dt = (V2 – V1) ÷ (T2 – T1) and, knowing the current applied, we can calculate the capacitance from C = I ÷ (dv/dt).

This technique is a bit limited because, with 10mA flowing, a capacitance of 50nF will have a dv/dt of 10^{-2}A ÷ 5 × 10^{-8}, which is 0.2V/µs. Our ADC has a full-scale voltage of 3.3V, which means that the total charging time is about 16µs. Our dv is actually 2.2V if you dig into the software, which means dt is 11µs. At 3.75Msa/s, this is only 41 samples.

Further, the power devices in the circuit have some pretty substantial self-capacitances that we have to calibrate out in software. So we have settled on 50nF as a practical lower capacitance limit.

What happens if a big capacitor is connected? Our software data buffer is 12,800 samples long, which means we can measure a capacitance with dv = 2.2V and dt = 12,800sa ÷ 3,750,000sa/s or 3.4ms. This gives a maximum capacitance of 15µF or so.

Luckily, our software can look at the measurement buffer and see that we have not achieved our preferred V2 threshold, then reduce the sampling rate and rerun the test. Sampling rates of 3.75Msa/s, 375ksa/s and 37.5ksa/s are used. If a big capacitor is being tested, we can then increase the test current to 100mA and then 1A. This gives us an upper measurement limit of 1.5F.

Discharging a huge capacitor from 2.2V down to 0V requires a little caution. The software does this by pulsing on the Mosfet across the DUT, starting with 1µs pulses and increasing them until the Mosfet is fully on. This is intended to discharge large capacitors without creating massive current spikes.

Similar to the resistance measurement system, the resolutions of these measurements are good. Parasitic capacitance, slew rate limitations and suchlike limit the precision below about 100nF. From there up, the meter will provide a measurement accuracy of a few percent.

Measuring inductance

The meter has two approaches to measuring inductance. Both use the property of applying a voltage to the DUT and measuring the rate of change of current.

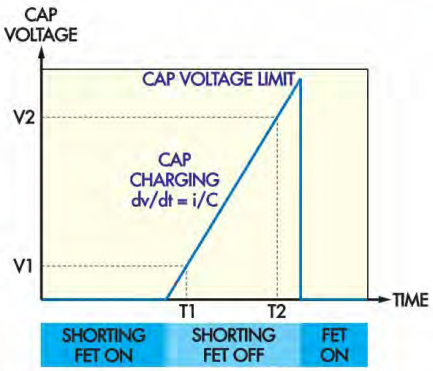


Fig.3: measuring a capacitor value involves first discharging it, then applying a fixed current and measuring the rate of voltage rise.

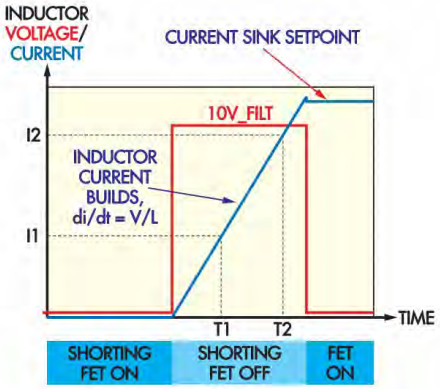


Fig.4: measuring inductance is similar to capacitance, except that we are applying a fixed voltage and measuring the rate of current rise.

The simple inductance measurement uses the constant current sink. As shown in Fig.4, it is similar to how we measure capacitance in that we start by setting up the constant current sink with the DUT discharge FET switched on. Then, when we are ready, we switch it off and monitor the voltage across the inductor (DUT) and also the current flowing through it.

Keep in mind that the "constant current sink" is really a current-limited constant voltage. That means the current sink is saturated and switched on hard right until the end of the test.

We have chosen this approach for initial inductance measurement as we know that the current will be controlled to the limit set by the constant current sink. If a user attempts to test an extremely low inductance, or a short circuit, the DUT will be subjected to a brief current pulse that grows to 1A and runs for no more than 12,800 samples at 3.75Msa/s, or 3.4ms.

An inductor's key property is that it 'resists' changes to the current flowing through it, hence that di/dt = V/L property. So what happens if we switch a constant current sink on across an inductor?

The current starts at zero, then immediately after the shorting Mosfet is switched off, the current is still zero. The constant current sink is on hard, applying the full 10V across the inductor, with no current (yet) flowing. Remember that equation, di/dt = V/L? Now V = 10V, and the current through the inductor grows at a rate set by the inductance.

This increase in current continues linearly. Once the current through the inductor reaches the current sink's set point, it starts throttling back to maintain the current at a constant value. So di/dt goes to zero, and the inductor current is constant, with notionally 0V across the inductor.

Our software in this test captures a series of readings of both the voltage across the inductor and the current being through it. The software switches the current measurement ADC to high-gain mode, which uses a 1Ω shunt. We start with the maximum sampling rate, which allows us to measure the smallest inductors.

On this test, the minimum practical measurement is about 50µH, which results from the 1A test current; di/dt = 10V ÷ (50 × 10^{-6}H), which is 0.2A/µs. Our cutoff current is 1A, so we get 5µs of data before the current limit is reached. The op amp takes a while to respond and the current overshoots quite a lot, so we actually get somewhat more than this to work with.

If you have a smaller inductor, the saturation test mode (see below) will measure down to about 10µH.

We capture two sets of data: the voltage across the inductor and the current through it. Similar to the capacitance test, when the voltage across the inductor transitions from close to 0V to full-scale on our ADC, we know the pulse has started. When this voltage falls again, we know the maximum current has been achieved.

If the software does not find the voltage falling before the end of the buffer, we know we need to reduce the sampling rate. The minimum sampling rate is 37.5ksa/s, which allows a minimum di/dt of 0.29A/s (0.1A ÷ [12800sa ÷ 37500sa/s]). This allows the device

to measure very high inductances, in the Henries range.

The software uses only the mid-section of the current vs time curve to calculate the inductance, between 25% and 75% of the buffer. This means this inductance test result is at about 0.5A.

The voltage across the inductor might not fall right down to 0V once the current through the inductor reaches the limit because real inductors have resistance.

If there is a DC resistance of say 1Ω, once we reach 1A, there is 1V across the inductor. We can easily get around this in software by changing our detection threshold voltages. However, the voltage drop across that internal resistance affects the measured inductance.

We need to consider the internal resistance of a component like an inductor as a property of the device. We can represent a real inductor as several ideal components, as shown in Fig.5.

We ignore R2 in our meter, as this is the equivalent of a resistance 'shorting' your windings. In real-world circuits, especially tuned LC filters, such a resistance may be intentionally added to dampen the circuit, but it is generally not significant in normal devices.

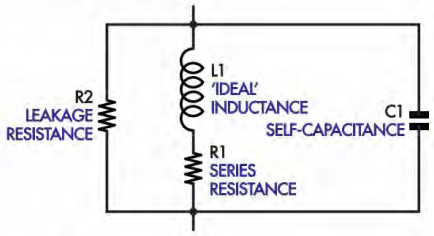


Fig.5: even ignoring core saturation, a real inductor can be modelled as four ideal components. It's the self-capacitance that is most troublesome.

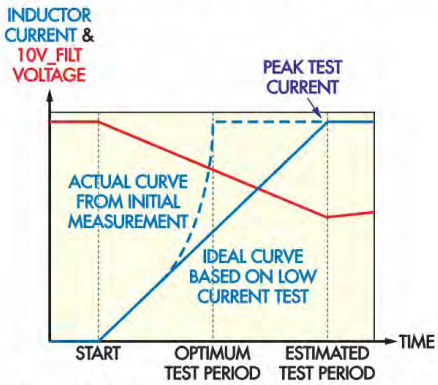


Fig.6: the current can rise higher than would be expected based on a low-current test due to core saturation. The software takes this into account.

C1 is a 'lumped parallel capacitance'. This is most commonly the result of capacitance between the windings in a coil; an iron-cored inductor can also have capacitance between the windings and the core. Our meter does not seek to correct for this in the measurement, as the errors resulting from it are not significant.

However, we see the impact of this when we apply a voltage across large coils, as the parallel combination of C1 and L1 causes visible ringing in the current in some cases. If you look at the data sheet for a commercial inductor, you will often see a 'self resonant frequency' figure; this capacitance plays in that characteristic.

You can see some of this ringing in Scope 1, right at the start. This plays havoc with inductance estimation!

R1 in Fig.5 is significant. This is the internal resistance we are concerned about. Our equation for di/dt = V ÷ L applies to only L1 in the figure; the voltage dropped across R1 is excluded from this. As the current flowing in the coil creates a voltage drop across the internal resistance, the effective voltage across L1 decreases.

For a real coil, di/dt reduces as the current increases. This is clearly visible for a large air-cored inductor, which has a DC resistance of 0.46Ω, shown in Scope 2.

When measuring inductance, the first thing our software does is to measure the DC resistance of the DUT. When calculating the inductance of the DUT, the software uses this as a correction factor; with some inductors, this correction is very significant.

Measuring inductor saturation

The saturation test will give you insight into the inductance's behaviour as a function of current.

This test uses the same principal as above, but this time we are not using a constant current sink. Instead, we will be connecting the DUT directly across the two 47,000µF capacitors using the main switching Mosfet, with our hardware current detector switching the Mosfet off when our pre-programmed limit is reached.

Again, the first thing the software does is to measure the DC resistance of the DUT. This is crucial, as the DC resistance tells us the maximum current that can flow through the DUT with our 10V across it. The software selects the current limit as 50% of the theoretical maximum, as for high-value inductors, we expect the 47mF capacitors to discharge significantly during the test.

The software then measures the inductance of the DUT using the constant current technique. This gives us a pretty optimistic value of inductance at high currents for all but air-cored inductors. The software uses this to calculate the time required for this inductor to reach the peak test current, and the sampling rate is adjusted to fit this into our sample buffer. Fig.6 shows what is happening here.

The software then checks to see if the user has a capacitor on the DUT terminals. While we never managed to damage anything during development, it isn't a great idea to suddenly apply 10V to a potentially large capacitor.

The ADC inputs are set to monitor the high-current (low gain) measurement, which can measure up to 33A, while the second ADC channel monitors the 10V rail, which we know will droop throughout the test.

From a first run, the software looks to see if the sampling rate was OK. If the DUT has saturated early, we increase the sampling rate to get a closer look at the saturation curve. So a second set of samples is taken with an optimised sampling rate, which fills our buffer with usable data.

Once we have this data, the software splits this into 10 sections and calculates the inductance for each of the 10 regions. This set of results is stored to allow the user to scroll through.

Our ADC has 12 bits of resolution, which gives a maximum of 4095 discrete current measurement values. If we had more than about 10 time slices, our current measurements would introduce quantisation errors and any noise on the measurements would become more significant. On the other hand, with fewer slices, we wouldn't get as good an idea of the inductor's behaviour. We decided that 10 readings is the best compromise.

We initially intended for the software to estimate the saturation current from this data. If you take a look at some of the sample curves we provide later on, you might get a sense of the challenge this presents. The shape, rate of collapse of inductance and all sorts of 'interesting' effects would need to be considered.

We were concerned that if the software throws up a guess, it will take

on a credibility it does not deserve. A human can scroll through the numbers and easily see where the inductance rolls off and how fast it falls. The software sets the second measurement to "100%" as it is normally the cleanest data point, and all others are relative to this.

Real measurement examples

Let's look at why saturation matters in inductors. We will include some plots of the inductors we used in testing this device, as the behaviours illustrate not only why this parameter is important, but also why it is hard for the meter to give a simple answer to this value.

In each of the following oscilloscope plots, the blue/cyan curve is the current, which is at a scale of 100mV per amp. So 3.3V is the full scale of the ADC at 33A. The yellow trace is the trigger signal. The duration of these sweeps all vary, as that is a function of the DC resistance of the coil, and thus the inductance and the sampling rate the software selects.

Scope 1 and Table 1 are for a Bourns 2200LL-470-V-RC rated for 20.9µH @ 10.3A. It starts with an inductance somewhat above its rating and falls relatively slowly.

Scope 2 and Table 2 are for a 1.8µH air-cored inductor. This shows current curving downwards as a result of the DC resistance of the inductor.

Scope 3 and Table 3 show the values for a 550µH air-cored inductor. There is some variation in the measured inductance; given this is an air cored inductor, this is due to measurement errors in the meter.

Scope 4 and Table 4 are for the secondary of a Dick Smith M-2156 transformer. This is typical of the saturation in a soft-iron-cored device. The measured values show the remarkable collapse in inductance past saturation.

Scope 5 and Table 5 are for an Altronics L6630 470µH 5A inductor. This shows that this device behaves as specified at the rated current, but the initial inductance is substantially higher, and the inductance rolls off in a reasonably controlled manner.

That step at the start of the curve is the INA181 differential amplifier slewing to 'keep up' with the rate of change; this forms a real limitation on the lower inductance we can measure.

Scope 6 and Table 6 are for a Bourns 2200HT-100-V-RC rated for 7.9µH @ 11.2A. This shows the inductance starts somewhat above the rating and falls relatively slowly. That glitch on the initial application of the test voltage is from inter-winding capacitance, and demonstrates why we do not try to measure inductance from the start of the sample set.

Circuit details

Now that we know how the device works, let's look at the full circuit. We have broken it up into three separate diagrams that perform logically distinct functions: the test circuitry (Fig.7), the control circuitry (Fig.8) and the power supply (Fig.9). We'll start by examining the test circuitry, which is where most of the complexity lies.

The DUT connects between CON5 and CON6, to the left of centre in Fig.7. This places it between the +10V_FILT high-current supply rail and Q4 & Q5.



Current	Inductance
2.2A	44µH
3.4A	41µH
4.7A	35µH
6.2A	32µH
7.9A	27µH
9.9A	22µH
12.4A	18µH
15.4A	14.2µH
19.3A	10.4µH
24.4A	7.7µH

Scope 1: testing a Bourns 2200LL-470-V-RC inductor rated for 20.9µH @ 10.3A.



Current	Inductance
1.0A	1.8mH
1.9A	1.8mH
2.8A	1.8mH
3.7A	1.8mH
4.5A	1.8mH
5.2A	1.8mH
6.0A	1.8mH
6.6A	1.9mH
7.2A	1.8mH
7.8A	1.8mH

Scope 2: testing a 1.8µH air-cored inductor with a relatively high DC resistance of 0.55Ω.



Current	Inductance
1.2A	555µH
2.4A	559µH
3.6A	570µH
4.7A	566µH
5.7A	559µH
6.6A	575µH
7.5A	544µH
8.4A	512µH
9.2A	582µH
10A	564µH

Scope 3; testing a 550µH air-cored inductor with a 0.44Ω DC resistance.

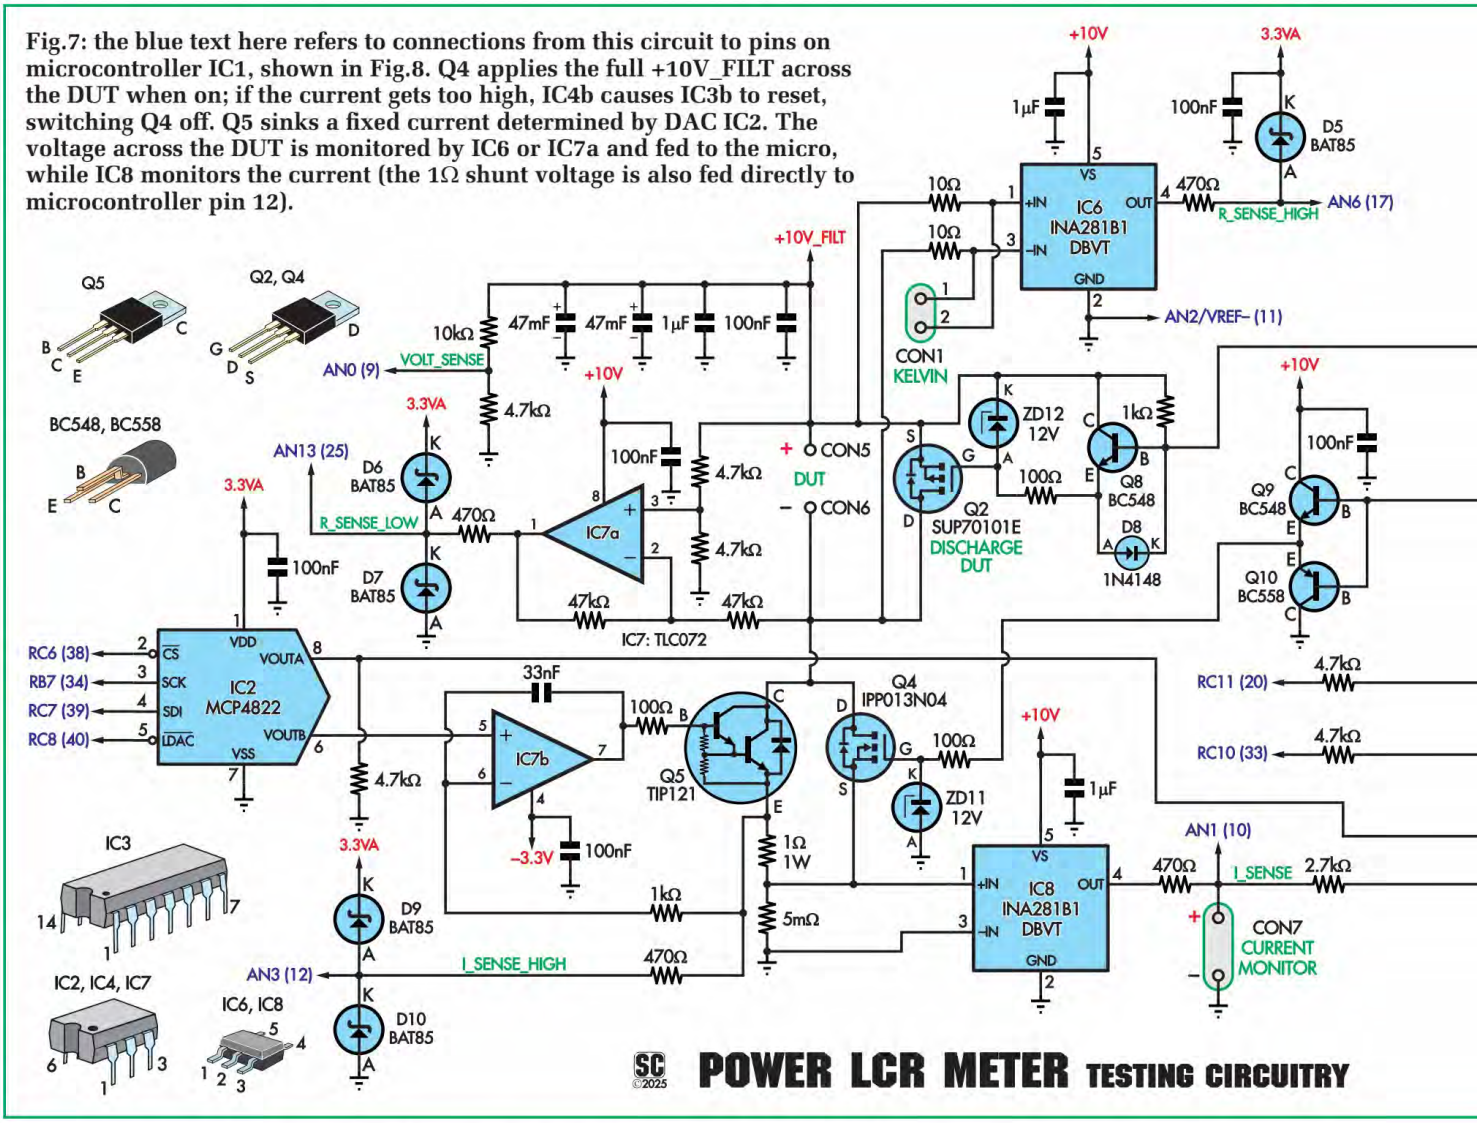


Fig.7: the blue text here refers to connections from this circuit to pins on microcontroller IC1, shown in Fig.8. Q4 applies the full +10V_FILT across the DUT when on; if the current gets too high, IC4b causes IC3b to reset, switching Q4 off. Q5 sinks a fixed current determined by DAC IC2. The voltage across the DUT is monitored by IC6 or IC7a and fed to the micro, while IC8 monitors the current (the 1Ω shunt voltage is also fed directly to microcontroller pin 12).

Q4 is the main switch that connects the bottom end of the DUT to ground via the low-value 5mΩ shunt, while Q5 is the constant current sink, in combination with an extra 1Ω shunt and op amp IC7b.

Q4 is an Infineon IPP013N04NF-2SAKMA1. These are relatively inexpensive but only have 1.3mΩ of on-resistance and can carry up to 197A continuously. We don't suggest you substitute this part, but if you must, most Mosfets with under 10mΩ of on-resistance and a continuous current rating of at least 50A should be OK.

The gate drive for Q4 comes from transistors Q9/Q10, which form a totem pole drive. This buffers the very high gate capacitance of the Mosfet, speeding up the switch-on and switch-off times. Q9 & Q10 are controlled by a 4013 D-type flip-flop (IC3b).

This enables us to switch off the Mosfet rapidly once the current through it exceeds a programmed threshold. To achieve this, we drive the reset pin of our flip-flop from a dedicated comparator (IC4b) which compares the current reading to a programmed threshold.

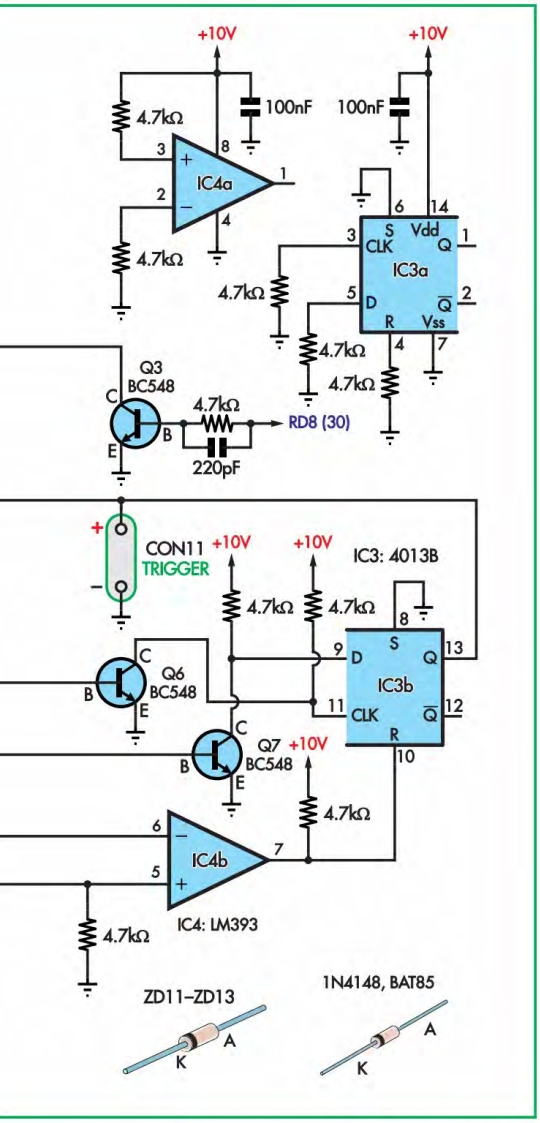
IC2 is an MCP4822 12-bit DAC. One of its outputs (output A, pin 8) drives the negative input of this LM393 comparator, while the current sense line from IC8 (INA281) drives the positive input. The INA281 has a 5% settling time of 1µs, the LM393 has a response time of 1µs, and the 4013 reset-to-Q time is 130ns (0.13µs).

The Q output of IC3b controls the Mosfet, so we have an overall delay from sensing the current to switching the Mosfet drive off in less than 5µs. We can program the MCP4822 to produce voltages from 0 to 2V, allowing us to set a current limit from 0 to 30A.

The microcontroller sets flip-flop IC3b using control signals from its RC11 and RC10 pins that are level-shifted by transistors Q6 & Q7 and applied to the CLK and D inputs of IC3b, respectively. This allows the micro to set the flip-flop and enable Mosfet Q4, but the reset signal from the over-current detection circuitry can always override this and switch the Mosfet off.

Mosfet Q2 is used to discharge the DUT. This is a Vishay SUP70101EL P-channel Mosfet. It is included to allow capacitors to be discharged before testing. It also provides a reverse current path for an inductive DUT at the end of tests. If the voltage across the DUT reverses, up to 30A can flow through the body diode in Q2.

This device needs to be able to handle the maximum saturation current, as at the end of a saturation test, Q2 is



switched on, and it will carry the full current through its body diode until it decays. This will result in energy being dissipated in Q2.

The SUP70101EL is rated at 100A, with an $R_{DS(ON)}$ of 10mΩ. If substituting, we would stick to similar ratings, but this is not an expensive part for what it does.

As Q2 is a P-channel device, its gate drive is relative to its positive source terminal, which connects to the +10V_FILT rail. Thus, there is a level-shifter based around transistors Q3 and Q8 that allows us to drive it from a 0-3.3V microcontroller output pin (RD8). Diode D8 speeds up its switch-off.

The programmable current sink is included to allow us to measure the resistance of the DUT. We generally want to run this at 10mA, 100mA or 1A. The actual current is programmed by the other DAC channel, output B (pin 6) of IC2.



Current	Inductance
1.1A	15.9µH
1.5A	10.1µH
2.1A	6.8µH
3.0A	4µH
4.4A	2.7µH
6.3A	2.0µH
8.4A	1.7µH
10.8A	1.5µH
13.1A	1.3µH
15.5A	1.1µH

Scope 4: testing the secondary of a Dick Smith M-2156 transformer with a soft iron core (DC resistance of 0.28Ω).



Current	Inductance
1.9A	718µH
2.6A	604µH
3.5A	476µH
4.7A	342µH
6.2A	255µH
8.4A	176µH
11.3A	131µH
15A	98µH
19.8A	73µH
25.5A	57.1µH

Scope 5: testing an Altronics L6630 470µH 5A toroidal inductor with a powdered iron core.



Current	Inductance
3.2A	11.8µH *
5.4A	13.6µH
7.5A	13.0µH
9.6A	12.0µH
11.9A	11.5µH
14.3A	10.9µH
16.9A	9.5µH
19.6A	9.3µH
22.4A	8.6µH
25.5A	7.8µH

Scope 6: testing a Bourns 2200HT-100-V-RC inductor rated for 7.9µH @ 11.2A (with a DC resistance of 0.01Ω).

**** this low measurement is likely due to inter-winding capacitance***

The current sink is fairly conventional, with IC7 amplifying the difference between the voltage from IC2 on its positive input with the voltage sensed across the 1Ω and 0.005Ω shunt resistors.

Running a wide-bandwidth current sink with various inductances as a load is a challenge with stability. The 33nF feedback capacitor across IC7b helps to keep it stable.

The differential current and voltage sense circuits are implemented with INA281B1 instrumentation amplifiers (when a higher gain is required) and TLC072 op amp for low-gain (unity) differential voltage measurements. These differential amplifiers all have

their ground references tied to the VREF– pin used by the ADC.

The TLC072 was chosen in preference to the more common TL072 due to its lower DC offset.

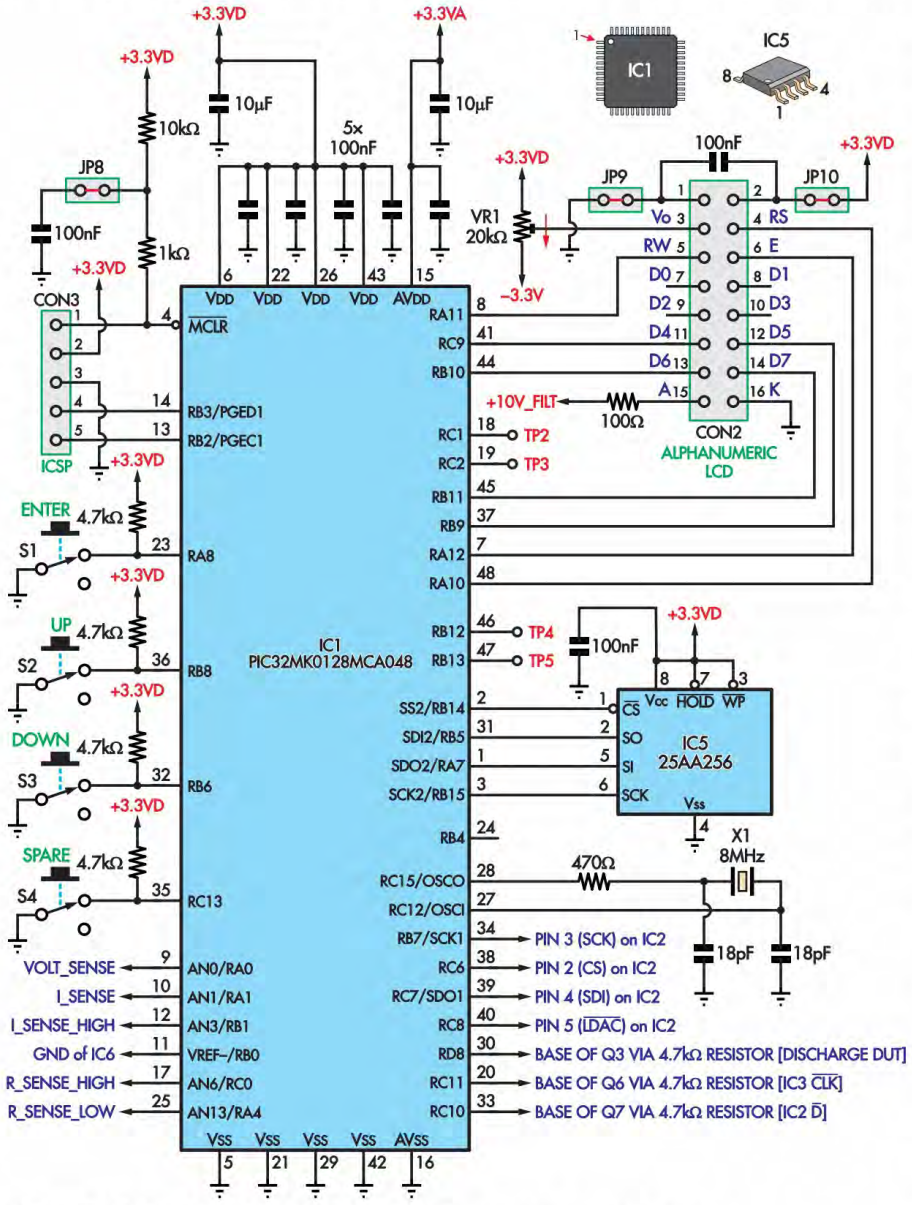
Using differential amplifiers allows us to remove common-mode signals on the +10V_FILT rail, which we know will droop during measurements, and also remove the DC offset on the DUT voltage. This allows us to make best use of the 0-3.3V input voltage range of the ADC.

CON1 provides a way to improve measurement accuracy by removing the voltage drop in the connectors and leads to the DUT from the equation. If CON1 is not used, IC6 measures the voltage across the DUT directly from CON5 & CON6 via 10Ω resistors.

If instead separate wires connect from pin 1 of CON1 to the DUT's negative lead and pin 2 to the positive lead, any voltage drop between the PCB connectors to CON5/CON6 and the DUT itself will not affect the measurements. The extra test leads will effectively short out the 10Ω resistors, providing a voltage measurement right at the DUT terminals.

Schottky diodes D6, D7, D9 & D10, in combination with 470Ω series resistors, protect the microcontroller's analog inputs from voltages outside the range of 0-3.3V that it can handle. At the same time, zener diodes ZD11 & ZD12 protect the Mosfets from voltage spikes at their gates that could damage them.

Half of IC3 (IC3a) and IC4 (IC4a) are not used, so their inputs are terminated to avoid them floating and possibly causing extra noise.



SC 2025 POWER LCR METER CONTROL CIRCUIT

Fig.8: microcontroller IC1 dominates the control circuitry. Results are displayed on an LCD screen connected to CON2, while pushbuttons S1-S4 allow the user to select modes. EEPROM IC5 stores calibration data, while crystal X1 provides accurate timing.

Microcontroller circuitry

The control circuit is shown in Fig.8. All the lower pins of microcontroller IC1 with arrows and blue labels connect to points in Fig.7. Those points in Fig.7 also have arrows and blue writing showing which microcontroller pin they connect to.

In Fig.8, besides the 48-pin microcontroller, we have an alphanumeric LCD screen connected via 16-pin header CON2, four pushbuttons (S1-S4) with associated pull-up resistors, an EEPROM (IC5) that's controlled by IC1 over an SPI serial bus, an 8MHz crystal for accurate timing, an in-circuit programming header (CON3) and numerous bypass capacitors.

We chose the PIC32MK0128MCA048 because it incorporates two 3.75Msa/s, 12-bit ADCs that we can operate synchronously. Well, it has three, but we only use two. These can sample voltages at various pins; we are using five inputs in this project (AN0, AN1, AN3, AN6 & AN13).

This microcontroller includes a direct memory access (DMA) subsystem that allows the ADC data to be moved into RAM quickly by hardware inside the microcontroller. This means the PIC32 is only interrupted after 128 samples are captured on each channel, allowing the PIC32 to capture a large buffer of samples at this full sampling rate for us to analyse.

This PIC has 32kiB of RAM and 128kiB of flash program memory. We use the majority of the RAM for a large data buffer. In inductor tests, this includes synchronised 10V rail voltages and current measurements.

This processor runs at 120MHz and includes an amazing array of peripherals for a chip that costs less than $7. The only thing that is more amazing is how complicated some of these modern processors can be to set up and write software for!

The 25AA256 EEPROM is used to store calibration data; just about any

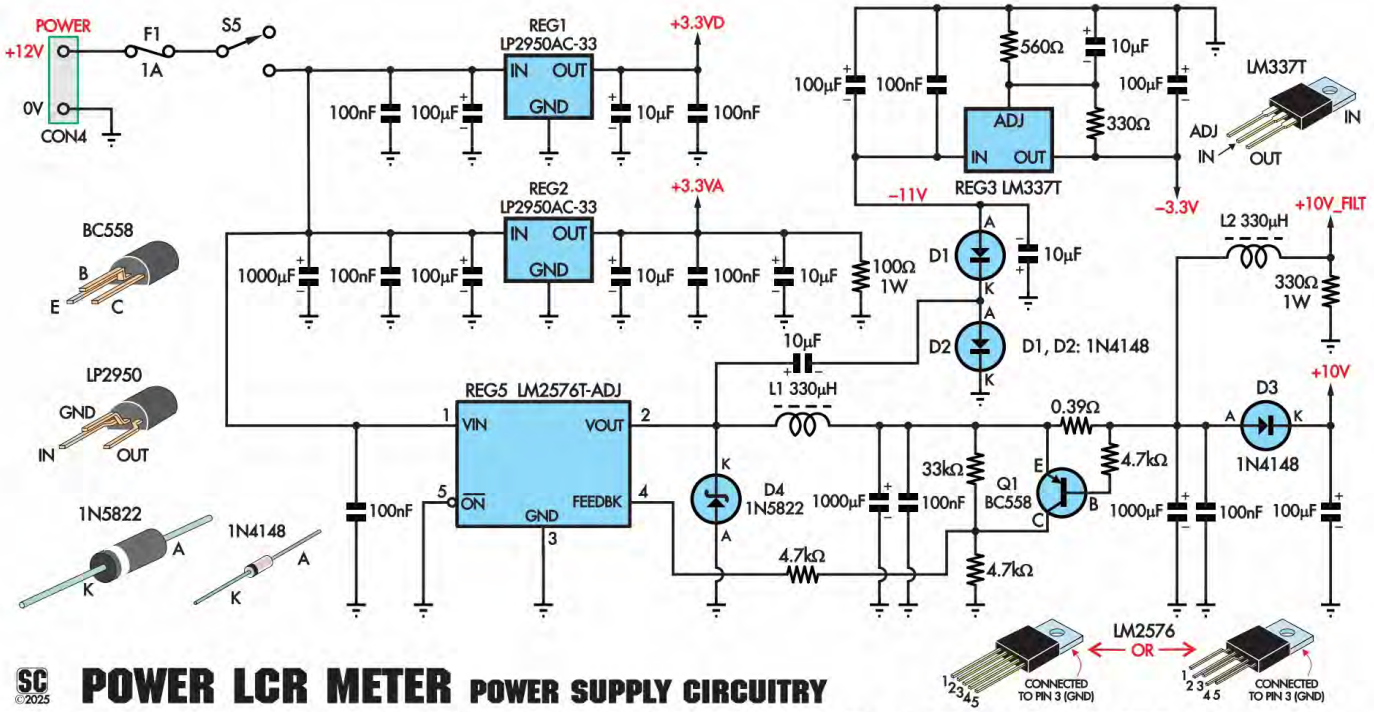


Fig.9: the power supply generates five rails from the 12-20V DC input: +3.3VD, +3.3VA, -3.3V, +10V_FILT and +10V. The +10V and -3.3V rails power devices like op amps, +3.3VD powers the micro, +3.3VA is the ADC reference voltage and +10V_FILT supplies current to the DUT.

pin-compatible SPI EEPROM will do, as we only use a handful of locations at the bottom of the address range.

The three pushbuttons (ENTER/UP/DOWN) are mounted on the rear of the PCB so they project through the case, near the standard 16×2 character LCD screen. We have stuck to a text display as these are commonly available, and aside from the power and LED connections being inconsistent between suppliers, they are interchangeable.

While a graphical display might be nice, it is not essential and would add cost and complexity. It would also complicate sourcing, as there are many similar but incompatible graphical displays available.

The support circuitry around the PIC microcontroller is mostly per application notes. There are separate 3.3V digital (+3.3VD) and analog (+3.3VA) rails. The analog rail forms the reference for the ADC, so keeping it clean is important.

Power supply

Turning now to Fig.9, the 10V rail is generated with an LM2576 switchmode buck regulator (REG5) for reasonable efficiency. Its switching output is filtered by inductor L1 and the 1000µF capacitor on its output to efficiently produce a smooth 10V rail.

The output current from the LM2576 is limited to about 1.5A by PNP transistor Q1 and the 0.39Ω resistor, which together pull the feedback pin of REG5 high if the current limit is exceeded. This is included to ensure that for extended high current pulses, the LM2576's internal over-current protection is not activated.

There are two 10V rails. The one labelled +10V drives the op amps. This is isolated by diode D3 and a 100µF filter capacitor to minimise disturbances on this rail during high-current pulses on the +10V_FILT rail.

The +10V_FILT high-current rail is filtered by a 330µH inductor which feeds two 47,000µF 16V capacitors (they are in Fig.7). This very substantial filter has been selected to ensure that even when testing large inductors, the +10V_FILT rail does not droop too much.

Diodes D1 & D2 form a charge pump with the two 10µF capacitors to generate a negative voltage from REG5's switching node, which is regulated down to -3.3V by REG3 (LM337). This allows our op amps to handle signals right down to 0V.

The +3.3VA and +3.3VD rails are generated from the 12-20V input by identically configured linear regulators REG1 & REG2. We can get away with this since there isn't a huge demand for current on those rails.

During calibration, we measure the 3.3VA rail voltage. This is important as any error in this voltage will translate to the measurement errors. The LP2950 regulators have an inherent accuracy of about ±1%. We recommended using the LP2950ACZ-3.3G but the non-AC version is available widely with only slightly reduced specs. Calibration will take care of that.

Bulk capacitance

Using two 47,000µF capacitors seems pretty generous but if we want to test a 10mH inductor at 20A, these capacitors need to deliver ½LI²J of energy, which is 2J (0.5 × 0.01H × 20A²). This energy can only come from the capacitors, as the LM2576 can only deliver an amp or so.

At 10V, our two capacitors store ½CV²J = 4.7J (0.5 × 0.047F × 2 × 10²V). If we take 2J from these capacitors and put it in the inductor, the voltage rail will drop from 10V down to 7.6V or so.

We can test much larger inductors than 10mH, but our design assumes that they will be tested at lower currents, which we think is reasonable as such a large inductor will have a fairly high DC resistance. This DC resistance will limit the maximum current.

By the way, 2J is quite a lot of energy. Do not touch the DUT during testing,

and do not disconnect it while the test is running, as the back-EMF from such a large amount of energy will generate a very high voltage if it is disconnected while current is flowing. The substantial P-channel Mosfet on the board handles this back-EMF at the end of each test.

If you disconnect the DUT from the tester during operation, this current path for back-EMF will be removed, which will result in a very high voltage spike. This could deliver a serious shock if you are touching the DUT terminals. An industrial electric fence delivers 2-5J, which is sufficient to dissuade the most stubborn animals. You really do not want to experience that!

So much software

The real functionality in this meter is all in the software. While we'll discuss it in broad strokes, you can download the complete source code from our website. So if you are inclined, you can take a deeper look, or perhaps could write your own functions or display routines.

The code is written in the relatively low-level C language, as the software needs to run very quickly during testing. C is not too hard to learn, but it does leave a lot of responsibility for the user to get things right.

I am a hardware engineer, and the code does kind of reflect this. Still, you will find a structured state machine controlling the meter, drivers for various hardware subsystems, and an effort to implement code in functions to improve its readability.

A brief look at the PIC32MK0128 data sheet is also enough to turn your hair grey. It is 562 pages of joy, and a lot of it directs you to sub-data-sheets for specific functions. This project uses multiple timers, ADC channels, direct memory access, SPI modules and general purpose I/Os. For different measurements and component values, the software changes the settings on many of these.

Microchip provides a development environment that integrates the compiler, code editor, programmer and debugger. There is also a "Code Configurator tool", which helps configure the bewildering array of modules in the microcontroller.

This is a mix of graphical representation of how parts of the microcontroller interconnect with text boxes

Parts List – Power LCR Meter

1 double-sided PCB coded 04103251, 156 × 118mm
1 Ritec or Hammond RP1285 186 × 146 × 75mm IP65 sealed ABS enclosure **[Mouser 546-RP1285]**
1 12-20V DC 1A+ power supply
1 3D-printed LCD bezel (see text next month)
2 PCB-mount M205 fuse clips (F1)
1 1A fast-blow M205 fuse (F1)
1 16 × 22mm TO-220 PCB-mounting heatsink (HS1)
2 330µH 3A toroidal inductors (L1, L2) **[Mouser 530-HCTI-330-5.2]**
3 vertical PCB-mount SPDT momentary pushbuttons plus small button caps (S1-S3) **[Mouser 612-800CSP1B6M2QE + 612-800C1BLK]**
1 vertical PCB-mount SPDT mini toggle (S5) **[Mouser 118-2MS1T1B1M2QES]**
1 8MHz low-profile crystal resonator, HC-49S (X1)
1 20kΩ top-adjust single-turn trimpot (VR1)
1 16 × 2 wide blue backlit alphanumeric LCD **[Displaytech 162KCCBC3LP]**

Cable/wire/tubing
2 200mm lengths of red/black heavy-duty hookup wire
2 200mm lengths of red/black medium-duty hookup wire
1 500mm length of light-duty figure-8 wire
1 200mm length of 16-way ribbon cable
1 100mm length of 3mm diameter black heatshrink tubing

Hardware
1 TO-220 silicone insulating washer and bush
4 M3 × 10mm tapped spacers
1 M3 × 10mm panhead machine screw
8 M3 × 6mm panhead machine screws
1 M3 hex nut
9 M3 flat washers

Connectors
3 2-pin vertical polarised headers, 2.54mm pitch, with matching plugs and pins (CON1, CON7, CON11)
1 2×8-pin vertical header, 2.54mm pitch (CON2)
1 5-pin vertical header, 2.54mm pitch (CON3)
1 2-way mini terminal block, 5/5.08mm pitch (CON4)
2 6.3mm PCB-mount vertical spade lugs (CON5, CON6)
3 2-pin vertical headers, 2.54mm pitch (JP8-JP10)
3 jumper shunts
1 16-way IDC inline socket
2 panel-mount binding posts, red & black
2 panel-mount banana sockets, red & black
2 chassis-mount BNC sockets
1 panel-mount DC socket (to suit power supply)

Integrated circuits
1 PIC32MK0128MCA048 32-bit microcontroller programmed with 0410325A.HEX, TQFP-48 (IC1) **[Silicon Chip SC7454]**
1 MCP4822-E/P 12-bit SPI DAC, DIP-8 (IC2)
1 4013B dual D-type flip-flop CMOS IC, DIP-14 (IC3)
1 LM393 dual single-supply comparator, DIP-8 (IC4)
1 25AA256-I/SN 32kiB EEPROM, SOIC-8 (IC5) **[Mouser 579-25AA256-I/SN]**
2 INA281B1 20V/V 1.3MHz current sense amplifiers, SOT-23-5 (IC6, IC8)
1 TLC072AIP dual low-noise JFET-input op amp, DIP-8 (IC7)
2 LP2950-3.3 3.3V 100mA low-dropout linear regulators, TO-92 (REG1, REG2)
1 LM337 1A adjustable negative linear regulator, TO-220 (REG3)
1 LM2576 3A integrated buck switch-mode DC/DC converter, TO-220-5 (REG5)

Other semiconductors
2 BC558 30V 100mA PNP transistors, TO-92 (Q1, Q10)
1 SUP70101EL 100V 120A P-channel Mosfet, TO-220 (Q2) **[Mouser 78-SUP70101EL-GE3]**
5 BC548 30V 100mA NPN transistors, TO-92 (Q3, Q6-Q9)

1 IPP013N04NF2SAKMA1 40V 197A N-channel Mosfets, TO-220 (Q4)
[Mouser 726-IPP013N04NF2SAKM]
1 TIP121 NPN Darlington transistor with integral diode, TO-220 (Q5)
2 12V 400mW/1W zener diodes, DO-35/DO-41 (ZD11, ZD12)
4 1N4148 75V 200mA signal diodes, DO-35 (D1-D3, D8)
1 1N5822 40V 3A schottky diode, DO-201AD (D4)
5 BAT85 30V 200mA schottky diodes, DO-34 (D5-D7, D9, D10)

Through-hole capacitors
2 47,000µF 16V snap-in electrolytic **[Mouser 598-81LX473M016A452]**
3 1000µF 25V low-ESR radial electrolytic
5 100µF 25V low-ESR radial electrolytic
6 10µF 50V low-ESR radial electrolytic
16 100nF 63/100V MKT
1 33nF 63/100V MKT
1 220pF 50V ceramic

SMD capacitors *(all X7R ceramic SMD M2012/0805 unless noted)*
2 10µF 10V
3 1µF 50V
8 100nF 50V SMD M2012/0805
2 18pF 50V NP0/C0G ceramic SMD M2012/0805

TH resistors *(all ¼W 1% axial unless noted)*

2 47kΩ	23 4.7kΩ	1 560Ω	1 330Ω	1 1Ω 1W 5%
1 33kΩ	1 2.7kΩ	4 470Ω	4 100Ω	1 0.39Ω 5W 5%
1 10kΩ	2 1kΩ	1 330Ω 1W 5%	1 100Ω 1W 5%	

1 0.005Ω open-air **[Welwyn OAR1-R005FI or similar]**

SMD resistors *(all ⅛W 1% SMD M2012/0805 unless noted)*

1 10kΩ	1 1kΩ	1 470Ω	2 10Ω



The PCB mounts on the rear of the lid. Except for the power input, the terminals on the side are optional.

that allow you to, for example, program timer periods. This code configurator generates device driver code for each subsystem that you can use in your program. This forms the device driver layer for the microcontroller, making our life much easier.

For the external parts like the SPI EEPROM, DAC and 16×2 LCD, we don't have the luxury of a code configurator, but we can use libraries that we or other people have developed in the past. For control of things like the Mosfets, we wrote simple 'drivers' to allow the main program to perform common functions without the complication of directly interfacing to hardware.

A good example is the "CORETIMER_DelayMs(X)" function call, which lets our program ask for a delay of X milliseconds. The driver looks after things like clock frequencies and suchlike.

Similarly, HDByteWriteSPI() is a function that allows us to write a byte to the SPI EEPROM at a defined location. We don't care what specific EEPROM it is, as long as we can save data to it and read data from it.

By using these high-level drivers in our program, the code is much easier to read and also, once we have tested them, we can treat them as 'black boxes'. Our top-level algorithm can command the hardware to perform functions such as "Pulse_Start()" without needing to bother with any of the hardware details.

You will have noticed there are actually four fairly independent measurement modes. We won't go into detail here, but from a high level, the main program file "Inductor_Tester.c" performs the main functions:

1. Configures the hardware
2. Loads and saves calibration data
3. Allow the operator to:
   a. Calibrate the meter
   b. Measure resistance, inductance, capacitance and inductor saturation
   c. Display results.

These functions all live in a simple state machine that allows the user to select the measurement desired and review the measurements.

Next month

This article has already become quite long, and we haven't gotten up to the construction, calibration or testing steps yet. We'll have all that next month, along with some hints on using the device effectively. PE

Teach-In 2026

by Mike Tooley



World of Wireless – An Introduction to Radio and Wireless Technology

Series 12, part 4: software-defined radio

In the last instalment of this series, we introduced the fundamentals of radio communication systems, provided an overview of Morse code and CW (continuous wave) as a simple method of communication, and discussed the importance of modulation and demodulation.

This month, our focus shifts to software-defined radio (SDR), an innovative technology that enables radio signal processing tasks traditionally handled by hardware to be managed via software. We examine different SDR solutions and show how a cost-effective SDR paired with powerful software can allow you to receive a wide range of radio signals at frequencies extending from HF to UHF.

This month's *Hands-On project* uses the VFO (variable-frequency oscillator) and 10MHz crystal-controlled reference oscillator modules from last month as the foundation for an amplitude-modulated RF (radio-frequency) signal source.

Digital frequency synthesis

Before introducing our main topic, it's important to outline how digital systems can be used to control analog oscillators, with particular emphasis on the application of phase-locked loops (PLL).

PLL technology was first incorporated into military communications equipment during the mid-1960s, addressing the need to generate a wide range of accurate and stable frequencies within multi-channel frequency synthesisers.

In those early applications, cost considerations were secondary, permitting the use of advanced circuit designs that utilised large numbers of discrete components and integrated circuits (ICs).

Phase-locked loop techniques did not arrive in mass-produced equipment until the early 1970s. Compared with today's equipment, such arrangements were crude, employing as many as nine or ten ICs. Complex as they were, those PLL circuits were more cost-effective than their comparable multi-crystal mixing synthesiser counterparts.

With the advent of large-scale integration (LSI) in the late 1970s, the frequency-generating unit in most radio equipment could be reduced to one, or perhaps two, LSI devices together with a handful of additional discrete components. The cost-effectiveness of this approach is now beyond question, and it is unlikely that, at least in the most basic equipment, much further refinement will be made.

In the area of more complex receivers and transceivers, we are now witnessing a further revolution in the design of synthesised radio equipment, with the introduction of dedicated processors that permit keypad-programmed channel selection and scanning with pause, search, and lock-out facilities.

The most basic form of PLL consists of a phase detector, filter, DC amplifier and voltage-controlled oscillator (VCO), as shown in Fig.4.1. The VCO is designed so that its free-running frequency is at or near the reference frequency. The phase detector senses any error between the VCO and reference frequencies.

The output of the phase detector is fed, via a suitable filter and amplifier, to the DC control voltage input of the VCO. If there is any discrepancy between the VCO output and the reference frequency, an error voltage is produced, which is used to correct the VCO frequency. The VCO thus remains locked to the reference frequency. If the reference frequency changes, so does the VCO's.

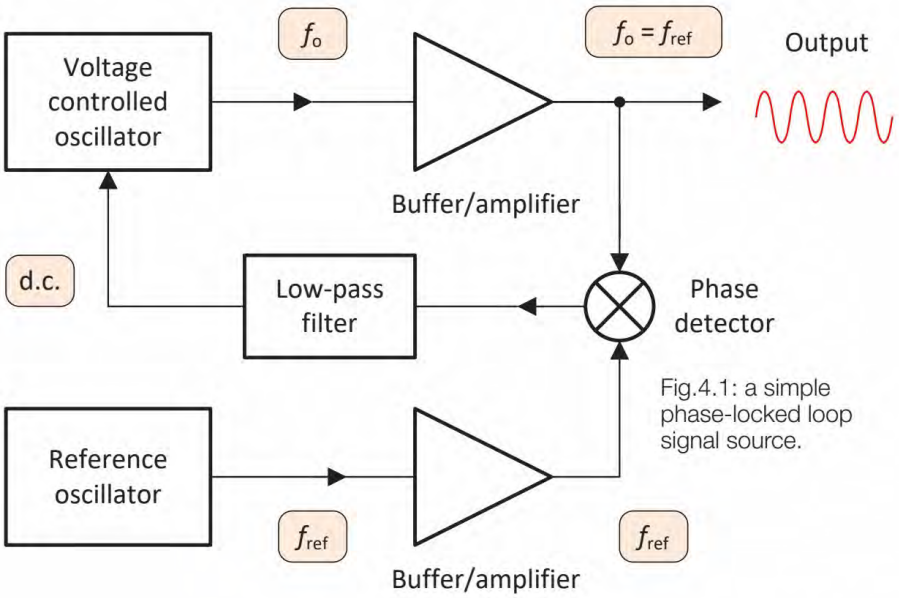


Fig.4.1: a simple phase-locked loop signal source.

www.xsava.xyz

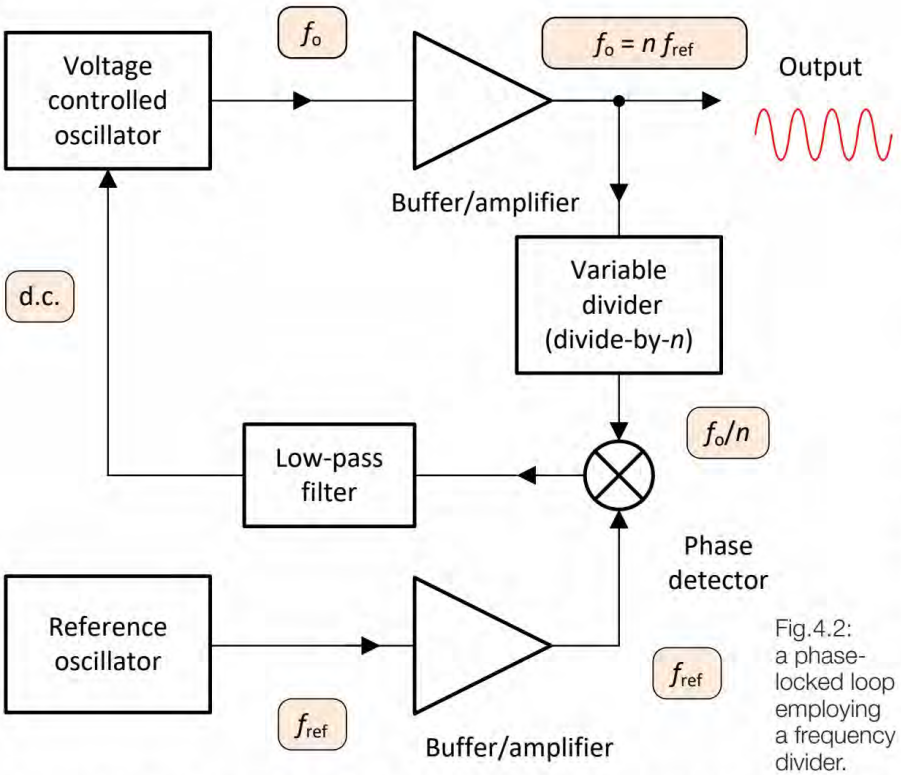


Fig.4.2: a phase-locked loop employing a frequency divider.

The bandwidth of the system is determined by the time constants of the loop filter. In practice, if the VCO and reference frequencies are very far apart, the PLL may be unable to lock.

The frequency range over which the circuit can achieve lock is known as the capture range. A PLL takes a finite time to achieve a locked condition, so that the VCO locks to the mean value of the reference frequency.

The basic form of PLL, shown in Fig.4.1, is limited in that the reference frequency is the same as that of the VCO and no provision is incorporated for changing it, other than by varying the frequency of the reference oscillator itself.

In practice, it is normal for the phase detector to operate at a much lower frequency than that of the VCO output; hence, a frequency divider is incorporated in the VCO feedback path, as shown in Fig.4.2. The frequency presented to the phase detector will thus be $f_o \div n$, where n is the divisor.

When the loop is locked (ie, no phase error exists), we can infer that $f_{ref} = f_o \div n$ or $f_o = n \cdot f_{ref}$.

A similar divider arrangement can also be used at the reference input to the phase detector, as shown in Fig.4.3. The frequency appearing at the reference input to the phase detector will be $f_{ref} \div m$ and the loop will be locked when $f_{ref} \div m = f_o \div n$ or $f_o = (n \div m) f_{ref}$.

Thus, if f_{ref} were 100kHz, n were 2000 and m were 10, the output frequency, f_o, would be (2000 ÷ 10) × 100kHz = 20MHz

If the value of n is made variable by replacing the fixed divider with a programmable divider, different output frequencies can be generated. If, for example, n were variable from 2000 to 2100 in steps of one, f_o would range from 20MHz to 21MHz in 10kHz steps. Fig.4.3 shows the basic arrangement of a PLL incorporating a programmable divider driven from the equipment's digital frequency controller.

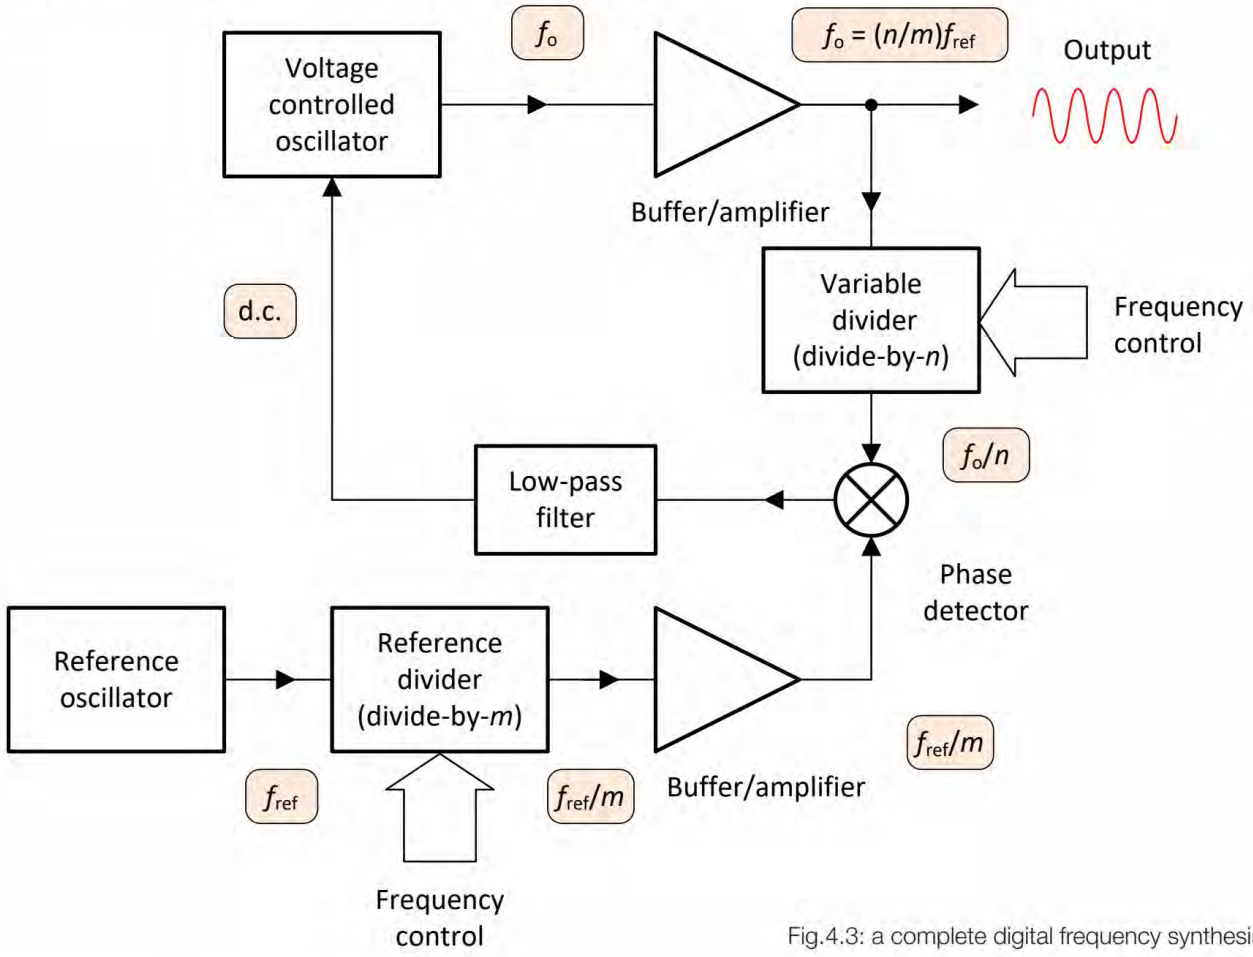


Fig.4.3: a complete digital frequency synthesiser.

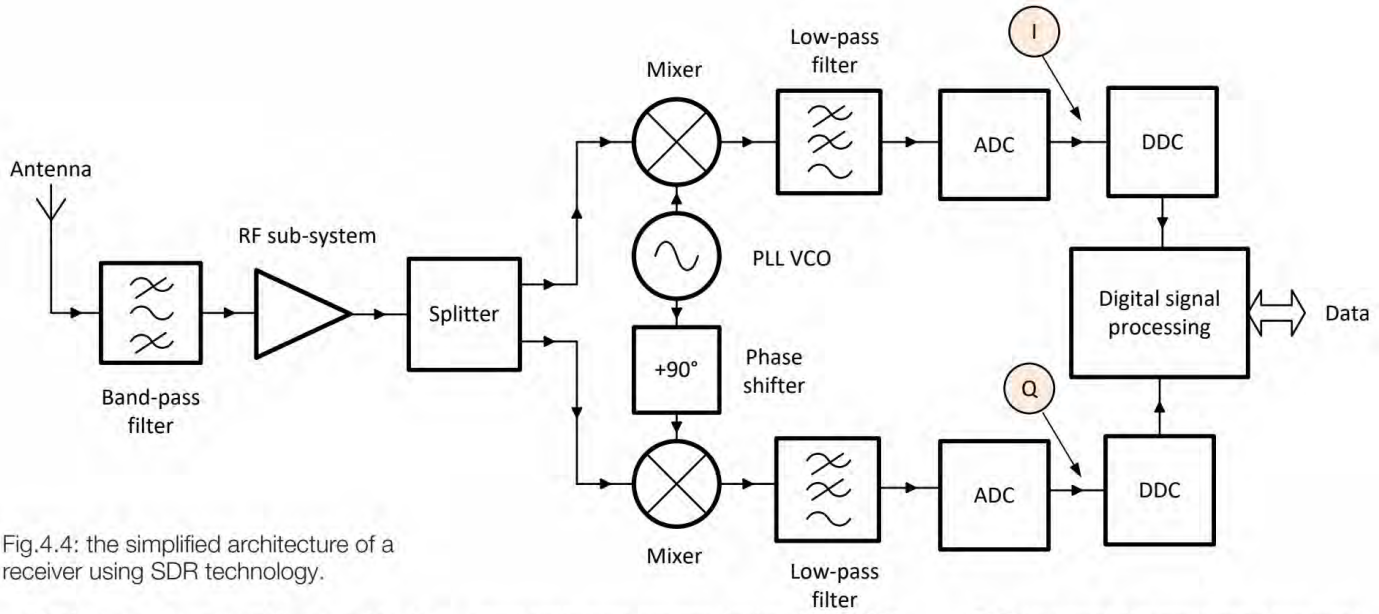


Fig.4.4: the simplified architecture of a receiver using SDR technology.

Problems can sometimes arise in high frequency synthesisers where the programmable frequency divider, or divide-by-n counter, has a restricted upper frequency limit. In such cases, it is necessary to mix the high-frequency VCO output with a stable locally generated signal derived from a crystal oscillator.

The mixer output (a relatively low difference frequency) will then be within the range of the programmable divider.

Software-defined radio

In the radio architecture that we met last month, the signal paths were implemented using a traditional approach based on application-specific hardware components such as resistors, capacitors, inductors and semiconductor devices.

In modern radio equipment, there is an increasing use of software both for controlling the hardware and for signal processing. This has led to the advent of two important technologies: software-controlled radio (SCR) and software-defined radio (SDR).

SCRs are now extensively used in current equipment, enabling the operating parameters of large and medium-scale ICs to be configured using digital techniques based on microprocessors, where software and data are stored in solid-state memory devices, such as flash memory.

In software-controlled radio, the signal path is implemented using hardware, with some functionality controlled by software. Parameters usually controlled by software include frequency selection, tuning, mode selection, gain control and transmission power.

Software-defined radio (SDR) takes this one step further, with most of the signal processing (including filtering, modulation/demodulation and encoding/decoding) being performed by software (ie, instruction code running on a general-purpose processor or digital signal processor [DSP]) rather than hardware.

SDR is a rapidly emerging technology, showing considerable promise with state-of-the art implementation in commercial and military radio equipment. An SDR may still use conventional radio architecture in the front-end RF and mixer stages.

With SDR technology, the signal path can be easily reconfigured without the need for costly changes to hardware. Most of the digital signal processing within an SDR is conventionally implemented using one or more field-programmable gate arrays (FPGA), DSPs or an equivalent embedded processing device.

The signal path in SDR equipment can be quickly and easily reconfigured by making changes to the software. This allows modification and upgrading without the need to change any hardware. It also permits the rapid cloning of operational parameters such as frequency, channel spacing and selectivity.

Figs.4.4 & 4.5 show the simplified arrangements of receivers and trans-

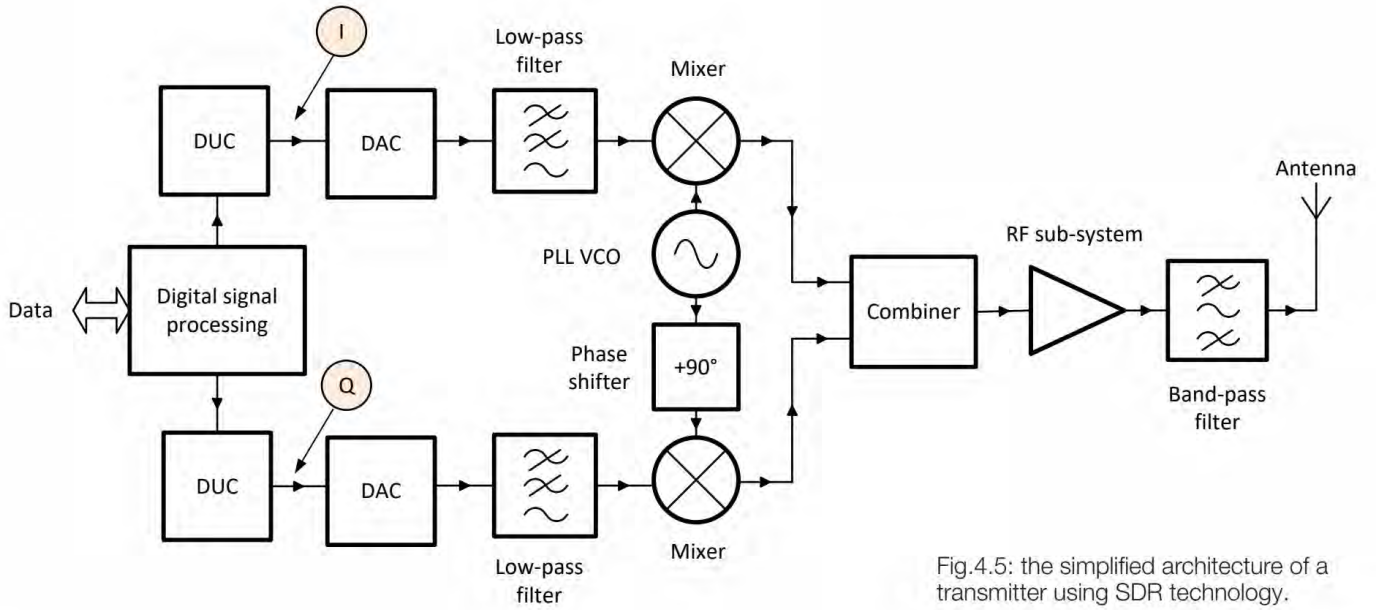


Fig.4.5: the simplified architecture of a transmitter using SDR technology.

mitters based on SDR technology, respectively. Note how analog-to-digital conversion (ADC) is used in the receiver, while digital-to-analog conversion (DAC) is employed in the transmitter.

Analog circuitry is still present in both the receiver and transmitter, the former having it in the low-level RF amplifier and mixer stages, while the latter uses conventional analog circuitry in its high-level driver and RF power amplifier stages.

In the SDR receiver arrangement shown in Fig.4.4, the RF subsystem (typically comprising band-pass filters and RF amplifiers) supplies an analog signal to the splitter with identical in-phase signals applied to the two mixer stages.

The local oscillator input to the two mixers is derived from a PLL VCO arrangement. The local oscillator input to one of the mixers is phase shifted by 90° so that the two local oscillator signals are in phase quadrature (two signals with a 90° phase shift between them).

The two mixer outputs go through low-pass filters before being applied to two ADCs, the outputs of which constitute in-phase (I) and quadrature (Q) components.

The I and Q signals then pass into a digital down converter (DDC) to reduce the sampling rate of the signal before being passed to the digital signal processor (DSP) where the modulation (AM, FM or PSK data) is recovered from corresponding pairs of down-sampled quadrature data.

The digital back-end of an SDR usually comprises an FPGA or embedded processor with onboard DSP functionality for modulation, demodulation, up-converting, down-converting, coding, decoding and protocol handling. This can also be implemented using a powerful enough general-purpose computer CPU. All this complex processing is achieved using easily reconfigurable software, rather than extensive hardware that uses conventional components and circuitry.

In the corresponding SDR transmitter arrangement shown in Fig.4.5, the digital data is processed before being applied to digital up-converters (DUC), from which the I and Q signals are derived. These are then fed to two DACs, each followed by low-pass filters.

The mixing process (like that used in the SDR receiver shown in Fig.4.4) produces two quadrature signals that are combined before application to the RF subsystem, which typically comprises a pre-driver, driver and power amplifier. After bandpass filtering, the final output is applied to the antenna.



Fig.4.6: inexpensive USB dongles are a great way to start experimenting with software-defined radio.

Getting started with SDR

There are several ways to get started with SDR. Fully integrated SDR receivers have hardware, software and displays integrated in the same package. They offer exceptional performance but can be expensive.

PC-based SDR adaptors are also available from several suppliers including Airspy, HackRF, and RS-Play. This alternative can be great if you know that you will be using SDR regularly for listening and experimentation. These mid-range SDR receivers require a PC or laptop and appropriate software.

Inexpensive DVB-T (digital TV) receiver dongles can also be used. This low-cost option will get you started with SDR, but will still allow you to experiment with a wide variety of fully featured software such as SDR#, HDSDR, and SDR Console. As it requires minimal outlay, this is the approach that we've adopted for this Teach-In series.

Fig.4.6 shows just a small selection of low-cost SDR receivers currently available.

Based on low-cost chips originally designed for use in DVB-T set-top boxes, the receivers shown in Fig.4.6 are widely available from online sellers in versions that support reception over the entire VHF and UHF range. In recent years, several manufacturers have added features that not only improve performance but also extend coverage to MF and HF.

Companion software can be freely downloaded from the web for running under most popular operating systems, including Windows, Linux and macOS. This means you can leverage this technology as an effective means of studying RF principles and exploring the electromagnetic spectrum.

The internal arrangement of a basic RTL-SDR is shown in Fig.4.7. RF input is via an MCX coaxial connector, and initial analog signal processing is handled by an R820T from Rafael Micro. Designed to function as a low-power digital TV tuner, this device comprises a low-noise amplifier (LNA), mixer, PLL, VCO, crystal-controlled reference oscillator and intermediate frequency (IF) filter.

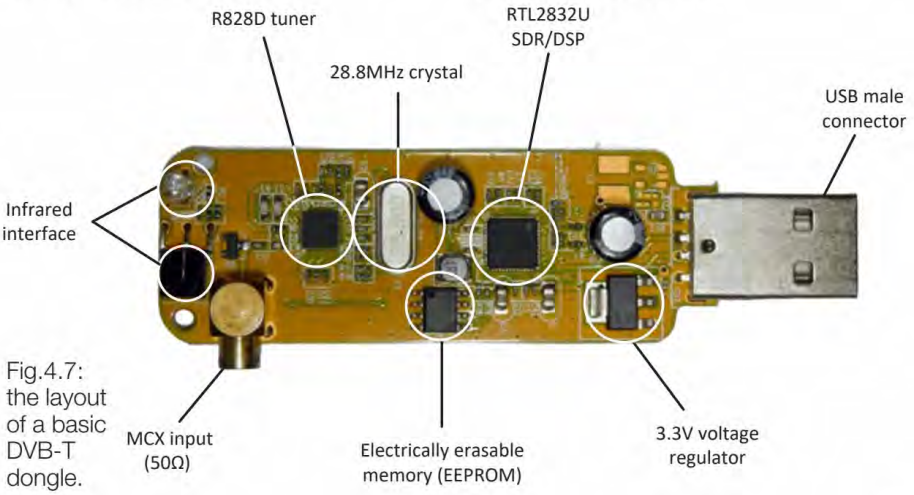


Fig.4.7: the layout of a basic DVB-T dongle.

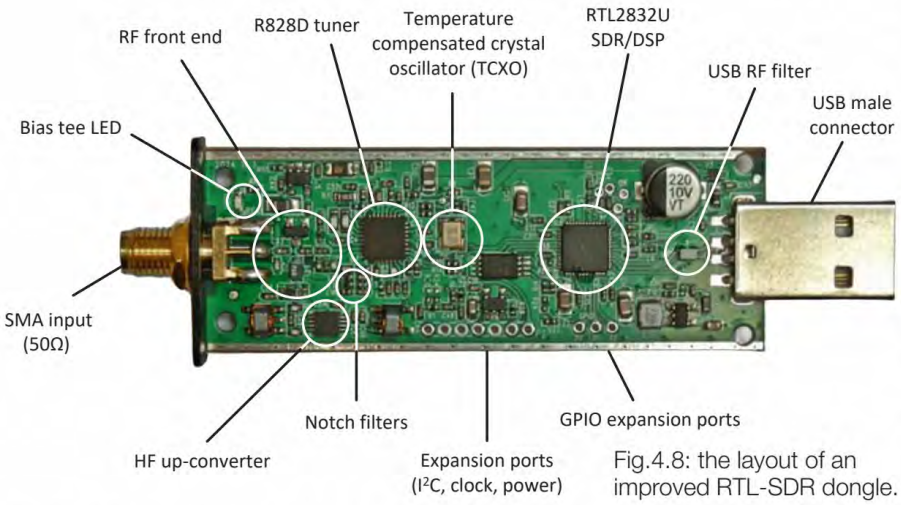


Fig.4.8: the layout of an improved RTL-SDR dongle.

The chip operates from a 3.3V supply, has a quoted frequency range of 42-1002MHz and a noise figure of 3.5dB. The R820T comes in a 24-pin QFN (quad flatpack, no leads) package.

Following IF conversion, the next stage is an RTL2832U IC that provides full SDR functionality with a USB interface. The chip incorporates a sampling clock, IF to baseband conversion and low-pass filters in the I and Q signal paths. Multiple IF input frequencies are supported, as well as a zero-IF input. The chip is supplied in a 48-pin QFN package.

An improved RTL-SDR is shown in Fig.4.8. This upgraded device is fitted with a more robust SMA input connector as well as an improved RF front end. The tuner chip has been replaced by an R828D and a temperature-controlled crystal oscillator (TCXO) fitted to provide a frequency accuracy of typically better than ±1ppm.

Additional filtering is applied to reduce noise from the USB interface. An HF converter has been added to extend frequency coverage to the spectrum below 30MHz.

Improving RTL-SDR performance

RTL-SDR performance can be improved in several ways. Some manufacturers have added extra RF filtering (see Fig.4.9), while others have focused on noise reduction and RF input protection. Yet others have incorporated improved heat dissipation for the tuner and SDR chips, which both tend to run hot when mounted in small plastic enclosures.

Even the performance of the most inexpensive dongles can be improved by mounting them in a small diecast enclosure and fitting external SMA or BNC adaptors, as shown in Fig.4.10.

It's also possible to connect an HF input directly to the RTL2823U chip via a suitable RF filter (Fig.4.10), allowing direct sampling that bypasses the tuner stage completely. Details of these and various other useful modifications can be found on the web (or see the article starting on page 32 of our November 2018 issue).

Up-converters

An alternative approach to extending the coverage of a basic RTL-SDR dongle is using mixing techniques to convert the desired frequency range into a range that lies within that covered by the RTL-SDR device.

A typical device such as the popular Ham-it-Up HF Upconverter uses a 125MHz local oscillator (LO) to mix an RF input between 100kHz and 50MHz, and produce intermediate frequency (IF) signals between 125.1MHz and 175MHz, thus adding effective MF and HF coverage.

Our SiDRADIO project, published in four parts starting with the October 2014 issue, used a similar approach. It combined a DVB-T dongle with a tuned front-end, RF amplifier and up-converter incorporating a 125MHz TCXO.

For example, if a signal is to be received on 1.6MHz, the SDR software will be tuned to 125MHz + 1.6MHz = 126.6MHz. The software can add a preset offset to the display frequency, so that although it is

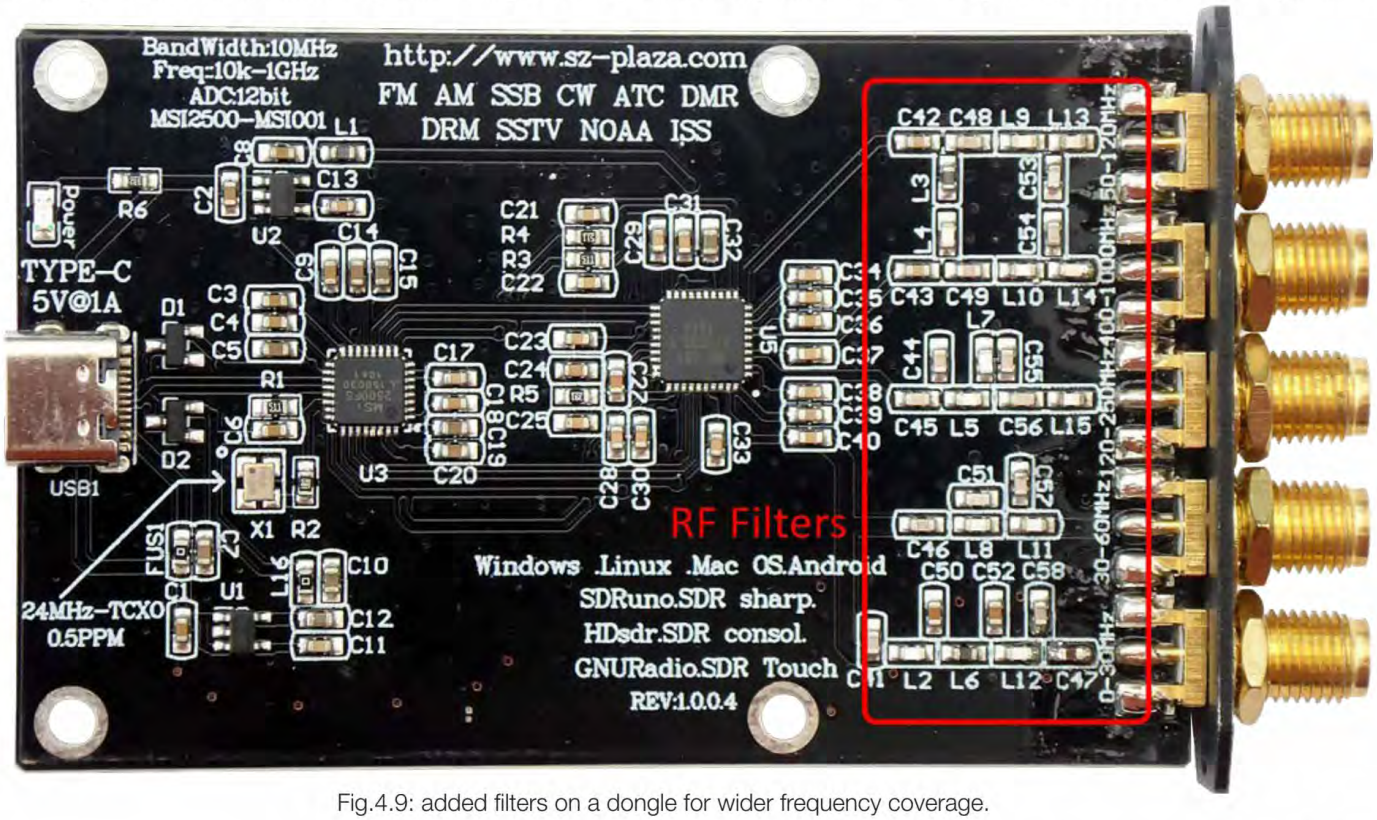


Fig.4.9: added filters on a dongle for wider frequency coverage.

receiving the signal at 126.6MHz, the display will read 1.6MHz.

SDR software

An extensive selection of SDR software is available. Depending on your operating system, you could consider several popular SDR software packages including SDR Console, SDR#, SDR++, HDSDR, Cubic SDR and GQRX. Here's a summary of their main features.

SDR Console is a comprehensive and feature-rich SDR program for Windows. SDR Console supports a wide range of SDR hardware. It offers advanced features such as multi-receiver support, recording and playback, remote operation, and a highly customisable user interface.

SDR Console is well-suited for beginners as well as more advanced users requiring powerful signal analysis and management tools. SDR Console has an attractive and reasonably intuitive interface, and is the package that we've used extensively to produce this article.

SDR# (also called SDRSharp) is an extremely popular SDR application that offers a user-friendly interface as well as extensive support for optional plugins. SDR# supports a wide range of SDR devices and works well when paired with a basic RTL-SDR dongle. Its functionality can be easily extended with the currently available library of plugins.

SDR++ is a cross-platform open-source SDR that aims to be "bloat free and simple to use". The package offers wide hardware support and has a straightforward user interface. It's available for Windows, Linux, macOS and BSD.

HDSDR: this excellent Windows-based SDR supports a wide variety of SDR hardware. It features a clean and attractive interface, powerful signal processing capabilities and multiple demodulation modes. The package provides excellent tuning control and incorporates a useful recording capability.

CubicSDR is an open-source, cross-platform SDR program available on Windows, macOS and Linux. CubicSDR supports a wide range of SDR devices and provides intuitive spectrum and waterfall displays.

GQRX is popular with Linux and macOS users (a Windows version is also available). Built on the GNU Radio framework, GQRX offers a lightweight and easy-to-use interface with frequency scanning and a range of audio output options.

Setting SDR Console

Setting up an SDR receiver requires several stages and can sometimes be confusing, so we've provided a few notes to help you with this task. Based on SDR Console and our recommended device, the popular and well-supported RTL-SDR V4, here's a five-step overview of the process (useful guides and walkthroughs are available at www.sdr-radio.com).

1. Connect the SDR dongle and install the drivers

Insert your RTL-SDR into an available USB port. Since it's unlikely that you will already have the required driver installed, your next task will be to install the correct driver for the SDR. We recommend using the Zadig utility on Windows systems, which can be obtained from https://zadig.akeo.ie/. Simply download it, open the file and then click on "Yes" to install it.

To replace the default driver, go to the "Options" menu in Zadig and select "List All Devices". This will show all USB devices connected to your computer, including your SDR dongle.

In the dropdown list, look for "Bulk-In, Interface (Interface 0)" (sometimes it may appear as "RTL2832U" or "RTL2832UHIDIR").

Installing the driver to the wrong device may cause problems, so ensure that you select "Interface 0" (not "Interface 1" or any other device). With the correct device selected, ensure that "WinUSB" is chosen as the target driver in the box next to the arrow. Click "Replace Driver" so that Zadig uninstalls the default driver and replaces it with WinUSB for your newly added device.

This process may take a little time. If Zadig doesn't manage to find your device, unplug and then reinsert it. Alternatively, simply restart your computer and run Zadig again.

2. Download, install and launch SDR Console

Having successfully installed the replacement driver, visit the SDR

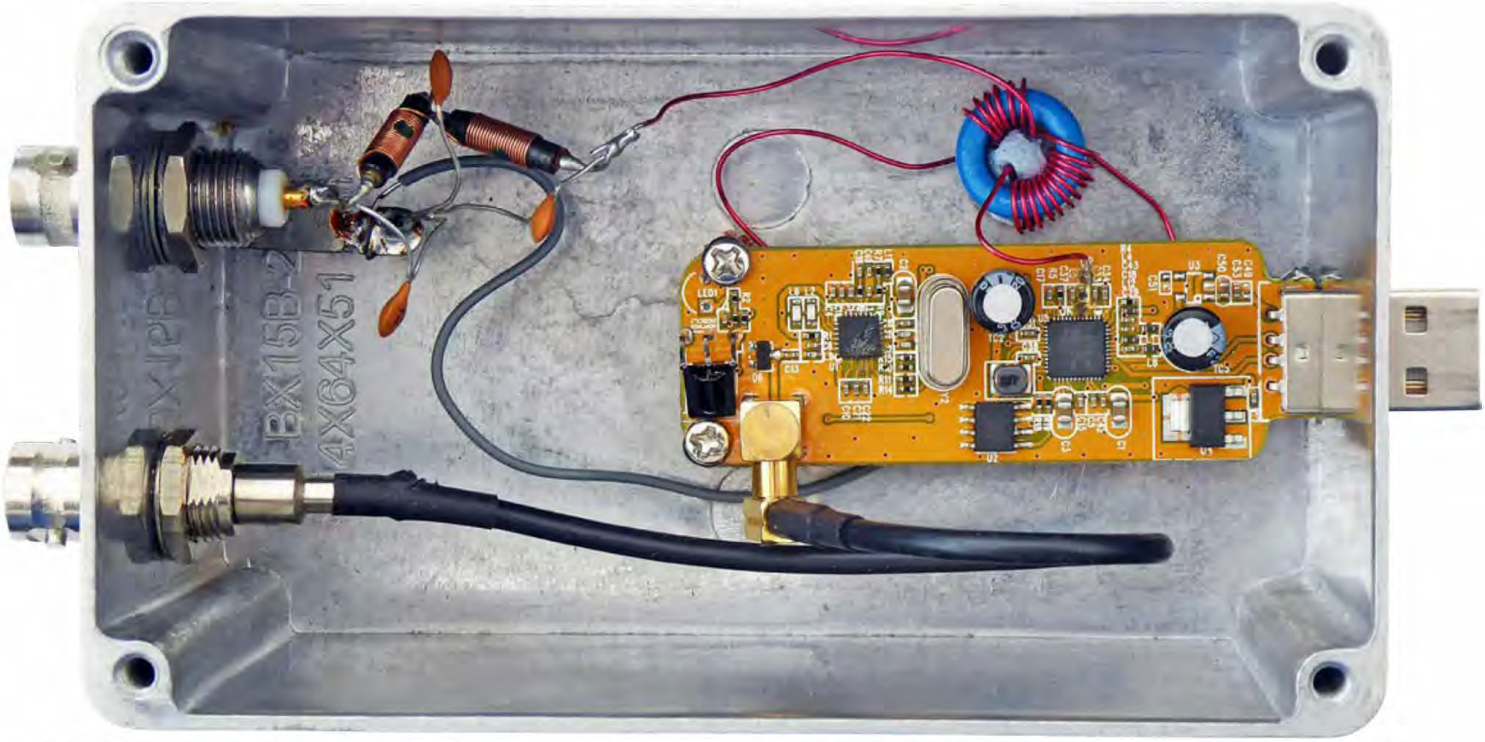
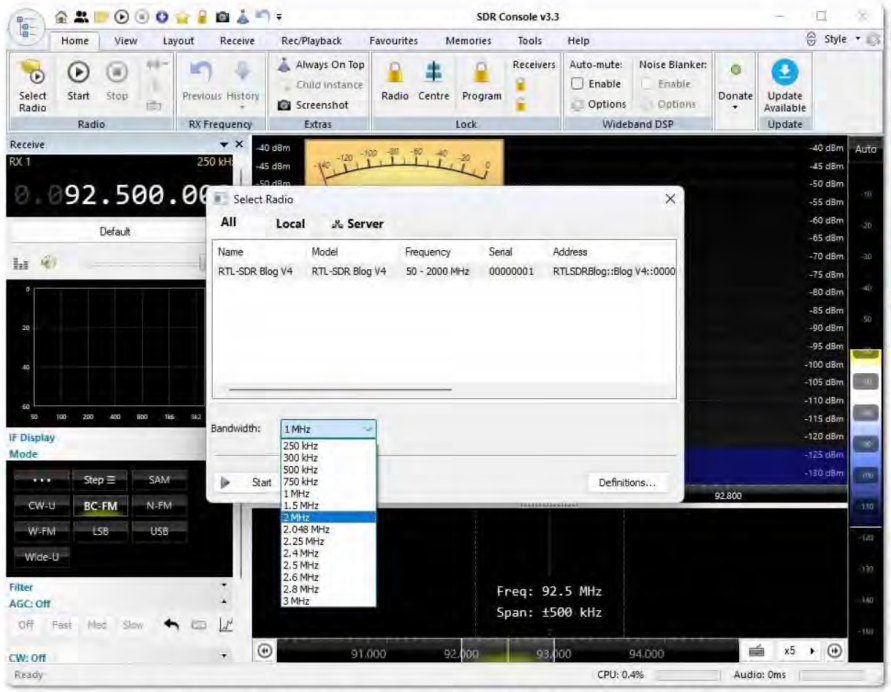


Fig.4.10: improvements made to a low-cost dongle: a screened enclosure, RF connectors, low-pass filter and direct sampling mod.



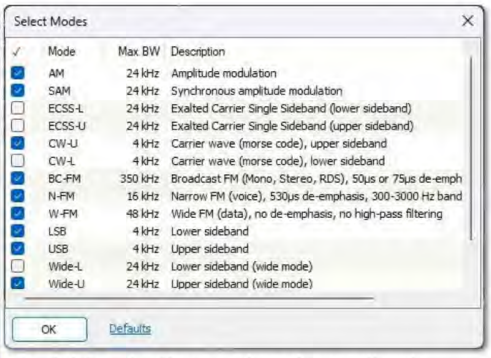


Fig.4.12: selecting modes of reception in SDR Console. These are added to the radio Definition and thereafter can be selected from the Mode dropdown menu.

Fig.4.11: setting the bandwidth before starting the RTL-SDR in SDR Console.

Console website at www.sdr-radio.com/download and download the latest version of the SDR-Radio software (we chose "Beta 3.4 Build 3818"). **Warning: there are many scam download ads on this site, so make sure you don't click on any of them.**

There are currently two options to access the installation package, either "Microsoft" or the "OneDrive" hosting service. We chose the former. Download the SDR-Radio package, open the downloaded file and run the installer.

Several additional C++ files may be required during the installation process. Follow the prompts to complete the installation and then open SDR Console from your Start menu or desktop shortcut.

3. Add your RTL-SDR device for use with SDR Console

On first running SDR Console, you will be prompted to add a radio. Click "Definitions" and then "Search". SDR Console will then scan for any connected SDR hardware. Select the RTL-SDR from the search list and click "Save". If an RTL-SDR dongle can't be found, ensure the drivers have been installed correctly (repeat Step 1).

4. Configure basic settings and connect the antenna

Click on "Select Radio" and highlight the saved definition for your SDR dongle. This will usually appear as "RTL-SDR Blog V4", or something similar. Next, select a bandwidth for your device (see Fig.4.11) and then click "Start" to activate the radio and return to the main window and start the SDR receiver.

Next, attach a suitable antenna. We don't recommend connecting the antenna that's usually supplied with the low-cost SDR dongles. Instead, some initial testing can be carried out using a small dipole antenna (more on this next month) or just a 1.5m length of insulated wire connected to the centre contact of the dongle's SMA connector.

If you wish to add reception modes other than those that are selected by default, these can be easily added by making appropriate changes to the radio's "Definition" settings (see Fig.4.12).

5. Check the SDR with broadcast FM radio signals

Set the receive mode to broadcast FM ("BC FM"), adjust the RF gain to around 30dB and search for local signals at 88-108MHz. This will allow you to verify the setup. Note that for optimum reception, you will need to move the antenna well away from any nearby sources of interference. You are now ready to enjoy your new SDR radio!

Hands-On: An AM signal source

This month's Hands-On project is an AM signal source that extends the functionality of the two modules described last month. The AM signal source comprises three modules:

1. The VFO (or the 10MHz crystal-controlled reference oscillator) described last month. This module generates the RF carrier as an input to the modulator module.
2. An AF oscillator that generates a constant 900Hz audio tone for the modulating signal.
3. An AM modulator that produces an amplitude-modulated (AM) output using the inputs from the other two modules.

We will now describe the construction of the two new modules before bringing them together in the complete AM signal source.

AF oscillator module

The circuit of the AF oscillator module is shown in Fig.4.18. The

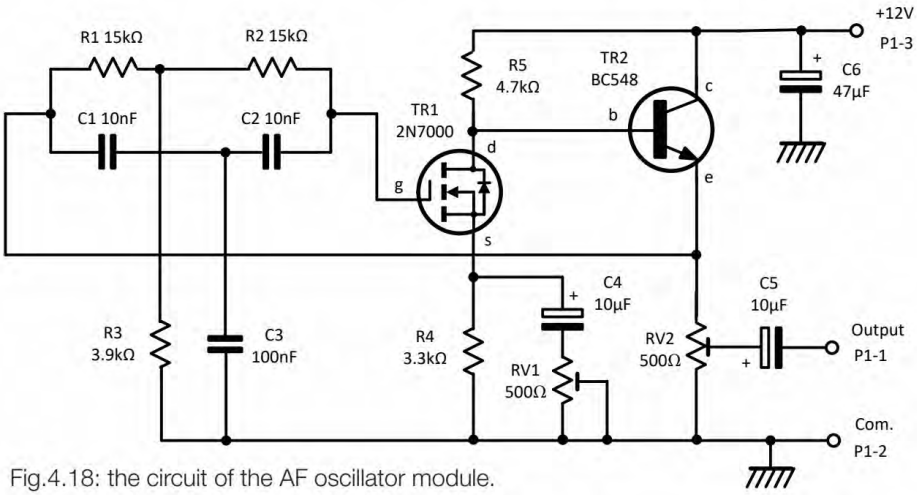


Fig.4.18: the circuit of the AF oscillator module.

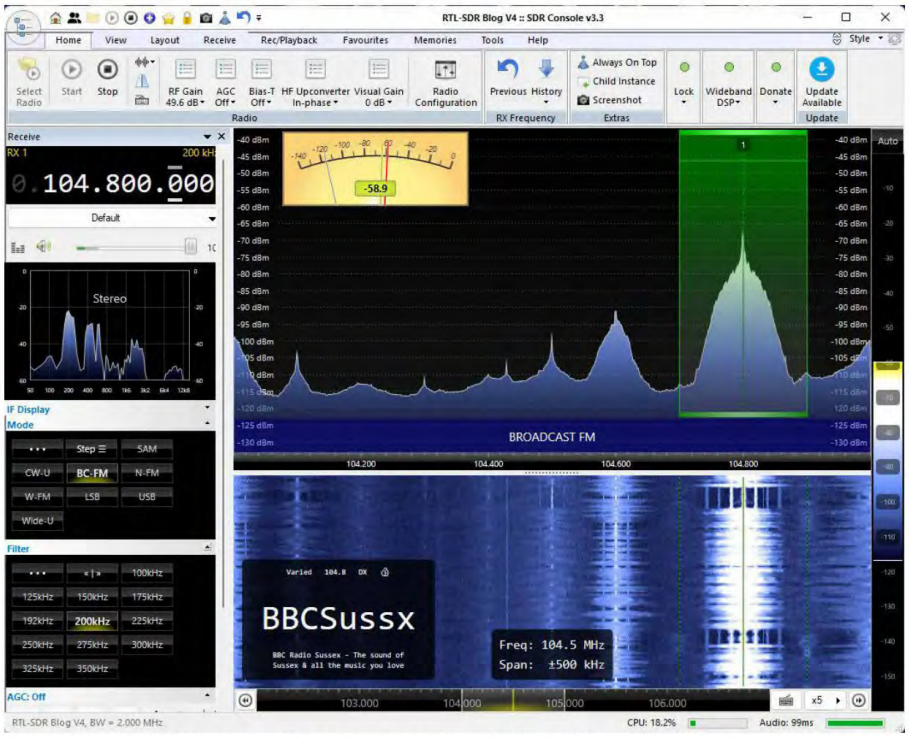


Fig.4.13: using SDR Console to receive a local FM broadcast station on 104.8MHz.

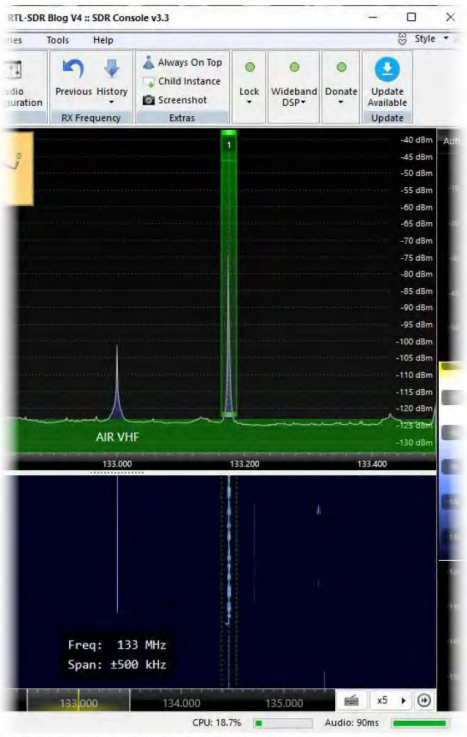


Fig.4.14: using SDR Console to receive Air Traffic Control (AM) on 133.175MHz.

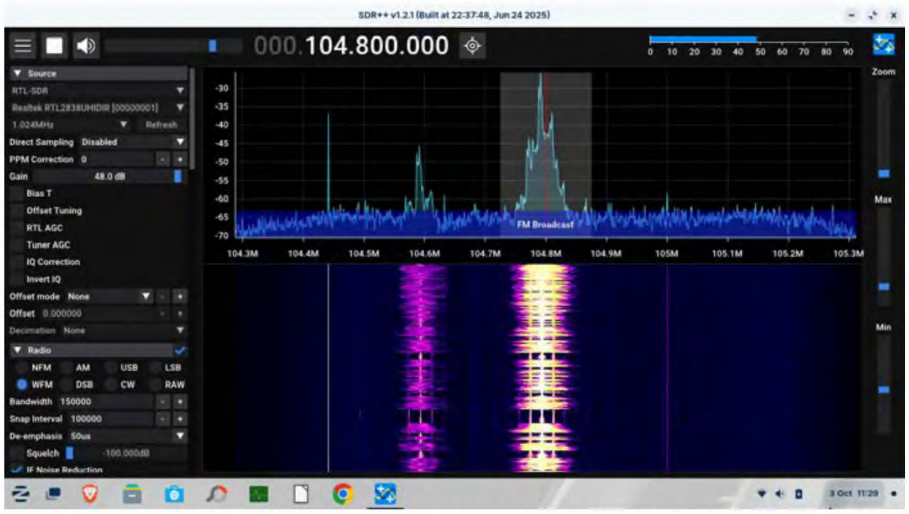


Fig.4.16: using SDR++ to receive the same local FM broadcast station as in Fig.4.13.

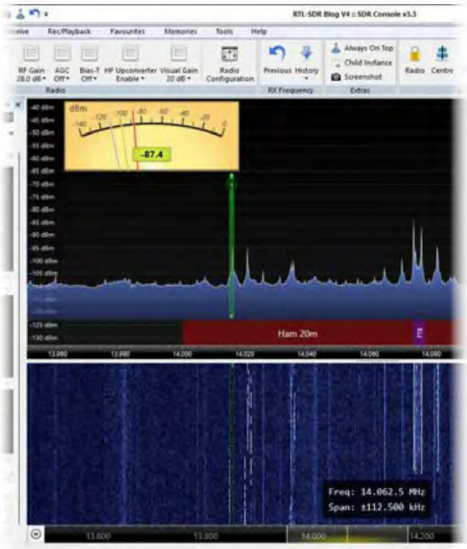


Fig.4.15: SDR Console receiving amateur radio CW (Morse) transmissions on the 20m amateur band.

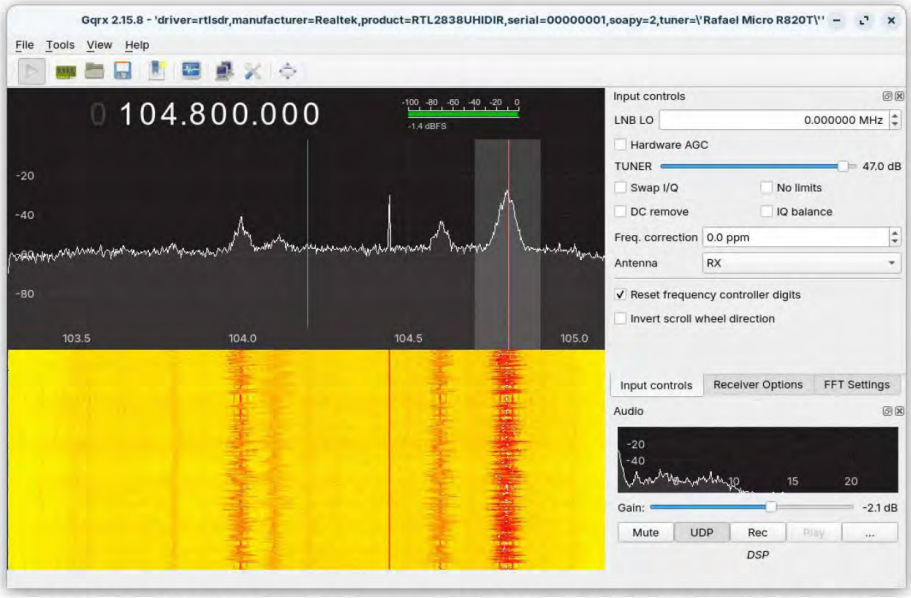


Fig.4.17: using GQRX to receive the same local FM broadcast station as in Fig.4.13.

oscillator is based on a simple twin-T arrangement with TR1 operating in common-source mode. A second emitter-follower stage, TR2, minimises the effects of loading on the oscillator circuit.

DC feedback from the emitter of TR2 to the gate of TR1 stabilises the operating conditions, while presets RV1 and RV2 provide gain adjustment and output level control, respectively.

To ensure the purity of the output signal, RV1 is adjusted for the minimum gain needed for reliable oscillation, while RV2 is adjusted for the required output amplitude. If TR1's swing is too large, it could generate undesirable harmonics.

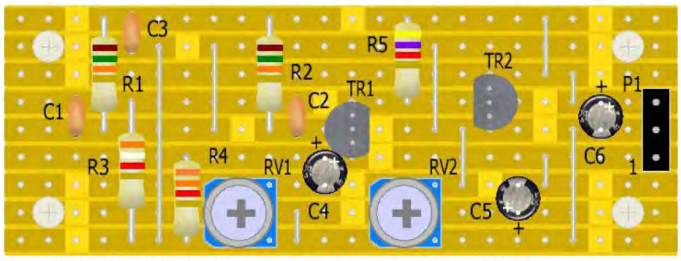


Fig.4.19: the AF oscillator module's component layout.

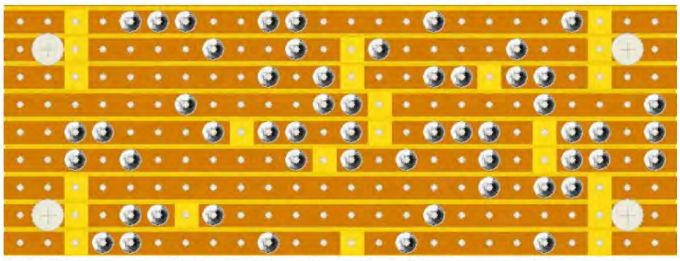


Fig.4.20: the track layout for the AF oscillator module.

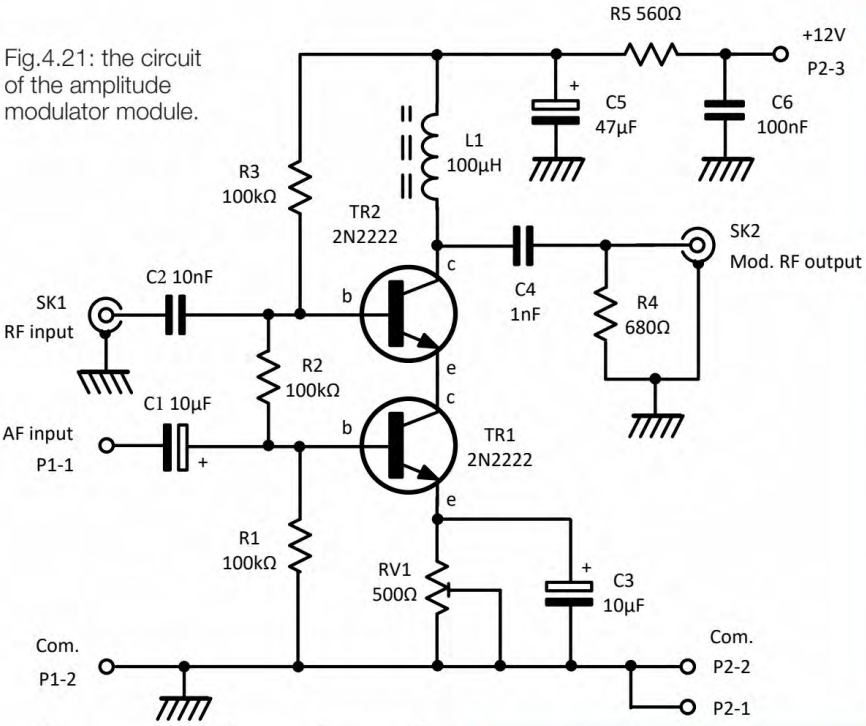


Fig.4.21: the circuit of the amplitude modulator module.

Parts List – AF osc.

1 25 × 64mm piece of stripboard (9 × 24 holes)
1 3-pin male 0.1in/2.54mm header (P1)

Semiconductors
1 2N7000 enhancement-mode N-channel Mosfet (TR1)
1 BC548 NPN transistor (TR2)

Capacitors
1 47µF 16V radial electrolytic (C6)
2 10µF 16V radial electrolytic (C4, C5)
1 100nF 50V ceramic (C3)
2 10nF 50V ceramic (C1, C2)

Resistors (all ¼W axial, ±5% or better)
2 15kΩ (R1, R2)
1 4.7kΩ (R5)
1 3.9kΩ (R3)
1 3.3kΩ (R4)
2 500Ω miniature preset potentiometers/trimpots (RV1, RV2)

The component layout for the AF oscillator module (viewed from the top) is shown in Fig.4.19, while the corresponding track layout (viewed from below) is in Fig.4.20. The required track breaks can be made using a spot face cutter or small drill bit, and the links on the upper side of the boards are made using short lengths of tinned copper wire.

Amplitude Modulator

The circuit of the amplitude modulator is shown in Fig.4.21. The module uses a direct-coupled cascode arrangement where the AF modulating and RF carrier signals are applied respectively to the bases of TR1 and TR2. The amplitude-modulated output is then extracted from the collector of TR2. An output of approximately 1.5V RMS is developed across the load resistor, R4.

The component layout for the AM module (viewed from the top) is shown in Fig.4.22, while the corresponding track layout (viewed from below) is given in Fig.4.23. As before, the required track breaks can be made using a spot face cutter

Parts List – Amplitude modulator module

1 25 × 64mm piece of stripboard (9 × 24 holes)
1 2-pin male 0.1in/2.54mm header (P1)
1 3-pin male 0.1in/2.54mm header (P2)
2 PCB-mounting female SMA connectors (SK1, SK2)
1 100µH axial RF inductor (L1)

Semiconductors
2 2N2222 NPN bipolar junction transistors (TR1, TR2)

Capacitors
1 47µF 16V radial electrolytic (C5)
2 10µF 16V radial electrolytic (C1, C3)
1 100nF 50V ceramic (C6)
1 10nF 50V ceramic (C2)
1 1nF 50V ceramic (C4)

Resistors (all ¼W axial, ±5% or better)
3 100kΩ (R1, R2, R3)
1 680Ω (R4)
1 560Ω (R5)
1 500Ω miniature preset potentiometer/trimpot (RV1)

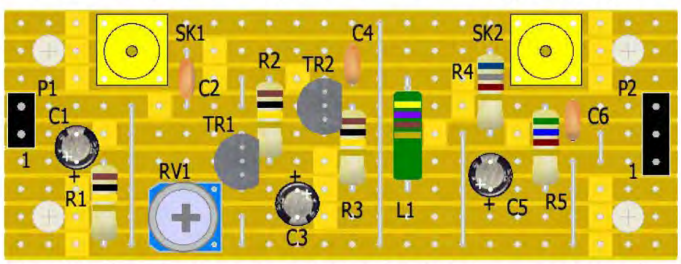


Fig.4.22: the amplitude modulator module's component layout.

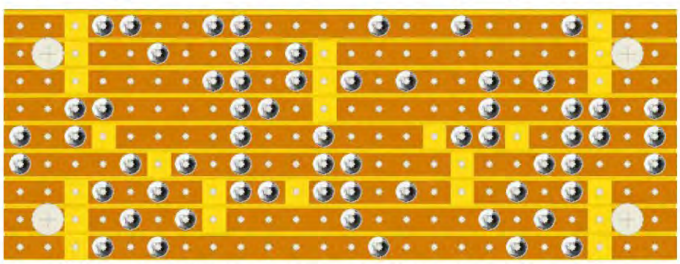


Fig.4.23: the track layout for the amplitude modulator module.

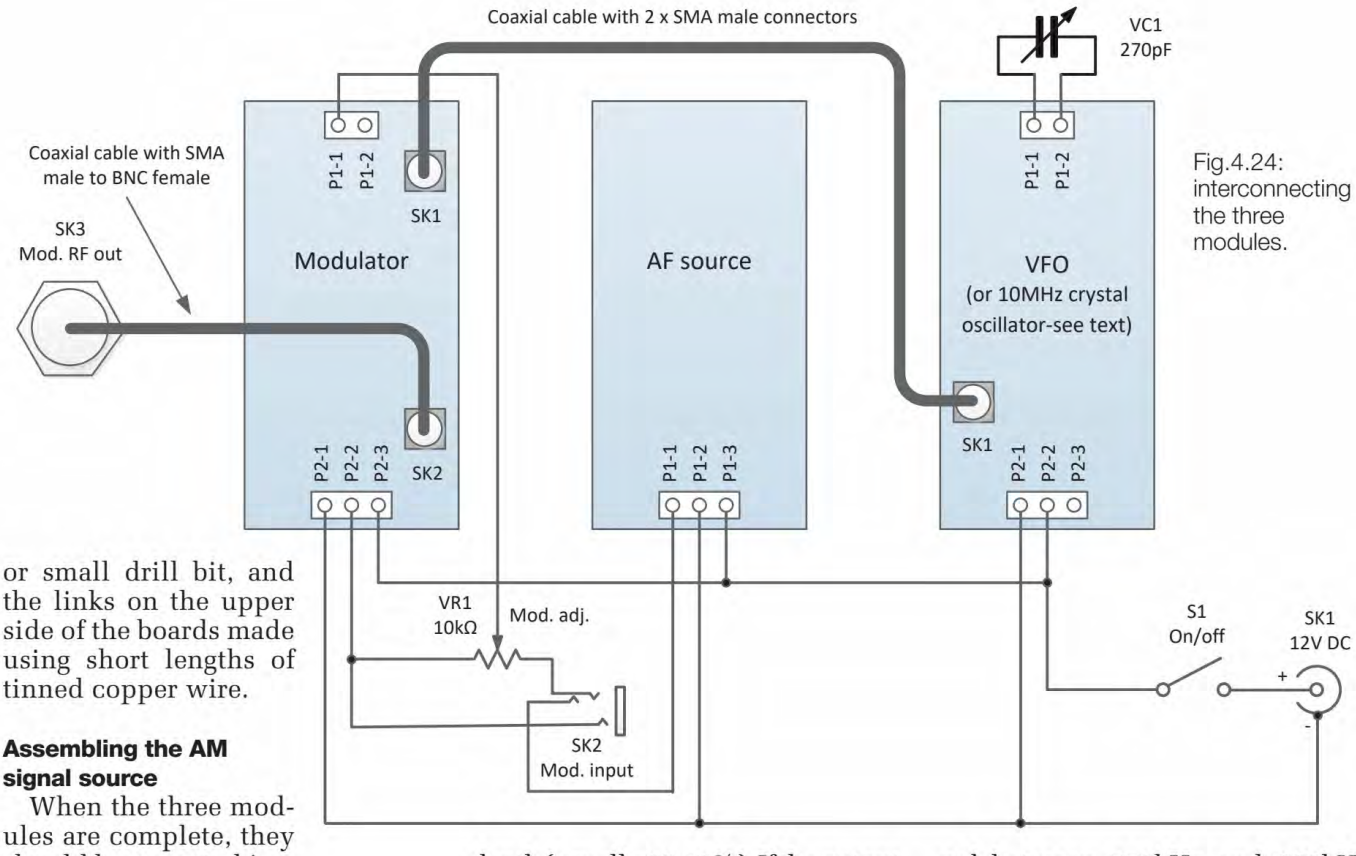


Fig.4.24: interconnecting the three modules.

or small drill bit, and the links on the upper side of the boards made using short lengths of tinned copper wire.

Assembling the AM signal source

When the three modules are complete, they should be mounted in a suitable ABS enclosure and interconnected as shown in Fig.4.24. Refer to the parts list on page 32. This diagram shows our VFO module from last month being used as the RF source, but if you prefer to use the 10MHz frequency standard (also described last month), the wiring from P1 to VC1 can be ignored.

To avoid the somewhat tedious task of manually fitting the SMA and BNC connectors, we recommend the use of ready-made coaxial cables to link the VFO to the modulator RF input and to link the modulated RF output to the panel-mounted BNC output connector. Both cable assemblies are available at a reasonable cost from online suppliers. Our prototype wiring is shown in Fig.4.25.

Testing

As always, it's important to check the stripboard and internal wiring before applying power. When these checks are complete, apply power to the module, adjust all presets as well as the modulation control and VC1 to mid-position and then switch it on. If you have an oscilloscope, it can be connected to the BNC connector so that you can observe the modulated RF waveform.

The modulation control can then be adjusted for the correct modulation depth (usually 30-50%). If the output appears unmodulated at all settings of VR1, check the AF oscillator and, if necessary, adjust RV1 just beyond the point at which oscillation starts. RV2 can then be adjusted to set the modulation depth in conjunction with the front-panel control, VR1.

For testing the VFO-driven modulator, you will require the services of an AM receiver covering the medium-wave band. This can be virtually any domestic portable receiver or either of the two simple receivers described previously in this series.

Connect a short length of hookup wire to the centre of the BNC connector and position the receiver about 1m away from the modulator. Tune the receiver to a vacant channel between 750kHz and 900kHz (ie, one without a broadcast AM station near it), then tune VC1 until a strong modulated signal is heard.

At this point, it's worth experimenting with the setting of the modulation depth control, RV1. If you increase the modulation depth slowly, you will notice that the amplitude of the received audio increases.

As the modulation depth reaches and then exceeds 100%, the signal will start to spread over a wider range of frequencies and begin to sound

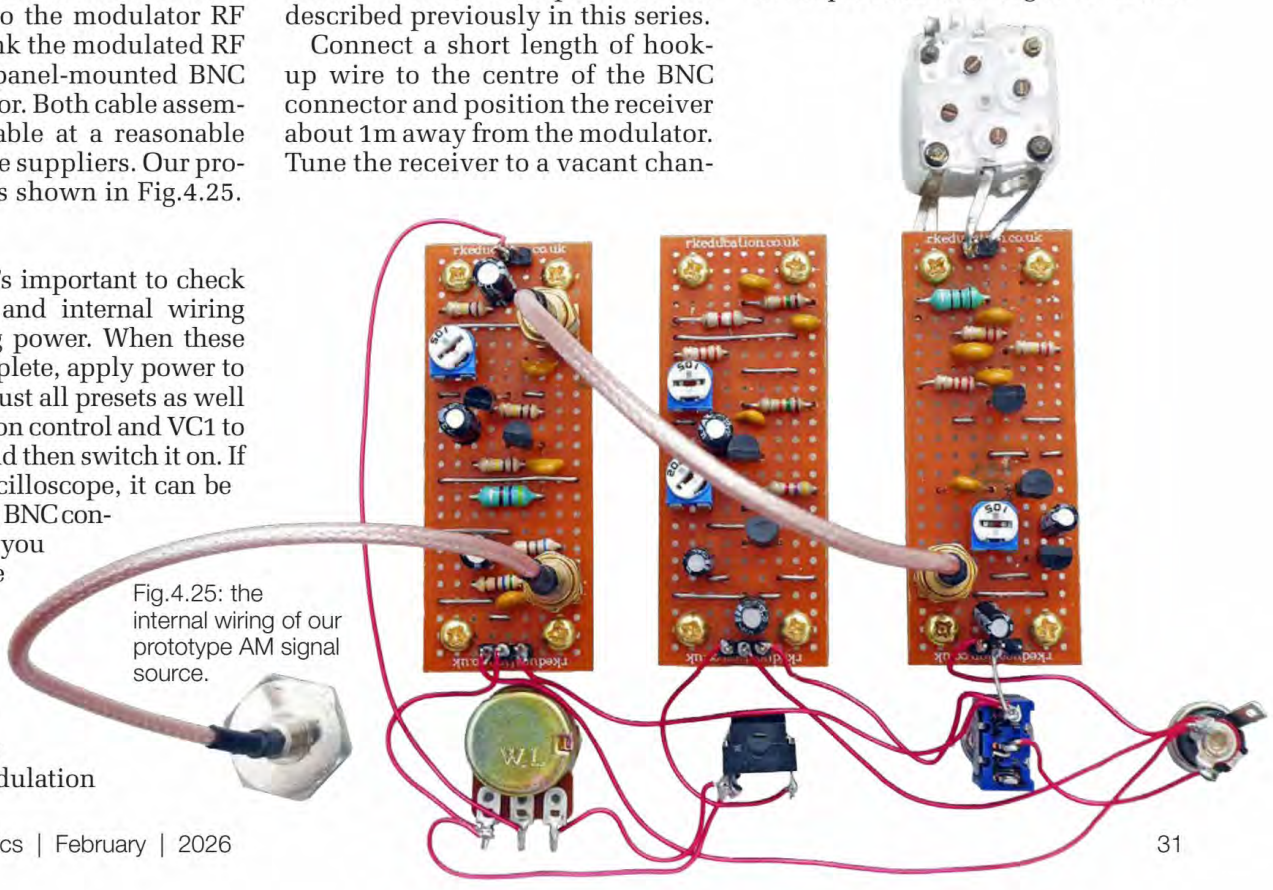
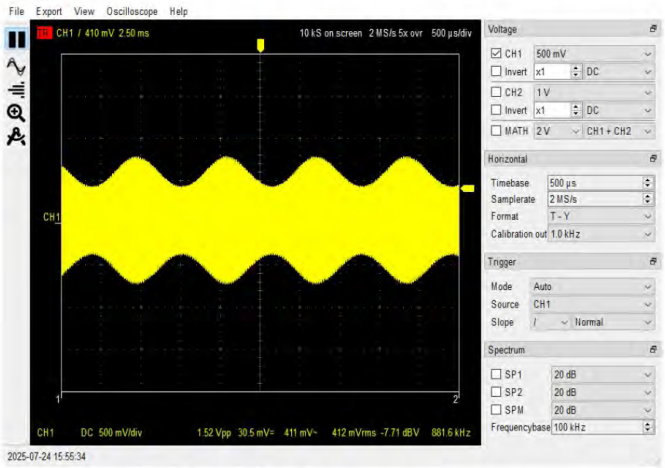
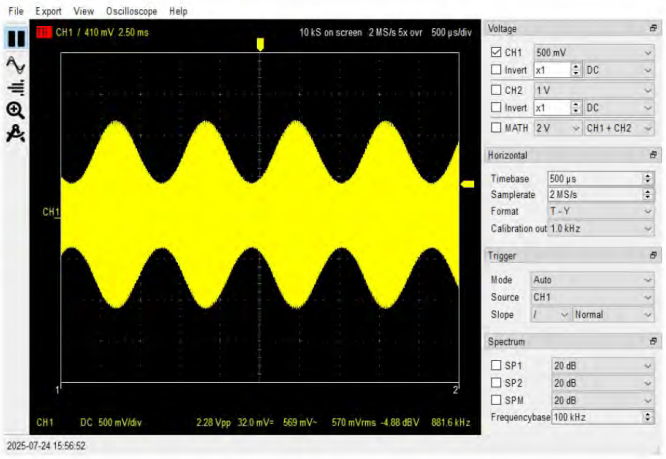


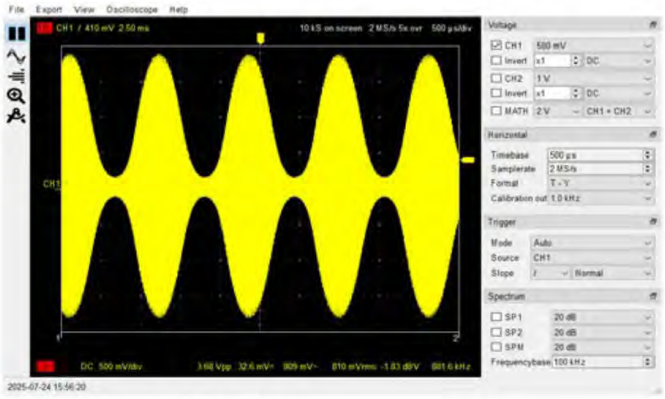
Fig.4.25: the internal wiring of our prototype AM signal source.



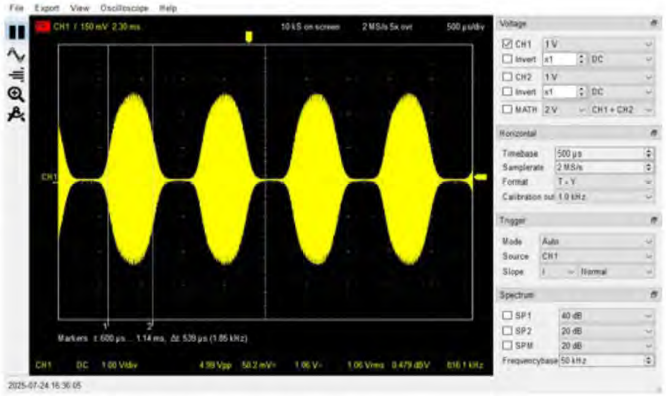
(a) 30% amplitude modulation



(b) 50% amplitude modulation



(c) 85% amplitude modulation



(d) >100% amplitude modulation

Fig.4.26: these examples shows the effect of increasing the modulation depth on the RF output from the amplitude modulator.

noticeably distorted (see Fig.4.26).

With the modulator working, the next step is to check its output over the full range of adjustment of VC1, tuning the radio at each stage to locate the signal and again noting the effect of varying VR1.

If you have been using the 10MHz crystal-controlled frequency standard instead of the VFO, you will require the services of a short-wave receiver, or you can use an SDR receiver like those described earlier.

If no output is detected from the modulator, disconnect the power and recheck the two new modules and off-board wiring. Tables 4.1 and 4.2 show the test voltages obtained from our AF oscillator and AM modulator modules, respectively.

Side frequencies and sidebands

If you have the RTL-SDR and companion software described earlier, it can be very instructive to examine the modulated signal in detail and see what happens to the AM waveform when the carrier is being modulated.

Tune VC1 to obtain an output as near to 900kHz as possible and set the SDR 'close-in' to display the spectrum on either side of the carrier frequency. Increase the modulation depth slowly and observe the display (see Fig.4.27).

Fig.4.27(a) shows the carrier in its unmodulated state. Note how this appears as a single frequency (vertical line) at 900kHz. Fig.4.27(b) shows the effect of increasing the modulation depth to around 30%. The 900kHz carrier is still present, but two new 'side frequency' components have appeared with amplitudes below that of the carrier.

The lower side frequency (LSF) appears at 899.1kHz, 900Hz below the carrier. The upper side frequency (USF) is evident at 900.9kHz, 900Hz above the carrier. This should make sense when you realise that the frequency from our AF source is 900Hz. The important outcome from this is that our amplitude modulated signal now occupies a range of RF frequencies extending from 899.1kHz to 900.9kHz (900kHz ±900Hz).

If you increase the modulation depth so it approaches and then exceeds 100%, you will have a spectrum that looks like Fig.4.27(c). Here, several more side frequencies have appeared on either side of the carrier. These correspond to the second, third, fourth and fifth harmonics of the modulating signal, and the total bandwidth has extended to around 9kHz (900kHz ±4.5kHz).

Parts List – off-board components

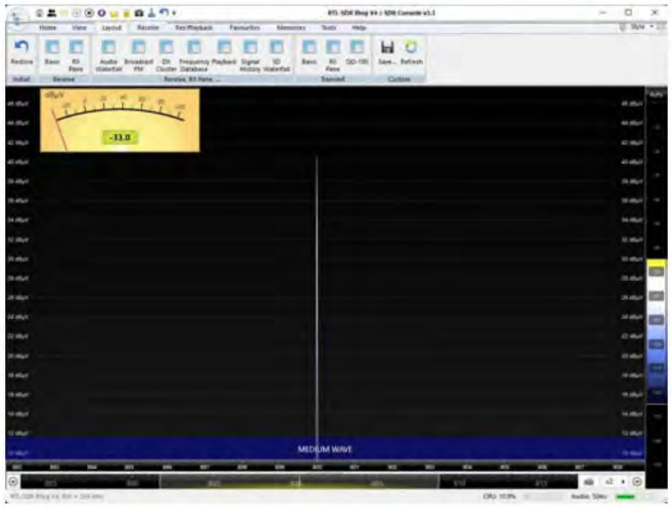
1 ABS enclosure
1 15cm bulkhead panel-mount BNC female to SMA male coaxial cable adaptor
1 15cm male-to-male SMA coaxial pigtail cable
1 panel-mount 5.5mm DC jack connector (SK1)
1 SPST miniature toggle switch (S1)
1 3.5mm miniature jack connector with switch contacts (SK2)
1 10kΩ linear variable potentiometer (VR1)
12 brass or nylon M3 × 10mm hex spacers/standoffs
24 M3 panhead machine screws

Table 4.1 – AF osc.

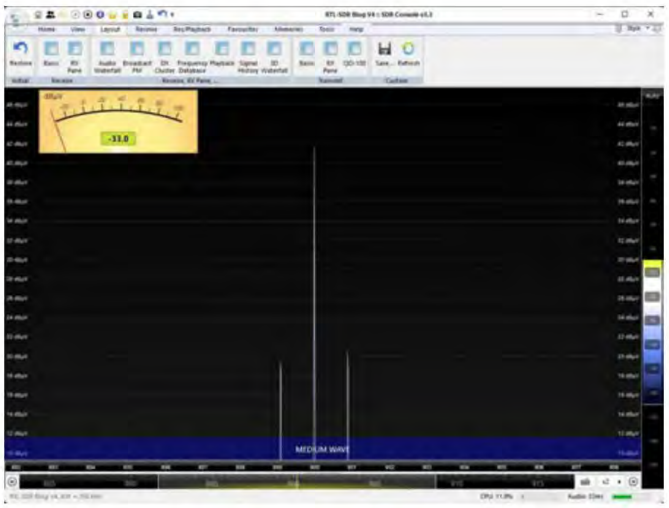
Dev.	Pin	Voltage
TR1	D	6.5V
	G	5.8V
	S	3.8V
TR2	C	12.0V
	B	6.5V
	E	5.9V

Table 4.2 – modulator

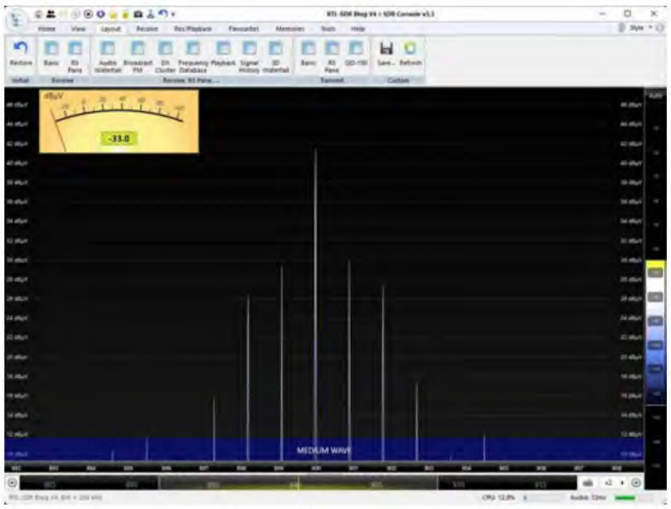
Dev.	Pin	Voltage
TR1	C	4.0V
	B	0.8V
	E	0.1V
TR2	C	10.5V
	B	4.7V
	E	4.0V



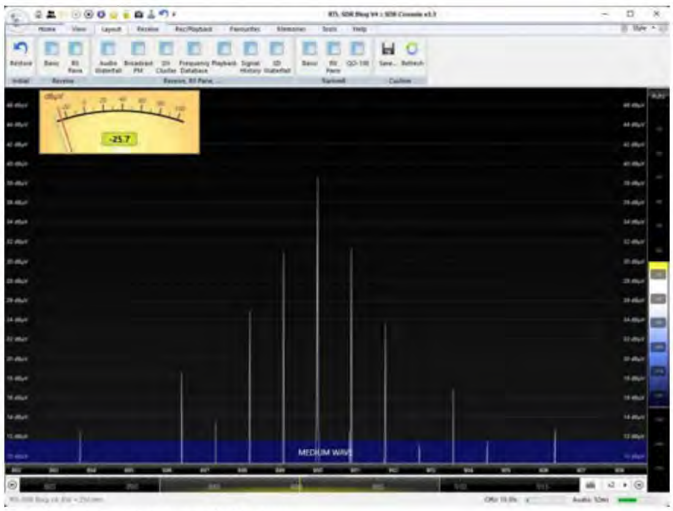
(a) Unmodulated carrier



(b) 30% amplitude modulation



(c) Modulation exceeds 100%



(d) Gross overmodulation

Fig.4.27: this shows effect of increasing the modulation depth on the output frequency spectrum of the amplitude modulator.

Increasing the modulation depth further causes the signal to spread even more, as in Fig.4.27(d). This grossly over-modulated signal occupies a bandwidth of around 12.5kHz, with up to the 7th harmonic appearing in the side frequencies.

Measuring modulation depth

Modulation depth can be easily determined from the waveform of the AM signal displayed on an oscilloscope. Fig.4.28 shows how this is done. If V_m is the peak-peak value of the modulating signal and V_c is the peak-peak value of the unmodulated carrier, the depth of modulation, M, expressed as a percentage is calculated as $M = (V_m \div V_c) \times 100\%$. Using the values observed in Fig.4.28, $M = 0.75 \div 1.5 \times 100\% = 50\%$.

External modulation

You might like to try an external modulation source instead of the 900Hz tone. This can be achieved by simply inserting a 3.5mm jack into the "Mod. Input" connector, SK2.

The audio signal (speech or music) can be taken from any media player or even a small dynamic microphone. A signal of around 100mV will produce 100% modulation with the "Mod. Adj." control, VR1, rotated to its fully clockwise position. Have fun with your AM modulator, but ensure that the signal is only detectable within the confines of your own property!

Coming up!

Part 5 of *Teach-In* will focus on antennas, covering how several popular types are built, how they work and how well they perform. We'll demonstrate easy ways to construct several basic antennas and discuss why matching them properly with your equipment matters.

We will also explore antenna feeders, cables, connectors, handy accessories and test equipment. As a practical Hands-On project, we'll build a useful variable frequency dipole antenna. PE

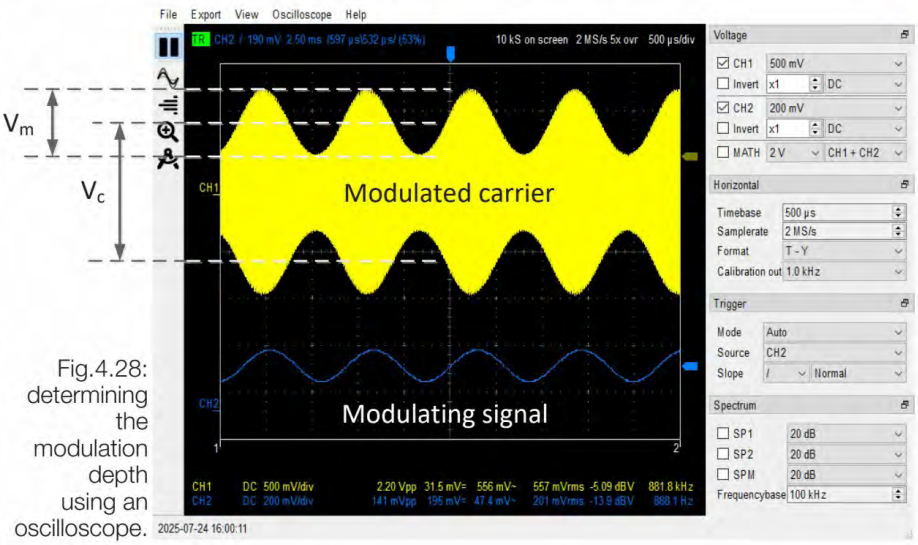


Fig.4.28: determining the modulation depth using an oscilloscope.

Project by Randy Keenan

Versatile Waveform Generator



This versatile waveform generator (also known as a function generator) is handy for a variety of uses, including audio equipment analysis, circuit development, displays and demonstrations and as a pulse source for developing switching and motor controls. It uses three op amps to deliver square, pulse, triangle, ramp and sine waves from 1Hz to 30kHz.

Waveform generators are often built around specialised ICs, such as the Exar XR2206, Intersil 8038 or the Maxim MAX038. However, I wanted to make a waveform generator using only generic components, like op amps, with these features:

- Output frequencies covering the audio range and more, from 1Hz to 30kHz.
- Waveform outputs of:
  - a. square/pulse, variable from 5% to 95% duty cycle, or wider
  - b. triangle/ramp/sawtooth, variable from positive to negative ramps
  - c. sinewave with a total harmonic distortion (THD) of around 1%
- Duty cycle/symmetry adjustments do not alter the frequency or amplitude appreciably
- Output amplitudes of the three waveforms can be matched, peak or RMS, from 0V to 6V peak-to-peak.
- An output impedance less than 200Ω.
- Battery-powered for portability and isolation.
- Compact size.

The design presented here is the result. It uses three op amps, two voltage regulators, six diodes, plus passive components. If any of the specified ICs become scarce, others of the same or better specifications could be substituted.

Operating principle

The circuit needs to generate the three basic types of waveform: square/pulse, triangle/ramp and sine. Since producing triangle/ramps and sinewaves from a pulse is complicated, the design begins with an op amp integrator producing a repeating triangle/ramp waveform.

Referring to the block diagram, Fig.1, the integrator at left produces the triangle/ramp waveform, with its frequency range set by switching in one of nine different integrator capacitor values. The triangle/ramp waveform is fed to a comparator that turns it into a square/pulse waveform, which is then fed back via the frequency adjustment pot to ensure oscillation. This gives us the triangle and ramp waveforms.

The two diodes and symmetry adjustment pot allow the positive and negative ramp rates to be varied to give square/pulse output waveforms.

Modifying (shaping) the triangular waveform by a separate circuit section converts it into a sine shape. While the result is not a perfect sinewave, it's pretty close, as demonstrated by its relatively low distortion/THD figure of about 1%.

The waveforms are selected by the middle switch, buffered and level-adjusted by IC3, and then fed to the outputs.

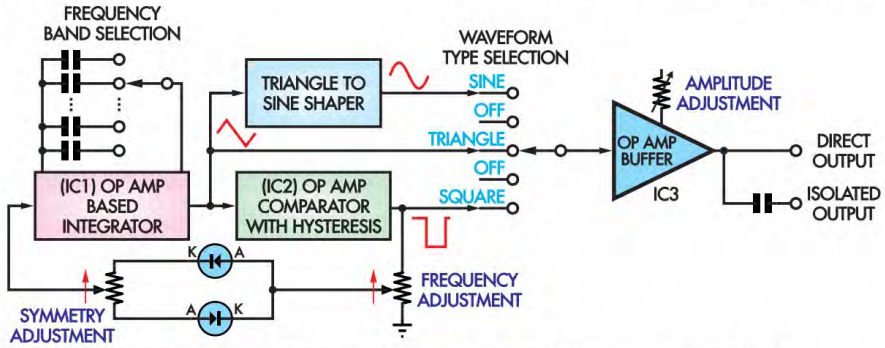


Fig.1: the Waveform Generator is designed around three op amps. IC1 is configured as an integrator and its output feeds into IC2, acting as a comparator, which feeds back into IC1. This feedback loop causes both to oscillate, with IC1 generating a triangular or sawtooth waveform and IC2 producing a square or pulse wave. A triangle-to-sinewave shaper produces the third waveform option.

Circuit details

The full circuit is shown in Fig.2. The heart of the circuit is the integrator composed of op amp IC1. It uses capacitors as the timing element and switched frequency range switch S1.

Versatile Waveform Generator

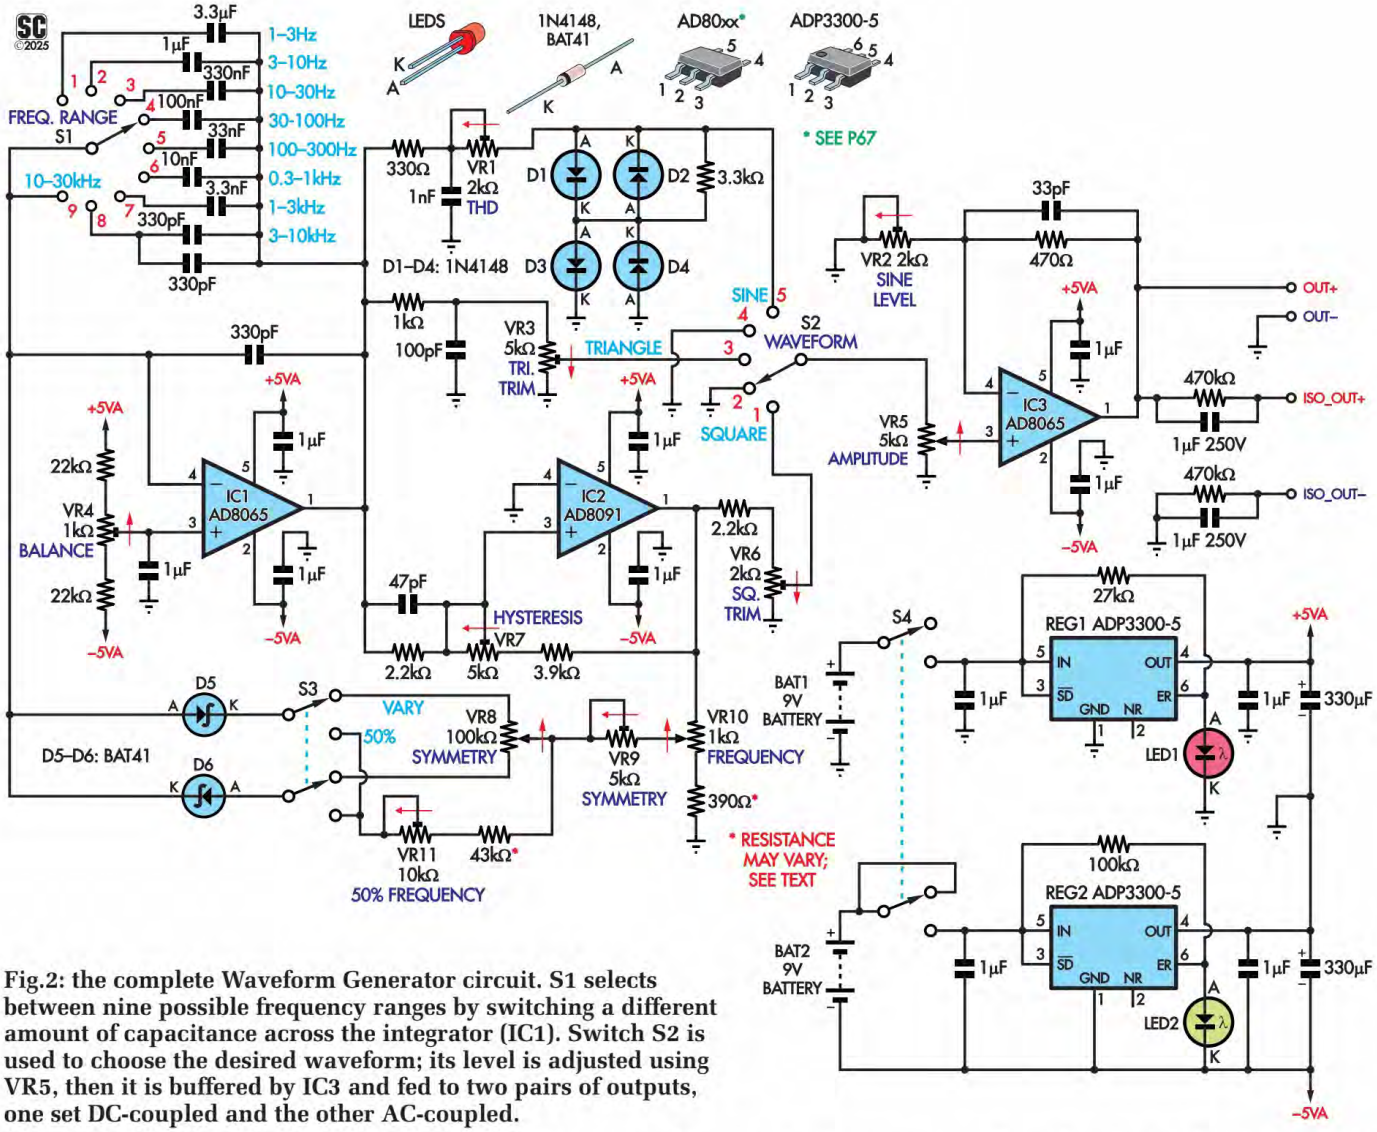


Fig.2: the complete Waveform Generator circuit. S1 selects between nine possible frequency ranges by switching a different amount of capacitance across the integrator (IC1). Switch S2 is used to choose the desired waveform; its level is adjusted using VR5, then it is buffered by IC3 and fed to two pairs of outputs, one set DC-coupled and the other AC-coupled.

The capacitor is charged and discharged via pot VR8, trimpots VR9 & VR10 and diodes D5 and D6. It works as follows.

Assume that initially the timing capacitor is discharged, and it is being charged by a current to pin 4 of IC1 through D6. IC1's output will be a linear negative-going ramp to counteract the increasing charge of the capacitor. The integration needs to be stopped at some point, so the op amp output is fed to a second op amp, IC2, configured as a comparator with hysteresis.

When IC1's output reaches the lower hysteresis voltage, set by trimpot VR7 and associated components, the comparator is triggered and its output goes negative, which is fed back to IC1's input via potentiometers VR10, VR9, VR8 and D5, which is now forward-biased. This causes the timing capacitor to start discharging, resulting in a positive-going linear output ramp from IC1.

This continues until IC1's output reaches the upper hysteresis voltage of the comparator, and the output of IC2 switches again, producing a negative-going ramp from IC1. Thus, the process of charging and discharging of the timing capacitor and switching of IC2's output continues indefinitely to produce an upward and downward ramp, plus a coincident square wave from the output of IC2.

Varying the duty cycle/ symmetry

The upward and downward slopes of the triangle or ramp are determined by the charging and discharging currents through the two arms of VR8. If VR8 is at its midpoint, the slopes are equal and a triangular wave is produced. If VR8 is off-centre, the currents through D5 and D6 are unequal, and a sawtooth waveform is produced.

Since the sum of the resistances to D5 and D6 and to IC1 is the same at any setting of VR8—equal to the total resistance of VR8—the period of the ramp, or triangle, will be constant regardless of its shape. (This is not quite true because of the non-ideal schottky diode characteristics and non-ideal characteristics of VR8, but it's pretty close.)

The setting of VR8 also determines the duty-cycle of the square wave/ pulse from IC2, since it depends on the periods of the upward and downward triangle wave ramps.

To vary the frequency, the square/ pulse output voltages from IC2 are adjusted by VR10 over a range of approximately 3:1. I chose this range to allow for precise setting of the frequency and to reduce non-ideal effects of the components.

To cover a wide range of frequencies, a series of nine charging/timing capacitors can be selected by rotary switch S1, as shown in Table 1.

Note that there is a 330pF capacitor always connected between pins 1 & 4 of IC1, and this is the only timing capacitor that is used on the highest (10-30kHz) range. It also adds to the switched-in capacitances on the 3-10kHz and 1-3kHz ranges, but for lower frequency ranges, its value is too small to have any real effect.

To obtain a precisely symmetric triangle or 50% duty-cycle square wave, the potentiometer's centre detent has to be pretty close to the point where the resistance from the wiper to each end of the track is identical. I have found that for a typical pot, the resistances of the two arms are not equal when set at the detent; furthermore, the detent generally has some 'wobble'.

Also, PCB-mounting potentiometers with a centre detent are not readily available. So, to ensure a symmetric waveform, the S3 "Symmetry" switch can be switched to its "50%" position, engaging VR11 and its 43kΩ series resistor for equal charging and discharging currents, and thus a fixed 50% symmetry. In the other position, S3 enables variable symmetry, as described earlier.

Table 1 – Timing capacitors

S1	Freq. range	Capacitance
1	1-3Hz	3.3µF
2	3-10Hz	1µF
3	10-30Hz	330nF
4	30-100Hz	100nF
5	100-300Hz	33nF
6	300Hz-1kHz	10nF
7	1-3kHz	3nF or 3.3nF *
8	3-10kHz	2 × 330pF
9	10-30kHz	330pF

** 3.3nF might make the 1-3kHz band too low in frequency*

Table 2 – Li-ion battery options

Type & size	Voltage	Capacity
6F22, "9V"	~8V (use two)	600-1300mAh
10440	~3.7V (use four)	350-1000mAh
14200/14250	~3.7V (use four)	~300mAh
14500	~3.7V (use four)	800-2500mAh

The final task is to produce a sinewave, and the method must work over the entire frequency range of the generator. In other words, it must be frequency-independent from 1Hz to 30kHz. This requires some non-linear circuit elements. There are various methods, but I chose a simple one.

Feeding the triangle wave to four diodes—two for positive and two for negative, plus a couple resistors—can reasonably approximate a sinewave. These diodes (D1-D4) should be closely matched, ideally from a single order and adjacent on a tape. This technique will never achieve a perfect sinewave, but it can come close (see Scope 3).

The waveforms square/pulse, triangle/ramp, and sine are selected by S2 and then buffered by op amp IC3 before being sent to the output terminals. Both direct and capacitor-isolated outputs are provided.

S2 is arranged with a pattern of square, off, triangle, off, sine for two reasons. Firstly, it provides some isolation among the waveforms, and secondly, having an off position or positions can be handy during use.

Because the sinewave from the shaper has the lowest amplitude of the three waves, the output op amp gain is adjusted, via trimpot VR2 ("Sine"), to accommodate the sinewave. Then the square/pulse and triangle/ramp amplitudes can then be adjusted via trimpots VR3 ("Tri") and VR6 ("Sq").

The wave amplitudes may be adjusted to either have equal peak amplitudes or equal RMS amplitudes, as desired. One reason for choosing equal RMS (root-mean-square) voltages is that each of the waveforms would deliver the same power to the load at the same setting.

Power supply

I wanted the waveform generator to be battery-powered for easy portability as well as electrical isolation.

The two batteries need sufficient voltage for the 5V voltage regulators (REG1 and REG2), meaning about 5.5V minimum, and preferably 7-8V. The specified batteries are "9V" (actually about 8V) lithium-ion rechargeable types.

Alternative rechargeable lithium-ion batteries are listed in Table 2, but check the capacity. I don't recommend using 14500 (AA-size) cells, as four are required, in two holders, and it's difficult to fit those into the specified enclosure.

Compared to 78L05 & 79L05 voltage regulators, the ADP3300-5.0s have a much lower dropout voltage and lower quiescent current use for lower battery drain. They also have the ability to drive dropout LED indicators (LED1 and LED2 in this circuit) and provide a more accurate regulated voltage.

The specified LEDs are high-brightness types for operation at low current and thus lower battery drain. The more accurate voltages, coupled with low-input-offset voltage op amps, reduces the need for compensation-adjustment circuitry. The ADP3300-5.0 is used for both the positive (IC4) and negative (IC5) voltage regulation. Thus, the batteries do not have a common connection.

If you use USB-rechargeable batteries with a double charging cable, be sure to remove the USB cables from the batteries before switching on the Waveform Generator as the circuit does not have a common battery connection, whereas the USB charging cables do have a common battery connection.

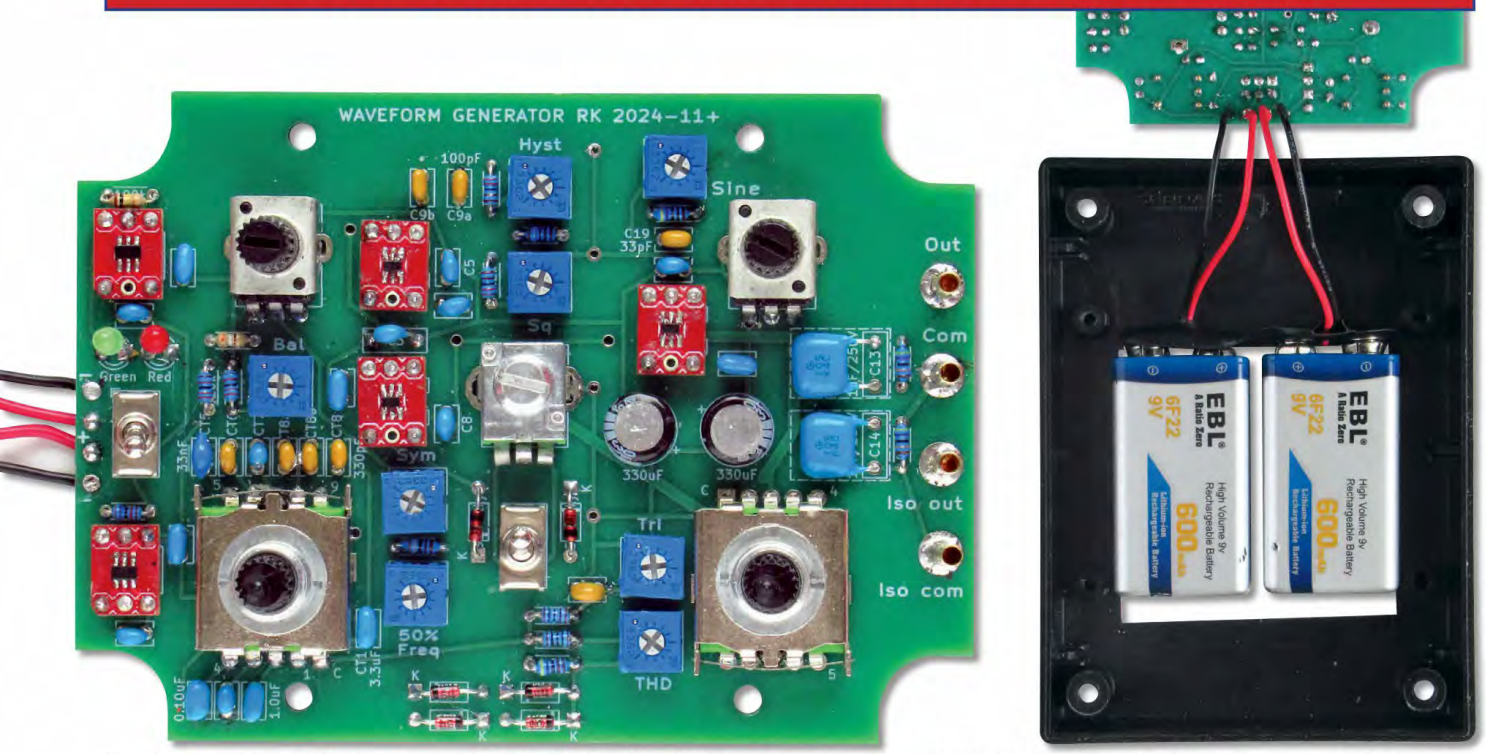
The current drawn from each battery is about 18mA each polarity, depending slightly on the frequency and waveform. Thus, the "9V" 600mAh batteries should provide about 20 hours (or more) of operation per charge, as confirmed by my trials, or twice as long for 1200mAh batteries.

A 220Ω load increases the current up to 26mA for a square wave output at 6V peak-to-peak, or several milliamperes lower for the other waveforms.

Part choices/variations

Two different parts are specified in the parts list for VR8, the Symmetry adjustment potentiometer. The P0915N version is better as it results in smaller frequency shifts at the extremes of symmetry/duty cycle, on the order of about 1-2%. Using the PTV09 version will probably result in larger frequency shifts.

However, if using the (better) P0915N version, its terminals will need to be reformed or trimmed and the two projections on the bottom—not the mounting tabs—will need to be removed so the pot will sit directly on the PCB. Since its shaft is smooth, you can drill out a knurled knob for a clean fit.



Photos 1 & 2: this PCB was assembled with the five SMDs on adaptor boards. Note how the miniature banana sockets on the right are soldered to the pads on the top of the PCB. I glued the 9V rechargeable batteries to the bottom of the enclosure and connected them to the PCB using standard battery snaps.

Unfortunately, potentiometers typically have a resistance tolerance of ±20%. Consequently, the values of some resistors may need to change depending on the actual resistance of the pots you get.

1. VR8's nominal value is 100kΩ. If yours measures above 100kΩ or below 92kΩ, you should ideally change the value of the 43kΩ resistor. Halve the measured value of VR8 and subtract 5kΩ, then pick the closest available value to use in place of the 43kΩ resistor.

2. VR10's nominal value is 1kΩ. If its value is below 935Ω or above 1.03kΩ, you should ideally change the value of the 390Ω resistor. Multiply VR10's actual resistance by 0.4 and then pick the closest available value to use in place of the 390Ω resistor.

A good alternative combination of op amps is AD8065 for IC1, either AD8051 or AD8091 for IC2, and AD8033 or AD8065 for IC3 (the AD8033 comes in a smaller package than the others, so will be more tricky to solder).

IC mounting

For the five surface-mount ICs, there are two mounting techniques: (a) directly on the PCB as surface mount, or (b) using adaptor boards with pins and receptacles. The main advantage of using adaptor boards is that you can unplug the ICs for testing and it's easy to replace them later (eg, for experimentation).

If you decide to use the adaptor boards, you can prepare them by first inserting five pins, long end down, in the appropriate pattern into a stably mounted DIL socket – see Photo 3. Then place an adaptor board, with the surface-mount pads upward, onto the pins and solder each pin (Photo 4).

With the pins attached, solder the IC to the pads using your preferred technique. There are a few ways to do it, either with a regular iron or hot air; the construction procedure below goes over our preferred method. Make sure that the orientation of the IC is correct (see Photo 5).

For the op amp ICs, finding the correct orientation is straightforward—they only have five leads. For the regulators, it's a bit more tricky as they are rotationally symmetrical; refer to the construction procedure below for instructions.

Inspect with a magnifying glass to verify that all leads have been soldered correctly. Pin sockets need to be inserted into the PCB to receive the adaptor board pins. It's best to temporarily attach the adaptor board, solder those socket pins to the main board,

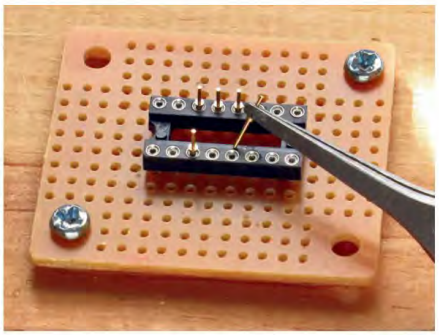


Photo 3: using a DIP socket as a jig to hold the PCB pins.

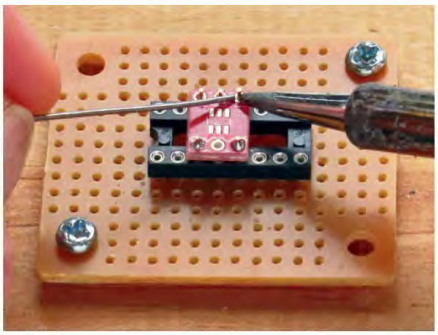


Photo 4: soldering the PCB pins to the SMD adaptor board.

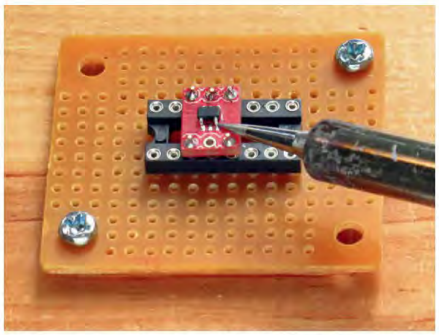


Photo 5: soldering the SMD IC to the adaptor board.

then unplug it before you power it up.

Construction

The Waveform Generator is built on a double-sided PCB coded 04104251 that measures 101.5 × 73.5mm. The following instructions assume you will be soldering the three op amp and two regulator ICs directly to the PCB pads.

If you want to use adaptors instead, the procedure is not terribly different except that you will be soldering those parts to the adaptors, then fitting the adaptors with pins and soldering matching sockets to the sets of five through-hole pads arranged around each chip location.

Start by soldering the five SMDs. In each case, spread a thin layer of flux paste over the PCB pads first. The op amps, IC1-IC3, each have five pins with two on one side and three on the other, so the correct orientation of each should be obvious. Place the part on the board, tack-solder one pin and check that the device is flat on the board and each lead is centred over its pad.

If not, remelt the initial solder joint and gently nudge the part into place. Repeat if necessary until it is nicely aligned, then solder the remaining pins. Add a small amount of flux paste to the first pin and touch it with a clean soldering iron tip to reflow the joint.

Given that these leads are quite close together, you may have accidentally bridged two or more pins. Use a magnifier to check.

If you have, it's quite easy to correct: simply add a small amount of flux paste to those pins, put the end of some solder-wicking braid on top and press it down onto the board and pins with your soldering iron. Wait for a few seconds until the solder melts, then drag the wick away from the pins and lift it and the iron off the board. That should leave behind just the right amount of solder.

REG1 and REG2 are similar to IC1-IC3, but they're a bit more tricky because they have three pins on each side. That means you'll have to figure out which of the two possible orientations is correct. The PCB is missing a pad on one side because pin 2 of these devices is not used.

Examine the chip under magnification and find the pin 1 indicator in one corner. Rotate it so that corner is next to the missing central pad, then tack-solder one pin. Proceed with soldering as for IC1-IC3 but of course you can skip the pin which has no corresponding pad. You should still check for bridges to pin 2 (however unlikely they are) and fix them if present.

If you manage to solder them in the wrong orientation, simply remove the middle pin and resolder it on the other side of the adaptor.

Now move on to fit the through-hole resistors and diodes. The orientations of the resistors do not matter but the diodes do, so make sure their cathode stripes face as shown in the overlay diagram (Fig.3). Also, don't get the similar-looking 1N4148 (standard silicon, D1-D4) and BAT41 (schottky, D5 & D6) diodes mixed up.

Note that the resistors used are smaller than the standard ¼W or ½W types generally used in our projects. As ¼W resistors won't fit in the specified case, we recommend you use ⅛W or ⅛W miniature body resistors.

There are many resistor values used, so refer to the colour code table in the parts list or use a DMM set to measure ohms to ensure they go in the right locations.

Follow with the capacitors, none of which are polarised except for the two larger electrolytics. Their longer (positive) leads face each other, as shown by the + marks on Fig.3. While many of the ceramic capacitors are 1µF types, there are quite a few different values, so don't get them mixed up.

The two larger 1µF 250V caps go near the output terminals as shown, laid over as otherwise they will be too tall to fit in the enclosure later.

Next, fit the trimpots. There are eight in four different values, so again, make sure the right ones go in the right locations. Note that the footprints for the trimpots on the PCB have four pads, while the trimpots have three pins. This is to allow you to use either the common 3362P types or the less-common 3362R reversed version.

Fig.3 shows the correct orientations for 3362P trimpots, and the PCB also has "P" and "R" labels on the two possible locations for the central pin. If using 3362R trimpots, rotate them 180° compared to what's shown in Fig.3, so the central pin goes into the pads marked "R" on the PCB.

Testing

If you are using adaptors for the op amps, you can test the board before connecting any of the expensive op amps to the circuit. Connect the batteries, plug in the two regulators

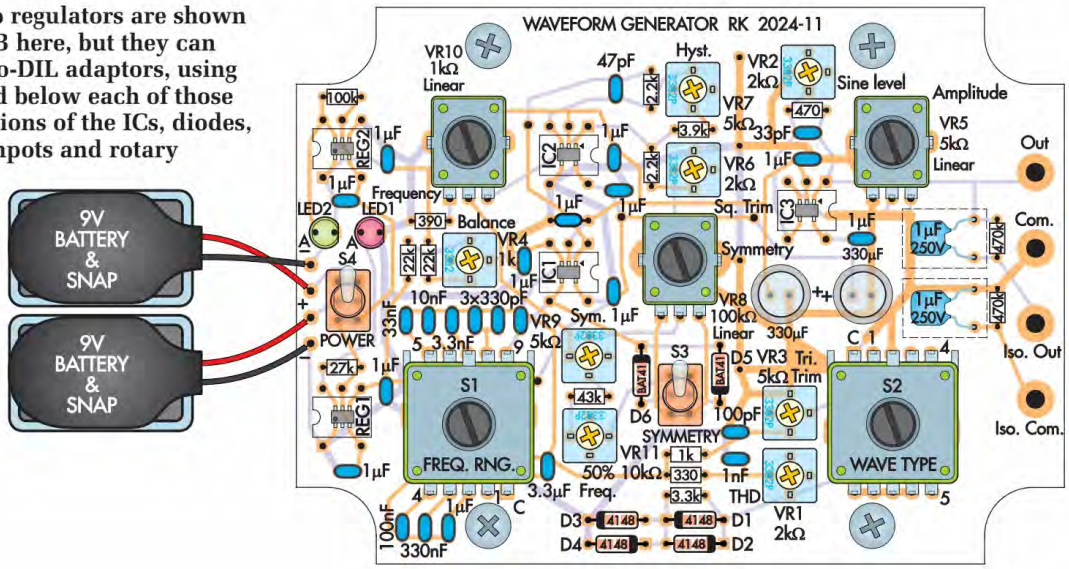


Fig.3: the three ICs and two regulators are shown soldered directly to the PCB here, but they can also be attached via SMD-to-DIL adaptors, using the rows of holes above and below each of those devices. Watch the orientations of the ICs, diodes, electrolytic capacitors, trimpots and rotary switches. The two LEDs indicate both when it is switched on and also whether the 9V batteries are still OK. Also note the way the batteries are wired – there is no reverse polarity protection!

and switch the power on; both LEDs should light up. When connecting the batteries, it is best to have the power switch off; otherwise, accidentally touching a connector with the wrong polarity could damage a voltage regulator.

Using the output ground ("COM.") as a reference, measure the voltages at pins 2 & 5 of IC1 (you can use the larger through-hole pads or sockets rather than trying to probe the SMD pads). Pin 2 is at top centre and should measure -4.98V to -5.02V, while pin 5 is at lower-right and should measure +4.98V to +5.02V. If not, switch off and check for faults.

If you've soldered these ICs directly to the board, you can still perform this test, but there is a risk of damaging the ICs if something is wrong with the regulators. So check the orientation of REG1 & REG2 carefully before switching on, as well as the polarity of the batteries and their wiring (you can do this by probing the battery terminals on the PCB with a multimeter).

If everything checks out, and you have socketed the ICs, switch the power off and plug in IC1, IC2 and IC3. Make sure they're all orientated correctly, with the sides with two pins facing towards the bottom of the PCB.

Set the Amplitude control (VR5) to maximum and the Waveform switch (S2) to square wave. Set Symmetry (S3) to the 50% position, and all trimpots to around midrange. When power is switched back on, there should be a square waveform—or nearly so—at the output, centred at 0V.

Troubleshooting

Are both LEDs on? If not, the batteries, voltage regulators and associated circuitry need attention. If they're on but there's no output, check that the Waveform switch (S2) is not at one of the off positions and that the Amplitude control (VR5) is not at or near minimum. Try adjusting trimpot VR7 ("Hyst").

As usual, if you run into any problems, check that the ICs and diodes are all in the correct orientations. Remove the ICs, if using adaptor boards, and verify the supply voltages again. Check that the resistors and capacitors are all the correct values. Look for unsoldered pins or wires, and for solder bridges on both sides of the PCB.

If you're still stuck, check the output of IC1 at pin 1 (upper right).

Parts List – Waveform Generator

1 double-sided PCB coded 04104251, 101.5 × 73.5mm
1 Serpac 131,BK plastic enclosure **[Mouser 635-131-B]**
1 panel label, 104 × 74mm
2 9-position vertical rotary switches, 18t split shafts (S1, S2)
[Taiwan Alpha SR1712F-0109-15K0A-N9-N-027]
2 miniature PCB-mount vertical DPDT toggle switches (S3, S4)
[Nidec ATE2D-2M3-10-Z]
4 miniature 2mm banana sockets **[Amazon B08M3BCSB1]**
5 SOT-23-6 to DIL breakout boards (**optional**) **[SparkFun BOB-00717]**
25 0.51mm diameter PCB pins (**optional**)
[DigiKey ED90325-ND, Mouser 575-90810001508]
25 0.51mm diameter PCB pin sockets (**optional**)
[Mouser 575-3016015152127]
2 9V rechargeable batteries [eg, 600mAh EBL6F22] (BAT1, BAT2)
2 9V battery snaps with flying leads (BAT1, BAT2)
5 knobs to suit the 18t spline shafts of S1, S2, VR5, VR8 & VR10
4 3mm inner diameter, 1mm-thick plastic or fibre flat washers
4 No.4 × 8mm self-tapping screws
4 stick-on rubber feet

Semiconductors
2 AD8065ART op amps, SOT-23-5 (IC1, IC3; see text for other options)
1 AD8091ART op amp, SOT-23-5 (IC2; see text for other options)
2 ADP3300ARTZ-5 low-dropout 5V linear regulators, SOT-23-6 (REG1, REG2)
1 high-brightness 3mm red LED (LED1) **[Kingbright WP710A10SRD/J4]**
1 high-brightness 3mm green LED (LED2) **[Kingbright WP710A10ZGDK]**
4 1N4148 or equivalent 75V 200mA signal diodes (D1-D4)
2 BAT41 or equivalent 70V 15mA schottky diodes (D5, D6)

Capacitors *(all 50V radial multi-layer ceramic, 2.5mm pitch unless noted)*
2 330µF 6.3V low-profile (5mm tall) radial electrolytic
[Panasonic ECE-A0JKS331]
1 3.3µF 25/50V X7R ±10% **[Murata RCER71E335K2DBH03A]**
2 1µF 250V X7R ±10% **[Murata RDER72E105K5B1H03B]**
12 1µF 25/50V X7R ±10% **[Murata RDER71H105K2M1H03A]**
1 330nF 25/50V X7R ±5% **[Kemet C333C334J5R5TA]**
1 100nF 25/50V NP0/C0G ±5% **[Murata RCE5C1H104J2A2H03B]**
1 33nF 25/50V NP0/C0G ±5% **[TDK FA14C0G1H333JNU00]**
1 10nF 25/50V NP0/C0G ±5% **[Kemet C315C103J3G5TA]**
1 3.3nF NP0/C0G ±5% **[Murata RCER5C1H332J0DBH03A]**
1 1nF NP0/C0G ±5%
3 330pF NP0/C0G ±5% **[Kemet C315C331J3G5TA]**
1 100pF NP0/C0G ±5% **[Vishay K101J15C0GH53L2]**
1 47pF ±5% **[TDK FG18C0G1H470JNT00]**
1 33pF NP0/C0G ±5% **[Vishay K330J15C0GF53L2]**

Potentiometers *(all 9mm vertical plastic pcb-mount 18t spline shaft types)*
1 5kΩ linear B-type (VR5) **[Bourns PTV09A-4030U-B502]**
1 100kΩ linear B-type (VR8) **[DigiKey 987-1708-ND – see text]**
1 1kΩ linear B-type (VR10) **[Bourns PTV09A-4020U-B102]**

Trimpots *(all 3362P-style miniature top-adjust)*
3 2kΩ (VR1, VR2, VR6)
3 5kΩ (VR3, VR7, VR9)
1 1kΩ (VR4)
1 10kΩ (VR11)

Resistors *(all ⅛W miniature axial 1%)*

2 470kΩ	1 3.3kΩ
1 100kΩ	2 2.2kΩ
1 43kΩ	1 1kΩ
1 27kΩ	1 470Ω
2 22kΩ	1 390Ω
1 3.9kΩ	1 330Ω

Value	4-band code	5-band code
330Ω		
390Ω		
470Ω		
1.0kΩ		
2.2kΩ		
3.3kΩ		
3.9kΩ		
22kΩ		
27kΩ		
43kΩ		
100kΩ		
470kΩ		

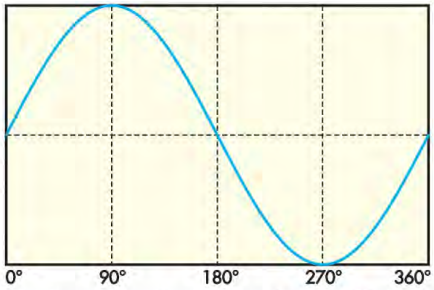


Fig.4: a pure sinewave shaped like this will have a low distortion figure, well under 1% THD. Try to get the output of your unit to match this as closely as possible.

If there is a triangle waveform, then IC1 & IC2 are working and IC3 may need attention.

If you're getting strange waveforms, verify that the schottky and regular diodes have the correct orientations. Check the values of the following components: the filter capacitors across VR3 and series diode pair D1 & D3, IC3's feedback capacitor, and compensation capacitor across the 2.2kΩ resistor from IC1's output to VR7.

Set-up and calibration

Calibration requires the following steps in sequence.

1. Set the Frequency Band switch (S1) to the 1-3kHz position. Set the Frequency pot (VR10) and all trimpots at approximately midrange.

2. Connect an oscilloscope to the lowest lead of a capacitor below S1, using the output common terminal as the reference.

3. Set the S3 Symmetry switch to the 50% position and apply power. A triangle wave should be displayed on the oscilloscope. Adjust trimpot VR7 (Hyst) so you get exactly 4V peak-to-peak. The triangle may be slightly asymmetrical; that will be fixed in step 5.

4. Connect the oscilloscope to the direct output terminal, set the Waveform switch (S2) to square wave mode and adjust VR5 for maximum amplitude. A square wave should be displayed on the oscilloscope.

5. Adjust trimpot VR4 (Balance) for an exactly symmetrical square wave. A multimeter with a duty-cycle measurement option would be useful here, or use a similar oscilloscope measurement. Adjust VR10 (Frequency) if necessary.

6. Set S3 to its alternative Vary position. Adjust trimpot VR9 ("Sym") so you get slightly less than 5% duty cycle with VR8 fully anti-clockwise and slightly more than 95% duty cycle with VR8 fully clockwise. The duty cycle can be pushed from 2% to 98%, but frequency shift may increase.

7. With S3 still in the Vary position, adjust VR9 (Sym) for an exactly symmetrical waveform. Note the frequency. Set S3 back to the 50% position and achieve exactly the same frequency by adjusting VR11 (50% Freq).

8. Set S3 back to the 50% position and S2 to sinewave. An approximate sinewave should be displayed.

Sinewave adjustment

9. Adjust trimpot VR1 (THD) to achieve the cleanest possible sinewave. You can trace Fig.4 onto tracing paper, baking paper or clear plastic and place it over the oscilloscope screen as a guide. Alternatively, if your 'scope has a spectrum analyser mode (or you have a spectrum analyser) adjust VR1 for minimum harmonics.

If you are not fussy, forming an approximation to a sinewave on a 'scope screen may be good enough.

If using a spectrum analyser, I suggest setting the Wave Generator frequency to 1kHz and the analyser frequency span to cover the audio range. Momentarily switch to triangle wave mode and adjust trimpot VR4 ("Bal") to minimise the second (2kHz) and all other even harmonics. This should only require a slight readjustment. Switch back to sinewave mode and adjust VR1 ("THD") to minimise the odd harmonics.

Then adjust trimpot VR1 (THD) to minimise the odd harmonics. VR7 (Hyst) may also be adjusted a slight amount, but this will also alter the frequency bands.

When you've finished, all even harmonics should be approximately 60dB lower than the fundamental and all odd harmonics (starting at 3kHz) should be at least 45dB lower than the fundamental. Adjust the amplitude setting as necessary to avoid overloading the spectrum analyser. A sinewave THD close to 1% should be achievable.

Wave amplitudes

10. Leaving the 'scope connected to the direct output and S2 in the sinewave position, set VR5 (Amplitude) to maximum. Now you have a choice of equal peak voltages or equal RMS voltages for the three waveforms. For equal peak voltages, decide on what maximum you want and adjust VR2 (Sine) to that maximum. I do not recommend greater than 6V peak-to-peak.

Next, set S2 to square wave mode and adjust VR6 to achieve the chosen maximum output level. Switch S2 to triangle wave mode and adjust trimpot VR3 (Tri) to achieve the same maximum level. Alternatively, to set the waveforms to equal RMS voltages, use Table 3 or an RMS-reading device (multimeter or oscilloscope).

11. Check that VR10 (Frequency) varies the frequency over a range of at least 3:1 and check the minimum

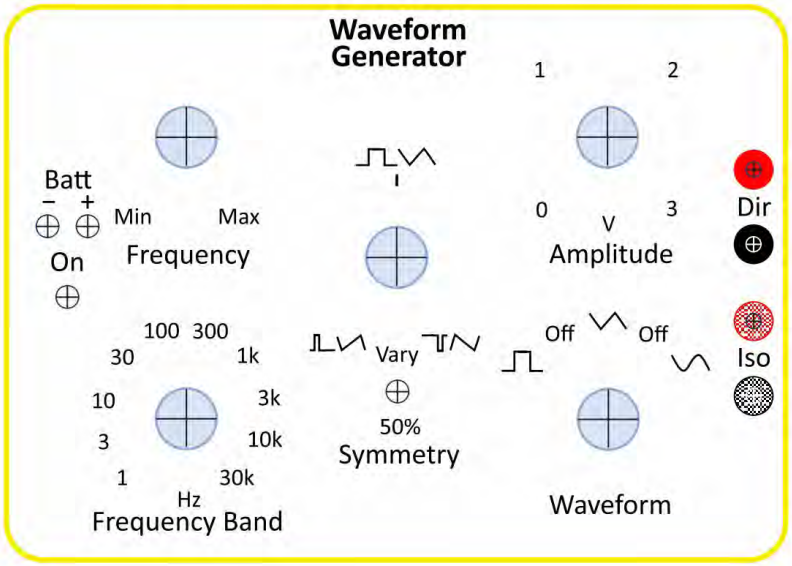


Fig.5: the controls are quite complicated so you'll need this panel label to understand what they all do. It will also help you locate the holes for the switch and potentiometer shafts, LEDs and banana sockets. You can download it as a PDF from our website and print it at actual size (1:1).

and maximum frequency for each band. The bands should overlap. If the minimums are not low enough, decrease the value of the 390Ω resistor. If the maximums are not high enough, adjust VR7 (Hyst) slightly and return to step 8.

The frequency bands will likely not track by exact factors because of the typical variations in capacitance of the timing capacitors. That's why these capacitors (all the ones that connect to pin 1 of IC1) should have a ±5% or better tolerance, if possible. In the worst case, you may need to replace one or two caps or parallel them with lower-value capacitors.

12. With S2 (Waveform) set to triangle wave and S3 (Symmetry) at the Vary setting, rotate VR8 (Symmetry) to both extremes to check that the triangle wave becomes a clean downward or upward ramp/sawtooth, and recheck that, on the square wave setting, the output becomes a pulse that varies in duty cycle between 5% and 95%.

Enclosure preparation

Fig.5 is a front panel label that can also be used as a drilling guide. You can download it from **siliconchip.au/Shop/11/1823**

We have instructions on preparing and attaching panel labels online, see: **https://pemag.au/Help/FrontPanels**

With the panel label attached, the holes can then be drilled through carefully. The final hole sizes are 3mm for the LEDs, 8mm for the potentiometers, 10mm for the rotary switches, 4mm for the toggle switches and 2.5mm for the banana sockets.

If possible, I suggest punching the small holes. I also suggest countersinking the small holes on the inside of the enclosure for easier insertion of the LEDs, switches and banana receptacles. The mounting post on the top part of the enclosure that is near rotary switch S2 needs to be trimmed back a bit to allow room for the switch. The anti-rotation tabs on the tops of the rotary switches and pots need to be removed.

Insert the LEDs and banana sockets into the PCB with the LEDs in the correct orientations, but do not solder them yet. Temporarily fit the PCB into the enclosure using a 1mm-thick non-conductive (eg, plastic or fibre) spacer or washer on each mounting post.

Top tip: use super glue to stick the washers in place temporarily (either to the enclosure or top of the PCB) so they don't slide out as you're trying to assemble everything.

Adjust the LEDs and banana receptacles as desired, then solder the LEDs, and tack-solder the sockets quickly to avoid melting the plastic. Remove the PCB and solder the sockets to the upper surface of the PCB, being careful to maintain their position. You can now screw the PCB into place in the enclosure on the 1mm spacers. Do not use panel-mount hardware on the rotary switches or VR8.

After considering several mounting methods for the batteries, I simply used a little epoxy to attach them to the lower part of the enclosure, with a piece of thick paper in between should I ever want to remove them. You could also consider foam-cored double-sided tape, although it may not be strong enough to hold them long-term.

Usage notes

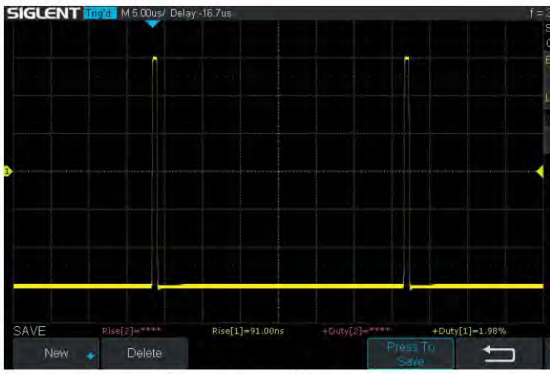
The square wave or pulse rise and fall times are approximately 90ns (see Scopes 1 & 2). There is a barely noticeable non-linearity in the triangle waves at the three lowest frequency bands. I attribute this to the capacitors, which are X7R for these bands.

The higher-frequency bands use C0G/NP0 capacitors and look perfectly linear to my eye. Using C0G or film capacitors for the higher-value timing capacitors would eliminate the slight non-linearity, but they are too large to realistically fit. For an explanation of capacitor types, see our detailed article on capacitors in the December 2024 issue.

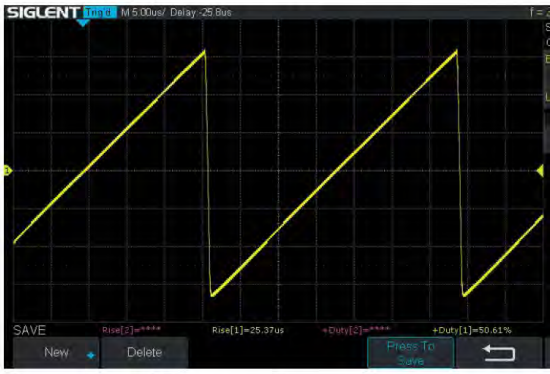
Scope 3 compares the Waveform Generator's quasi-sinewave (mauve) to a pure sinewave (yellow) at 1kHz; the pure sinewave was generated by sending the Waveform Generator quasi-sinewave through a three-stage RC filter.

Table 3 – peak vs RMS voltages

Waveform	RMS formula	Peak for 1V RMS	Peak for 2V RMS
Square/pulse	$V_{RMS} = V_{PEAK}$	1V	2V
Triangle/ramp	$V_{RMS} = V_{PEAK} \div \sqrt{3}$	1.73V	3.46V
Sine	$V_{RMS} = V_{PEAK} \div \sqrt{2}$	1.41V	2.83V



Scope 1: a 30kHz pulse with a duty cycle of 2%, from setting "Waveform" to square/pulse and the "Symmetry" control fully anti-clockwise.



Scope 2: a 30kHz ramp, from setting Waveform to triangle/ramp and Symmetry control fully anti-clockwise.



Scope 3: a pure sinewave (yellow) with the generator's output overlaid (mauve) at 1kHz. The total harmonic distortion (THD) is around 1% if it's properly adjusted. There is a slight phase shift between the two waveforms.

There is a frequency shift, up to 1-2%, as the symmetry/duty cycle is varied between 5% and 95%. This appears to be a peculiarity of the potentiometers; in particular, carbon-element potentiometers. Cermet pots have much less shift, but they are considerably more expensive. A likely additional contributor is the non-ideal characteristics of the schottky diodes. PE



Circuit Surgery

Regular clinic by Ian Bell

Measuring the frequency response of a circuit using a PC sound card, part 5: op amp frequency response & stability

In recent articles, we have been looking at measuring frequency responses. This was motivated by the need to measure the response of a practical implementation of an example digital filter (from our series on DSP) without advanced test equipment.

We covered the general principles of frequency response measurement and the use of a "sound card" (audio interface) in a PC or laptop together with a free app called REW (Room EQ Wizard).

Recently, we have been discussing the signal conditioning that might be needed to interface the circuit under test to the sound card's line inputs and outputs.

These circuits contain operational amplifier ICs (op amps), so this month we will take a detailed look at op amp frequency response and its impact on amplifier performance, particularly the possible problem of instability.

The interface circuit can be connected to a wide range of loads, and there is potential for some loads to cause problems. This month, we will consider op amps in general, so our discussion will have wider application than just to the sound card interface.

Feedback structure

Fig.1 shows the structure of a generic feedback system. This is not a circuit diagram – it represents the relationship between signals, rather than showing specific components and wiring. It can apply to both inverting and non-inverting op amps, but there are differences in some details. We can use this diagram to help define and analyse key concepts relating to op amp feedback.

The signals in Fig.1 could be either voltages or currents, but we have labelled them as voltages. The input signal (v_{inp}) may pass through an input network, which multiplies the input by a factor N. For the non-inverting op amp (Fig.2), the input goes straight to the op amp, so N=1. Fig.1 is labelled on the assumption that N=1, as this is the simplest situation, but it is not always the case.

For the inverting op amp (Fig.3), the resistors form a potential divider for the input signal, so the input network is not unity gain.

After the input network, the signal passes to the mixing network, the output of which is the feedback signal (v_{fb}) subtracted from the input (v_{inp}) to give the input to the amplifier block (v_{ai}). This function is performed by the op amp, so it is not a separate component in the implementation.

For the non-inverting amplifier, the op amp's differential amplifier operation achieves the subtraction. For the inverting amplifier, the 'summing junction' shown in Fig.3 adds signals at the inverting input (it's effectively subtraction because the feedback is from the output that is obtained via the inverting input).

The amplifier block represents the open-loop gain of the op amp (A_O), that is, the gain of the op amp without any feedback. The op amp amplifies the difference between its two inputs by A_O. A_O is typically in the order of 10^5 to 10^6 times (100dB to 120dB), although some devices have much higher gains, and others lower, related to requirements in specific applications.

It is not practical to use this high gain directly as a linear amplifier, but it provides the potential for very good performance, and straightforward circuit design, when negative feedback is applied.

The output of the amplifier block is $v_O = A_O v_{ai}$ for the non-inverting amplifier and $v_O = -A_O v_{ai}$ for the inverting amplifier. The feedback signal (v_{fb}) is obtained from the output of the amplifier (v_O) by passing it through the feedback network (a resistor-based potential divider in the op amp circuit), which multiplies it by a factor of β, so $v_{fb} = \beta v_O = \pm\beta A_O v_{ai}$.

β is known as the feedback factor, and typically $\beta \leq 1$. The feedback is either added to (inverting amplifier) or subtracted from (non-inverting amplifier) the input signal by the mixing block. The overall effect, taking into account the sign of A_O, is **subtraction** – hence **negative** feedback.

There are three different gain values that can be used when discussing an amplifier with feedback. These relate to the three paths in the block diagram shown in Fig.4:

1) The amplifier block on its own has an open-loop gain of $A_O = \pm v_O \div v_{ai}$

2) The whole circuit (amplifier with feedback) has a closed-loop gain of $A_C = \pm v_O \div v_{inp}$

3) $-\beta A_O$ is the 'loop gain', the gain around the closed feedback loop

Gain calculations

Referring to Fig.1 (with N=1 for the non-inverting amplifier), we can calculate the closed-loop gain in terms of

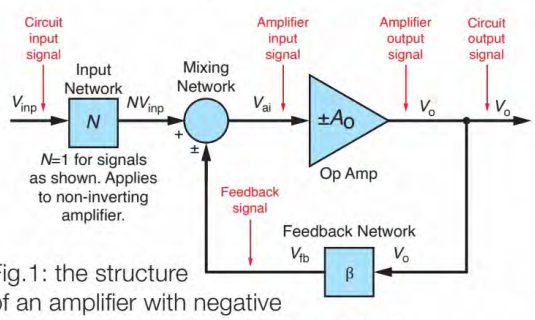


Fig.1: the structure of an amplifier with negative feedback (shown here as non-inverting).

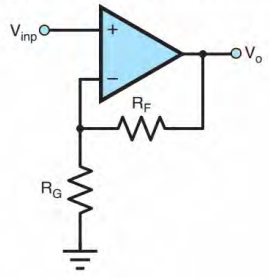


Fig.2: a non-inverting op amp based amplifier circuit.

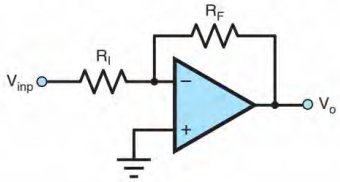


Fig.3: an inverting op amp based amplifier circuit.

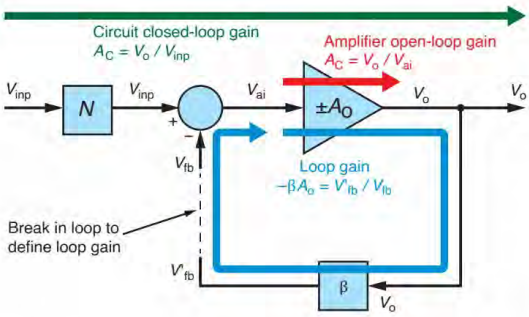


Fig.4: the relationship between open-loop gain, closed-loop gain and loop gain, for N=1 as in Fig.1.

the open-loop gain and feedback factor. Subtracting the feedback signal (βv_O) from the circuit input gives the amplifier input as $v_{ai} = v_{inp} - \beta v_O$. The amplifier (and circuit) output is the amplifier input multiplied by the amplifier open-loop gain, $v_O = A_o v_{ai}$.

Substituting in v_{ai} from the previous equation, collecting output terms and rearranging to find the closed-loop gain, we get:

$$A_C = \frac{v_o}{v_{inp}} = \frac{A_O}{(1+\beta A_O)}$$

In a typical op amp circuit, A_O is very large and β is a moderate fraction, so βA_O is usually much larger than 1 and we can approximate $(1 + \beta A_O)$ to βA_O. Then the A_O terms cancel to give:

$$A_C = \frac{1}{\beta}$$

This means that the gain of the op amp circuit is set by the feedback resistor values (which determine β), independent of the op amp's open-loop gain – as long as βA_O is much larger than one (notated as $\beta A_O \gg 1$).

This is important because the use of negative feedback means that manufacturing variations in individual op amp gains have minimal impact on the circuit, and we can swap between different op amp types without changing the gain-setting calculations.

Negative feedback to a very high-gain amplifier has other advantages, including a reduction in distortion effects due to non-linearities, obtaining constant gain over a large bandwidth, and improvements in input and output impedances.

For both amplifiers (Figs.2 & 3), the feedback fraction is determined by the potential divider formed by the two resistors. For the non-inverting amplifier in Fig.2, we have:

$$\beta = \frac{R_G}{R_G + R_F}$$

This leads to $A_C = 1 \div \beta = 1 + (R_F \div R_G)$, the gain formula for this circuit.

I expect people have often first used the well-known op amp formulae such as $A_C = -R_F \div R_G$ without wondering why there is nothing related to the op amp itself in the calculations. Of course, we can't be too complacent – the op amp does play its part; in particular, variation of the open-loop gain with frequency means that $\beta A_O \gg 1$ does not always apply.

Inverting amplifiers

In the preceding discussion, we used N=1 in Figs.1 & 4, which applies to the non-inverting amplifier. For the inverting amplifier, the input block does not have unity gain. The circuit input passes through the potential divider formed by the two resistors.

It may not be obvious that we have a potential divider in this circuit because we don't have a simple case where one of the resistors is grounded (as R_G is in the non-inverting amplifier). However, when analysing linear circuits with multiple voltage sources, we can treat each source separately and then add up their contributions to the voltage on the node of interest.

To calculate the contribution of one source, we set all the others to zero (this is the superposition theorem in linear circuit theory). So, in the inverting amplifier (Fig.3), the contribution of v_{OUT} to the total voltage at the summing junction is found by setting $v_{inp} = 0$.

Now we have a potential divider in which R_G is the lower resistor. From this, we get effectively the same formula for β as the non-inverting amplifier, but with R_I in place of R_G (because we used different resistor names).

The input voltage also makes a contribution to the total voltage at the summing junction. This time, we set $v_{OUT} = 0$, which gives us a potential divider with the resistors swapped with respect to the potential divider seen by the output (R_F is now the lower/grounded resistor). Thus, the contribution to the summing junction voltage is $(1 - \beta)v_{inp}$.

If we swap the resistors in a potential divider, the sum of the output fractions for the two cases must add up to one.

It follows that for the inverting amplifier, $N = 1 - \beta$ in Figs.1 & 4. Calculating the closed-loop gain for this amplifier using the block diagram gives us both the full and approximate (for $\beta A_O \gg 1$) expressions for closed-loop gain as:

$$A_C = \frac{-(1-\beta)A_O}{(1+\beta A_O)} \approx -\frac{(1-\beta)}{\beta} = 1 - \frac{1}{\beta}$$

≈ means "approximately equal to". Plugging the potential divider formula for β into the approximate version and rearranging gives the inverting amplifier gain formula, $A_C = -R_F \div R_I$.

Loop gain

We can define loop gain by breaking the loop, as shown in Fig.4. This allows us to obtain two different signals, v_{fb} and v'_{fb}, which are really at the same point in the loop, giving a gain for the complete opened loop path as $v'_{fb} \div v_{fb}$. We trace the entire path of the loop in the system diagram from and back to the break, so it does not matter where the break is.

For the system diagram in Fig.4, the loop contains three elements. For the non-inverting amplifier, these are gain (A_O), feedback (β) and the subtraction in the mixer block (a gain of -1). The loop gain for this system is these multiplied together, giving $-\beta A_O$. For the inverting amplifier, we have $-A_O$, β and addition in the mixer block (unity gain). Again, the loop gain is $-\beta A_O$.

Here, the minus sign shows we are subtracting the feedback from the input (negative feedback). We have not yet considered any delay in the feedback signal.

Phase shift

The output of a circuit does not respond infinitely quickly to changes at its input, so any signal fed back from the output to the input will be delayed in time with respect to the original input.

Phase shift represents the delay of a sinewave signal through a circuit at a given frequency. If the delay is constant, the phase shift will increase linearly with frequency, but often things are not this simple, particularly over wide frequency ranges.

Consider the total phase shift around the feedback loop (in Figs.1 & 4), assuming a constant delay for simplicity. As the loop gain is negative, this is equivalent to a phase shift of 180° at sufficiently low frequencies. "Equivalent to" because negative gain (signal inversion) is not an actual phase shift as it is not due to delay.

At "sufficiently low" frequencies means where the cycle time of the waveform is very long compared with the delay, so the phase shift due to any delay is close to zero at these frequencies.

Pure delay causes signals to lag the input, which is a negative phase shift. The 180° equivalent phase shift of signal inversion could be taken as either positive or negative, because shifting a steady-state sinewave by +180° and -180° gives the same result.

By convention, intentional signal inversion (eg, by an inverting amplifier) is regarded as a positive phase shift because this makes it more distinct from delay-induced phase lag. LTspice and other tools plot inversion as +180° on graphs, such as Bode plot frequency responses.

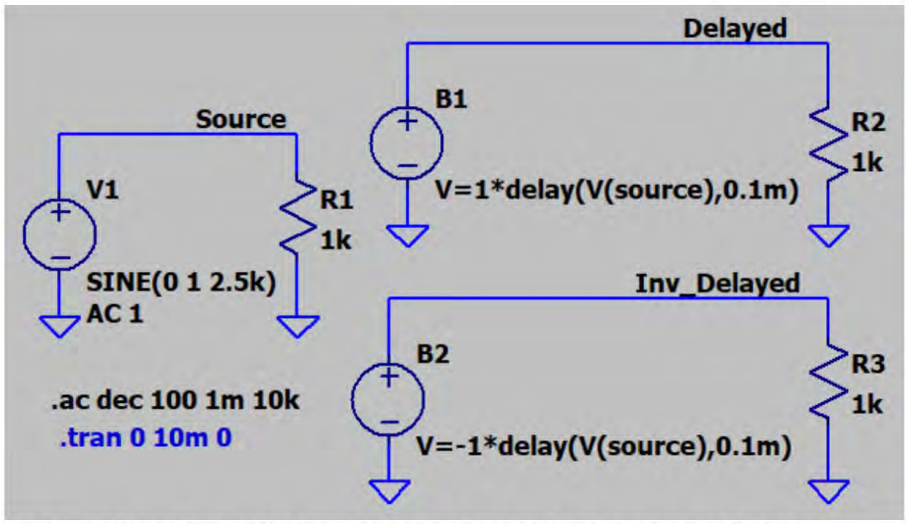


Fig.5: an LTspice simulation for illustrating phase shift for inversion and delay.

Fig.5 shows an LTspice schematic to illustrate the ideas discussed above and help interpret the results for op amp circuits discussed shortly. This is an abstract circuit in which a signal is delayed by a fixed time using the delay function of a behavioural source.

The function `delay(v(n),t)` delays the voltage on node `n` by time `t`. In the schematic, there are two delays of 0.1ms, both fed from the signal `Source`. Signal `Delayed` is just delayed, while signal `Inv_Delayed` is inverted by multiplying the delay function by -1.

Fig.6 shows the phase for the signals in the circuit in Fig.5 up to 10kHz. The `Source` signal (green trace) has zero phase shift for all frequencies because it is directly from the voltage source (V1) used for AC analysis – it is the reference from which phase shift of other signals in measured.

This plot uses a linear frequency axis, rather than the more usual logarithmic one, as this allows us to see the linear increase in phase shift with frequency produced by the ideal constant delay.

The phase shift of the `Delayed` signal (red trace) is zero at very low frequencies and becomes increasingly negative as frequency increases. The phase shift is straightforward to calculate due to the constant delay used here. For example, at 2.5kHz, the cycle time is 1 ÷ 2500 = 0.4ms, so the 0.1ms delay imposed by source B1 is one quarter of the cycle time, or -90° (360° ÷ 4).

The phase shift is negative for delay from the definition of phase shift.

The `Inv_Delayed` signal (cyan trace) follows the same pattern of phase shift, linearly getting more negative as frequency increases – this is because it is subject to the same delay as the `Delayed` signal. However, the `Inv_Delayed` signal starts with a phase shift of +180° at very low frequencies due to the signal inversion.

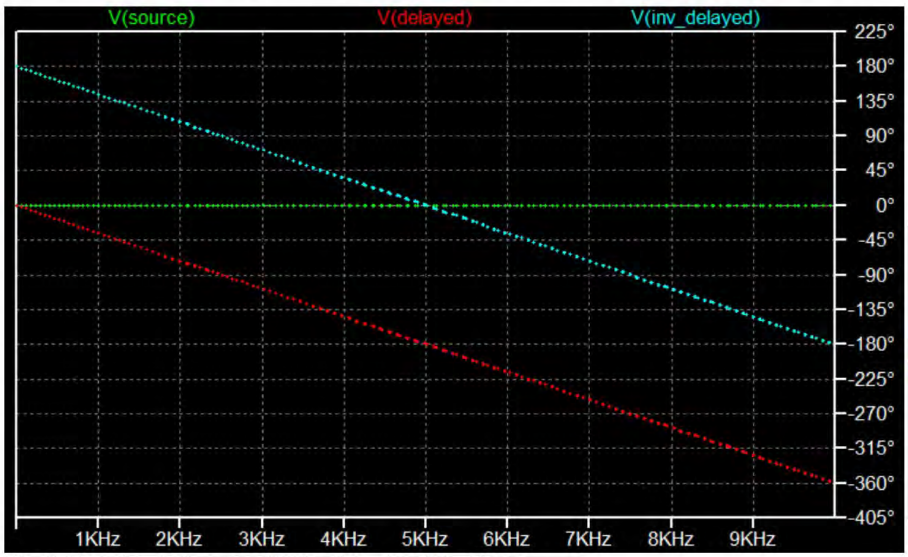


Fig.6: the simulation results for the circuit in Fig.5.

Simulating open-loop gain

It is straightforward to simulate closed-loop gain in LTspice – it is just the normal simulation of the circuit. However, both measuring and simulating open-loop frequency response is more challenging because the DC operating point of the op amp may not be calculated correctly if the op amp is connected in a simple open-loop configuration, such as shown in Fig.7.

Wired open-loop, the op amp output may (and almost certainly will) saturate at one of the supplies, so the bias conditions are not correct for amplifier operation. In terms of simulation, the DC operating point analysis runs prior to an AC analysis may result in parameters used in the AC simulation not being correctly configured to show the true open-loop behaviour.

This may not be obvious in the AC analysis results, because it is likely to run OK and produce a graph. However, the results of such an AC analysis could differ significantly from the datasheet graphs.

Fig.8 shows a schematic for simulating open-loop gain with the OPA2134 op amp (U1) used in the sound card interface circuit. A closed-loop circuit is also included (U2). The trick here is in the use of separate DC and AC analysis values for the resistors.

You can add `ac=value` to a resistor's value as shown. The first (normal resistor)

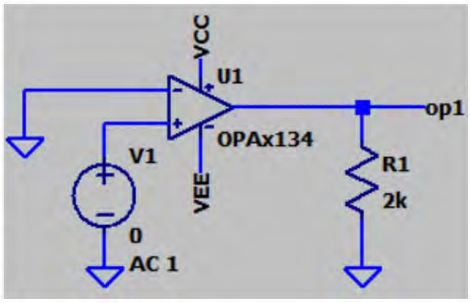


Fig.7: a simple open-loop op amp circuit is likely to give poor simulation results.

5-year collections
2019-2023
£49.95
2018-2022
£49.95
2017-2021
£49.95
2016-2020
£44.95
Purchase and download at:
www.electronpublishing.com

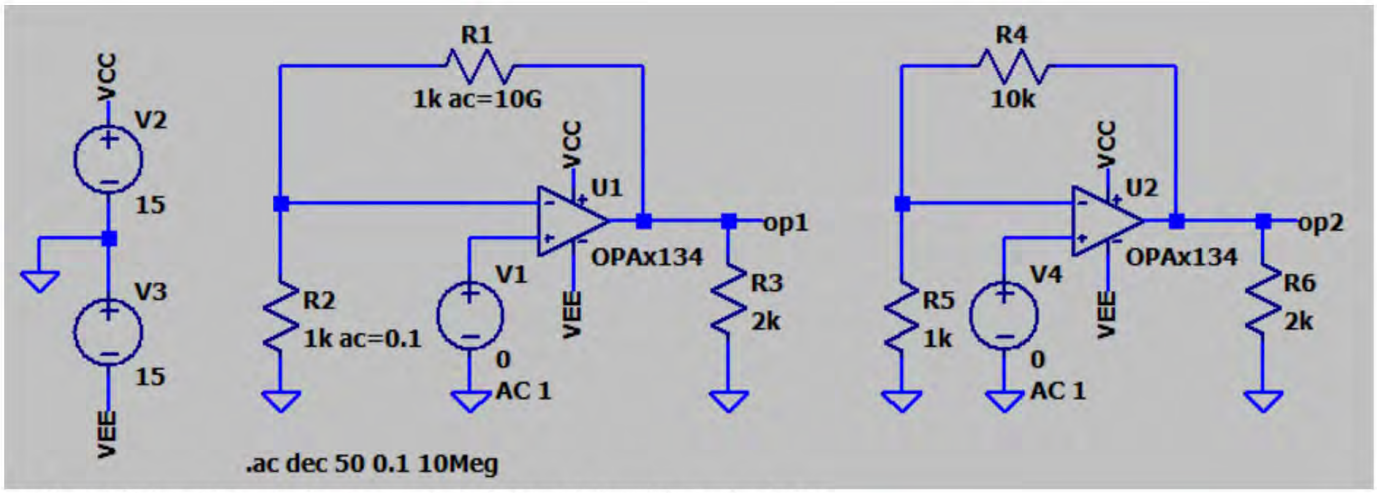


Fig.8: an LTspice schematic for op amp open-loop and closed-loop simulation.

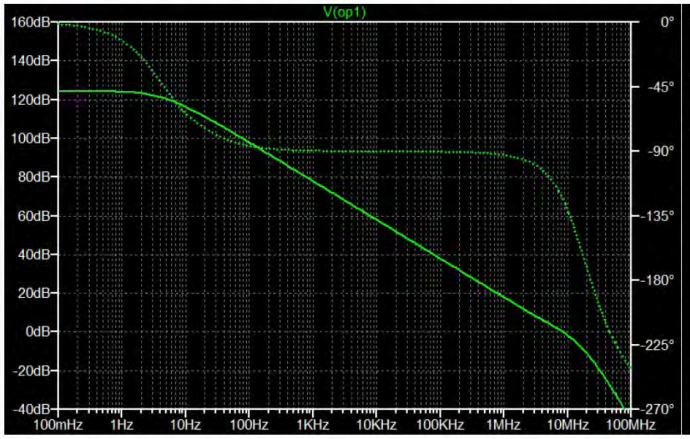


Fig.9: the results from simulating Fig.8, giving the open-loop response of the OPAx134.

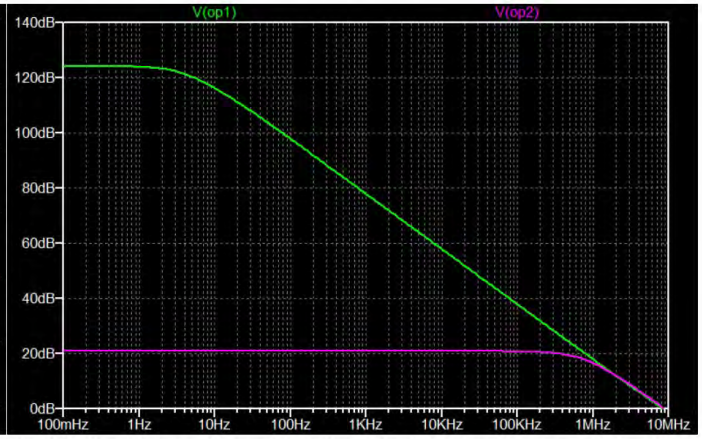


Fig.10: the results from simulating Fig.8, comparing open-loop and closed-loop gain.

value will be used in the DC operating point calculation, and the specified AC value is used in the AC analysis that plots the frequency response.

In Fig.8, the U1 circuit is configured as a non-inverting amplifier with two 1kΩ resistors, to be used in the DC operating point calculation. The specific values do not matter too much; we just use typical circuit values.

The key thing is that the AC resistor values that set the formula circuit gain value ($1 + (R_1 \div R_2) = 1 \times 10^8$ as shown) are much larger than the op amp's DC open-loop gain (about 125dB (1.8×10^6) for this op amp). Therefore, the resistor ratio does not set the circuit gain – the op amp cannot provide this gain, and what we see is the maximum op amp gain, which is the open-loop response.

To achieve this, the resistor ratio should be ten or a hundred times the open-loop gain, but excessively large and small resistor values or excessive ratios may not work due to numerical limitations in the simulator calculations.

The results from Fig.8 are shown in Fig.9. This plot uses slightly extended axes compared to the open-loop response graph in the OPAx134 data sheet but shows a good match to the shape of the curves.

Fig.10 shows the open-loop gain of the OPAx134 together with the closed-loop gain of a ×10 (20dB) non-inverting amplifier using the same op amp (U2 in Fig.9). The open-loop gain varies significantly with frequency but remains much larger than the closed-loop gain until the frequency reaches hundreds of kilohertz.

For low frequencies, the approximation $\beta A_O \gg 1$ discussed above is valid, and the gain is equal to $1 \div \beta = (R_4 + R_5) \div R_5 = 10$ (20dB) from the potential divider formula. At sufficiently high frequencies, the approximation is not valid, and the closed-loop gain decreases with the open-loop gain, following the $A_O \div (1 + \beta A_O)$ formula discussed above.

Poles and zeros

We discussed poles and zeros in depth last month. These terms come from the analysis of a circuit's frequency response in the Laplace domain.

Although this is based on advanced mathematics, we can use the concept of poles and zeros in a practical sense to describe frequency responses and help decide how they might need to be manipulated to achieve our circuit design aims.

As explained last month, poles and zeros produce breakpoints in a circuit's frequency response where the rate-of-change of gain and phase-shift change. At frequencies above a pole frequency, gain decreases at a rate of 20dB per decade more compared to below it, and phase shift becomes 90° more negative – assuming there are no other poles or zeros at similar frequencies.

Above a zero, the gain decreases at a rate of 20dB per decade less and the phase shift becomes 90° more positive.

In the gain plot in Fig.9, a couple of breakpoints are visible. There is a pole at about 5Hz; the gain is flat below the pole, at very low frequencies, and eventually falls by 20dB per decade, above the first pole, until frequencies of several megahertz. There is an increase in phase shift from 0 to -90° due to this pole.

There is another breakpoint at around 20MHz. There are two poles here at the same or similar frequencies because the gain eventually drops at 60dB per decade (40dB more than in the mid-range) and the phase shift increases by -180° (2 × -90° for two poles) to -270° (3 × -90° for three poles total in Fig.9).

Simulating loop gain

To plot the loop gain of an op amp in a simulation, we need to somehow break into the loop to inject a signal. This can be done with a voltage source having a DC voltage of 0, as shown in Fig.11. The

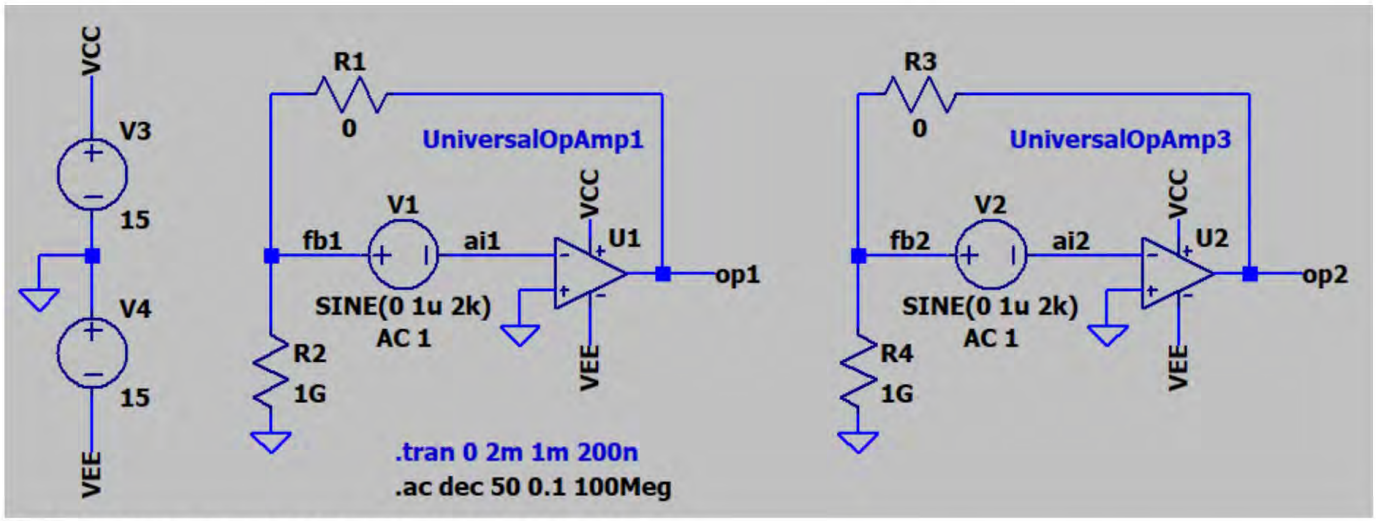


Fig.11: an LTspice schematic for op amp loop gain simulation.

two sides of the voltage source will be at the same DC voltage, so are effectively shorted together at DC.

Inserting the voltage source provides separate nodes in the simulation model that relate to v_{fb} and v_{ai} in Figs.1 & 4, and are labelled correspondingly in Fig.11. We can therefore plot $v_{fb} \div v_{ai}$ in an AC analysis to show the loop gain of the amplifier.

With the presence of feedback resistors (this is a version of a normal amplifier circuit, not open-loop as discussed above), the op amp will be biased normally. The circuit in Fig.11 is not specifically either an inverting or non-inverting amplifier. It could be either, because the feedback loop is the same in both cases.

In this simulation, we are just looking at the loop gain, so the normal input signal (which would either be to the inverting input or to resistors R_2 and R_4) is zero.

The two op amps in Fig.11 are the UniversalOpAmp1 and UniversalOpAmp3 idealised behavioural models provided by LTspice. The feedback resistors R1 and R3 are set to zero, and the grounded resistors to a very large value (1GΩ). This feeds back the whole output signal ($\beta = 1$), so the loop gain (βA_O) of this circuit will be equal to the open-loop gain, at least with a relatively ideal circuit.

Using resistors like this easily allows us to change the circuit gain without having to redraw it.

Fig.12 shows the results of the AC analysis from Fig.11 and highlights the difference between the UniversalOpAmp1 (U1, green trace) and UniversalOpAmp3 (U2, red trace) models, specifically that UniversalOpAmp1 has one pole and UniversalOpAmp3 has two poles.

The single-pole model provides a simple idealised op amp, whereas the two-pole version is closer to real devices (note the similar high-frequency poles in Fig.9), but may not cover their full complexity.

The specific gain and pole frequencies are set by right-clicking the op amp and adjusting the model parameters. We will look at the details of this later. The op amps were configured to give similar response to the OPAx134, within the limits of the models.

The plot in Fig.12 shows that phase shift is +180° at low frequencies, which fits with the preceding discussion on inversion and phase shift.

Fig.13 shows the results of a transient simulation using a 1μV, 2kHz signal from the V1 voltage source in the loop. This feeds a signal directly into the op amp. From Fig.12, we expect the loop gain to be about +72dB (4000 times).

We expect the phase shift to be -90° because Fig.12 shows +90° at 2kHz, which comprises the +180° equivalent shift from the inversion and -90° phase shift from the op amp's first pole.

Fig.13 shows the feedback signal (v_{fb1}, red trace) multiplied by -1 to remove the inversion so that we can more easily see the phase shift due to delay round the loop. The input signal (v_{ai1}, green trace) is multiplied by 4000 to scale it to the same level as the feedback to facilitate comparison and confirm the loop gain value.

The cursors show a time difference of 125μs, which is a quarter of the cycle time of 1 ÷ 2000 = 500μs, representing a 90° phase shift. The feedback signal is lagging (delayed), so the phase shift is -90°.

Fig.14 shows two more versions of the circuits from Fig.11, both using UniversalOpAmp3 (U3 and U4). In both cases, β=0.1; this is set by the potential divider formed by the resistors, for example $R_6 \div (R_5 + R_6)$ = 1kΩ ÷ (1kΩ + 9kΩ) = 0.1. The circuits would have a gain of 10 if used as an inverting amplifier, or -9 in the inverting configuration.

Fig.15 shows the results from the circuit using U3 (cyan trace) compared with the response of U2 from Fig.11 (red trace). The only difference is the β value (1 for the U2 circuit and 0.1 for the U3 circuit).



Fig.12: the results from simulating Fig.11, showing the loop gain of LTspice's UniversalOpAmps.

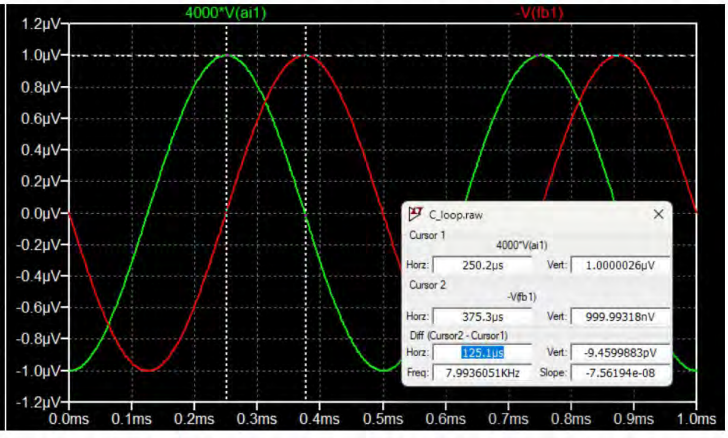


Fig.13: the results from simulating Fig.11,showing phase shift in the feedback loop.

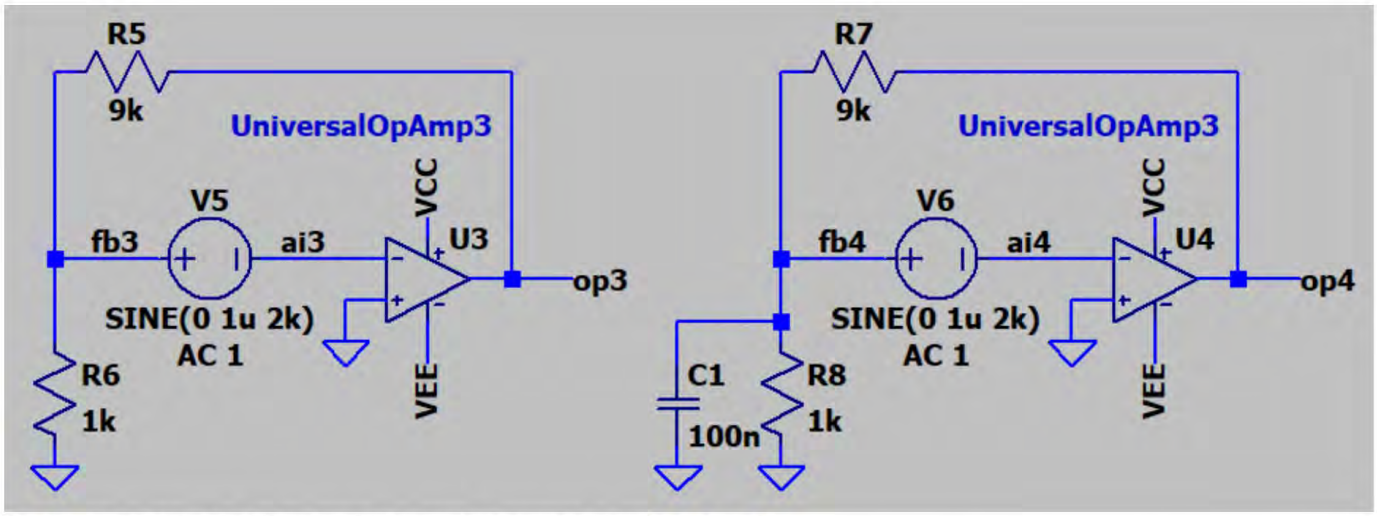


Fig.14: more loop simulation examples, part of same schematic as Fig.11.

The loop gain for the U3 circuit is 10 times smaller than for U2 circuit, so its gain curve is exactly 20dB [$20\log_{10}(0.1) = -20$] below the curve for the U2 circuit. Both circuits have the same phase response.

Fig.16 shows the results from the U4 circuit in Fig.14 (yellow trace). This has an extra capacitor compared to the circuit for U3 (also plotted, cyan trace). The capacitor is frequency-dependent, so it changes the shape of the frequency response, rather than just shifting on the dB axis as different resistor values do.

It also changes the phase response significantly, adding -90° compared with the op amp alone. Adding this capacitor introduces an additional pole in the feedback loop. The value of the capacitor was chosen to place this pole in the middle of the graph, so it was clear to see.

Instability

Consider the total phase shift around the feedback loop (in Figs.1 & 4). We have +180° equivalent phase shift at DC due to the signal inversion, which is required to achieve negative feedback.

As we increase the signal frequency, the total phase shift will tend to become more negative in a similar way to what we have seen in loop simulation examples above. However, for real op amps, the shape of the curves may be more complex.

Once the phase shift due to the loop delay reaches -180°, the total effective phase shift around the loop is 0°. At this point, the -180° delay phase-shift cancels the +180° effective phase-shift from the inversion. We effectively multiply the loop gain by -1, giving $+\beta A_O$. What was negative feedback has become positive feedback.

Positive and negative feedback produce profoundly different behaviours in circuits, so it is possible that the presence of positive feedback at high frequencies will have a significant impact.

Positive feedback at DC causes bistable operation, as observed in circuits such as Schmitt triggers and latches, where the output tends to saturate at one of two values, often at or close to the supply rails. Such circuits are unstable with intermediate DC voltages, but stable at voltages close to the supply rails.

An op amp linear amplifier will typically not have strong (or any) positive feedback at DC and low frequencies, so it will probably not exhibit unstable behaviour related to DC (eg, the output stuck at a supply voltage). Instead, positive feedback of AC signals will add to the signal level of any existing signal at that frequency.

If the gain around the loop is sufficient, this could potentially cause the signal amplitude at that frequency to continuously increase until it is limited by circuit nonlinearities (the op amp's maximum output level is reached).

The result is oscillation at a frequency where the positive feedback is strong. The positive feedback reinforces the signal and so sustains the oscillation even if the input at the relevant frequency is removed.

Typically, there will be some initiating signal at the relevant frequency that is amplified and built up by the feedback to sustain the oscillation. This can be due to noise in the circuit or disturbances; for example, from the power supply.

Looking at the denominator of the closed-loop gain equations above ($1 + \beta A_O$) gives further insight into this. If the denominator is zero, the gain will be infinite in an idealised circuit.

The gain is infinite for $\beta A_O = -1$; remembering that the loop gain (βA_O) has both phase and magnitude, the conditions for this are that the magnitude of βA_O is unity ($|\beta A_O| = 1$) and the total effective phase shift (including the +180° inversion) is zero or an integer multiple of 360°. That is positive feedback with enough gain for it to be possible to sustain oscillation.

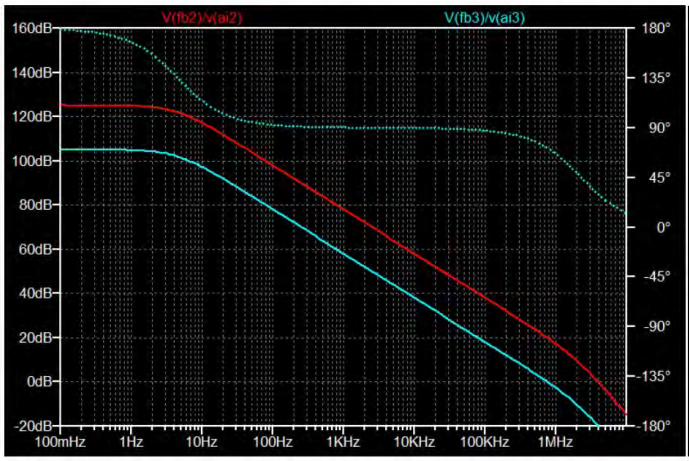


Fig.15: the results from simulating Figs.11 & 14, giving loop responses (βA_O) for $\beta = 1$ (U2, red) and $\beta = 0.1$ (U3, cyan).

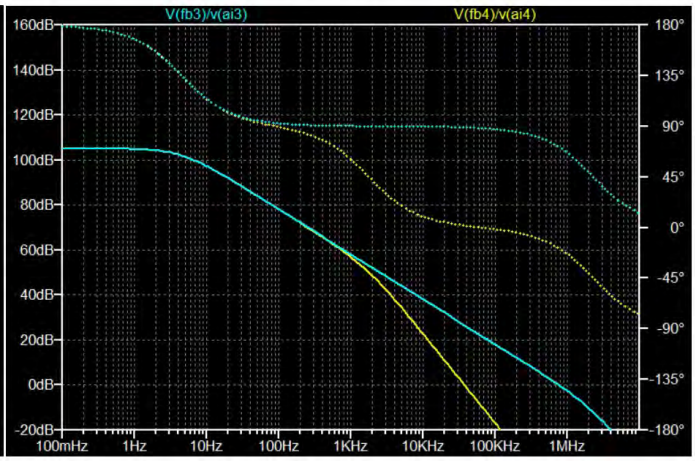


Fig.16: results from Fig 14: loop responses (βA_O) without (U3, cyan) and with (U4, yellow) a capacitor across the grounded resistor.

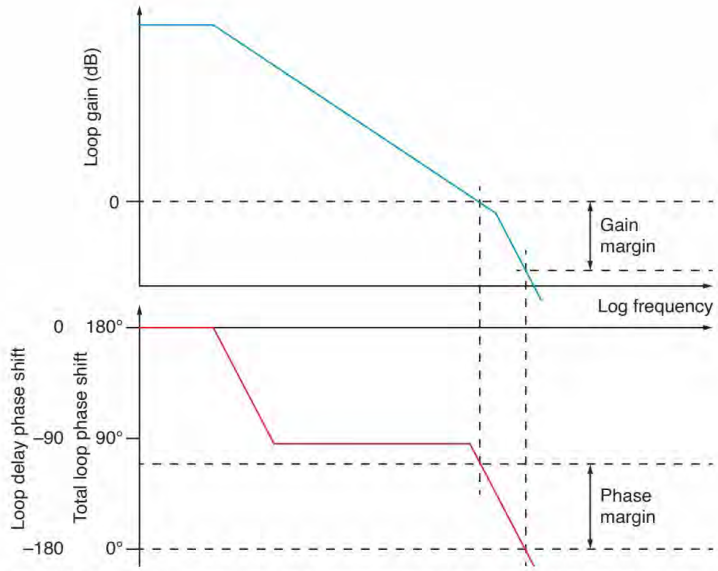


Fig.17: how phase margin and gain margin relate to the open-loop response of an op amp.

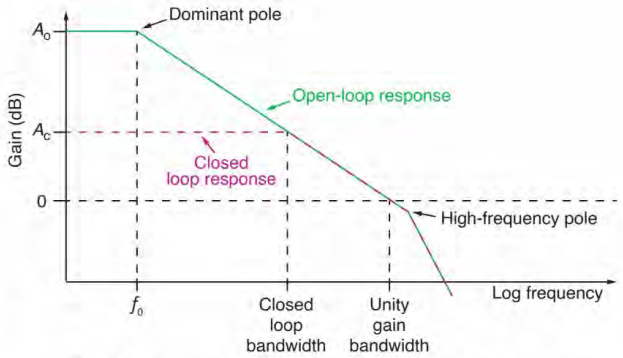


Fig.18: the relationship between open-loop and closed-loop frequency responses.

These two conditions for oscillation are referred to as Barkhausen's criteria, which is commonly used in discussions of amplifier stability, however, general conditions for stability, or conversely, sustained oscillation are more complex, and Barkhausen's criteria are not sufficient to guarantee oscillation.

In unstable amplifiers, a small (eg, noise) signal will initiate an oscillation with increasing amplitude due to high loop gain. If sustained oscillations occur, nonlinearities will reduce the gain to settle at Barkhausen's criterion, $|\beta A_O| = 1$.

If these (or other) conditions for sustained oscillation are not met, a circuit may produce decaying oscillations (severe ringing) after a trigger event, such as an input transient. Such circuits may not be strictly 'unstable' in the sense of oscillating continuously, but they will often be described as such if the amount of ringing is unacceptable.

Compensation

Lower closed-loop gain corresponds with higher β (β closer to or equal to one) and hence higher loop gain. The higher loop gain means a low-gain op amp circuit is more likely to be unstable because it is easier to meet the criteria for oscillation. This may seem counterintuitive if we think of something with high gain as being 'unstable' because it is very sensitive to its input signals etc.

Adding poles and zeros to a circuit's feedback loop will change the response (see the two responses in Fig.16) and hence the stability of the circuit. For example, adding a pole may make a circuit more unstable because it causes the phase shift to increase; however, it also may make the circuit more stable because the gain is reduced at high frequencies.

The actual situation will depend on the relationship of all the poles and zeros in the frequency response.

A key example of adding a pole is found in the open-loop response of op amps – for example, in Fig.9 we see a low-frequency pole in the op amp's frequency response. This is deliberately part of the design of most op amps, with the pole being typically in the low hertz range. It is chosen so that the op amp's high-frequency gain is sufficiently low for it to be stable with 100% feedback (unity gain configuration).

Modifying an amplifier's open-loop response to ensure stability with feedback applied is called compensation. Adding a low-frequency pole to achieve this is called dominant-pole compensation, because the pole dominates the overall shape of the response. Most, but not all, op amps use this approach.

Phase margin and gain margin

Due to variations in individual components, component ageing and temperature variation, it is not a good idea to design an amplifier that is stable but close to being unstable. If so, some individual builds of the circuit, and others over time, or as environmental conditions change, will start to oscillate and fail as amplifiers.

We can measure how close a circuit is to being unstable using the concepts of gain margin and phase margin.

Figs.12, 15 & 16 show loop frequency responses for op amps. As frequency increases, the gain decreases and the phase shift becomes more negative, reducing the total phase shift from +180° at low frequencies towards zero as the phase shift due to delay moves from nearly zero towards -180°.

For the amplifier to be stable, as the loop gain magnitude approaches unity (0dB), the negative phase shift (lag) in the loop due to delay must be less than -180°, so the total including inversion is greater than zero. The difference between the delay phase shift at this point and -180° is the phase margin.

Similarly, as the phase shift due to delay around the loop approaches -180°, the magnitude of the gain must be < 1 to prevent oscillation. The difference between the loop gains when its phase shift reaches this point and unity gain is the gain margin (usually expressed in dB). Fig.17 illustrates these definitions.

It is common to regard 45° as the minimum acceptable phase margin. This is to allow sufficient headroom for component and environmental variations, but it is a rule of thumb and not universally used. Usually, the acceptable range is between 30° and 60°, but it depends on who you ask and the trade-offs they are considering.

A lower phase margin will tend to mean the circuit produces more ringing, while higher values tend to produce slower responses to fast-changing inputs.

Gain-bandwidth product

Fig.18 shows the frequency response of a typical dominant pole compensated op amp for both open-loop and closed-loop conditions (this is a generalised version of Fig.10). The op amp has full open-loop gain (A_O) at DC, but the gain starts dropping at the low-frequency dominant pole (f_0), typically from 1-10 Hz.

Above f_0, the gain falls off at 20dB/decade. As A_O is very large, it does not drop to unity (0dB) until a high frequency (the unity gain frequency, often many megahertz or more). The closed-loop gain (A_C) remains flat until it approaches the open-loop gain, as discussed above.

If we multiply closed-loop gain by closed-loop bandwidth, we get the gain-bandwidth product (GBP). A feature of dominant-pole compensated frequency responses is that the GBP is constant. The gain drops by 20dB per decade at frequencies above f_0. A drop of 20dB is a factor of 10 decrease in gain, but this is for a factor of 10 increase in frequency. Thus, the product of gain and frequency is constant.

The shape of the open-loop curve up to any high-frequency poles is set by

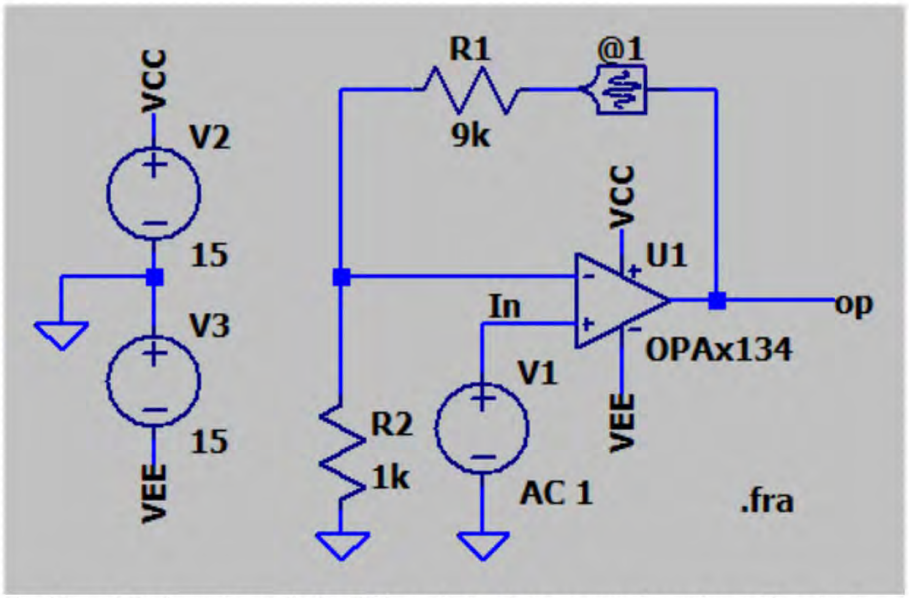


Fig.19: an LTspice schematic for FRA analysis of an op amp based amplifier circuit.

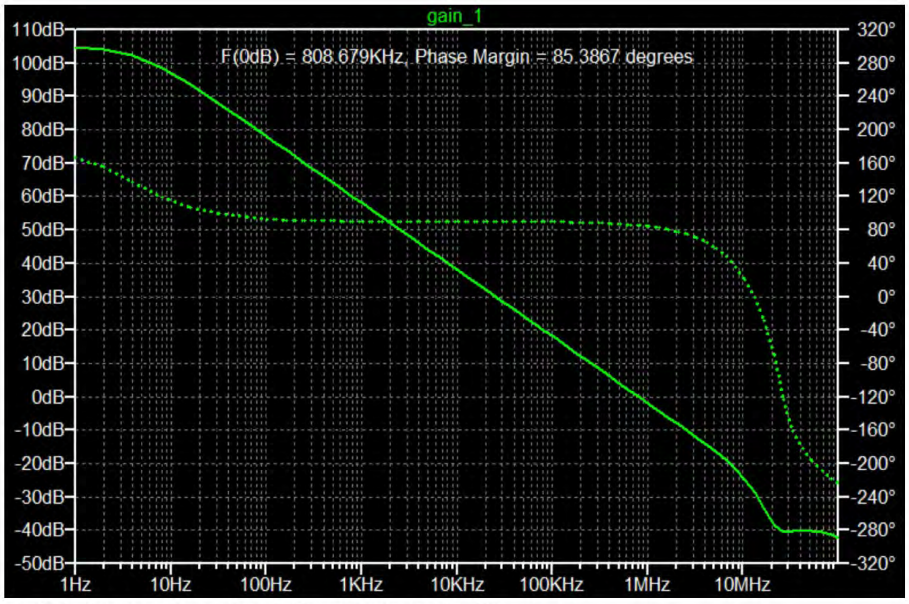


Fig.20: the results of the FRA analysis in Fig.19.

A_O and GBP. These values can be set by the user for LTspice's UniversalOpAmp devices. For example, for the circuits in Figs.11 & 14, I used $A_O = 1.778 \times 10^6$ (125dB) and GBP = 8MHz to provide a reasonable match to the OPA2134 used in the soundcard interface.

UniversalOpAmps numbered 3 and above also have a high-frequency pole, which can be set via their phase margin parameter.

Some commercial op amps are not designed to be stable in low closed-loop gain configurations (eg, for gains below 2 or 5 etc). This is done specifically to allow higher-gain closed-loop gain circuits to have higher bandwidth (the open-loop gain A_O can be higher at higher frequencies with smaller β values because the loop gain βA_O is still sufficiently small there to achieve a good gain margin).

This is a potential pitfall for those who do not read datasheets carefully enough to notice the minimum closed-loop gain specification.

LTspice FRA example

You can measure gain margin and phase margin from loop simulations such as in Figs.11 & 14, but LTspice also has an automated method to do this called a Transient Frequency Response analysis (`.fra` command), which uses a stepped-frequency sinewave.

This has the advantage that is uses the full op amp model, not the linearised version employed in AC analysis. A special FRA element (a component with a name starting with @) is inserted into the loop instead of a voltage source, as shown in Fig.19.

There are several settings required, which are made by right-clicking the FRA element. The setting for the FRA in Fig.19 were: 10 points per decade, coarse steps to 100Hz, amplitude 0.5V up to 50kHz, reducing to 0.1V by 500kHz, analysis start at 100µs, minimum time 100µs, settling time 200µs.

The results are shown in Fig.20 and include generated text stating the unity gain frequency (808.679kHz) and phase margin (85.3867°). PE

Simulation files

Most *Circuit Surgery* columns feature the use of the free circuit simulation software LTSpice. It is used to support descriptions and analysis in *Circuit Surgery*.

The examples and files for this issue are available for download: https://pemag.au/link/ac9b

HAMMOND®
1591 ABS flame-retardant enclosures
Learn more:
hammondmfg.com/1591
uksales@hammfg.com • 01256 812812



Raspberry Pi Pico 2

Review by Tim Blythman

The Raspberry Pi Pico 2 microcontroller board is similar to the original Pico (released in 2021) except that it features the new RP2350 microcontroller chip. This makes it considerably faster but it has another interesting feature: you can select between its two ARM cores or two RISC-V cores.

The Raspberry Pi Pico 2 microcontroller board was released in August 2024. We reviewed the relatively new Raspberry Pi 5 single-board computer (SBC) from the same company in the May 2025 issue.

The original Pico was released in 2021, followed by the WiFi and Bluetooth equipped Pico W in 2022. Both these boards are based on the RP2040 microcontroller, the first microcontroller designed by the Raspberry Pi Foundation. The Raspberry Pi 5 introduced the RP1 microcontroller, acting as an I/O controller.

Like the Raspberry Pi SBCs, the Pico was designed to be low cost and easy to use, with a target price of US$4 (about £3). Not long after its launch, it could be programmed in the C language, with the Arduino IDE and MicroPython; PicoMite BASIC was released soon afterwards, around December 2021.

About a year later, the Pico W was released. It shares the same form factor and processor as the Pico but includes an Infineon CYW43439 radio module, adding WiFi and Bluetooth support.

The bare RP2040 microcontroller later became available for purchase at around 50 pence, from the likes of DigiKey and Mouser. That led to its incorporation into many third-party boards.

We created the Pico BackPack, which adds features like an LCD touchscreen, microSD card socket and audio output to a Pico or Pico W. That was detailed in the March 2023 issue, with the related Pico W BackPack introduced in January 2024.

We have used the Pico and Pico W in various projects, including the WebMite (August 2024), VGA PicoMite (December 2024) and Pico Gamer (March 2025). So we were very interested to see what the Pico 2 has to offer.

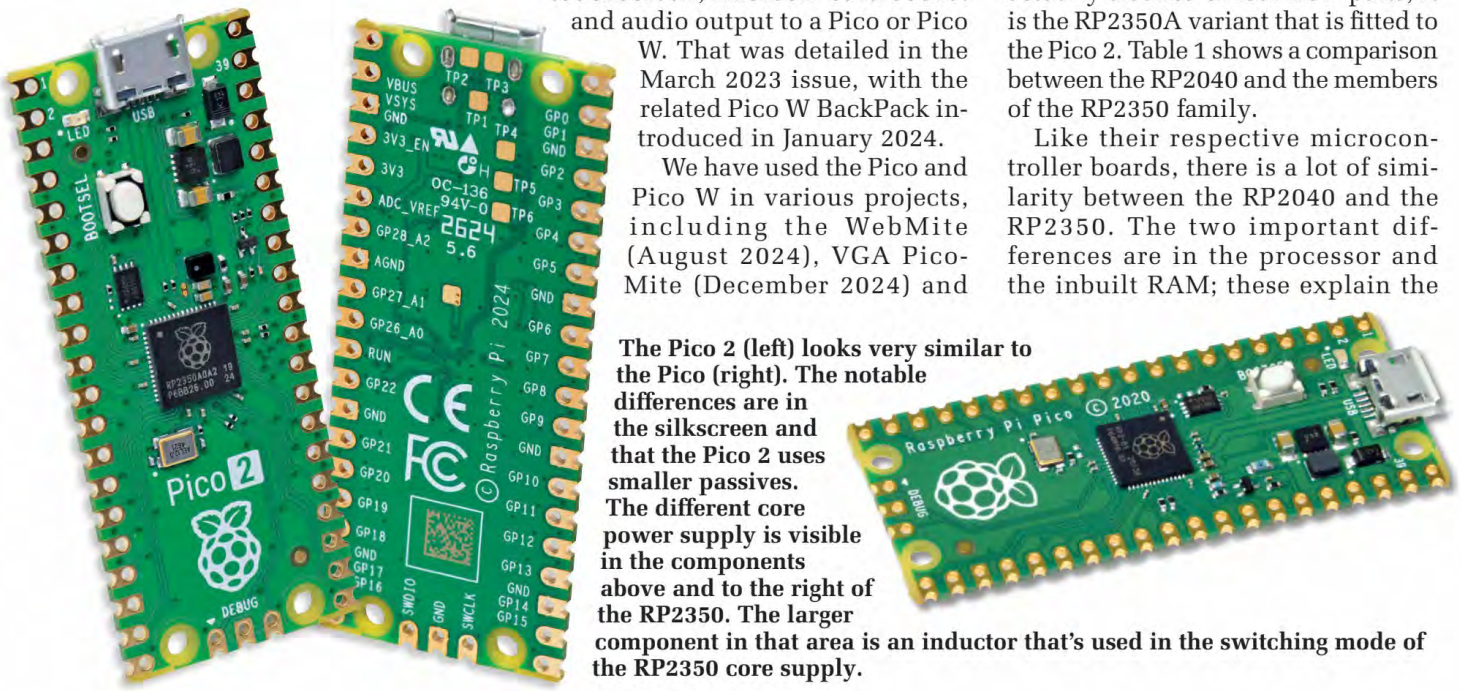
There are a lot of similarities; it has the same layout and footprint as the Pico & Pico W. Apart from the silkscreen being marked as a Pico 2, you might not even know it was a different board! It appears the Pico 2 is backwards compatible with the Pico; we shall investigate that later.

The Pico 2 was intended to be sold for US$5, and we purchased our test boards for about £4 (excluding delivery), which is much the same price at the time of writing.

The RP2350

The new RP2350 microcontroller is actually a series of four new parts; it is the RP2350A variant that is fitted to the Pico 2. Table 1 shows a comparison between the RP2040 and the members of the RP2350 family.

Like their respective microcontroller boards, there is a lot of similarity between the RP2040 and the RP2350. The two important differences are in the processor and the inbuilt RAM; these explain the



The Pico 2 (left) looks very similar to the Pico (right). The notable differences are in the silkscreen and that the Pico 2 uses smaller passives. The different core power supply is visible in the components above and to the right of the RP2350. The larger component in that area is an inductor that's used in the switching mode of the RP2350 core supply.

differences between the part numbers.

The RP2350 has a dual ARM Cortex M33 processor compared to the RP2040's dual ARM Cortex M0+ (hence the '3' in RP2350), while the '5' indicates that it has twice as much RAM (see Fig.1). Its data sheet can be found at https://pemag.au/link/ac1u

The QSPI controller (which is used to communicate with an external flash memory chip) has been provided with a second interface. This can be used to connect a second flash chip or a PSRAM (pseudo-static random access memory), to expand the memory available to the system.

8MiB (64Mbit) PSRAM chips are available for a few dollars. That is a phenomenal amount of RAM for a microcontroller, but note that the Pico 2 board does not have provision for a PSRAM chip to be fitted.

The RP2350 also has a dual Hazard3 RISC-V (pronounced 'risk five') processor that can be selected at boot time. RISC-V is an open RISC (reduced instruction set computer) architecture that is gaining traction as an alternative to other proprietary architectures. In theory, one core can be a RISC-V processor and the other, an ARM processor.

The new M33 ARM processor has native floating-point instructions that the M0+ processor in the RP2040 lacks; floating-point support for the RP2040 is provided by software routines in ROM. That means a big uplift in performance when performing floating-point calculations.

The M33 also includes Arm TrustZone and secure boot, using an OTP (one-time programmable) memory to store an encryption key. The M33 processor also performs better (at the same processor clock speed) than the M0+ in tests such as the Dhrystone benchmarks. The security features are not available when the RISC-V cores are used.

The RP2350 has a nominal maximum clock speed of 150MHz, although we have already read reports that it can be overclocked (much like the RP2040). There are reports of operation up to 300MHz. Such overclocking is also subject to the limits of the flash memory chip.

Two of the RP2350 variants boast a larger chip with more I/O pins; those have the 'B' suffix. These have 48 general-purpose I/O pins, compared to just 30 on the RP2040 and 'A' variants.

Table 1 – RP2040 and RP2350 family comparison

RP2040	RP2350A	RP2350B	RP2354A	RP2354B	
Dual ARM Cortex M0+	Dual ARM Cortex M33 and Hazard3 RISC-V				Processor (CPU)
External only			2MiB internal		Flash memory
264kiB	520kiB plus external PSRAM				RAM
133MHz	150MHz				Clock
56	60	80	60	80	Pins
30	30	48	30	48	GPIO
2					UART
2					SPI
2					I²C
16	24				PWM
4	4	8	4	8	ADC channels
Full-speed host or device					USB
8	12				PIO state machines
–	HSTX peripheral, secure boot with OTP storage, hardware random number generator				Other

Then there are the RP2354 variants, which bond a 2MiB (16Mbit) Winbond W25Q16JVWI QSPI NOR flash memory chip to the RP2350 processor die; this die is otherwise identical to a bare RP2350A or RP2350B chip. Thus, the four variants of the RP2350 are the 60-pin 'A' versions and 80-pin 'B' versions, either with (RP2354) or without (RP2350) an attached flash memory chip.

Having only four ADC (analog-to-digital converter) channels on the RP2040 saw the Pico falling short compared to many other microcontrollers' analog abilities. The larger RP2350B variants now have eight ADC channels, which means that the Pico 2 is still stuck with only four channels.

Some errors have been identified in the analog-to-digital converter (ADC) silicon hardware of the RP2040. The RP2350 data sheet indicates that those have been fixed in the newer chip.

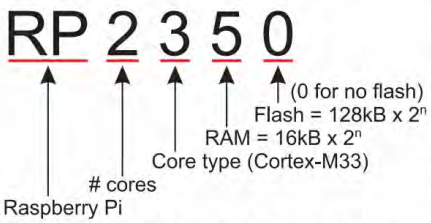


Fig.1: the part naming of the RP2350 (and RP2040) is based on this scheme. The RP2354 parts have 2MiB (2^4 × 128kB) of non-volatile storage in the form of a flash memory chip bonded to the processor die.

The novel PIO (programmable input output) peripheral saw a lot of attention, and has been put to good use in emulating all sorts of peripheral functions. That includes SPI, USB and even the protocol that is used to control WS2812 programmable LEDs.

The RP2350 provides 12 PIO state machines, up from the RP2040's eight. There are also some minor updates to the PIO peripheral itself.

The RP2350 also has a new HSTX peripheral; this stands for 'high-speed serial transmit'. It can stream data out on eight I/O pins at up to 300MHz (using double-data-rate output registers). There is example code to use the HSTX to generate DVI-compatible video.

The RP2350 data sheet notes that each processor core implements a TMDS (transition minimised differential signalling) encoding algorithm. TMDS is an encoding used with HDMI and DVI video, so clearly there is an intention for the RP2350 to be able to directly produce video output.

Power management on the RP2350 has been improved by splitting the power domains and allowing some parts to be selectively powered off,

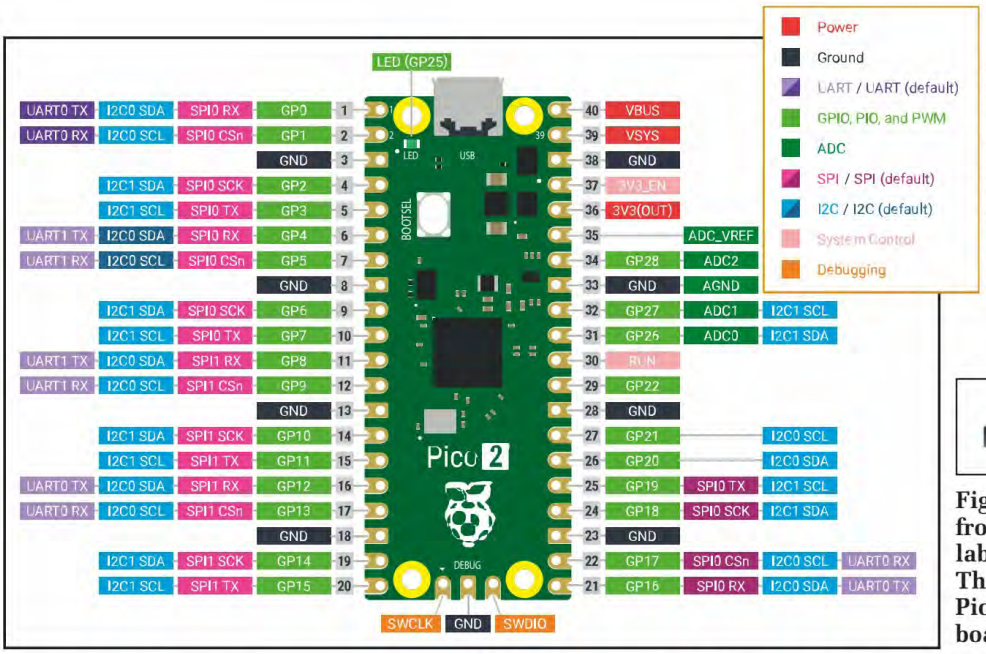


Fig.2: the Pico and Pico 2 share this pinout diagram, meaning that I/O and peripheral mappings are identical. The new HSTX peripheral is not shown; it uses the GP12-GP19 pins of the Pico 2. *Source:* *www.raspberrypi.com/documentation/microcontrollers/pico-series.html*

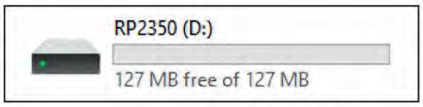


Fig.3: an easy way to tell the Pico from the Pico 2 is the drive volume label displayed by the bootloader. The RP2350 label indicates that it's a Pico 2. A Pico or other RP2040-based board would show this as RPI-RP2.

thus potentially using less power than the RP2040 in sleep mode.

The Pico 2

Unsurprisingly, the biggest difference between the Pico and Pico 2 is the new processor chip. As well as doubling the RAM, the Pico 2 has double the available flash memory, with a 4MiB (32Mbit) flash memory chip onboard. The data sheet for the Pico 2 can be downloaded from https://pemag.au/link/ac1v

That's about the extent of the changes between the two boards. The same RT6150 buck/boost regulator allows the Pico 2 to operate from anywhere between 1.8V and 5.5V. Similar to the Pico, the Pico 2 also has a diode between the VBUS and VSYS pins.

The Pico 2 appears to use smaller passive components, and there is some extra circuitry related to the RP2350's core 1.1V power supply, which has a regulator that can operate in both linear and switching modes, allowing it to achieve better efficiency.

The rear of the Pico 2 has test points in the same place as the Pico, with the addition of an extra test point in the area of the switching regulator's circuitry. Otherwise, a 2024 copyright notice is the most prominent difference.

From what we can see, there isn't even a new pinout diagram for the Pico 2; the Pico diagram has simply been annotated to include the Pico 2. So it appears that there are no electrical or mechanical reasons that rule out using a Pico 2 in place of a Pico.

Fig.2 shows the pinout. It does not note the HSTX-capable pins, presumably to retain the consistency between the Pico and Pico 2 diagrams. The HSTX pins are fixed to GPIOs 12-19.

A fault in the silicon

While the Pico 2 may appear to be better in all ways than the Pico, there was a severe erratum in the original revision of the chip that has now been resolved by the release of a new stepping. The data sheet notes this as erratum RP2350-E9, and it applies to stepping A2.

An excessive leakage current is sourced from a digital input pin if its voltage is in the undefined input voltage region, between valid high and low levels.

When connected to a high impedance source, this could result in erroneous readings. It is especially a problem if the internal pull-down is active, since the weak pull-down cannot overcome the leakage and the pin remains stuck in the undefined input voltage region (around 2.2V for a 3.3V supply).

Software fixes can be applied to some but not all situations. The general advice is to use an external pull-down resistor of no more than 8.2kΩ instead of the internal pull-down when required.

The good news is that basically all Pico 2 modules and RP2350 chips you will find on the market today are the A3 stepping or later, so they do not suffer from this leakage current problem.

The bug that caused poor ADC performance in the RP2040 has also been fixed in the RP2350.

Security

We aren't surprised that security was a low priority for the Raspberry Pi Foundation in creating a cheap and easy to use board in the Pico. The Arm TrustZone and secure boot features of the RP2350 intend to address one of the claimed weaknesses of the RP2040: a lack of security for the program flash memory.

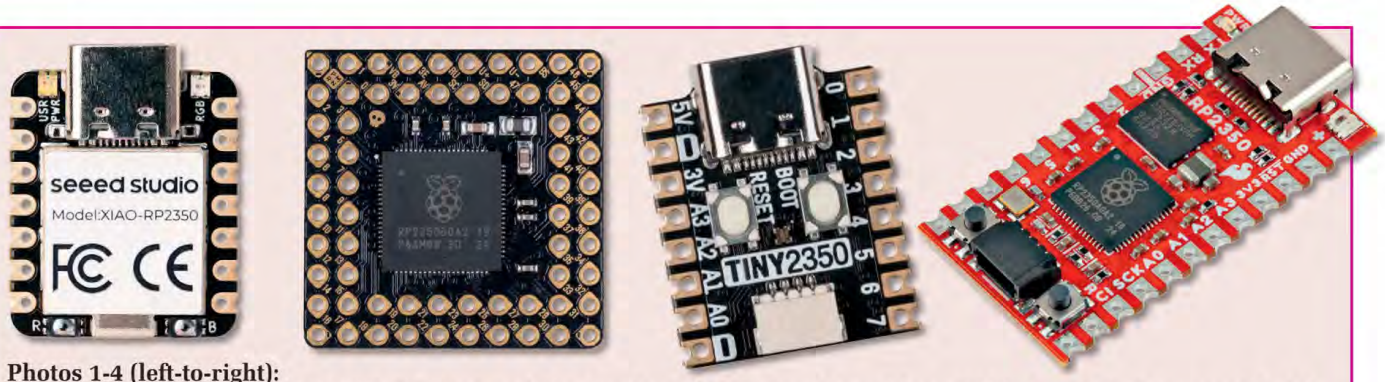
For example, reading or modifying the program in the flash chip (on the original Pico) would be as easy as accessing the flash chip and performing read or write commands.

The security on the RP2350 depends on the flash memory contents being encrypted and signed. The encryption means that the data stored on the chip is meaningless until the processor decrypts it. The signing process is a way to tell if the data has been modified, and generally involves creating a hash or checksum of the data that can indicate if it has been changed.

The signing is necessary as the encryption only means that the data cannot be easily read. It would still be possible, for example, to write random data to the flash chip in the hope of provoking insecure behaviour. The signing prevents any modified data from being run.

The OTP (one-time-programmable) memory of the RP2350 can be used to

The Raspberry Pi Pico 2 module



Photos 1-4 (left-to-right):

- **the Seeed Technology XIAO RP2350 is one of the smaller RP2350 boards and has a USB-C socket. It appears that it will not cost much more than a Pico 2.**
- **Pimoroni's PGA2350 RP2350B is a compact but comprehensive breakout board for the 80-pin RP2350B. It includes 16MiB of flash memory and an 8MiB PSRAM chip.**
- **the Pimoroni Tiny 2350 appears to be pin-compatible with their Tiny 2040. We noted the Tiny 2040 in our original review of the Pico; it was one of the early RP2040 boards.**
- **Sparkfun's Pro Micro RP2350 has a USB-C socket and incorporates a PSRAM chip, giving access to over 8MiB of random access memory. It also has a 16MiB flash memory chip.**

store the keys needed to decrypt and check flash data, among other things. The OTP can be locked and hidden by programming specific bits.

To test the security, the Raspberry Pi Foundation launched a competition with a $20,000 prize to see if anyone can break into the locked OTP memory.

One of the great features of the RP2040 on the original Pico is the ROM bootloader, which makes it almost impossible to 'brick'. The OTP provides a means to permanently modify the RP2350's behaviour, so it's possible that a wrong OTP operation could brick the RP2350. However, we understand that has been deliberately made difficult to do.

Hands-on testing

We are in the process of doing some detailed testing of the Pico 2 with our previous Pico projects, including the Pico Audio Analyser, which should hopefully improve its performance.

To summarise what we've found, the Pico 2 works just about seamlessly in all cases where we had previously used a Pico! Of course, the differing architectures mean that code recompilation is required, but we generally have not had to make any changes to the code itself.

For example, we fitted a Pico 2 to the prototype of one of our projects under development and compiled the exact same Arduino sketch files (without any changes whatsoever) and the Pico 2 worked exactly as expected. Similarly, the example MicroPython program and libraries that we created for the Back-Pack worked without any changes on the Pico 2.

The process for setting up the Pico-series C SDK (software development kit) on a Windows machine has changed substantially. Still, apart from that, we had little trouble in compiling the exact same code as we used with a Pico.

With just one click, we were able to create a separate project to use the RISC-V processor (instead of the ARM processor) and that too compiled flawlessly and worked identically. Curiously, the compiled RISC-V code is about half the size of the ARM code.

Elsewhere in this issue, you will find an article describing a version of PicoMite BASIC for the RP2350. There is an HDMI video version, using the HSTX peripheral (in addition to VGA), and PicoMite BASIC has been bumped to version 6.0.x.

Pico 2 W

The related Pico 2 W was launched in November 2024, featuring the same Infineon CYW43439 radio module that was used in the Pico W. So, as you would expect, it brings the features of the Pico W with the benefits of the new, faster and more powerful RP2350 chip.

Bare RP2350 chips in all four variants are available, including from DigiKey, Mouser and LCSC. As LCSC is owned by the same company as JLCPCB, you can also get PCBs assembled using RP2350 chips from JLC.

Other RP2350 boards

Other companies have already announced RP2350-based products. It appears some firms have had access to the RP2350 for some time before the launch, allowing them to develop a range of products, test out the chips and their software.

It was the makers of the Bus Pirate (https://buspirate.com) who identified the erratum mentioned earlier. Bus Pirate is an open-source digital tool for working with microcontrollers and other digital ICs.

Photos 1-4 show some of the new boards that have been announced. At the time of writing, we have not seen any of these boards available to purchase.

Conclusion

The Pico 2 is better than the Pico in pretty much every way, as long you get the A3 stepping (or later) that fixes the input pin leakage current problem that was identified in the A2 versions.

The extra RAM and improved ADC are well worth the extra couple of dollars to buy the Pico 2 module over the Pico (unless you have an application in mind which simply doesn't require them).

While it might appear that the Pico 2 could easily obsolete the Pico, there is a note on the Pico's product page that it will be available until January 2036. The Pico 2 is similarly noted as being available until January 2040.

The Pico 2 is available from Farnell, TME, DigiKey, Mouser and numerous smaller electronics outlets. PE

PICOMITE FOR THE RASPBERRY PI PICO 2

WORDS AND MMBASIC : GEOFF GRAHAM
FIRMWARE : PETER MATHER

This new MMBasic interpreter for the Raspberry Pi Pico 2 and Pico 2 W takes advantage of the new features of the Raspberry Pi Pico 2. It is a comprehensive programming environment that converts the Pico 2 into an easy-to-use and powerful platform for beginners and experts alike.

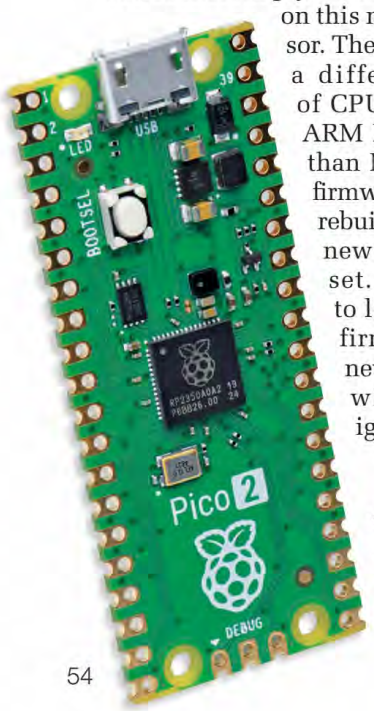
The Raspberry Pi Pico is a complete package with its own power supply, USB interface and more. It is sold at an extremely good price, making it the perfect drop-in microcontroller for many applications.

As described in the review elsewhere in this issue, the new Raspberry Pi Pico 2 has more memory, better performance and more features. The highlights of the Pico 2 are:

- A faster base clock speed, up from 133MHz to 150MHz.
- More efficient CPU cores, up to 50% faster.
- More on-chip RAM, up from 264kiB to 520kiB.
- More flash memory for program storage, up from 2MiB to 4MiB.
- New features, such as the HSTX peripheral for HDMI output and support for external PSRAM.

The original PicoMite was released in January 2022. It is essentially a port of MMBasic from the Micromite/Maximite to the RP2040 chip used on the Pico modules.

You cannot simply use that firmware on this new processor. The Pico 2 uses a different type of CPU core (the ARM M33 rather than M0), so the firmware must be rebuilt to suit the new instruction set. If you try to load the old firmware, the new processor will simply ignore it.



The top side of the Raspberry Pi Pico 2.

With our new release (version 6.00.xx) of the PicoMite firmware, we now support both the Raspberry Pi Pico 2 and the original Pico. However, this firmware is much more than just a recompiled version of the original. We have changed it substantially to make the most of the speed and additional features of the new processor.

Headline features include HDMI video output in various resolutions up to a wide screen resolution of 1280 × 720, and up to 32,768 colours in other resolutions. The video interface has extensive support for sprites, multiple layers, BLIT and other features used in creating detailed graphics for applications such as games.

A new feature in v6.00.01 is a USB interface for connecting a USB keyboard, USB mouse and game controllers. This interface includes support for a USB hub, so you can connect up four devices simultaneously.

The new firmware also supports the extra memory available on the Raspberry Pi Pico 2, which provides BASIC programs with up to 256kiB of program space and 228kiB of general-purpose RAM. We exploit the extra speed of the processor with a default CPU clock rate of 150MHz. It can also be overclocked to nearly 400MHz.

At the core of the PicoMite firmware is the MMBasic interpreter. This is a fully featured BASIC interpreter that is mostly compatible with Microsoft BASIC. It includes features such as long variable names, multiple data types (float, integer and string) and modern structures, such as multi-line IF and CASE statements.

New Raspberry Pi microcontrollers

The Raspberry Pi Pico 2 uses a new microcontroller called the RP2350A, developed by the Raspberry Pi Foundation. There are three other variations of this chip. The first, the RP2350B, is the same as the RP2350A except that it has 20 more pins. This allows for a total of 48 GPIO pins, with eight capable of operating as analog inputs.

The PicoMite firmware for the RP2350 will work with either the A or B variants and will automatically recognise the extra I/O pins when it is running on the RP2350B. They will be available to the BASIC program as GP30 to GP47.

Currently, only a few modules use the RP2350B, primarily supplied by Pimoroni (https://pimoroni.com). However, it is likely that other suppliers will soon follow with their own versions.

The RP2354 A and B chips are the same as the RP2350 versions, but they have 2MiB of flash memory integrated in the package. The PicoMite firmware may support these in the future, but currently they are not readily available for purchase.

The RP2350 includes some security features intended to prevent third-parties from accessing the program and interfering with its operation. These features are not supported in the PicoMite firmware, as we doubt users will be that concerned with security.

Upgrades for the original Picos

This release still supports the RP2040 microcontroller used in the original Raspberry Pi Pico and many other third party modules. If you use the new firmware on a board with an RP2040 chip, you will gain many of the extra features listed here, such as USB keyboard support.

While the original Raspberry Pi Pico is a little slower than the Pico 2 and has a less memory, it is still more than enough for most projects, so you don't need to throw away your old modules.

- More flash and RAM for user programs
- HDMI video output, up to 1280 × 720 pixels and up to 32,768 colours
- VGA video output, up to 640 × 480
- USB keyboard, mouse, Wii controller and hub support
- Improved TCP/IP stack for WiFi boards
- High-speed frequency counter input
- Overclocking up to 400MHz

Load them with this firmware and they will still be very useful.

The Raspberry Pi Pico 2 W

The Pico 2 W has also been recently released, with the same RP2350A processor and the addition of a WiFi interface. This, and the RP2040 version of the MMBasic firmware that we called the WebMite (published in the August 2024 issue), are also supported by this new firmware.

The WebMite firmware running on the RP2040 suffered from a problem that caused the processor to reboot intermittently for no reason. The cause of this was buried deep in the TCP/IP protocol stack that is used to communicate with the wireless interface and, despite a lot of effort, it proved impossible to eliminate.

In this new firmware version, we have completely rebuilt the networking features using a different protocol stack, eliminating the annoying reboots. Thus, we strongly recommend that any designs based on the previous WebMite be upgraded to this version.

The Raspberry Pi Pico 2 W using the RP2350A processor (also called the WebMite) also uses this new protocol stack, so we are confident it will not suffer from the same problem.

The E9 erratum

As explained in the Pico 2 review article, the Raspberry Pi Foundation issued an erratum called E9 for the RP2350. This describes a hardware fault that affects the GPIO and PIO pins that interferes with the use of internal pulldown currents when they are used as digital inputs.

We have implemented some workarounds in the firmware, so you can continue to use the pulldown option for pins configured as digital inputs. However, ideally, an external resistor of 8.2kΩ or less should be used instead.

Beta testing of the PicoMite firmware revealed that this error also affected the ability of MMBasic to communicate using the 1-wire protocol that's used to measure temperature and humidity with DHT22 sensors. However, workarounds added to the firmware for the RP2350 mean these functions are now unaffected.

HDMI support

The RP2350 includes an internal peripheral called HSTX. This is a high-speed serial transmission circuit that streams data to up to eight output pins in parallel. It balances the delay between these outputs to within 0.3ns, making it perfect for generating the signals required for DVI/HDMI video.

To produce such a signal, the PicoMite firmware builds the video image in a reserved portion of RAM (the video buffer) and then configures a DMA (direct memory access) channel using the second CPU core to rapidly push that data to the HSTX peripheral.

The firmware supports three screen resolutions: 640 × 480, 1280 × 720 (wide-screen/720p) and 1024 × 768 pixels. Within each resolution, there are several modes (set by the MODE command) that can trade resolution for more colours and features.

The MODE command can save memory by doubling or quadrupling the size of each pixel, both horizontally and vertically. The monitor will still see the same resolution (ie, the same pixel rate). However, since there will be fewer pixels in the video buffer, the memory saved can be used for more colours.

For example, the resolution can be set to 1280 × 720 using the RESOLUTION command. Following this, the MODE 1 command can be used to generate an image of 1280 × 720 pixels in monochrome, or MODE 3 can be used for a 640 × 360 image in 16 colours, while MODE 4 will provide a 320 × 180 pixel image in 256 colours.

The memory saved by doubling and quadrupling each pixel can also be put to other uses. For example, MODE 4 also releases enough memory for two optional video layers that can be used for an independent overlay. A typical use of this would be to create an image of a moving vehicle overlaid on a background image of a stationary road.

The TILE command

Another handy feature is the TILE command, which allows you to colour individual characters in otherwise monochrome text. So, using the 1280 × 720 HDMI resolution in MODE 1, you can colour each character in one of 16 colours.

This is used by the built-in editor in MMBasic, which uses cyan for keywords, yellow for comments, green for constant numbers and so on.

Screen 1 shows a screen grab of the editor running in the 1280 × 720 resolution with colour coding turned on. This gives you a productive development environment with a colourful wide-screen program editor.

In Screen 1 (shown on page 56), the editor was using the default font 3, which gives 80 characters by 30 lines. If you want more, you can switch to font 1 and have an expansive editing window of 160 characters by 60 lines, still with colour coding.

HDMI overclocking

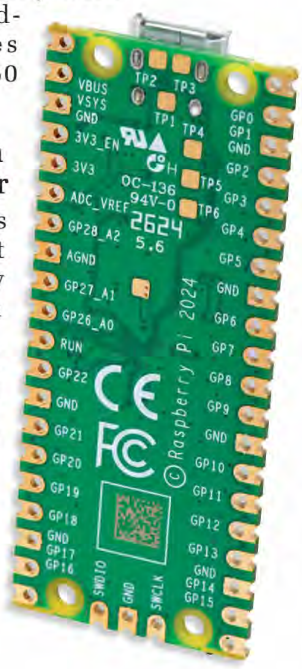
To generate the DVI/HDMI signal, the firmware needs to overclock the RP2350 to as high as 372MHz. Overclocking means running the CPU clock at a higher frequency than the maximum stated in the data sheet. The firmware automatically does this to accommodate the requirements of the video output.

All the Raspberry Pi Pico 2 modules that we tested work perfectly at these speeds. However, overclocking also depends on other components that accompany the processor, and manufacturers might decide to use components that are less tolerant.

For this reason, the HDMI capability won't necessarily work across all third-party modules using the RP2350 processor.

Connecting an HDMI monitor

Fig.1 illustrates how to connect the Raspberry Pi Pico 2 to an HDMI socket. At the high frequencies used by DVI/HDMI, it is important



The underside of the Raspberry Pi Pico 2.

```
Const OrangeLevel = 1000
Const BonusLevel = 3000

Const ManR = 9
Const ax = 18, ay = 18

Dim i, TestMode
If MM.Device$ = "MMBasic for Windows" Then TestMode = 1

If TestMode = 1 Then
  Const xMargin = 200, yMargin = 200
  Dim   ManS = 25          ' Speed of the PacMan
  Dim   GhostS = 35        ' Speed of the Ghosts
  Dim   GhostSlow = 45     ' Speed of the Ghosts when the power ball is active
Else
  Const xMargin = 10, yMargin = 2
  Font 2
  Dim   ManS = 18          ' Speed of the PacMan, Examples: 10 = fast, 18 = normal
  Dim   GhostS = 19        ' Speed of the Ghosts
  Dim   GhostSlow = 23     ' Speed of the Ghosts when the power ball is active
  ' init the button inputs
  For i = 8 To 15
    SetPin MM.Info(PinNo "GP" + Str$(i)), Din, PullUp
  Next i
  FRAMEBUFFER Create
Endif

ESC:Exit  F1:Save  F2:Run  F3:Find  F4:Mark  F5:Paste     Ln: 29  Col: 1     INS
```

Screen 1: using the HDMI output in wide-screen format, you have an excellent editing experience with the built in MMBasic editor. The text is clear and the colour coding identifies elements in the program with cyan for keywords, yellow for comments, green for constant numbers etc.

to keep the signal lines short and of the same length. To minimise reflections in the signal path, it is also recommended that surface-mount resistors be used.

We have seen poor quality cables that exhibited significant crosstalk, ruining the signal. So, if you get a poor image on your monitor, check your HDMI cable as well.

The signal generated by the PicoMite firmware is actually a DVI signal, but HDMI transparently supports DVI and, because HDMI monitors are more common, we recommend using an HDMI connector. However, keep in mind that the PicoMite 2 does not support the transmission of audio in the HDMI signal.

VGA video output

If you do not want to use HDMI, VGA video output is another option. VGA-capable versions of the firmware are available for both the Raspberry Pi Pico 2 (RP2350 processor) and the original Raspberry Pi Pico (RP2040).

It is simpler to connect it to a VGA connector compared to HDMI, and that the processor does not need to be overclocked (although it still can be).

VGA works the same as it did in the previous versions of the PicoMite firmware. It provides two video resolutions: 640 × 480 in monochrome and 320 × 240 with 16 colours. The RP2350 version supports a third mode, with a resolution of 640 × 480 in 16 colours. Refer to the VGA PicoMite article to see how the VGA socket is wired up (see the December 2024 issue).

The tile feature described above also works with the built-in editor using the VGA video output. As a result, you can still edit your program in a reasonable resolution (640 × 480 pixels) while enjoying a colourful editing experience, even though the output is nominally monochrome.

USB keyboard support

The USB connector on the Raspberry Pi Pico (both the original and Pico 2) is normally used to load the firmware and to access the MMBasic command prompt as a virtual serial interface over USB.

However, the Pico's USB connector and electronics are USB OTG (On The Go) compliant, similar to the connector on many mobile phones. This means that it is possible to switch the connector from a USB client (when loading firmware) to a USB host, which is required for communicating with a USB keyboard, mouse or similar.

When you load a PicoMite firmware image with USB capability (available for both the Pico and Pico 2), this switch will be made automatically when the upload is complete and MMBasic starts running. Using an adaptor cable, you can then plug in a USB keyboard and it will be immediately recognised and start operating normally with auto key repeat, function keys, etc.

The Raspberry Pi Pico even supplies the 5V necessary to power the keyboard.

You need the adaptor cable because keyboards usually have a Type-A host plug, while the Pico has a micro Type-B USB socket. These adaptors are common (see Photo 1 for an example) and you can find them online or in local stores.

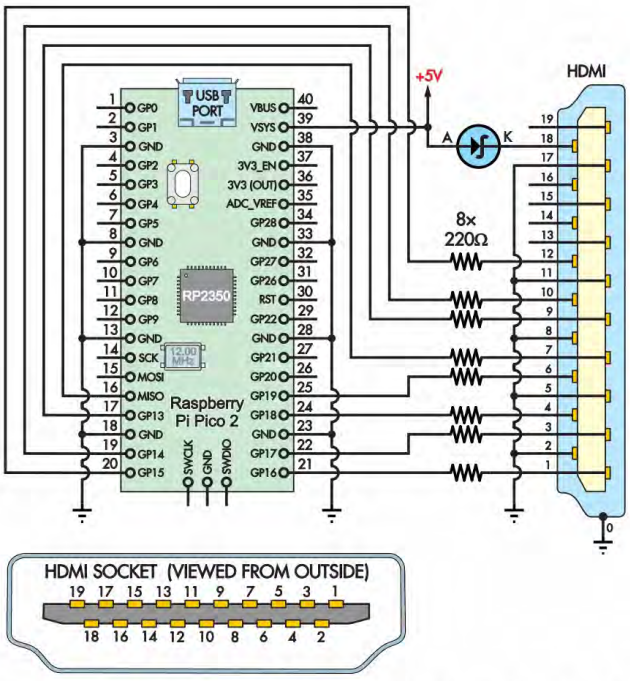


Fig 1: this is how you connect a Pico 2 to an HDMI monitor. It is important to keep the PCB traces short and the same length. The adjacent table shows the function of each pin on the HDMI socket and how they are connected.

Because the USB interface on the Pico is now used for a keyboard/mouse, you will not be able to power the Pico via this connector, so the Pico must be powered via 5V applied to the VSYS pin. Another consequence is that you won't be able to use the serial console over USB.

In a self-contained computer with a keyboard and HDMI or VGA video, this is not normally a problem, as the MMBasic console output will be available on your monitor.

For users who wish to retain access to the serial console, MMBasic automatically switches the console to pin 11 (GP8) for the serial transmit signal and pin 12 (GP9) for receive. It will also set the baud rate to 115,200.

To access this console, you will need a USB to serial bridge that provides a TTL serial interface on one side and a USB interface on the other. These are cheap and commonly available on eBay and similar sites (search for modules using the keyword CP2102).

Using a USB hub

The PicoMite USB capability supports a USB hub and, by using one of these, you can connect up to four USB devices, including keyboards, mouse and Wii game controllers.

You can even plug in multiple keyboards if you wish, and they will all operate in parallel, although why you would want to do that will remain a mystery!

It is better to use an unpowered hub (ie, one that is powered by the Pico's USB connector). This is because the USB protocol stack running on the Pico cannot reset the hub so, if the power on the Pico is cycled without powering down the hub, the hub will keep its previous connections and be confused when the Pico tries to reconnect.

This phenomenon can also cause the hub to be confused if devices are swapped while the hub is powered. If this causes trouble, the simple solution is to cycle the power on the Pico followed by the hub, then plug in the USB devices one by one.

USB mouse support

The USB interface also supports a computer mouse. The main use for this is within the MMBasic program editor, but you can also use it within a program.

If you use the editor with VGA/HDMI video, colour coding turned on and a mouse connected, the mouse position will be indicated by a character in red on a white background. When you move the mouse, this highlight will move accordingly.

HDMI Pin	Function	To Pico Pin
1	TMDS Data 2+ (Red)	GP16 (pin 21) via 220Ω resistor
2	Shield	Ground
3	TMDS Data 2- (Red)	GP17 (pin 22) via 220Ω resistor
4	TMDS Data 1+ (Green)	GP18 (pin 24) via 220Ω resistor
5	Shield	Ground
6	TMDS Data 1- (Green)	GP19 (pin 25) via 220Ω resistor
7	TMDS Data 0+ (Blue + Sync)	GP12 (pin 16) via 220Ω resistor
8	Shield	Ground
9	TMDS Data 0- (Blue + Sync)	GP13 (pin 17) via 220Ω resistor
10	TMDS Clock+	GP14 (pin 19) via 220Ω resistor
11	Shield	Ground
12	TMDS Clock-	GP15 (pin 20) via 220Ω resistor
13	CEC (Consumer Electronics Control)	NC (no connection)
14	ARC (Audio Return Channel)	NC
15	DDS Clock (I²C Clock)	NC
16	DDC Data (I²C Data)	NC
17	Ground	Ground
18	+5V	+5V via schottky barrier diode
19	HPD (Hot Plug Detect)	NC



Photo 1: you need this kind of converter to connect a standard USB keyboard, mouse, game controller or hub to the Pico or Pico 2. This example is Jaycar Cat WC7725.

Clicking on the left mouse button will move the edit cursor to that position (like if you had used the arrow keys on the keyboard), while clicking the right button is the same as pressing F4 on the keyboard (ie, select and cut to the clipboard). Finally, clicking the scroll wheel is the same as using F5 (copy and paste).

This means that, within the editor, you can use the mouse to position the edit cursor, cut or copy text to the clipboard, then paste it in a different location, all without touching the keyboard. This is similar to using an editor in a desktop computer's graphical interface (such as Windows) and makes for a very productive environment.

The mouse position and button states can also be read from within a program by using the DEVICE(MOUSE) function. Similarly, one or more USB Wii Classic game controllers can be used within a program using the DEVICE(WII) function to determine the position of the joysticks and buttons (you may need a USB adaptor to connect them).

Support for external PSRAM

New in the RP2350 is support for PSRAM (pseudo-static RAM). This is a type of RAM chip that sits on a quad SPI bus (similar to flash memory) that can be used to increase the amount of RAM accessible by the RP2350.

For MMBasic, this feature has limited application, as the RP2350 already has plenty of internal RAM (520kiB) for BASIC programs. Because it is accessed via a serial bus, PSRAM is slower than the internal RAM. However, there are programs that might need to create very large arrays and would not mind the slower access.

An example of a module that includes PSRAM is the Pimoroni Pico Plus 2, which comes with 8MiB of PSRAM, a dramatic increase on the internal RAM of the RP2350.

The PicoMite firmware supports PSRAM with the OPTION PSRAM command. When this is enabled, MMBasic will simply add this extra RAM to the general memory used for I/O buffers, strings and arrays. This is transparent to the BASIC programmer, who can then define truly enormous arrays.

Clock speed

We have mentioned before that the Pico's processor can run at various clock speeds. It turns out that the RP2040 and RP2350 processors are quite tolerant of overclocking to above the stock frequency listed in the data sheet.

The standard clock speed for the RP2040 is 133MHz, while for the RP2350, it is 150MHz. These are the defaults used by MMBasic. Most chips will run fine at speeds up to 400MHz, and will only experience a temperature rise of 5-6°C, which is hardly significant, so additional heatsinking is not required.

When the clock speed is increased, it is also necessary for the CPU core voltage (supplied by an internal voltage regulator) to be increased in a balanced manner.

The PicoMite firmware does this automatically; the programmer only needs to use the command CPU SPEED to set the clock rate. For example, with the RP2350 at its base clock speed (150MHz), the CPU core voltage is set to 1.1V, but above 300MHz, it is automatically increased to 1.4V.

At these higher clock speeds, programs run proportionally faster. However, it might not be as easy as that. The main limitation on overclocking is not so much the RP2040 and RP2350 CPUs, but the layout of the Pico's board and the memory (flash and PSRAM) attached to the quad SPI bus.

All the official Raspberry Pi Pico and Pico 2 modules we have tested ran at high speeds without a problem (some as high as 400MHz), but other manufacturers might decide to use components that are less tolerant.

For this reason, the degree of overclocking cannot be guaranteed and there is no way of knowing beforehand how a module may perform. The only certain way of discovering this is to test it yourself.

Versions of the PicoMite firmware that support video output (VGA or HDMI) need to run at a specific clock frequency to generate the correct video timing and this is enforced by the firmware.

For example, the HDMI firmware is fixed at 315MHz for a 640 × 480 pixel resolution and 372MHz for 1280 × 720. VGA defaults to a clock speed of 126MHz and you can only select integer multiples of this, such as 252MHz or 378MHz, if you wish to run it faster.

The internal drive A:

One very useful feature in the firmware is drive "A:", an internal filesystem created when the firmware is loaded. This feature was released in the previous version of the PicoMite firmware, but it is so handy that it deserves to be mentioned here.

Drive A: is a portion of the flash memory on the Pico that is reserved to create a pseudo drive that looks like an SD card or hard disk. It has a normal file system with subdirectories and long filenames, and acts much the same as an SD card, except that you cannot remove it.

Within MMBasic, you can open files on drive A: for reading/writing, rename files, create subdirectories, search for a file, list files and so on. By opening a file for random access, you can even create and operate a miniature database, all within the Pico.

On the Raspberry Pi Pico 2 (with 4MiB flash), its size is just over 2MiB and this will increase if more flash memory is available. For example, with third-party boards using the RP2350 or RP2040 that have 16MiB of flash memory, the size of drive A: is almost 10MiB.

That is a lot of storage (relatively speaking) and allows you to store many images, music tracks, configuration files, log files and more on the Pico without needing to connect an SD card. You can even store multiple versions of your program as you edit and experiment with it.

Once you get used to it, you will find this feature invaluable. The amazing thing is that it is all internal to the Pico – nothing extra is required.

Library support

Another feature that was available in the previous firmware version but is worth mentioning is the library facility. This allows you to add your own commands, functions and features to MMBasic so that they are a permanent part of the BASIC language.

To install components in the library, you write them as normal MMBasic subroutines and functions and use the LIBRARY SAVE command to transfer them to the library. They are then permanently added to MMBasic and will be available to any BASIC program running on the Pico.

For example, you might have written a series of subroutines and functions to retrieve data from a specialised sensor. You could also add them to the library to perform similar functions to those that are already part of the language.

This feature is very handy as, from time to time, you can find yourself thinking that it would be nice if MMBasic implemented some feature that you often need. Now you



We recommend using a controller like this clone of a SNES controller which has a USB Type-A connector, so you don't need to worry about adaptors.

can easily add that feature yourself.

More RP2350 features

Some additional features available on the RP2350 versions include the ability to play MP3 audio files, so you can create your own MP3 player or employ high-quality music as a background to your games. Other audio formats that are supported are WAV, FLAC and MOD.

For high-quality audio, you can use a VS1053 CODEC module or a MCP4822 DAC, so building your own music player is a possibility.

MMBasic also includes support for the extra features in the PIO (Programmable Input/Output) peripherals provided by the RP2350. There are three in the RP2350, and each acts like a miniature CPU controlling a set of I/O pins. Within MMBasic, you can load program routines into the PIOs, set them running, and pass data to/from them.

The VGA output in the PicoMite firmware uses one PIO to generate the video. This is a good example of what you can achieve using this feature.

The ability to convert GP1 into a high-speed frequency counter input is another new feature introduced with the RP2350 that MMBasic supports. This allows you to accurately measure frequencies up to half the CPU clock frequency.

Firmware files

When you unpack the Ver 6.00.01 firmware zip, file you might be surprised to find there are 12 firmware images in it.

These are needed to cover the variations between the CPU (RP2040 or RP2350), the keyboard support (PS/2 or USB), the video output (none, VGA or HDMI) and whether it has WiFi capability or not.

The features provided by the various firmware images fall into one of three categories. The first is a general embedded controller. This is where the Pico might be the brains inside a heating controller, burglar alarm etc. For this application, you might, for example, select a firmware image that supports an attached LCD panel.

The second category is the self-contained, boot-to-BASIC computer reminiscent of the home computers in the 1970s and 80s, such as the Apple II, Tandy TRS-80, Commodore 64 and so on.

This is where you turn the computer on and it boots straight to the MMBasic prompt, at which point you can enter a program, edit it and run it (no operating system is required).

For this, you would select a firmware image that supports a PS/2 or USB keyboard and VGA or HDMI video output.

The third category is the web/internet capable controller (ie, the WebMite) and you have two choices, using either the Raspberry Pi Pico W or the Pico 2 W. These can have an attached LCD panel for displaying data, but their best feature is that they can connect to your WiFi network to serve web pages, access the internet, send emails etc.

This does not mean that you cannot use firmware optimised for one job in another role. The above categories are simply to help make sense of the available options and ultimately the choice will depend on what works best.

A typical filename for a firmware image is "PicoMiteRP2350-VGAUSBV6.00.01.uf2", where:

- RP2350 is the processor that the firmware is compiled for.
- VGAUSB is the feature set supported (VGA and USB).
- V6.00.01 is the version number. This will be incremented in future releases.
- .uf2 is the extension, indicating a loadable Raspberry Pi Pico firmware image.

Table 2 makes it easy to identify the feature set you need and the corresponding firmware image file.

Conclusion

The PicoMite firmware is a comprehensive BASIC programming environment for the Raspberry Pi Pico and Pico 2 that converts the Pico hardware into an easy-to-use platform for beginners and experts alike. It is completely free to download and use.

In this introduction, we have covered many features of the firmware but, in reality, we have just skimmed the surface. There are many more features that are both useful and amazing. For the full story, download the PicoMite User Manual and work your way through that.

This manual runs to over 200 pages and covers all the features of the PicoMite firmware in detail. It even includes a beginner's tutorial in programming in BASIC, so it is easy to get started.

Both the firmware and user manual are available for download from:

- https://geoffg.net/picomite.html
- siliconchip.au/Shop/6/833 PE

Table 2 – firmware variations

Filename prefix	CPU	LCDs	Keyboard	Video	WiFi	Flash	Default clock	Max. clock
PicoMiteRP2040	RP2040	Yes	PS/2	None	None	128kiB	133MHz	420MHz
PicoMiteRP2350	RP2350	Yes	PS/2	None	None	256kiB	150MHz	396MHz
PicoMiteRP2040USB	RP2040	Yes	USB	None	No	128kiB	133MHz	420MHz
PicoMiteRP2350USB	RP2350	Yes	USB	None	No	256kiB	150MHz	396MHz
PicoMiteRP2040VGA	RP2040	No	PS/2	VGA	No	100kiB	126MHz	378MHz
PicoMiteRP2350VGA	RP2350	No	PS/2	VGA	No	180kiB	126MHz	378MHz
PicoMiteRP2040VGAUSB	RP2040	No	USB	VGA	No	100kiB	126MHz	378MHz
PicoMiteRP2350VGAUSB	RP2350	No	USB	VGA	No	180kiB	126MHz	378MHz
PicoMiteRP2350HDMI	RP2350	No	PS/2	HDMI	No	180kiB	315MHz	372MHz
PicoMiteRP2350HDMIUSB	RP2350	No	USB	HDMI	No	180kiB	315MHz	372MHz
WebMiteRP2040	RP2040	Yes	PS/2	None	Yes	88kiB	133MHz	252MHz
WebMiteRP2350	RP2350	Yes	PS/2	None	Yes	208kiB	150MHz	252MHz

Welcome to the sensorium

Techno Talk

From breadcrumb-sized accelerometers to microwatt neural nets and backscattered wireless links, ultra-low-power technologies are turning everyday objects into intelligent, connected companions.

Max the Magnificent

I read a lot of science fiction. Some of these stories go in one ear and out the other, while others are a little 'stickier'. One that keeps floating back to the top of (what I laughingly call) my mind is *The Last Human* by Zack Jordan (**https://pemag.au/link/ac9s**).

One facet of this far-future tale is that every citizen has a brain implant that connects them to "The Network," a super-AI-based information and governance layer that links everyone and everything together.

These implants host their own local, personal companion AIs—digital confidantes that converse with their users and act as tireless assistants. They help you remember things, schedule tasks, filter incoming messages, look up information and perform all sorts of other useful duties.

But wait, there's more. In this universe, just about everything—doors, light switches, toilets, rubbish chutes, ventilation controllers, safety monitors, and goodness only knows what else—is kitted out with its own AI-enabled sensors. All of these are wirelessly connected to "The Network", and hence to each other... and to you.

This may sound a bit dystopian at first blush, but it's not all gloom and doom. Also, if you're worried about ubiquitous AI-infused fixtures, I don't know what you'll make of the gigantic, highly intelligent, killer spider creatures. Still, they belong to a different part of the tale.

Obviously, we're still a long way from this hypothetical future, but we're certainly starting to dip our toes into the sensorium waters. For example...

CMOS MEMS sensors

When I was young, sensors were big, bulky and horrendously expensive. Take accelerometers, for example. These were serious pieces of hardware in the domain of 'cost is no object' applications like airplanes, missiles and rockets.

Fast forward to today's micro-electro-mechanical systems (MEMS), in which microfabrication processes—similar to those used for integrated circuits—are employed to create tiny mechanical structures and sensors on a silicon substrate.

A MEMS accelerometer, for example, involves creating a microscopic cantilever that's anchored at one end and free to bend at the other. There's a gap between the free end and a fixed electrode. The cantilever, gap, and fixed electrode form a tiny capacitor. Under acceleration, the free end of the cantilever deflects ever so slightly, thereby changing the capacitance.

By measuring this change, we can determine the acceleration causing it. The result is that an accelerometer, which used to be the size of a shoebox and consume watts of power, is now smaller than a breadcrumb, consumes only milliwatts, and is cheap enough to appear in everything from smartphones and smartwatches to cat toys.

You may think that MEMS sensors are 'as good as it 'gets', but as the ancient Greek philosopher Heraclitus famously noted, "the only constant in life is change". I recently got to chat with Dr Josep Montanyà, one of the founders of Nanusens (**www.nanusens.com**). Headquartered in Edinburgh, this company's claim to fame is that it can create MEMS-like sensors using the same standard CMOS technology we use to build silicon chips.

The difference between these sensors and traditional MEMS sensors is that they are a thousand times smaller and consume only microwatts of power, while offering higher sensitivity, higher bandwidth, higher yield, higher reliability, and lower cost. What's not to love?

Analog neuromorphic processors

Today's artificial intelligence systems are modelled on the functioning of biological brains. Layers of "neurons" are connected to form artificial neural networks (ANNs). Each of these neurons receives and integrates signals from multiple 'upstream' neurons. If the total activity exceeds a specified threshold, the neuron 'fires' by sending a signal to multiple 'downstream' neurons.

It's tempting to think of all this as being radically new, but such is not the case. Neurophysiologist Warren McCulloch and logician Walter Pitts proposed a mathematical model of a neuron as far back as 1943.

The first physical electronic implementation of an artificial neural network was the SNARC (stochastic neural analog reinforcement computer), which was built in 1951 by Marvin Minsky and Dean Edmonds. The first practical, trainable electronic neural network was the Perceptron, which was developed by Frank Rosenblatt in 1958.

Today's ANNs can be implemented using digital or analog, synchronous or asynchronous, and spiking or non-spiking technologies. Although the vast majority of ANNs are digital, synchronous and non-spiking, other technologies are finding niche applications.

For example, I had a nice chat with Aleksandr Timofeev, the CEO of POLYN Technology (**https://polyn.ai/**). This company, which is headquartered in Bristol, has developed an asynchronous analog ANN based on its NASP (neuromorphic analog signal processing) technology.

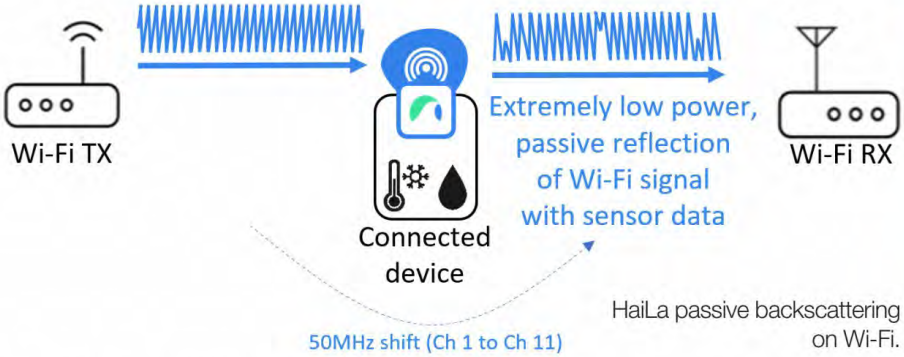
They can take an existing, trained ANN defined in TensorFlow or PyTorch and automatically convert it into an ultra-low-power analog realisation for implementation in silicon as part of a larger device.

For example, a neural network that can detect a human voice requires only around 500 neurons. When implemented in NASP, this neural net consumes only microwatts of power. Colour me impressed!

Working wireless magic

In my previous column, I mentioned a company called HaiLa (**www.haila.io**). Their claim to fame is the ability to convey sensor information over Bluetooth or Wi-Fi while using only a fraction of the power.

"How big a fraction?" I hear you cry. You are wise beyond your years. You obviously know that even 99/100 is a valid fraction in the mathematical sense. This prompts me to ask if you've read *How to Lie with Statistics* by Darrell Huff (**www.amazon.co.uk/dp/0140136290**), but we digress...

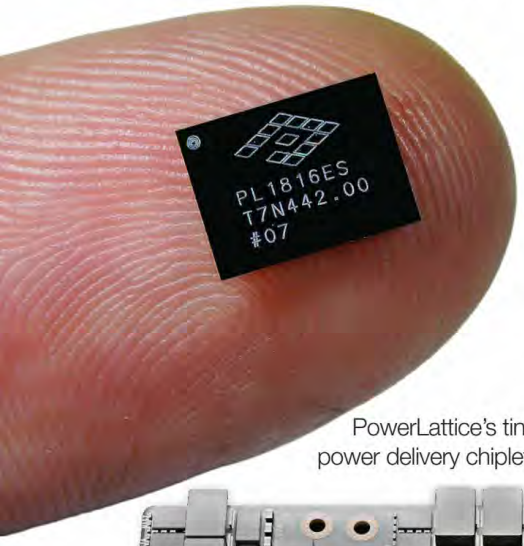


HaiLa passive backscattering on Wi-Fi.

Sustained Bluetooth transmissions typically draw around 10–50 mW, while sustained Wi-Fi transmissions can draw 0.2–1.0 W. By comparison, HaiLa can transmit data while consuming only 10 microwatts. That is low enough that you could run a sensor for 15 to 20 years on a single CR2032 coin cell!

How is this magic achieved? HaiLa's chip doesn't actually transmit information, per se. It uses bistatic passive backscattering to 'reflect' packets as they pass from an external transmitter to a receiver, modulating them as they whiz by.

Since the reflected signal is extremely weak, HaiLa's chip employs a clever 50MHz channel shift to move the backscattered frame away from the original downlink channel. This provides sufficient adjacent-channel rejection to allow the receiver to detect the reflected signal as distinct from the original.



PowerLattice's tiny power delivery chiplet.

Meanwhile, at the data centre

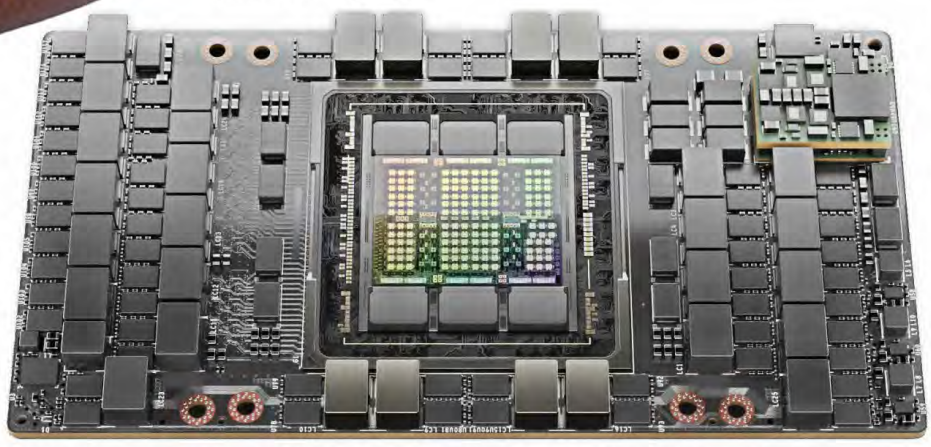
'XPU' refers to any high-end AI processing unit, such as a CPU, GPU, NPU, TPU, FPGA or a custom AI ASIC (application-specific integrated circuit). These beasts can easily consume 1kW. If we assume a core voltage of 1V (which is high), we are talking about 1000A flowing through the accelerator card carrying the XPU!

As transistor counts rise and core voltages fall, pundits are predicting that data centre accelerator cards will carry up to 3000A in the not-so-distant future.

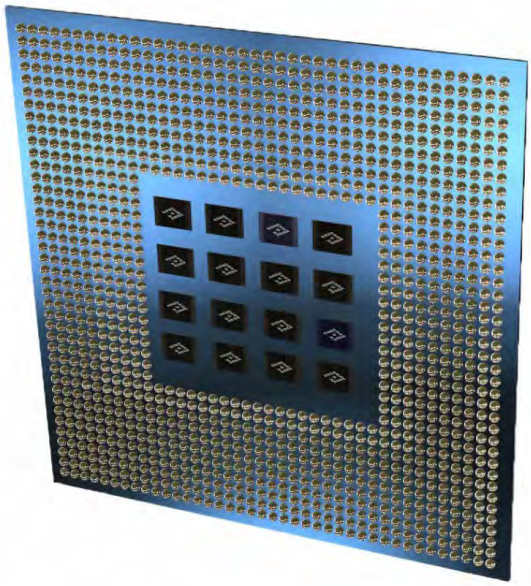
A typical accelerator card has the XPU in the centre of the board. The XPU actually consumes only around 25% of the board's real estate. The rest of the board is devoted to powering it via numerous DC-DC converters. Each of these converters has an associated inductor. The bottom of the board under the XPU is packed with hundreds of capacitors.

This scenario, which is known as horizontal power delivery, has many problems.

For example, the power dissipated in the board's traces increases quadratically with current. Also, because the DC-DC converters are so far from the XPU, they struggle to respond to voltage fluctuations caused by current surges from billions of transistors switching simultaneously.



An H100 accelerator card, covered in power delivery components. Source: Nvidia.



16 power delivery chiplets as part of an XPU package. Source: PowerLattice.

I recently found myself chatting with Dr. Peng Zou, who is the CEO and President of PowerLattice (www.powerlatticeinc.com). PowerLattice is attempting to address these problems and has just emerged from stealth mode with a solution.

The folks at PowerLattice have created a tiny power delivery chiplet so small and thin (a few hundred micrometres thick) that arrays of them can be built into the same package as the XPU.

Each of these power delivery chiplets can handle 200A sustained and 250A peak (which blows my socks off). In addition to control and power logic, they include all necessary capacitors and inductors, with everything fabricated into a single monolithic package.

This means that the remaining 75% of the board is now free. Also, an accelerator card using this vertical power delivery technology consumes only 50% of the power of a card using traditional horizontal power delivery, which will make data centre operators (especially the people responsible for power delivery and cooling) very happy indeed.

Wakey wakey

We may not have brain implants, personal companion AIs, or polite, sentient toilets just yet, but I think we're headed that way.

Sensors are shrinking, wireless links are sipping power, and intelligence is migrating to the very edge of reality. It may not be long before everything around us starts to wake up. I'm not sure whether that prospect should leave us thrilled or terrified. Perhaps both! PE

AUDIO OUT



By Jake Rothman

The Penta-Fuzz 'bender' PCB



This new PCB can be used to build many different fuzz boxes.

This month's pedal offering includes a PCB that can accommodate five different fuzz box designs: a simple silicon diode circuit, a buffered output version, one with extra output gain, a generic Big Muff, and a negative-earth germanium output transistor fuzz pedal.

Again, they draw on my Colorsound experience, and hopefully clear up the incorrect circuits online. I've even had one US company trying to sell me a schematic for a Colorsound circuit I designed!

No more pot hassles

And finally, the problem of fitting standard, low-cost 16mm potentiometers to a flat PCB has been solved after 20 years of head and pot scratching. It's all thanks to our PCB designer Mike Grindle having an epiphany.

Having been brought up on PCB tape-up and hand drilling, my mind was totally closed to the idea of drilling three 16.5mm holes in a small board. Grindle knew that modern PCB fabrication can deal with any sized hole, and this led to his simple solution, shown in Photo 1.

He managed to fit the back of the pot into a hole routed in the PCB. If you use Tayda pots and they were supplied with dust covers, remove them before installation. Remember to cut off the anti-rotation lug, as shown in Photo 2. Use old, strong side cutters (or grab it with chunky pliers and flex it until it snaps off). Don't use your posh Swiss cutters!

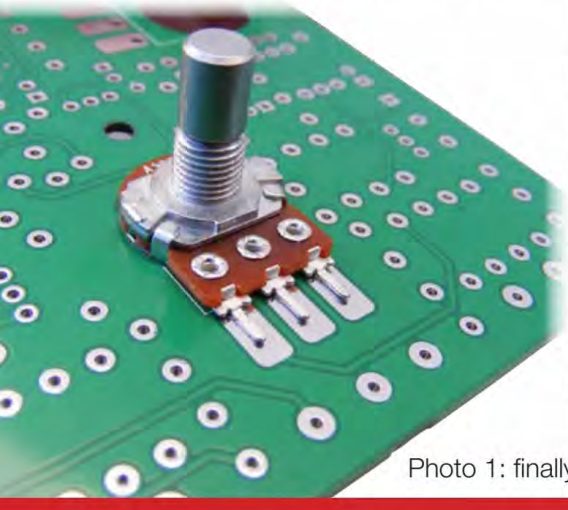


Photo 1: finally, a cheap solution for pot mounting!

The Penta-Fuzz board

The diagram of the available circuit blocks on the board is shown in Fig.1. This is the same as the basic fuzz pedal block diagram shown in Fig.3 of the November 2025 issue, but with an extra output stage after the volume control.

The circuit diagram in Fig.2 shows all the component positions for the five designs. Of course, not all are required for each version, and there are value differences, but it's useful for the experimenter to see the equivalent circuit of the whole PCB to allow one's own circuits to be developed. There are three outputs on the PCB: normal, low-Z and 'supa'. (Is it line level? Is it +4dBm? No, it's supa!)

The Silverback

The simplest design is basically a Colorsound Silver Tonebender, but we can't call it that – "Silverback" sounds suitably silly for a guitar pedal! This circuit is based on a single parallel back-to-back IN4148 diode clipper (diodes D4 and D5). This gives a smooth, soft square-wave fuzz.

The AC coupling of the signal to the diodes provided by C9 reduces at low frequencies, which compensates for the increased output

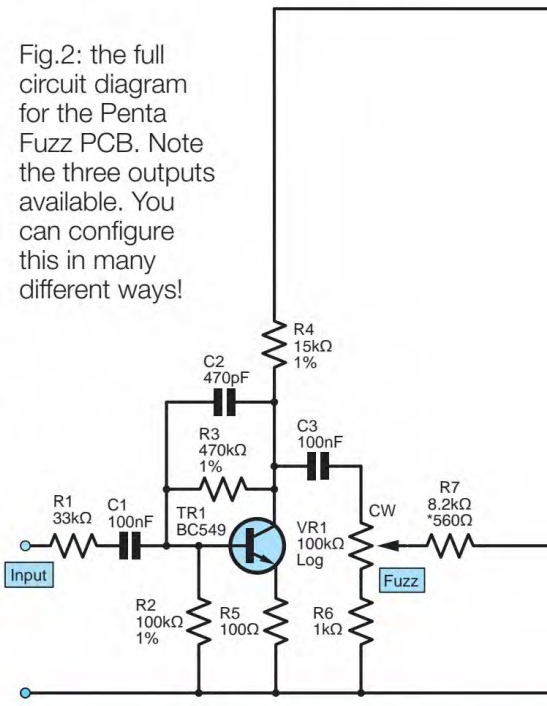


Fig.2: the full circuit diagram for the Penta Fuzz PCB. Note the three outputs available. You can configure this in many different ways!

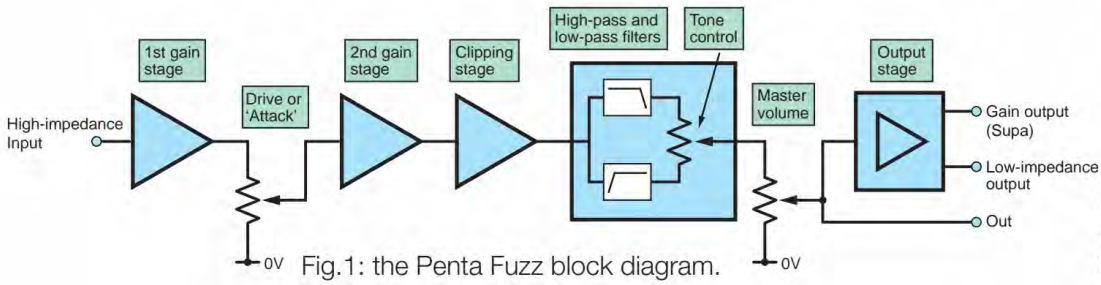


Fig.1: the Penta Fuzz block diagram.



Photo 2: cutting off the pot lug.

of the guitar at lower frequencies, giving more uniform clipping. The inherent harshness of square waves is moderated by the tone control, called a 'tone stack' in guitar parlance. This puts a dip in the midrange when set flat, called 'scooping'.

The sound is akin to that of a kazoo (obviously the instrument that all the best guitar shredders aim to imitate), which sounds smooth and integrates well with electronic sounds. So it is a safe fuzz box for beginners. The passive high-pass and low-pass filters are set to 1kHz and 500Hz, respectively.

There is around 10dB attenuation here, giving a low output level, around 300mV peak-to-peak. This is sufficient to drive most guitar amplifiers, which have plenty of gain. The normal output of the board is used.

Tone control

On this board, it may appear to some that the tone control VR2 goes the wrong way, with the treble end being anti-clockwise and the bass end in the clockwise direction. This is because it was this way on the original Tonebender pedals, going from a 'thin' to a 'full-bodied' sound.

For the more electronically minded folk who prefer clockwise rotation to emphasise high frequencies, the high-pass and low-pass filters will have to be reversed. R17 will become C12 and C13 will become R18.

The Bufftone

The output impedance of the filter/volume circuit is high (25kΩ), and this can interact with long guitar leads that may have high capacitance (up to 20nF!). In this case, you can add a buffer stage using transistor TR4. R21 has to be linked out and R20 is increased in value, to 470kΩ.

The result sounds slightly fizzier with a typical guitar lead than the Silverback and is also better at driving other effects units. The low-Z output is used.

Supa trouper

The stage around TR4 can also be configured as a common-emitter amplifier to give an extra boost of up to four times, to compensate for losses in the filter/volume circuit as well as the buffer stage.

This is for those who want to overdrive their valve amps or need a higher output level (0.9V peak-to-peak, 318mV RMS, -7.7dBm, approaching line level status) for studio effects and power amplifiers. The circuit is basically the Colorsound Supa Tonebender; hence, the "supa" output is used.

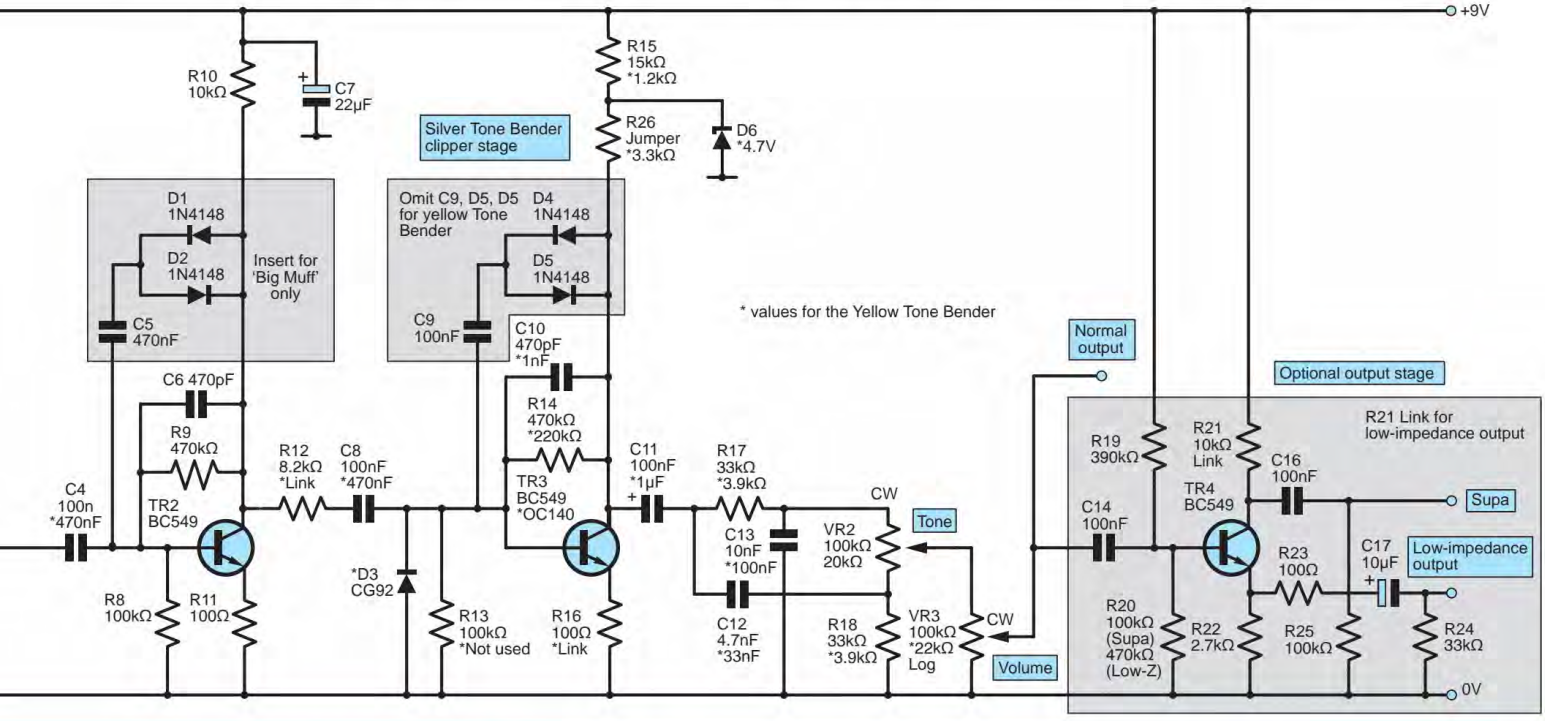
Wave Grinder

I couldn't think of a name worthy of surpassing Big Muff (the famous pedal first released around 1970), but the general circuit is simple enough. It was arguably a derivative of the Tonebender, but all fuzz designers copy each other anyway, so it is difficult to identify the true originators of these circuits.

The transistors have been changed to PNP types (specifically, 2N5087s), plus an extra clipping stage and an output stage have been added.

There are around 20 variations of the Big Muff, and they are well detailed on Kit Rae's π page (www.bigmuffpage.com). The circuit board will accommodate all these circuits. The circuit I use on this board is based on the 1973 design and values shown in Fig.2. There's no need to use PNP transistors; good old BC549s work perfectly.

During testing, I left C5 off. This was so I could compare the



difference between using two clipping networks rather than one. I preferred the more aggressive Big Muff sound with the extra network, especially with C5 selected as 470nF.

A lot of Big Muffs have bigger coupling capacitors throughout, but I found that just increasing C5 made a difference, except with bass guitar, and then it becomes a Bass Fuzz. (The Colorsound Bass Fuzz is the Silver Tonebender circuit with all the coupling capacitors increased to 220nF).

If you do not want to experiment, build the standard Big Muff/Wave Crimper circuit with the output gain stage. This is the best one, in my opinion.

Germanium OC139 fuzz

This germanium design is equally good but more expensive to make, offering a softer fuzz at low levels for jazzy sounds.

A well-known trick among fuzz freaks is to use a flat battery. Fuzz pedals often sound better as the battery wears out; voltage starvation always increases non-linearity and clipping. In the Yellow Tonebender, I emulated this effect by clamping the 9V power rail with a 4.7V zener diode.

The clipping stage (TR3) is a germanium common-emitter stage. It has to be run at twice the current (0.7mA) of the Silverback silicon stage because of the higher leakage current and lower impedance of the filter/volume stage.

In the Yellow Tonebender, an NPN OC140 computer transistor selected for low leakage (<0.1mA) allows for a high-value (220kΩ) bias resistor, R14. To get maximum gain, emitter resistor R16 is linked out. A germanium diode (D3) is also used for leakage cancellation and DC bias-shift clamping.

If you can't get a low-leakage NPN germanium transistor, a PNP can be used, but then a complete polarity flip is required where the other transistors are changed to PNPs, eg, BC559s, and all the diodes and electrolytic capacitors are reversed.

I've found the old glass-cased (SO-2/SB3-2) germanium transistors, such as the OC series, to

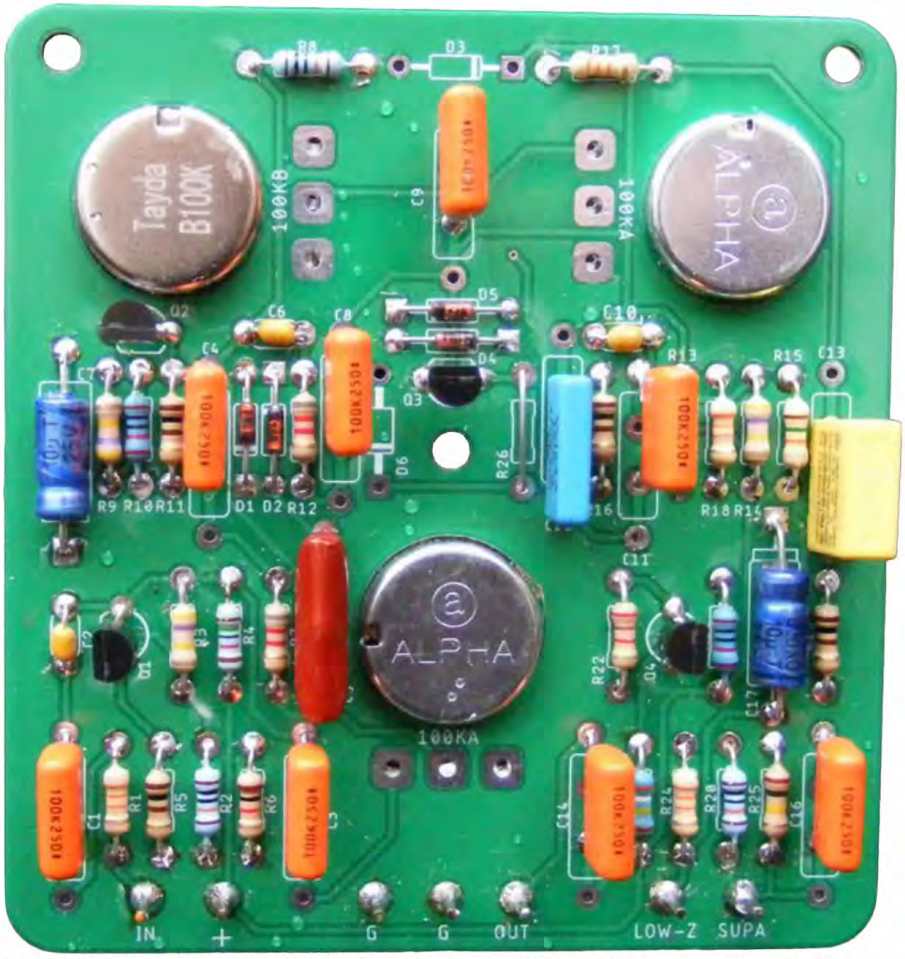


Photo 3: the assembled Wave Grinder board. The Phillips orange polyester capacitors are the modern replacement for the Mullard C280s. Not quite as good-looking, but at least the leads don't fall off.

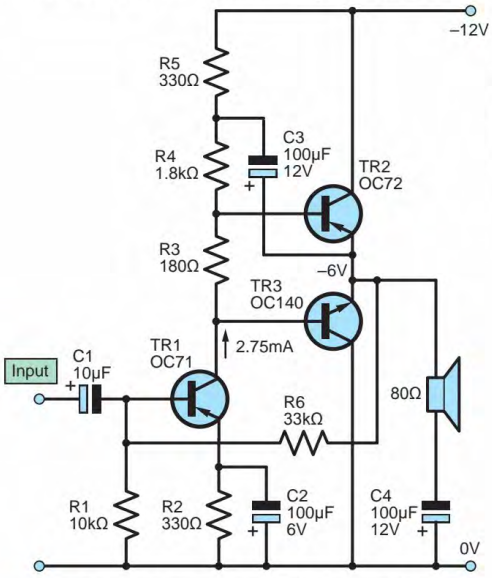


Fig.3: an interesting OC140 circuit from "Transformerless Circuits for Broadcast Receivers" by Macario and Broadberry, Wireless World, March 1960.

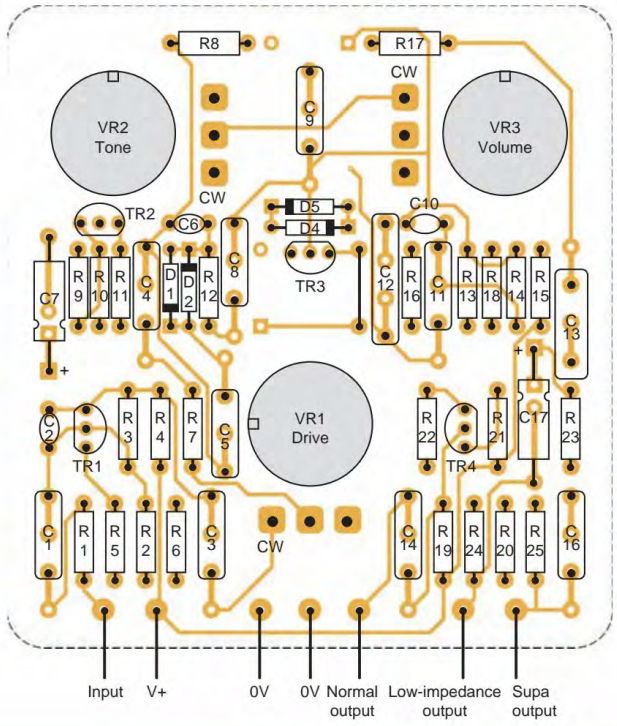


Fig.4: the main PCB overlay (component placement) diagram for the Penta Fuzz pedal.

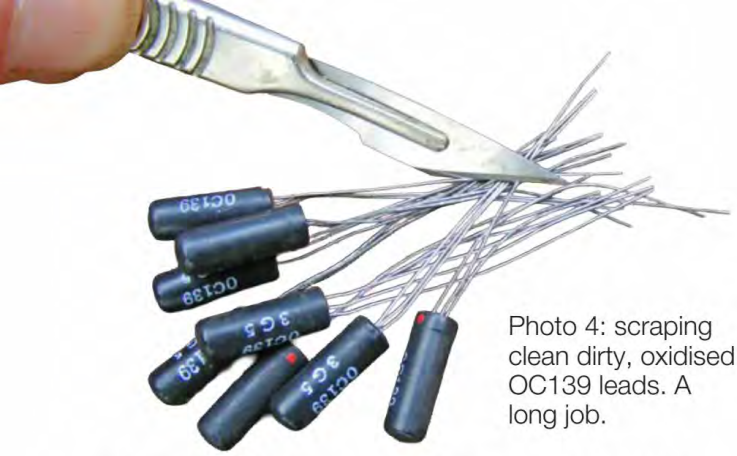


Photo 4: scraping clean dirty, oxidised OC139 leads. A long job.

have lower leakage than the later European TO-1 metal-can type, such as the NKT214 and AC127. Proper hermetic (with a glass ring around each lead) TO-5 and TO-18 cases also seem to be good. The epoxy-filled base versions are not.

Interestingly, the OC139/140 transistors are symmetrical in that the collector and emitter can be swapped with little change to the Hfe. This was revealed when the Peak transistor tester could not identify which was the collector and which was the emitter lead on some examples. The OC139 has half the Hfe of the OC140 (Hfe range 50 to 150) and is perfectly usable.

[Editor's note – early germanium transistors sometimes behave symmetrically due to constraints of the manufacturing processes of the time. Modern devices are deliberately asymmetrical, since optimising each region for its specific role yields far higher gain, lower noise and better overall performance.]

Historical diversion

The Mullard OC139/40 was one of the first European NPN transistors, introduced in 1959. An OC140, along with the OC72, made the first complementary-pair push-pull class-B audio amplifier published in the UK, shown in Fig.3. It delivered 150mW into an 80Ω loudspeaker.

Photo 5: mounting an OC139 germanium transistor for TR3.



I will have to build one to see if it has a good fuzz sound! It is likely to have a very asymmetrical output waveform and might make a unique fuzzy practice amp.

Output details

By lowering the impedance of the tone stack and volume control, and increasing the output level from TR3, the circuit can drive a typical amplifier input directly. In other words, if using the tone section, a separate output buffer (TR4) is no longer required.

The filter frequencies are widened slightly, to 400Hz and 1.2kHz. I used a 20kΩ W-law pot for the tone control. This law is specially designed for tone controls, having a softening of the slope at each end of the rotation, and often a centre detent.

Assembly

I'm a bit naughty; I often solder resistors from the component side of the board these days. This always ensures they lie flat on the board. When I use those foam-backed PCB holders, I usually find a few sitting up in the air when I've finished.

I used to prevent this by bending the leads over. With plated-through holes there's no need for bending, and doing it makes it much more difficult to remove components if that becomes necessary. Fig.4 is the overlay diagram that indicates which components go where, while Photo 3 shows the assembled board.

Golden oldies

When using old components such as germanium transistors and "tropical fish" C280 capacitors, it's essential to scrape the oxide off the leads with a scalpel, as shown in Photo 4. Also, because the leads have lead (Pb) based tinning, 3% activated rosin flux 60/40 leaded solder should be used (for home construction). There was no RoHS back in the 1960s...

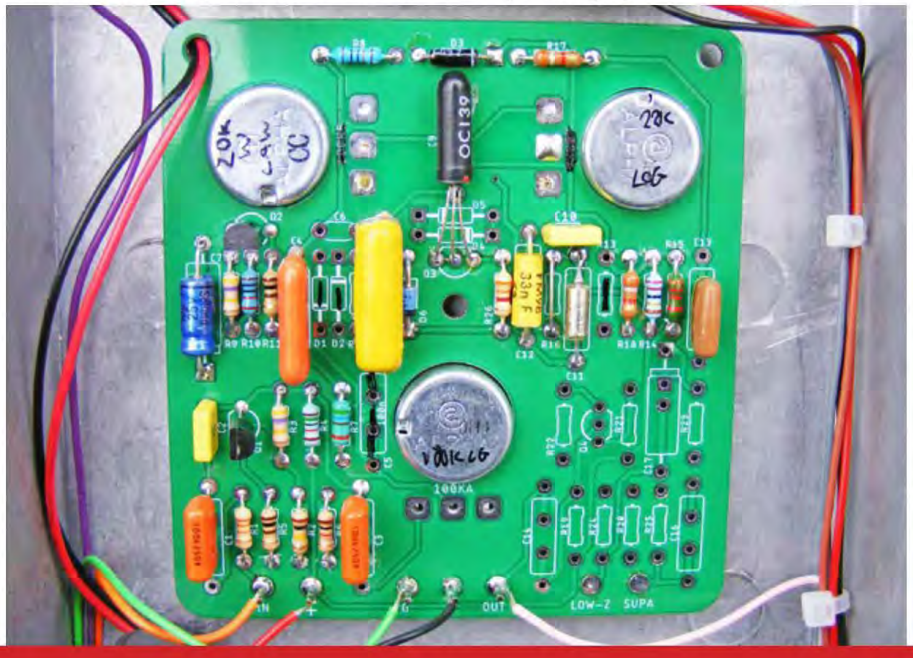
My stock of OC139s and CG92 diodes is from 1963, and the leads are almost black with oxide. In the industry, an ultrasonic solder bath is used, which is so effective that the components can also be converted to RoHS lead-free status, assuming there is no lead inside. They then pass the external X-ray spectroscopy analysis that's used to check for conformity when exporting to the EU.

Mounting germanium devices

These devices are heat sensitive, so solder them quickly, with long mounting leads. It's another reason to avoid lead-free solder, which has to be hotter for longer to achieve satisfactory wetting. Photo 5 shows the mounting of the OC139. A small piece of double-sided tape on the board underneath stops it from flapping about on its long leads.

Because clipping components D4, D5 and C9 are not used in this configuration, they provide a convenient space in which to bend the transistor over. Photo 6 shows the different components used

Photo 6: the PCB built for the germanium transistor fuzz box. There are quite a few different components fitted, and the output stage parts are left off.



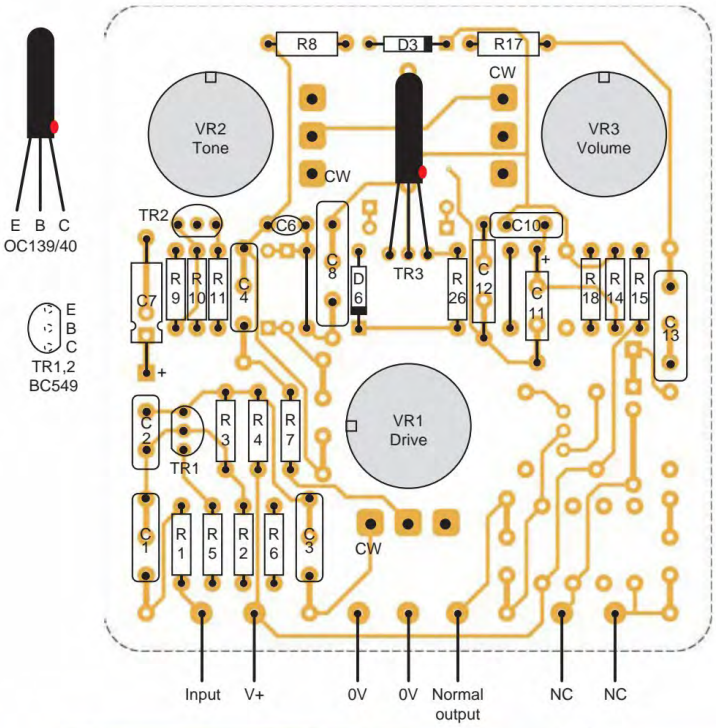


Fig.5: the components used for the OC139 fuzz pedal option.

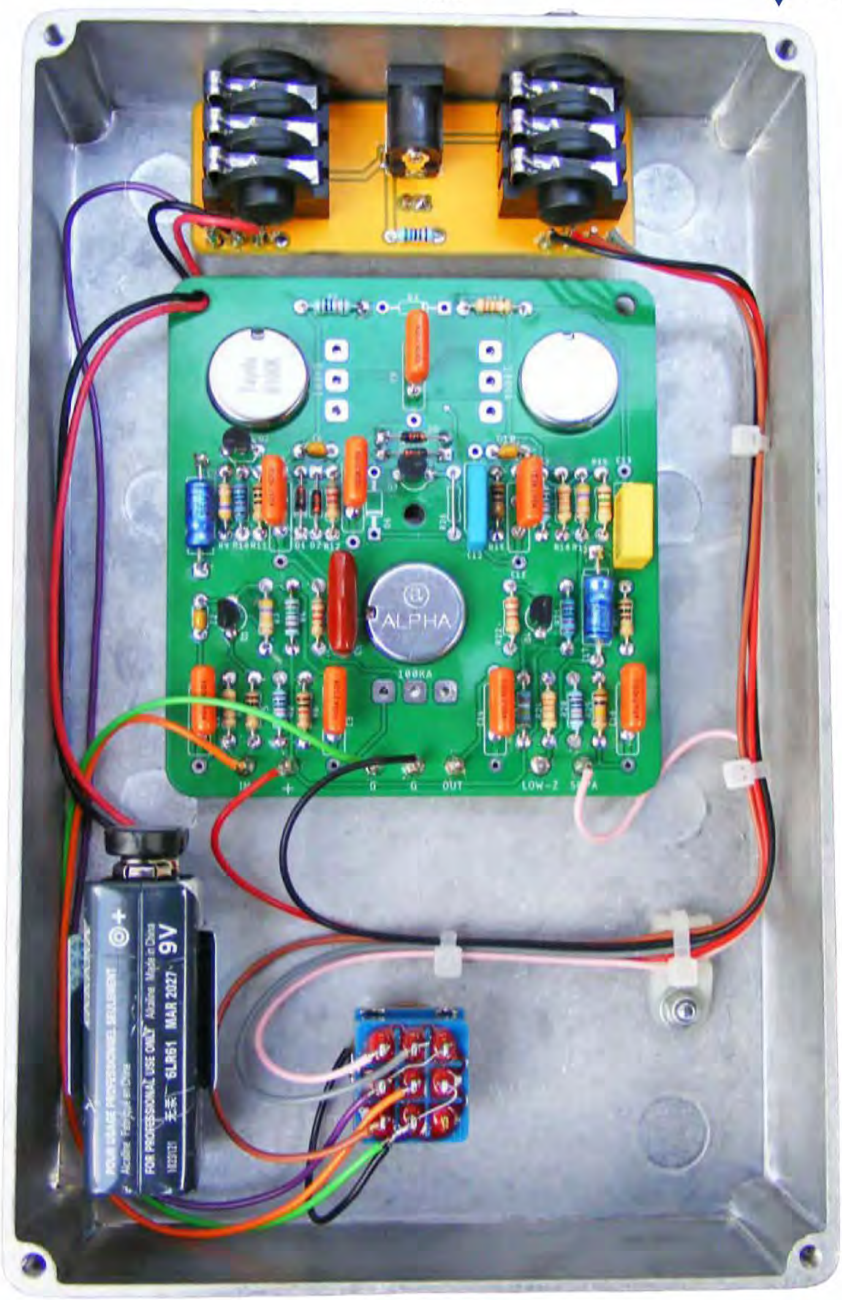


Photo 7: the internal wiring of the Wave Grinder. Notice I've moved the battery clip to the other side so that the wires can go straight to that side of the jack PCB.

compared to the standard board, while Fig.5 is the component overlay diagram.

Boxing it up

As in the previous pedals, I've used a Tayda 1590DD 'Hammond clone' box with virtually the same layout, apart from the battery clip, which I've swapped to the other side to get more lead length. This wiring is shown in Photo 7.

The drilling is the same as before, except that the new potentiometer holes are smaller, at 7.5mm. The jack board is difficult to drill, so I have had a drilled fibreglass template made by the PCB company. It's much easier to mark out the drilling positions using this.

There is little leeway for error when dealing with PCB-mounted connectors and pots. For this reason, I've also made a pot template, which also acts as an under-board insulator if you forget to cut the component lead outs very short. These templates are shown in Photo 8.

Place them in the desired position on your box, draw around the holes, centre punch, make a 3mm pilot hole and then drill to full size. Pfaffing about with a ruler is eliminated.

Wiring

The jack board is the same as last month, but is repeated in Fig.6 for completeness.

Some early boards may be marked for positive earth with the LED square pad denoting the cathode (-Ve). Ignore that and follow this overlay diagram.

The wiring diagram and switching are also the same (Fig.7) bar the three selectable outputs available on the board, and the Earth wire to the switch washer. I've standardised on a colour code for wiring pedals as follows:

- **Red**: battery positive and DC connector output to board. Don't use red for anything else!
- **Black**: battery negative Earth, jacks to board Earth, switch washer Earth
- **Violet**: guitar input signal from the jack socket to the switch
- **Grey**: output signal going to the amplifier socket
- **Pink**: output signal from the main board to the switch
- **Orange**: input signal to the board from the switch
- **Green**: Earth wire from the board to the switch
- **Brown**: LED ground wire to the switch

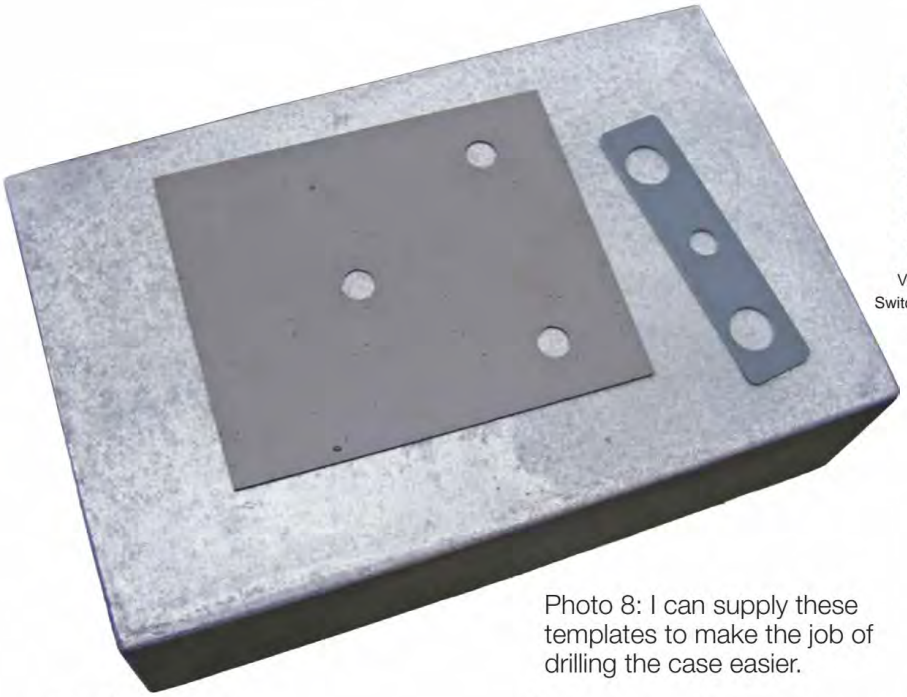


Photo 8: I can supply these templates to make the job of drilling the case easier.

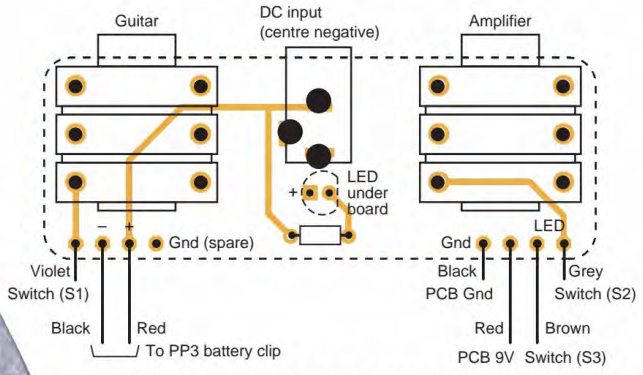


Fig.6: the jack board overlay diagram.

Earthy topics

Since there is usually a multitude of Earth wires in any audio circuit, they can be differentiated from each other by using "earthy" colours. I use dual-colour green wire with various tracers. Occasionally I use blue, except where it is used for negative power rails (if present).

Brown can be employed for signal Earths, but this may cause trouble if confused with the ludicrous EU colour for Live/Active in mains wiring. *[Editor's note – there are good reasons why brown was chosen, but if it weren't for historical baggage, it's likely we'd have a more sensible colour for Live wires, like red or orange.]*

I always use tinned 7/0.2 equipment wire for general hook-up jobs. It was once cheap, but it is not any more in Britain's austere de-industrialised economy (or perhaps this just reflects the increased cost of copper). Because it doesn't degrade with time, it's worth going to radio rallies where the wares of dead electronic engineers are sold off. No widow wants her house cluttered with dusty cable reels.

The old pot-mounting system also served to Earth the box. The new pot-mounting method doesn't do this, because the pot bushes don't go through Earthed holes on the board. Now the Earth point is an M12 washer on the stomp switch. This is connected by a black wire to the Earth tags on the switch, as shown in Photo 9 and Fig.7.

Soldering the wire to a standard zinc-plated steel washer is difficult, so it may be better to use a crimp tag. I found some copper

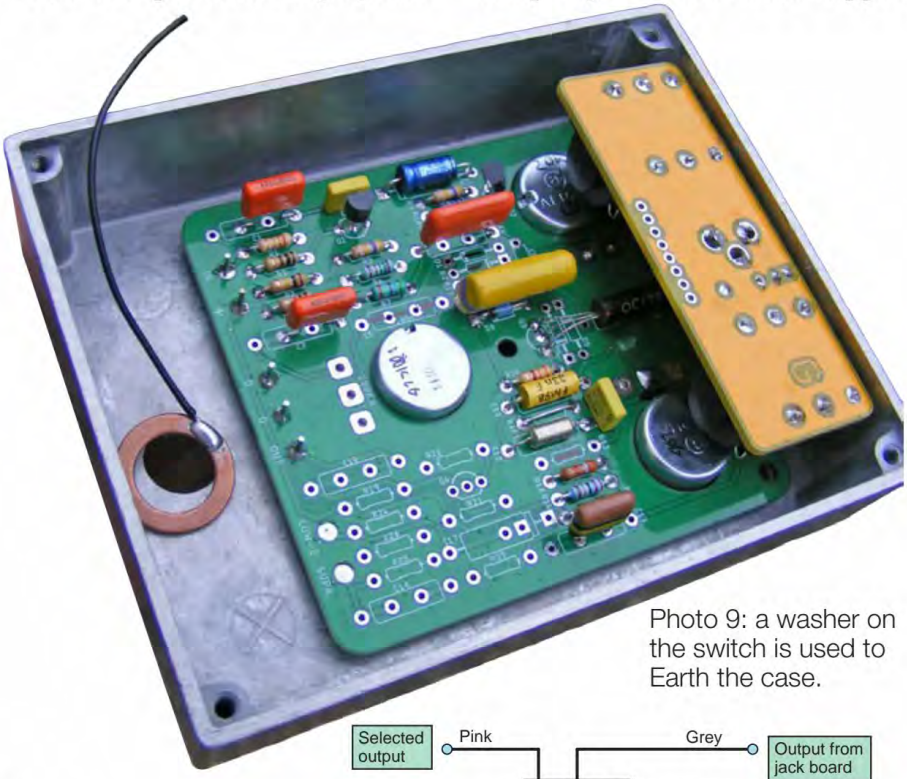


Photo 9: a washer on the switch is used to Earth the case.

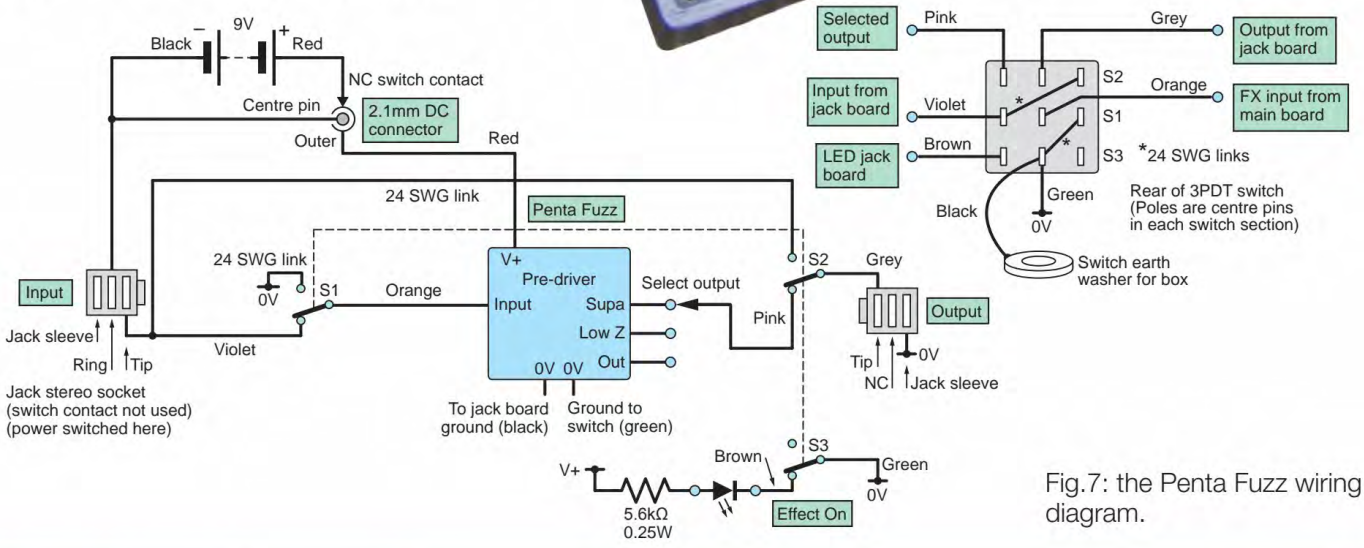


Fig.7: the Penta Fuzz wiring diagram.

roofing washers in the local hardware shop, which were easy to solder, shown in Photo 12.

Incidentally, the word Earth is capitalised in PE, partly to distinguish it from dirt, and partly for consistency with Live and Neutral. We capitalise these because want to make sure our descriptions of mains wiring are as clear as possible (so that there's no ambiguity that we're referring to specific mains wires).

The finished unit is shown in Photo 10.

Testing

The current consumption for each version is:

- Silverback: **2.6mA**
- Buffered Silverback: **4.0mA**
- OC139 germanium fuzz: **5.8mA**
- Wave Crimper: **2.9mA**

These are all very low, ensuring a long battery life. This allows cheap zinc-carbon batteries to be used.

A common problem with testing guitar pedals is getting the two jack leads mixed up. With bypass mode enabled, the pedal works fine because the hard bypass switch conducts in both directions. It then doesn't work when the effect is switched on, because the guitar is connected to the output. So if it only works in bypass mode, try swapping the leads.

Electronic problems are usually revealed by checking the transistor collector voltages. The signal voltage also appears on the collectors, enabling oscilloscope tracing (or using an audio probe). Loss of signal is usually due to a missing component or link.

The worst fault of all is the unsoldered PCB pad. With holes that are plated through or pads with hard bent-over leads, things can sometimes work fine for three months before failure of the contact due to oxidation occurs. We once sent a pedal to the USA with this unknowingly self-inflicted fault. 5p of soldering work became $80 of postage, duty and hassle. Never again!



Photo 10: the finished unit, one of many. All fuzzes sound different but also the same.



Photo 11: the miniaturised OC139 fuzz box.



Photo 12: copper washers are the easiest to solder. I can supply one to readers if needed.

Miniaturisation

Mike Grindle has dispensed with batteries in his pedal designs, just going for the 2.1mm DC connector, avoiding the use of its unreliable battery switching contact. This, along with the use of rear-connected jack sockets with metal barrels for Earthing, allows the case size to be reduced to the absolute minimum.

It does risk an absolute mess of cables if you're using a lot of pedals, though, which is why axemen often preferred batteries. The choice is yours.

I had a go at miniaturising the OC139 germanium fuzz into a 95 × 120 × 32mm Eddystone 29830/P diecast box, shown in Photos 11 & 13. It looks great, but is harder to build and service compared to the one in the large case.

The jack board has to be mounted upside-down and the plastic jacks are touching the pots on the board below. The tags have to be cut off close to the board to stop them from shorting against the lid. Finally, foam has to be stuck to the unused output section of the board to clamp the battery when the lid's screwed down.

Another thing to look out for if powering a pedal with a plugpack is that almost all modern plugpacks are switch-mode designs. Many are fine for driving audio circuits, but some can inject high-frequency noise, especially in a circuit that's dealing with low-level signals like from a guitar.

Later on

A reader asked me about making a phaser pedal, which sounds like a good idea, tying in with a column I'll be doing on all-pass filters for audio circuits. Part one will be theoretical, but a practical project will follow. PE

Photo 13: the miniature fuzz box internals. It is quite a squash.



Parts List – Penta-Fuzz / "Big Muff"

1 double-sided main PCB coded AO-FEB26-1, 80 × 85mm
1 double-sided jack PCB coded AO-JAN26-2, 72.5 × 25mm
1 optional pedal insulation board/drilling template, coded AO-FEB26-2
1 optional jack board rear panel/drilling template, coded AO-FEB26-3
1 Tayda 1590DD diecast metal case, 190 × 121 × 37mm
1 Tayda 3PDT latching stomp switch (S1)
1 battery snap with 200mm-long flying leads **[Rapid 18-0093]**
1 9V PP3 style zinc-carbon battery
1 9V battery mounting clip **[Rapid 18-3480]**
1 2.1mm PCB-mount DC connector **[Tayda A-4118]**
2 6.35mm/¼-inch PCB-mount stereo switched jack sockets **[Tayda A-5079]**
1 M3 × 12mm countersunk machine screw, hex nut and washer
1 12.5mm diameter washer (must be solderable, eg, copper or bright steel)
Various lengths and colours of medium-duty hookup wire
A few small table ties and clips

Semiconductors
4 BC549B/C NPN transistors (TR1-4)
4 1N4148 or similar small-signal diodes (D1, D2, D4, D5)
1 3mm low-current/high-brightness LED **[orange Tayda A-264 recommended]**

Capacitors (all 63V ±20% unless noted)
1 22µF 16V radial or axial electrolytic (C7)
1 10µF 16V low-ESR radial or axial electrolytic (C17)
1 470nF polyester film (C5)
8 100nF polyester film (C1, C3, C4, C8, C9, C11, C14, C16)
1 4.7nF ±10% polyester film (C12)
1 10nF ±10% polyester film (C13)
3 470pF 50V ceramic (C2, C6, C10)

Potentiometers (PCB-mount 16mm single-gang Alpha, 6.35mm/¼in shaft, M7 bush)
2 100kΩ log (VR1, VR3) **[Tayda A-2427]**
1 100kΩ linear (VR2) **[Tayda A-5636]**
(Spare nuts are Tayda A-5037)

Resistors (all ¼W ±5% carbon film or 1% metal film as noted)

2 470kΩ (R9, R14)	4 33kΩ (R1, R17, R18, R24)	1 2.7kΩ (R22)
1 470kΩ 1% (R3)	1 15kΩ (R15)	1 1kΩ (R6)
1 390kΩ (R19)	1 15kΩ ±1% (R4)	4 100Ω (R5, R11, R16, R23)
4 100kΩ (R8, R13, R20, R25)	2 10kΩ (R10, R21)	1 0Ω (R26)
1 100kΩ 1% (R2)	2 8.2kΩ (R7, R12)	

Substitutions/additions/deletions for Yellow Tonebender
1 OC139/40 germanium NPN small-signal transistor (replaces TR3)
1 CG92, OA91 or similar small-signal germanium diode (D3)
1 4.7V 400mW zener diode (D6)
1 1µF 16V tantalum (replaces C11, 100nF)
2 470nF ±20% polyester film capacitors (replace C4 & C8, both 100nF)
1 100nF ±10% polyester film capacitor (replaces C13, 10nF)
1 33nF ±10% polyester film capacitor (replaces C12, 4.7nF)
1 1nF 50V ceramic or polyester film capacitor (replaces C10, 470pF)
1 20kΩ W-law potentiometer (replaces VR2) **[Tayda A-1959]**
1 22kΩ log potentiometer (replaces VR3) **[Tayda A-3558]**
delete C5 (470nF)

Resistors (all ¼W ±5% carbon film)
1 220kΩ (replaces R14, 470kΩ)
2 3.9kΩ (replacing R17 & R18, 33kΩ)
1 3.3kΩ (replaces R26, 0Ω)
1 1.2kΩ (replaces R15, 15kΩ)
1 560Ω (replaces R7, 8.2kΩ)
1 0Ω (replaces R12, 8.2kΩ)
delete R13 (100kΩ)

Using Electronic Modules with Jim Rowe

Mini Uninterruptible Power Supply (UPS)



If there's a blackout when using your computer, it might keep running (eg, off its internal battery or a UPS) but what about your WiFi router? It will likely drop out and not come back until power is restored. This low-cost UPS module can keep it going as well.

Most consumer-grade uninterruptible power supplies (UPSs) have similar configurations, with a storage battery that's charged when mains power is available and switched to running an inverter to replace mains power when it fails. Many use a sealed lead-acid (SLA) battery to store the energy.

In most cases, the switchover takes only 10-25ms, which usually doesn't cause problems with loads like PCs or LCD monitors. When delivering power from the battery via the inverter, most UPSs can do so for at least 20 minutes, even when the load requires its full rated output power. That is generally enough to allow you to save your work and shut down the computer safely.

The mini UPS module we're looking at here is a bit different from that. It is intended to provide continuous 9V or 12V DC power to small electronic devices like WiFi routers while being powered from 5-12V DC. It can supply up to 12W of output power continuously, making it suitable for powering most WiFi routers and many other small devices.

Instead of a sealed lead-acid (SLA) battery, it uses a small lithium-ion battery like a single 18650 cell, which is much smaller than just about any lead-acid battery.

All of the mini UPS module's circuitry is on a PCB measuring 50 × 20mm. It doesn't have an onboard battery holder; the Li-ion battery (which is not supplied) is intended to be connected alongside it.

We obtained the module pictured from an AliExpress supplier called ACELEX, which had it available for only £1 plus shipping. Another supplier on AliExpress called MOKCUM seemed to have an identical module for £2 plus shipping – twice the price, but still surprisingly low.

From the supplier's photos, the MOKCUM module is set to produce a 9V DC output, whereas the ACELEX module produced an output of 12V DC as received. However, as we'll explain shortly, the modules can be easily changed to produce either output voltage.

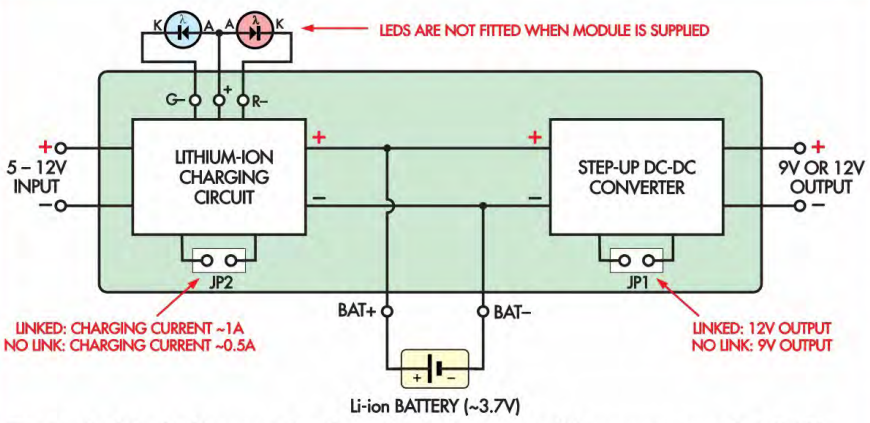


Fig.1: the block diagram for the mini uninterruptible power supply (UPS) module. It is a straightforward design with only two main sections.

How it works

After examining the module's PCB, I was able to glean enough information to produce the basic block diagram shown in Fig.1. There are two main circuit sections; on the left is the lithium-ion charging circuit, while on the right, there is a DC/DC step-up ('boost') converter. The offboard Li-ion cell connects to the lines between the two sections.

The charging circuit accepts the incoming 5-12V DC input power and produces a regulated 4.2V DC output to charge the Li-ion cell while also driving the step-up converter to provide either 12V or 9V to the load on the right.

Link JP1 lets you switch the step-up converter's output between 12V and 9V. When a solder bridge links its pads, the module delivers 12V to the load; when they are not linked, it delivers 9V instead.

Link JP2 changes the maximum charging current for the Li-ion battery. If the pads are not joined by a solder bridge, the maximum charging current is limited to 500mA (0.5A); if they are linked, the maximum charging current is 1A. Most 18650 cells can happily charge at 1A (well under 1C for their typical capacity), but if you are unsure, you can leave it at the safer 500mA setting.

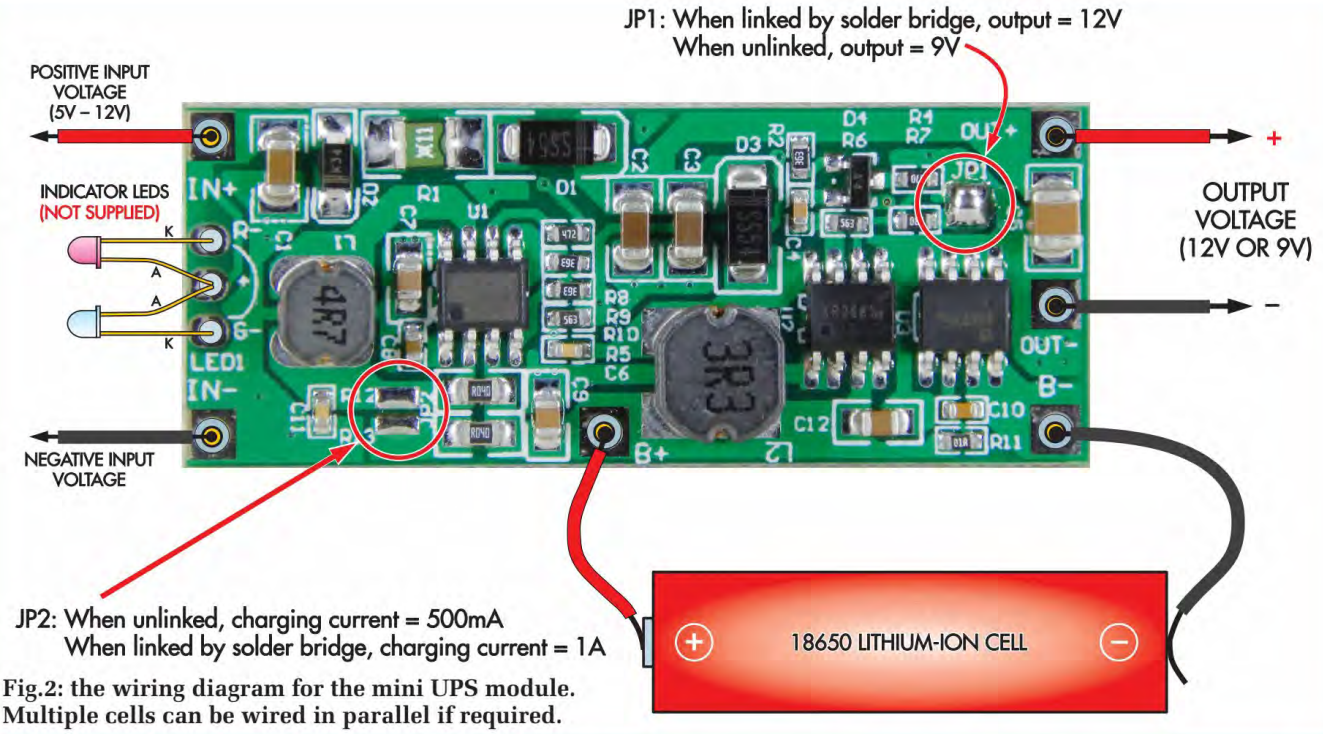


Fig.2: the wiring diagram for the mini UPS module. Multiple cells can be wired in parallel if required.

For small LiPo cells like those used in mobile phones, it's best to leave the JP2 pads open. If you want to use a large cell or several cells in parallel, you will probably want to go for the higher charging current.

The LEDs shown at upper left in Fig.1 are not supplied with the module, but are regarded as an 'optional extra'.

The sketchy data provided with the modules suggests that you should fit a common-anode dual red/blue LED (even though the legends on the PCB show R–, + and G–), but of course, you can use a red/green LED or even two separate 3mm LEDs.

The blue (or green) LED indicates whether a load is connected to the output of the module, while the red LED indicates the charging state of the Li-ion battery. If the red LED is flashing, no battery is connected; if it is on continuously, the battery is being charged; if it is off, it is fully charged.

Fig.1 shows no circuitry to perform the switchover to battery power when the mains-derived input power fails. That's because there is no switchover as such. The Li-ion battery is already connected to the input of the step-up converter, so it will provide current and power when needed. No switchover time at all!

Setting it up

Hooking up the module to a low-voltage power source, a lithium-ion battery and a low-voltage load like a WiFi router is quite straightforward, as shown in Fig.2. The incoming DC supply connects to the IN+ and IN– pads on the left, the output load to the OUT+ and OUT- pads on the right, and the Li-ion battery to the B+ and B– pads at bottom middle and bottom right.

If you want to add a couple of LEDs (or a dual LED), these can be added at centre left, as shown. Just make sure you use high-efficiency LEDs because the driving currents are low.

Link JP1 is just to the left of the output pads, as indicated by the red circle. It's shown linked by a solder bridge, so the boost converter provides a 12V DC output. If you want 9V instead, simply remove the solder bridge with a soldering iron and some solder-wicking braid.

However, note that diode D1 connects the input to the output, so if you set the unit up for a 9V output, you can't use a 12V supply.

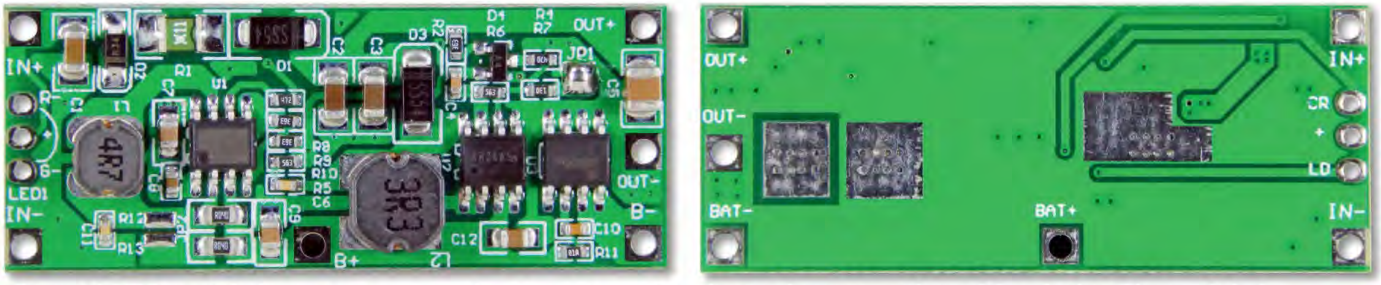
Link JP2 at lower left is indicated by the second red circle. As shown in Fig.2, it usually comes without a solder bridge, limiting the battery charging current to 500mA. It's best to leave it this way unless you know your battery can handle charging at 1A.

By the way, the B–, OUT– & IN– terminals are not all connected together, so make sure your supply, load and battery have independent grounds or else the circuit will not work.

Trying it out

To check out the module, I powered it from a standard 5V DC, 1A plugpack and connected its output to a programmable DC load. I then fired up my bench DMMs, connecting one to the module's output and the other to measure the Li-ion battery voltage.

5-year collections
2019-2023
£49.95
2018-2022
£49.95
2017-2021
£49.95
2016-2020
£44.95
Purchase and download at:
www.electronpublishing.com



The mini UPS module is compact, measuring 50 × 20mm; the photos above are enlarged for clarity. The module is typically supplied as shown with JP1 bridged, JP2 unbridged and no LED(s).

After making sure the sole Li-ion 18650 cell was fully charged, I switched off the input voltage and tested its performance at both output voltages, with load currents of 100mA, 200mA and 300mA. These tests took a few hours, and the results are summarised in Fig.3.

The red/mauve and cyan/blue lines show the module's output voltage at either voltage setting and for the tested load current levels for up to three hours from the removal of input power.

For the lightest loads, 100mA in both cases, the output voltage at either setting remained essentially constant for more than two hours after input power removal.

That corresponds to a load power of 1.2W at the 12V voltage setting and 0.9W at the 9V setting. There was no significant voltage droop over this time. In fact, the voltage on both settings remained within ±2mV for the duration of the tests.

However, it did not last quite as long with a load drawing more current. On the 9V setting, with the load drawing 200mA, the cell voltage fell to 3.2V and I terminated the test after around 2.5 hours.

I repeated the test at 300mA, which naturally gave a shorter runtime, and also with the output set to 12V, which also reduced the runtime.

With the UPS module fed with 12V from a big bench supply (rated at 5A), and two charged 18650 cells in parallel, the module delivered 600mA to the load at 12V for about 10 minutes before the battery voltage dropped to 3.095V.

With a third 18650 cell in parallel and the load current increased to 800mA, even with fully charged cells, the unit could only supply 12V to the load for about 5 minutes before the cell voltage dropped to 2.97V and I turned it off. The small inductor in the output boost converter became very hot in that short time.

So the Mini UPS module is really only really suitable for loads up to 600mA, even with three 18650 cells in parallel. It may be rated to supply 1A, but it wouldn't be able to do so for a useful time.

That's probably enough to power the average WiFi router; many are supplied with a 1A plugpack, although I doubt they draw anywhere near that upper limit unless they are going 'flat out'.

This UPS should be able to power your WiFi router in a blackout for long enough to make it worthwhile with sufficient battery capacity, although that is the kind of thing you should test if you are going to rely on it.

Conclusion

This module is nicely made, low in cost, has no switchover time and performs reasonably well, with the ability to power low power (<12W) DC loads like WiFi routers for about 10-60 minutes, depending on how much current they draw. PE

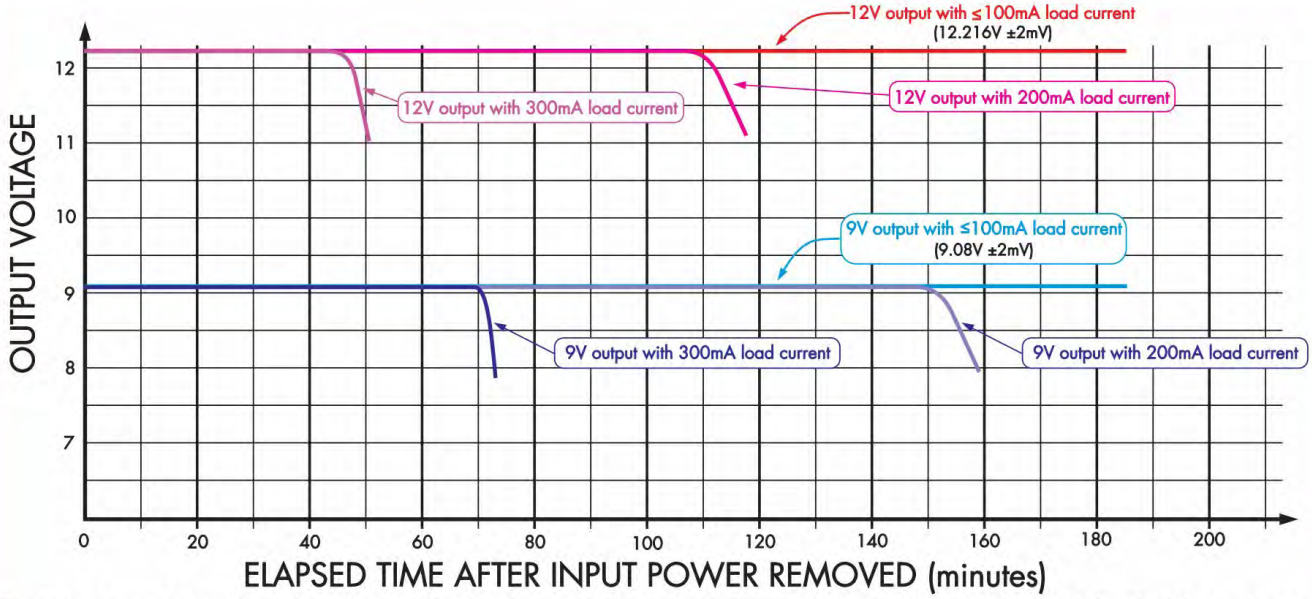


Fig.3: test runs to see how long it would take the module to discharge at 100mA, 200mA & 300mA loads. The unit can deliver up to about 600mA (a little short of the 1A advertised) with reduced runtime unless larger/more cells are used.

Practical Electronics

UK readers SAVE £1 on every issue

SUBSCRIBE NOW!

01202 087631



Take out a one-year subscription and save more than £10 over the year.
Even better – save over £1 per issue if you subscribe for two years – a total saving of £26.

Overseas rates represent exceptional value

You also:
- **Avoid any cover price increase for the duration of your subscription**
- **Get your magazine delivered to your door each month**
- **Guarantee a copy, even if your local newsagent sells out**

Order by:
- **Phone or post with a cheque, postal order or credit card**
- **Or order online at: www.electronpublishing.com**

Practical Electronics

SUBSCRIPTION PRICES

Subscriptions for delivery direct to any address in:

UK: 6-months £33.99, 12-months £59.99, 24-months £114.99

Europe Airmail: 6-months £38.99, 12-months £69.99, 24-months £129.99

Rest Of The World Airmail: 6-months £46.99, 12-months £84.99, 24-months £164.99

Cheques (in **£ sterling only**) payable to ***Select Publisher Services Ltd*** and sent to: Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Tel: 01202 087631
Email: pesubs@selectps.com

Also via our online shop at**: www.electronpublishing.com**

Subscriptions start with the next available issue. We accept MasterCard or Visa.
(For past issues see the Back Issues page.)

DIGITAL SUBSCRIPTIONS

Electronic subscriptions are available from £29.99 for 12 months, more details at: **www.electronpublishing.com**

SUBSCRIPTION ORDER FORM

☐ **6 Months:** **UK £33.99, Europe £38.99 (Airmail), Rest of the World £46.99 (Airmail)**

☐ **1 Year:** **UK £59.99, Europe £69.99 (Airmail), Rest of the World £84.99 (Airmail)**

☐ **2 Years:** **UK £114.99, Europe £129.99 (Airmail), Rest of the World £164.99 (Airmail)**

To: **Practical Electronics Subscriptions**
PO Box 6337
Bournemouth BH1 9EH
United Kingdom

Tel: 01202 087631

Email: pesubs@selectps.com

I enclose payment of £ .....
(cheque/PO in £ sterling only)
payable to ***Select Publisher Services Ltd***

☐ Please charge my Visa/Mastercard

My card number is: .....
Please print clearly, and check that you have the number correct

Card Security Code ..... Valid From Date.....
(The last 3 digits on or just under the signature strip)

Card Ex. Date .....

Name .....

Address .....

Post code ..... Tel. .....

Email .....

Subscriptions start with the next available issue.

The Fox Report

Barry Fox's technology column

Pure audio

It is a perennial problem: one half of a couple is serious about audio and has a top end hi-fi system with complicated controls; the other half just wants to listen in the kitchen, bedroom or home office at the press of a couple of buttons.

There are countless kitchen DAB and FM radios, some also working as Internet radios, but very few of them play CDs. Those that do it all for around £100 sound awful and fall apart after a short period of use.

So there is a real need for something simple and affordable, which is solidly made and does not offend the ears.

I found one that looked hopeful from Pure, the British company originally owned by chip designer Imagination Technologies that kick-started consumer DAB.

But after around ten years at the top, Pure went downhill with OLED displays that stopped displaying and Internet radio portals that stopped porting radio. Pure was then bought by Austrian company AVenture AT and, after initial marketing confusions, has been trying to rebuild its reputation.

So would Pure's Classic C-D6i do-it-all box turn back the clock to the good old days, or turn out to be more junk from China?

I am pleased to report that the "designed in the UK" C-D6i really does do it all, and does it pretty well for an affordable price not too far over £200.

The basic features include DAB+ and FM radio, CD playback, Bluetooth, USB connectivity, Internet radio and podcast listening. The electric basics are a stereo pair of built-in 2 × 3 inch (50 × 75mm) full-range speakers, rated at 15 watts each.

What matters more than the specifications is that the speakers deliver enough clean mid-fi sound for a small room. Put a hand over the speaker bass ports at the rear of the box and you get the satisfying feel of pumping air.

What the company describes as "an intuitive, user-friendly interface" could also be described as "other half" friendly. The user is led through initial setup by a menu that makes entering a WiFi SSID and password without a keyboard as easy as it is possible to be.

The displays shows a side-scrolling alphabet which, unusually, has its upper-case and lower-case character options side by side; AaBbCc and so on. Usually, the user has to keep jumping between separate lists of all upper-case and all lower-case options.

The display is clear bright colour. It's only when you have tried to set up and use a radio with a failing OLED display or weedy LCD backlight that you realise just how important it is for any system to have a decent display. Let's hope these display components last longer than those previously used by Pure.

Searching through the world's gazillion Internet radio stations is always a pain, but Pure eases the pain with an iterative system that lets the user enter short word search strings for "Starts With", "Ends With" and "Contains".

Bluetooth 5.3 allows streaming from a smartphone; a USB socket plays music already ripped from other sources, which is the modern equivalent of making a cassette mix tape. Inconveniently, the USB socket is on the rear. So I added a short USB male-to-female fly lead to make front access easier.

There's no option to feed audio out to anything else. This makes sense because no-one is buying £200 all-in-one boxes to connect them to a serious hi-fi system.

But there is an Aux mini-jack stereo socket input for an external analog line level signal. I tried connecting an old tape deck, and it should also work with a vinyl turntable with built-in pre-amp; or a TV's audio output; or the headphone output of pretty much anything else.

Timers wake you with your favourite morning station and switch the box off for sleep.

As expected, there is an option to create a list of favourite radio stations. Not-so-expected, the list can be a seamless mix of DAB, FM and Internet radio.

A nice touch is that the remote control has six buttons that can be long-press programmed to store six stations of any type. So pressing button 1 calls up an Internet station from California, button 2 switches to BBC Radio 2 on DAB, button 3 gets BBC Radio 3 by FM and so on.

Having both DAB and FM reception makes good sense. In one semi-rural location, auto-scan found 97 DAB stations and only 3 FM stations; in other locations, the results could be just the opposite, with next to no DAB reception but plenty of FM choices.

The Classic C-D6i has an option to listen to podcasts, which involves searching through the jungle of 'casts that are now available, using the same scrolling list used to enter a WiFi password for filtering 'casts by Language, Country of Origin, Category and keywords.

Subscribing to a podcast makes it a Favourite, and there are options to Listen Again and look back at your History of searches.

In practice, I often could not find 'casts that I knew (from using a computer, tablet or phone) to be available. Those that I did find were more by luck rather than judgement.

I reckon it's easier all-round to use a mobile phone or tablet with Apps to search for podcasts and then connect the mobile by Bluetooth to play through the Pure speakers.

The C-D6i comes with a bare 15V DC/2.67A power supply with a twist-fit socket for a choice of separate two-pin European or three-pin UK plugs.

This creates an unecessary opportunity for impatient consumers to get things wrong. I've suggested to Pure that they make things easier for everyone by pre-assembling the PSUs to suit the country where the box is sold.

I didn't hear back, so don't hold out much hope that Pure will do something so simple. *[Companies love it when they can create one 'worldwide' SKU rather than having to keep track of several different versions, and one can understand why that would be easier to deal with – Editor]* PE

DOWNLOAD YOUR PE/EPE BACK ISSUES

COMPLETE-YEAR collections – all 12 issues for each year!

From just £5.95 each

VOL 28 – January to December 1999 – **£5.95**
VOL 29 – January to December 2000 – **£6.95**
VOL 30 – January to December 2001 – **£7.95**
VOL 31 – January to December 2002 – **£7.95**
VOL 32 – January to December 2003 – **£7.95**
VOL 33 – January to December 2004 – **£7.95**
VOL 34 – January to December 2005 – **£8.95**
VOL 35 – January to December 2006 – **£8.95**
VOL 36 – January to December 2007 – **£8.95**
VOL 37 – January to December 2008 – **£8.95**
VOL 38 – January to December 2009 – **£10.95**
VOL 39 – January to December 2010 – **£11.95**
VOL 40 – January to December 2011 – **£12.95**
VOL 41 – January to December 2012 – **£14.95**
VOL 42 – January to December 2013 – **£15.95**
VOL 43 – January to December 2014 – **£16.95**
VOL 44 – January to December 2015 – **£17.95**
VOL 45 – January to December 2016 – **£19.95**
VOL 46 – January to December 2017 – **£22.95**
VOL 47 – January to December 2018 – **£27.95**
VOL 48 – January to December 2019 – **£27.95**
VOL 49 – January to December 2020 – **£29.95**
VOL 50 – January to December 2021 – **£32.95**
VOL 51 – January to December 2022 – **£34.95**
VOL 52 – January to December 2023 – **£34.95 – NEW!**
VOL 53 – January to December 2024 – **£34.95 – NEW!**

FIVE-YEAR collections – 60 issues!

From just £14.95 each

PE2832: Jan 99 to Dec 03 – **£14.95**
PE2933: Jan 00 to Dec 04 – **£14.95**
PE3034: Jan 01 to Dec 05 – **£14.95**
PE3135: Jan 02 to Dec 06 – **£14.95**
PE3236: Jan 03 to Dec 07 – **£14.95**
PE3337: Jan 04 to Dec 08 – **£14.95**
PE3438: Jan 05 to Dec 09 – **£19.95**
PE3539: Jan 06 to Dec 10 – **£25.95**
PE3640: Jan 07 to Dec 11 – **£25.95**
PE3741: Jan 08 to Dec 12 – **£25.95**
PE3842: Jan 09 to Dec 13 – **£33.95**
PE3943: Jan 10 to Dec 14 – **£33.95**
PE4044: Jan 11 to Dec 15 – **£33.95**
PE4145: Jan 12 to Dec 16 – **£35.95**
PE4246: Jan 13 to Dec 17 – **£37.95**
PE4347: Jan 14 to Dec 18 – **£38.95**
PE4448: Jan 15 to Dec 19 – **£39.95**
PE4549: Jan 16 to Dec 20 – **£42.95**
PE4650: Jan 17 to Dec 21 – **£44.95**
PE4751: Jan 18 to Dec 22 – **£44.95 – NEW!**
PE4852: Jan 19 to Dec 23 – **£44.95 – NEW!**

20 YEARS OF PE!

AMAZING BUNDLE OFFER
Every issue from 1999 – 2018

240 magazines for just

£69.95

PE2832: Jan 99 to Dec 03
PLUS!
PE3337: Jan 04 to Dec 08
PLUS!
PE3842: Jan 09 to Dec 13
PLUS!
PE3447: Jan 14 to Dec 18

Download your collections at: www.electronpublishing.com

Practical Electronics PCB SERVICE

PROJECT	CODE	PRICE
FEBRUARY 2026		
Power LCR Meter PCB	04103251	£7.99
Versatile Waveform Generator PCB	04104251	£5.99
JANUARY 2026		
High-Bandwidth Differential Probe	9015-D	£5.99
NFC-Programmable IR Remote Keyfob	15109231	£2.99
Pre-driver Overdrive Pedal main PCB	AO-JAN26-1	£4.99
Updated jack PCB for the Pre-driver	AO-JAN26-2	£3.99
DECEMBER 2025		
Variable Speed Drive (black)	11111241	£11.99
Digital Capacitance Meter	04111241	£5.99
Battery-Powered Model Train transmitter	09110241	£3.99
Battery-Powered Model Train through-hole receiver	09110242	£3.99
Battery-Powered Model Train SMD receiver	09110243	£3.99
Battery-Powered Model Train battery charger	09110244	£3.99
Transmanium Fuzz Box jack PCB	AO-DEC-25	£3.99
NOVEMBER 2025		
Capacitor Discharger	9047-01	£4.99
5MHz, 40A Current Probe	9049-01	£5.99
Transmanium Fuzz Box main PCB	AO-NOV-25	£4.99
OCTOBER 2025		
Seaside Sound Simulator PCB (blue)	01111241	£7.99
Fender Jazz Bass pickguard	23109241	£5.99
J&D Luthiers T-Style pickguard	23109242	£5.99
Music Man Stingray pickguard	23109243	£5.99
Fender Telecaster pickguard	23109244	£4.99
3D Printer Filament Dryer control PCB	28110241	£4.99
SEPTEMBER 2025		
Automatic LQ Meter PCB	CSE240203A	£5.99
Automatic LQ Meter front panel (black)	CSE240204A	£5.99
Compact Hi-Fi Headphone Amplifier PCB (blue)	01103241	£4.99
Discrete Ideal Bridge Rectifier TH PCB	18108241	£2.99
Discrete Ideal Bridge Rectifier SMD PCB	18108242	£2.99
Dual Rail Supply Protector PCB	18109241	£3.99
Audio Out Cascode Amplifier PCB	CASCODEA	£4.99
Audio Out Cascode Amplifier Power Supply PCB	CASCODEP	£4.99
AUGUST 2025		
Micromite Explore-40 PCB	07106241	£2.99
Micromite Explore-40 Audio Breakout PCB	07101222	£2.99
8-Channel Learning IR Remote Receiver PCB	15108241	£5.99
FlexiDice PCB set	08107241-2	£4.99
JULY 2025		
Compact OLED Clock & Timer PCB	19101231	£4.99
180-230V DC Motor Speed Controller PCB	11104241	£9.99
Mains Power-Up Sequencer PCB (same as Feb '25)	10108231	£11.99
JUNE 2025		
Styloclone (case version)	23106241	£6.99
Styloclone (standalone version)	23106242	£7.99
Dual Mini LED Dice (TH LEDs, white solder mask)	08103241W	£2.99
Dual Mini LED Dice (TH LEDs, black solder mask)	08103241B	£2.99
Dual Mini LED Dice (SMD LEDs, white solder mask)	08103242W	£2.99
Dual Mini LED Dice (SMD LEDs, black solder mask)	08103242B	£2.99
Dual Hybrid Supply regulator board	18107211	£5.99
Dual Hybrid Supply front panel control board	18107212	£2.99
Dual Hybrid Supply CPU board	01106193	£3.99
Dual Hybrid Supply LCD adaptor	01106196	£2.99

PROJECT	CODE	PRICE
MAY 2025		
Skill Tester 9000	08101241	£9.99
USB-C Serial Adaptor (black solder mask)	24106241	£2.99
DC Supply Protector (SMD adjustable version)	08106241	£3.99
DC Supply Protector (TH adjustable version)	08106242	£3.99
DC Supply Protector (TH fixed version)	08106243	£3.99
APRIL 2025		
ESR Test Tweezers PCB set	04105241+	£5.99
Fan Speed Controller Mk2	10104241	£4.99
Calibrated Microphone (SMD version)	01108231	£3.99
Calibrated Microphone (TH version)	01108232	£3.99
MARCH 2025		
Pico Gamer	08104241	£9.99
Compact Frequency Divider	04112231	£4.99
FEBRUARY 2025		
Laser Communicator transmitter	16102241	£3.99
Laser Communicator receiver	16102242	£3.99
Model Railway Points Controller transmitter	09101241	£3.99
Model Railway Points Controller receiver	09101242	£2.99
Mains Power-Up Sequencer	10108231	£11.99
JANUARY 2025		
USB to PS/2 Converter for VGA PicoMite	07111231	£3.99
ps2x2pico USB to PS/2 Converter	07111232	£3.99
Compact Microphone Preamplifier (case version)	01110231	£5.99
Compact Microphone Preamplifier (integrated version)	01110232	£5.99
DECEMBER 2024		
Raspberry Pi Clock Radio PCB pair	19101241/2	£12.99
VGA PicoMite (blue)	07107221	£5.99
Secure Remote Switch receiver	10109231	£4.99
Secure Remote Switch transmitter (module version)	10109232	£2.99
Secure Remote Switch transmitter (discrete version)	10109233	£2.99
NOVEMBER 2024		
Ideal Diode Bridge Rectifiers (28mm square)	18101241	£3.99
Ideal Diode Bridge Rectifiers (22mm square)	18101242	£3.99
Ideal Diode Bridge Rectifiers (4-pin SIL)	18101243	£3.99
Ideal Diode Bridge Rectifiers (tiny SOT-23)	18101244	£3.99
Ideal Diode Bridge Rectifiers (standalone SMD)	18101245	£4.99
Ideal Diode Bridge Rectifiers (standalone TH; 2oz copper)	18101246	£6.99
Multi-channel Volume Control (volume PCB)	01111221	£4.99
Multi-channel Volume Control (control PCB)	01111222	£4.99
Multi-channel Volume Control (OLED PCB, black)	01111223	£4.99
Coin Cell Emulator (black)	18101231	£3.99
OCTOBER 2024		
TQFP-32 Programming Adaptor	24108231	£4.99
TQFP-44 Programming Adaptor	24108232	£4.99
TQFP-48 Programming Adaptor	24108233	£4.99
TQFP-64 Programming Adaptor	24108234	£4.99
2m VHF FM Test Signal Generator	06107231	£4.99
SEPTEMBER 2024		
PIC Programming Adaptor	24105231	£4.99
Salad Bowl Speaker crossover	01109231	£7.99
AUGUST 2024		
Arduino ESR Meter combined L/C-ESR	04106182	£9.99
Arduino ESR Meter separate ESR	04106183	£7.99
Watering System Controller	15110231	£11.99

PCBs for most recent *PE/EPE* constructional projects are available. From the July 2013 issue onwards, PCBs with eight-digit codes have silk screen overlays and, where applicable, are double-sided, have plated-through holes, and solder mask. They are similar to photos in the project articles. Earlier PCBs are likely to be more basic and may not include silk screen overlay, be single-sided, lack plated-through holes and solder mask.

Always check price and availability in the latest issue or online. A large number of older boards are listed for ordering on our website.

In most cases we do not supply kits or components for our projects. For older projects it is important to check the availability of all components before purchasing PCBs.

Magazine back issues are available – see **https://www.electronpublishing.com/product-category/pe-back-issues-print-and-digital/**

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
JULY 2024		
Dynamic NFC Tag (22 × 31mm)	06101231	£5.95
Dynamic NFC Tag (86 × 54mm)	06101232	£7.95
Reciprocal Frequency Counter	CSE230101C	£8.95
Reciprocal Frequency Counter front panel (black mask)	CSE230102	£7.95
Pi Pico-based Thermal Camera	04105231	£6.95
JUNE 2024		
Loudspeaker Test Jig	04106231	£12.95
MAY 2024		
GPS-Disciplined Oscillator	04103231	£12.95
Dual RF Amplifier	CSE220602A	£5.95
Songbird	08103231	£7.95
APRIL 2024		
Basic RF Signal Generator	CSE221001	£8.99
Basic RF Signal Generator front panel (black solder mask)	CSE220902B	£9.99
Silicon Chirp Cricket	08101231	£7.99
Test Bench 'Swiss Army Knife'	04110221	£9.99
MARCH 2024		
Digital Volume Control Potentiometer SMD version	01101231	£8.95
Digital Volume Control Potentiometer through-hole version	01101232	£9.95
Advanced SMD Test Tweezers body	04106221	£6.95
Advanced SMD Test Tweezers legs (pair)	04106212	£7.95
FEBRUARY 2024		
Active Mains Soft Starter	10110221	£10.95
JANUARY 2024		
Q Meter	CSE220701	£8.95
Q Meter front panel (black solder mask)	CSE220704	£7.95
Raspberry Pi Pico W BackPack	07101221	£8.95
DECEMBER 2023		
Digital Boost Regulator	24110224	£8.95
Dual-Channel Power Supply for Breadboards	04112221	£11.95
Display Adaptor for the Breadboard PSU	04112222	£9.95
NOVEMBER 2023		
LC Meter Mk3	CSE220503C	£9.95
DC Supply Filter for vehicles	08108221	£8.95
Discrete Audio Op Amp PCB	AO1-JUL23	£9.95
Discrete Audio Op Amp PCB + essential components	AO2-JUL23	£17.95
OCTOBER 2023		
Buck/Boost Charger Adaptor	14108221	£9.95
PIC Breakout Board for SOIC parts	24110225	£5.95
PIC Breakout Board for DIP parts	24110222	£5.95
AVR64DD32 Breakout board	24110223	£5.95
Automatic Train Controller	09109221	£6.95
Chuff Sound module	09109222	£6.95
SEPTEMBER 2023		
Mini LED Driver	16106221	£7.95
New GPS-Synchronised Clock	19109221	£9.95
AUGUST 2023		
110dB RF Attenuator	CSE211003	£9.95
Universal Battery Charge Controller (2023 update)	14107192	£9.95
Wide-Range OhmMeter	04109221	£12.95
JULY 2023		
Multimeter Checker/Calibrator	04107221	£10.95

PROJECT	CODE	PRICE
JUNE 2023		
Arduino Programmable Load	04105221	£8.95
Buck-Boost LED Driver	16103221	£8.95
MAY 2023		
Precision AM-FM DDS Signal Generator	CSE211002	£14.95
Improved SMD Test Tweezers programmed PIC	0410621PIC	£12.95
AO 2x-dual-to-quad through-hole (pack of five)	AO1-MAY23	£6.95
AO 2x-dual-to-quad SMD (pack of five)	AO2-MAY23	£6.95
AO 2x-single-to-dual through-hole (pack of five)	AO3-MAY23	£6.95
AO 2x-single-to-dual SMD (pack of five)	AO4-MAY23	£6.95
APRIL 2023		
500W Amplifier Module	see p.22, April 2023	
Clipping Indicator (per channel)	01112211	£7.95
CD Welder Power Supply (one needed)	29103221	£9.95
CD Welder Controller (one needed)	29103222	£9.95
CD Welder Energy Storage module (several needed)	29103223	£7.95
AO Universal Dual Op Amp Board	AO1-APR23	£9.95
AO Stereo RIAA precision passives kit for dual op amp	AO2-APR23	£8.95
MARCH 2023		
Pico BackPack	07101221	£9.95
Semaphore Signal (controller)	09103221	£7.95
Semaphore Signal (blade)	09103222	£5.95
FEBRUARY 2023		
Tesla Coil driver board	26102221	£9.95
Tesla Coil potentiometer board	26102222	£5.95
Cooling Fan Controller & Loudspeaker Protector	01102221	£8.95
Remote Gate Controller	11009121	£12.95
JANUARY 2023		
Classic LED Metronome – 8-LED	23111211	£7.95
Classic LED Metronome – 10-LED	23111212	£8.95
Multi-Channel Speaker Protector – 6 channel	01101221	£9.95
Multi-Channel Speaker Protector – 4 channel	01101222	£7.95
Remote Control Range Extender – IR-to-UHF	15109212	£5.95
Remote Control Range Extender – UHF-to-IR	15109211	£7.95
AO Universal Single Op Amp Board	AO1-JAN23	£7.95
DECEMBER 2022		
Hummingbird Amplifier	01111211	£9.95
SMD Trainer PCB	29106211A	£8.95
NOVEMBER 2022		
USB Cable tester – main PCB	04108211	£12.95
USB Cable tester – front panel	04108212	£5.95
USB Cable tester – optional panel	SC5970	£5.95
Model Railway Carriage Lights – PCB	09109211	£6.95
AO trnsfmr PCB – standard VTX-A range	VTX-101-007	£6.95
AO trnsfmr PCB – dual-outline VTX102-3001/101-3002	VTX-Dual	£6.95
OCTOBER 2022		
SMD Test Tweezers – PCB and pair of tweezer arms	04106211-2	£11.95
SMD Test Tweezers – programmed PIC12F1572-I/SN	0410621A	£7.95
Tele-com	12110211	£12.95
SEPTEMBER 2022		
Touchscreen Digital Preamp – main board	01103191	£12.95
Touchscreen Digital Preamp – adaptor board pair	01103192	£5.95
20A DC Motor Speed Controller	11006211	£9.95
AUGUST 2022		
Multi-purpose Battery Manager – I/O Expander	11104212	£5.95
Multi-purpose Battery Manager – Switch Module	11104211	£8.95
Simple MIDI Music Keyboard (for 8 switches)	23101213	£6.95
Nano Pong	08105212	£7.95

All prices include VAT and UK p&p. Add £4 per *project* for post to Europe; £5 per *project* outside Europe.

Orders and payment should be sent to:

Practical Electronics, Electron Publishing Ltd
113 Lynwood Drive, Merley, Wimborne, Dorset BH21 1UU
Tel 07918 614662 Email: pe@pemag.au
Online Shop: www.electronpublishing.com

Sorry, cheques cannot be accepted as we are based in Australia and our UK bank account does not allow cheque deposits.

NOTE: Most boards are in stock and sent within seven days of order receipt. Please allow up to 28 days delivery if we need to restock.

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
JULY 2022		
Silicon Labs AM/FM/SW Radio	CSE210301C	£10.95
Level Crossing Controller	09108211	£6.95
JUNE 2022		
Full-wave Motor Speed Controller	1010221	£8.95
PIC Programming Helper for 8-pin PICs only	24106211	£7.95
PIC Programming Helper for 8, 14 or 20-pin PICs	24106212	£10.95
Advanced GPS Computer	05102211	£9.95
MAY 2022		
Bus board PCB for Analogue Vocoder	AO1-MAY22	£10.95
Complete set of 14 PCBs for Analogue Vocoder	AO2-MAY22	£97.95
Programmed EEPROM for Digital FX Unit	FX1-MAY22	£10.95
Programmed PIC for Digital FX Unit w/ potentiometer	FX2-MAY22	£8.95
APRIL 2022		
64-key MIDI Matrix shield	23101211	£8.95
64-key MIDI Matrix switch board	23101212	£11.95
High-current Battery Balancer	14102211	£10.95
Digital FX Unit – using potentiometer	01102211	£9.95
Digital FX Unit – using BCD switch	01102212	£9.95
Universal Audio PSU	AO1-APR22	£11.95
MARCH 2022		
Mini Isolated Serial Link	24102211	£5.95
Busy Loo Indicator	16112201	£5.95
Analogue Vocoder – Band-pass filter board	AO1-MAR22	£9.95
Analogue Vocoder – HP/LP filter board	AO2-MAR22	£9.95
FEBRUARY 2022		
Arduino-based Power Supply	18106201	£9.95
Battery Monitor Logger	11106201	£10.95
Electronic Wind Chimes	23011201	£10.95
Analogue Vocoder – Driver Amplifier	AO-FEB22	£8.95
JANUARY 2022		
Vintage battery Radio Li-ion Power Supply	11111201	£9.95
MiniHeart: A Miniature Heartbeat Simulator	01109201	£8.95
DECEMBER 2021		
AM/FM/SW Digital Receiver	CSE200902A	£13.95
Balanced Input and Attenuator for USB CODEC	01106202	£11.95
NOVEMBER 2021		
Dual Battery Lifesaver	11111202	£6.95
OCTOBER 2021		
Mini Wi-Fi LCD BackPack	24106201	£8.95
SEPTEMBER 2021		
USB SuperCodec PCB	01106201	£14.95
Audio DDS Oscillator PCB	01110201	£5.95
Audio DDS Oscillator rotary encoder	01110201-ENC	6.95
Programming Adaptor Board for Audio DDS Oscillator	01110202	£5.95
High-power Ultrasonic Cleaner main PCB	04105201	£14.95
High-power Ultrasonic Cleaner front-panel PCB	04105202	
Night Keeper Lighthouse PCB	08110201	£5.95
AUGUST 2021		
OI' Timer PCB	19104201	£11.95
OI' Timer 8x8 RGB LED module using WS2812B	19104201-88	£8.95
OI' Timer set of acrylic case pieces and spacer	19104201-ACR	£8.75
OI' Timer DS3231 RTC IC wide SOIC-16	19104201-RTC	£5.95
Wideband Digital RF Power Meter	04106201	£9.75
Switchmode 78xx regulators ***(PACK of 5!)***	18105201	£7.95
Cool Beans SMAD display	CB-AUG21	£11.95

PROJECT	CODE	PRICE
JULY 2021		
ATtiny816 Breakout / Dev Board w/ Capacitive Touch	24110181	£9.75
IR Remote Control Assistant (Jaycar version)	15005201	£8.95
IR Remote Control Assistant (Altronics version)	15005202	£8.95
PIC18F Development Board	PNM-JUL21	£12.95
Microphone Preamplifier	AO-JUL21	£11.95
JUNE 2021		
Roadies' Test Signal Generator (surface-mount ver.)	01005201	£8.95
Roadies' Test Signal Generator (through-hole version)	01005202	£9.95
Touchscreen Wide-range RCL Box (Resistor module)	04104201	£18.95
Touchscreen Wide-range RCL Box (Ind/Cap module)	04104202	
KickStart Part 3 – Gyrator-based Audio Filter	KS3-2021	£7.95
MAY 2021		
7-Band Equaliser (Mono)	01104201	£8.95
7-Band Equaliser (Stereo)	01104202	£10.95
Car Altimeter	05105201	£7.95
APRIL 2021		
Reflow Oven – DSP Active Crossover (CPU)	01106193	£19.95
Reflow Oven – DSP Active Crossover (Front panel)	01106195	
Reflow Oven – DSP Active Crossover (LCD)	01106196	
Frequency Reference Signal Distributor	CSE200103	£8.95

For the many earlier PCBs that we stock please see the PE website: www.electronpublishing.com

PE/EPE PCB SERVICE

Order Code Project Quantity Price

.....

.....

.....

.....

.....

.....

.....

Name .....

Address .....

.....

Tel .....

Email .....

Card No .....

Valid From ..... Expiry Date .....

Card Security No .....

MasterCard VISA

You can also order PCBs by phone, email or via the shop on our website: www.electronpublishing.com

No need to cut your issue – a copy of this form is just as good!

Silicon Chip kits & parts for projects in this issue (posted from Australia):

PIC32MK0128MCA048 programmed for the Power LCR Meter https://siliconchip.au/Shop/9/7454 AUD $24 (about £12.00) + P&P

20% VAT is included in the above price (it will be removed for customers outside the UK). Postage & packing cost is typically around £12-15. Please allow 1-2 weeks for delivery (longer in some cases).

CLASSIFIED ADVERTISING



Practical Electronics

If you want your advertisements to be seen by the largest readership at the most economical price then our classified page offers excellent value.

The rate for semi-display space is £10 per centimetre high, with a minimum height of 2·5cm. All semi-display advertisements have a width of 5.5cm. The prepaid rate for classified adverts is 40p per word (minimum 12 words).

For rates, further information on display and classified advertising or to book an ad, please email us at **pe@pemag.au**

Practical Electronics reaches more UK readers than any other UK monthly hobby electronics magazine. Our sales figures prove it.

We have been the leading monthly magazine in this market for more than thirty years.

W5SWL Electronics

Technical Products and Small Component Parts, RF and Audio Connectors, Hobby Electronics, Ham Radio to Microwave Products, General Hobby Supplies, Technical Books, Tools, and More. **Located in the United States.**

Visit **www.W5SWL.com**

BOWOOD ELECTRONICS LTD

Suppliers of Electronic Components

www.bowood-electronics.co.uk

Unit 10, Boythorpe Business Park, Dock Walk, Chesterfield, Derbyshire S40 2QR. Sales: 01246 200 222

Send large letter stamp for Catalogue

COAST ELECTRONICS

BREAKOUTS-COMPONENTS-CONTRACT DESIGN-3D PRINTER PARTS-MUSICAL-MICROCONTROLLERS

WWW.COASTELECTRONICS.CO.UK

MISCELLANEOUS

VALVES AND ALLIED COMPONENTS? For free stock list and/or advice, please contact me: geoffdavies337@gmail.com Telephone: 01788 574774

PIC DEVELOPMENT KITS, DTMF kits and modules, CTCSS Encoder and Decoder/Display kits. Visit www.cstech.co.uk

Electrical Industries Charity (EIC)
We help people working in the electrical, electronics and energy community as well as their family members and retirees.

We use workplace programmes that give the industry access to financial grants and a comprehensive range of free and confidential services.

www.electricalcharity.org

electrical industries charity

ADVERTISING INDEX

ENIGMA COMPONENT SHOP. 3
HAMMOND MANUFACTURING. 49
MICROCHIP TECHNOLOGY . . . . . Outer back cover
PCBWAY. . . . . Inner front cover
PDF BACK ISSUES 44, 71, 75
TAG CONNECT 3

Advertising mailing address

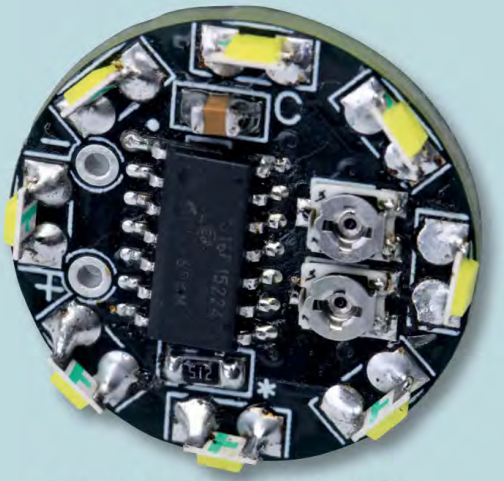
Electron Publishing Australia
PO Box 194
Matraville NSW 2036
Australia

Email pe@pemag.au

Web www.electronpublishing.com

For editorial contact details, see page 2.

Next Month – in the March issue

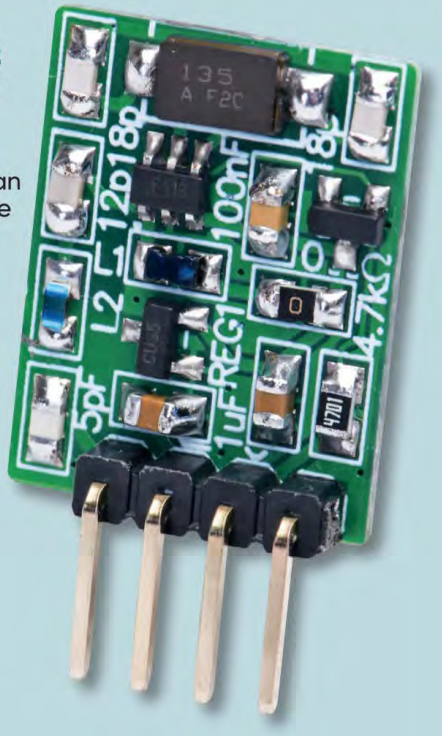


Rotating Light for Models

Make a realistic lighthouse or siren with this small board that sequences 8 LEDs in a circle.

DIY 433MHz Transmitter Module

With better performance than many pre-made modules.



USB-Programmable Frequency Divider

It has hundreds of useful division ratios and operates to over 70MHz.



On sale 5th of February, 2026

Content may be subject to change

Practical Electronics

Subscriptions hotline

To subscribe or renew online, paying by credit/debit card or PayPal, please visit our shop at: www.electronpublishing.com

For those wishing to pay by direct debit, cheque or over the telephone, please call 'Select Publisher Services' on 01202 087631, or send a cheque (payable to: 'Select Publisher Services Ltd') with your details to:

Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Remember, we print the date of the last issue of your current subscription in a box on the address sheet that comes with your copy.

Digital subscribers, please call 07918 614662 or visit: www.electronpublishing.com

Published on approximately the first Thursday of each month by Silicon Chip Publications, PO Box 194 Matraville NSW 2036 Australia. Printed in England by The Manson Group, Reynolds House, 8 Porters Wood, Valley Road Industrial Estate, St Albans AL3 6PZ. Distributed by Seymour, 86 Newman St, London W1T 3EX. Subscriptions UK: £33.99 (6mo); £59.99 (12mo.); £114.99 (2yrs). EUROPE: airmail service, £38.99 (6mo); £69.99 (12mo); £129.99 (2yrs). REST OF THE WORLD: airmail service, £46.99 (6mo); £84.99 (12mo); £164.99 (2yrs). Make payments to 'Practical Electronics', Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH United Kingdom. Email: pesubs@selectps.com. PRACTICAL ELECTRONICS is sold subject to the conditions that it shall not, without the written consent of the Publishers first having been given, be lent, resold, hired out or otherwise disposed of by way of Trade at more than the recommended selling price on the cover, and that it shall not be lent, resold, hired out or otherwise disposed of in a mutilated condition or in any unauthorised cover by way of Trade or affixed to or as part of any publication or advertising, literary or pictorial matter.

Did you know our online shop now sells the current issue of *PE* for £5.99 inc. p&p?



You read that right! We now sell the current issue of your favourite electronics magazine for exactly the same price as in the High Street, but we deliver it straight to your door – and for UK addresses we pay the postage. No need to journey into town to queue outside the newsagent. Just go to our website, set up an account in 30 seconds, order your magazine and we'll do the rest.

www.electronpublishing.com



Enabling AI at the Edge

Optimized Hardware for Machine Learning Models

Artificial Intelligence (AI) and Machine Learning (ML) are transforming industries by enabling smarter, more autonomous systems. However, deploying AI at the edge—where real-time decision-making is essential—requires highly efficient, low-power solutions tailored for embedded environments. Microchip Technology is leading the way in this transformation with a comprehensive suite of processors, FPGAs, and development tools that make it easier than ever to bring intelligence to the edge.

Across critical sectors like **smart factories, medical diagnostics, automotive systems, construction equipment, agriculture, and smart cities**, edge AI is unlocking real-time insights and automating tasks where cloud connectivity is limited or latency is unacceptable. Whether it's predictive maintenance on a factory floor, patient monitoring in a hospital room, object detection on a construction site, or energy-efficient traffic management, Microchip's solutions empower developers to deploy ML where it's needed most.

Microchip simplifies the path to edge AI integration with our **ML Development Suite**, a robust toolkit that streamlines model training, quantization, and deployment on our microcontroller and microprocessor platforms. For high-performance edge inferencing using FPGAs, our **VectorBlox® Accelerator Software Development Kit** enables fast and efficient implementation of neural networks on PolarFire® FPGAs—delivering industry-leading performance-per-watt.

Our solutions are engineered with scalability, security, and long-term reliability in mind, ensuring that edge AI systems can operate seamlessly in demanding environments. With Microchip, developers can confidently build intelligent edge devices that are efficient, responsive, and built for the future.



MICROCHIP AI

MICROCHIP

microchip.com/EdgeAI

The Microchip name and logo and the Microchip logo are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners. © 2025 Microchip Technology Inc. All rights reserved.
MEC2628A-UK-09-25