PE Practical Electronics

The UK's premier electronics and computing maker magazine

WIN!
Microchip dsPIC33CK Low Voltage Motor Control Board
WIN!
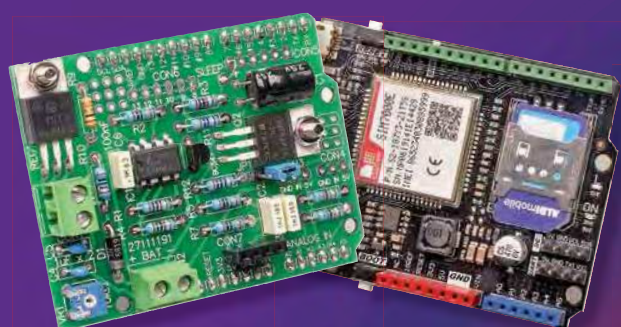
4G Arduino Remote Monitoring Station

4G

ARDUINO
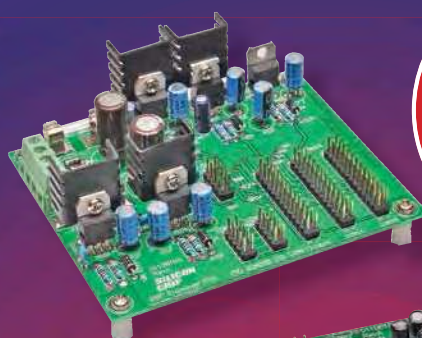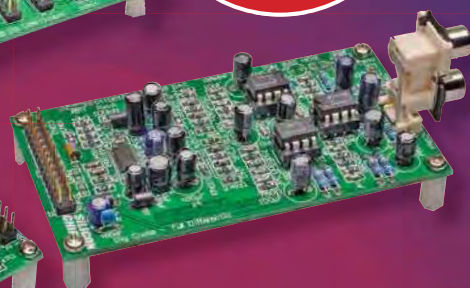
SUPERB Low-distortion DDS Signal Generator

53 ppb VOC

Indoor 'Poor Air Quality' Monitor

BARGAIN! Logic Analyser

PLUS!
Techno Talk – Dare we laugh about scams?
Cool Beans – Programming NeoPixel effects
Net Work – Updates to Amazon Fire, GoogleTV

Feb 2021 £4.99

02

9 772632 573016

MCUs Made For You

Versatile and Easy-To-Use In Any Design

Whether you're designing a smart building sensor network or the latest wearable tech, the PIC16F15244 family of MCUs provides the perfect starting point for your next design. With a 10-bit ADC and an essential feature set, these MCUs are ideal for sensor interface, simple real-time control applications. All in a small, affordable package that enables you to deploy intelligent systems at a large scale.

Key Features

- Up to 28 KB of Flash for ample program storage
- Essential peripheral set for signal generation, sensor interface, communication and real-time control applications
- 5.5V-capable for robustness in noisy and harsh environments

www.microchip.com/MCUsforYou

MICROCHIP

Contents

Projects and Circuits

Series, Features and Columns

Regulars and Services
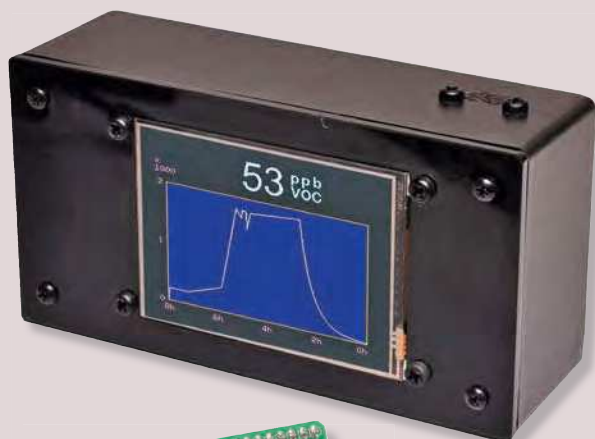
Made in the UK.

Written in Britain, Australia, the US and Ireland.

Read everywhere.

The March 2021 issue of *Practical Electronics* will be published on Thursday, 4 February 2021 – see page 72.

WIRELESS FOR THE WARRIOR

by **LOUIS MEULSTEE**

THE DEFINITIVE TECHNICAL HISTORY OF RADIO COMMUNICATION EQUIPMENT IN THE BRITISH ARMY

The *Wireless for the Warrior* books are a source of reference for the history and development of radio communication equipment used by the British Army from the very early days of wireless up to the 1960s.

The books are very detailed and include circuit diagrams, technical specifications and alignment data, technical development history, complete station lists and vehicle fitting instructions.

Volume 1 and *Volume 2* cover transmitters and transceivers used between 1932-1948. An era that starts with positive steps taken to formulate and develop a new series of wireless sets that offered great improvements over obsolete World War I pattern equipment. The other end of this timeframe saw the introduction of VHF FM and hermetically sealed equipment.

Volume 3 covers army receivers from 1932 to the late 1960s. The book not only describes receivers specifically designed for the British Army, but also the Royal Navy and RAF. Also covered: special receivers, direction finding receivers, Canadian and Australian Army receivers, commercial receivers adopted by the Army, and Army Welfare broadcast receivers.

Volume 4 covers clandestine, agent or 'spy' radio equipment, sets which were used by special forces, partisans, resistance, 'stay behind' organisations, Australian Coast Watchers and the diplomatic service. Plus, selected associated power sources, RDF and intercept receivers, bugs and radar beacons.

cebek
velleman

Official Main Dealer stocking the full range of Kits, Modules, Robots, Instruments, Tools and much, much more...

Santa's Workshop

Electronic Kits & Modules

We have a massive selection of self-assembly electronic kits and pre-assembled modules. Please see the full range on our website or call for details.

LED Buddy / LED Tester Kit
Hold any type of LED to the contact pads to see it's polarity, forward voltage & the recommended series resistor value. Adjustable target current & forward voltage. Great design aid. 1x PP3 battery powered.
Order Code: MK198 - £13.92

3-in-1 All Terrain Robot Kit
Multifunction tracked mobile robot. Transform into 3 amazing models using different track modules: Forklift, Rover (shown above) and Gripper. The wired controller commands the robot to make it move forward, backward, turn, grip or lift. Solderless assembly. Requires 4x AA batteries.
Order Code: KSR11 - £34.40

Snowman Flashing LED Kit
Have some educational festive fun with this animated LED snowman gadget. 69 multicoloured LEDs (great for soldering practice). Snowflake effect with PWM controlled LEDs. Random generator for a more realistic effect. 1x 9V PP3 battery or 9-12Vdc wall adapter powered. Makes an ideal gift.
Order Code: MK200 - £20.34

Brightdot Clock Kit - BLACK Edition
Brighten any room or space with this fully Arduino® compatible, ESP32 controlled BrightDot clock kit. This designer black edition features 60 bright RGB LEDs that reflect against the surface on which you mount the clock, hence telling you what time of day it is. ESP32 data cable & power supply included.
Order Code: K2400B - £117.43

DIY Electronic Watch Kit
Make your own DIY, Arduino compatible electronic wrist watch! 24 amber coloured LEDs are bright enough to be clearly visible in broad daylight! Pre-programmed with an addictive reflex game and of course with a basic time view. You can easily re-program it to your liking by using open-source Arduino® library and the K1201 Custom Cradle Kit or a USB to UART module (neither included).
Order Code: K1200 - £23.94

Stereo Ultrasonic Bat Detector Kit
Converts high frequency sounds (20 - 90kHz) normally imperceptible to humans like bat signals into audible noise. Can also help detect failures in machines, engines, etc. Stereo feature adds the possibility to pinpoint the source. Requires 3xAA batteries. 3.5mm jack output for headphones.
Order Code: K8118 - £21.59

LED Christmas Tree Kit
Ho! Ho! Ho! The classic Christmas kit for the budding electronics enthusiast. 15 blinking blue LEDs. Requires 1xPP3 battery. Get cosy…
Order Code: MK100B - £7.19

Digitally Controlled FM Radio Kit
Build your own modern, high quality FM receiver project with excellent sensitivity powered by a simple 9V PP3 battery (not included. Auto-seeking button. 4 station presets. Volume control. Excellent learning project for schools and colleges.
Order Code: MK194N - £20.39

Audio Analyser Display Kit
Small, compact LCD display, ideal for panel mounting. Give your homemade audio gear a high-tech look. Upgrade existing equipment. Provides Peak Power, RMS Power, Mean dB, Peak dB, Linear Audio Spectrum And 1/3 Octave Audio Spectrum. Auto / Manual range selection. Peak-hold function. Speaker impedance selection.
Order Code: K8098 - £39.54
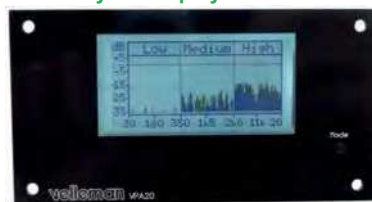
Electronic Component Tester Kit
Build your own versatile component tester. Shows value and pin layout information for resistors (0.1 Ohm resolution, max. 50 MOhm), coils (0.01mH - 20H), capacitors (28pF - 100mF), diodes, BJT, JFET, E-IGBT, D-IGBT, E-MOS & D-MOS.
Order Code: K8115 - £44.34

LCD Oscilloscope Educational Kit
Build your own LCD oscilloscope with this exciting new kit. Learn how to read signals. See the electronic signals you learn about displayed on your own LCD oscilloscope. Despite the low cost, this oscilloscope kit has a lot of features found on expensive units like signal markers, frequency, dB, true RMS readouts and more. A powerful auto-setup function will get you going in a flash!
Order Code: EDU08 - £48.54

Card Sales & Enquiries | 01279 467799
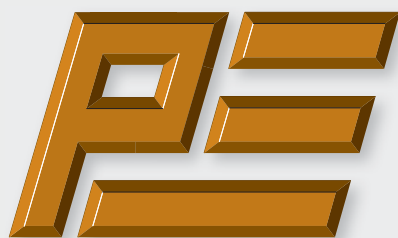
Practical Electronics

UK readers SAVE £1 on every issue

Take out a one-year subscription and save **£10** over the year.
Even better – save **£1** per issue if you subscribe for two years – a total saving of **£24**.

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 the newsagents sell 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 £26.99, 12-months £49.85, 24-months £94.99

Europe Airmail: 6-months £30.99, 12-months £57.99, 24-months £109.99

Rest Of The World Airmail: 6-months £37.99, 12-months £70.99, 24-months £135.99

Cheques (in **£ sterling only**) payable to *Practical Electronics* 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 £24.99 for 12 months, more details at: **www.electronpublishing.com**

SUBSCRIPTION ORDER FORM

☐ **6 Months:** UK £26.99, Europe £30.99 (Airmail),
Rest of the World £37.99 (Airmail)

☐ **1 Year:** UK £49.85, Europe £57.99 (Airmail),
Rest of the World £70.99 (Airmail)

☐ **2 Years:** UK £94.99, Europe £109.99 (Airmail),
Rest of the World £135.99 (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 *Practical Electronics*

☐ 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.

Editorial

Sourcing parts

When it comes to projects, if there is one thing that can drive me – and doubtless you too – to distraction, it is sourcing parts. It's an endless game of cat and mouse. First you have to find the bits and pieces you need, then work out who does what the cheapest, next calculate if X is really cheaper than Y after shipping is factored in… then realise X is only cheaper if you buy 100… but Y is out of stock, so in the end, over-priced Z gets the order because they actually have some in stock, even though they charge an unconscionable amount for putting a single capacitor in an envelope. And so it goes on.

I'm thinking about parts because I recently had a very useful email conversation with subscriber Geraint Williams, who quite reasonably asked why we regularly list Jaycar as a supplier for *PE* projects. This made me crystalise my reasoning, which I will now share with you. For those of you who are wondering what all the fuss is about, *PE* works in partnership with *Silicon Chip* (*SC*) magazine in Australia. We publish a selection of their projects, which are mostly sourced in Australia, and not surprisingly, their authors sometimes quote parts from local suppliers, especially Jaycar (and Altronics).

We spend a lot of time going through the parts lists of *SC* projects to make sure they are UK friendly. Most parts are so generic that neither *SC* nor *PE* will bother to tell you where to go (resistors, capacitors, common silicon devices and connectors). Such items are widely available and you don't need our help. But – for the larger projects there can be obscure ICs, strange inductors, very specific hardware (boxes or switches) and so we do check that such items can be sourced from a UK perspective. Often, in fact the majority of the time, the references *SC* provide for these trickier parts are from international suppliers like Farnell/element14, Digi-key or Mouser, so they are just as easy to buy in the UK as in Australia. Sometimes, though, the reference is Jaycar, which brings us to Geraint's question – and my answer.

There are three reasons for a UK-based magazine to keep the Jaycar references in its projects, despite the apparent geographical disadvantage. First, some readers simply want to buy the author-recommended parts and since the articles are supplied with precisely that, it would be foolish not to include it.

Second, buying from Australia is not as expensive or slow as you might think. Yes, there is a premium, but if you purchase selectively and possibly for several projects in one hit then it really isn't such an inconvenience. And of course, air freight makes a huge difference. (Don't forget that unfortunately you may get hit with VAT at UK customs.)

Third, and by far the most important reason, Jaycar's website gives a pretty good definition and description of the parts in question. So, if readers wish to look elsewhere for a part then they can do so with a good technical spec and hence a fair degree of confidence that they are looking for the right part.

Keep well everyone

Matt Pulzer
Publisher

The **Fox** *Report*

Barry Fox's technology column

Smart meters follow-up

The contentious TV adverts for smart meters (see *Fox Report*, December 2020), which suggest that putting a smart meter under your stairs will help the UK have more wind turbines, were followed by full-page newspaper adverts pushing the non-sequitur even further by picturing a wind turbine and claiming that: 'by having that odd, plastic box (a smart meter) you'll be helping Britain's energy system use more renewable resources… meaning more wind and solar, and less coal'. (See right, the advert in London's *Evening Standard*, 16 November 2020.)

As previously reported, I complained to the UK's ASA (Advertising Standards Authority) about the TV adverts and followed through with a fresh complaint about the newspaper ads.

I think both complaints have been rejected. But why can I only 'think' that, rather than know so?

No reply or explanation

It's because the ASA has now adopted a new policy for dealing with all complaints. If your complaint is rejected, you simply hear nothing. And you don't know what length of nothing heard means rejection, because the ASA does not give a timescale to help decode silence. Neither does the ASA offer any way for complainants to check progress or non-progress.

What the complainant does get, though, is a message from market research company IPSOS, hired by the ASA to follow through complaints with a customer satisfaction survey. But of course, complainants cannot usefully complete the survey without knowing whether they are still waiting for news from the ASA, or not.

The Covid crisis has made life difficult for everyone, and many companies and public service bodies have done a remarkable job of coping. Less so, it seems, the ASA.
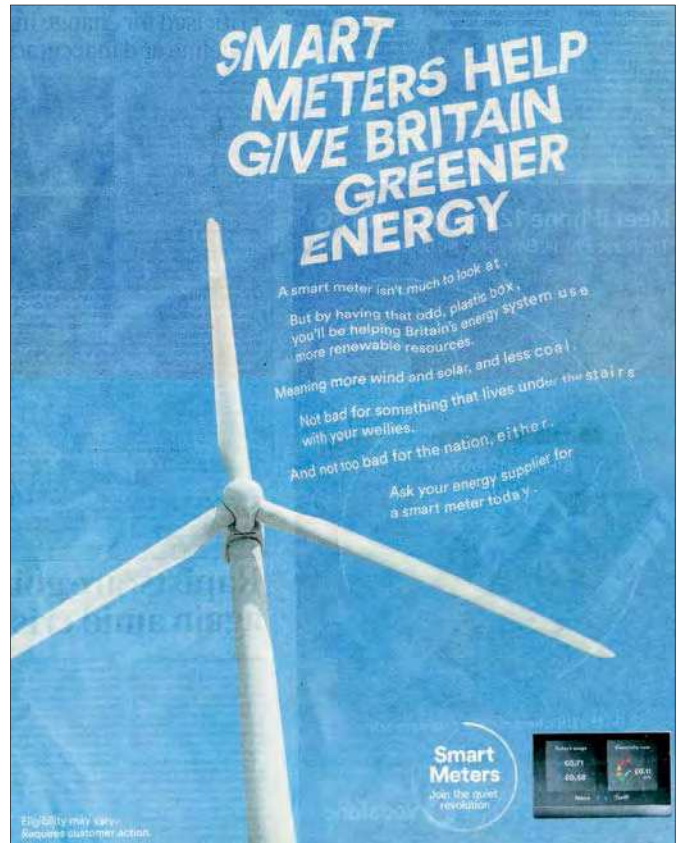
Legal, decent, honest and truthful? – the ASA seem to be shy about stating its position.

Rejected

Only because I asked for press comment, I can now report that the ASA rejected my complaint about the TV adverts because it had previously rejected complaints on earlier adverts (April 2019) for smart meters that included the honest acknowledgement that 'smart meters can't solve climate change on their own but with the smarter, more energy efficient grid they help to create, they're a start'.
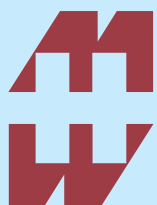
The ASA had said: 'We understood that the data captured and reported by smart meters was intended to help the national grid to anticipate demand and plan energy provision more effectively, with less reliance on fossil fuels. We acknowledged that take up of smart meters was one of many factors that contributed to more efficient energy. Nevertheless, we considered the points Smart Energy GB had made, and the BEIS, the Committee on Climate Change and Ofgem comments that they had cited, supported the view that take up of smart meters would play a part, and would therefore contribute to the environmental benefits associated with more efficient energy use'.

Emboldened by the ASA's previous baptism, the new ads even more brazenly link wind power with smart meters, picturing a wind turbine and saying, 'by having that odd, plastic box (a smart meter) you'll be helping Britain's energy system use more renewable resources… meaning more wind and solar, and less coal' (**https://bit.ly/3pe-feb21-asa**).

'very helpful'

An ASA spokesperson says my rocket science suggestion that the ASA should at least give complainants a timescale clue to when they should assume their complaint has been rejected is a 'very helpful suggestion (and) something that struck a chord immediately and we will look to implement something like this.'

But as it is now some three weeks since the ASA acknowledged my new complaint, and I've heard nothing, I can only deduce that: a) the complaint has been silently rejected and b) my 'very helpful suggestion' has not been implemented.

I've now asked the ASA's chief executive Guy Parker how it came to pass that on his watch the body which is paid big bucks to protect us from adverts that aren't 'legal, decent, honest and truthful' judged it decent to leave complainants dangling with no idea whether their complaint has been understood and is being acted on, or has been misunderstood, refuted, rejected or just plain lost in the system.

I await his reply with interest – but not much hope – just as I await a judgment on whether the smart meter under my stairs really is the reason that there are now so many more wind turbines working for us.

Dare we laugh about scams?

Techno Talk

Mark Nelson

This month, we turn to online scams involving electronic gizmos. To *PE* readers, these scams are self-evidently fraudulent; they're totally ludicrous and could never work as described. However, non-technical customers might easily take the scammers at their word – otherwise the fraudsters would not waste their time flogging their rubbish. If you get a kick from pseudo-scientific nonsense, just read on.

Do you enjoy a good one-liner? My all-time favourite is 'Start every day with a smile – and get it over with.' I chuckle every time I hear it; and for this absolute gem we must thank the American actor and comedian WC Fields (1898–1946). Each of his witticisms was a cracker, such as my second favourite, 'There's no such thing as a tough child – if you parboil them first for seven hours, they always come out tender.'

That's the darker side of his humour, and as we launch into the serious side of this article, here are just two more. 'It's morally wrong to allow a sucker to keep his money', also 'You can fool some of the people some of the time – and that's enough to make a decent living.'

FUD and FOMO

FUD is a combination of 'fear, uncertainty and doubt', while FOMO is a 'fear of missing out'. Both are powerful motivators in marketing methods, to convince punters into spending money on PDPs (potentially dodgy products) that they don't need. Here's one to consider: the Blushield Tesla Cube (**https://bit.ly/pe-feb21-bs**). It's not something you drop in the toilet cistern and at £480, it's not as affordable as a blue brick for the WC. Described as an EMF Protection device, it's an electronic gadget for home and office use that protects the human body against 5G radiation.

As if simply shielding you against microwave radiation were not enough, it also claims to promote emotional stability and restful sleep, while it 'may' also reduce fatigue associated with EMF and increase your energy level. Helpfully, the write-up also advises: 'Spring mattresses induce EMF through the coils and can even amplify EMF fields nearby such as wiring in walls. Get a latex or foam mattress and make sure your head is not near any electrical outlet or wiring near the wall.'

Smart meters are 'extremely dangerous', it continues. and Wi-Fi routers and DECT cordless phones are 'very dangerous'. I now realise how little my electronics education had taught me about how physics really works. Never again will I trust 'experts'! But the people who sent in these, umm, remarkable testimonials (**https://bit.ly/pe-feb21-more-bs**) are adamant that their lives have been changed as well as having their wallets emptied. Blushield can even make your neighbours' dogs less hostile, after all.

All hail the Accelerator Frequency Generator

Just as the coronavirus pandemic has stimulated research into multiple anti-COVID vaccines, several prophylactic measures have been devised for combatting 5G radiation. So, if Blushield is not your cup of tea, maybe you need an Accelerator Frequency Generator (AFG). Unlike the Blushield product mentioned above, the AFG Smart 5G Shield is not mains-powered, using instead its attached solar cell chip is the power source. This makes it much cheaper to buy ($39 from **https://bit.ly/pe-feb21-different-bs** – Amazon UK did sell it but has run out of stock). By the way, don't be taken in by the mean-spirited customer reviews at **https://bit.ly/pe-feb21-donotbuy**; I am quite sure that Amazon would never sell a product that was 'ineffective', 'pure snake oil' and 'possibly haunted'.

But how does it work? By exploiting the effects of frequency acceleration, obviously. However, for reasons of commercial confidentiality, all the manufacturer is prepared to reveal for now is that the product is 'infused with anti-radiation frequencies from the AFG that aids in protecting you from harmful RF radiation for one year'. The form factor is fortunately very modest: a one-inch round chip made of a crystalline carbon material with a strong adhesive back that adheres to any surface. It is designed to block harmful RF radiation from common sources such as 4G and 5G cellphones, tablets, laptops and all other electronic devices.

The ultimate scam

Perhaps the prize for the most notorious 5G scam product should go to the fabulous 5G BioShield. Despite being denounced by the City of London Police's Action Fraud Squad and UK Trading Standards, it is still on sale, now reduced by 40% to £165 (**https://bit.ly/pe-feb21-home-of-bs**). Unlike rival products, it employs a quantum nanolayer as a portable radiation balancing system to secure the ultimate protection for you, your family, and your home. It's ingenious too. Without being plugged into a computer, the device actively protects the human body 24/7 within an 8m-diameter sphere, but when plugged into a power source, the sphere of active protection expands from the original 8m up to a massive 40m in diameter, protecting all in its field, both individuals and whole households. No larger than a USB stick, it can be carried on your person or placed near a mobile phone (or any other radiation-producing electrical device).

The reason why this amazing device is no larger than a memory stick is down to the fact that it is, well, just an unenhanced memory stick worth about £5. Two scientific teardown tests have proved this. The last word goes to the splendid Jim Al-Khalili OBE FRS, professor of physics at the University of Surrey and presenter of BBC Radio's *The Life Scientific* programme, who tweeted: 'As a renowned world expert on quantum biology, quantum entanglement and relativistic time dilation, I can say, categorically, that if you spend £339 on a 5GBioShield then you're an utter numpty.'

And finally…

A quick follow-up to last month's article about dodgy 'deco' lightbulbs. The editor helpfully reminds me that the CE symbol sometimes has a third meaning: c*aveat emptor*. Thanks, Matt! Many thanks go also to fellow contributor Alan Winstanley for launching me on this journey of investigation.

Win a Microchip dsPIC33CK Low Voltage Motor Control Board

Practical Electronics is offering its readers the chance to win a dsPIC33CK Low Voltage Motor Control Board (DM330031) – and even if you don't win, receive a 20%-off voucher, plus free shipping for one of these products.

The dsPIC33CK Low Voltage Motor Control (LVMC) Development Board is a cost-effective rapid development platform for brushless DC (BLDC), permanent magnet synchronous motor (PMSM) and internal permanent magnet (IPM) motor control applications. The LVMC development board is ideal to explore and prototype motor control applications that operate from 12 to 48V and up to 10A of continuous current. The board can also support higher current, up to 20A with an optional fan or heatsink.

The board supports application development for sensored six-step trapezoidal or sinusoidal control with hall sensors,

sensorless BEMF-measuring phase voltages, or field-oriented control (FOC) with single, dual or triple-shunt resistor current-sensing circuits. This flexibility in architecture allows you to evaluate various design techniques that suit your application requirements, balancing the cost and performance.

The LVMC board features the high-performance, single-core dsPIC33CK256MP506 Digital Signal Controller (DSC) with 256KB Flash, 8-pairs of advanced motor control PWMs, and three integrated 3.5 Msps ADCs for implementing maths-intensive motor control algorithms, such as sensorless FOC, initial-position detection (IPD) or field weakening (FW). The LVMC board comes with two mikroBUS headers to expand the board features to meet your application needs. It features a high-speed USB-to-UART interface for debug purposes and to monitor key parameters in your control algorithm using a graphical PC tool like X2C Scope or RTDM.

Worth $95.00 (approx £72.00) each

Free-to-enter competition
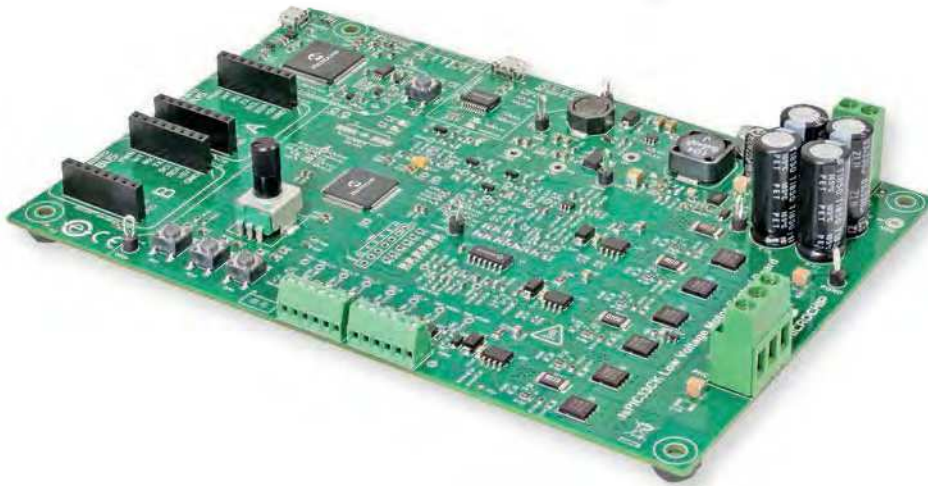
Microchip dsPIC33CK Low Voltage Motor Control Board

How to enter
For your chance to win a Microchip dsPIC33CK Low Voltage Motor Control Board or receive a 20%–off voucher, including free shipping, enter your details in the online entry form at:

https://page.microchip.com/PE–dsPIC.html

Closing date
The closing date for this offer is 31 January 2021

October 2020 winner
Ian Williams

He won a Microchip MPLAB PICkit 4 In-Cirucit Debugger

Net Work

Alan Winstanley

This month, *Net Work* looks at the latest updates to Amazon Fire, GoogleTV/Chromecast and Devolo's Wi–Fi extender; plus the charging trials of being an EV early adopter.

Consumer demand for Internet-based services is soaring in Britain, and the race is on, somewhat belatedly compared with some other countries, to deliver fibre-based Internet comms directly to the home (FTTH). Even in the author's rural location, village roads are being dug up by fibre providers eager to 'pipe' high-speed data directly into our homes and offices. About 20 years earlier, Diamond Cable (later, NTL then Virgin Media) was rolling out cable TV for the first time and coax cable was ducted underneath the same pavements in the expectation of delivering cable TV to towns and villages in the future. Virgin Media would offer a hybrid fibre coax (HFC) network followed by its £3bn Project Lightning superfast fibre network, but for funding and technical reasons much of the original cable network was never switched on locally: the cable simply ran straight through towns, but the scale of civil works made it non-viable to connect every home in every street to cable TV, and coax cable would be no substitute for FTTH services anyway. Those legacy 'CATV' access covers, which perforate thousands of miles of pavements, bear witness to an early endeavour, but thanks to a new crop of service providers homeowners are finally reaping the benefits of superfast fibre broadband fed directly into their property.

As predicted some years ago, British Telecom (BT) has now moved its phone line subscribers to flat-rate quarterly tariffs that bundle, say, 500 or 700 'free'

Virgin Media Project Lightning engineers lay fibre optic cable (Image: Virgin Media)

minutes. Thanks to cellphones, increasing numbers of households don't use a landline at all anymore. Eventually, as networks are switched over to optical fibre to carry voice calls and streaming media, many local telephone exchanges will become redundant, along with their copper-wire networks.

Content is king

In the early 2000s I attended a trade event where BT was the guest speaker, and BT implored web designers like us to exploit the then-new ADSL services by adding more bandwidth-intensive media than had previously been feasible with dial-up. The theory was, that by tantalising web users with more immersive content, this would stimulate the demand for ADSL. As FTTC and FTTH services continue to expand, connectivity is becoming a commodity and content is still king. Major telcos and entertainment providers such as BT Sport, Netflix and Amazon are falling over themselves to feed our insatiable appetites for more streaming media. Netflix was recently roundly criticised over its series *The Crown*, which depicts life in the Royal Family. The controversial series has drawn fire in Britain from royal friends, commentators and politicians who worry that non-discerning viewers might actually believe that the heavily dramatised and fictional programme is really true! Netflix rejected calls to include a disclaimer in the opening titles.

Fire-powered TV

Amazon continues to add more streaming content as part of its Prime package, including a vast array of movies (some free, some paid-for), specially commissioned series, catch-up TV and millions of songs for free, coupled with unlimited cloud-based photo storage and, of course, they bundle in free next-day deliveries on most product lines.

It's easier than ever to access streaming video on an ordinary 'non smart' TV with a spare HDMI port, and a

Amazon's Fire TV Stick HD dongle in situ, fitted with an HDMI extension. The 5V power cord hangs separately: you could try powering it from a spare TV USB port.

number of platforms are all jostling to lock you into their services. Google's new Chromecast (plus remote control) with Google TV (see last month) has been soft-launched and, at the time of writing, is stubbornly stuck at £59.99 online. It could be a good choice for Android users.

During last November's Black Friday, I spotted that rival Amazon had dropped the price of their new-generation Fire TV Stick devices to £24.99 (usually £39.99), making them a viable way of upgrading a couple of TVs. Included with the HDMI dongles are a remote control with microphone that interacts with Amazon's Alexa at the push of a button (no eavesdropping here!). Note the cheapest Amazon 'Lite' Fire TV (normally £29.99) omits the Alexa and TV power/volume buttons, while the costlier 4K UHD type is £49.99.

Setting up an Amazon Fire TV Stick on an HD TV was impressively simple. Amazon warns that only the supplied mains adaptor (5.25V 1.0A) mains adaptor should be used, but as it's a USB type I tried a spare USB port on the TV itself (well you would, wouldn't you). In

any case, the Fire TV dongle checks the power supply during setup and warns if insufficient current is available. Regardless, the setup allowed me to continue using my TV's USB port anyway – this idea worked with one TV but not the other. A sure sign of power problems is when the Fire TV dongle keeps rebooting, which the official adaptor would cure, but I also tried a powerbank instead with success. A short HDMI extension is also provided if HDMI ports are inaccessible on the TV chassis (see photo). Once powered up, after entering Wi-Fi and Amazon logins the Fire TV updated its firmware and it was easy enough to install popular streaming TV apps. Users of Amazon Photos can also set up a TV screensaver that rotates their favourite shots with various blend effects.

Generally, there was nothing at all to dislike about Amazon's Fire TV Stick: setup was smooth and flawless and, of course, it picked up watchlists and Amazon movies that had been 'paused' on other devices. This model's remote control also includes infra-red control that learns the TV's power and volume codes (which worked perfectly on both TVs). Plus, Alexa can now broadcast via the TV, although I found it hard to discern its voice on a flat screen's tinny speakers.

Overall, Amazon has got it right and the Fire TV upgrade will add a new lease of life to any TV with a spare HDMI port. A cheaper choice for Android fans might be a classic Chromecast, controlled via apps, or if you want to 'mirror' your tablet or phone onto a TV screen then a multitude of screen-player dongles is available (search online for Miracast display dongle or similar).

Candid camera

In recent columns I suggested some ideas for IP security cameras, most of which require a dedicated app from Google Play or the App Store to view footage or manage the camera's setup. They may offer (paid-for) cloud storage too, and hopefully the apps will be updated by their designers when Android or iOS (or Windows) get future upgrades. The HeimVision HM311 3MP Outdoor Security Camera (see *Net Work*, December 2020) was spotted at under £25 during Black Friday, which would make it a bargain. A dedicated UCam 247 HD Wi-Fi bullet camera has lasted for five years on the author's network and is accessed via an app or web browser, but with a total of three of my TVs now sporting Amazon Fire TV dongles, could I check the camera on a TV screen as well? One idea was to try Onvifer, which claims to be the world's first Android app supporting IP cameras compatible with ONVIF (originally, the Open Network Video Interface Forum)

standards. Although the Wi-Fi camera made no claims about ONFIV-compatibility, happily the Onvifer app did connect to it and streamed the video on an older 7-inch tablet. (Separately, as a sign of the times, the UK importers of this good quality IP camera tell me that it is now unavailable due to production problems and sadly it's not clear if they can offer bullet cameras again.)

As it turns out, Onvifer is also available in Amazon's app store (ASIN B007OW-2WZI) for direct download onto Fire Sticks, so a demo version was soon up and running on a Fire-powered TV. Again, it's potluck whether a camera is compatible, but it worked in the author's case though the app's choice of image resolution had to be downsized (the available image resolution is set in the camera app). Also, do select 'Transport by TCP' instead of the default HTML. The app also offers multi-screen viewing and other options. The paid-for Pro version of Onvifer costs just £3.79 per platform, so Fire Stick and Android phone users will need to buy it twice.

Devolo's new Wi-Fi Repeater

When Wi-Fi coverage is a problem around the home and it's not feasible to run Ethernet cable everywhere, one workaround is to use powerline communications (PLC) to transmit data through the ring mains. This can be very convenient, especially if the mains wiring is of decent quality and not electrically 'noisy': Devolo PLC network products have been covered in the past and they claim a range of 500m for their 'Magic' adaptors. Other brands, including TP-Link and Netgear, also offer PLC ranges. Devolo recently sent *PE* their new Wi-Fi Repeater+ ac to try. (NB this model 8703 is *not* a PLC adaptor as it doesn't carry mains-borne data: it's to extend Wi-Fi coverage instead.)

Devolo says its new 802.11ac repeater offers a wireless LAN with a maximum data rate of 1200 Mbps and crossband repeating for faster performance, utilising both the 2.4GHz and 5GHz frequency bands. Built-in beamforming through four antennae targets terminal devices for better data transmission, and wireless Mesh operation provides seamless networking around the property [ie, whole home coverage, when used with other suitable Devolo Magic wireless adaptors]. Apart from Wi-Fi, there are two 10/100 Ethernet ports to connect static equipment. The model supplied is larger than many plug-in repeaters, but it has a pass-through mains socket, and it's worth noting that Devolo configures it for British households where mains outlets are frequently at floor level. Hence the UK mains plug is near the *bottom* of the ventilated housing and the Ethernet

ports sprout from the top. I noticed that the continental versions with Schuko sockets are the other way round.

When installed in default 'repeater' mode it will join and extend an existing wireless network, or it can be set as a wireless access point instead, which needs an Ethernet cable to hook to the router. As readers doubtless know, devices like these often take either five minutes or five hours to install and the printed setup guide appeared very simple. The repeater can be installed using Devolo's Home App on, for example, a smartphone. Simply plug into the mains and change the phone's network SSID to the one broadcast by the Wi-Fi repeater. After a minute, a list of available Wi-Fi networks appears – choose which one to extend (in my case, either my router's Wi-Fi, or my PLC-based wireless network), type in the password and eventually the Home and Wi-Fi LEDs will indicate success. The setup initially performed flawlessly and after a couple of minutes the repeater rewarded me with two solid white LEDs. Push-button or PIN number WPS setup is also offered, something I never find successful and I could not make it work this time either, perhaps due to my router.

As a test I could connect tablets, phones and the two Amazon Firestick TV dongles mentioned earlier, and everything worked very well indoors. My enthusiasm was

Devolo's Wi-Fi ac Repeater+ has a pass-through mains socket and extends existing wireless networks by up to 1200Mbps.

dimmed, though, as next day I found the repeater had disconnected and was stubbornly blazing a red LED for 'pairing mode'. There followed the usual frustrations of downloading manuals, FAQs, reboots and factory resets. After some perseverance the repeater was finally up and running again. A white/red blinking LED indicated a sub-optimal connection to a distant router, but it has operated reliably ever since.

Configuration can be accessed more easily using Devolo's Cockpit software, which worked well on a W10 PC: a new 'Next network' button in the software threw me initially before the new repeater made an appearance on screen. Compared with some arcane setups found with other repeaters, this one is very good indeed.

If you already have a Devolo Magic Wi-Fi system, or any other wireless network that you wish to extend, then the new Devolo Wi-Fi Repeater+ ac is a simple and elegant extender solution with built-in features that improve coverage and reliability. Some might use it as the first stage of gradually upgrading to a Devolo Magic Mesh whole-home network. The repeater is listed by CPC (order code CS33940, £57.90 inc VAT). More details are at: **https://bit.ly/pe-feb21-dev** and the useful 'missing manual' is available from: **https://tinyurl.com/yy83yg4n**

Driven to despair

In a hands-on road test, an English couple who own an electric Porsche Taycan sports car (a snip at £83,000 or $107,000) reportedly embarked on a 130-mile journey home in their battery electric vehicle, hoping to stop for a top-up charge somewhere along the way. In the event, the journey took nine hours – an average of 14 mph – as they struggled to find any suitable charging points or even one that actually worked. Eventually, they found a working, available fast charger and limped home with 11% capacity to spare. Next day, they located three charging points in town but none of them was working either.

The availability of electric vehicle charging points in Britain is fairly dire, but with petrol and diesel cars falling out of favour with environmentalists and the government, fossil-fuelled cars will probably be taxed off the road or petrol pumps themselves will become scarce, overtaken by EV charging points instead. It will become the norm to drive to a supermarket and recharge for free while popping inside to shop. At the moment, poor chargepoint availability (less than 20,000 across the UK) is a major obstacle towards adopting electric vehicles (cost is another), and the UK Government's Competition and Markets Authority (CMA) has launched a market study into the EV charging market, to make sure this new and fast-growing sector works well for UK drivers, they say. The National Grid reckons there could be between 2.7 and 10.6 million EVs on Britain's roads by 2030, and could reach 36 million by the year 2040. Today there are about 38 million cars on Britain's roads, says the RAC (Royal Automobile Club).

Britain's first all-electric EV charging station has now opened in Braintree, Essex which operator Gridserve claims to be the world's first 'Electric Forecourt'. It charges up to 36 cars at a time at a cost of 24p (0.32¢) per kWh, with AC and DC chargers ranging from 22kW to 350kW. Services for buses and trucks will also be offered. More information is at: **https://gridserve.com/braintree-overview/**. One hopes that electric forecourts work better than their website: the newsroom and press office web pages had blown a fuse and were 404 when checked.

Other news

France has pressed ahead with a 3% Digital Sales Tax (DST) after a temporary stay since announcing it last year. The tax, aimed at major digital operators including Facebook, Amazon, Google and eBay, is levied on sales of 'soft' digital services such as search engine advertising and online marketplace fees. Until an internationally agreed tax model arrives, several countries are targeting these primarily US digital services to bolster falling tax revenues. The UK imposed its own 2% DST last April (eBay is swallowing the cost, while Amazon is passing it on to traders). In Britain, Facebook UK's last filed accounts showed a gross turnover of £1.6 billion, a net operating profit of £96 million and a £28 million tax bill. Last year America's US Trade Representative (USTR) policy body started an investigation of its own, citing the supposedly negative and unfair impact that DSTs may have on these American tech giants.

In Europe, Nokia, Ericsson and others have embarked on a joint project called Hexa-X to research the next level of 6G mobile connectivity (see **https://hexa-x.eu/**). Meantime Finland has joined the list of countries, including Australia, banning China's Huawei and ZTE from its 5G infrastructure. The UK has also changed its mind and is removing Huawei hardware from its networks, a task that will take seven years to accomplish. After trying to keep Huawei's 5G network security 'manageable', the UK National Cyber Security Centre (NCSC) had deep concerns about Huawei's future system integrity since US component and chip manufacturing sanctions had impacted on supply chains. Meantime, despite the apparent trade ban, Google apps on UK Huawei phones continue to be updated, for now anyway, but future Huawei phones are destined to use their Harmony OS rather than Google's Android. The author's phone is backing up to a well-run Huawei Cloud as usual, with 50GB storage costing about £8 a year.

Several Chinese-made video doorbells have been taken off sale after UK consumer body *Which?* discovered they could be easily hacked, or they sent unencrypted usernames or passwords data back to China. All eleven doorbells posed risks, says *Which?*, who criticised the Victure brand in particular. More details at **https://tinyurl.com/y3fztu7z**

Micro-satellite maker Swarm (**www.swarm.space**) has launched 24 more ¼U-sized SpaceBEE satellites (*Net Work*, October 2020) as part of its low-cost network aimed purely at serving IoT devices. Part government-owned satellite network OneWeb (**www.oneweb.world**) which had been posited as Britain's post-Brexit satnav replacement for the EU Galileo GPS system, has emerged from bankruptcy and is set to resume satellite launches. OneWeb aims to lob more than 600 of them into low-earth orbit. The UK government will likely bid *adieu* to GPS and Galileo as, moving forward, the UK's Space-Based Positioning Navigation and Timing Programme (SBPP) is exploring alternatives for building a home-grown satnav and timing network instead of relying on foreign sources.

See you next month for more *Net Work!*

The author can be reached at: **alan@epemag.net**

Britain's first all-electric EV charging station in Essex which Gridserve claims to be the world's first 'Electric Forecourt'.

Introducing the new
PicoVNA™ 108

See down deeper to -124 dB

See out further to 8.5 GHz

Only **$8255**

Vector Network Analyzer
PicoVNA 108 300 kHz - 8.5 GHz Quad RX

Port 1

S21

S11   S22

S12

DC bias, 15 V
250 mA max.

Port 2

DC bias, 15 V
250 mA max.

Visit www.picotech.com/A753 to discover **MORE**

Remote Monitoring Station

4G

If you have an expensive car, boat, caravan, holiday house... virtually anything... you need to know what's going on when you are away. Is the battery going flat? Is your boat taking on water? You need to find out about these ASAP. All you need to do this is a couple of Arduino shields and a little software. You can even remotely trigger actions and avoid a serious problem!

We have to admit that the *raison d'être* for this project originally had nothing to do with monitoring expensive cars or boats, remote holiday houses, farm water tanks or anything so esoteric. It was all to do with wombats (see right).

For the benefit of non-Australian readers, wombats a somewhat threatened species, are cute, (usually) slow-moving furry animals that inhabit the Australian bush. A colleague of mine actually lives way out in the bush and is a member of his local wombat protection society.

Part of their remit is to re-home wombats in areas where they are less likely to be attacked by other animals (eg, dogs). They do this by trapping them and relocating them.

The problem is/was that wombats are very easily stressed and will die if they are trapped for too long.

What was needed was a means of letting society members know, as soon as possible, that one of their traps had been sprung... and our *Wonderful Wombat Warning Whatchamacallit* is the result. Of course, what *you* use it for is entirely up to you!

2G, 3G and 4G
In my part of the world, the 2G (GSM) mobile network has already been essentially shut down, and some telcos are starting to threaten to shut down their 3G network. So to do this kind of job reliably for the next few years at least, you need a 4G device.

When we found a locally-available Arduino 4G shield at a reasonable price, we jumped at the opportunity to design a *Remote Monitoring Station* around it.

Since this Station is based on an Arduino board, it can be easily programmed to suit your specific requirements. It can monitor the state of switches, voltages, sensors – just about anything, provided there is an Arduino library to interface with it (and there usually is).

Similarly, you can send commands to the Arduino from your mobile phone or PC to do things like switch mains power on or off, using a simple add-on device, such as our *Opto-isolated Mains Relay* (*PE*, October 2019).

This is an Arduino-based project, so naturally you'll need to be familiar with the Arduino IDE (integrated development environment) software. This is a free download available from: **www.arduino.cc/en/software**

4G Shield
This project revolves around a 4G Shield designed by DFRobot. It is

by Tim Blythman

based on a SIMCom SIM7000E module, which provides the 4G capability. Its circuit diagram is shown in Fig.1.

The SIM7000E module is fed power from the Arduino's VIN pin via an MP2307 buck regulator. This produces a steady 3.3V with at least 4.75V at its input.

While this eliminates the option of powering the shield from a 3.7V lithium-ion or LiPo cell, it will work with most Arduino applications powered from the VIN pin or DC barrel socket.

Below the regulator is the power control section. The PWRKEY pin on the SIM7000E is pulled low to signal that it should power on or off.

Pushbutton S1 connects this pin momentarily to ground, while NPN transistor Q1 allows Arduino pin D12 (driven high) to achieve the same effect.

Communication between the host Arduino and the shield is with a serial TX/RX pair, via level-shifting transistors Q8 and Q9. Slide switch S2 routes the signals to either D0/D1 (which is usually a hardware serial port on Arduino boards) or D8/D7 on the Arduino.

The SIM7000E's USB port is broken out to a micro-USB connector. This does not supply power to the shield, but can be used by a PC to communicate with the SIM7000E module.

We didn't investigate this in detail, but it appears that many features of the module are useable via the USB connection. It may even be able to act as a USB 4G modem.

There are also sockets for a full-sized SIM card, 4G antenna and GNSS (Global Navigation Satellite System) antenna, which is used for GPS and GLONASS.

The completed assembly is compact, needing only a few flying leads. Even the smallest 12V SLA battery dwarfs it. The Arduino board (here, a Duinotech Leonardo) is on the bottom, the SIM7000E board in the middle and our power control board on top. The antennas for both the GPS module and the main antenna (left) should be mounted with a clear view of the sky for best reception.
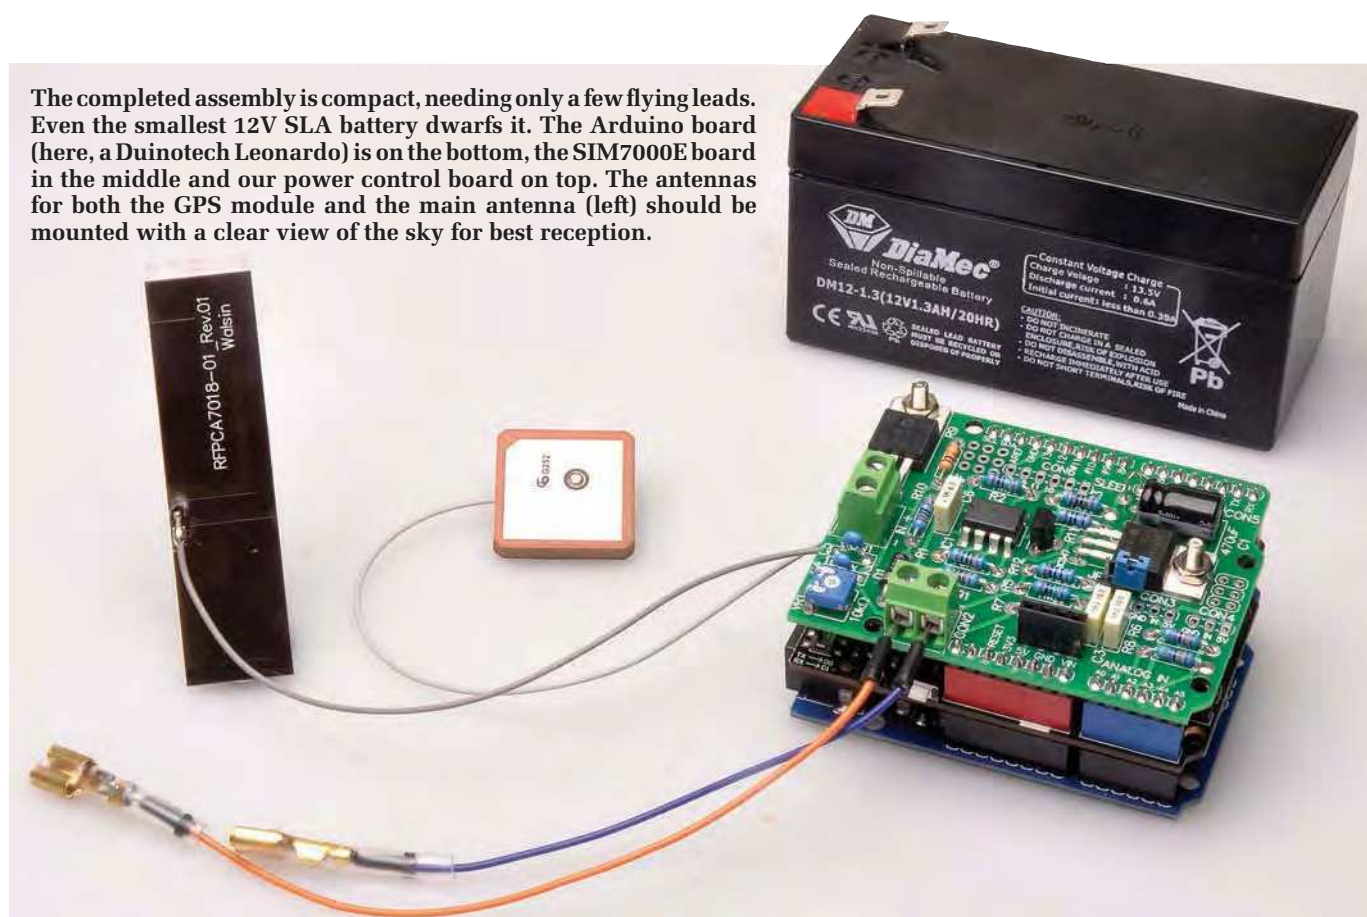
One common use for a remote monitoring station is vehicle tracking, and in this case, a GNSS receiver is practically mandatory.

We don't need to add any extra hardware to implement tracking into our *4G Remote Monitoring Station*.

The two antennas are included when you purchase the shield. The mobile network antenna is a simple, self-adhesive PCB type.

Some photos of the shield show a small whip-style antenna, but it appears this has been replaced by the PCB type. A ceramic patch antenna is supplied for GNSS use.

Also on the shield is a BME280 temperature, pressure and humidity sensor. We covered modules built from similar sensors back in *PE* December 2018. This is a great addition as it adds even more sensor data to our *4G Remote Monitoring Station* without needing extra hardware.

We found that the temperature read by the sensor was higher than ambient, probably due to the heat generated by the surrounding circuitry; think of how hot some mobile phones get!

SIM7000E module

The SIM7000E module is touted as an NB-IoT/LTE/GPRS/GPS module. LTE and GPRS are longstanding mobile data transmission technologies, but NB-IoT is a newer standard.

NB-IoT is a low-power, narrowband variant of mobile phone technology, designed to be used by IoT (internet of things) devices. We aren't using the NB-IoT feature in this project; at this stage, it appears the technology is still being rolled out, and an NB-IoT-specific SIM card is required.

The SIM7000 module comes in several variants which support different mobile frequency bands. We are using the SIM7000E, which is designed for the European market and supports bands 3, 8, 20 and 28. There is also a SIM7000C, which is designed for the frequencies used in China.

We used a cheap supermarket SIM card to test our SIM7000 shield. It cost a few pounds and didn't need extra credit, even after two months of testing. The shield also has sockets for external mobile and GNSS antennas.

Because not all frequencies are offered in all areas, we suggest that you thoroughly research what frequencies are used where you plan to deploy the *4G Remote Monitoring Station*, to make sure this shield supports them.

This module does not support voice calls. Most monitoring stations typically use SMS (text messages) or data packets for communication. The SIM7000E module does support mobile data, and this is a great way to communicate lots of small snippets of monitoring data. Our design uses both the SMS and mobile data features of the shield.

ThingSpeak data logging

We're also using Thing-Speak for this project. It has a simple API (application programming interface) for uploading data, which is great for resource-constrained devices like Arduino microcontrollers. It also provides simple graphical visualisations of the recorded data. The data can also be downloaded as a CSV (comma-separated value) file.

These files can be opened into a spreadsheet program to allow more advanced analysis to take place. Creating charts is also an option in many spreadsheet programs.

Uploading data to the ThingSpeak website requires mobile data, so the SIM card used needs to support this.

Here are the three boards used in this project. From left to right, they are the *Power Control Shield*, the DFRobot 4G shield and the Arduino Leonardo microcontroller unit.

For the low cost, longer-expiry pre-paid mobile phone plan that we tried, it was typically cheaper to send weeks of data to the ThingSpeak website than to send a single text message.

Our *4G Remote Monitoring Station* is ideally suited to providing continuous logging of data via 4G as well as sending text messages for raising alerts for unusual situations that need to be acted on promptly.

Power control shield

In addition to the pre-built 4G shield, our *4G Remote Monitoring Station* also uses a custom-designed shield to provide for battery power, solar charging of that battery and some power-saving techniques to give it a long runtime when using a small battery.

Most Arduino boards have poor power efficiency; they have simply not been designed with this in mind. Even with the processor set to sleep mode, other components such as linear voltage regulators and LEDs have quiescent currents in the tens of milliamps.

Our shield reduces the standby battery draw to microamps, which it does by completely disconnecting the Arduino board (and SIM7000 shield) from the battery using a MOSFET, and only powering those components up periodically.

The shield provides a reasonably efficient way to charge the battery, and also monitors the supply and battery voltages via the Arduino's analogue inputs.

Most of the unused pins are broken out to headers, which allows other sensors or peripherals to be easily connected. There's even a small prototyping area on it, for extra components.

Shield circuit

The custom shield circuit is shown in Fig.2. Unregulated DC power is fed into CON1 (from, for example, a solar panel or plugpack), while the battery is connected via CON2.

The battery needs to operate in the range of 7-15V, so a 12V lead-acid or SLA battery is suitable. We used a 1.3Ah SLA with our prototype.

Power from CON1 feeds REG1, an LM317 adjustable regulator. The 220Ω fixed resistor and 10kΩ variable resistor VR1 allow you to set the output voltage.

REG1 maintains about 1.25V between its OUT and ADJ pins, so around 5mA flows through the 220Ω fixed resistor. This current mostly also flows through VR1, so by adjusting its resistance, you change the voltage between ADJ and GND.

Hence, you can set the voltage at VOUT, since this will be the voltage across VR1 plus the 1.25V.

The output from the regulator is filtered by a 1µF capacitor and fed to the battery via 1A

The electrolytic capacitor and 555 timer used on the *Power Control Shield* have been carefully chosen for low leakage and low quiescent current, to extend battery life. Note the jumper wire connecting the Arduino's D7 pin and the SLEEP terminal.

schottky diode D1. This prevents the battery discharging into the power source, for example, if it is a solar panel in darkness.

The 1Ω resistor that's between the output of REG1 and the anode of D1 reduces the output voltage as the current drawn from REG1 increases.

Hypothetically, if the current through this resistor reached 1.25A (which would not be possible in practice), the voltage across this resistor would rise to 1.25V, cancelling out REG1's reference voltage, so the output would drop to 0V.

Thus, the output voltage drops approximately 1V for every 100mA of load current.

So, if a battery is heavily discharged and its terminal voltage is low, the regulator output current is moderated until its voltage rises to the normal range, at which point virtually no current will flow into the battery.

In practice, the charging current is limited by dissipation to about 160mA for a 12V solar cell (with nominal 18V open-circuit voltage) feeding into a discharged 12V battery.

While the range of VR1 allows a terminal voltage from 1.25V up to 56V to be set, it shouldn't be set any higher than around 15V as this may damage the regulator on some Arduino boards, as well as IC1.

If you don't need to use a battery, power can instead be fed directly into CON2. D1 will prevent back-feeding into the charge circuit.

Arduino and SIM7000 shield power control

Power control is provided by 7555 CMOS timer IC1 and P-channel MOSFET Q2. Q2 is simply used as a high-side switch. Q2 can handle much more current than is required (hundreds of milliamps at most), so it does not need heatsinking. We have chosen the CMOS variant of the 555
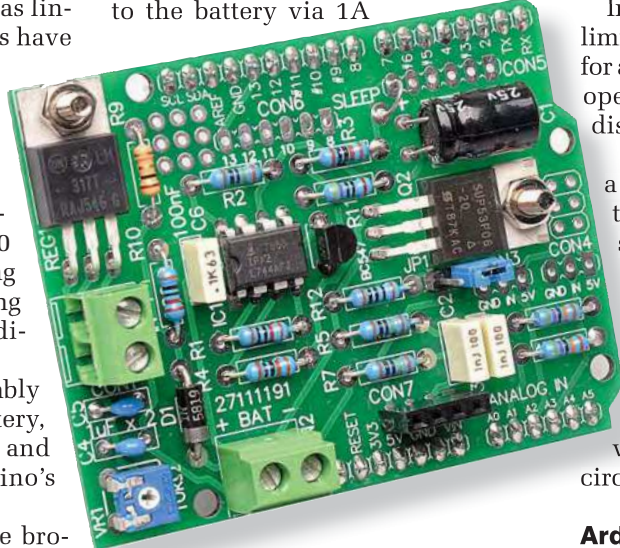
Fig.1: the SIM7000 shield circuit. It has several nice features, including an efficient buck regulator and a BME280 temperature/pressure/humidity sensor. The slide switch (at upper left) allows the serial data to be rerouted; we are using pins D0 and D1.

SIM7000 Arduino Expansion Shield

for its low quiescent current of around 60µA, compared to about 10mA for the bipolar version. This is because it is active and drawing current from the battery at all times.

IC1 is configured as a monostable timer. When power is first applied, the 470µF timing capacitor is discharged, and the threshold pin (pin 6) is below 2/3 of the supply voltage. The trigger pin (pin 2) is held high by the 10kΩ resistor. The transient conditions during power-up result in output pin 3 being high, and discharge pin 7 is in a high-impedance state.

With output pin 3 high, Q2's gate is high and so it is off, and the Arduino is not powered. The 470µF capacitor slowly charges up through the 1MΩ resistor. This capacitor needs to be a low-leakage type; otherwise, the leakage would prevent it from charging up fully.

The time constant of this circuit is nominally 470 seconds (just under eight minutes). Due to the 555's tripping point not being exactly 63% of the supply voltage, it actually takes around 10 minutes for the timer's state to change.

Once the trigger pin voltage reaches about 2/3 of the supply voltage, output pin 3 goes low, pulling down Q2's gate, switching it on and connecting the battery to the Arduino board's VIN pin.

This powers on the Arduino board and attached 4G shield. IC1's discharge pin, pin 7, goes low at the same time, discharging the 470µF capacitor quickly via the 470Ω resistor.

Being a monostable circuit, it remains in this state until the Arduino decides that it needs to power down.

To do this, it drives the base of NPN transistor Q1 positive, pulling the trigger pin (pin 2) of IC1 low. IC1's flip-flop toggles, output pin 3 goes high (switching off Q2 and the Arduino) and the discharge pin (pin 7) goes back to a high-impedance state, and allowing the timing capacitor to charge.

When the Arduino is shut down, it can no longer keep Q1 switched on, so there is no chance of this state latching.

Thus the cycle continues where it began. The Arduino has no way of turning itself on at a particular time; it just shuts down for the period of the monostable timer.

It's not exactly high precision, but it allows very low power consumption while ensuring that the Arduino is powered up periodically to do whatever it needs to do.

Jumper JP1 allows the monostable circuit to be bypassed. If JP1 is shorted, IC1's threshold pin is pulled above 2/3 of its supply, so MOSFET Q2 is forced on. As long as this jumper remains in place, the Arduino is unable to shut itself down.

This can be used to bypass the sleep mode during testing, or to force the *4G Remote Monitoring Station* to operate when deployed.

Sensing external voltages

Two six-way headers, CON5 and CON6, are provided to make connections to the Arduino's digital pins. A small prototyping area with nine pads is also provided.

A pad connecting to SLEEP is placed nearby. This is intended to be connected with one of the digital pins via a short jumper wire, meaning the pin used for shutting down the Arduino is not fixed in hardware, but can be altered. For our prototype, we used D7.

A small four-way header is also broken out for 5V, VIN and GND, since connected sensors or peripherals will need access to power.

Two analogue pins are connected to resistive dividers to sense the battery voltage (A0) and incoming supply voltage (A1). The 1MΩ/470kΩ divider means that voltages up to 15.6V can be measured. These high values are chosen to minimise loading (to around 10µA), especially on the battery.

The two 1nF capacitors provide a low source impedance for the analogue inputs, as otherwise, these voltage readings would be inaccurate.

Two more analogue pins (A2 and A3) are broken out to separate three-way headers (CON3 and CON4), along with ground and 5V. These allow common three-wire analogue sensor modules to be connected.

Note that there is nothing about this shield which ties it specifically to the 4G shield. Any application which requires battery charging, monitoring and low power consumption could use this shield.

Building the shield

Use Fig.3, the PCB overlay, as a guide during construction. The shield is built on a double-sided PCB coded 27111191, which measures 53.5 × 68.5mm and is available from the *PE PCB Service*.

We built our shield with simple headers to plug into the 4G shield below it. If you intend to add another shield above this one, you could use stackable headers instead, but that would make it difficult to access the pin headers on top of this board.

Start construction by fitting the resistors. There are several different values, so check the resistance of each with a multimeter. Then solder one of the lead off-cuts between the pad marked 'SLEEP' and the Arduino digital pin that you want to use for the shutdown function. We used D7, simply because it is close to the SLEEP terminal and it is not usually used for any specific purpose.

Next, mount the three rectangular MKT capacitors, which are not polarised. The 100nF part may be marked with the code 104 or possibly 0.1µF, while the 1nF parts may be marked 102. Follow with the two 1µF ceramic capacitors, which are also not polarised.

The final capacitor is the low-leakage electrolytic type. It is polarised, and we have left space on the PCB for it to be mounted on its side so that another shield can be fitted above. The negative lead (usually shorter and indicated by a stripe on the can) goes into the pad closer to Q2. If you want to use a larger capacitor for a longer delay, we have left a bit of extra room.

You may wish to apply a small dab of hot glue or neutral-cure silicone sealant to help hold it in place in case the unit is subjected to vibration.

Fit the semiconductors next. D1 is the only diode and goes near CON2, with its cathode stripe closest to CON2. Mount Q1 near the middle of the PCB, oriented as shown. Carefully bend its leads to suit the PCB footprint, push down firmly onto the PCB and solder it in place.

Q2 and REG1 are both in TO-220 packages that are mounted flat against the PCB, to keep the overall height low. Don't get them mixed up. Bend the leads back on each part, around 8mm from where the body meets the leads. Push the leads into the PCB pads and make sure the tab hole lines up with the PCB, then use an M3 machine screw and nut to secure the regulator before soldering its pins.

Next, fit timer IC1. You may need to gently bend the leads inwards to fit the IC to the PCB. Check its orientation to make sure it matches the PCB overlay diagram, then solder two diagonally opposite pins and check that it is flat against the PCB. If not, remelt the solder and adjust, then solder the remaining pins.

To solder the Arduino headers, plug them into another Arduino board (such as the Leonardo) to keep them straight. Place the shield over the pin headers and once you are happy that the headers are straight, solder each pin to the PCB. Plain headers are soldered on top while stackable headers are necessarily soldered underneath. Then remove the shield from the Arduino board.

Now mount CON1 and CON2, the screw terminal connectors. They are identical, and care should be taken that the wire entry holes face out from the PCB.

You can now fit headers for JP1 and CON3-CON7 for breaking out the various Arduino pins. They aren't needed for the most basic usage of the *4G Remote Monitoring Station*, but they are handy for adding extra sensors if and when necessary.

We used a four-way female header strip for CON7, to allow the VIN voltage to be monitored easily.

Fig.2: our control shield's circuit. Adjustable regulator REG1 sets the charge termination voltage and voltage/current characteristics for charging a battery connected via CON2. IC1, Q1 and Q2 shut down the Arduino for a fixed time before powering it back up again, as a power-saving measure. The effect is that the Arduino is powered intermittently, for as long as necessary to do its monitoring tasks.

SIM7000 Power Control Shield

To help test whether the shield is feeding power to the VIN pin, we rigged up a test LED by soldering a 1kΩ resistor to one lead. We then plugged this LED/resistor combo into the VIN/GND pair on CON7, with the LED anode to VIN.

Testing

We can do a few tests to check everything is working as expected. The shield must not be connected to any boards during these tests. Wind VR1 fully anti-clockwise before powering it up.

The first step is to adjust VR1 for the correct battery-charging voltage. This is done without a battery connected. To do this, connect a power supply to CON1, which supplies at least 3V more than the fully charged battery voltage. Adjust VR1 until the correct maximum charge voltage is reached at CON2.

For an SLA or similar 12V nominal type battery, set it to around 14.4V. In practice, this voltage is only reached at zero current, so the actual charge voltage is a bit lower than this.

If you are unable to set the voltage correctly, check the components relating to REG1. Otherwise, connect the battery and check that it is charged. You should see the battery voltage rising slowly.

Next, check the voltage between VIN and GND, using the LED we mentioned earlier or a voltmeter. These pins are easily accessible on CON7. You will probably get a zero reading, meaning that Q2 is off.

In this case, if you check the voltage across the electrolytic capacitor, you should find that it is slowly rising. This can be measured at pin 6 of IC1 referred to GND. Note that the load presented by your multimeter might affect this reading.

Now bridge JP1, to force Q2 on, and re-check the voltage at VIN. It should be close to the battery voltage and the voltage at pin 6 of IC1 should be low.

To simulate the Arduino activating the switch-off, momentarily connect the SLEEP pad to the 5V pin of CON2 using a 1kΩ resistor. VIN should drop to zero, and the electro should start charging. If it isn't, that could indicate that your capacitor is leaky.

If you're fussy about the exact timing of the sleep period, you can measure the time and change the values of the timing components to tweak it. Keep in mind that the Arduino needs to operate for at least 30 seconds to update its state, so sleep periods shorter than two minutes are not that useful, as the Arduino will spend much too much time starting up.

Once testing is complete, disconnect the power supply and batteries.

Building the Remote Monitoring Station

Having built the shield, now we can put it all together. We chose to use an Arduino Leonardo board for our prototype. It uses the ATmega32U4 micro rather than the Uno's ATmega328. The 'U' indicates that this IC supports USB.

Their specs are otherwise quite similar, but the Leonardo has the advantage that the hardware serial port on D0/D1 is not shared with the USB host serial interface used for programming. We can therefore use this to communicate with the SIM7000. The Leonardo also has an extra 512 bytes of RAM; this can be handy for remote monitoring because we need to store and process data before sending it.

Fig.3: fit the components to the *Power Control Shield* PCB as shown here. It shouldn't take you long to assemble. Just watch the orientation of D1, IC1, Q1 and the electrolytic capacitor. Also, ensure that the wire entry holes for CON1 and CON2 face outwards. You can use standard male headers (fitted to the underside of the board), or stackable headers, depending on how you plan to use the shield.

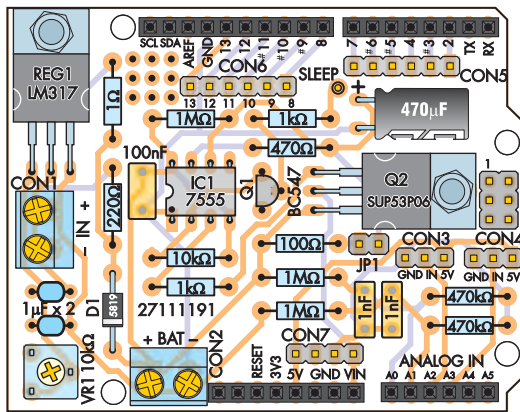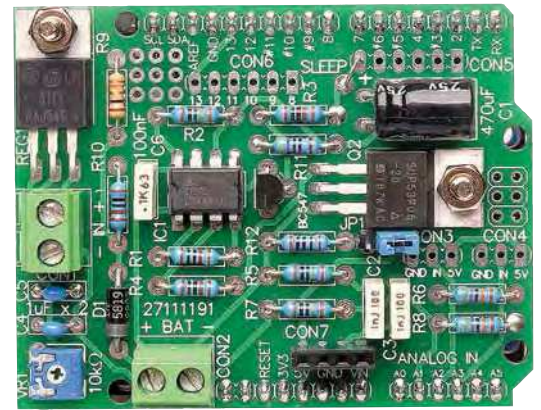If we had used an Arduino Uno, we would have been forced to choose between using the hardware serial port (D0/D1) to communicate with the SIM7000, which would interfere with programming and debugging, or using a software serial port which is slow and has a lot of overhead.

So we set the switch on the SIM7000 shield to the D0/D1 position, and as mentioned above, we used D7 as the sleep control pin.

To set up the 4G shield, fit the two antennas and a working SIM card. As with many of these sorts of applications, a prepaid SIM is preferred in case the microcontroller 'goes nuts'. With a prepaid limit in place, there is no chance of accidentally racking up huge data or call charges.

Now plug the 4G shield into the Leonardo and then plug the power control shield on top. Check for any fouling between the shields; if you have not trimmed all the leads closely, they may short together.

Our sample software simply logs data from the onboard sensors. We've also marked some places in the code to add your own tests or actions. For example, you could monitor a voltage and send an SMS if it gets too low or high. Or similarly, you could send an SMS if a switch is opened or closed.

You will need a ThingSpeak account to make full use of our sample code.

Setting up a ThingSpeak account

ThingSpeak can be accessed for commercial use with a time-limited free period, but a free license is available for personal use and offers four 'channels' and up to three million updates per year. If we were to send an update every ten minutes, then we would only need around 50,000 updates per year.

Go to **https://thingspeak.com/users/sign_up** and enter the information as shown on Fig.4. You may be prompted to confirm that you wish to use a personal email address, and also to click on a link sent in an email to verify that email address.

Once this is done, create a user ID and password, accept the Online Services Agreement and click continue, as per Fig.5. You will be prompted to select how you will use ThingSpeak. To be able to use the free license, you should choose 'Personal, non-commercial projects'.

The next step is to create a channel. Each channel consists of up to eight fields, so in theory, you could have up to four *4G Remote Monitoring Station*, each writing to their own independent channel.

Click on 'New Channel' and fill out the information as shown in Fig.6. You don't need to use all eight fields, but we have set the Arduino software to use all eight as shown. You should use the same fields unless you plan to modify the software.

Click Save, and you are shown the Channel data on the Channel management page, as seen in Fig.7.

Note that we did not create fields for latitude and longitude. ThingSpeak has hidden fields for this information. It can't be seen on the graphs, but is downloaded in the CSV data. Our Arduino code logs latitude and longitude to these hidden fields.

API keys

To allow our device (and only our device) to upload data to our channels, we need an API key. It must be programmed into the Arduino code for your *4G Remote Monitoring Station* to work with your ThingSpeak channel. Copy the 16-character alphanumeric code under 'Write API Key' to somewhere safe; we'll add this to the Arduino code soon.

You can test that your channel is working by copying the text after the word 'GET' in the 'Write a Channel Feed' box. Paste this into a web browser and hit Enter; you should see a blank page with the number '1'.

This indicates that this is the first update, and shows how the *4G Remote Monitoring Station* uploads data to ThingSpeak. This only updates one field; if you are familiar with HTTP, you might want to experiment with this.

With stackable headers, the three shields (PCBs) simply connect together via their header pins and sockets, as shown here.

Browse back to the 'Private View' of the created channel, and you should see some activity in the first field; this is the data you sent from the web browser. You can leave this window open while testing, as it will update in near-realtime and you can see the results.

Arduino libraries

There are four libraries needed for the software we have written; two are included with most Arduino IDE distributions. We used version 1.8.5 of the Arduino IDE.

The avr/sleep and Wire libraries are the two usually included. The first library provides functions for low-power modes, while the second provides an I²C interface for communicating with the BME280 sensor.

The third library, which we created, is named 'cwrite'. It lets us read and write from a character array as though it is a stream object, so we can use the print function's ability to format floating-point numbers to form a URL.

The resulting datum can then be sent to the 4G Shield in one go.

This library can be seen as two extra tabs in the Arduino project. If you make a copy of the project (by using File -> Save As...), then this library is copied too.

The final library is to make the actual temperature, humidity and pressure readings from the BME280 sensor. It is written by a company called SparkFun and can be installed via the Library Manager utility of the Arduino IDE. Search for 'sparkfun BME280' under the Library Manager and click Install.

We have included this library in our software bundle for this project, in case you can't find it.

Arduino software

We set up the Arduino software to work with the eight fields that we have just created, plus three hidden fields of latitude, longitude and height above sea level.

These three fields are from by the GNSS receiver on the SIM7000 module, plus the BME280's atmospheric pressure sensor to determine altitude.

The software is modularised in such a way that proficient Arduino users can modify it to behave differently, if needed.

In any case, you will need to edit the software to suit your API key. Around line 28, change the text API_KEY_HERE to the API key that you copied earlier. You should end up with a 16-character sequence surrounded by double quotes.

Below this, on lines 29 and 30, are entries for phone numbers. Any

incoming text messages have their number checked against the AUTH_NUMBER string. The sequence AUTH_NUMBER_HERE should be replaced by the trailing digits of your phone number.

We have done it this way to allow matching of both national and internationally formatted numbers. Thus for an Australian mobile number, the first digit should be the '4', meaning the leading '0' is dropped. Your local national standard for numbers may well be different.

The sketch simply matches whatever digits are present. So if this were changed to '693', then any number ending in '693' would be accepted. If you don't wish to use this feature, leave it as the default string, as this is highly unlikely to match an incoming number.

The outbound number should be a fully qualified international mobile number; eg, an Australian mobile phone number including country code would start with '+614' (+61 for Australia) followed by eight digits. This is used for outgoing text message alerts.

Many of the other 'defines' are set to set things like the analogue input voltage measurement ratios, and how much memory is available. There is no reason to change these unless you are modifying the hardware.

The software performs basic initialisation of the hardware in the `setup()` routine. More initialisation happens in the `loop()` function, particularly for the 4G shield.

The code sends the shield some data to see if it is powered up and if not, toggles the power line.

A set of fixed sequences are sent to put the 4G shield in a known state. The shield is given ten seconds to get a GNSS fix. If this is successful, the unit's speed is checked, and a message is sent if it is higher than 100km/h. This is a basic demonstration of how easily an action can be performed based on sensor state.

The code then sends an update to ThingSpeak; this is wrapped up in a single function which validates the GNSS data and only sends that data if it is valid. The Arduino then checks for text messages from the authorised number. If one is found, a canned response is sent.

You can modify the code to check the message content and perform different actions (and supply different responses) depending on it.

If the GNSS data is not valid, then instead of powering off, the Arduino goes to sleep and leaves the 4G shield running to allow it to get a fix. This does not reduce power as much as switching the Arduino off,

but it does give it a chance to get a position fix.

If the GNSS data is valid, then the modem's power pin is toggled (to perform a controlled shutdown), and D7 is driven high to power down everything else.

Next time the Arduino powers back up, the sequence repeats, resulting in ThingSpeak updates about every 10 minutes.

Each update uses around 2kB of data, which, according to our mobile plan, costs around $0.0001. During our testing, we sent around 6000 updates (over a month's worth of updates) for a total cost of $1.08. *Your plan might vary.*
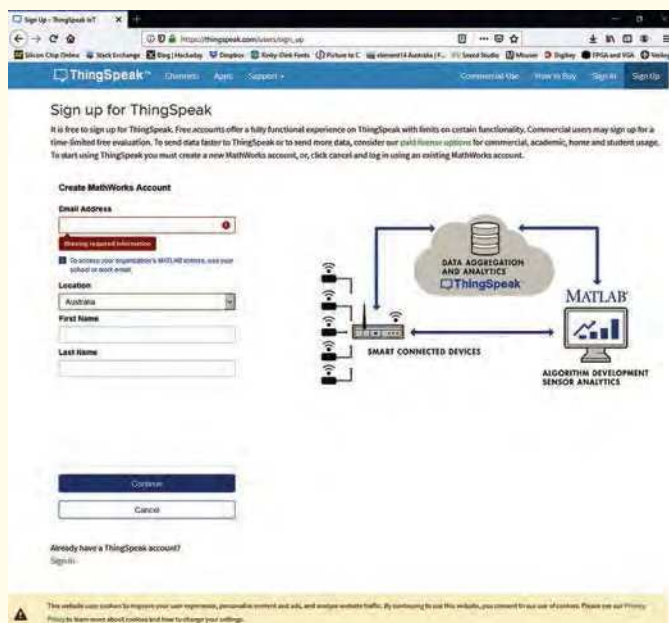

Fig.4: apply for a ThingSpeak account via the web page shown here. This is needed to use the software we've written, as ThingSpeak lets you upload data to 'the cloud'. MATLAB users can use their existing account for ThingSpeak.
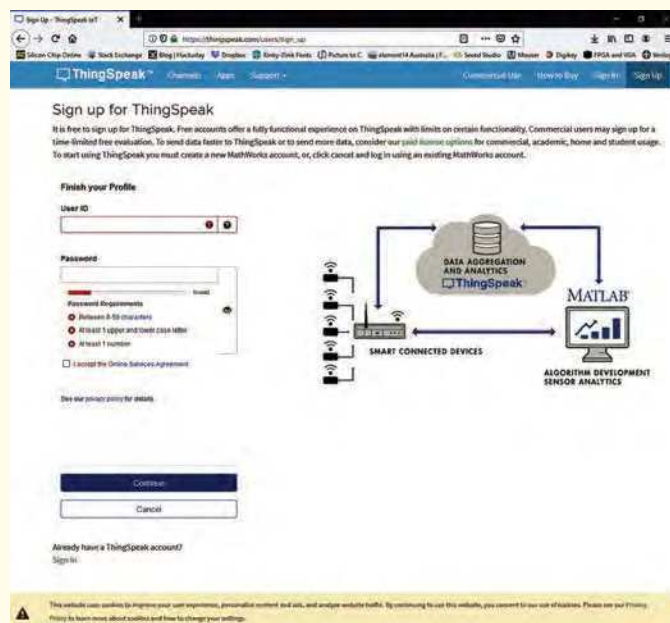
Fig.5: as with many online services, you need to create a username and password for ThingSpeak. This page indicates if your chosen username is free, and how strong it thinks your password is.

Finishing it off

Connect the Arduino board to your computer and select the Leonardo board and its corresponding serial port from the Arduino IDE menus.

Compile and upload the **4G_Monitoring_Station.ino** sketch. Unplug the Leonardo and attach the two shields.

Connect the battery to CON2 and the power source to CON1. Briefly short out JP1 and check that the whole assembly powers up.

The upload to ThingSpeak should take less than a minute. If it does not, then you may need to do some debugging to find out what's wrong.

Re-connect the Leonardo board to the computer and open a serial terminal program to monitor the output. It should look like that shown in Fig.8. Look for a 200 HTTP code and 'ThingSpeak Success' message.

If you get this, then uploads to ThingSpeak are working correctly.

You might find that the Arduino Serial Monitor does not behave well when the Leonardo powers down. We had success with a program called TeraTerm, as this automatically re-connects to the serial port if it disconnects and reconnects.

Unfortunately, the USB lead will also power the Leonardo, so the power down functions may not work as expected while connected to a computer.

A trick for testing

To test our prototype, we needed a way to allow USB communication but without powering the Leonardo, as this interferes with the power control hardware.

To achieve this, we used one of our *USB Port Protector* PCBs described in *PE*, May 2019.

If the *USB Port Protector* PCB is wired up with no components except the USB plug and USB socket (CON1 and CON2), then it connects GND, D+ and D–, but not 5V.

Thus, this 'dongle' can be used to connect a USB device to allow data but not power to be transferred. The Leonardo is then powered via its onboard 5V regulator fed from the VIN pin.

Take care that the ground of your computer is not at a different potential to the ground of the *4G Remote Monitoring Station*; for example, if you are powering it from a bench supply or similar, make sure the outputs are floating.

You can use a battery-powered computer for testing if you are not sure about this.

Debugging

While the code is quite complex, we did not run into many problems with it. But in case you do, we'll run through some of the error messages the Arduino might display.

If you don't see the 'GNSS on' or 'Format set' messages, your Arduino is probably not communicating with the 4G shield.

According to the shield's data sheet, it communicates at 115,200 baud, but our unit was set to 19,200 baud. You can change this setting at line 5 in the Arduino code.

After the 'GNSS on' message, you should then see 'Success' and a network name.

If you see 'Fail' here, the 4G shield is not registering with the network. This generally happens when the 4G shield has no reception. It could be due to the shield not supporting your telco's frequency band, or you may be out of range of a cell tower. Check that the antennas are connected correctly.

You will occasionally see 'GNSS fix fail' as the *4G Remote Monitoring Station* compromises getting a fix at all times for saving power.

The code tries to retrieve the APN (access point name) from the 4G shield and use it to connect to mobile data. If you see a message referring to APN, CSTT or bearer failing, then this is not being set up correctly. Check the APN name that your provider uses.

The URL that the *4G Remote Monitoring Station* uses should be displayed, followed by an HTTP result code. If it is not 200 (HTTP success), check **https://bit.ly/pe-feb21-err** to see what other error codes mean.

If it still isn't working, numerous extra debugging lines in the code have been commented out (by adding '//' to
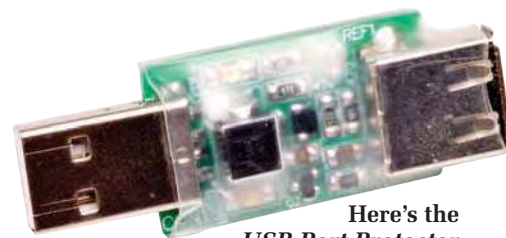
Here's the *USB Port Protector* from our May 2019 issue – we used one of these without components, except for the USB plug and socket, during testing. This allows data to be transferred, but not power.
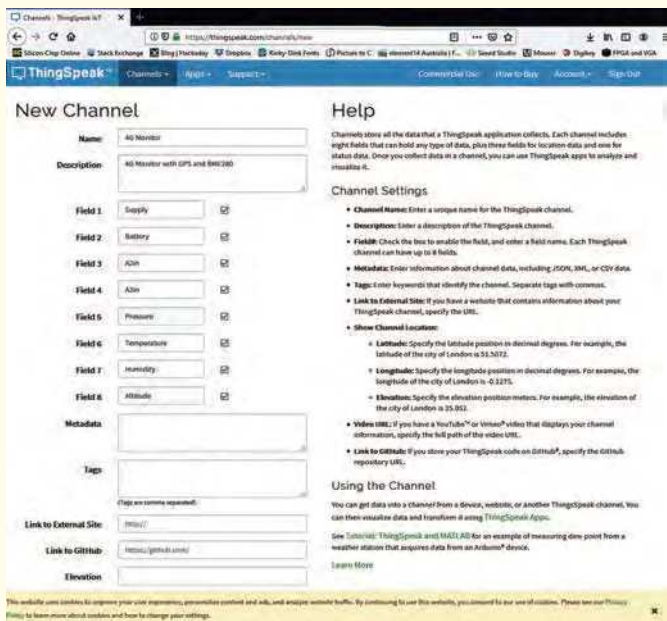
Fig.6: we recommend that you (at least initially) create a ThingSpeak channel – set up its fields as shown here. These fields suit the data produced by the *4G Remote Monitoring Station* software. They can be changed later if necessary.
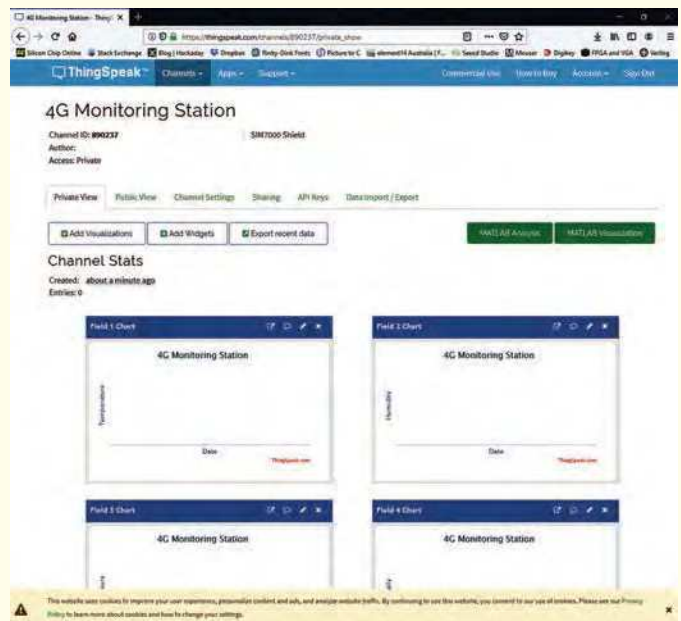

Fig.7: once the channel has been created, you can go to its overview, which defaults to a series of charts. You can add more with the 'Add' buttons. By default, the channel data is private, but you can set it to be visible to others.

the start). You can enable the extra messages by removing these and compiling and uploading the code again.

You can also try our **Leo_FTDI_ with_passthrough.ino** sketch. This configures the Leonardo to allow direct communications between the serial port and the 4G Shield.

You can try different baud rates to see what works and send commands directly to the 4G Shield. Upload this to the Leonardo and short JP1 on the power control shield. You may need to press the 4G shield's BOOT button to power it up manually. Once you have confirmed the correct baud rate, upload **4G_Monitoring_ Station.ino** to the Leonardo again.

Conclusion
We've deliberately left this as an open-ended project; we expect that readers will customise the hardware and code to suit different applications.

For outdoor use, we recommend housing everything in an IP-rated plastic enclosure, with both antennas mounted on the underside of the lid.

Including some vent holes, facing down, can help to drain any condensation which may form, and allow the outside air to be sampled by the BME280 sensor.

Reproduced by arrangement with *SILICON CHIP* magazine 2021.
www.siliconchip.com.au
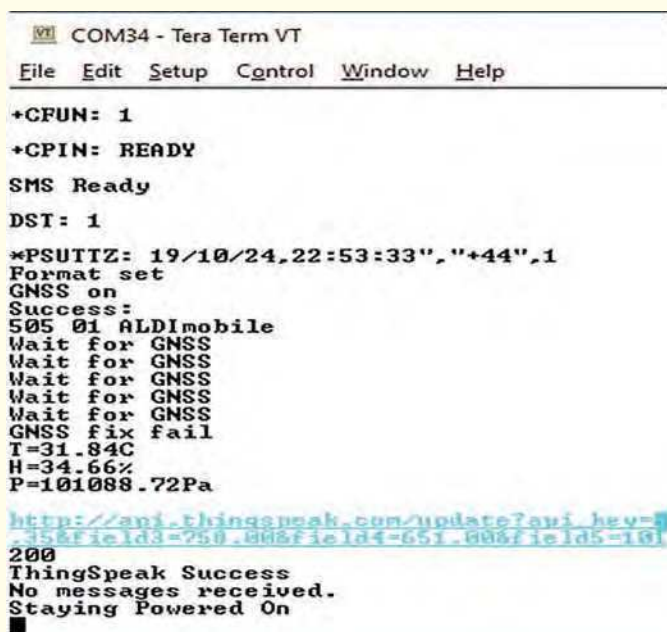
Fig.8: here is some sample debugging data from the serial port of the *4G Remote Monitoring Station* during normal operation. Your data may differ, if, for example, you have a different telco.

Fig.9: as the *4G Remote Monitoring Station* is only powered up infrequently, it may not respond to an SMS command immediately. The two-minute delay shown here occurred during testing, when we had the Interface Shield set to power down for about one minute at a time. The canned response can be changed by editing the Arduino code.

Low-distortion DDS Signal Generator

This two-channel audio-signal generator produces very low distortion sinewaves as well as triangle waves, square waves, pulse trains and noise. It has adjustable output frequency, amplitude and phase, plus sweep and pulse modes. Control is via a rotary encoder knob and graphical LCD screen with an intuitive graphical user interface). It's ideal for testing amplifiers, loudspeakers and all sorts of audio equipment, as well as for general purpose use.

by Phil Prosser

On my work bench, I find the most useful test equipment – after my multimeter – are a signal generator and an oscilloscope.

When testing analogue circuitry, I find it very important to be able to 'stimulate' a circuit and then look to see how it responds.

In the past, I used simple oscillators for this job, usually Wien Bridge types due to their high performance at low cost.

But when testing speakers, it is very useful (in fact, almost essential) to be able to sweep the generated tone frequency. This allows multiple drivers and crossovers to be tested.

And when testing amplifiers and speakers, it is very handy to be able to generate short bursts of a tone with a silent gap. A simple oscillator can't easily do either of those things.

Burst tones serve a couple of purposes. First, when testing a power amplifier driving a low impedance, it can be quite stressful on heatsinks, dummy loads and power supplies using a continuous waveform.

If you can generate, say, two cycles at 1kHz, followed by a second of silence, you can see what the amplifier does at and near clipping and at high currents without really stressing things.

Also, when building loudspeakers, it is really useful to be able to generate tone bursts at and around the crossover point. This lets you set up a microphone to measure the time delay for the tone burst through a bass driver and a tweeter.

This is essential if you want to 'time-align' drivers in a speaker cabinet. For some time, I used the free Audacity software on my PC to do these jobs. But that set-up was a bit clunky, so I set about designing a more convenient hardware device that could do all this for me.

The device that I came up with, presented in this article, uses the same hardware as my *DSP Active Crossover/Parametric Equaliser* project that

was published in the January, February and March 2020 issues of *PE*.

If you've already built that, it's simply a matter of reprogramming the microcontroller to perform these waveform direct digital synthesis (DDS) functions.

If building it from scratch, you can build it as described in those earlier articles, although you only need one of the stereo digital-to-analogue converter (DAC) boards (not two) and you don't need the analogue-to-digital converter (ADC) board at all.

The January 2020 article gave an overview of the hardware and described how all the separate boards worked, except for the CPU control board and the front panel controls. Those were covered in February, along with the PCB assembly details. The March article gave programming and final assembly instructions, along with

usage instructions that are not directly relevant to this project, as the software is different. However, if you have seen that, the user interface of this new software will be familiar to you.

Those were long and detailed articles so we won't reproduce all that information here.

We'll just give a quick overview of how the hardware works and then jump into describing the new software.

The hardware

There are four PCBs involved in this project, and the basic arrangement is shown in Fig.1: first is the CPU board which hosts the powerful PIC-32MZ2048 32-bit processor, a couple of regulators and a crystal to provide an accurate clock source.

This connects to a power supply and signal routing board which derives the DC supplies required to power the various other boards from a 9V AC plugpack or transformer.
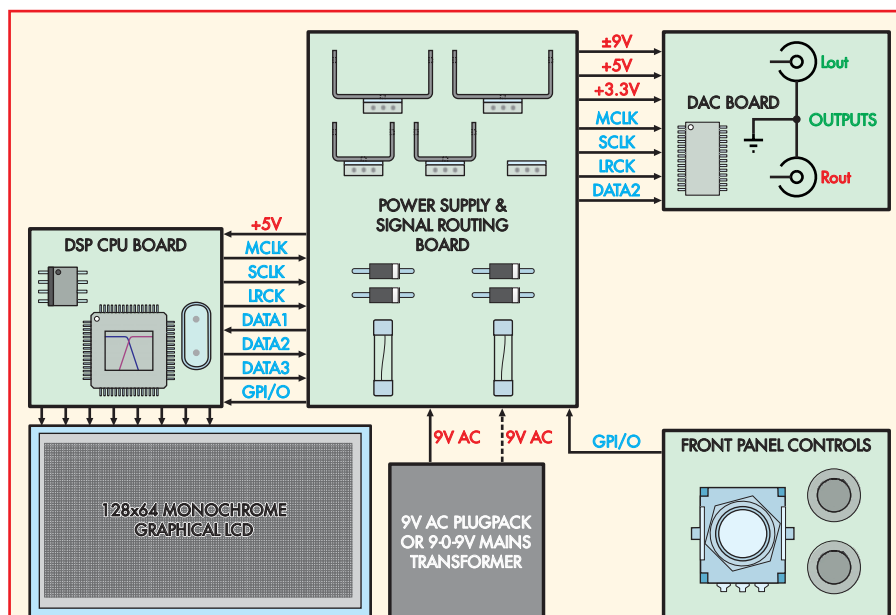
Fig.1: the powerful 32-bit PIC processor on the DSP CPU board generates a stereo digital audio signal which is routed to the DAC board (at upper right) through the central power supply/signal routing module (upper centre). The DAC board (upper right) converts the audio data into two analogue output signals. It's all controlled via an LCD screen, one knob and two buttons.

Features & specifications

Output channels:	Two, fully independent
Output waveforms:	sine, triangle, square, pulse (1-99%), noise
Output frequency:	0.1Hz to 20kHz in 0.1Hz steps
Output amplitude:	-123dBV to +5dBV
Frequency sweeps:	linear or log, phase-continuous
Pulse modulation:	1-1000 cycles on, 1-1000 cycles off, zero cross switching
Phase offset:	0-360° per channel in 1° steps
Distortion performance:	below -100dBc (0.001%) – often well below (see Table1)

Digital audio signals from the CPU are routed through this central board to a stereo DAC board which provides the two analogue outputs via onboard RCA connectors.

The fourth board is a front panel control board with a rotary encoder (which may have an integral pushbutton) and one or two separate buttons. The CPU board drives the graphical LCD module directly.

Fig.1 only differs from Fig.3 on page 16 of the January 2020 issue, which shows the *DSP Active Crossover/Parametric Equaliser* configuration, in that we've removed the unnecessary ADC input board and the second stereo DAC output board. Otherwise construction is identical.

Principle of operation

In case you aren't already familiar with how a DDS works, we'll give a quick description and describe why they are so much more useful than basic oscillators.

Direct digital synthesis is a process for mathematically generating a waveform. The system comprises a clock source, a 'phase accumulator' and a lookup table that determines what analogue output is generated at any given time. Fig.2 shows this configuration.

The clock source runs at a high frequency compared to the output waveform frequency; the clock frequency is often hundreds of times more than the waveform frequency. On each clock cycle, the DDS system adds the phase increment to the phase accumulator. The lookup table can hold any waveform, though conventionally it would be a sinewave or something similar.

As the content of the phase accumulator increases in value, the system 'steps through' the table, feeding subsequent values in this table through to the output, reconstructing the waveform stored within. The rate at which it steps through the table (determined by the phase increment) determines the frequency of this reconstructed waveform, ie, how many times it runs through the table each second.

For example, if the clock frequency is 48kHz and you want an output frequency of 4800Hz, then you

would need to produce one full sinewave every 10 clock cycles. Thus the phase increment needs to be 1/10th of the maximum phase value (equivalent to 36°).

If, instead, you want to produce a 1Hz output with the same clock frequency, the phase increment needs to be 1/48,000th of the size of the phase accumulator (0.0075°). You can see then that for good low-frequency performance, a high-resolution phase accumulator is desirable. The PIC Microcontroller has a natural word size of 32 bits. That's suitable for moderate clock frequencies (up to a few MHz).

Also, to make the software simple, you want to use numbers in the DDS that the processor can divide easily. Using 32 bits makes things easy, since dividing by powers-of-two is very easy and fast (it can be done with simple shift/bit masking operations).

So we have a phase accumulator that is used to look up the waveform. For high precision, we need the sinewave to be very accurate. The limited memory in the PIC makes this a little

Reproduced from the *DSP Signal Processor* project in our January, February and March 2020 issues, here are the four main PC boards used in this *DDS Signal Generator*. Top left is the main CPU board, bottom left the power supply board, top right is the digital-to-analogue (DAC) board, while bottom right is the front panel board. Not shown here is the graphical LCD module nor the AC transformer.
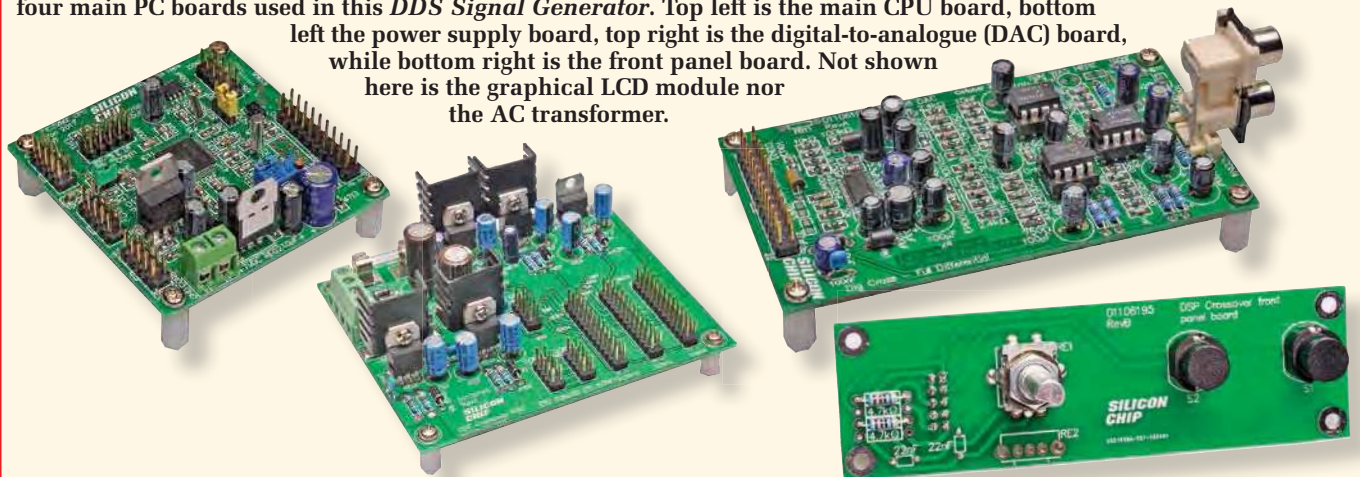
Fig.2: the basic configuration of a direct digital synthesiser. The phase increment value (chosen for a specific output frequency) is added to the phase accumulator on each pulse from the clock source. The accumulator is then used to index a waveform lookup table, and the values looked up feed the DAC to producing a varying analogue waveform at the output. Its shape is determined by the values in the lookup table.
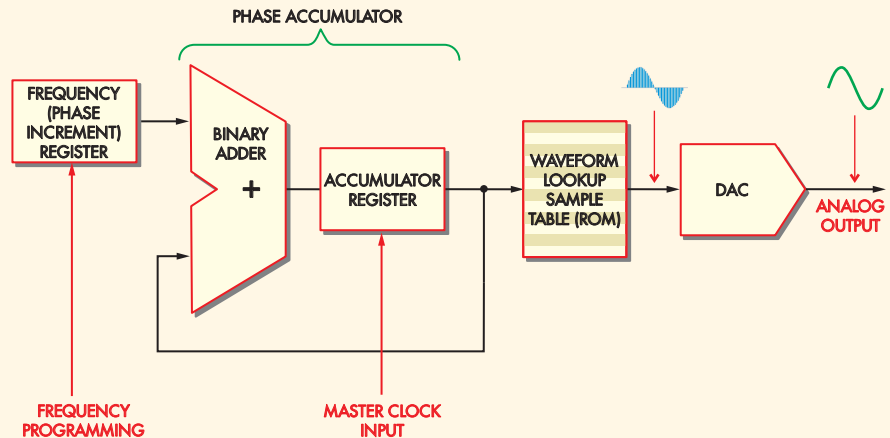
tricky, as the largest practical sine-wave table is around 4096 entries long. That's quite good, but not good enough for extremely low distortion.

The configuration of a basic DDS using a 4096-entry (12-bit) table is shown in Fig.3. This works reasonably well but, as shown, there are rounding errors in the values produced which makes it not quite a pure sinewave.

To solve this, we use the top 12 bits of the 32-bit accumulator to look up values from the 4096 entry (2^{12}) table, with the remaining (bottom) 20 bits determining where the current point in the waveform falls *between* samples in the table. This value can then be used to linearly interpolate between the two nearest (adjacent) entries in the table, to emulate having a much larger table.

The error in the resultant sinewave is vanishingly small. So the DDS configuration is now as shown in Fig.4. It's just a little more complicated than the one shown in Fig.3 and the extra calculations are well within the capabilities of a PIC32. The result is now very close to a perfect sinewave at just about any frequency.

The DDS needs to do all these sums and interpolations 48,000 times per second. But this is what computer chips are really good at. The net result is that the data fed into the DAC is correct to within 1-2 least-significant bits.

Without interpolation, spurs in the output are between −70dBc and −80dBc, which makes sense as the errors are about 1/4096 of the full amplitude ($20 \times \log_{10} (4096) = 72$dB). Fig.5 shows the frequency spectrum of the output without interpolation, and Fig.6 shows the same output, but with interpolation.

The DDS running without interpolation gives a creditable performance; those spurs (they are not harmonics) at −75dBc are around 0.02% of the waveform amplitude. For the version with interpolation, the largest signal by far is the hum being picked up by the test set at −100dBc or about 0.001% of the amplitude, with the distortion products being 1/10th of this, or around 0.0001%!

So with interpolation, we can get a result that is close to the limits of the DAC's performance.

Along with getting the correct amplitude values to the DAC, we need to send them at the right time. Jitter in the clock, which determines when data is sent to the DAC, can also distort the resulting waveform. So we set up the PIC32MZ and the SPI port going to the DAC so that the notional DAC clock is an integral fraction of the PIC32MZ system clock of 252MHz.

By choosing the DAC clock this way, there is no jitter or error on the timing of the DAC signals. This is critical in all the clocks, phases and precisions being controlled. It does result in a rather unusual DAC clock rate of 49,218 samples per second, but that doesn't really matter, except to slightly complicate our phase increment calculations.

The PIC32MZ can generate a very accurate 48,000kHz clock, but this uses a 'trim' on the DAC clock, which introduces jitter. We don't want that!

Construction and programming

As mentioned earlier, the construction steps were detailed in the February and March 2020 issues of *PE* on the *DSP Active Crossover*. My only suggested change, other than leaving out the redundant second DAC board and the ADC board, is that you may wish to mount the DAC board with its two output RCA sockets coming through
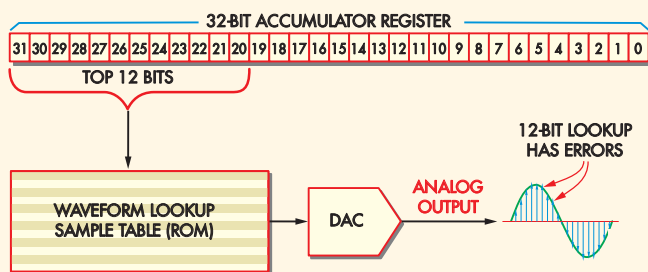
Fig.3: here's how the contents of a 32-bit accumulator is used to index values in a 4096-entry (12-bit) lookup table. The bottom 20 bits are ignored (although they are still needed to achieve the correct output frequency) while the top 12 bits are used directly as the table index value. This results in an output waveform with an accurate frequency but varying amplitude errors, leading to increased (but not necessarily excessive) waveform distortion.
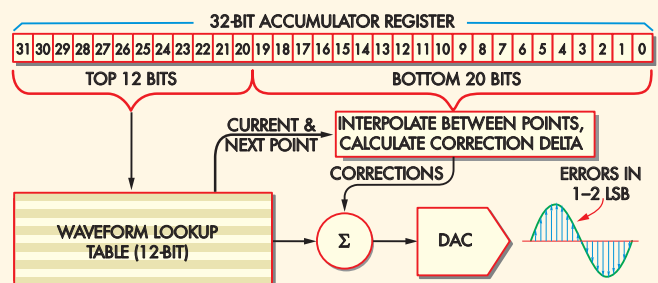
Fig.4: by adding a little extra complexity to the table lookup scheme, we can dramatically reduce the waveform distortion. The table lookup now retrieves the current and next values, and the bottom 20 bits of the accumulator are no longer ignored. Instead, they are multiplied with the difference between the two values from the table, to produce an error correction term which means that the values fed to the DAC are linearly interpolated between the table values. This reduces the instantaneous output amplitude errors to a tiny fraction of full-scale, and they essentially become negligible.

the front panel, rather than the rear. That makes it more practical to use as a test instrument.

The software for this project, both as a compiled HEX file and an MPLAB X IDE C project, can be downloaded from the Febrauary 2020 page of the *PE* website. You can then upload that HEX file to the PIC32MZ chip using the procedure described on pages 22-24 of the March 2020 issue.

You can use that same procedure to reflash a *DSP Active Crossover* so that it can perform the DDS role outlined here. There is absolutely nothing stopping you from flashing it back to the *Active Crossover / Parametric Equaliser* firmware when you have finished with that.

Remember that these chips eventually *do* wear out if you keep reflashing them, but given the PIC32MZ's specification of a minimum of 10,000 erase cycles, you're unlikely to wear it out this way.

User interface
When the unit is first powered on, you can cycle through the following four main screens using the rotary encoder:
- Load settings
- Save settings
- Channel 1 options
- Channel 2 options

Once you've selected one of the channels, by pressing the Select button or pushing the rotary encoder, it jumps to the output frequency setting screen (see Screen 1).

Turning the dial does the obvious thing – ie, increments and decrements the frequency. The Exit button returns to the channel selection screen, and the Enter button advances to the next menu page.

As you change frequency, if you turn the dial continuously for about a second or so, the step size increases by a factor of 10, allowing you to use both fine and larger frequency steps from the one screen.

This is one of a few data entry screens that are 'sticky'; if you leave the controls untouched for a long while, the interface will remain on this screen, rather than reverting to the initial screen. The reasoning here is that if you are fine-tuning a speaker or a filter, you may need to make the occasional frequency adjustment and it would be annoying to have to go through the menus again each time.

If you press the select button/knob on the frequency setting screen, it takes you to a step adjustment screen (Screen 2), which lets you increase the base step size from 0.1Hz to 1Hz, 10Hz or 100Hz. This may be useful if you are making lots of rapid changes to the output frequency and don't need to be extremely precise.

The behaviour of the Exit button on this screen is different; exiting from this screen takes you back to the previous frequency setting screen, rather than the initial screen. Pressing Enter on this screen takes you to the output level adjustment screen (Screen 3).

The output level is set in dB volts (dBV). 0dBV = 1V RMS, −20dBV = 0.1V RMS, −40dBV = 0.01V RMS and so on. The resolution is 0.1dB, which is quite a small step. As you rotate the control, if you rotate for more than a second or so, the increment size increases (as with the frequency adjustment), allowing you to make larger changes in output level reasonably quickly.

The output range goes from +5dB volts (1.78V RMS) to −123dB volts (0.7mV RMS), but note that −123dB volts is essentially the noise floor of the DAC. The Exit button takes you back to the main screen, while pressing Enter takes you to the waveform selection menu (Screen 4).

On this page, rotating the encoder cycles through the various waveform types, including sine, triangle, square, pulse and noise. If you select pulse, a secondary menu pops up, allowing you to set the duty cycle in 1% increments.

Having chosen a new waveform, press Select to reprogram the waveform lookup table and change the output signal. This also takes you to the next screen, which is the sweep selection menu (Screen 5). As usual, pressing Exit will take you back to the channel selection menu.

Sweeping and pulses
Sweep options include None, Linear and Log (logarithmic). Logarithmic sweeps (actually exponential) are very useful for wideband tests and audio testing. Your ear is much better 'tuned' to a logarithmic sweep than linear, giving the sense of a constant rate of pitch increase or decrease.

This table shows how the *DDS Signal Generator's* sine-wave performance varies with the output frequency and amplitude. You can see that the harmonics are very low, only increasing to a significant level when the output amplitude is below −20dBV.

Frequency	Amplitude (on DDS)	Amplitude (measured)	Harmonics
110Hz	0dBV	0dBV	−105dBc (0.00056%)
110Hz	−3dBV	−3dBV	−116dBc (0.00016%)
1kHz	0dBV	0dBV	−104dBc (0.00063%)
1kHz	−6dBV	−6dBV	−113dBc (0.00022%)
1kHz	−10dBV	−10dBV	−109dBc (0.00035%)
1kHz	−20dBV	−20dBV	−103dBc (0.00071%)
1kHz	−30dBV	−30dBV	Approx −90dBc (0.003%)
1kHz	−50dBV	−50dBV	Approx −70dBc (0.03%)
1kHz	−80dBV	−80dBV	Approx −40dBc (1%)
2kHz	−3dBV	−3dBV	−116dBc (0.00016%)
5kHz	−3dBV	−3.1dBV	−116dBc (0.00016%)
Table 1 – measured performance with various sinewaves			

Screen 1: this screen lets you set the output frequency of the currently selected channel.

Screen 2: here, you can change the size of the frequency steps in the previous screen.
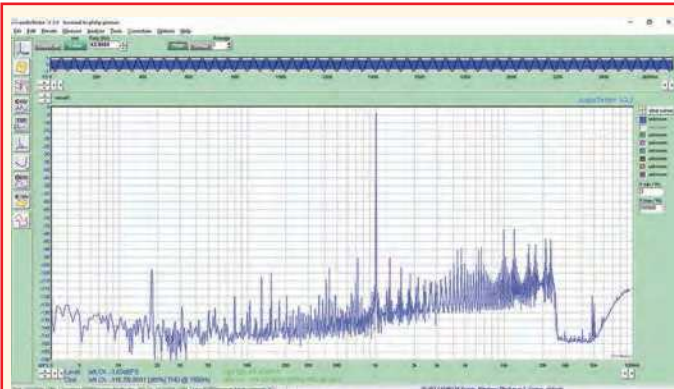
Fig.5: to prove the advantages of the linear interpolation scheme, here is a spectrum analysis (FFT) of the output of the unit using just the basic 12-bit lookup scheme shown in Fig.3 (ie, no linear interpolation). You can see that there are a large number of spurious signals at various frequencies present in the waveform, although most of them are not harmonics of the output sinewave. This leads to a distortion level of around 0.02% (the level shown in the screengrab is not accurate).
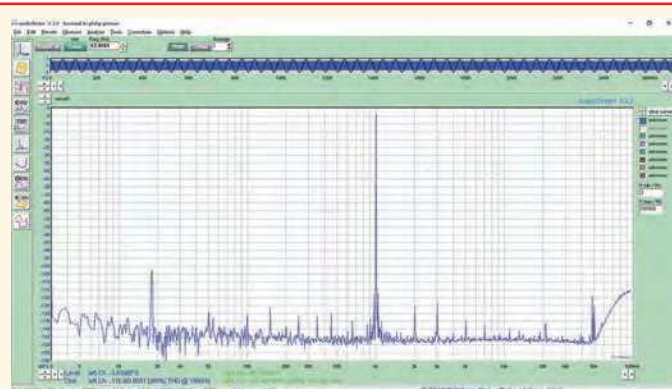
Fig.6: this is a spectrum analysis of exactly the same signal as in Fig.5, but this time, with linear interpolation. You can see that there is a lot less noise in the signal, and the actual harmonics are now visible, along with some mains interference at 50Hz, 100Hz and related frequencies.

The frequency at which the sweep starts is the frequency set on the main frequency menu. Having chosen the sweep type, you can then adjust the sweep end frequency. This may be above or below this start frequency. This allows sweeps upward and downward without complicating the interface.

A sweep is implemented by the DDS frequency being changed in 4096 steps throughout the sweep. Without

Screen 3: the next screen lets you set the output level from −123dBV (0.7mV RMS) to +5dBV (1.78V RMS) in 0.1dBV steps.

Screen 4: as you rotate the knob on the waveform selection screen, you can select Sine, Triangle, Square or Pulse. For Pulse, you can set the duty cycle from 1% to 99%.

getting into too much detail, the sweep involved repeatedly changing the phase increment value.

This has the benefit of the output phase from the DDS always being continuous (ie, no phase 'jumps'). A sample of the sinewave output in sweep mode is shown in Fig.7.

Having adjusted the end frequency and pressed Select, you can then adjust the sweep time from 0.1 seconds to 60 seconds in 0.1-second steps using the knob. Press Select again to move onto the next screen, or Exit to go back to the main selection screen.

The next screen configures pulse-train mode (Screen 6). The pulse train can be set to off or on. If you turn pulses on, you need to set the number of pulses on and off for the train. Both can have values of 1-1000 pulses. So with a 1kHz tone, you could have a single 1kHz cycle followed by one second of silence.

As noted earlier, this is very handy for high power amplifier testing (especially stability testing into very nasty or low impedances), and for testing loudspeaker driver phase centres. Fig.8 shows an example of the output using the pulse mode, with a 1kHz waveform set for five pulses on and five pulses off.

Pressing Select after the last pulse mode screen takes you to the phase shift screen (Screen 7), which allows you to change the relative phase between the two output channels. This only really makes sense where both channels are set to identical frequencies. This is particularly

Screen 5: you can disable sweeping, or have a linear or log (exponential) sweep. If enabled, you set the end frequency and sweep time in the next couple of screens (the start frequency is the channel's set frequency). It can sweep up or down.
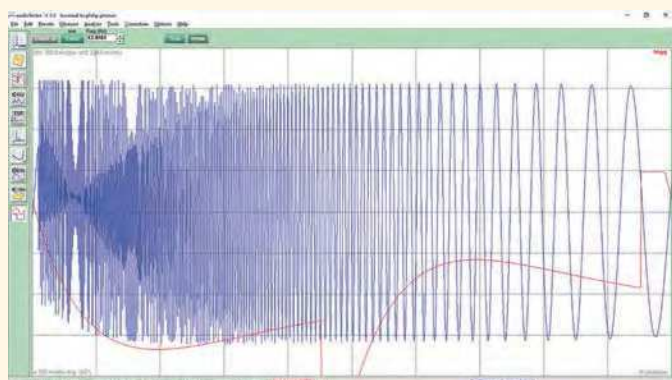
Fig.7: an example output waveform from the unit in sweep mode (blue trace). It is sweeping from a high frequency down to a low frequency.
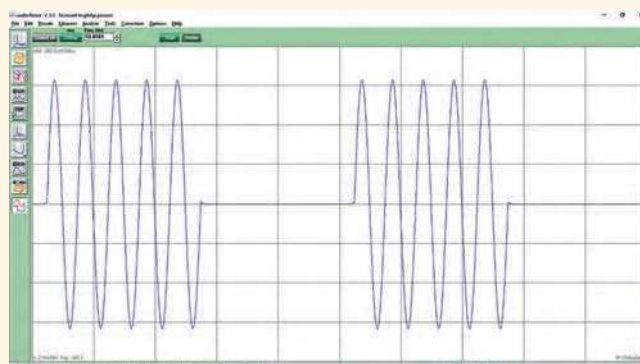
Fig.8: a sample pulse train output, with five pulses on (at 1kHz), followed by five pulses off. This sequence then repeats endlessly.

handy when you want to use one channel to trigger an oscilloscope. The phase shift lets you move the waveform of the second channel on the oscilloscope screen.
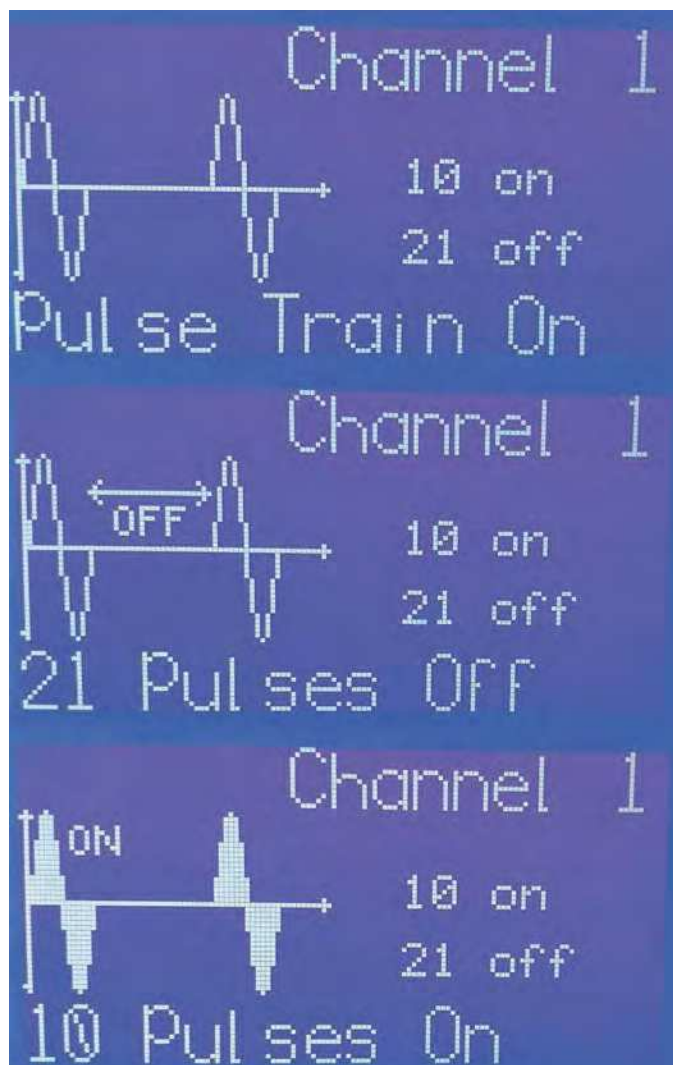
Screen 6: the unit can be set to produce a continuous waveform or a pulse train. With a pulse train, you can have 1-1000 pulses followed by 1-1000 gaps. Each pulse or gap is the set waveform period (the reciprocal of the frequency), eg, 1ms for 1kHz, 10ms for 100Hz...

Conclusion

All the PCBs to build this project are already available from the *PE PCB Service.* As we mentioned earlier, they were also used in the *DSP Active Crossover/Parametric Equaliser* project. They are coded 01106192-01106196.

Many of the parts are SMDs, some of them fine-pitched, so this is not an ideal project for beginners.

Building this will provide you with a test instrument with excellent performance, that can carry out many important audio tests, especially for Hi-Fi.

You could even combine this *Low-distortion DDS Signal Generator* with a high-quality ADC connected to a computer, feeding into spectrum analysis software, as an audio analyser capable of measuring distortion in devices like amplifiers, preamplifiers and filters, down to very low levels.

It can also be used in conjunction with our *High-resolution Audio Millivoltmeter* project, published in the October 2020 issue of *PE*, for making signal-to-noise ratio measurements.

That would be a great way to make and save frequency response plots, with the *Low-distortion DDS Signal Generator* in sweep mode and the *High-resolution Audio Millivoltmeter* connected to a USB port, for logging the results, which could then be fed to a plotting program.
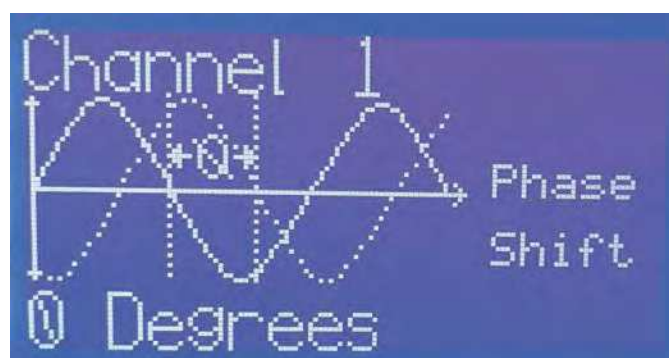
Screen 7: finally, if you wish you can set a phase offset between the two channels, although this only makes sense if they are set to the same frequency.

INDOOR 'POOR AIR QUALITY' MONITOR

By Geoff Graham

Is the air in your office or home 'stuffy'? That can be subjective – but there is an objective way to measure air quality – using a volatile organic compound (VOC) meter. The Micromite BackPack and a cheap module make building one dead easy! Now you really can find out if you are being hit by 'Sick Building Syndrome'.
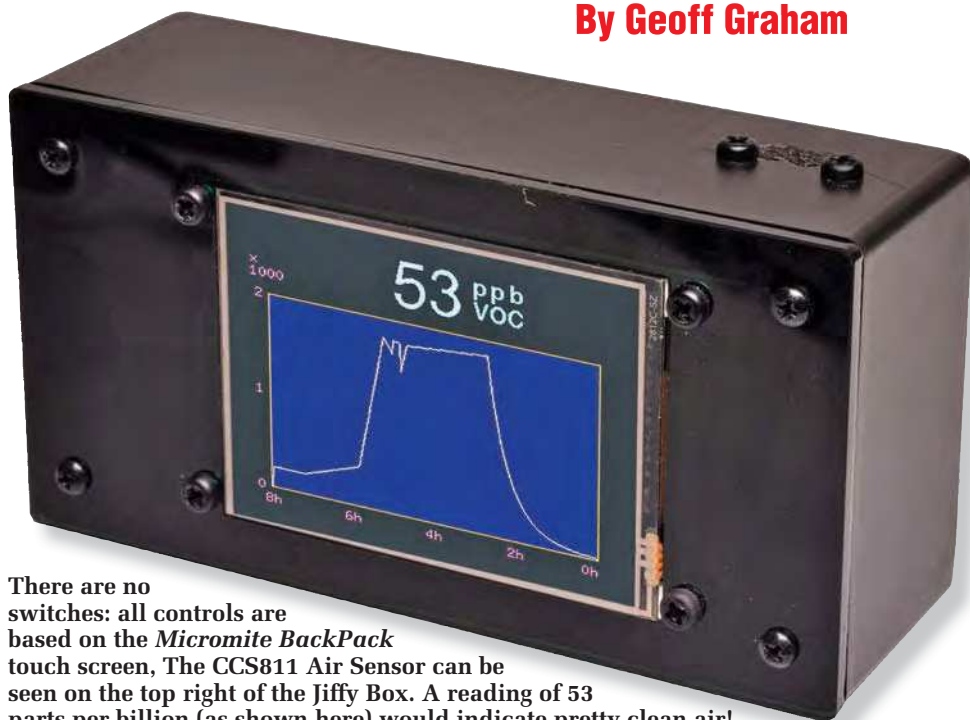
There are no switches: all controls are based on the *Micromite BackPack* touch screen, The CCS811 Air Sensor can be seen on the top right of the Jiffy Box. A reading of 53 parts per billion (as shown here) would indicate pretty clean air!

The amount of volatile organic compounds (VOCs) in the air is a fundamental measure of air quality. VOCs include thousands of chemicals that can be present in the air, many of which are recognised to have a significant impact on the health of people breathing them in.

This monitor uses a CCS811 MOX (metal oxide) sensor made by ams AG, Austria, to measure the total VOC level in the air. This sensor comes in a tiny surface-mount package that is very difficult to solder.

Fortunately, it can be purchased as part of a low-cost (£7.50) fully assembled module which, when coupled with a *Micromite LCD BackPack*, makes a capable *Air Quality Monitor*.

The VOC reading varies with temperature and humidity, so the sensor module also includes temperature and humidity sensors, allowing it to compensate for variation in both.

Once per second, the Micromite's BASIC program reads these values and feeds them to the CCS811 sensor, which then uses them to adjust its VOC reading to maintain accuracy.

The BASIC program then extracts the VOC reading and displays it as a number, expressed in parts-per-billion (ppb). It also draws a graph on the LCD

MOX sensors

MOX stands for 'Metal OXide' which, in a sensor such as the CCS811, is a thin film of a metal oxide such as tin oxide (SnO_2) in a porous granular state.

In clean air, oxygen is adsorbed on the surface of the metal oxide and this attracts free electrons in the material to the surface which, in turn, has the effect of increasing the resistance of the MOX layer.

In the presence of a reducing gas (eg, a volatile organic compound), the oxygen is pulled from the MOX surface to react with this gas. This frees previously trapped electrons, causing a measurable decrease in resistance.

To assist in this process, the MOX layer is raised in temperature by a built-in heater.

In a modern sensor, the whole structure is built on a semiconductor substrate using normal manufacturing processes.

By varying factors such as the type of oxide, the grain size and the heating temperature, the manufacturer can adjust the sensitivity to suit various reducing gases. The CCS811 sensor used in our Air Quality Monitor is especially sensitive to alcohols, aldehydes, ketones, organic acids, amines, aliphatic and aromatic hydrocarbons. These are typically produced by humans and building materials.

Humidity has a strong influence on the performance of metal-oxide gas sensors. Reactions between the surface oxygen and water molecules cause a reduction in the resistance of the MOX layer, reducing its sensitivity. Temperature also has an effect.

This is why the module used in our Air Quality Monitor includes a temperature/humidity sensor.

The program running on the Micromite reads the temperature and humidity values once per second and transfers them to the CCS811 sensor. The internal algorithms of the CCS811 then use these to adjust the readings accordingly.

The sensor needs an initial burn-in period of 48 hours to remove manufacturing contaminants from the surface. It also needs a 20-minute warm-up period whenever power is applied.

The sensor can become contaminated, so it has a limited lifetime. According to the manufacturer of the CCS811, this is at least five years
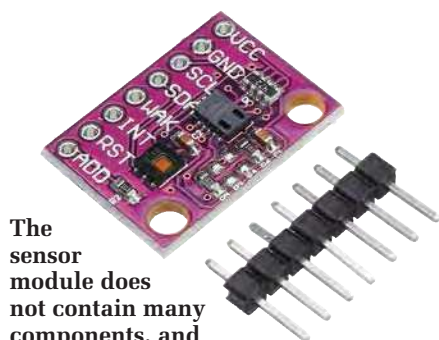
The sensor module does not contain many components, and so is quite tiny (it's shown here about twice life size). Search eBay or AliExpress for the two keywords CCS811 and HDC1080. Many of the available modules will match either of the keywords but not both, so make sure that the module you purchase matches this photograph and has both the CCS811 and HDC1080 sensor ICs.

screen, so that you can see the trends in the reading.

Other functions of the BASIC program allow you to set the baseline for the VOC reading (ie, essentially set the zero reading) and even upgrade the firmware running in the CCS811 sensor.

We've described a few different versions of the *LCD BackPack* in past issues of the magazine. You can use any of them for this project. However, we recommend that you build the V2 or V3 BackPack for this project, as those versions allow the screen brightness to be controlled by the BASIC program.

But note that the software is written with the 2.8-inch 320×240 LCD screen in mind, so if you build V3 with a larger 3.5-inch 480×320 display, you will have to modify the software to suit.

The program controlling the *Air Quality Monitor* is written in the easy-to-use BASIC programming language, so you are free to get in there and modify it to suit your preferences. However, we don't suggest that you undertake the job of modifying the software for the 3.5-inch screen unless you have some prior MMBasic programming experience.

By the way, if you are designing on another project based on the *Micromite LCD BackPack*, you may wish to extract segments of the *Air Quality Monitor* program for your own needs. For example, the graph drawing algorithms could come in handy for many other tasks.

Volatile organic compounds

The side box describes how a MOX sensor works; essentially, it measures organic carbon molecules that are in vapour form suspended in the air. Many are given off by humans, and the CCS811 is particularly sensitive to these.

In scientific literature, the term 'VOC' (volatile organic compound)

refers to any one of the thousands of organic (ie, carbon-containing) chemicals that are present in the air. These are mostly gases at room temperatures. The list includes both man-made and naturally occurring chemicals.

The amount, or concentration, of VOCs present is expressed in a variety of units, but in the case of our sensor, it is reported in parts per billion (ppb). For example, if the concentration is 10ppb, then for every billion molecules of air, there are ten volatile organic compound molecules.

Sources of VOCs include plants, manufactured products (such as plastics) and animals (including humans).

An important subset of VOCs is semi-volatile organic compounds, which come from building materials, furnishings, cleaning compounds, air fresheners, pesticides and activities such as tobacco smoking and cooking with a gas stove.

Some of the key indoor sources of semi-volatile organic compounds are pesticides, building or decorating materials containing flexible plastics such as vinyl wallpaper or vinyl flooring and materials containing flame retardants.

One common VOC that has been recognised as having significant health impacts is formaldehyde, which is widely used in the manufacture of building materials and household products. It is also a by-product of combustion and other natural processes.

The CCS811 sensor used in our *Air Quality Monitor* reports on total volatile organic compound concentrations. This term refers to the concentration of many different VOCs that are present simultaneously in the air, and the CCS811 is more sensitive to the subset of VOCs that are typically caused by human activity.

Taken together, the total VOC reading indicates the quality of the air that we breathe.

For example, a reading of zero indicates clean, fresh air. A reading of 4000ppb to 8000ppb (4-8ppm) would indicate a stuffy room, while a reading of over 16,000ppb (16ppm) would indicate a particularly bad environment.

Note that the CCS811 does not claim to be highly accurate; in fact, the data sheet talks in terms of 'indicated' levels of VOCs in the air.

Fig.1: the *Air Quality Monitor* circuit consists of just two modules. These are the *Micromite LCD BackPack* and the sensor module, with the CCS811 and HDC1080 ICs onboard. The sensors communicate via I^2C, so the sensor module connects to the SCL and SDA pins on the BackPack.
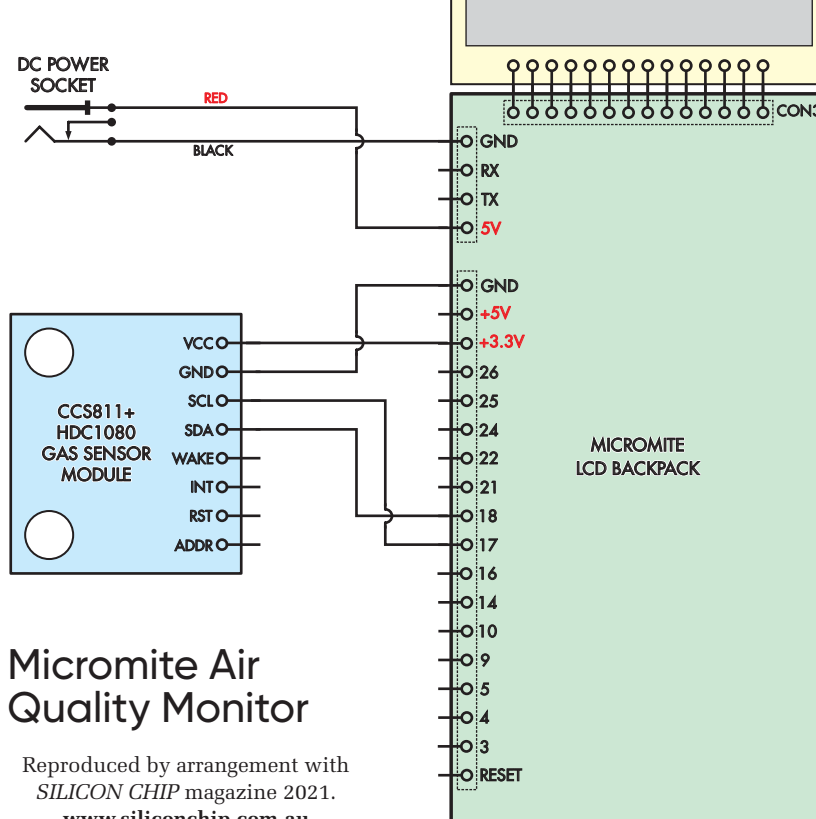
Micromite Air Quality Monitor

Reproduced by arrangement with *SILICON CHIP* magazine 2021.
www.siliconchip.com.au

So, if you are thinking of suing your employer over a sick building environment, you would need to employ much more accurate equipment that could measure specific chemicals (and hire some pretty good lawyers!).

Regardless, the CCS811 is quite sensitive. We tested it in a large room without much ventilation and with two people present, the VOC reading climbed to more than 3000ppb over a couple of hours. Opening a window just a crack caused the reading to drop back to near zero within half an hour.

Circuit description

The circuit of the *Air Quality Monitor* basically consists of just two modules connected together, as shown in Fig.1. These are the *Micromite LCD Back-Pack* and the sensor module (with the CCS811 and HDC1080 chips).

Both the CCS811 and HDC1080 communicate via I^{2}C, so the sensor module's SCL (clock) and SDA (data) pins are wired to the corresponding I^{2}C pins on the *Micromite LCD Back-Pack*. The module runs from 3.3V, so its VCC pin goes to the 3.3V output on the *BackPack*, and the GND pin is the common ground.

The other pins on the sensor module (for example, WAKE or ADDR) are not used in our application so they can be left unconnected.

The sensor module does not contain many components besides the CCS811 and HDC1080 ICs; just a few pull-up resistors and bypass capacitors. It is also surprisingly small, barely large enough to cover the tip of a finger.

The module that we used can be found on eBay or AliExpress by searching for the keywords CCS811 and HDC1080 together.

There are many modules on offer that will match either of the keywords, so you need to make sure that your module includes both sensors, and preferably looks identical to the one shown here.

This last part is important as there are many sensor modules circulating that claim to incorporate both sensors, but the images displayed by the vendor show that the module does not include the HDC1080 temperature and humidity sensor. This is critical to the correct operation of the device.

Construction

Construction is quite simple and consists of just assembling the *Micromite LCD BackPack* (which should take less than an hour) then mounting and connecting the sensor module.

All three versions of the *Micromite LCD BackPack* are available. V1 and V2 are available in kits from our preferred Micromite retailer: **https://micromite.org/product-category/kits/**

A PCB for V3 is available from the *PE PCB Service* (we no longer sell V1 or V2 PCBs).

The *Micromite LCD BackPacks* comprise about a dozen components, and in each case, the PCB is printed with the component placement and values. So you simply populate the board and attach a suitable LCD panel.

We have reproduced the *Micromite LCD BackPack* V1 and V2 PCB overlay diagrams below in Fig.2 and Fig.3 in case you need them.

Full descriptions of building V1, V2 and V3 appeared in *PE*, May 2017, May 2018 and August 2020 respectively.

Note that CON1 and CON2 go on the underside of the board. If in doubt when it comes to assembling the *Micromite LCD BackPacks*, refer to the relevant constructional article referred to earlier.

All three versions of the *Micromite LCD BackPacks* fit neatly into a standard UB3 plastic box (see parts list).

For the neatest appearance you can purchase a laser-cut acrylic lid with cutouts for the LCD and its mounting screws. This lid is available from the SILICON CHIP ONLINE SHOP for A$5 plus shipping. It's supplied in matte/gloss

The Micromite-based Air Quality Monitor before it is put into its Jiffy box,
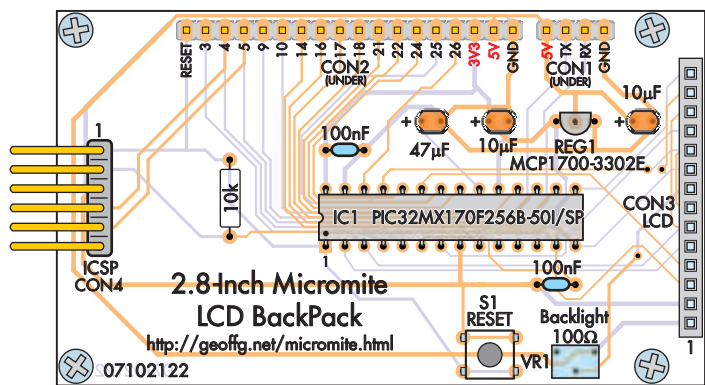
Fig.2: if you're building the *Air Quality Monitor* around the original *Micromite LCD BackPack*, this overlay shows how to fit the components. The orientations of IC1 and REG1 are critical. You may have polarised or non-polarised high-value capacitors. Also, a 10µF capacitor can be used in place of the 47µF capacitor if it's a ceramic type. CON1 and CON2 are fitted to the underside of the board.
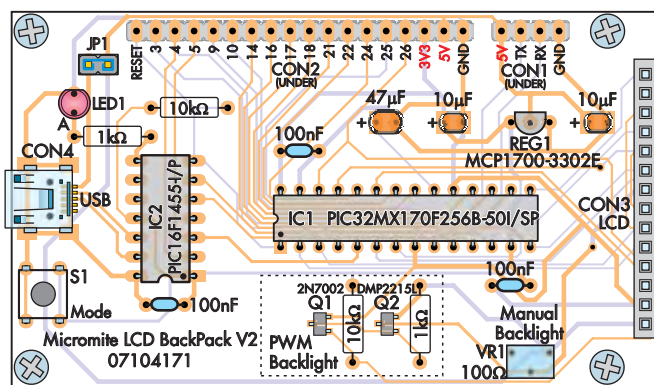
Fig.3: use this overlay diagram instead if you're building the *V2 BackPack*. The main difference is the addition of the Microbridge, IC2, which also must be oriented correctly. You may receive three 10µF ceramic capacitors (their orientations are not important) and one can be used in place of the 47µF. If you're building a *V3 BackPack*, refer to the August 2020 issue or just follow the PCB silkscreen.
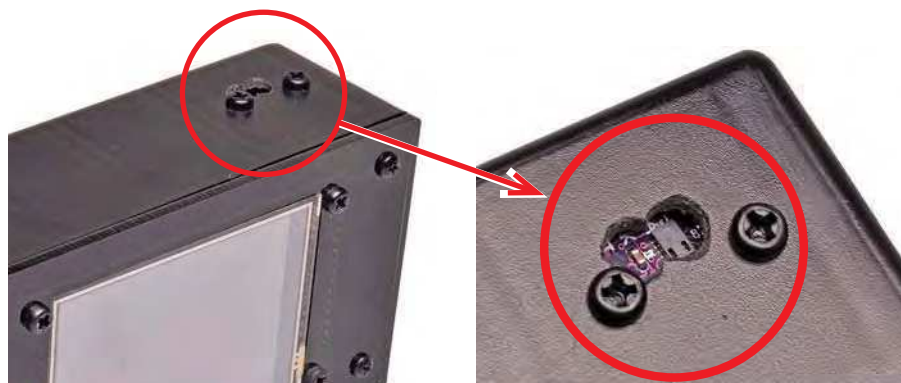
If using the laser-cut lid (see text), there are only five extra holes required in the UB3 Jiffy box – one for the power socket, as seen at left, and the others for mounting the CCS811 sensor module. Two of these holes are drilled to form a 'figure 8' hole which allows plenty of air circulation to strike the sensor module (see right). The project can be powered from a 5VDC regulated plugpack, from a computer's USB socket or even from a 5V 'powerbank' to allow fully portable operation.

black (Cat SC3456; **www.siliconchip. com.au/Shop/19/3456**) or clear (Cat SC3337; **www.siliconchip.com.au/ Shop/19/3337**).

The laser-cut black lid can be fitted either way around, so you can have either a matte or gloss finish on the lid's outside.

Fig.4 provides the box mounting details. The result is a neat looking assembly with the display and Back-Pack securely fastened. The laser-cut panel is thicker than the lid supplied with the UB3 box (3mm), and it lacks recesses, so the self-tapping screws supplied with the box may or may not be long enough. If they're too short, replace them with four 10mm-long 4G self-tapping screws.

On the other hand, there is nothing to stop you (carefully) cutting your own holes in the lid supplied with the UB3 box. It is tricky to make a clean cut around the LCD screen, but can be done. If you choose this route, then make sure that no part of the box is pressing on the surface of the LCD as that will upset the touch sensitivity of the panel.
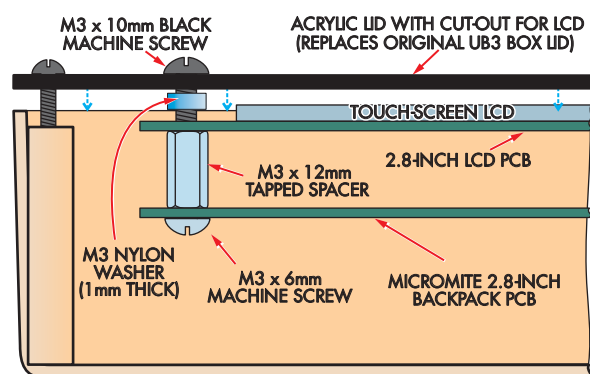
Final assembly

Use an M3 × 10mm machine screw on each corner with a 1mm thick plastic M3 washer and an M3 × 12mm tapped spacer to attach the LCD panel to the (acrylic) lid.

This ensures that the surface of the LCD will be flush with the acrylic lid. Then, the *BackPack* can be plugged into the LCD and fastened by M3 × 6mm machine screws to each spacer, as shown in Fig.4.

The module with the CCS811 and HDC1080 sensors typically comes with an eight-pin header that is not soldered to the board. So you need to fit this, but make sure that the header pins point out from the back (non-component) side of the board.

After you have cut two holes in the box (for the two sensors), this arrangement will allow you to mount the component side of the sensor board close to the wall of the box. The sensors will be next to the holes and therefore sampling the freely circulating air outside the box.

The sensor module should be held in place using two nylon M3 machine screws, with a nut between the sensor PCB and the inside wall of the case, as shown in Fig.5.

This will space the components on the sensor module slightly away from the case wall, but still keep them close enough that they are exposed to the outside air.

We are currently experimenting with several other air quality sensors (ie, CO_2 and particulate sensors). If they work out, we will update the BASIC program to incorporate them in conjunction with the VOC sensor.

The VOC sensor module should be mounted to one side of the case, leaving space for the other modules if later required.

The best method for connecting the sensor module to the *BackPack* is to use 'DuPont' jumper leads. These are lengths of wire with single pin female header sockets on each end. They are designed to slip onto header pins like those of the sensor module and *Micromite LCD BackPack*.

Example leads are those from Altronics (Cat P1017) or Jaycar (WC6026), or just search eBay for 'dupont jumper'.

Using these not only makes assembly easy, but you can also replace the sensor module for testing or if the CCS811 chip becomes contaminated.

Loading the software

The program running on the Micromite consists of three parts: the MMBasic interpreter, the BASIC program for the *Air Quality Monitor* and the configuration settings for the LCD panel.

The easiest method of loading all three at once is to program the PIC32 chip with the file **AirQuality.hex** which can be downloaded from the February 2021 page of the *PE* website.

This is easy if you are using either a V2 or V3 *Micromite LCD BackPack*, as both of these include the capability to program the PIC32 chip (ie, Micromite) with new firmware. If you are using a V1, then you will need a

Fig.4: this shows how the touchscreen, *BackPack* PCB and laser-cut lid go together before the lid is attached to the UB3 Jiffy box base using four self-tapping screws. The screws supplied with your Jiffy box may or may not be long enough to go through the thicker laser-cut lid; if not, you will need four No.4 × 10mm (or thereabouts) self-tappers.
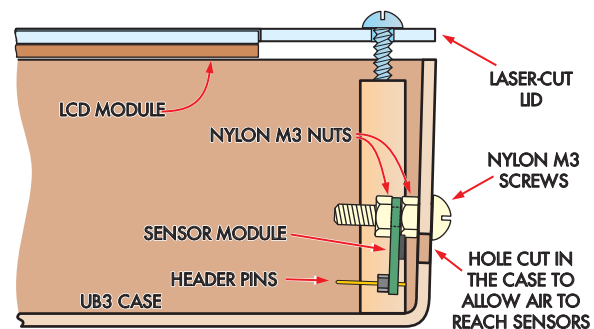
Fig.5: the sensor module is mounted separately on the side of the box and wired to the *BackPack* via its eight-pin header and flying leads. Make sure the two sensor ICs sit adjacent to holes drilled in the side of the box, so they can sample the air outside. Note the nut between the sensor PCB and inside surface of the case, so those sensors are not too close to the holes.

PIC32 programmer such as the PICkit 3 or PICkit 4.

This firmware file contains everything that you need, including the MMBasic interpreter, the LCD configuration and the BASIC program for the *Air Quality Monitor*. So as soon as you load it, you're ready to go.

But be aware that the touch calibration in the combined firmware and BASIC program file was done using a standard LCD panel. Therefore, your unit might require display recalibration if it is significantly different from the one that we used. Unfortunately, we've seen panels with the touch sensor rotated 180° from others, and they are not easy to tell apart!

Luckily, recalibration can be easily performed by connecting the *Back-Pack* to a desktop or laptop computer via USB, opening the serial port in a terminal emulator, halting the program with CTRL-C and re-running the calibration routine.

For further information on this procedure see the *Micromite User Manual* (which can also be downloaded from: **https://geoffg.net/micromite.html**

Alternatively, if you have a *Back-Pack* that's already programmed with the Micromite firmware, you can set up the LCD screen (if that hasn't already been done), then load the *Air Quality Monitor* BASIC code into it.

This file is named **AirQuality.bas** and is part of the same download package for this project, from the February 2021 page of the *PE* website.

Power supply

The completed *Air Quality Monitor* requires a 5V power supply with a minimum capacity of 500mA.

You can use a 5V plugpack or a USB charger. If you are using a plugpack, make sure that it is regulated and that its unloaded output does not rise above 5.5V, as that could cause damage.

We built the prototype with a DC power socket (the barrel or 'concentric' type) for the incoming power, mounted on the side of the UB3 box. The two flying leads from this socket were fitted with female header sockets and slipped over the *BackPack's* header pins.

The centre pin of the socket should go to the 5V pin on CON1, while the sleeve should connect to the pin marked GND (ground).

That matches up with the most common plugpack wiring scheme, where the inner barrel conductor connects to the positive supply wire.

If you are using a USB charger as the power supply, you can make up a power cable by cutting off one end of a standard USB cable while retaining the Type A socket on the other end, as shown in Fig 6. Then solder the free end to a suitable DC power plug. Most sockets have either a 2.1mm or 2.5mm inner pin, so make sure your plug matches it; 2.5mm plugs will go into 2.1mm sockets but won't make good contact!

The red wire in the USB cable (+5V) should go to the centre pin of the plug and the black to the sleeve. The other two wires (the signal wires) can be cut short as they are not used.

Alternatively, suitable ready-made cables are available, such as a Jaycar PP1985 or Altronics Cat P6701 (USB Type A Male to 2.1mm DC Plug).

Testing

Before connecting the sensor board to the *Micromite LCD BackPack*, you should confirm that the *BackPack* itself is working correctly. The testing procedure for this is described in the *Micromite User Manual* and the relevant *PE* articles referenced above.

It's then simply a matter of connecting the sensor board and powering up the whole gadget.

If it does not work straight away, carefully check and re-check each connection. Then measure the voltage across the pins marked VCC and GND on the module; you should get a reading very close to 3.3V.

Using the device

Using our *Air Quality Sensor* is quite straightforward. You plug it into a source of 5V DC power, and after a warm-up period, it displays the air quality as a number and draws a graph showing how it changes over time.

The VOC reading is displayed in parts per billion (ppb), in large digits

---

---

at the top of the screen, with the time-based graph below (see Screen1).

The graph's vertical axis can be configured in steps from a very sensitive 500ppb full scale to 64,000ppb, while the horizontal time scale can be set to cover from 15 minutes to 32 hours.

When you build the *Air Quality Sensor* and turn it on for the first time, there is a burn-in period of 48 hours that you need to observe.

This is necessary as the CCS811 sensor's readings will drift considerably as surface contaminants from the manufacturing process are burnt off from the sensitive metal-oxide layer.

So, when you first turn it on, leave it powered up and running for at least two days before taking any readings. You might be tempted to interpret

---

Fig.6: here is our suggested method for powering the *Air Quality Monitor* from a USB port or charger. You can make up the USB-to-DC-plug cable, or you can buy one from Altronics (Cat P6701) or Jaycar (PP1985). Alternatively, use a 5V DC regulated plugpack which will most likely already have a concentric plug with the right polarity (+ to centre).

USB CONNECTOR TYPE A MALE
DC PLUG
DC INPUT SOCKET (ON END OF BOX)
4-PIN FEMALE HEADER CONNECTOR (OR DUPONT POWER CABLE)
MICROMITE CON1 POWER AND CONSOLE CONNECTOR
5V / Tx / Rx / GND

Screen1: the main display shown at power-up, with the current VOC reading at the top and a graph of VOC over time at the bottom. The graph's vertical scale can be configured from 500 to 64,000ppb, while the horizontal scale can be adjusted to cover 15 minutes to 32 hours.

Screen2: touching the main screen takes you to this setup screen. Here you can adjust the graph parameters, set the VOC baseline (the zero point for the CCS811 sensor), upgrade the CCS811's firmware and set the screen brightness. The raw data from the CCS811 and HDC1080 sensors is shown at the top of this screen.

Screen3: this screen allows you to set the baseline (zero point) for the CCS811 sensor. Initially, this should be done once a week, but after a couple of months, the baseline will only need to be set once a month.

its readings during this period, but it will be meaningless until the burn-in period has run its course.

Another requirement of the CCS811 is that it needs a 20-minute warm-up period every time power is applied. This is different from the initial burn-in period and applies whenever the *Air Quality Monitor* is turned on.

During this period, the BASIC program displays a countdown. You can skip it if you wish, but like the original 48-hour burn-in, it is much better to let the sensor stabilise. Otherwise, the readings could be nonsense.

Setup screen

To access the setup screen from the main screen, touch anywhere on the LCD panel.

This will take you to a display where you can see the raw readings from the sensor and configure things like the graph's parameters (see Screen2).

Raw data from the sensor module is shown at the top of this screen. The VOC reading is the same as that on the main screen, but there is also a CO_2 reading. This is an estimate of the amount of CO_2 in parts-per-million (ppm) that would be present in the air if the measured VOCs were created by human respiration.

The CCS811 sensor calculates the CO_2 reading, but it does not necessarily relate to the actual level of CO_2, because the VOC reading used for this estimate could be partly or wholly due to other processes (eg, paint drying). Regardless, the CCS811 sensor produces this reading so we display it on this screen for you.

The temperature (°C) and humidity (%RH) readings come from the HDC1080 temperature/humidity sensor that is included on the sensor module. They are used by the CCS811 to give an accurate VOC reading, as mentioned earlier.

The temperature reading made by the HDC1080 is supposedly accurate to within ±0.2°C, but we found that in practice, it was high by a few degrees. This is likely due to its proximity to the CCS811 sensor, which has an on-board heater.

Despite this, it is accurate enough for its primary purpose, which is to allow the CCS811 to compensate for variations in temperature and humidity.

Graph parameters

There are two buttons below the raw readings on the setup screen that allow you to change the parameters for the graph on the main screen.

The 'Vert Scale' button shows the current vertical scale for the graph in parts-per-billion (ppb). Repeatedly touching this button will step you through a sequence of full-scale values from 500ppb to 64,000ppb.

This setting is automatically saved by the BASIC program and will be reloaded when the *Air Quality Monitor* is next powered up.

Similarly, the 'Time Scale' button selects the horizontal time scale for the graph. Repeatedly touching this button will step you through various time scales from 15 minutes (full scale) to 32 hours. This setting is also saved for the next power-up.

There are also buttons on the setup screen to set the baseline (zero) reading, update the firmware in the CCS811 and set the screen brightness. Finally, to exit the setup screen, touch the 'Exit' button to return to the main screen.

Setting the baseline

The CSS811 documentation refers to 'Manual Baseline Correction', which in effect means determining the zero point for the VOC reading.

The MOX sensor used in the CCS811 (see the earlier panel for a description)

can be contaminated over time, causing the zero point to drift.

The manufacturer recommends that the baseline should be set once a week for the first couple of months of use, and from then on, the baseline will only need to be set once a month.

To set the baseline, place the device outside in clean air and touch the 'Baseline' button on the setup screen (see Screen3).

The BASIC program will step through this process which involves waiting for the sensor to stabilise from the power-on condition (20 minutes), then allowing the sensor to determine the baseline over a 10-minute period.

This baseline is saved in non-volatile memory by the BASIC program and copied to the CCS811 every time the power is turned on.

This is necessary because, without this bit of information, the CCS811 will essentially be forced to 'guess' the baseline.

At the conclusion of this process, the *Air Quality Monitor* will return to the main screen showing the reading and graph.

Updating the CCS811 firmware

The CCS811 sensor is quite a complicated device, and it includes a microcontroller, which is used to measure the resistance of the MOX sensor, control the heater and many other functions.

At the time of writing, the latest version of the firmware for the CCS811 is V2.0.1. However, many modules manufactured in China are still using sensors running V1.1.0 firmware or even earlier versions.

The V2.0.1 firmware incorporates an improved algorithm for the VOC calculation, and the range of readings has been extended to 64,000ppb VOC (the old firmware limited the sensor to a maximum of 1187ppb VOC).

This firmware can be updated by the BASIC program running on the

Screen4: you can update the CCS811 sensor firmware to the latest version (V2.0.1) using this screen. Touching the 'Update' button initiates the upgrade (takes under 30s). You don't need a firmware file or anything else for this operation; everything is contained within the Micromite BASIC program.

Screen5: this screen allows you to set the display brightness and also the auto-dimming feature, which has the benefit of reducing the unit's power consumption. If your *BackPack* only supports manual display brightness adjustment (ie, using a trimpot) then these settings will do nothing.

Micromite. This is done by touching the 'Firmware' button on the setup screen. You will then be taken to a screen which displays the current version of the firmware running on the CCS811 and an offer to update it (see Screen4).

Touching the 'Update' button will initiate the upgrade process, which takes less than 30 seconds. Note that you do not need a firmware file or anything else for this operation; everything is contained within the BASIC program running on the Micromite.

After this process, the *Air Quality Monitor* restarts with the new firmware running in the CCS811.

Setting the screen brightness

As mentioned above, recent versions of the *Micromite LCD BackPack* (V2 onwards) include the ability to control the brightness of the screen from within the BASIC program.

You can control this by selecting the 'Brightness' button on the setup screen. If your *LCD BackPack* only supports manual adjustment of the brightness (ie, a trimpot), this setting will do nothing.

On the brightness screen (shown in Screen5), repeatedly pressing the 'Brightness' button will step you through a range of brightness levels from 10% to 100% in 10% steps. If you enable the 'Auto Dimming' checkbox,

you can set a time period and a second brightness level which applies after no activity.

This is useful if you are using the *Air Quality Monitor* in a bedroom at night, or if you are running it from a battery.

The auto-dimming function operates when the main screen (with graph) is displayed and if it has dimmed, touching anywhere on the screen will restore full brightness. A second touch will then take you to the setup screen.

Updating the BASIC program

One of the great features of the Micromite is that it is easy for you to get in there and modify or update the BASIC program that provides this instrument with its unique functions. This program is stored on the chip in clear text, so you can do things like change colours, menu choices and other features as you wish.

If you are using V2 or V3 of the *Micromite LCD BackPack*, this is as easy as plugging your desktop PC or laptop into the USB socket and running a terminal emulator on your computer.

Typing CTRL-C into the terminal editor will interrupt the running program and display the command prompt at which point you edit the program using the EDIT command. This is covered in detail in the *Micromite User Manual*, so we will not go into detail here.

If you have an earlier version of the *Micromite LCD BackPack*, you will need a separate USB-to-serial converter (they are cheap). All the details on this are also in the *Micromite User Manual*.

Battery operation

You might want to power the *Air Quality Monitor* from a battery, so it's truly portable. This would allow you to make a quick survey of a large office space or house.

The best option for this is to use a USB 'power bank' as sold for recharging mobile phones while on the go. These have everything that you need in a portable power source, including a charging circuit, protection circuits and a regulated 5V output.

Even better, because they are a common item in mobile phone shops, they are quite cheap. They cost much less than the parts that you would need to build a similar device yourself!

The most significant power drain in the *Air Quality Monitor* is the LCD screen backlighting. That alone can consume up to 250mA at full brightness.

This is one of the reasons for the auto-dimming feature described earlier; with that enabled, you can reduce the brightness of the display to (say) 10% after a short period of inactivity.

This reduces the current drawn by the backlight to about 25mA, essentially halving the unit's power consumption.

For the record, the Micromite draws about 26mA and the CCS811/HDC1080 sensors combined about the same. So, with the displayed dimmed to 10%, the total drain on the battery should be about 80mA, or about 2Ah/day. A 5000mAh power bank should last for two days of continuous operation.

However, consider that the actual energy delivered is lower than rated, due to the difference between the average battery voltage of 3.7V and the output voltage of 5V, and the voltage conversion is less than 100% efficient.

Don't forget that many cheap power banks grossly overstate their capacity!

For truly portable use, you would ideally incorporate the power bank into the case. That would require you to use a larger box and to add an on/off switch.

Some of the smaller cylindrical power banks could fit into the UB3 box and still provide sufficient capacity for many hours of use. The details of these modifications are something that we will leave as an exercise for the reader.

Firmware updates

For firmware updates for the Micromite and the BASIC program for the *Air Quality Monitor*, check the author's website at **http://geoffg.net**

<div style="border:2px solid red; background:linear-gradient(pink,yellow);">

Micromite LCD BackPack versions

There are three generations of the Micromite LCD BackPack, and all will work in the Air Quality Monitor using the same software. The main difference in this application is that Version 1 only has manual brightness control, so the Air Quality Monitor firmware can not control its brightness.

The others (V2 and V3) have optional software control of the LCD backlighting, so if the appropriate components are installed, you can adjust its brightness via the settings screen. This also enables the auto-dimming feature.

The V2 and V3 BackPacks also have an onboard USB/serial and PIC32 programming interface called the Microbridge.

</div>

USB Logic Analysers

This 8-channel USB logic analyser is essentially a clone of the original version of the well-known and respected Saleae Logic unit. It's completely compatible with the Saleae design, but you can get this one for as little as £8 – less than one-fortieth the cost!

Some of you may be familiar with the range of Saleae USB Logic Analysers. Although most models only have eight channels, on the whole, they are excellent performers and probably suitable for 95% of the analysis work done by hobbyists, technicians and designers.

They start from several hundred pounds – not nothing, but still only a fraction of the cost of getting a 'professional' logic analyser, which makes it an attractive purchase – particularly since Saleae allow you to download the matching control and analysis software for free.

I bought an early 24 megasamples per second (MS/s) eight-channel 'Logic' model, and when Saleae brought out a 16-channel version, I bought one of those too. And I've been very happy with them both, especially as Saleae has kept their software up to date.

But recently, USB logic analysers like the original Saleae Logic began to appear on the world market, at a much lower cost. They all seem to be made in China, and many of them are virtual 'clones' of the original Saleae Logic. They're all 8-channel units with a maximum sampling rate of 24MS/s, and in most cases, they work exactly the same as the Saleae Logic when hooked up to your PC.

Even Saleae's software can't tell the difference. What was that old saying about, 'imitation being the sincerest form of flattery'? However, Saleae's 'Terms of Service' says that you may not use their software with unauthorised clones of their products (**www.saleae.com/terms-of-service**). But there is a alternative, which we'll get to a bit later.

At the time of writing, there are quite a few of these Saleae Logic clones available on various online sites, at prices varying between £6 and £20 depending on the supplier and their location. Sources include Banggood in China, Sparkfun in the US, Sparkfun's suppliers in the UK and quite a few others you can find on eBay (item 174534651401) and AliExpress

I should point out that Saleae no longer makes the original 8-channel 24MS/s logic analyser, the one I bought. Instead, they offer the Logic 8, which samples at up to 100MS/s, priced at £299 direct from Saleae in the US (**www.saleae.com**). They also offer the Logic Pro 8, which provides 500MS/s sampling (£529). Plus they sell the Logic Pro 16, offering 500MS/s sampling of 16 channels (£749). (The above prices are plus VAT and shipping). So Saleae has moved up-market, leaving the lower end to the clones.

Cheapest and most popular

Unsurprisingly, the most popular of these clone logic analysers is one of the cheapest. It's available from Banggood for approximately £8 (Banggood 1177821; **https://bit.ly/pe-feb21-bg**).

Inside the logic analyser (shown enlarged for clarity), we found a board with a 56-pin Cypress Semiconductor USB-capable micro, a 24C02 EEPROM, octal bus transceiver chip, regulator, crystal, transistor, two LEDs and a handful of passive components. The bus transceiver operates as a level shifter and also protects the micro inputs.

For comparison, here's the original Saleae Logic, purchased a few years ago, along with the supplied (functional!) USB cable and high-quality test clips.

This includes a USB type-A to mini type-B cable (the analysers are powered via the USB cable) and a set of 10 ribbon cable test leads, with female pin sockets at each end.

The US$19.95 Sparkfun clone appears to be identical (**www.sparkfun.com/products/15033**), apart from a more attractive label on the top.

This device is housed in a plastic box measuring 55 × 27 × 14mm, which appears slightly smaller than the original Saleae Logic unit. That came in a nicely finished machined metal box measuring 41 × 41 × 8.5mm. Apart from this, the main physical difference is that the cheap unit has a 10-pin DIL connector for the test leads, whereas the original Saleae Logic had a 9-pin SIL connector.

Another difference is that, while the original Logic (and the later Saleae models) came with a set of nine high-quality E-Z-Hook Micro test clips, the clone units just come with female-to-female ribbon cable leads. If you want test clips, you have to buy them separately – more about this a little later.

Finally, whereas the original Saleae Logic (and later versions) came with a sturdy little storage case for itself, its USB cable, test leads and clips, the clones come without a case.

Electrically, the clones seem to be virtually identical to the original Saleae Logic. When you connect them up to a USB port of a PC, they are recognised as a Saleae unit. And if you fire up Saleae's software, it behaves in precisely the same way as it does with a genuine Logic (but you're not really allowed to do that...)

Alternative software

If – like me – you're not willing to violate Saleae's terms of use and use their software with one of these clone units, there is an alternative.

The open-source software website **http://sigrok.org** has a software package for download called PulseView. This comes in three versions, to suit Windows, Linux and macOS.

For PulseView to be installed on a Windows machine (XP or later), the WinUSB driver must already be installed. If this isn't already installed, you can install it using software called Zadig (**http://zadig.akeo.ie/**), which we have used in the past for driver installation (eg, in the Banggood SDR kit, reviewed in *PE*, November 2018).

I downloaded and installed Pulse-View, and found using it to be quite similar to the Saleae software in operation. You can also download a 24-page PDF user manual from the sigrok website, to help you figure out how to use it (**https://sigrok.org/doc/pulseview/unstable/manual.html**).

PulseView worked equally well with both the Saleae Logic and the clone, and offers an even larger range of protocol analyser options (54 in all), including all the popular ones like I^{2}C, SPI, UART, CAN, I^{2}S, 1-wire, S/PDIF, HDMI CEC, JTAG, MIDI, PS/2 Keyboard/Mouse, USB LS/FS, RGB LED (WS281x) and PWM.

Most importantly, there is no limitation on using PulseView with the clone logic analysers! That, and the fact that it has more features, makes PulseView the way to go.

You can see a typical display from the PulseView software in Screen 1, which again shows the SPI signals on one of my *Arduino-based Audio Millivoltmeter* prototypes (*PE*, October 2020).

As this software works with both the original Logic and the clones (as shown in Screen 2), I was able to make a direct comparison when analysing

An end-on view of the clone logic analyser showing its 10-pin DIL box header and the supplied USB cable, which was missing a pin.

The genuine E-Z-Hook test clips, shown at right, are much easier to use and more reliable than the cheap test hooks I initially purchased, at left. I eventually found a much better-looking set of low-cost test clips at Banggood (see text).
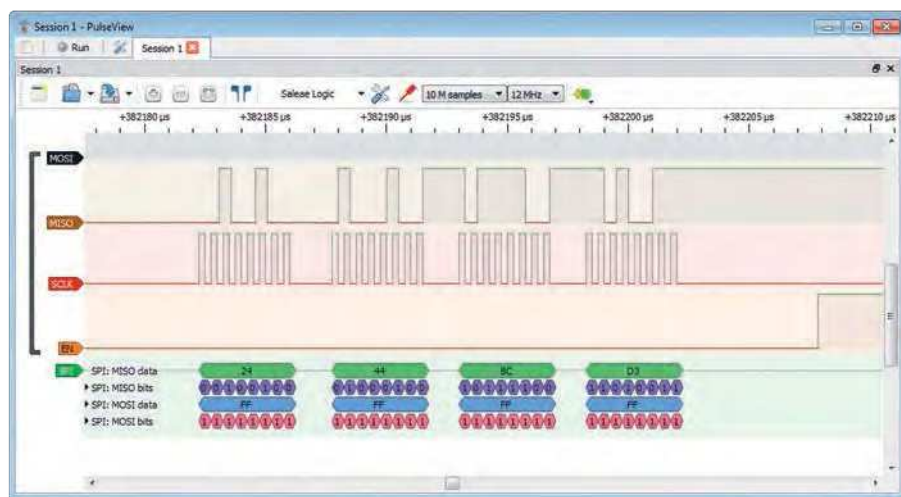
Screen 1: the free PulseView software working with a clone logic analyser to decode SPI serial data. It's easy to use and can decode many different digital protocols.

both I²C and SPI signals on one of the prototypes of my *Audio Millivoltmeter*. I'm glad to report that the results were identical.

Sparkfun have a short tutorial on using their USB logic analyser with PulseView at: **https://bit.ly/pe-feb21-sf**

USB cables and test clips

I ordered two of these low-cost analysers to check them out for this review, and when they arrived, I tried connecting them to a PC using the supplied USB cables. But neither of them were identified by Windows when plugged in. It was only when I tried using another, known-good USB cable that they sprang into life.

When I looked inside the type-B mini plugs on the end of the supplied cables, they were both missing the pin 4 contact! Presumably, that's why the analysers wouldn't work with those cables. That was a bit disappointing, but luckily this type of cable is readily available for only a few dollars (and you're likely to have a few already, as they're pretty common).

As mentioned earlier, these ultra-low-cost USB logic analysers don't come with any test clips – just a set of female-to-female ribbon cable leads. So unless you're only going to be testing signals already routed to pin headers, you will need a set of test clips as well.

The E-Z-Hook XKM-S micro test clips originally supplied by Saleae are available separately, but at a fairly steep price. For example, you can buy a set of 10 from Digi-key (461-1012-ND), but they'll set you back about £20 plus shipping. This doesn't seem all that attractive, since it's around three times the price of the USB logic analyser itself!

I found some likely-looking sets of 10 ribbon cable leads with micro hook test clips on eBay for just over £3, so I ordered a couple of them. But when they arrived, I found that the micro hook test clips were not nearly as well made as the genuine E-Z-Hook clips.

The female connector on the end of the lead would not attach securely to the pin at the back of the clip, and the hook tips didn't seem to have the right shape to mesh properly and attach securely to an IC pin.

So it was 'back to the drawing board' – or more accurately, back to the web. On the Banggood website, I found sets of 'logic analyser test leads with hook clips' that looked more promising.

These have the lead wires soldered to the rear of the hook metalwork inside the plastic housing, removing the problem of connecting the clips reliably to the leads (Banggood 956251; **https://bit.ly/pe-feb21-bg2**). These sets of leads cost just over £5 for 12, or £7.50 for 20. If you don't want to invest in a set of genuine E-Z-Hook clips, these would be a good alternative.
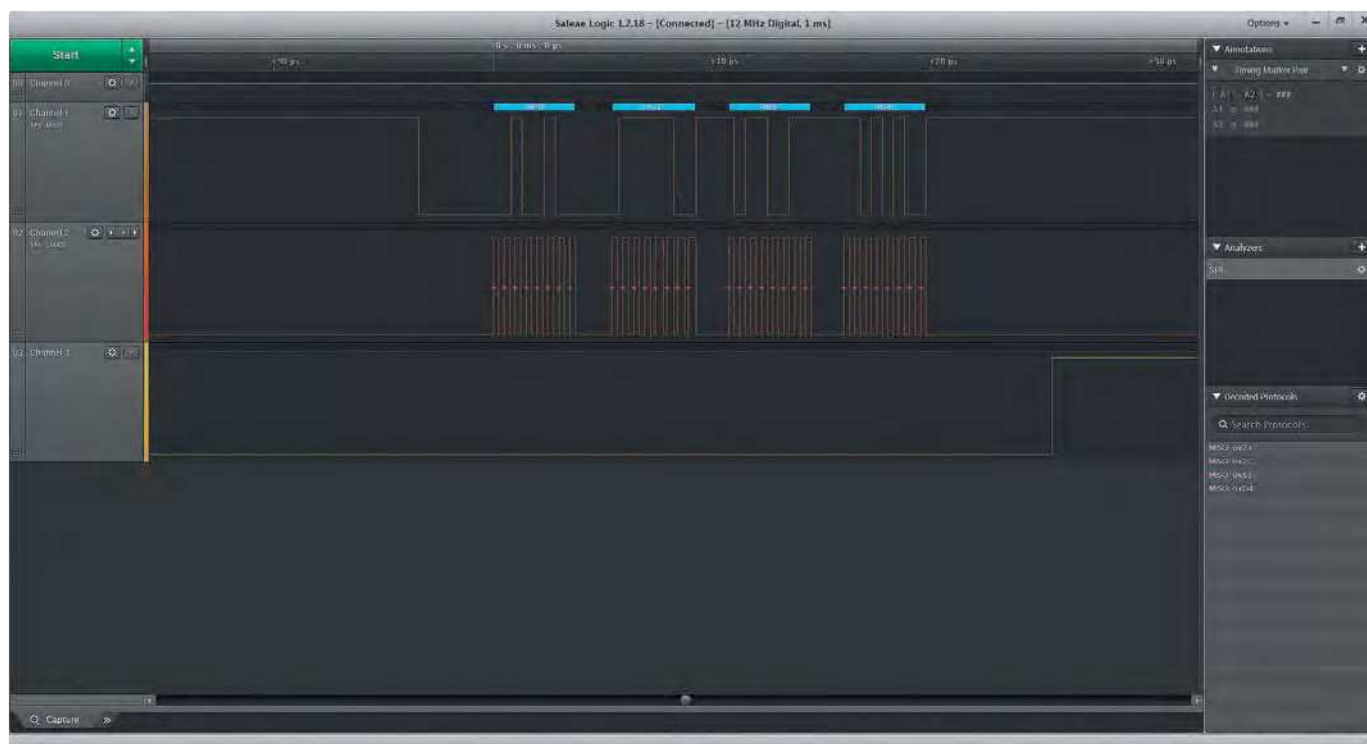
Screen 2: the official Saleae logic software operating with one of their logic analysers to decode the same bus as shown in Screen 1, although there's slightly different data this time. While there are some visual differences, the two pieces of software work similarly.

AUDIO OUT
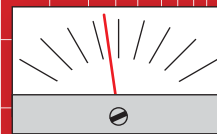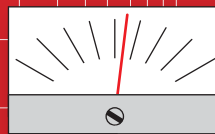
By Jake Rothman

Theremin Audio Amplifier – Part 4

In Part 3 last month, we looked at some germanium transistor versions of the *PE Theremin Amplifier*. This month, we'll detail the component options and see how the different design strategies performed.

Heatsinks

For the silicon amplifier, standard clip-on heatsinks are adequate. Germanium devices, however, need much more heat-sinking. They should run barely warm to obtain a long life. Bolting the devices to the base of a metal box or chassis is sensible. The completed boards and their heatsinking arrangements are shown in Fig.12 and Fig.13.

Specifications

The gain of both amps was 16.7 (24dB). This is about right for running a radio tuner. For CD use, R1 should be increased to 27kΩ. The germanium amplifier worked down to a supply of 5V while it was about 6.5V for the silicon.

Power output

As expected, the germanium amplifier gave more power, providing 880mW RMS ($7.5V_{pk-pk}$). The silicon circuit gave 720mW ($6.5V_{pk-pk}$). This was less than the Siemens circuits because of the addition of an emitter follower (TR2). This stage does allow considerable mismatching of the output transistors due to the low impedance drive. The class-A emitter follower is run at a high current of 17mA to allow for the modest H_{fe} of germanium and old TO5 silicon devices, which can often be as low as 40. Thus, the total quiescent current is around 20 to 30mA, which precludes most battery operation. The maximum current just clipping was 190mA for the silicon circuit and 220mA for the germanium one.

Using a germanium output stage, but with TR1 and TR2 being silicon, gave 810mW ($7.2V_{pk-pk}$). As an experiment, the germanium amplifier was found to give 1.32W ($6.5V_{pk-pk}$) into 4Ω, but only with good heatsinks.

Fig.12. Completed silicon amplifier. Note how bias transistor TR3 is tied to the heatsink of output transistor TR4, providing a degree of thermal stabilisation of the quiescent current.

Distortion

The crossover glitch was softer and shallower with the germanium output stage and less quiescent current was needed (3mA) to smooth it out. The silicon devices needed around 10mA. With no negative feedback the germanium output stage would in theory produce less distortion. However, the lower open-loop gain of the germanium circuit meant there was less negative feedback factor to reduce distortion. This also gave the effect of softer
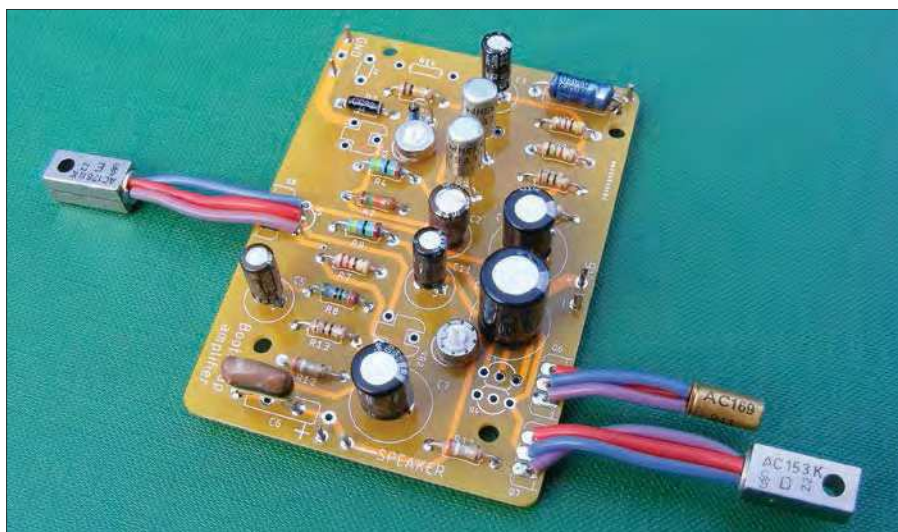
Fig.13. The completed germanium amplifier. The output transistors must be bolted to sufficient metal to keep them cool, almost at ambient. If this is done, the bias transistor can be left unbolted to the heat sink to track the ambient temperature.
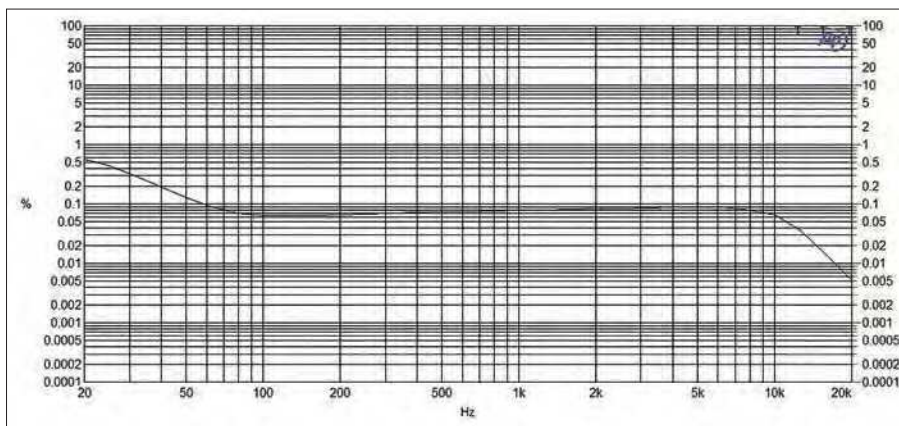
Fig.14. Distortion for the silicon amp: this is at 3V$_{pk-pk}$ into 8Ω.
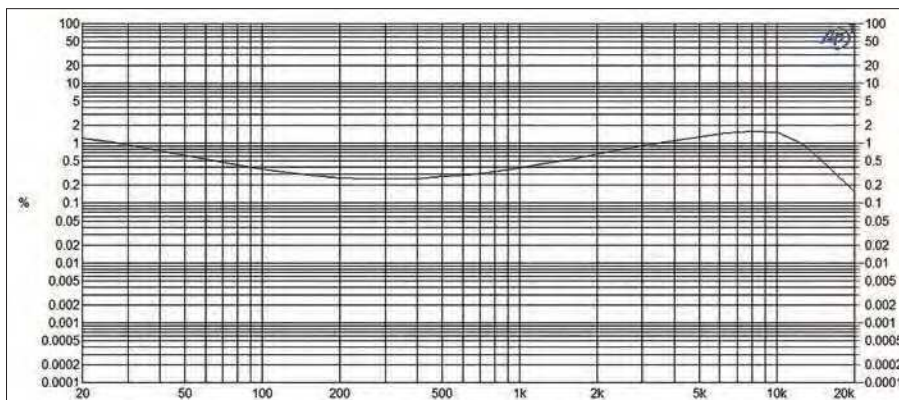
Fig.15. Distortion for the germanium amplifier under the same conditions. It can be seen this is much higher than the silicon amplifier. This is using slow NKT214 transistors for TR1 and TR2.
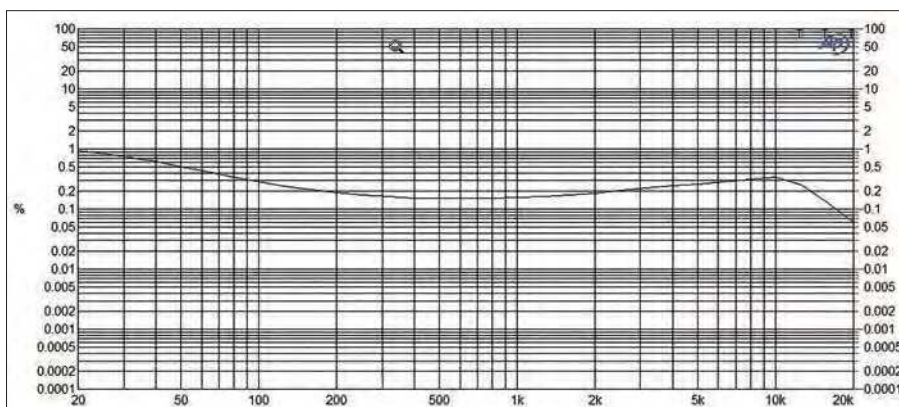
Fig.16. Using fast 2SA12 RF transistors for TR1 and TR2 improved the distortion of the germanium amplifier. This was a trick Leak used in their Stereo 30 amplifier design.

clipping. The loop gain was mainly reduced by the lower value of R2 used on the germanium amplifier to reduce the effect of leakage currents. The result was that the total harmonic distortion was 2.5-times higher in the mid-band (500Hz) and 10-times higher at high frequencies (8kHz) in the germanium amplifier.

The audible effect of these differences could be described as the germanium amplifier sounding 'softer and fuzzier', with the silicon being 'clearer but harsher' when clipped. The germanium amplifier was better for guitar practice and the silicon version for Hi-Fi. The distortion curves are shown in Fig.14 and Fig.15.

In Fig.16 you can see the effect of using fast 2SA12 transistors for TR1 and TR2, dropping the 10kHz distortion from 1.7% to 0.35%. Technically, the germanium amplifier was 'less linear' than the silicon. However, a silicon amplifier with the same level of total harmonic distortion would sound worse because of the higher proportion of high harmonics.

In the end, the physics of bipolar transistor action is the same, an exponential voltage-to-current conversion process, whether you use germanium or silicon. It's the second-order differences which affect the overall results. When I get the spectrum analyser to work we will look

into this deeper. Interestingly, guitarists are prepared to pay £700 for small germanium amplifiers such as the Deacy amps. I suspect it's because the lack of clarity exhibited by imperfect amplifiers hides their mistakes!

Mix it up
The best result was obtained with a BC557B silicon input transistor for TR1, while keeping the rest germanium (TR2, 2SA12). The distortion curve for this set-up is shown in Fig.17. Making TR2 silicon made the mid-band distortion a bit worse and gave less output. Leak did the right thing when they replaced the OC44 input transistor with a BC153 in their Stereo 30 amplifier.

Bass response
The capacitors limit the amplifier's low-frequency −3dB point to around 65Hz. This is fine for the small paper-coned loudspeakers that are used in low-power applications (guitar practice amplifiers) that typically cut-off around 100Hz. For a true vintage guitar/radio sound I suggest using the Celestion and Philips 4072 units shown in Fig.18. The frequency response of the silicon and germanium amplifiers are shown in Fig.19 and Fig.20.

Noise
The noise on both circuits was 1mV$_{pk-pk}$ with the input shorted to ground. There was more high-frequency noise in the silicon design and more low-frequency noise in the germanium one. The odd germanium device was very noisy and had to be swapped out from the input stage. Both units worked well as headphone amplifiers.

Overload protection
The total forward-voltage drop of the short-circuit protection diode network (D1 and D2) needed to be 1V. This was obtained by combining the forward voltage drop of a Zener diode (0.75V) with a germanium diode (0.25V).

Parts list

PCB – this is the same board used in the November/December 2020 *Audio Out*, part AO-1220-01, available from the *PE PCB Service* at: **www.electronpublishing.com**

Resistors
All 0.25W 5% carbon film (silicon amplifier resistors in brackets)

R1	4.7kΩ
R2	15kΩ (47kΩ for silicon TR1, such as BC557B)
R3	100kΩ
R4, R9	47Ω
R5, R6	1.5kΩ
R7	220Ω (1kΩ)
R8	160Ω

R10, R13 10Ω
R11, R12 0.39Ω
R14 not used

VR1 500Ω TO5 outline preset
VR2 5kΩ TO5 outline preset

Capacitors
All 10V or higher radial electrolytic un-
less otherwise stated

C1 10µF
C2 100µF
C3 22µF
C4 470µF
C5 100µF
C6 100nF ceramic or polyester
C7 470µF
C8 1000µF
C9 not used
C10 33pF ceramic (Not normally used,
   see text.)
C11 10µF

Semiconductors
TR1,2 For Ge amp, use small-signal
   high-frequency PNP germanium
   OC44, 2SA12, GET872 or audio
   devices such as OC71, OC75,
   NKT214; for Si amp, use BC549
TR3 For Ge amp, use AC153, AC128
   or any medium-power low-fre-
   quency germanium device. A
   Brimar/Thorn/Mazda AC169 bias
   transistor can also be used; for Si
   amp, use BC337
TR4 For Ge amp, use AC153K, AC188
   PNP germanium power transistor;
   for Si amp, use BC138, BC140,
   BFY51 TO5 medium-power sil-
   icon NPN transistors
TR5 For Ge amp, use AC176K, AC187
   NPN germanium power transistor;
   for Si amp, use BC143, 2N2905 TO5
   medium-power PNP transistors
D1 For Ge amp use low-voltage Zener
   3.3V to 12V, 400mW; eg, BZY-
   88C5V6; for Si amp use standard
   high-brightness red LED
D2 For Ge amp use germanium
   small-signal diode; eg, CG92,
   OA91; for Si amp use BAT86
   small Schottky diode

Germanium parts available from the author:
Tel 01597 829102
Email jrothman1962@gmail.com

Why do it?

For the audio historian and component
obsessive it's definitely worth the work.
It shows 1970s components did the job,
and five decades on they are still viable.

Fig.20. Frequency response of the silicon
amplifier. The germanium amplifiers using
fast transistors for TR1 and TR2 gave the
same response.
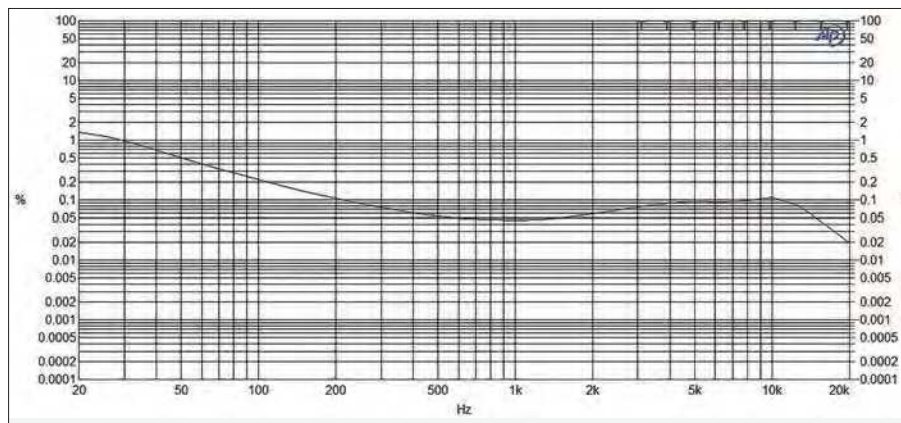
Fig.17. A further improvement to the design was to use a BC557B silicon transistor
for the input stage TR1 while keeping the rest germanium. TR2 was a 2SA12 device
and R2 is 47kΩ. This configuration was the best, producing just 0.045% mid-band
distortion. Again, Leak did this as a later mod to the Stereo 30 amplifier.

Fig.18. Suitable speakers (available from the author) to complement the germanium
transistor amplifier. They are particularly suitable for guitar.
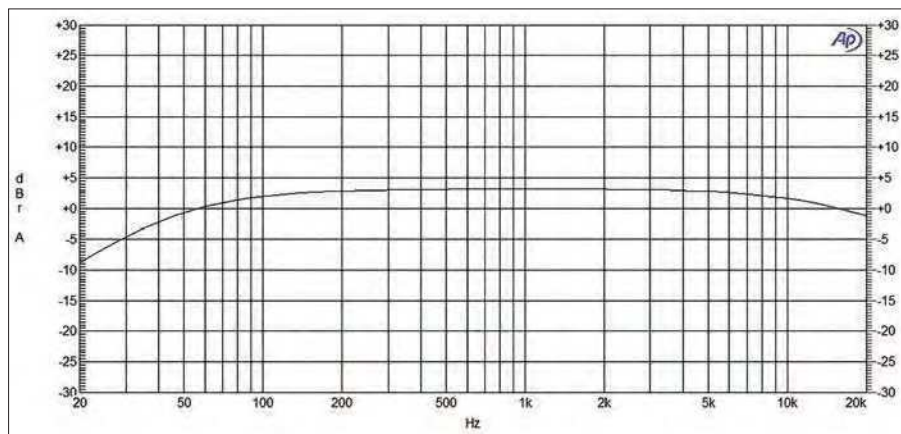
Fig.19. Frequency response of the germanium amplifier with slow input transistors
(NKT214) and C10 connected.

Circuit Surgery

Simulation issues with a digital divide-by-two circuit

In response to the Circuit Surgery article on Micro-Cap 12 (December 2020), Ken Wood decided to try to simulate a circuit he had problems simulating years ago. He writes: 'I was delighted to be alerted to the availability of Micro-Cap 12, and duly downloaded it. I have not toyed with simulators since my professional life circa 1988, where we had a huge Unix machine at work.

'I have had a little problem in the back of my mind for years: recreating standard logic register functions using just gates. I had no success simulating those circuits back then, so I figured it was worth another go. I designed a simple divide-by-two (output toggles at half the rate of the input) using eight gates, and input it into MC12 (see Fig.1).

'Running the simulator results in instability (see Fig.2). There are a couple of cross-coupled gates, and with the other signals in the appropriate state this creates a positive feedback loop with delay – which oscillates in the simulator. Experience says this won't oscillate, it will latch into one state or the other – so I built the circuit and proved it works. The photo shows the breadboarded circuit (built from two CD4011 quad-NAND ICs), powered from 5V and using the 1kHz 5V square wave calibration signal from the 'scope as a clock.

'The 'scope traces are the input at the top, and the output at the bottom – it worked first time, exactly as designed.
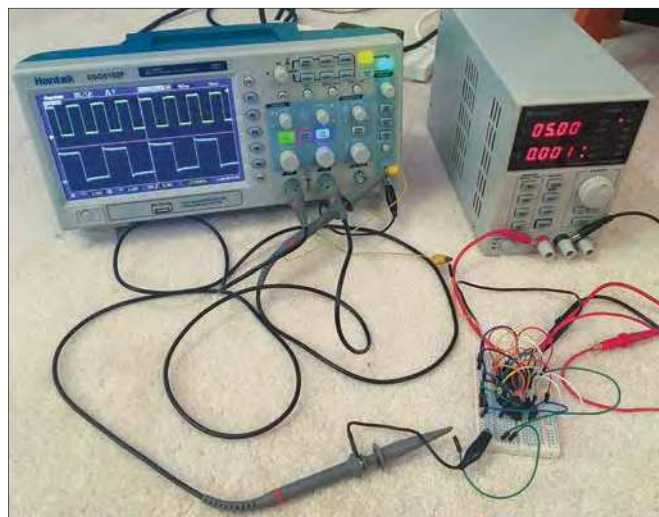
Fig.3. The Fig.1 circuit built on a breadboard, operating correctly.

'Is there any answer to this simulation problem? I think it might be due to the models being perfect and creating a timing race hazard, whereas real devices are never perfect. I tried tweaking some of the gate delays, but I haven't hit on a working combination.

'So now I know why my simulations didn't work in 1988. The simulator was regarded as the be all and end all, and it's a good job I never needed a circuit like this in a professional project, because it would have been difficult getting it past a pre-production review without simulator confirmation it would work!'

'Before looking at Ken's circuit and what might be going wrong in the simulation we will go back to basics and look at how flip-flops are built from logic gates, how they work and the problems that can occur with their operation.
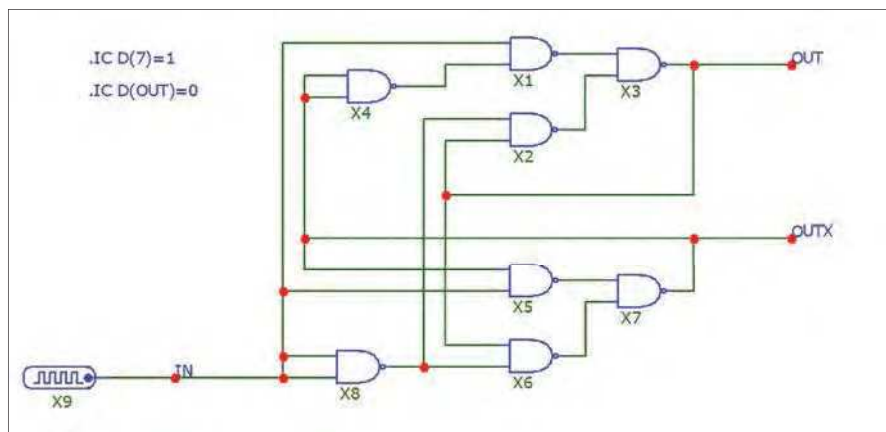
Fig.1. Ken's Micro-Cap 12 schematic for his divide-by-two circuit.
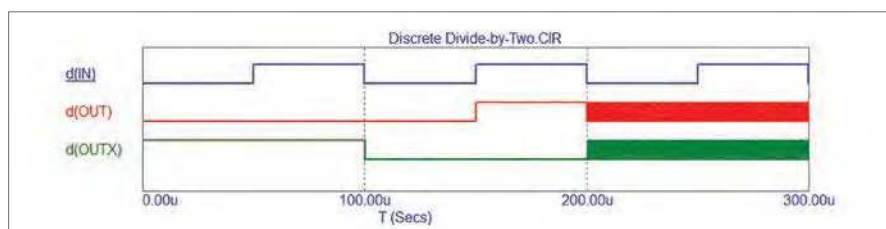
Fig.2. Micro-Cap 12 simulation results from the circuit in Fig.1 – the solid blocks of colour on the outputs are oscillations which are too fast to see on this timescale.
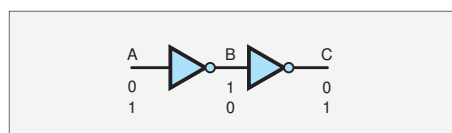
Fig.4. Two inverters in series.
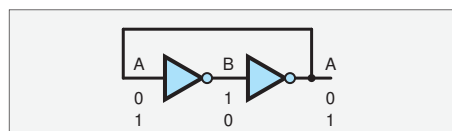
Fig.5. Two inverters in a loop – a bistable circuit, the basis of flip-flops.
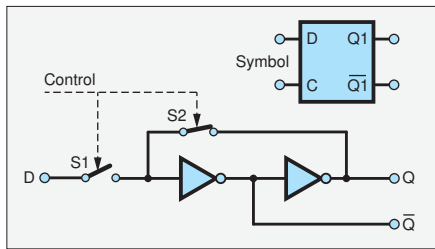
Fig.6. Latch circuit and symbol.
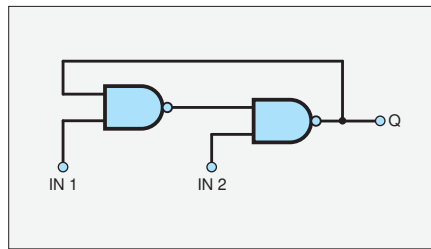
Fig.7. Set-reset flip-flop.
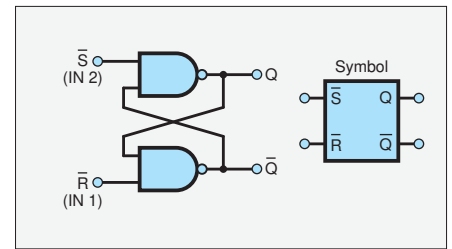
Fig.8. Set-reset flip-flop circuit and symbol.

Memories are made of this

Think of two inverters in series (see Fig.4). The 1 in (at A) gives a 1 out (at C) and a 0 in gives a 0 out. Consider what happens when we connect the output back to the input (see Fig.5). There is no conflict since the 'input' and output are both at the same logic level (there is not really an input anymore). If point A (now the same as C) is made 1, it will stay there indefinitely. If we set A to 0, it will stay at 0. These two conditions are referred to as 'stable' states, and the circuit is described as being 'bistable'.

This ability to indefinitely hold one of two possible states is the basis of the static digital memory, where the term 'static' refers to the fact that the state is held permanently (as long as power is applied) and does not 'fade' or need refreshing. The memory function provided by the circuit in Fig.5 is not particularly useful as there is no input by which it can be given a value to remember; however, we can modify the circuit to achieve this in a couple of ways. First, we could break the loop using a switch while employing another switch to connect the input. Second, gates can be used to modify the logic of the loop in order to enable the state to be set.

The way in which switches may be used to set the state of the loop in Fig.5 is shown in Fig.6. This circuit can be implemented in CMOS, where MOSFETs can be used for the switches. The two switches (S1, S2) are controlled together and are arranged so that when S1 is open S2 is closed and vice versa. When S1 is open and S2 is closed the input is isolated and the loop is closed (like Fig.5). The loop stays in its current state and is unaffected by the input. When S1 is closed and S2 is open the input is connected to the output via the two inverters (like Fig.4, and the output will follow the input. When the switches change to S1 open and S2 closed the loop is formed, storing the data. The gate capacitance of the inverters will hold the value while the changeover occurs until the loop locks the stored state in place.

The set-reset flip-flop

The circuit in Fig.7 replaces the loop inverters of Fig.5 with two NAND gates. If IN1 = IN2 = 1 the circuit is equivalent to the two-inverter loop. However, if one of the inputs is zero the Q output is forced to either 1 or 0, irrespective of the previously stored value. If IN1 = 0 and IN2 = 1 an output of Q = 0 is produced. This is retained when IN1 returns to 1. If IN1 = 1 and IN2 = 0 an output of Q = 1 is produced. This is retained when IN2 returns to 1.

The IN1 input going low forces a 0 to be stored and so this input is called reset (R). The IN2 input going low forces a 1 to be stored and so this is called set (S). This type of circuit is called a 'set-reset' flip flop (SR or RS flip-flop) and its circuit is often drawn in the form shown in Fig.8 and can be represented by a block symbol, also shown in Fig.8. The set and reset inputs are 'active low', hence the bars drawn over them in the diagram. If you connect two NOR gates in the same configuration you get an SR flip flop with active high inputs. There are two SR flip-flops in Ken's circuit one formed by X2 and X3, and the other by X5 and X7.

The outputs of either NAND gate in Fig.8 may be used and are *complementary* (opposite logic levels to each other) and so are labelled Q and $\overline{Q}$. However, if we set both inputs to zero then both outputs will go to 1. If we require that the outputs are complementary then this is an 'illegal' condition, perhaps implying that our flip-flop is not being used properly – we are effectively asking it to set and reset at the same time, which does not really make sense. However, if we remove one of the active input conditions the remaining one will be applied correctly. For example, if we apply $\overline{R}$=0 and $\overline{S}$=0 to the flip-flop in Fig.8 and then change to $\overline{S}$=1 while keeping $\overline{R}$=0 the flip-flop will reset. If the condition where both outputs are 1 does not cause problems elsewhere in the circuit then using the flip-flop in this way may be alright.

Unpredictable

A far more serious problem occurs if we return both inputs to 1 simultaneously. The flip-flop effectively remembers the last instruction it received (set or reset). If we apply both together and remove them simultaneously, which should it remember, given that it can only store one state? The circuit is symmetrical and there is nothing to indicate what *should* happen – the response is not defined. It is not a good idea to have a circuit with undefined behaviour, but the SR flip-flop is useful otherwise, so we can make a rule to say, *never do this* and try to design to ensure it will never occur. However, a real circuit will do *something* if we remove the set and reset simultaneously and it is instructive to understand what this might be – in case it does occur in something we design.

When we simultaneously change both inputs of a NAND SR flip-flop from 0 to 1 we effectively create a situation like the looped inverters in Fig.5, where both have an output of 1. Both inverters will react to their inputs (also both 1) by changing their outputs toward logic 0. If we consider the gate to behave as a Boolean NOT function with a delay, then both inverters will output 0 after their delay time. If they have identical delays, then both will change to 0 simultaneously. We now have the looped inverters in Fig.5 both outputting 0, they will both change to 1 at the same time and we are back to were we started. The process will continue indefinitely – the looped inverters will oscillate as happens in Ken's simulation.

In a real circuit it is unlikely (near enough impossible) that the two inverters (actually NAND gates in the SR flip-flops of Fig.8) will have exactly the same delay. In this situation, the fastest inverter will switch to zero first, at which point the inverter loop will be in one of the stable states shown in Fig.5. The state is stable, so it will stay there – no oscillation will occur. The problem is that in general we will not know which gate is fastest, so the outcome is unpredictable – the flip-flop may set or reset, but we do not know which.

This unpredictability may mean that a physical circuit behaves differently at different times, or different copies of the circuit do different things – both could have very serious consequences. However, it is also possible that (almost)
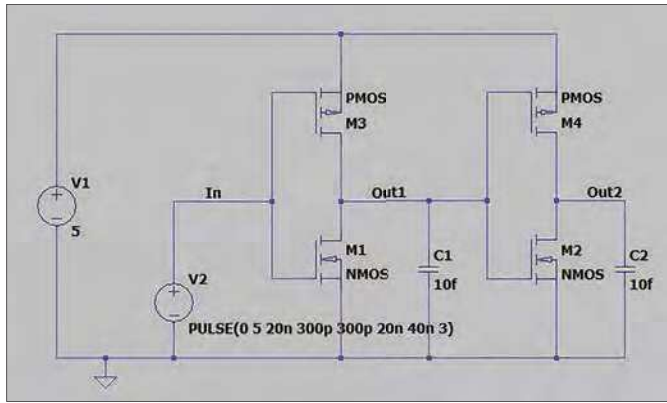
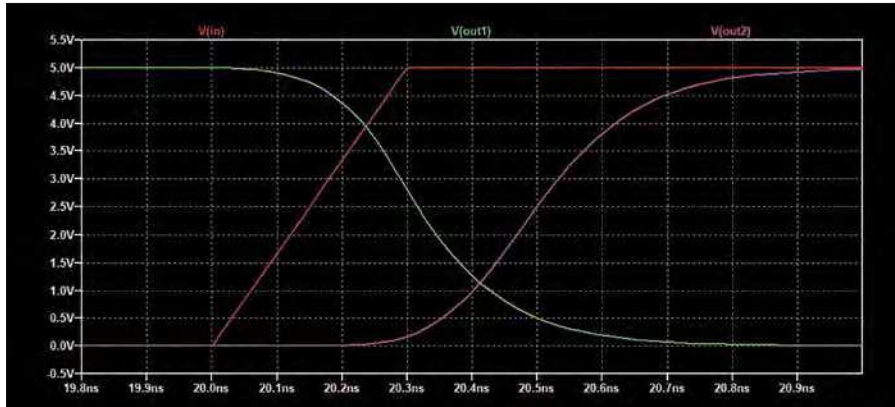Fig.9. LTspice schematic of two CMOS inverters in series.

Fig.10. Simulation results for the circuit in Fig.9 showing an input change from 0 to 1 and the response of the two gates.

at the loop behaviour. The 10fF (10 femtofarad, 10×10^{-15}F) capacitors represent the capacitance of the wiring linking the inverters (as it may be on an integrated circuit). The simulation results in Fig.10 show that the average propagation delay of the two inverters is about 175ps (175×10^{-12}s).

Fig.11 is the circuit from Fig.9 reconfigured to form a loop of inverters, as in Fig.5. We no longer have the input signal from V2 and we need to specify the voltages on the inverter inputs/outputs at the start of the simulation using the SPICE .ic (initial condition) directive. By setting initial conditions where both inputs/outputs are at logic 1 (5V in this case) we simulate what happens directly after we change both inputs of a NAND SR flip-flop from 0 to 1 simultaneously.

Fig.12 and Fig.13 show the results of running two simulations with the two outputs at initially 5V and 1nV less than 5V (4.999999999 V) – trying both ways round. This tiny difference is sufficient to determine the final state of the circuit (which output goes to logic 1 (5V) and which to logic 0 (0 V)). The output which starts at the very slightly lower voltage ends up at 0. There are a couple of key things to note here. First, there is no oscillation, unlike that predicated by the digital gate model with equal delays (and Ken's simulation). Second, the two voltages remain very close for a long time compared to the delay of the gates seen in Fig.10. It is more than ten times the 175ps propagation delay of the inverters before there is any significant voltage difference visible in Fig.12 and Fig.13.

every copy of the circuit we build behaves consistently, for example, because the circuit configuration ensures that one of the looped gates is (almost) always slowest. If this behaviour is in line with the design intention it could lead to a false sense of security – but the circuit may fail as a relatively rare event, or under certain conditions.

If we look a little deeper at what is happening when we change both inputs of the NAND SR flip-flop from 0 to 1 simultaneously we find that the Boolean-function-plus-delay view of the logic gate does not really describe what happens. This may mean that a digital simulation (based on Boolean plus delay) fails to predict the behaviour of the circuit under the conditions we have been discussing. Some people may say that this means you should not trust the simulator – build the circuit instead. But building a prototype in which the design is inherently unpredictable may be equally misleading. If this copy behaves in a particular way, it does not mean that other copies will do the same thing.

Analogue simulation

Digital simulation simplifies the behaviour of logic-gate basic circuits to allow very large and complex circuits to be simulated in a reasonable time. We can of course simulate relatively small digital circuits more accurately using an *analogue* simulator and a full transistor circuit schematic. We will look at an LTspice simulation of the circuits in Fig.4 and Fig.5 to gain some insight into what happens when we change both inputs of the NAND SR flip-flop from 0 to 1 simultaneously. We will just use inverters built from generic MOSFETs rather than attempt to model the full SR flip-flop used by Ken – this will be sufficient to illustrate typical behaviour.

The circuit in Fig.9 is two CMOS inverters in series, which we will use as a reference point when looking

Metastability

The behaviour seen in Fig.12 and Fig.13 is a form of what is known as 'metastability'. As mentioned earlier, the circuit has two stable states, but it also has a metastable state in which the two inverters' input/output voltages are equal. This is a point of perfect balance, like balancing a ball on the point of a needle – in theory it could stay there forever, but the situation is inherently unstable – a miniscule air current will cause the ball to fall off the needle, and likewise any small voltage perturbation will cause the circuit to fall into one of its two stable states.

The response of metastable flip-flops can be analysed by solving differential equations describing the circuit behaviour. The results show that the more perfect the initial balance is, the longer it takes to exit from the metastable state (known as resolution time), and also that brief oscillations may occur for some circuits before the output stabilises.
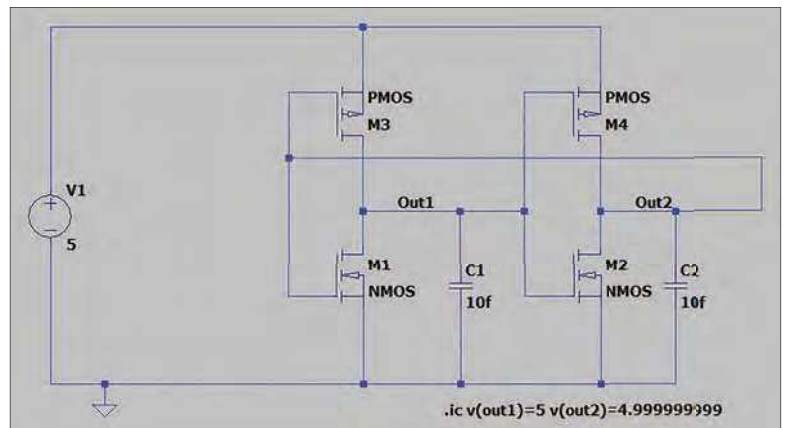
Fig.11. LTspice schematic of two CMOS inverters in a loop. Note the .ic SPICE directive to set the initial voltage on the outputs.
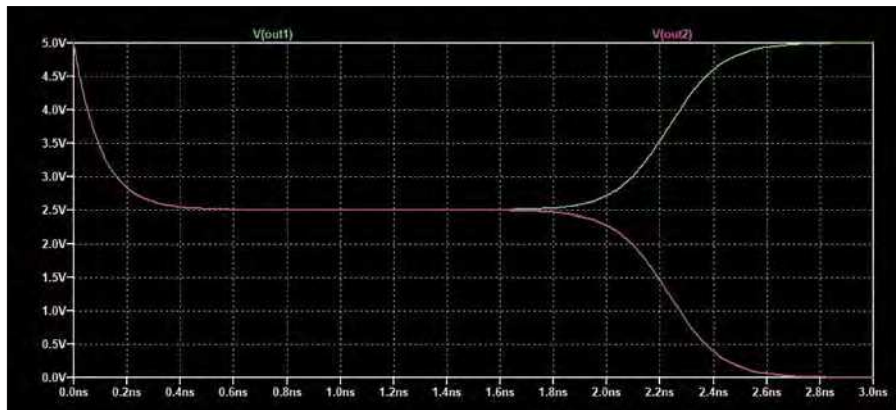
Fig.12. Simulation of the circuit in Fig.11 with initial conditions `v(out1)=5` and `v(out2)=4.999999999`
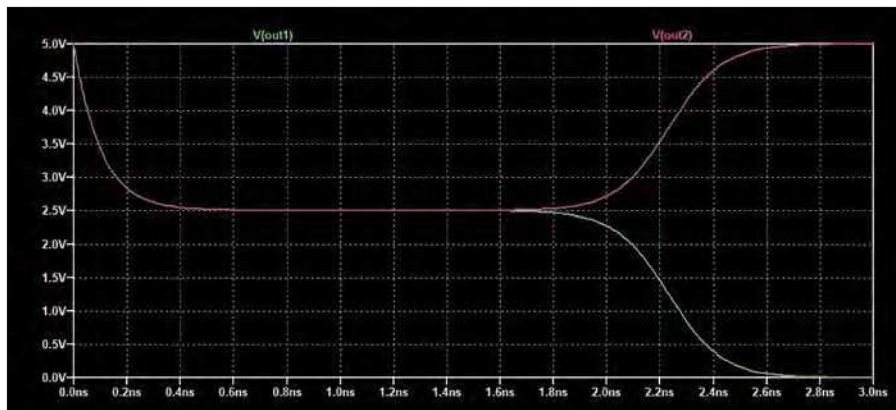
Fig.13. Simulation of the circuit in Fig.11 with initial conditions `v(out1)= 4.999999999` and `v(out2)=5`

Metastability may cause digital systems to fail. Unpredictable outcomes and sustained intermediate voltage levels may result in incorrect logic values propagating through the circuit. Even if the logic level from a metastable flip-flop finally ends up correct, if the resolution time is too long (eg, longer than the clock cycle) errors will occur (eg, the wrong value is sampled by a flip-flop connected to the metastable output).
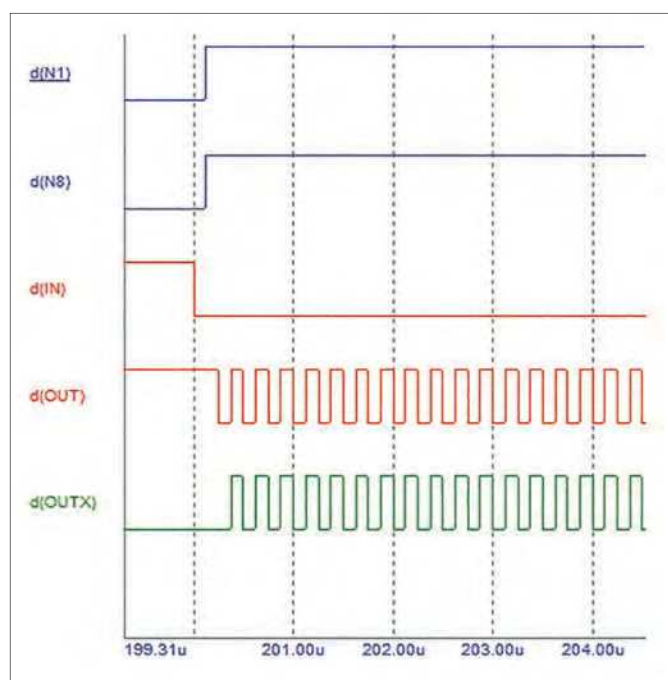
Fig.14. Detailed look at the oscillation in Ken's circuit.

Tracing the problem

The reason that the Micro-Cap 12 simulation of Ken's circuit shows an oscillation is because the inputs to an SR flip-flop can both change from 0 to 1 simultaneously and by default the simulation uses exactly equal delays for all the gates. Consider the SR flip-flop formed by X2 and X3. When the clock is 0 the set and reset inputs are both 1 via the inverter X8 and the NAND gate X1, which are driven by the clock. When the clock is 1 the inputs to the SR flip-flop are $\overline{S}$=OUTX (via X4 and X1) and $\overline{R}$=0 (via X8). If OUTX is 0 when the clock changes from 1 to 0 the flip-flop's inputs both change from 0 to 1 simultaneously. This initiates an oscillation due to the equal delays for X2 and X3, as described above. This situation is shown in Fig.14. This is a zoom-in on the results shown in Fig.2 with additional traces to show the outputs of gates X1 (N1) and X8 (N8). We see the clock (IN) changing from 1 to 0 followed by N1 and N8 both changing from 0 to 1 simultaneously one gate delay later. The oscillation starts in the SR flip-flop formed by X2 and X3 after another gate delay (trace for OUT).

In the physical implementation of Ken's circuit, the metastable condition in the X2/X3 flip-flop may not be triggered because it depends on the delays of X1 and X8 being close to equal. Also, as we have seen, sustained oscillation is not likely to occur in a real circuit if a metastable condition occurs – so we would not expect to see this on the oscilloscope even if there was metastability. If we assume the X1 and X8 delays are sufficiently different to not trigger the problem we can simulate the circuit under these conditions. We can do this in Micro-Cap 12 by setting `PARAMS:=MNTYMXDLY=1` for X8 to select the minimum timing delay model for the gate, rather than the default value. Double click X8 on the schematic to see and set the parameter settings. The results of this are shown in Fig.15 – the circuit operates as Ken intended and matches the results obtained from the real circuit.

Questions

These results lead to questions about the simulation. First, is the simulation shown in Fig.2 wrong? We can say 'no' – the simulator has done exactly as was asked and simulated the circuit with all gates having exactly equal delays and the results are correct at the level of abstraction used. It does not accurately represent the analogue behaviour of the metastable flip-flop, but we would not expect this to be possible given the logic gate models used to ensure fast digital simulation. Second, the specific situation shown in Fig.2 is unlikely to occur in a real circuit, so we might also ask if the result is useful? Here we can say, 'yes'; it alerts us to the fact that the circuit may suffer from a metastability problem. Knowing this we can decide about whether the risk of the problem occurring is acceptable or not, something which is dependent on the context in which the circuit will be used. Simulators are not the best tools for detecting timing problems, but software suites aimed at advanced digital design may include timing analysers, which detect possible timing problems, including metastability.
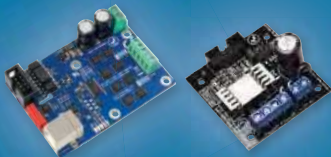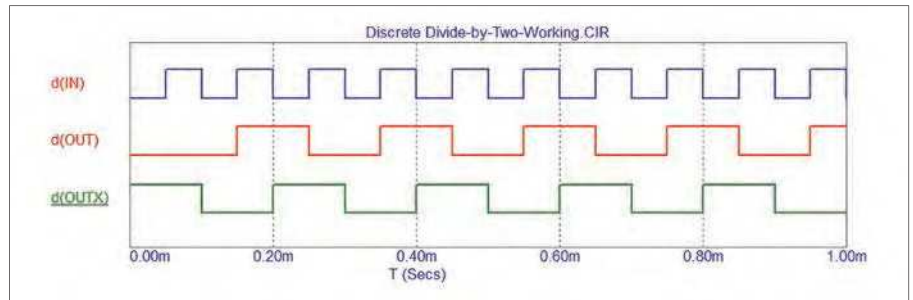
Fig.15. Micro-Cap 12 simulation results from the circuit in Fig.1 with the delay of X8 adjusted so that it does not match that of X1. (Compare with the original simulation of Fig.2 and working circuit of Fig.3.)

Another question that arises is can we build a similar circuit which does not suffer from the problems discussed above? One possibility is to make use of the data latch shown in Fig.16. This also uses a NAND SR flip-flop to store data, but does so under the control of a clock. It is configured so that the S and R inputs of the SR flip-flop cannot be active at the same time. When the clock (Clk) is low, the two NAND gates connected to the clock force the S and R inputs to the inactive state (1). When the clock is high the D input is passed through to the S and R inputs, but the NOT gate ensures that only one of S or R is active. The flip-flop sets or resets depending on the D value.

Using two of these latches in series with the output of the second latch inverted and fed back to the first creates a divide by two circuit, which is equivalent to wiring the $\overline{Q}$ output of a negative edge-triggered D-type flip-flop back to its D input. A
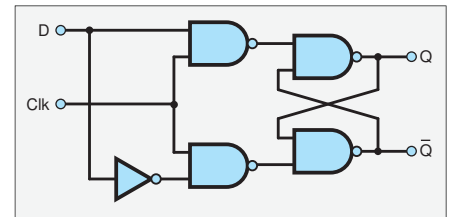
Fig.16. Data latch using a NAND SR flip-flop.

schematic for the suggested circuit is shown in Fig.17. This doesn't include the NOT gates from Fig.16 because we can use the complementary outputs from the SR flip-flops. Simulation results are shown in Fig.18, where the divide-by-two operation can be seen. The circuit is not immune from metastability – it can be caused by changing the latch clock and data at the same time, but to achieve this we would have to clock the circuit fast enough for the delay through the latches to match the clock period.
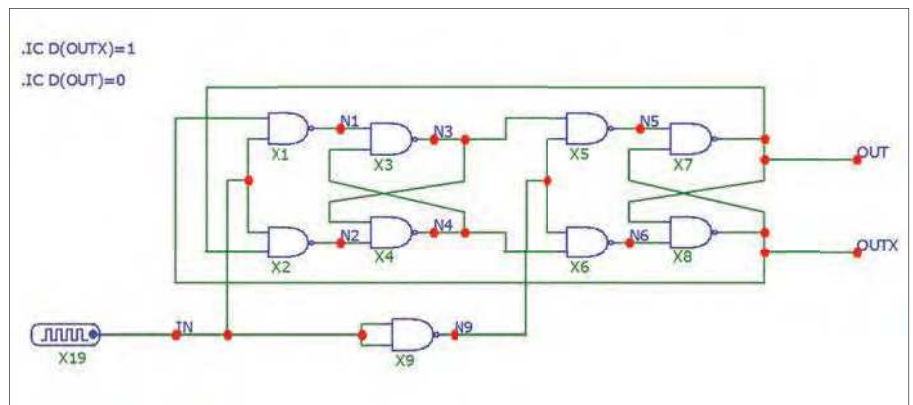
Fig.17. Divide-by-two circuit based on two D latches.
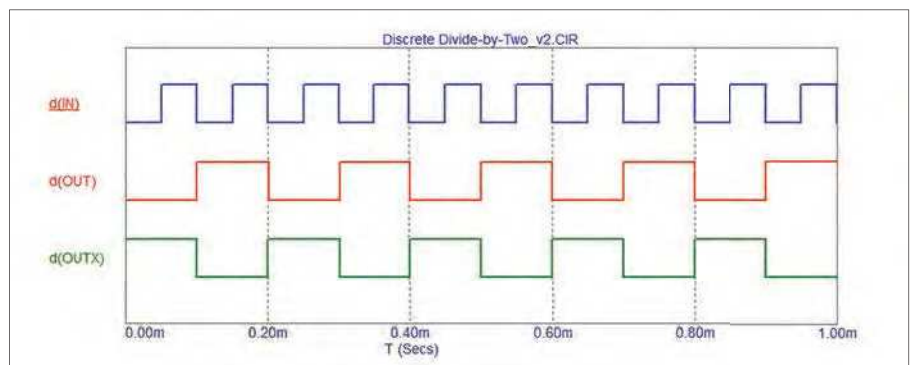
Fig.18. Micro-Cap 12 simulation of the circuit in Fig.17 (all gate delays equal).

Make it with Micromite

Phil Boyce – hands on with the mighty PIC-powered, BASIC microcontroller

Part 25: MMBASIC, Micromite and Maximite varieties

For those of you who have followed this series over the last two years, I hope that somewhere along the journey you have been encouraged to get involved and try some of the topics presented. If so, you should be able to see just how easy it is to implement an idea, and program a microcontroller to do something intelligent and/or useful.

From the feedback we get, many of you have built the *Micromite Keyring Computer* (*MKC*), and lots of you have told us that you have been inspired enough to include a Micromite in your own projects. This was always the main intention of writing these articles, so it is fantastic to hear about the variety of Micromite-based projects that you have created. However, we have noticed that we are being asked on a more frequent basis whether there is something more powerful than the standard *MKC*, but still as easy to use.

This is a good sign, as it implies that you have really grasped how to use the *MKC* to its maximum capabilities. The good news is that 'yes', there are other Micromites available that are more powerful than the *MKC*, and this month we will discuss them. Note that these alternative Micromites may require a different version of MMBASIC to be able to provide the additional features they have to offer. We'll highlight the main differences between the various versions of MMBASIC.

Finally, we will explore some of the many different Micromite and Maximite modules that are currently available.

For those of you that are relatively new to this series, let's start with a quick recap.

So why use a Micromite?
The fundamental purpose of this series has been to demonstrate how software can be used to control hardware. To be more specific, how simple it can be (with the right tools) to use a low-cost microcontroller at the heart of an electronics project.

Unfortunately, many people are put off using microcontrollers due to the complexities of setting up a computer-based development environment, and then having to learn how to use the development tool (that's before you even start to write any code). Next comes learning to use the required programming language; 'compiling' the code in the right format; uploading the relevant file(s) to the microcontroller; and then – finally – seeing the result of the running code. As you can see, this whole process involves many steps, so there is no debate that there is a learning curve to master. Overall, it is often regarded as a 'big chore' to go through all these required steps. The unfortunate result is that a lot of people are probably put off before they begin, and hence they miss out on the huge benefits that a microcontroller can bring to literally *any* electronics project.

We saw at the very beginning of this series that MMBASIC is a solution which avoids having to work through all the above steps. Instead, MMBASIC makes it a simple task (for anyone) to be able to control external hardware with just a few lines of program code. To use MMBASIC, it is loaded as firmware (code) into a Microchip PIC chip, the end result is called a 'Micromite' (or a 'Maximite' – more on this later).

MMBASIC is the creation of Geoff Graham (**geoffg.net**) and what he has created in MMBASIC should be regarded as a real game-changer for those wanting a quick, pain-free route to including a microcontroller in a project.

With the Micromite, there is no development environment to learn, no complex coding language, and no need to compile or upload code. Instead, you just need a computer (or Raspberry Pi) onto which a standard terminal application is loaded (eg, TeraTerm, PuTTY or Terminal). You also need a single USB lead. By using a remarkably simple programming language (MMBASIC), you can immediately start to control external hardware connected to the Micromite.

For anyone that may be thinking BASIC is an 'old language that died years ago', or even, 'why use a Micromite at all?', then consider the following. A benchmark I use with beginners is to show them how to physically connect a red LED to the Micromite, and then how to use software to turn the LED on and off. Next, I show them how to make the LED flash by introducing a couple more commands. After that, they add two more LEDs (yellow and green) by themselves. On paper, I show them the sequence of UK traffic lights and leave it to them to write code to simulate the traffic-light sequence on the three coloured LEDs. The youngest person I have taught this to was five years old, and the oldest 104. Now the impressive part. Bearing in mind they had near zero knowledge of either coding or electronics at the start of the exercise, the above is always completed within 15 minutes – something I have never seen achieved on any other development platform.

In summary, MMBSAIC allows you to be creative and achieve impressive results very quickly, rather than requiring you to be an experienced system-trained user with a good understanding of a complex coding language (which typically means a thorough knowledge of command syntax). For this reason alone, I consider MMBASIC to be a *creative* tool that can be used by anyone, rather than a *system* tool that is used by a system-minded person. That said, do not underestimate MMBASIC's power and capability – it has been used in some very impressive solutions (a possible topic for a future article).

Questions? Please email Phil at:
contactus@micromite.org

Micromite vs Maximite

We are often asked what the difference is between a Micromite and a Maximite, and the simple answer is: a Maximite is like a standalone computer, whereas a Micromite is an embedded controller chip.

To expand on this, a Maximite is a standalone device to which you connect a keyboard and a VGA monitor, and then insert an SD card, to create a standalone computer. When switched on, it can instantly be used to experiment with the many MMBASIC commands, or used to directly control external hardware, or even used to play a good old-fashioned retro game (that is displayed on the VGA monitor). A Maximite can be thought of as a very powerful 1980s home computer.

A Micromite on the other hand is a single chip (IC) that you embed into an electronics project (hence, 'embedded controller'). They are typically used when all you need is an intelligent controller to interact with some connected electronics, all under the direction of the program code.

Maximites and Micromites are both fully capable of controlling hardware under software control, and each has its own version of MMBASIC. This makes it easy to migrate code from one to the other. However, the features and connectivity that each offer do vary considerably, and that is why they have different versions of MMBASIC.

Something to consider, and also something we often advise newcomers to do, is to start with a Maximite and use it to explore what MMBASIC is capable of. It can then be used to write some program code to prototype a new design. If all you need to do is control some external hardware, it is then a simple matter to migrate the Maximite code into a Micromite, and then embed the Micromite in a PCB or some stripboard, resulting in a nice compact solution.

In fact, you can eliminate the Maximite and work directly with a Micromite to test and explore a prototype design; however, this requires the use of a computer or laptop. This is exactly what we have done throughout this series, although my personal preference is to use the Maximite first as they don't hinder with an operating system that suddenly decides that it is time for an update (or locks up due to something else happening within the computer). Also, sometimes a computer's operating system (OS) decides it can't recognise the USB port to which your Micromite is connected. OK, these are minor things in general, but when you do a lot of development work, these 'little things' do become noticeable.

Ultimately, if all you are doing is controlling some electronics, you will eventually need to have your program code on a Micromite, and the Maximite route is just my preferred method for getting to that position.

MMBASIC variants

We have just highlighted that there are two versions of MMBASIC, one for the Maximite, and another for the Micromite; each catering for the significant differences in functionality between the two devices. However, when it comes to the Micromite, there are actually three different versions of MMBASIC, referred to as: Standard, Plus, and eXtreme, which offer more features and functionality, respectively. Thus, eXtreme is the most powerful version for the Micromite, followed by Plus, and then Standard, which is the most basic. Note that they are all backward compatible, so the eXtreme version has everything that the Plus version has along with additional features and functionality. Likewise, the Plus version has everything that the Standard version has, along with more features.

Before we highlight the main differences between each Micromite MMBASIC version, it's worth pointing out that each one is designed to run on a specific Microchip PIC microcontroller. Because the versions of MMBASIC have different features, this translates into the Standard version running on PIC chips that have fewer pins than the Plus version, and in general, the eXtreme version runs on PIC chips with the most pins. The specific PIC chips the various MMBASICs run on will be referenced below.

Device	Standard Micromite (MM)		Micromite Plus (MM+)		Micromite eXtreme (MMX)		
Package	28-pin DIP	44-pin SMD	64-pin SMD	100-pin SMD	64-pin SMD	100-pin SMD	144-pin SMD
Max CPU Speed (MHz)	48	48	120	120	252	252	252
Max BASIC program size (KB)	59	59	100	100	540	540	540
RAM memory size (KB)	52	52	108	108	460	460	460
Clock speed (MHz)	5 to 48	5 to 48	5 to 120	5 to 120	200 to 252	200 to 252	200 to 252
Number of I/O pins	19	33	45	77	46	75	115
Number of analogue inputs	10	13	28	28	24	40	48
Number of serial I/O ports	2	2	3 or 4	3 or 4	3 or 4	3 or 4	3 or 4
Number of SPI channels	1	1	2	2	3	3	3
Number of I²C channels	1	1	1+RTC	1+RTC	1+RTC	2+RTC	2+RTC
Number of 1-Wire I/O pins	19	33	45	77	46	75	115
PWM or servo channels	5	5	5	5	6	6	6
Serial console	✓	✓	✓	✓	✓	✓	✓
USB console			✓	✓	✓	✓	✓
PS2 keyboard and LCD console			✓	✓	✓	✓	✓
SD card interface			✓	✓	✓	✓	✓
Supports ILI9341 LCD displays	✓	✓	✓	✓	✓	✓	✓
Supports 10 LCD panels			✓	✓	✓	✓	✓
Supports VGA displays						✓	✓
Sound output (WAV/tones)			✓	✓	✓	✓	✓
Supports PS2 mouse input					✓	✓	✓
Floating-point precision	Single	Single	Single	Single	Double	Double	Double
Current requirements at 3.3V	30mA	30mA	80mA	80mA	160mA	160mA	160mA

Table 1: Comparison of the key parameters of the Standard Micromite, Micromite Plus and Micromite eXtreme

Finally, there are also versions of MMBASIC designed to run on STM32 processors, and also a version that runs on Raspberry Pi(s). However, because the Pi continually updates its firmware, it is difficult to keep MMBASIC compatible with the latest versions of the Pi OS. For the purposes of this article, we will just focus on *Micromite* MMBASIC for the PIC chips.

What follows are simple summaries of the differences; for more specific details, please see the summary shown in Table 1.

Standard MMBASIC

This is the version of MMBASIC that is used in the *MKC*. It is the base version of the MMBASIC firmware and will run on either a 28-pin PIC32MX170F256B PIC, or a 44-pin PIC32MX170F256D PIC. This version of MMBASIC provides a powerful set of commands and is perfect for small-to-medium projects. It can drive TFT touchscreens up to 3.5-inch in size (via a serial SPI interface) and implements a simple set of commands to write text and draw basic graphics on the display. Interfacing to other hardware is via I²C, SPI, UART and one-wire protocols (see Table 1 for the number of interfaces). The 28-pin Micromite has 19 I/O pins, while the 44-pin Micromite has 33. It runs at a maximum speed of 48MHz. Note that the 28-pin PIC is the only Micromite available as a through-hole DIP package (as used in the *MKC*). It's ideal for users that wish to avoid soldering SMDs (surface-mount devices).

MMBASIC Plus

This version runs on either a 64-pin PIC (PIC32MX470H) or a 100-pin one (PIC32MX470L). In addition to all the features of the Standard version, extra I/O pins are available (45 on the 64-pin PIC, and 77 on the 100-pin PIC). It also operates at a faster speed of up to 120MHz. There is provision for adding an SD card to store program code files (.BAS), data files (.CSV), image files (.BMP) and sound files (.WAV).

It can drive larger touchscreens (up to 9-inch) with an SSD1963 parallel interface. These operate much faster than the serially driven displays of the Standard MMBASIC. Thanks to this faster communication to the display, MMBASIC Plus has additional powerful graphical commands: more font sizes, higher resolution (800×480 pixels), true 24-bit colour, and the ability to load bitmap files stored on the SD card. Improved touch-control functionality includes built-in touch keyboards, animated touch buttons and advanced touch-interrupts.

Improved functionality for sound control makes it possible to generate very precise sinewaves, as well as playing WAV sound files that are stored on the SD card.

A PS2 keyboard can also be added, making it possible to use an MMBASIC Plus setup as a standalone computer when used with a larger display. MMBASIC Plus is perfect for building solutions that require a touch-control user interface, or data logging (to SD card).

MMBASIC eXtreme

This version runs on 64-pin, 100-pin, and 144-pin versions of the PIC32MZ PIC chips. These run at speeds of up to 250MHz (five-times faster than Standard MMBASIC), have up to ten-times more memory, and also have many more I/O pins available (up to 115 on the 144-pin PIC).

The MZ PIC has built-in hardware floating-point capability, resulting in much faster double-precision calculations. A hardware random-number generator is also implemented that generates 'true' random numbers. More I²C, SPI, and UART ports make it easier to add multiple hardware modules that use these interfaces.

Interfacing to displays can use 16 bits (rather than 8, as with MMBASIC Plus) making them faster still to refresh the screen image. Sprite functionality, complete with collision detection is possible, along with screen scrolling. A VGA output is also available as an alternative to a TFT display (on the 100-pin and 144-pin variants), and this provides a resolution of 640×480 pixels in eight colours.

Sound playback is improved with the ability to play FLAC audio files and output them via I²S for playback via an external DAC module (giving superior sound quality).

It is possible to add a USB or PS2 keyboard; plus, a PS2 mouse can be connected that can mimic a mouse pointer (useful when used with a non-touchscreen).

A high-speed frequency counter input is implemented which can operate up to 10MHz. There are also several commands that interface directly to more advanced hardware; for example, a VGA CMOS camera module with 640×480-pixel resolution (the OV7670 module); commands to easily extract data return from GPS modules; and commands that make it easy to calculate angles (pitch, roll, and yaw) from accelerometers and magnetometers.

JSON parsing is implemented, making it relatively easy to extract useful data from
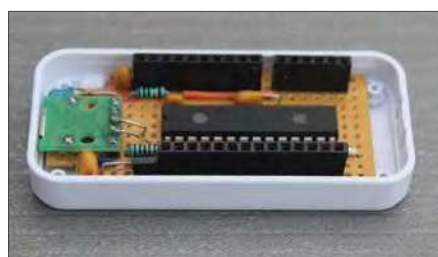

Fig.1. Micromite Keyring Computer (MKC): MM. A DIY kit based on the 28pin DIP Micromite. £12.50 (from **micromite.org**)

Fig.2. Explore 28 Module (E28): MM. An SMD version of the MKC and Development Module, all on one compact PCB. £17.50

Fig.3. BackPack v1 (BPv1): MM. A 2.8-inch touchscreen plugs directly into this module – as used in several PE projects. £19.95
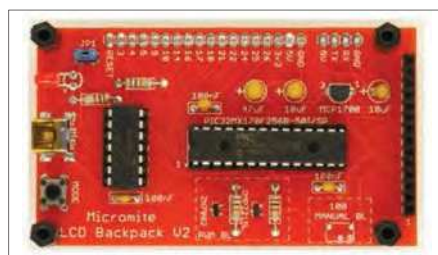

Fig.4. BackPack v2 (BPv2): MM. As v1, but also supplied with an onboard USB-to-serial module (MicroBridge). £24.95

Fig.5. Pi-croMite (MM-PHat): MM. An E28 that fits inside a RPi-Zero case. RPi PHat modules plug into the 40-way GPIO connector. £28.95
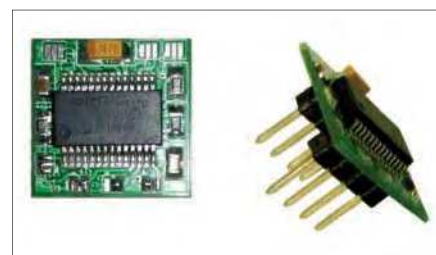

Fig.6. NanoMite: MM. A tiny module with the footprint of an 8-pin DIP. Perfect for compact applications with minimal I/O pin requirements. £14.00

appropriate websites (this requires extra hardware; eg, an ESP Wi-Fi module).

In summary, MMBASIC eXtreme is perfect for building a powerful self-contained standalone computer based on a touch TFT panel, or for when you require a large number of I/O pins, or many hardware interface ports.

Micromite modules

The various versions of MMBASIC make it an extremely versatile platform. This means the Micromite/Maximite can be used very effectively to enable simple (and when necessary, advanced) software control of a wide range of hardware. However, MMBASIC is just the firmware, and the PIC chip is just the microcontroller, onto which MMBASIC is installed – on their own, they are not useable. To make them 'useable', a Micromite only needs a few (less than ten) support components for it to become very powerful. Typically, a voltage regulator plus two smoothing capacitors are required, and also a good quality 10µF tantalum or ceramic capacitor on the PIC's vCap pin. A pull-up resistor is needed on the PIC's reset pin, and on PICs that have more than 44 pins, a crystal (or oscillator module) along with associated oscillator capacitors provide the PIC with a 'heartbeat'. Beyond this, everything else is really just a matter of the required connections between the relevant PIC pins to assorted styles of connectors (USB, PS2, VGA, audio, SD socket, header pins/sockets). You may also need the occasional resistor, capacitor, diode, or LED, and all this is what leads us to the multitude of different Micromite (and Maximite) modules that are available.

It's worth noting that some Micromite modules are referred to as a *BackPack*, and others are called *Explore* modules. In general, a *BackPack* is designed to directly accept a touchscreen, and an *Explore* module (with the exception of the E100) is for plugging directly into a breadboard, or into a custom PCB, without the need to worry about the power circuit and the required support components. The *Explore* modules effectively presents each relevant PIC pin to two rows of header pins – they are ideal for users who don't want to solder SMDs, or simply want a quick and easy way to begin exploring MMBASIC.

Below is a gallery of the 14 varieties of Micromite and Maximite hardware.

I hope I've provided a useful overview of the various versions of MMBASIC, Micromite and Maximite. Together, they offer solutions for virtually any project, from flashing an LED and controlling an attached electronic module, to building a control panel with an advanced graphical user interface, or creating a powerful easy-to-use standalone computer.

Next month

We've seen the different versions of MMBASIC that are available. Next month, we'll show you a fun project that uses some of the extra features available in MMBASIC Plus and MMBASIC eXtreme to help inspire you with your own designs. Until then, have fun!
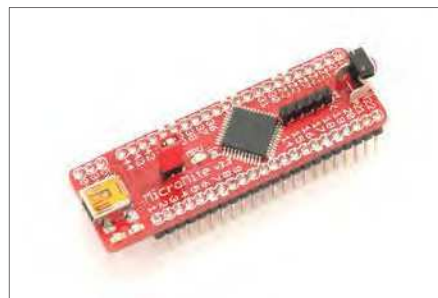
Fig.7. Explore 44 Module: MM. Useful if you need a few more I/O pins than the 28-pin Micromite supplies. £23.95
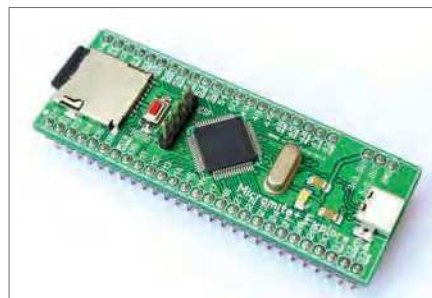
Fig.8. Explore 64 Module (E64): MM+. Includes an onboard uSD socket – perfect for data-logging projects. £29.95

Fig.9. Explore 100 Module (E100): MM+. This module has the footprint of a 5-inch TFT display. Perfect for building a touch-screen control panel. £75.00

Fig.10. Extreme 64 (MMX64): MMX. The entry module for the powerful MMBASIC eXtreme. £75.00

Fig.11. Extreme 100 (MMX100): MMX. Onboard GPS and gyroscope/compass headers make this a useful development platform. £80.00
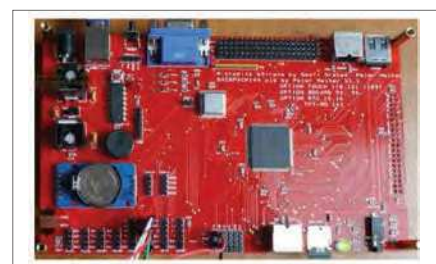
Fig.12. Extreme 144 (MMX144): MMX. Onboard VGA makes it suitable for a standalone computer with VGA output. £85.00

Fig.13. (left) Colour Maximite 1 (CMM): The original PIC-based Colour Maximite computer, as featured in *PE* (November 2019 – February 2020). £68:95

Fig.14. (right) The new Colour Maximite 2 (CMM2): An extremely powerful update to the CMM; based on a blisteringly fast STM32 processor. £85.00.
Soon to be featured as a *PE* project.

Practically Speaking

Hands-on techniques for turning ideas into projects – by Jake Rothman

Restoring old equipment – Part 2

Fig.1. Leak Stereo 30: the UK's best-selling, first-generation transistor amplifier.

Last year, in the November issue, we discussed some of the challenges and rewards of salvaging vintage electronic equipment. This month, we start a short series on actually restoring old electronics. The focus will be on a germanium transistor Hi-Fi amplifier, in this case a Leak Stereo 30 lent to me by loyal reader Andy Martell, as illustrated in Fig.1. (It was so old neither of us dared power it up!) The important point to make from the start is that this is just an example to demonstrate techniques. The ideas presented here could be applied to most older transistor-based electronics from the 1960s.

The Internet is now the main resource for those restoring any old technology. Most amplifier manuals and circuits, including the Leak ones, can be easily downloaded; for example, at **leakaudio.com** and other resources at: **https://bit.ly/pe-ps21-01**

Another good source of information on old Hi-Fi units are the *Hi-Fi Year Books*, for example those shown in Fig.2.
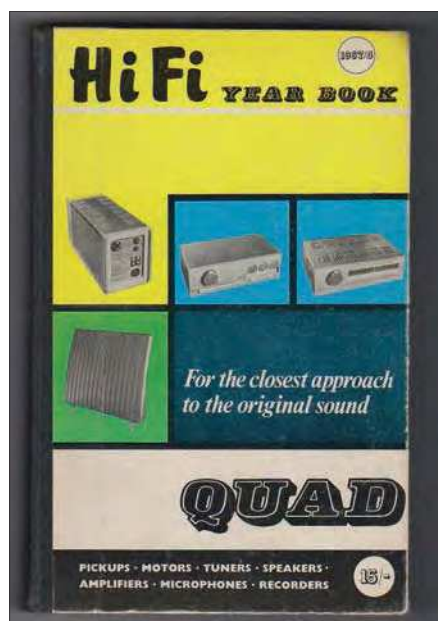
Leak Stereo 30

The venerable Leak Stereo 30 was introduced in 1964, and is recognised as the first successful transistor Hi-Fi amplifier in the UK. A Leak advert from the *1967-68 Year Book* is shown in Fig.3. The price back then was £49. Multiply that by 40 to factor in inflation, and the modern equivalent price is nearly £2000! By comparison, *PE* was 12.5p (two n' six) back then and it's £4.99 today. Of course, this is an unfair comparison, print media has inflated more in real terms and the cost of electronics has deflated. The new Stereo 130 now costs around £700 and *PE* is printed in high-quality colour with superior paper.

The Stereo 30's sound quality is not highly regarded by the Hi-Fi fraternity. The *Hi-Fi News Vintage Review* in October 2010 only gave it a 60% rating. This is because it used germanium transistors, which have poor high-frequency response. Also, the output stage was quasi-complementary, using single-polarity devices, because only PNP types were available in 1963. This circuit topology resulted in a high distortion figure with a spectrum of high odd-order harmonics. Malcolm Jones (of KEF and Falcon Acoustics fame) said the Stereo 30 destroyed many T27 tweeters, because the distortion was about 10% at 10W. Nevertheless, this did not stop 50,000 germanium Stereo 30s being sold as the standard Hi-Fi amplifier, often with a pair of KEF speakers.

Around 1967, the small-signal transistors in the pre-amplifier and power amplifier input stage were replaced with silicon types. In 1969, a fully silicon derivative, the Stereo 30+ appeared, which had double the output power. *Hi-Fi News* gave this 75% in September 2020. This then morphed into the Delta 30 which had a new front panel and then the Delta 70 (Fig.4) which had yet more power. Production continued for at least 17 years. Thus, all variants are plentiful and still cheap, unlike the earlier valve designs.

Double trouble!

While repairing Andy's Leak Stereo 30, I thought it would be a good idea to acquire

Fig.2. Old *Hi-Fi Year Books* are an excellent source of information on early Hi-Fi. (Mike Henslow Year Books Ltd. Later issues IPC).

Fig.3. The advert for the Stereo 30 in the *1967-68 Hi-Fi Year Book* in Fig.2.
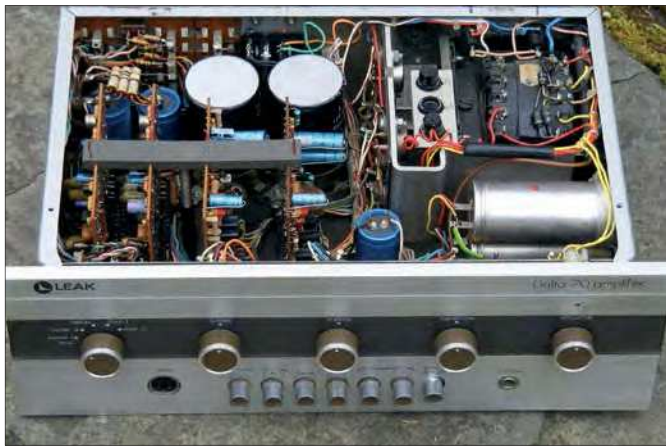
Fig.4. (left) The Delta 70, the final incarnation of the Stereo 30. Excellent, thick front panel, shame about the edge connectors on the plug-in PCBs. Fig.5. (Right) My poor-condition (but cheap) Stereo 30 from eBay.

a really beat-up example and bring it back to life at the same time. Typical 'sold-item' prices on eBay are £20 for poor condition ('spares only') examples to £60 for working ones. I bought a 'poor condition for spares etc' 1967 Leak Stereo 30 for £26.98 including postage especially for this series from eBay seller 'classic-things'. This is shown in Fig.5. It's going to need a massive clean-up and I've decided to do a radical repair/improvement job to see how far I can take it. Note that the only way I could accurately date this example was from date codes on the rectifier diodes. I usually add a year on to allow for delays in manufacturing/distribution.

You might think that the lukewarm reviews, complexity and amount of work involved in restoring a not very

valuable item make the Stereo 30 an unappealing candidate for restoration. For me, though, it was a desire to have a well-preserved example of early transistor history (1965). For Andy it was personal. His Leak amplifier had passed from his grandad to his dad, who had then left it in a loft for years. Also, they are a good project on which to hone one's restoration skills before moving onto pricier electronic classics.

Scratches

Before we get into the Leak, I forgot to show the Novus scratch remover kit discussed in the November 2020 *Practically Speaking* article – see Fig.6. This kit is great for restoring Perspex dials or a retro smoked acrylic turntable lid from a charity shop find. Also, Autosol cream from the car shop is fantastic for burnishing metal, such as the spun aluminium knob-caps shown in Fig.7. They often fall off and are missing. As part of the clean-up, take them off, polish and re-glue with Evostik.

Starting restoration

Let's be clear from the get-go; restoration is time consuming. It's often more work to restore than to build from new because the item must first be very carefully taken apart (see Fig.8). I actually enjoy pulling things to bits – it was how I got into engineering in the first place.

The dust in some 50-year-old equipment must be seen to be believed (Fig.9). A good brushing and a blow with the air duster does the trick (Fig.10). It's best done outside – seriously!

Like all Leak products, the Stereo 30 is well engineered with proper industrial wiring, sound mechanical construction and an excellent layout. I've never had a dry joint in a Leak amplifier. The knobs, although they look like plastic were die-cast metal on Andy's unit; but the later eBay model exhibited 'accountant attrition' and were plastic. Even the PCB is a work of hand-drawn art, much better than normal for the era. A full circuit diagram and service details are given in the manual and the PCB is labelled properly; no numbers *under* components and each one with a left or right channel suffix. Why do we have stupid numbering today? It's a skill that's been lost, like designing toilets that flush properly! These features mean they can be repaired forever. (There was even a company who rebuilt them to a modern complementary silicon design. I don't think that's a good idea though, buy an old NAD amplifier and restore that instead.)

Don't turn it on!

The cast iron rule for newly acquired or loft-liberated vintage electronics is, *don't turn it on!* Relegation to the loft usually happens when equipment breaks, and something better is bought. The huge

Fig.6. Scratch remover kit for plastic dials – a lot of rubbing is involved.

Fig.7. Chrome polish for metal parts; yet more rubbing and polishing.

Fig.8. The eBay Stereo 30 pulled apart. Sometimes equipment is so dirty it's the only way to clean things up to get to the components.

Fig.9. The dust build-up in old gear can be horrendous, and a fire hazard.

Fig.10. The difficult to reach bits can be blown with an air duster.

temperature extremes in roof spaces accelerate components ageing. With something over 56 years-old like this, it is guaranteed not to work. Indeed, if it's just switched on, electrolytic capacitors that have de-formed through lack of use may blow-up. The resulting leakage currents may then take out other components and the speakers. This leads us to basic rule number two – never connect speakers to untested amplifiers.

Voltage selector

The Stereo 30, like most amplifiers of the time, was designed to give optimum results into 15Ω speakers. Back then, this was considered *the* 'Hi-Fi' speaker impedance. The other normal value was 3Ω for radios and other 'lo-fi' products. This meant the main power rail was set to –42V, high for a transistor amp at the time. Note the main power rail is *negative*, and the earth is *positive*, normal with PNP germanium circuits. Although the

Fig.11. Old equipment often has rotary mains voltage selectors since there were wide local variations in mains voltages at the time. Use the 250V setting in the UK to minimise stress on old parts.

specs stated that the amplifier could drive down to 4Ω, doing so could easily result in damage because four-times more current was drawn through the output transistors. Eventually, Hi-Fi designers settled on a value of 8Ω for speaker impedance. 15Ω speakers are now rare, and if the Stereo 30 is to be used with normal 8Ω speakers it's sensible to drop the power rail to –38V. This is easily done by setting the mains voltage selector on the back to 250V, as shown in Fig.11. It would normally be on the 230V setting, for the UK. Of course, the real mains voltage in the UK is 240V, not 230V as commonly supposed. In France, it is 220V. In the interests of EU conformity, it was decreed mains voltage in EU countries would be 'harmonised' at '230V' but with a tolerance of +10% / –6%.

Speakers

If you can find good 1960s 15Ω speakers, such as the Goodmans Maxim or the Leak Mini Sandwich (shown in Fig.3) then they will give good results at low volumes. I sometimes use two EMI 30Ω elliptical speakers in open-back cabinets, which are very 'laid-back' sounding. (See *Audio Out*, *EPE*, June 2016, p.64).

Fig.12. Old gear often has lots of exposed mains wiring. The rear terminal on the fuse holder (far rear left) bit me!

Electric shock

In the 1960s an electric shock every few months was considered par for the course for electrical work. Not anymore – in today's safety-conscious world a product such as the Stereo 30 would be illegal. There was bare mains on the back of the volume control, the fuse holder, multiple tags on the mains transformer and the voltage selector. All this is shown in Fig.12 and Fig.13. I knew it was dangerous and would have to be properly insulated. However, in my enthusiasm to get it going, I forgot to do this vital preparatory work. When feeling the heatsink on the driver transistor during testing, the stud at the back of the mains fuse holder got me – Ouch! – I dropped the whole amp off the bench onto the floor and bent the corner of the front panel. Complacency comes to us all eventually. The good thing is that I can show you how to repair a bent front panel later.

Fig.14 shows how the mains wiring was made safe. The moral? Insulate with rubber sleeves before starting work. Even mains plugs in those days had unsleeved pins, as shown in Fig.15. However, I kept the 1960s plug because I liked its historical solidity. (Only idiots grab the pins as they pull the plug out – please don't be an idiot!). One interesting feature is the pre-1971 mains-cord colour-coding, with red for live. None of that brown for live nonsense! Interestingly, the earth wire was green with a spiral yellow tracer line. A minor concession for the colour-blind.

Circuit history

Before we turn on ancient gear, it's informative to have a look at the history behind its circuit design to aid understanding and fault finding. (There is also a salient tale here that is a warning to all aspiring circuit designers in *PE*). Most 1960s transistor amps used driver and sometimes output transformers. These were expensive, and due to their restricted frequency response and associated phase shift, the amount of negative feedback that could be applied was limited. However, if there is one thing needed to reduce the distortion

(Left) Fig.13. The mains on-off switch was often ganged with the volume control on old audio gear. I don't think it's a good idea, but it does mute the turn on/off thumps from the pre-amp. (Right) Fig.14. The fuse holder is now safe (too late in my case) with its (pink) sleeve, also on the mains switch.
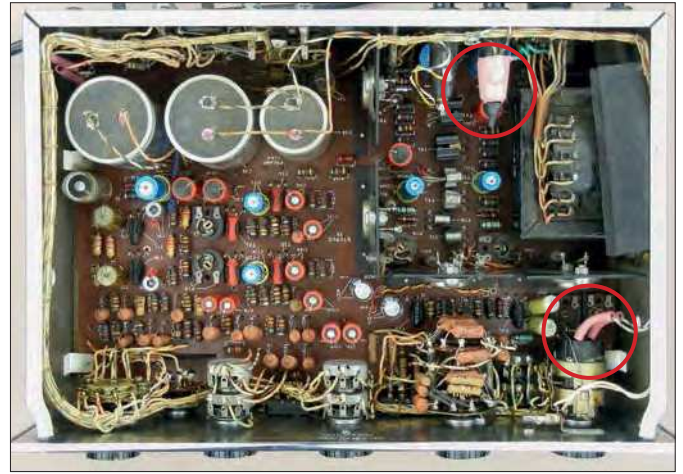
of transistors down to the levels of valves, it's lots of negative feedback. A classic driver-transformer-based Hi-Fi amp is the 1966 Philips TRA 2506 (**www.audio.nl**) which was my parents' first amp. The Rogers amps, such as the Ravensbrooke use the same basic circuit. They are reliable because they are AC-coupled. The first DC-coupled, transformerless amplifier was the RCA 1956 HC Lin design (an interesting man – see: **https://en.wikipedia.org/wiki/Hung-Chang_Lin**). Two Plessey engineers described a 2W version using Mullard transistors in *Wireless World*, March 1960. This led to the Tobey-Dinsdale amplifier published in *Wireless World* in November and December 1961, and which was later updated in *Wireless World*, January 1965. A good quality PDF can also be found at: **https://bit.ly/pe-ps21-02**

This design set the template for the modern transistor amplifier. It was described by Paul Kemble in his website as the 'Grand-daddy of them all' (see: **https://bit.ly/pe-ps21-03**).

Every audio engineer of a certain age must have built one. I was too young, though I have bought a few from radio rallies. They didn't sound as good as the valve amplifiers they replaced, but they

Fig.15. Unsleeved pins on a 1960s MK mains plug; it's built like a tank and good for another 50 years.

did sound reasonably Hi-Fi, unlike most of the other early transformer-coupled amps. However, just like all 'modern' DC-coupled amps they suffered from the 'domino effect', where one transistor going short-circuit killed the ones connected to it. *Hi-Fi News* covered a Dinsdale amplifier in the June to September 1964 issues (updated in March 1966). *PE* also did one in its second issue in December 1964!

Now to the story, Harold Leak invited Dick Toby and Jack Dinsdale over after seeing their article. After giving them the grand factory tour and a gourmet meal, he asked them to leave their sealed amplifier with him to take some measurements. It was later returned having been opened up, with the excuse 'his staff hadn't followed his instructions'. Later, when the Stereo 30 came out to great acclaim, it transpired Leak had copied most of their design. When challenged at its launch at the 1964 Audio Fair, he said to the 26-year-old Dinsdale, 'I've got more money than you, rock my boat and I will destroy you'. Jack Dinsdale tells the whole story in a letter published in the October 1996 *Hi-Fi News* magazine available on the Hi-Fi History Society website (**https://bit.ly/pe-ps21-04**). (I source back issues of *Hi-Fi News* from eBay, via companies such as Zip Audio).

It turned out that the low-pass filter of Dinsdale's circuit had been patented by his employer at the time, Elliott Brothers. Of course, it is easy to get round patents in circuit design. Dinsdale's unpatented idea of using a heatsink-mounted compensating diode to control the quiescent current was circumvented by using a thermistor. The scratch filter part on the tone control was also copied. Leak improved the Dinsdale circuit in conjunction with Mullard application engineers, mainly P Tharma, (see *Wireless World*, July 1963) by using a high-bandwidth video transistor in the voltage amplifying stage. They then claimed that 60dB of negative feedback had been applied. Leak

should be credited with 'productionising' the Dinsdale circuit. Harold Leak then sold his company to the Rank Organisation for £41m in 1969 (using the *PE* inflation index) just before the Japanese moved in. Only recently was Dinsdale's contribution grudgingly acknowledged by Leak's designer Ray Whitcombe as an inspiration in Steven Spicer's book, *Firsts in High Fidelity: The Products and History of H.J Leak and Co Ltd.* (currently out of print). Dinsdale also did okay though, becoming the first professor of mechatronics at Dundee University. He is regarded as the father of the subject, the science of controlling mechanics with electronics. There is a good interview with him in the June 2010 issue of *Hi-Fi News*.

Another Dinsdale rip-off design was the Pye HFS30T. Introduced slightly before the Leak, but not as successful. It used Newmarket transistors, except for the NPN, which was a Texas D451. The circuit is in the *1964-65 Radio and TV Servicing Book* on p.227. Henry's Radio of Edgware road in London used to do lots of Dinsdale power-amp modules, such as the MPA 12/15. He was properly credited and paid for these (see ad on back cover *PE*, Feb 1968). The Leak in turn was then copied by the Truvox TSA100. The **Vintage-radio.net** forum has some good info on these.

Next month

That's enough history! Next month, we'll move on to the restoration.

WARNING!

When restoring electronic equipment, do ensure you work with:
• Eye protection
• Fume extraction
• Safe wiring/earthing

These are not nice-to-have optional extras – you must follow all safety guidelines to protect yourself and those around you.

Max's Cool Beans

By Max the Magnificent

Flashing LEDs and drooling engineers – Part 12

I **must admit to feeling a tad smug** as I pen these words, because I just received an email from a devoted *PE* reader called David Humrich (a.k.a., 'The LED Baron') who hails from Down Under. As part of this antipodean communication, David was kind enough to say, 'I just received the November issue of *PE*. I particularly enjoyed reading your 'drip plus splash' fun. I must confess that I start reading the magazine from the back so that I find your columns sooner!'

'You're making me blush,' I thought. 'Don't stop,' I thought. Suffice it to say that, as a result of David's message, I'm currently bouncing around with a great big smile on my face. Wait until I tell my mother that we've finally found someone who (intentionally) reads my columns!

Prognostication procrastination

If I have one failing (and I'm not admitting to anything, you understand), it's that I'm easily distra… SQUIRREL!!! As a result, many of my hobby projects take years to come to fruition because something else appears on the scene to attract my attention ('Ooh, shiny!'). In the case of my *Pedagogical and Phantasmagorical Inamorata Prognostication Engine*, for example, I've probably been dabbling around with this for close to 15 years at the time of writing.

Some of my friends are prone to making unkind remarks, even going so far as to cast aspersions (which is something I would never do because my throwing arm isn't what it used to be) and telling me that I shouldn't commence any new projects until I finish some of the existing ones. Sometimes my lack of progress makes me feel sad but – for reasons I shall shortly explain – now I don't feel so bad.

On a trip to England to visit my mother in 2018, I went to see an exhibition of Leonardo da Vinci's work, which was being presented in downtown Sheffield. All I can say is that the man was a genius. I'm currently reading his biography by Walter Isaacson (**https://amzn.to/33D0PGR**) who – as always – does a brilliant job. Previously, I read Walter's biographies on Steve Jobs and Einstein. Regarding the former, I think he was a brilliant guy who I wouldn't have liked very much. Regarding the latter, even though we didn't get into the math, it wasn't until I read this book that I fully comprehended just how impressive an achievement Einstein's theory of General Relativity truly was (and still is, for that matter).

The point is that, reading this biography on Leonardo, I discovered we share a lot in common. It goes without saying (although I'll say it anyway) that both us are renowned for our ready wit, dapper dress, and sporty sense of style. Furthermore, I learned that the little rascal left unfinished many more things than he actually completed. He was constantly making sketches with the intention of turning them into pictures, and ceaselessly taking notes with the purpose of writing treatises, but before he got around to actually doing something, he became preoccupied with something else.

Now, I'm certainly not comparing myself with Leonardo da Vinci (I'm happy to leave such comparisons to my dear old mom, who would doubtless opine that Leonardo came in a distant second), but – still and all – it's nice to know that I'm not the only one who… SQUIRREL!!!

The final countdown

I currently have the words and tune to The *Final Countdown* rattling around in my head. This song, which was released by the Swedish rock band Europe in 1986, reached number one in 25 countries. I couldn't help myself; I just took a break to listen to it on YouTube (**https://youtu.be/NNiTxUEnmKI**).

This trip down memory lane was triggered by thoughts of the scene toward the end of the 1987 American science fiction action movie *Predator*, which starred Arnold Schwarzenegger. I'm thinking of the part where Dutch (Arnold) disables the alien's cloaking device and then crushes the short-tempered creature under a trap's counterweight, thereby making it a very unhappy alien indeed. In response, laughing maniacally, the beast activates a self-destruct device, whose countdown sequence is presented as a series of rotating symbols on an alien-arm-mounted display.

I was thinking we might implement something similar on the NeoPixel rings flaunted by the *Prognostication Engine*, as discussed in my previous column (*PE*, January 2021). Since the engine sports a Big Red Button ('Do not press!'), our countdown sequence could be associated with the activation of this button.

Of course, we've got only a single ring to play with on our prototype, but that will provide a starting point. I'm thinking of implementing what I'm calling a 'plodding pixel' and a 'racing pixel' (Fig.1). Remember that the pixel numbers shown here are the way we decided we wanted to visualise and manipulate things, with pixel 0 in the 'south-west' position and the pixel numbers increasing in a clockwise direction. The actual pixel numbers on the physical ring are completely different, so we employ a simple cross-reference operation using an array of integers called `RingXref[]` to translate from our ideal world into the real world (see last month's column and code for more details).
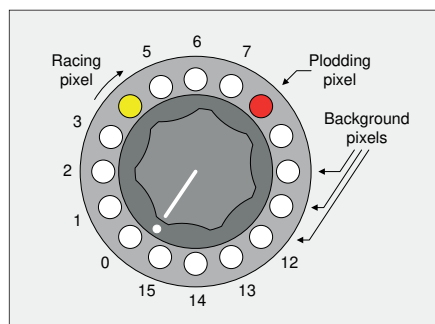
Fig.1. Countdown pixels.
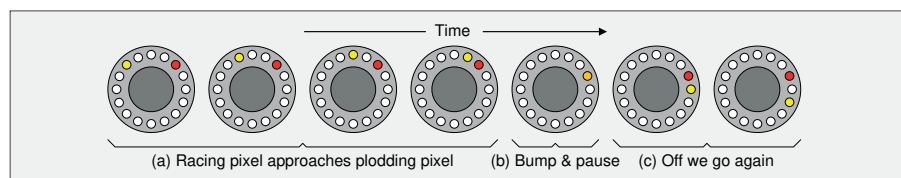
Fig.2. First-pass countdown sequence

The idea is that the racing pixel proceeds in a clockwise direction (Fig.2a) until it reaches the plodding pixel, at which point it 'bumps' the plodding pixel over into the next clockwise spot (Fig.2b). We then pause for a short time before the racing pixel sets off again (Fig.2c).

In an earlier column (*PE*, October 2020), we noted that it's a good idea to make our code and functions as simple and as general-purpose as possible. As part of this, the way I've written the sketch (program) for this particular effect is that we define four colours – one for the background pixels (eg, white), one for the racing pixel (eg, yellow), one for the plodding pixel (eg, red), and one we use when the racing pixel bumps into the plodding pixel (eg, orange). Doing things this way allows us to experiment with different scenarios quickly and easily, as illustrated in the video I just took (**https://bit.ly/3gXpueJ**).

When it comes to the code itself, we've done things a little differently this time (the full sketch is presented in file **CB-Feb21-01.txt** – which is available on the February 2021 page of the *PE* website, and also the *PE* general downloads page: **https://bit.ly/3oouhbl**).

For example, let's consider our bump colour (the one we use when the racing pixel crashes into the plodding pixel). Based on our most recent programs (*PE*, January 2021), we know that we are going to use our GetGammaCorrectedColor() function to correct the gamma value of our initial colour (COLOR_BUMP_PIXEL, in this example). Next, we are going to use our ModifyBrightness() function to control the intensity of this colour as defined by our BRIGHTNESS constant. Previously, we might have broken this sequence out into a number of discrete steps, as follows:

```
tmpColor = GetGammaCorrectedColor(COLOR_BUMP_PIXEL);
tmpColor = ModifyBrightness(tmpColor, BRIGHTNESS);
ColorBumpPixel = tmpColor;
```

We implemented things this way in order to make what we were doing easier to understand, but now it's time to 'take off the training wheels' (well, just one wheel to start with) and make things a tad more concise, as follows:

```
ColorBumpPixel =
ModifyBrightness(GetGammaCorrectedColor(COLOR_
BUMP_PIXEL), BRIGHTNESS);
```

I know this sort of thing can appear a little scary at first, but it's really not as bad as it seems. We can think of this as being like the layers of an onion. First, let's start at the outside and work our way in. As we see, we are assigning the value returned from our ModifyBrightness() function to our ColorBumpPixel variable. Now, our ModifyBrightness() function requires two arguments – a colour to be modified and the amount by which it should be modified (the difference between arguments and parameters was discussed in my *Tips and Tricks* column in the October 2020 issue of *PE*). In the case of the first argument, as opposed to specifying a colour directly, we are using the colour value returned from our GetGammaCorrectedColor() function. Now, that wasn't too bad, was it?

Another way to look at this is to think about how the C/C++ compiler sees things, which is to start by looking at the inner layers of the onion and working its way out. In this case, it will first determine the return value from our GetGammaCorrectedColor() function and subsequently use this as the first argument to our ModifyBrightness() function.

Another thing we are going to do is define two global integer variables called PtrPloddingPixel and PtrRacingPixel, and we are going to use these to 'point' to the relevant pixels

(you'll see what I mean in a minute). In the case of the racing pixel, we want to turn the new (next) pixel on and the old (last) pixel off, then update our PtrRacingPixel variable to point at the next pixel, pause for a short delay, and then do the whole thing again until our racing pixel bumps into our plodding pixel.

In my previous column, we introduced a way of determining the number of the old pixel using the % modulo operator so as to get around the boundary conditions when we transition from pixel 15 back to pixel 0 (or from 0 to 15 if our racing pixel were travelling in an anticlockwise direction). We could use a similar trick this time, but instead we're going to employ a slightly different technique.

You may recall that we introduced the concepts of typedef (type definitions), enum (enumerated types), and struct (structures) in my *Tips and Tricks* column in the December 2020 issue of *PE*. In this case, we'll start by declaring a struct called NextLastPair and then declare an array of these elements called PixelPtrs.

```
typedef struct
{
    int next;
    int last;
} NextLastPair;

NextLastPair PixelPtrs[NUM_NEOS];
```

Remember that the number of pixels in our ring, represented by NUM_NEOS, is 16. The idea is that, for any element in this array, which we will assume corresponds to the current pixel of interest, the next value will be the number of the adjacent pixel in the clockwise direction, while the last value will be the number of the adjacent pixel in the anticlockwise direction.

The clever part is when we initialise this array in our setup() function, as shown below. This initialisation is based on the same modulo operator tricks we introduced in my previous column.

```
for (int iNeos = 0; iNeos < NUM_NEOS; iNeos++)
{
    PixelPtrs[iNeos].next = (iNeos + 1) % NUM_NEOS;
    PixelPtrs[iNeos].last = (iNeos + (NUM_NEOS – 1))
        % NUM_NEOS;
}
```

I know this can be a little demanding on the little gray cells, so I've generated the results for you to peruse and ponder (Fig.3). In the case of the next values, for each value of iNeos, we add 1 and then divide the result by 16 (NUM_NEOS) using the modulo operator to give us the remainder, which – by the magic of maths – is the number we want. Similarly, in the case of the last values, for each value of iNeos, we add 15 (NUM_NEOS – 1) and again divide the result by 16 (NUM_NEOS) using the modulo operator to give us the

iNeos	+1	%16	iNeos	+15	%16
0	1	1	0	15	15
1	2	2	1	16	0
2	3	3	2	17	1
3	4	4	3	18	2
4	5	5	4	19	3
5	6	6	5	20	4
6	7	7	6	21	5
7	8	8	7	22	6
8	9	9	8	23	7
9	10	10	9	24	8
10	11	11	10	25	9
11	12	12	11	26	10
12	13	13	12	27	11
13	14	14	13	28	12
14	15	15	14	29	13
15	16	0	15	30	14
(a) The next values			(a) The last values		

Fig.3. Generating the next and last values.

remainder, which – once again – provides us with the value we require.

The rest is easy-peasy. All we have to do in our main `loop()` function is to cycle around turning the racing and plodding pixels on and off as required, including (possibly) changing the colour of the plodding pixel when the racing pixel bumps into it. Once again, you can see all of this code in file **CB-Feb21-01.txt**, which you can access from the February 2021 page of the *PE* website.

Of course, our experiments here covered a single ring. The *Prognostication Engine* has five such rings, which opens the door to many more possibilities. For example, every time the first (least-significant) ring experiences a 'bump,' this could increment the racing pixel on the next ring, and so on up the chain. My head is buzzing with ideas here, but I'm afraid they will have to wait their turn because – you guessed it – SQUIRREL!!!

But where's the GOL?

I fear I've let you down and led you astray. At the close of my previous column, I made mention of the fact that in this month's column we would be using our 12×12 ping pong-ball array to implement a version of Conway's famous *Game of Life* (*GOL*) (**https://bit.ly/pe-jan21-cgol**). Sad to relate, I got sidetracked by the countdown sequence discussed above, so I'm afraid that we will have to leave the *GOL* until next time (I hang my head in shame). Until then, as always, I (really do) welcome your comments, questions, and suggestions.

Want to build your own amazing array? All the details are in previous *Cool Beans* columns, starting in March 2020.

Cool bean Max Maxfield (Hawaiian shirt, on the right) is emperor of all he surveys at **CliveMaxfield.com** – the go-to site for the latest and greatest in technological geekdom.

Comments or questions? Email Max at: **max@CliveMaxfield.com**

Max's Cool Beans cunning coding tips and tricks

The late great Stephen Hawkins (RIP) wrote a ripping yarn called *God Created the Integers: The Mathematical Breakthroughs That Changed History* (**https://amzn.to/3lKfV3e**). It's sitting on the bookshelf here in my office. One day I intend to read it.

The reason this book just sprung into my mind is that a lot of things in C/C++ programming in general, and programming for the Arduino in particular, are, at their most fundamental level, integers in one form or another.

Are you high?

Let's consider the pin levels `LOW` and `HIGH`, for example. If we've defined a pin as a digital output, we might assign it a value of `LOW` or `HIGH` as follows:

```
digitalWrite(PinA, LOW);
digitalWrite(PinB, HIGH);
```

In the real world, assuming an Arduino Uno, these will appear on the pin as values of 0V and 5V, respectively. But did you know that we can achieve exactly the same effect using 0 and 1? For example:

```
digitalWrite(PinA, 0);
digitalWrite(PinB, 1);
```

How can this be? Well, if the truth were told, behind the scenes ('Pay no attention to that man behind the curtain'), the little scamps who created the Arduino used something like `#define` statements to equate the keywords `LOW` with 0 and `HIGH` with 1.

Knowing this tidbit of trivia opens the door to a lot of things when we come to think about it. For example, consider the following snippet of code:

```
int tmpValue;

tmpValue = digitalRead(PinA);

if (tmpValue == HIGH)
{
   // Do some stuff
}
```

We now know that we could replace `HIGH` with 1, or with the result of an integer expression, or with the integer value returned from a function, or... the list goes on.

We all have truths

As Pontius Pilate famously says (or sings) in the 1970 rock opera *Jesus Christ Superstar*, with music by Andrew Lloyd Webber and lyrics by Tim Rice: 'But what is truth? Is truth unchanging law? We both have truths – are mine the same as yours?' I know what he means and, in a moment, you will too.

Way back in the 1850s, a largely self-taught English mathematician, philosopher, and logician called George Boole (1815-1864) developed a form of mathematics that we now call Boolean algebra. Bool's intention was to use mathematical techniques to represent and test logical and philosophical arguments. Sad to relate, the significance of his work was

not fully appreciated until the late 1930s, when a graduate student at MIT, Claude Shannon, revolutionized electronics by submitting a master's thesis showing how Boolean algebra offered an ideal way for representing the logical operation of digital systems.

In the case of the Arduino, we have a data type called `bool` (or `boolean`). Variables declared using this type can be assigned values of `false` or `true`. Suppose we declare a Boolean variable called `tmpBool` and assign it a value of `true` or `false`. Later, we might perform some tests like:

```
if (tmpBool == true)
{
   // Do some stuff
}

if (tmpBool == false)
{
   // Do some stuff
}
```

Note that the way I'm writing this code – as opposed to using an `if else` construct – is to illustrate a point. Also, as a point of interest, we could make the expressions above more concise and achieve exactly the same effect by writing them as follows:

```
if (tmpBool)
{
   // Do some stuff
}

if (!tmpBool)
{
   // Do some stuff
}
```

In the latter example, we understand `!tmpBool` to mean 'Not tmpBool.'

It's all about the integers

Under the hood, as it were, the C/C++ compiler – and hence the Arduino – tends to treat Boolean operations as though the variables were integers. Let's start with the `false` value, which equates to 0. Based on this, many people think that the `true` value equates to 1. This is true as far as it goes (no pun intended), but in fact any non-zero integer is considered to be `true`, so 1, 2, 6, 10, 100... are all `true` in the Boolean sense.

In turn, assuming we've declared an integer variable called `tmpInt` and a Boolean variable called `tmpBool`, this allows us to do things like:

```
tmpBool = true;
tmpInt  = tmpBool;
tmpInt  = false;

tmpInt  = 6;
tmpBool = tmpInt;
tmpBool = 12;
```

Obviously, this is nonsense code, but it serves to give us an idea what we can do. Now, let's return to consider an `if()` statement as follows:

```
if (condition)
{
   // More statements
}
```

The statements encompassed by the `if()` statement will only execute if the condition returns a `true` (ie, non-zero) value. For example, suppose we have something like the following:

```
if (tmpInt == 6)
{
   // Do some stuff
}
```

If `tmpInt` does contain a value of 6, then the condition will return a value of `true` (1) and the statements associated with the `if()` will execute; otherwise, the condition will return a value of `false` (0), which will prevent the statements associated with the `if()` from executing. The reason I mention this here is that a common mistake is to use a single '=' (assignment) rather than a double '==' (comparison); for example:

```
if (tmpInt = 6)
{
   // Do some stuff
}
```

In this case, the compiler will first evaluate the expression and assign a value of 6 to `tmpInt` (which is going to mess up the rest of the program). Since the processor will see the end result of the 'comparison' as being 6, it will treat this a being true (non-zero), and subsequently execute the statements associated with the `if()`.

We will continue to consider the ramifications of all this in my next *Tips and Tricks* column. In the meantime, as always, please feel free to reach out to me with comments, questions or suggestions.

Electronic Building Blocks

By Julian Edgar
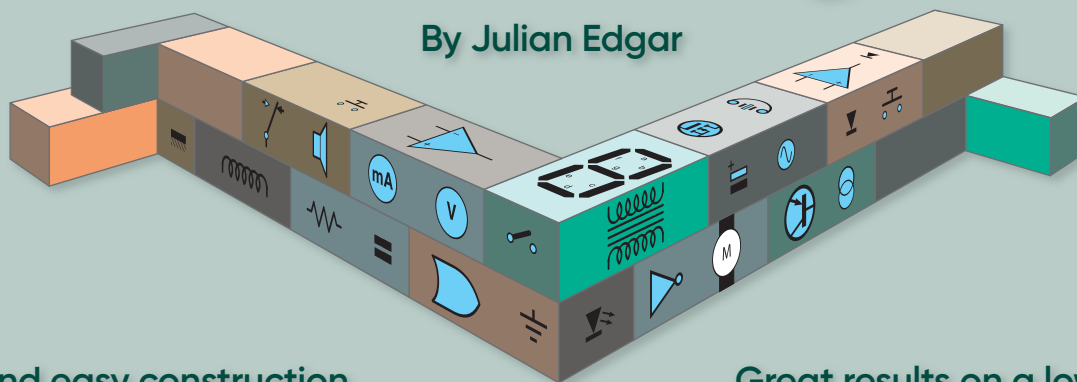
Quick and easy construction

Great results on a low budget

Wireless phone charging pad

Something a little different this month – I recently bought a handy item that was so cheap it didn't matter if I broke it, which I 'did' – more on that below! I wanted to experiment with it and I suspect many of you would enjoy it too.

If you have an up-to-date mobile/cellular phone, you'll be aware that it probably has a wireless charging function. You simply lay the phone on the flat pad and, as if by magic, it charges. (Of course, not much magic is involved – it's being charged inductively.)

Commercial inductive chargers are available, but if you want to have a play with the technology, this little module lets you explore to your heart's content – and at a cost of under £4, delivered. To find it, do an eBay search under 'PCB Circuit Board Wireless Charging DIY Transmitter Module Micro USB Port Stable'. (For example, at the time of writing, eBay item 254688491954 costs £3.68 including delivery.)
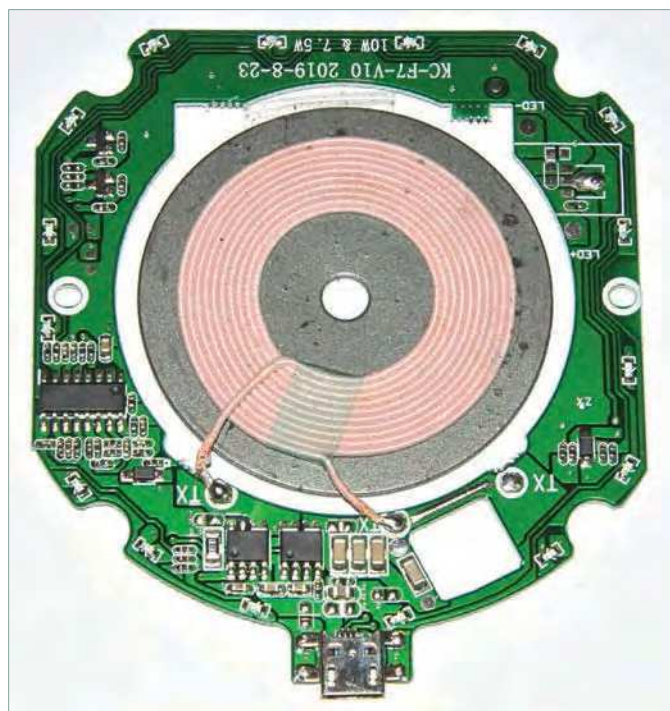
Fig.1. The inductive charging module comprises the charging coil and a surrounding PCB. LEDs around the outside of the PCB make the charger quite colourful in action. Cost is under £4, delivered.

The assembly is 83 × 71mm and about 2.5mm thick. While some of these modules use separate inductive coils, this one places the coil in an opening in the centre of the PCB. The module is powered via a micro-USB port – cable not supplied. (You probably have one in your spares box – lots of equipment comes with them.) Surface-mount LEDs are positioned around the outside of the PCB. These glow green when power is first turned on and then blue when charging. The board has over-temperature protection (60°C) and is said to have intelligent over-voltage and over-current protection.

Input current is specified as 2A and maximum charging current is 1A. However, transfer efficiency is rated as 75% – so make of those numbers what you will! In use, I measured a maximum current draw of 0.8A.

The module uses the QI wireless charging standard. Manufacturers using the standard include Apple, Asus, Google, HTC, Huawei, LG Electronics, Motorola Mobility, Nokia, Samsung, BlackBerry, Xiaomi, and Sony. However, that said, this module didn't work with my first-generation Apple watch, but did work with a current iPhone.

So – and here's the big question – at what distance will the module charge the phone? The specifications state that the best distance is 2-5mm, but that charging can occur at up to 8mm. Of course, place anything between the coil and the phone and that distance decreases.

Mounting

I must admit that mounting the device initially had me a bit flummoxed. For an electronics workbench or similar, I'd be tempted to just glue it to a baseboard made from a piece of wood. However, that risks scratching the back of the phone each time it is placed on the charger, so I took a different approach.

Using 3mm screws, nuts and nylon washes I attached the module to the underside of a plastic lid taken from a cheap kitchen storage box. The plastic lid is semi-translucent, so the operational-status indicator LEDs can be easily seen. However, mounted in this way, the coil sags away from the lid – not wanted. I drilled another hole so that the coil (which has a central hole) could be spaced a short distance away from the lid with the same screws/nuts/washers approach as was done for the PCB. However, I found that using a steel screw stopped the inductive coil from working. I then replaced it with a nylon screw, and then finally, removed the spacers on this screw and drew the coil up flush against the underside of the lid.

Fig.2. The charging module mounted to the underside of a lid taken from a translucent kitchen storage container.
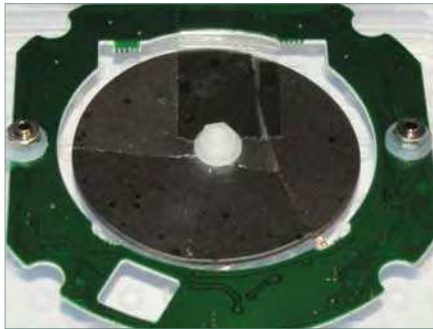
Fig.3. The magnetic induction disc is held in place with a single nylon screw and nut. Be very careful when tightening the nut as the disc can easily crack – as shown here. (The device still worked OK, even with the cracked disc.)
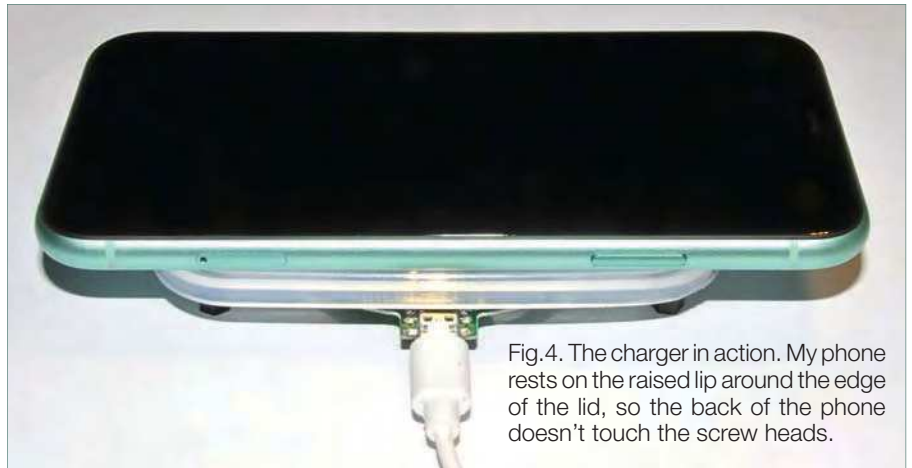
Fig.4. The charger in action. My phone rests on the raised lip around the edge of the lid, so the back of the phone doesn't touch the screw heads.

Then I made a mistake – and tightened this screw. The magnetic disc cracked in multiple places – so be careful not to tighten this screw too much! (The module still worked fine though.)

Four more holes were drilled to allow for standoffs to be screwed into place, one each corner of the plastic lid and so providing four feet on which the charging pad stands.

Note that the plastic lid has a raised ridge around its periphery, and the phone rests on this lip. That is, the phone's back never touches the heads of the screws. This mounting approach gives a distance between the charging coil and the phone of about 4mm, including the thickness of the plastic lid. Mounted in this way, the module charged an iPhone without issues.

Summary

This module is cheap enough to experiment with. But be aware that if you, say, have the cute idea of sticking the module to the bottom of your desk and then just laying your phone in the right place to charge, it won't work – not unless the top of your desk is *very* thin indeed!

Note also that this charger will be slower to charge a phone than a wired connection, and also a little more wasteful of power. That said, it's also very convenient to just lay your phone on the pad and watch it charge!

Practical Electronics PCB SERVICE

PROJECT	CODE	PRICE
FEBRUARY 2021		
4G Remote Monitoring	27111191	£9.95
JANUARY 2021		
Nutube Valve Preamplifier	01112191	£12.95
Arduino DCC Controller	09207181	£10.95
DECEMBER 2020		
Pseudo-Random Sequence Generator	16106191	£7.95
Clever Charger	14107191	£11.95
PE Theremin Amplifier	AO-1220-01	£7.95
NOVEMBER 2020		
LED Christmas Tree (1 off)	16107181-1	£6.95
LED Christmas Tree (4 off)	16107181-2	£14.95
LED Christmas Tree (12 off)	16107181-3	£24.95
LED Christmas Tree (20 off)	16107181-4	£34.95
USB/SPI Interface Board	16107182	£5.95
45V/8A Power Supply PCB *plus* acrylic spacer	18111181	£14.95
45V/8A Power Supply front panel five-way display bezel	18111181-BZ	£3.95
Five-way LCD Panel Meter/Display	18111182	£7.95
OCTOBER 2020		
Digital Audio Millivoltmeter	04108191	£8.95
Precision Signal Amplifier	04107191	£6.95
SEPTEMBER 2020		
PE Theremin PSU	AO-0920-01	£5.95
PE Theremin PSU transformer	AO-0920-02	£7.95
Micromite Explore-28	07108191	£5.95
Ultrabrite LED Driver	16109191	£5.95
AUGUST 2020		
Micromite LCD BackPack V3	07106191	£7.95
Steering Wheel Audio Button to Infrared Adaptor	05105191	£7.95
JULY 2020		
AM/FM/CW Scanning HF/VHF RF Signal Generator	04106191	£11.95
Speech Synthesiser with the Raspberry Pi Zero	01106191	£5.95
PE Mini-organ PCB	AO-0720-01	£14.95
PE Mini-organ selected parts	AO-0720-02	£8.95
High-current Solid-state 12V Battery Isolator – control	05106191	£6.95
High-current Solid-state 12V Battery Isolator FET (2oz)	05106192	£9.95
JUNE 2020		
Arduino breakout board – 3.5-inch LCD Display	24111181	£6.95
Six-input Audio Selector main board	01110191	10.95
Six-input Audio Selector switch panel board	01110192	
MAY 2020		
Ultra-low-distortion Preamplifier Input Selector	01111112	11.25
Ultra-low-distortion Preamplifier pushbutton Input Selector	01111113	
Universal Regulator	18103111	7.95
433MHz Wireless Data Repeater	15004191	8.50
Bridge-mode Adaptor for Amplifier	01105191	7.95
iCEstick VGA Terminal	02103191	4.95
Analogue noise with tilt control	AO-0520-01	7.95
Audio Spectrum Analyser	PM-0520-01	8.95
APRIL 2020		
Flip-dot Display black coil board	19111181	£14.95
Flip-dot Display black pixels	19111182	
Flip-dot Display black frame	19111183	
Flip-dot Display green driver board	19111184	

PROJECT	CODE	PRICE
MARCH 2020		
Diode Curve Plotter	04112181	£10.95
Steam Train Whistle / Diesel Horn Sound Generator	09106181	£8.50
Universal Passive Crossover (one off)	UPC0320	£12.50
FEBRUARY 2020		
Motion-Sensing 12V Power Switch	05102191	£5.95
USB Keyboard / Mouse Adaptor	24311181	£8.50
DSP Active Crossover (ADC)	01106191	
DSP Active Crossover (DAC) ×2	01106192	
DSP Active Crossover (CPU)	01106193	£29.95
DSP Active Crossover (Power/routing)	01106194	
DSP Active Crossover (Front panel)	01106195	
DSP Active Crossover (LCD)	01106196	
JANUARY 2020		
Isolated Serial Link	24107181	£8.50
DECEMBER 2019		
Extremely Sensitive Magnetometer	04101011	£16.75
Four-channel High-current DC Fan and Pump Controller	05108181	£8.75
Useless Box	08111181	£11.50
NOVEMBER 2019		
Tinnitus & Insomnia Killer (Jaycar case – see text)	01110181	£8.75
Tinnitus & Insomnia Killer (Altronics case – see text)	01110182	£8.75
OCTOBER 2019		
Programmable GPS-synced Frequency Reference	04107181	£11.50
Digital Command Control Programmer for Decoders	09107181	£9.95
Opto-isolated Mains Relay (main board)	10107181	£11.50
Opto-isolated Mains Relay (2 × terminal extension board)	10107182	
AUGUST 2019		
Brainwave Monitor	25108181	£12.90
Super Digital Sound Effects Module	01107181	£6.95
Watchdog Alarm	03107181	£8.00
PE Theremin (three boards: pitch, volume, VCA)	PETX0819	£19.50
PE Theremin component pack (see p.56, August 2019)	PETY0819	£15.00
JULY 2019		
Full-wave 10A Universal Motor Speed Controller	10102181	£12.90
Recurring Event Reminder	19107181	£8.00
Temperature Switch Mk2	05105181	£10.45
JUNE 2019		
Arduino-based LC Meter	04106181	£8.00
USB Flexitimer	19106181	£10.45
MAY 2019		
2× 12V Battery Balancer	14106181	£5.60
Deluxe Frequency Switch	05104181	£10.45
USB Port Protector	07105181	£5.60
APRIL 2019		
Heater Controller	10104181	£14.00
MARCH 2019		
10-LED Bargraph Main Board	04101181	£11.25
+Processing Board	04101182	£8.60
FEBRUARY 2019		
1.5kW Induction Motor Speed Controller	10105122	£35.00
NOVEMBER 2018		
Super-7 AM Radio Receiver	06111171	£27.50

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
OCTOBER 2018		
6GHz+ Touchscreen Frequency Counter	04110171	£12.88
Two 230VAC MainsTimers	10108161	£12.88
	10108162	
SEPTEMBER 2018		
3-Way Active Crossover	01108171	£22.60
Ultra-low-voltage Mini LED Flasher	16110161	£5.60
AUGUST 2018		
Universal Temperature Alarm	03105161	£7.05
Power Supply For Battery-Operated Valve Radios	18108171	£27.50
	18108172	
	18108173	
	18108174	
JULY 2018		
Touchscreen Appliance Energy Meter – Part 1	04116061	£17.75
Automotive Sensor Modifier	05111161	£12.88
JUNE 2018		
High Performance 10-Octave Stereo Graphic Equaliser	01105171	£15.30
MAY 2018		
High Performance RF Prescaler	04112162	£10.45
Micromite BackPack V2	07104171	£10.45
Microbridge	24104171	£5.60
APRIL 2018		
Spring Reverberation Unit	01104171	£13.95
DDS Sig Gen Lid	Black	£8.05
DDS Sig Gen Lid	Blue	£7.05
DDS Sig Gen Lid	Clear	£8.05
MARCH 2018		
Stationmaster Main Board	09103171	£17.75
+ Controller Board	09103172	
SC200 Amplifier Module – Power Supply	01109111	£16.45
FEBRUARY 2018		
GPS-Synchronised Analogue Clock Driver	04202171	£12.88
High-Power DC Motor Speed Controller – Part 2		
+ Control Board	11112161	£12.88
+ Power Board	11112162	£15.30
JANUARY 2018		
High-Power DC Motor Speed Controller – Part 1	11112161	£12.88
Build the SC200 Amplifier Module	01108161	£12.88
DECEMBER 2017		
Precision Voltage and Current Reference – Part 2	04110161	£15.35
NOVEMBER 2017		
50A Battery Charger Controller	11111161	£12.88
Micropower LED Flasher (45 × 47mm)	16109161	£8.00
(36 × 13mm)	16109162	£5.60
Phono Input Converter	01111161	£8.00
SEPTEMBER 2017		
Compact 8-Digit Frequency Meter	04105161	£12.88
AUGUST 2017		
Micromite-Based Touch-screen Boat Computer GPS	07102122	£10.45
Fridge/Freezer Alarm	03104161	£8.00
JULY 2017		
Micromite-Based Super Clock	07102122	£10.45
Brownout Protector for Induction Motors	10107161	£12.90

PROJECT	CODE	PRICE
JUNE 2017		
Ultrasonic Garage Parking Assistant	07102122	£10.45
Hotel Safe Alarm	03106161	£8.00
100dB Stereo LED Audio Level/VU Meter	01104161	£17.75
MAY 2017		
The Micromite LCD BackPack	07102122	£11.25
Precision 230V/115V 50/60Hz Turntable Driver	04104161	£19.35
APRIL 2017		
Microwave Leakage Detector	04103161	£8.00
Arduino Multifunctional 24-bit Measuring Shield	04116011	£17.75
+ RF Head Board	04116012	
Battery Pack Cell Balancer	11111151	£9.00
MARCH 2017		
Speech Timer for Contests & Debates	19111151	£16.42
FEBRUARY 2017		
Solar MPPT Charger/Lighting Controller	16101161	£17.75
Turntable LED Strobe	04101161	£7.60
JANUARY 2017		
High-performance Stereo Valve Preamplifier	01101161	£17.75
High Visibility 6-Digit LED Clock	19110151	£16.42

For the many pre-2017 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 . . . .

I enclose payment of £ . . . . (cheque/PO in **£ sterling only**)

payable to: **Practical Electronics**

Card No . . . .

Valid From . . . . Expiry Date . . . .

Card Security No . . . .

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!

Teach-In 2017
Introducing the BBC micro:bit

PYTHON CODING ON THE BBC MICRO:BIT
Jim Gatenby

Python is the leading programming language, easy to learn and widely used by professional programmers. This book uses MicroPython, a version of Python adapted for the BBC Micro:bit.

Among the many topics covered are: main features of the BBC micro:bit including a simulation in a web browser screen; various levels of programming languages; Mu Editor for writing, saving and retrieving programs, with sample programs and practice exercises; REPL, an interactive program for quickly testing lines of code; scrolling messages, creating and animating images on the micro:bit's LEDs; playing and creating music, sounds and synthesized speech; using the on-board accelerometer to detect movement of the micro:bit on three axes; glossary of computing terms.

This book is written using plain English, avoids technical jargon wherever possible and covers many of the coding instructions and methods which are common to most programming languages. It should be helpful to beginners of any age, whether planning a career in computing or writing code as an enjoyable hobby.

118 Pages   Order code PYTH MBIT   £7.99

GETTING STARTED WITH THE BBC MICRO:BIT
Mike Tooley

Not just an educational resource for teaching youngsters coding, the BBC micro:bit is a tiny low cost, low-profile ARM-based single-board computer. The board measures 43mm × 52mm but despite its diminutive footprint it has all the features of a fully fledged microcontroller together with a simple LED matrix display, two buttons, an accelerometer and a magnetometer.

Mike Tooley's book will show you how the micro:bit can be used in a wide range of applications from simple domestic gadgets to more complex control systems such as those used for lighting, central heating and security applications. Using Microsoft Code Blocks, the book provides a progressive introduction to coding as well as interfacing with sensors and transducers.

Each chapter concludes with a simple practical project that puts into practice what the reader has learned. The featured projects include an electronic direction finder, frost alarm, reaction tester, battery checker, thermostatic controller and a passive infrared (PIR) security alarm.

No previous coding experience is assumed, making this book ideal for complete beginners as well as those with some previous knowledge. Self-test questions are provided at the end of each chapter, together with answers at the end of the book. So whatever your starting point, this book will take you further along the road to developing and coding your own real-world applications.

108 Pages   Order code BBC MBIT   £7.99

THEORY AND REFERENCE

PRACTICAL ELECTRONICS HANDBOOK – 6th Ed
Ian Sinclair
440 pages   Order code NE21   £33.99

STARTING ELECTRONICS – 4th Ed
Keith Brindley
296 pages   Order code ELSEV100   £18.99

ELECTRONIC CIRCUITS – FUNDAMENTALS & APPLICATIONS – Updated version
Mike Tooley
400 pages   Order code TF43   £32.99

FUNDAMENTAL ELECTRICAL AND ELECTRONIC PRINCIPLES – 3rd Ed
C.R. Robertson
368 pages   Order code TF47   £21.99

A BEGINNER'S GUIDE TO TTL DIGITAL ICs
Robert Penfold
142 pages   OUT OF PRINT BP332   £5.45

UNDERSTANDING ELECTRONIC CONTROL SYSTEMS
Owen Bishop
228 pages   Order code NE35   £36.99

MICROPROCESSORS

INTERFACING PIC MICROCONTROLLERS – 2nd Ed
Martin Bates
298 pages   Order code NE48   £34.99

PROGRAMMING 16-BIT PIC MICROCONTROLLERS IN C – LEARNING TO FLY THE PIC24
Lucio Di Jasio (Application Segments Manager, Microchip, USA)
496 pages + CD-ROM   Order code NE45   £38.00

INTRODUCTION TO MICROPROCESSORS AND MICROCONTROLLERS – 2nd Ed
John Crisp
222 pages   Order code NE31   £29.99

THE PIC MICROCONTROLLER YOUR PERSONAL INTRODUCTORY COURSE – 3rd Ed
John Morton
270 pages   Order code NE36   £25.00

PIC IN PRACTICE – 2nd Ed
David W. Smith
308 pages   Order code NE39   £24.99

MICROCONTROLLER COOKBOOK
Mike James
240 pages   Order code NE26   £36.99

BOOK ORDERING DETAILS

All prices include UK postage.

For postage, add £3 per book to Europe, £4 for rest of the world per book.

CD-ROM prices include VAT and/or postage to anywhere in the world.

Send a cheque, (£ sterling only) made payable to: **Practical Electronics** or credit card details (Visa or Mastercard) to:

Electron Publishing Limited,
113 Lynwood Drive, Wimborne, Dorset BH21 1UU

Books are normally sent within seven days of receipt of order.

Please check price (see latest issue of Practical Electronics or website) before ordering from old lists.

For a full description of these books please see the shop on our website.
Tel: 01202 880299 – Email: shop@electronpublishing.com

Order from our online shop at: www.electronpublishing.com

TEACH-IN BOOKS

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 (+VAT) per centimetre high, with a minimum height of 2·5cm. All semi-display adverts have a width of 5.5cm. The prepaid rate for classified adverts is 40p (+VAT) per word (minimum 12 words).

Cheques are made payable to 'Practical Electronics'. **VAT must be added**. Advertisements with remittance should be sent to: Practical Electronics, 113 Lynwood Drive, Wimborne, Dorset, BH21 1UU.

Tel 07973518682 Email: pe@electronpublishing.com

For rates and further information on display and classified advertising please contact our Advertisement Manager, Matt Pulzer – see below.

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 the last twenty-seven years.

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

MISCELLANEOUS

VALVES AND ALLIED COMPONENTS IN STOCK. Phone for free list. Valves, books and magazines wanted. Geoff Davies (Radio), tel. 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

COAST ELECTRONICS

BREAKOUTS-COMPONENTS-
CONTRACT DESIGN-3D PRINTER PARTS-
MUSICAL-MICROCONTROLLERS
WWW.COASTELECTRONICS.CO.UK

Andrew Kenny – **Qualified Patent Agent**
EPO   UKIPO   USPTO
Circuits Electric Machinery Mechatronics

Web: www.akennypatentm.com
Email: Enquiries@akennypatentm.com
Tel: 0789 606 9725

ADVERTISING INDEX

Advertisement offices

Matt Pulzer

Electron Publishing Ltd
1 Buckingham Road
Brighton
East Sussex BN1 3RA

Tel   07973 518682

Email pe@electronpublishing.com

Web   www.electronpublishing.com

For editorial contact details see page 7.

Next Month – in the March issue

Programmable Thermal Regulator
Precision temperature control is an integral part of many industrial processes, but it is not easy to implement for the home designer. Next month, we introduce a fascinating Arduino/Peltier-device-based control system.

Nutube Guitar Overdrive and Distortion Pedal
Do you long for that true 'valve sound' in a guitar and distortion pedal? How about this one – it uses a unique low-voltage twin triode valve, so you know it's the real deal! The 'Nutube' dual triode needs no transformers and its cool plate glow is visible! It offers two distortion stages and the high input impedance suits most pickups.

Tunable HF Preamplifier with Gain Control
This simple tunable preamplifier greatly improves SDR HF performance. It has (optional) gain control and can run off a 5V supply or phantom power.

KickStart: Part 2
This second instalment shows you how to use low-cost RFID devices and work with cards and tags. In keeping with the KickStart philosophy, we've provided sufficient information for you to be able to incorporate RFID in your own projects.

PLUS!
All your favourite regular columns from *Audio Out, Cool Beans and Circuit Surgery*, to *Make it with Micromite, PIC n' Mix* and *Net Work*.

On sale 4 February 2021

Content may be subject to change

Did you know our online shop now sells the current issue of *PE* for £4.99 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