MAKE | **BUILD** | HACK | **CREATE**

HackSpace

TECHNOLOGY IN YOUR HANDS

hsmag.cc | March 2024 | Issue #76

ULTIMATE WORKSHOP

GET THE MOST OUT OF YOUR SPACE, WHATEVER THE SIZE

PRUSA XL
The multi-extruder printer on test

LEATHER WORK
Get started with this hardy material

CYBERDECK
Building a portable RF workstation

Mar. 2024
Issue #76 £6

PCBS **DINOSAURS** MUSIC **POWER**

Welcome to HackSpace magazine

Workshop is a flexible term. A workshop can be a chunk of a former warehouse packed with photogenic industrial machinery – you know, the sort of thing that half of YouTube seems to have – or it can be a corner of a room in your house with space for a few drawers, a soldering iron, and maybe a 3D printer. It can even be a temporary space – for years, I had to get the kids off to nursery and then set up the workshop on the kitchen table.

This month, we're looking at how to set up your workshop to be as good as it can be. That doesn't mean drooling over some fantasy list of features, but real steps to optimise your space. Whether it's in your bedroom, shed, garage, or mountain lair, we'll help you stress less and make more.

BEN EVERARD
Editor ben.everard@raspberrypi.com

Got a comment, question, or thought about HackSpace magazine?

get in touch at **hsmag.cc/hello**

GET IN TOUCH

hackspace@raspberrypi.com

hackspacemag

ONLINE

hsmag.cc

Available on the App Store

GET IT ON Google Play

PAGE 26
FREE PICO W WHEN YOU SUBSCRIBE

EDITORIAL
Editor
Ben Everard
ben.everard@raspberrypi.com
Features Editor
Andrew Gregory
andrew.gregory@raspberrypi.com
Sub Editors
David Higgs, Nicola King

ADVERTISING
Charlotte Milligan
charlotte.milligan@raspberrypi.com
+44 (0)7725 368887

DESIGN
Head of Design
Jack Willis
Designers
Sara Parodi, Natalie Turner
Illustrator
Sam Alder
Photographer
Brian O Halloran

CONTRIBUTORS
Marc de Vinck, Andrew Lewis, Jo Hinchliffe, Rob Miles, Nicola King, Tom Mladenov, Phil King

PUBLISHING
Publishing Director
Brian Jepson
brian.jepson@raspberrypi.com
Director of Communications
Liz Upton
CEO
Eben Upton

DISTRIBUTION
Seymour Distribution Ltd
2 East Poultry Ave,
London EC1A 9PT
+44 (0)207 429 4000

SUBSCRIPTIONS
Unit 6 The Enterprise Centre
Kelvin Lane, Manor Royal,
Crawley, West Sussex, RH10 9PE
+44 (0)1293 312193
hsmag.cc/subscribe
hackspace@subscriptionhelpline.co.uk

This magazine is printed on paper sourced from sustainable forests and the printer operates an environmental management system which has been assessed as conforming to ISO 14001.

HackSpace magazine is published by Raspberry Pi Ltd, 194 Cambridge Science Park, Milton Road, Cambridge, England, CB4 0AB.

ISSN: 2515-5148.

Contents

06

Cover Feature

18

28

Tutorial

Music

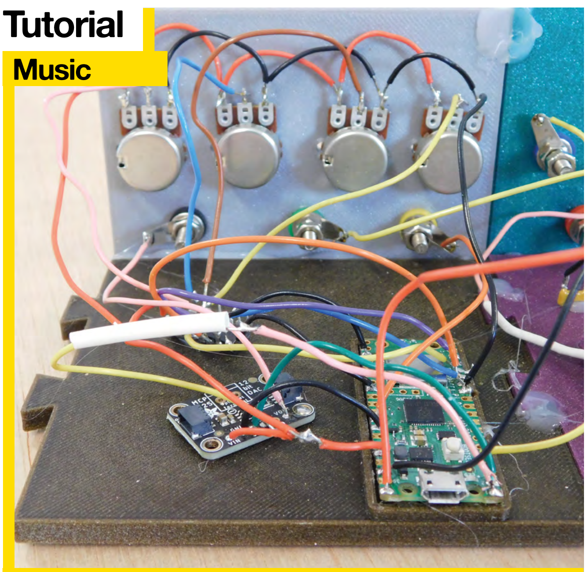

62 What on earth could this jumble of wires be doing?

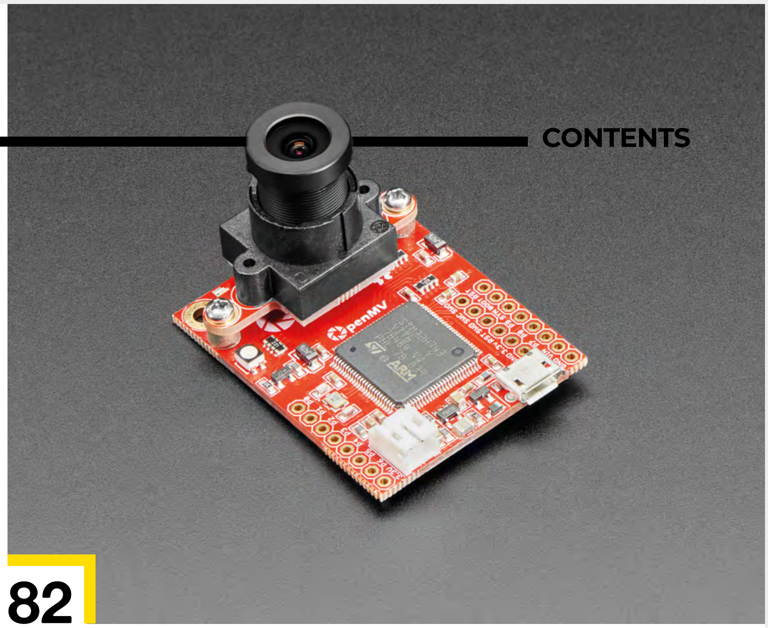

Interview

Tom Ranson

Review

Prusa XL

47 FORGE

81 FIELD TEST

 HackSpace magazine is published monthly by Raspberry Pi Ltd, 194 Cambridge Science Park, Milton Road, Cambridge, England, CB4 0AB, United Kingdom. Publishers Service Associates, 2406 Reach Road, Williamsport, PA, 17701, is the mailing agent for copies distributed in the US and Canada. Application to mail at Periodicals prices is pending at Williamsport, PA. Postmaster please send address changes to HackSpace magazine c/o Publishers Service Associates, 2406 Reach Road, Williamsport, PA, 17701.

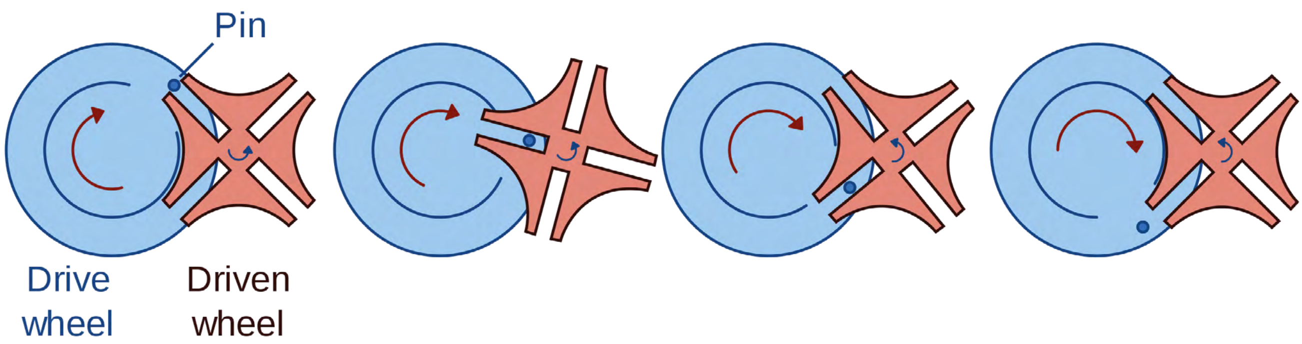

Geneva drive clock

Above ◈
On the clock, the three driven wheels all rotate around the same drive wheel which remains stationary in the middle (image by Booyabazooka CC-BY-SA)

By **Greg Zumwalt**

hsmag.cc/geneva

Clocks, they're pretty simple, right? Spin hands around at three different speeds, pop some numbers on the outside, and you're done, right? Well, that would do for most people, but not Greg Zumwalt. When his youngest son showed him a video of a watch with an unusual mechanism and asked if he knew how it worked, Greg didn't just explain it, he designed and 3D-printed a replica.

The strip of numbers along the top represent the minutes, and the number inside this strip is the hour. Despite it rotating in four different axes, it's powered by just a single motor and a Geneva drive.

Not only does it look undeniably cool but, unlike most clocks that redesign the mechanical system, it's actually easy to tell the time on this one. □

Right ◈
The clock showing the transition from 11.59 to 12.00

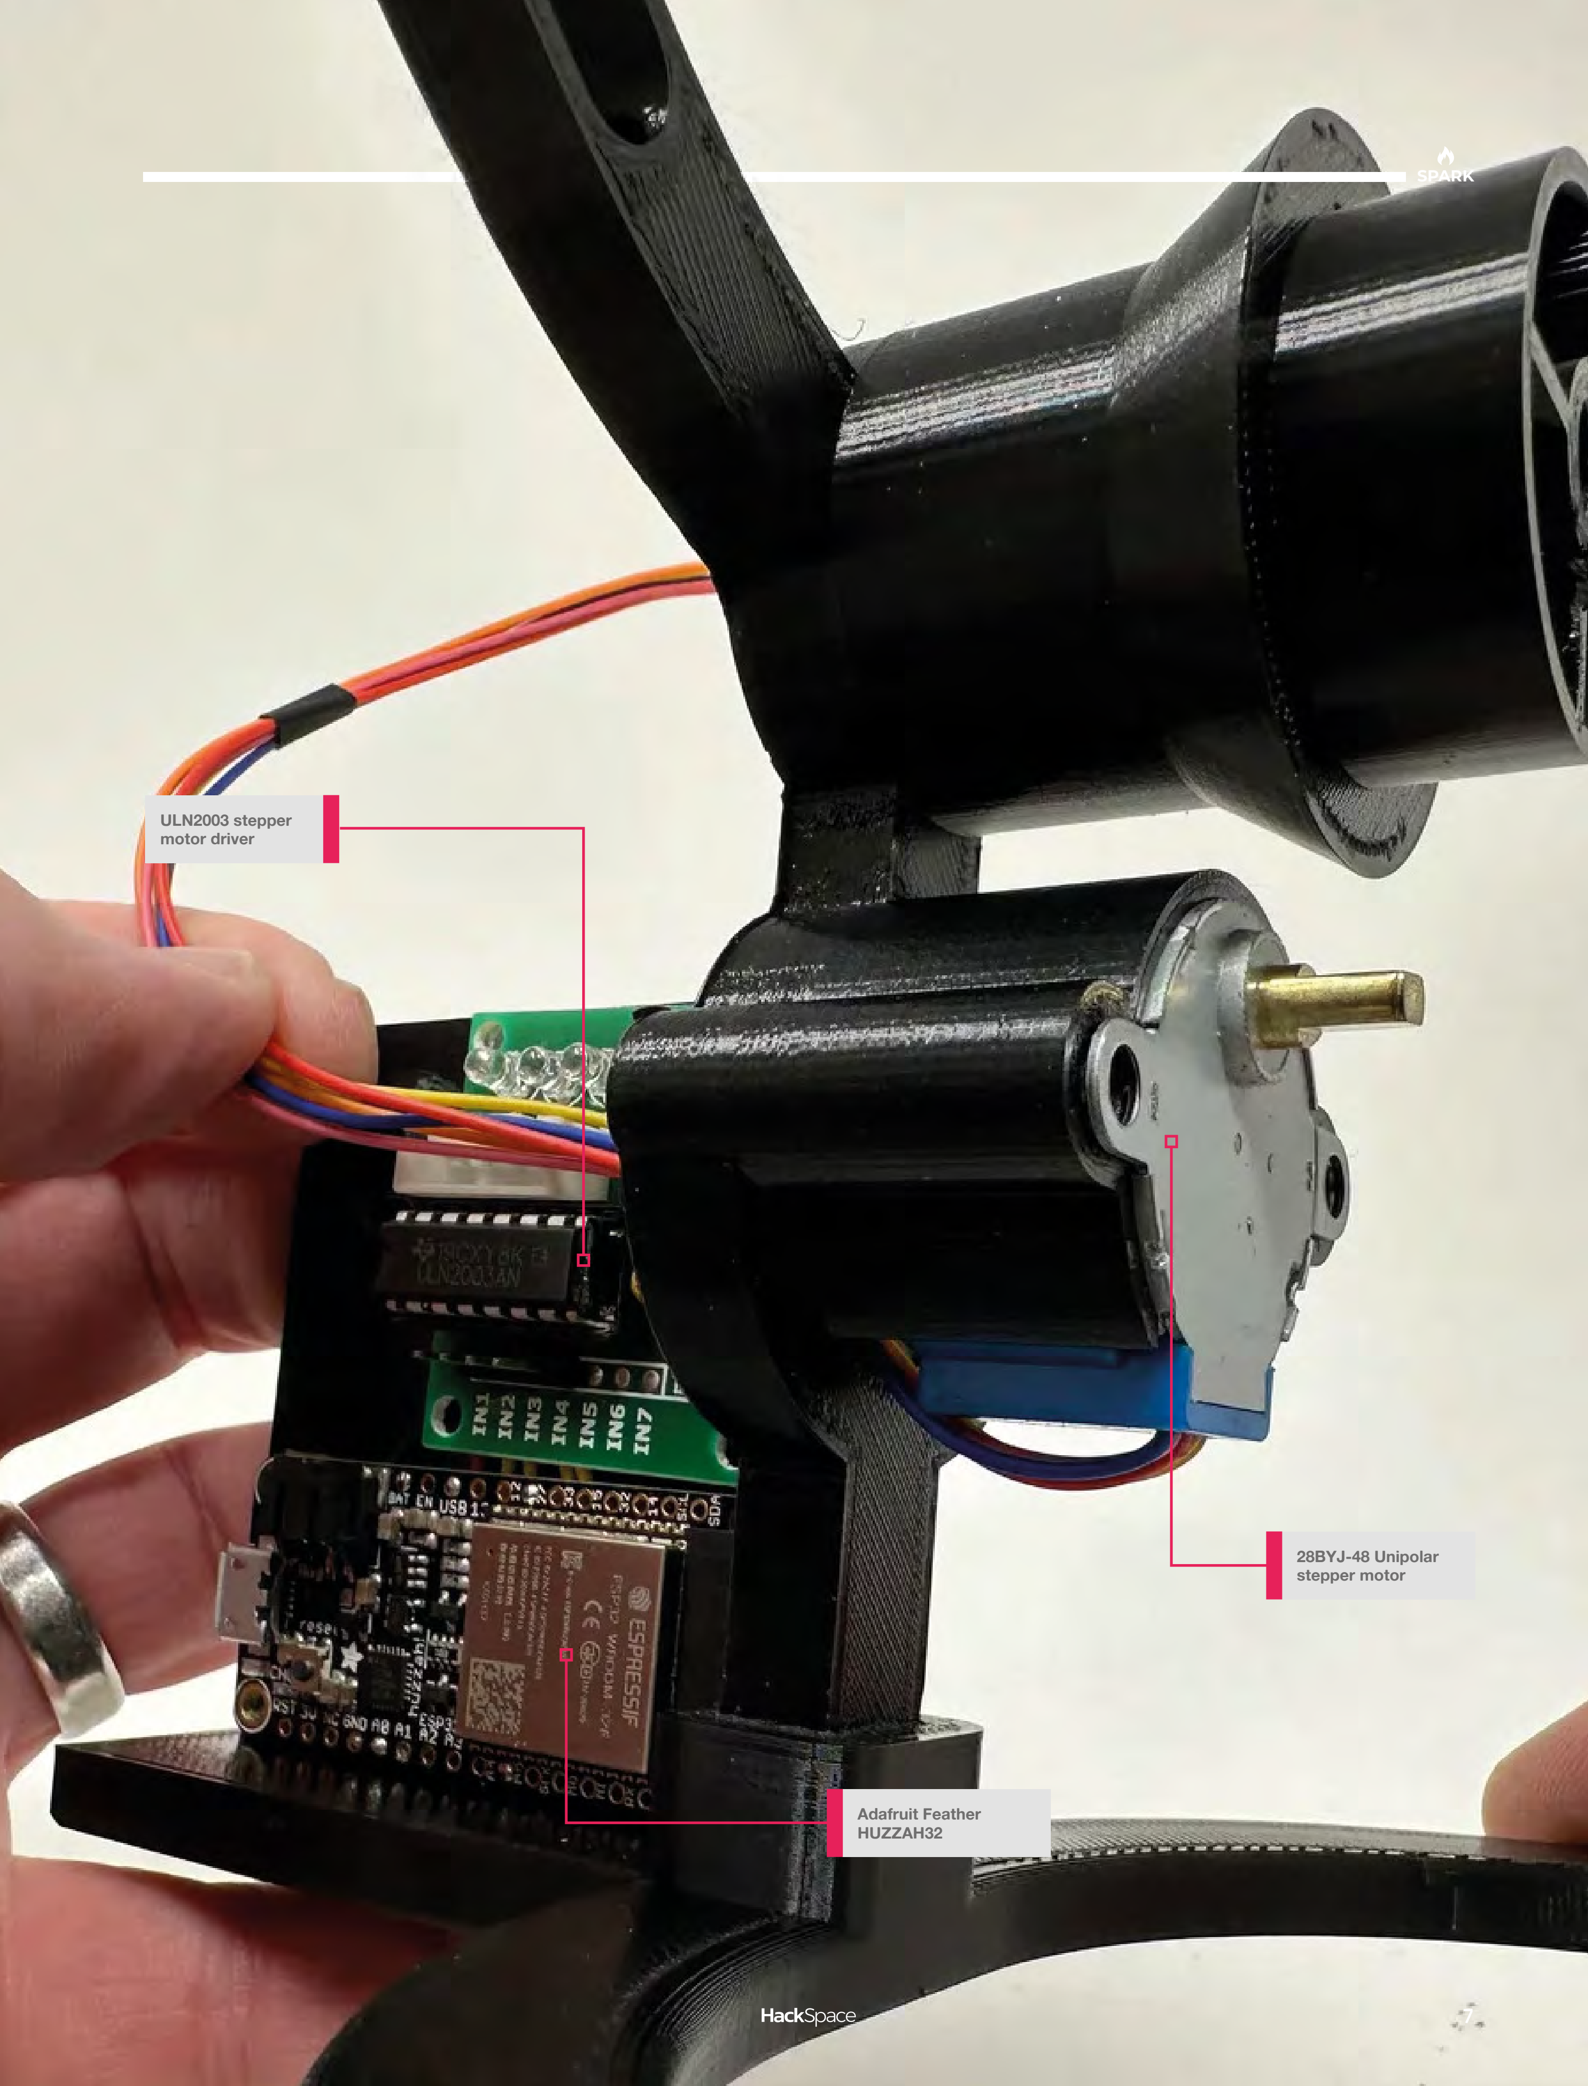
ULN2003 stepper motor driver
28BYJ-48 Unipolar stepper motor
Adafruit Feather HUZZAH32
ULN2003AN
IN1
IN2
IN3
IN4
IN5
IN6
IN7
ESPRESSIF

Mini PC

By **David Li**

hsmag.cc/RetroDesktop

If there are two things that we love, they're big versions of little things and little versions of big things. In both cases, they're made twice as cool if the new version works as the original does.

Here, we've got a tiny version of a desktop PC, and it actually works exactly like the original. Well, almost. The original was a 286 from David's childhood, while this version is based on a Raspberry Pi. That obviously means that the CPUs run completely different instruction sets but, thanks to the magic of emulation, this tiny PC still runs old DOS programs.

Cool-factor alone is enough to make this project worthwhile but, to David, nostalgia was also a motivator. "A big driver in my personal projects is the connection I have with the subject matter I'm working with. While retro-themed toys are nothing new (pun intended), I wanted to create something that was a close representation of an item that was a major part of my childhood. For me, building something like this was more than just a gaming nostalgia trip; it's also a physical representation of other memories from that time, like school, the house that we lived in, and family. It's also amusing knowing that the Raspberry Pi running it is about 7500% more powerful than the original 286 PC it's pretending to be." □

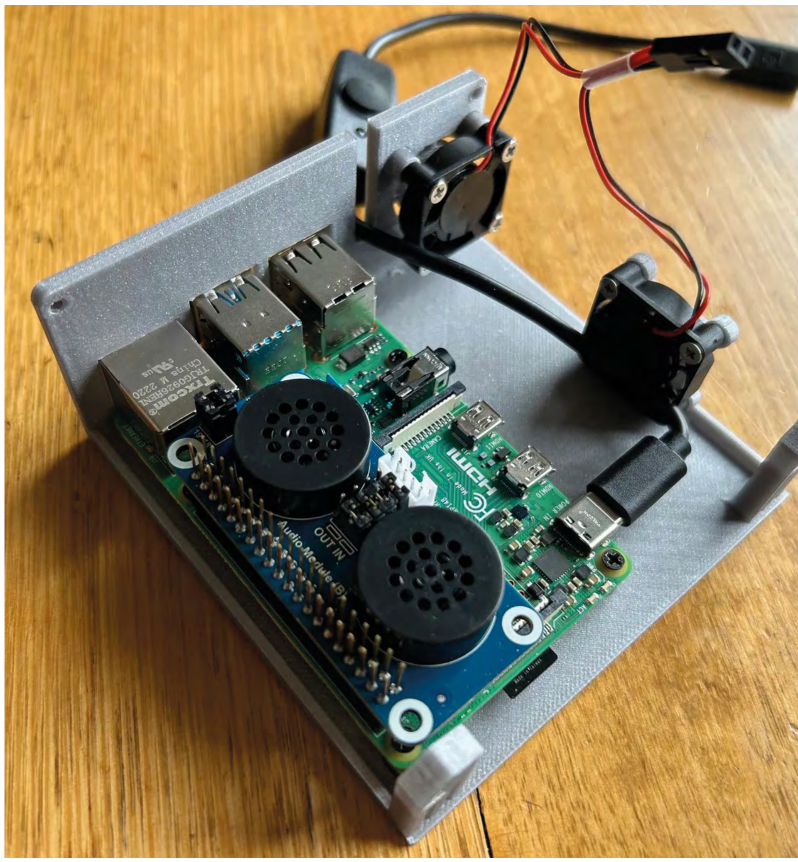

Above A dual speaker expansion board fits in the base to provide sound

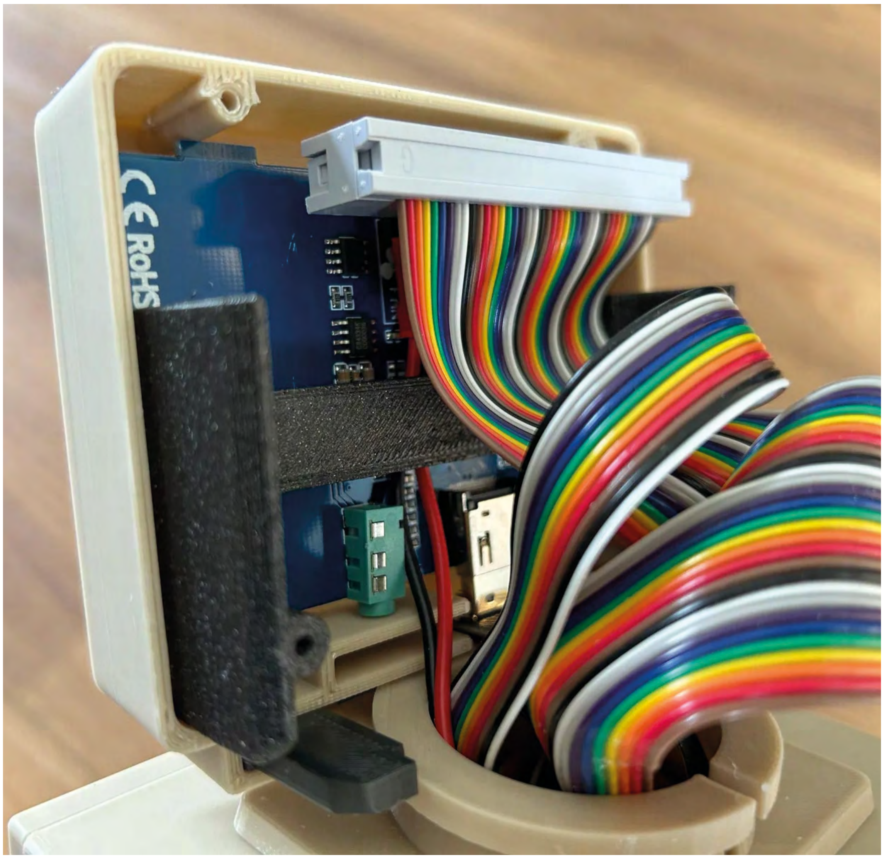

Above We particularly love how the SD card slots in the floppy drive

Left Graphics come out via a 3.5-inch Waveshare HDMI screen

Solder Sustainer

By **Justin Atkins** hsmag.cc/solderer

Soldering is one of those key maker skills. Typically, you take a soldering iron in one hand, some solder in your other hand, then hold a part with … hang on, we've run out of hands. Sometimes you can use clips or vices to hold wires and components together, but wouldn't it be very handy if you had a spare hand?

The Solder Sustainer gives you this extra hand, not by increasing the number of appendages you have, but by letting you control both the soldering iron and the solder wire with a single hand. Solder is fed out of a tube next to the soldering iron itself. Crucially, this tube is slightly flexible so you can bend it with one finger to direct it where you want to.

We're not sure if this is really the ultimate solution for all things solder-related, but it's great to see people creating tools to help them make stuff. Justin says that the Solder Sustainer works well for through-hole and larger surface-mount soldering, but he's working on version 2, which he hopes will be better for smaller surface-mount parts. □

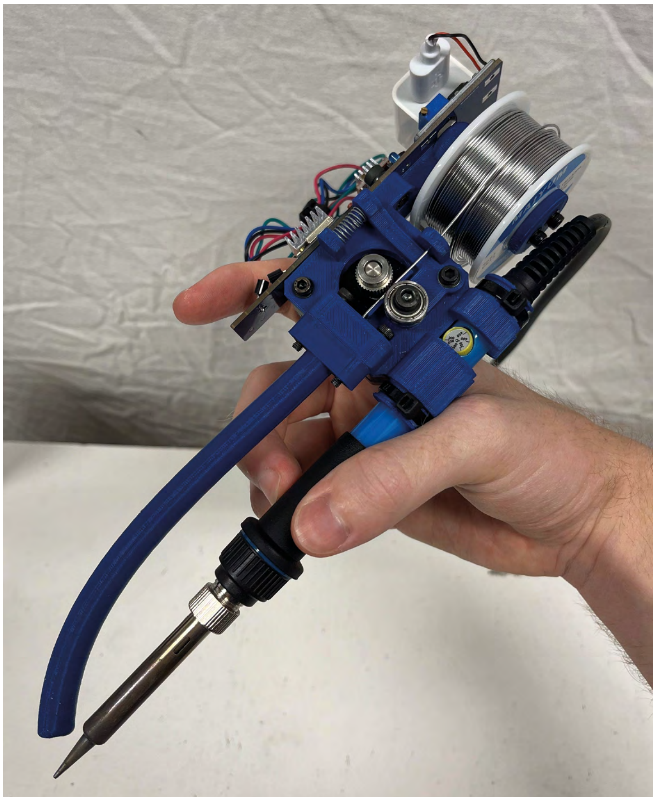

Right It might look a little cumbersome, but freeing up a hand when soldering can make things quicker and easier

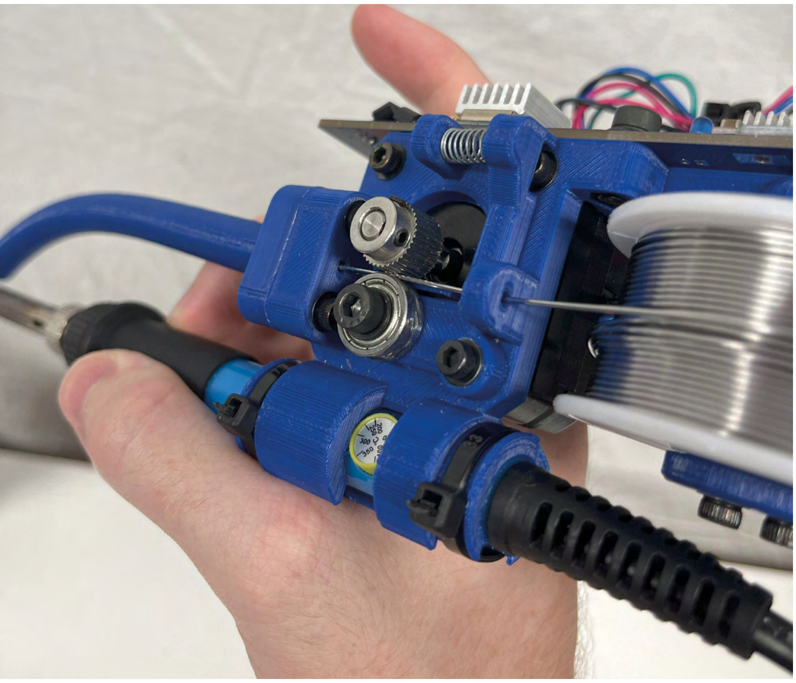

Left The stepper motor feeds solder forwards and backwards so you can get just the right amount of solder in your joint

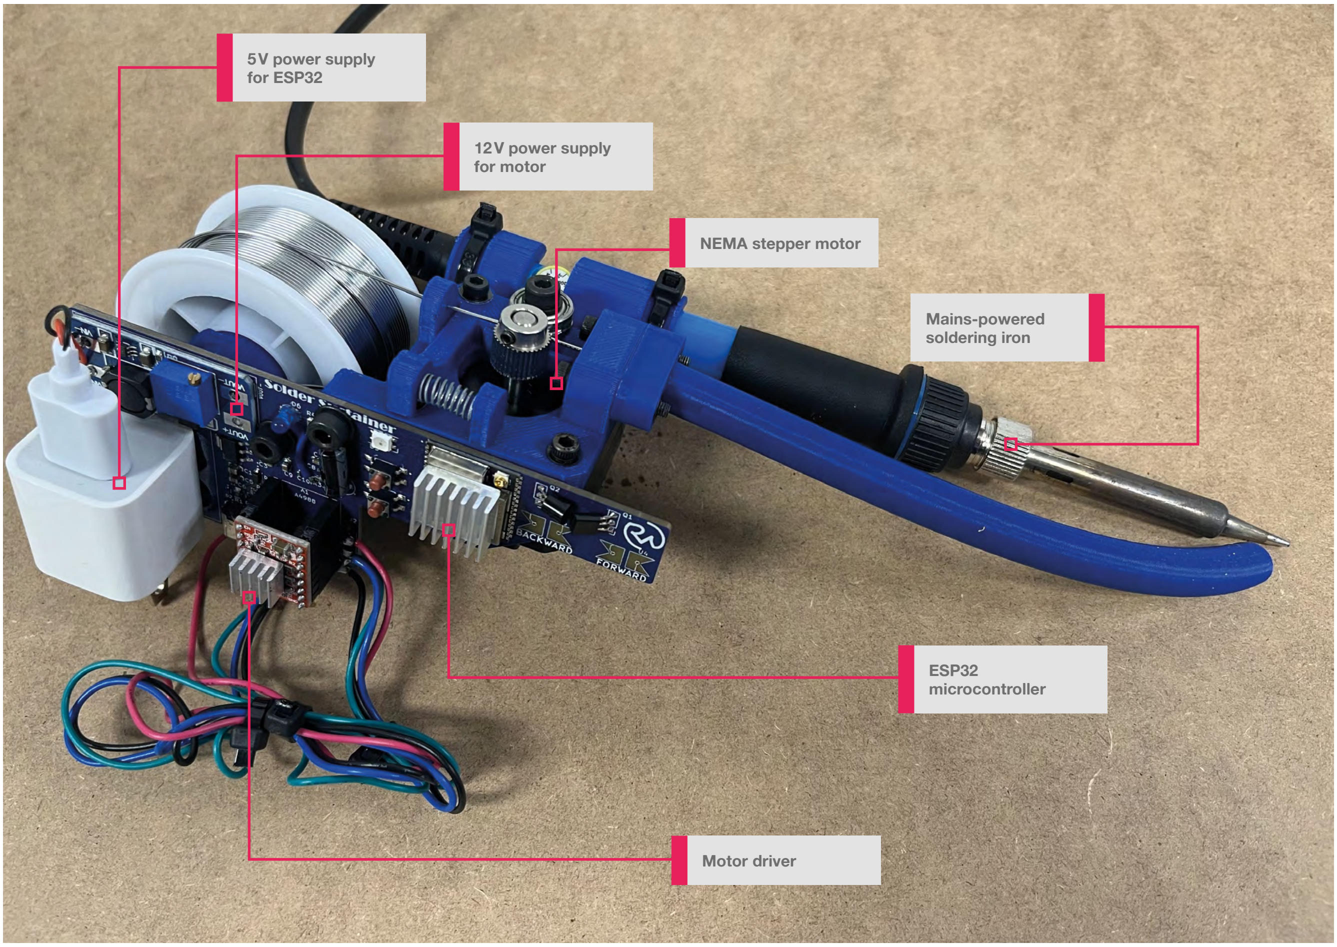

Ocreeb MK2

By **Salim Benbouziyane** hsmag.cc/ocreebmk2

Most people in the modern world use computers for their job – we spend hours every day interacting with our machines via a QWERTY keyboard – a Victorian design that was never meant to be used to program computers, control robots, manipulate enterprise resource planning software, or any of the myriad of other uses that we give it. Why do we stick with this input system designed in the time before computers?

Salim, for one, isn't going to. He's designed an adaptable system of modules that you can plug together to create whatever layout you want. You can add buttons, twiddly knobs, turn-y-handles, sliders, and more. Ocreeb MK2 runs the KMK keyboard firmware, so you can map whatever actions you want to the various buttons, knobs, and dials.

The Instructable guides you through how to assemble your own macropad, but it is quite an involved build including some surface-mount soldering. Nothing that should put off an experienced solderer, but probably not ideal for your first project.

We love the idea of having a reconfigurable workstation setup, and it helps that this one just looks gorgeous. □

Above These edge connectors allow you to join different modules together easily

Right Add whatever combination of modules you need for your workflow

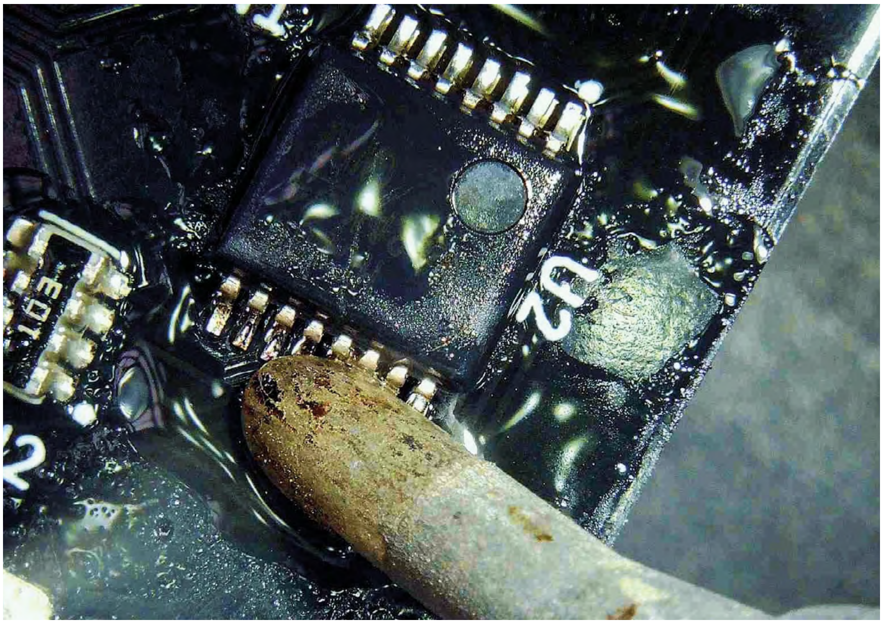

Left
This project does involve some surface-mount soldering

Stained glass succulent

By Jason Satira hsmag.cc/succulent

Stained glass work might sound like an exotic form of making, but when you get down to it, it's really just fancy soldering. Once you've formed the glass into appropriately shaped pieces, you pop some copper tape on the edges, then solder them together. Yes, we're glossing over the details a bit here, but the basic technique isn't too complex.

Jason's been experimenting with stained glass and noticed that plant designs are common at the moment. However, he couldn't find a suitable downloadable design to start with, so he decided to create one and share it on the Instructables website.

The design is based on a succulent – a type of plant that grows in arid regions and its drought tolerance makes it popular with people who forget to water their plants. The succulent's leaves sit in concentric rings, getting gradually smaller. To assemble it, you just have to start with the largest (bottom) ring of leaves and work upwards. The final result is a stunning plant that shimmers and shines in the light. □

Below After you cut the leaves out of glass, you need to tape the edges with copper tape before soldering them together

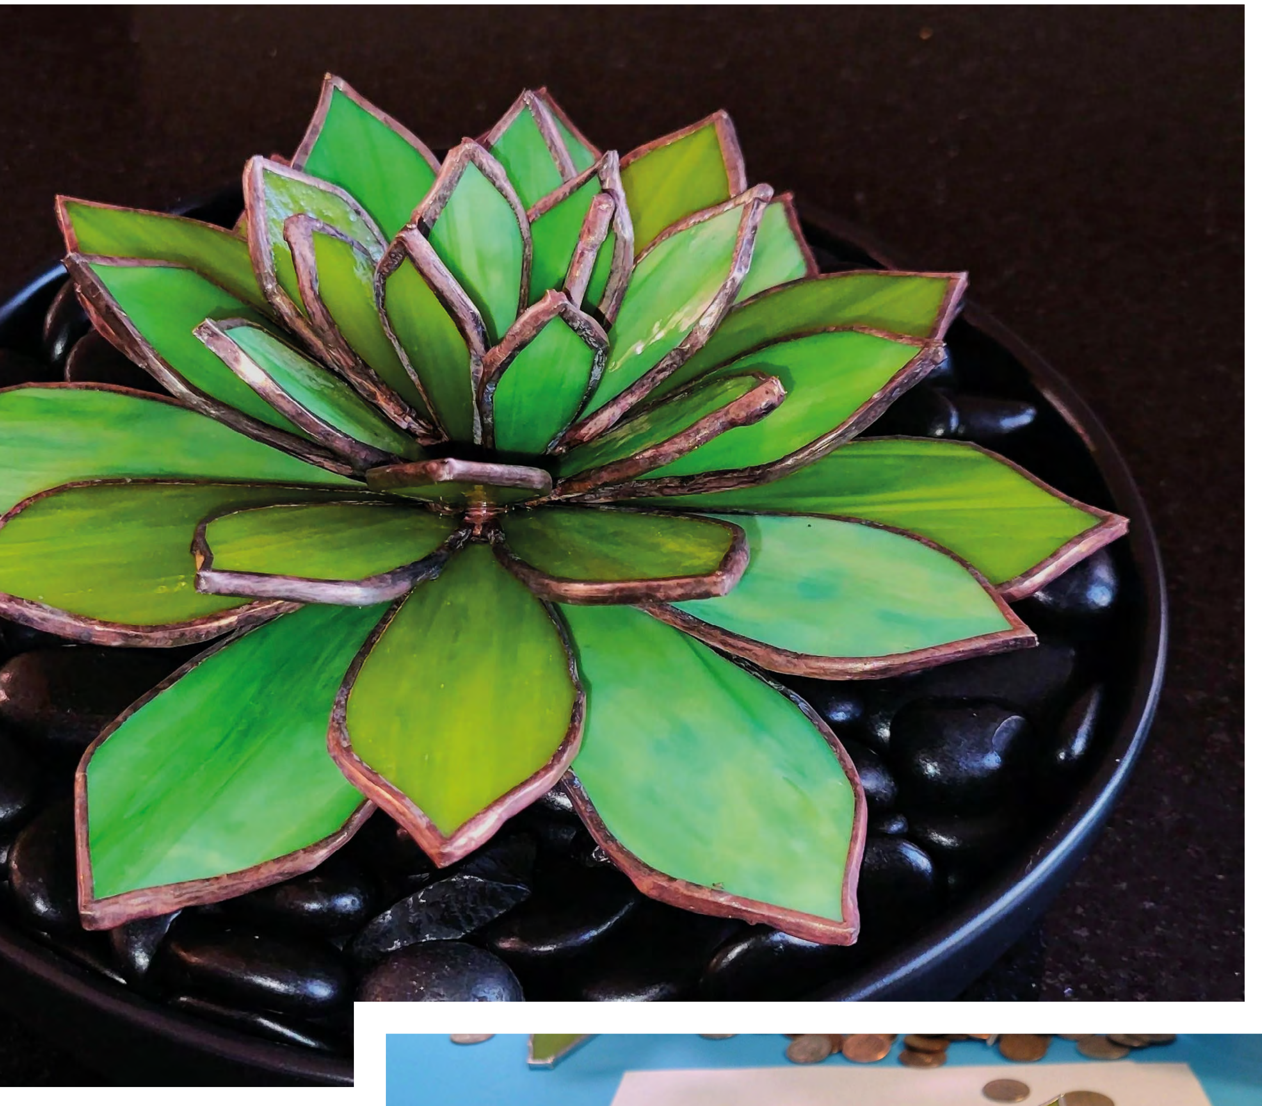

Above Develop your soldering skills and decorate your home at the same time

Right Jason used piles of coins to hold the leaves in place while he soldered them

HackSpace
TECHNOLOGY IN YOUR HANDS

LENS

HACK | MAKE | BUILD | CREATE

Uncover the technology that's powering the future

PG 28

HOW I MADE: RF CYBERDECK

The perfect radio workstation didn't exist, so Tom Mladenov built it

PG 34

INTERVIEW: TOM RANSON

An army of scientists live under the museum. We caught one and found out about 3D scanning

FEATURE

ULTIMATE WORKSHOP

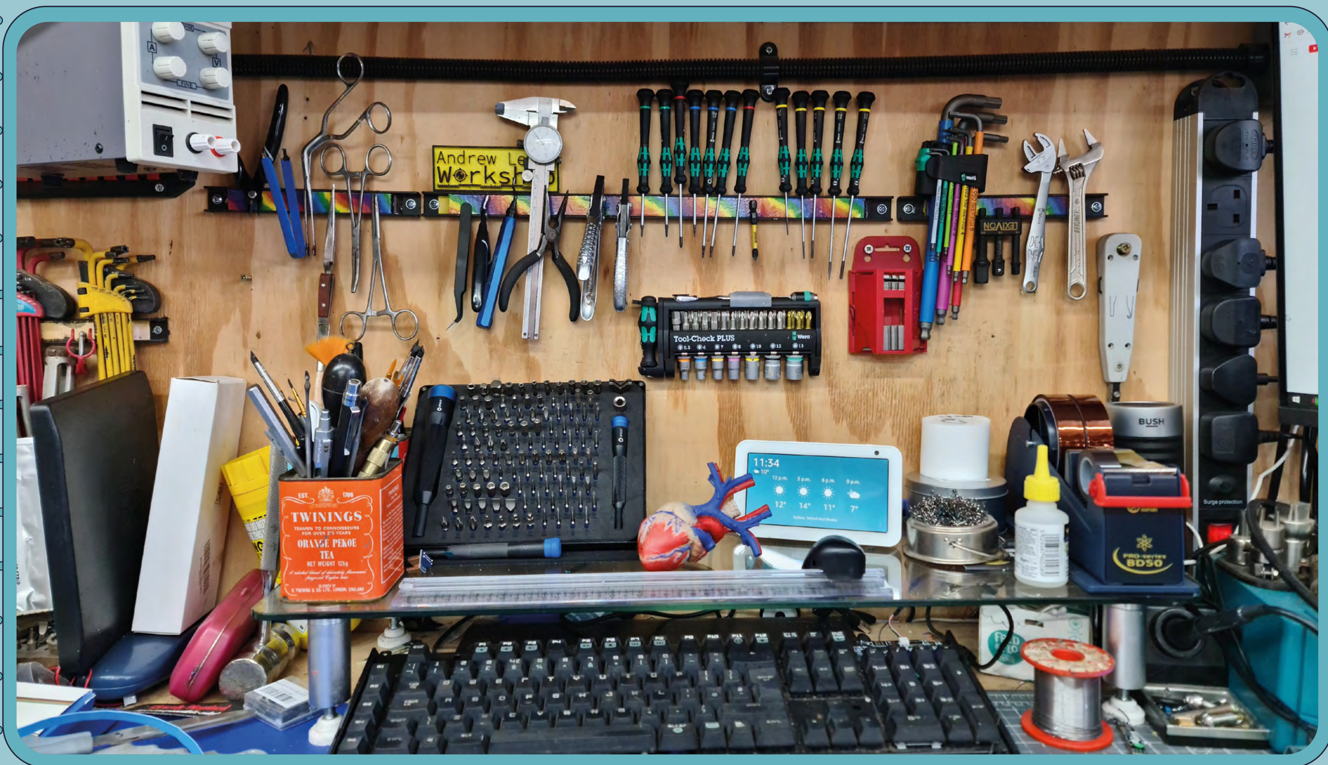

By **ANDREW LEWIS**

Every maker dreams of having their own workshop. Even those fortunate enough to already have a space of their own wish that they could change a few things to make it perfect. It turns out that a perfect workshop isn't a static thing. It's a living space that changes as your skills advance, and as the focus of your projects evolve. Even though a leather-working studio is going to look a lot different to a metalworking space or an electronics lab, there are areas of overlap between the tools and methodologies employed by almost every maker, and that's what this article is about.

There are two ways to get the perfect workshop – you can do it by accident or by design. In the former case, you start with a largely empty space and, as you complete projects, the space fills with the tools and materials that are most useful to you. You'd typically see this type of workspace in an old, established firm producing handmade goods. It's the sort of workspace where you don't touch anything, and you can't find anything unless it's your space. It can be a comfortable way to work, but it's almost impossible to move to a different space if you need to, and it's impossible to work with someone else because you'll have to keep finding the tools for them because only you know where everything is. Chances are, safe working practices in a space like this involve squinting when you're welding, and pouring a cold cup of tea over any fires that break out. Alternatively, you can take a more structured approach to your workshop design, take a look at other people's workshops, think about how you want to work, decide on the tools that you need, and then design a workshop accordingly. This is by far the most sensible approach to creating a workshop.

BUILD A MODEL

Assuming you have acquired a suitably sized space to work in, the most important thing you should do is make a 3D model. It doesn't matter if the model is made from cardboard, 3D-printed plastic, or is modelled virtually. Getting an accurate feel for the empty space will let you make good decisions about layout in a way that a 2D map can't.

The most valuable commodity in a workshop is space. It's very easy to fill your workshop with tools and equipment to the extent that you don't actually have any space to work in – so think of the largest item you're likely to want to make, and make sure that you leave enough space to move around that object with any equipment you need. Consider dividing your space into virtual zones or work areas so that you know what sort of task will happen in a particular place. That way, you'll be able to make better decisions about the routing of cables, ventilation, and storage for that zone. Your 3D model can be very helpful for this.

Do not keep infrequently used specialist tools or large amounts of scrap materials in your workshop. Either rent tools as you need them, or store them off-site. If you have leftover materials from a project, dispose of them unless you are absolutely sure you will need them in the near future. It's helpful to remember that empty space costs money, and disposal doesn't mean destruction. It might be possible to sell or recycle any leftovers from a project, or even store them off-site if you have space somewhere else.

Above
Splitting your workshop into clean and dirty zones can help prevent contamination of sensitive equipment with dirt or oil. Some tasks, like grinding or welding, should never be done in a workshop with a clean area – they should only be done outdoors or inside a dedicated workshop

QUICK TIP

Put wheels on everything. Your needs might change in the future, so put wheels on everything that you can. Being able to move a surface or storage wall easily can be very useful.

CONTROL THE POWER

Most people with a home workshop have a list in their head of what pieces of equipment they can use together without tripping the breaker or blowing a fuse. It's not unusual for a shed or outbuilding to have a single 16A circuit for the electrical sockets and lights, and that isn't a lot to work with when you consider things like electric heaters or kettles, laser cutters, compressors, and welders. Ideally, you'll have an electrician connect 30- or 40-amp circuits to your workshop with a choice of single- and three-phase power, and you won't have to worry about what tools you run. However, unlimited control over your power delivery isn't always possible, and there are a couple of things that you can do if you're struggling.

Fit larger tools with soft-start circuits if possible. For a lot of large machines like air compressors, the biggest power draw occurs when the machine first starts, as motors overcome inertia and friction in the system. Once the machine is running, the power draw is less intense. A soft-start module slowly ramps up the power delivery to the machine, avoiding the initial spike that might trip your power breaker. You can separate out non-essential electrical devices like electric heaters and wire them all onto a circuit with an isolator switch. That way, you can temporarily switch off all of the non-essential devices at once, maximising the amount of available power for other machines in the workshop. Critical machines like laser printers and computers can be given their own UPS (uninterruptible power supply) so that jobs won't fail if the power breaker does trip.

In the case of something like an air compressor, a larger tank will mean that the motor needs to run less often. A large tank makes it possible to complete a job with the compressor motor turned off, then refill the tank at a later date. This can also reduce noise levels if you're forced to work late. →

QUICK TIP

Always have emergency lights in a workshop. You don't want to be stuck in the dark with potentially dangerous equipment all around you.

Above
In the UK, messing with your consumer box or electrical systems isn't acceptable and will probably invalidate your household insurance unless you're a qualified professional. If you have to run wires to your workshop, get a professional to check the work and connect it to the consumer unit

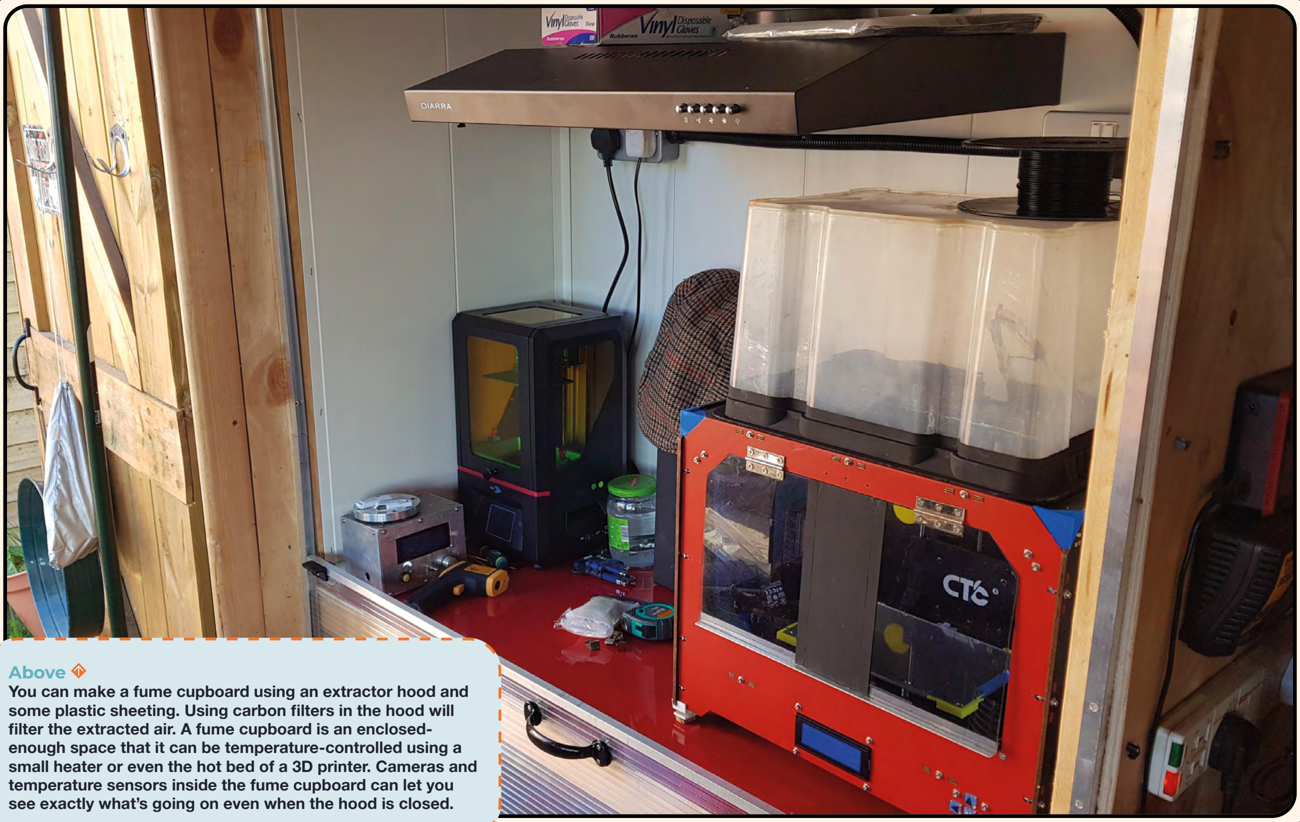

Above
You can make a fume cupboard using an extractor hood and some plastic sheeting. Using carbon filters in the hood will filter the extracted air. A fume cupboard is an enclosed-enough space that it can be temperature-controlled using a small heater or even the hot bed of a 3D printer. Cameras and temperature sensors inside the fume cupboard can let you see exactly what's going on even when the hood is closed.

LET THE ROOM **BREATHE**

A workshop should be a comfortable space for you to work in, but it should also be comfortable for your equipment and consumables. Air quality, temperature, and humidity are all important and need to be controlled in a workshop. The amount of volatile organic chemicals (VOCs), dust, and other particulates must be kept to a minimum to reduce the chance of long-term health problems like uncontrolled explosive defenestration, nausea, headache, and death.

Use a dedicated fume cupboard with active extraction for hazardous jobs, and use extractor fans with ducts to remove dust, smoke, and fumes while you work. Be careful about venting your fumes into the open air. Fit activated carbon filters in your system and change them regularly to reduce airborne pollutants. Check that the extraction system is working effectively with an anemometer, and regularly check the VOC levels inside and outside your workshop.

A shop-vac is a good way to reduce airborne particulates, and many power tools have adapters that allow you to fit a vacuum hose to them. You can reduce the need to empty the shop-vac by adding a cyclonic filter or dust separator. The cyclonic filter sits between the vacuum cleaner and the machine you're using, and separates larger particulates into a sealed container. While they're more suited to built-in vacuum systems, it is possible to add them to portable vacuum cleaners. To reduce particulates even further, a dedicated filtration system like the Scheppach 100W (**hsmag.cc/Scheppach**) can filter the air. If you're working regularly with MDF or powdered materials, a filter system like this is essential.

A workshop should always be kept near room temperature. Letting some equipment/consumables get down near freezing temperatures can spoil them. Pipes can crack, oil will thicken, and paint will separate. Laser tubes can shatter if they are started too cold, and metal rails can buckle. A radiator with a frost setting should be enough to keep most insulated spaces out of the danger zone. An active temperature controller can keep the coolant in a laser at a preset temperature, ensuring that the tube won't freeze.

QUICK TIP

Lighting is one environmental factor that gets ignored. You need more light. Good overhead lighting, desk lamps, and work lights make a huge difference.

LET'S PLAY: **KEEP, STORE, DISPOSE**

There's an old saying about storage in the workshop: "A workshop isn't a storeroom, and a workbench isn't a shelf." While it's OK to store some commonly used tools and consumables in your workshop, don't make the mistake of letting the space become something that it isn't supposed to be. The more things you store in the workshop, the less space you'll have to work. Adam Savage talks about storage in his Tested series, specifically the idea of "first-order retrievability" – where the tools and materials you need to complete a task are visible and easily retrieved. To paraphrase Adam further: "drawers are where tools go to die."

This is true, but can also be thought of as "If you can't find it, you don't have it." In a smaller workshop, it's necessary to hide things from view to maximise your available space. It isn't a big problem if you indicate clearly where things are stored using labels.

Storing tools vertically is a great strategy in a workshop, provided it doesn't cause a safety problem. Having tools on hooks above chemical tanks or machine tools is an industrial accident waiting to happen.

There is no one-size-fits-all solution for materials storage, but the Stanley Fatmax Waterproof Pro Organisers (**hsmag.cc/SF_Organiser**) are great for smaller components, and the 'deep' versions of the organiser can hold larger items like stepper motors and batteries. The Recycled versions of the Really Useful Box (**hsmag.cc/UsefulBox**) are great for large items, and can be labelled easily with a DYMO labeller. For portable storage, the Hobbycraft three-tier storage trolley is excellent (**hsmag.cc/StorageTrolley**) and is low enough to slide under most desks when it's not in use. →

QUICK TIP

If you don't want to use a tool because it's too difficult to get to, then you might as well get rid of it.

PLAY IT SAFE

The most important part of every workshop is safety. Even if the workshop is designed entirely for one person to use, you still need to consider the safety of others when you are working. The use of goggles, gloves, earplugs, and machine guards is well described elsewhere and doesn't need to be repeated here. Hopefully most readers will understand that working alone is dangerous, and remember to tell someone when you should be finished working. There are other considerations that you should make to improve the safety of your workshop. Adding a door alarm or a bell to your workshop door will let you know if someone comes in unexpectedly, and adding a warning light outside the door will warn people when you're doing something dangerous and probably shouldn't be disturbed. Adding a lock to the door is a risky proposition if you are the only person inside the workshop that can unlock the door. If you are injured, help might not be able to get to you through a locked door.

You should always have a well-stocked first aid kit with an eye-wash kit, and there should be enough fire extinguishers, fire blankets, sand, absorbent pads or sawdust to handle fires and chemical spills. Important safety equipment needs to be easily accessible from the place that it will be used, and you should have spare items to cover breakages and possible visitors.

Something often overlooked in the workshop is the idea of keyed access or interlocks to larger machines. If you add an electronic key to a machine, you can effectively control who can start and use that machine, and you can also use physical locks to mechanically isolate the start buttons of a machine. Adding a key to a machine isn't usually difficult, and the safety benefits you will get from the addition are enormous.

Danger signs are an important defence against injury through ignorance. If you have dangerous tools or materials, make sure that you use the appropriate signs to warn people of potential dangers. Remember that signs and other safety features aren't just there for your safety. They are there for the safety of anyone who enters the workshop.

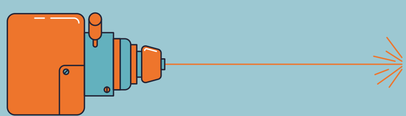

TOOLS GLORIOUS TOOLS

It's almost impossible to recommend the perfect toolkit, because different interests require different tools, and there are many ways to get the same result with different tools. However, there are quite a few tools that should be considered as essentials for everyone's workshop and a few that are considered as 'nice to have.'

As a general rule, having two sets of every important tool is good for several reasons. If you break a common tool, you've got a spare of that tool to work with and you can carry on. You can also keep a 'dirty set' of tools that you use for messy jobs, while the other set remains clean.

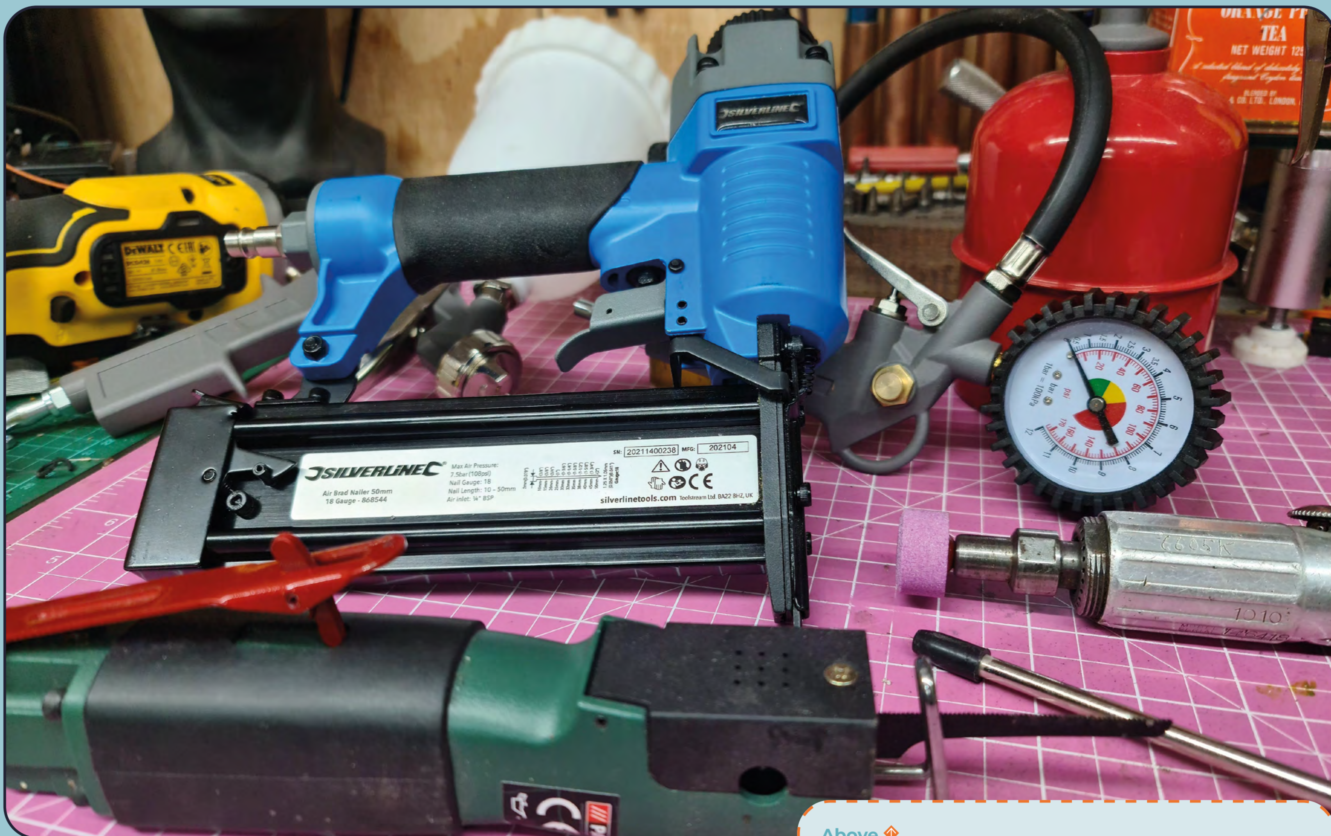

Above
Not every power tool is a battery tool. Don't discount the idea of air tools – they can be cheaper, faster, more powerful, and quieter than their electric equivalents if you already have a compressor

YOU'VE GOT THE POWER

A basic set of power tools in combination with the right sort of jigs can duplicate most of the capabilities of a workshop full of large-scale machine tools. Note that the key word here is capability, and not capacity. If you are running a factory making wooden benches, it makes sense to have a floor-mounted thicknesser and jointer because they can process industrial quantities of material very quickly. For general making, simple power tools are usually enough.

It's good to get tools that share a common branding so that batteries can be interchanged without needing an adapter and multiple chargers. Every maker will have their own favourite range of tools, with Triton, Milwaukee, Bosch, and DeWalt being popular names in the field. DeWalt kits usually represent good value, and spare parts are easy to get. With batteries that slide onto the bottom of the tool's grip, many DeWalt tools have slightly smaller and better-contoured handles than their competitors. This makes them a more comfortable choice for makers with smaller hands. For the average workshop, a drill with a reverse and hammer function, a jig-saw, and a multi-tool will be the essential kit that you'll use to do the majority of your making. If you do more metalworking, an angle grinder and portable bandsaw will be useful. If you do more woodworking, a router and circular saw will probably be more useful. It's worth noting that tools like routers and circular saws can often be used to cut aluminium with the appropriate blades. Cutting steel usually requires a separate saw with a lower-speed motor.

Belt and orbital sanders are good for shaping and finishing materials quickly, and can often be bench-mounted, giving them a similar functionality to a belt sander. Additional power tools will give you very specific advantages in certain situations but probably aren't necessary for everyday working. An impact driver is a common power tool these days, and it's excellent for driving in screws. However, a drill can do the same job and, in certain materials like concrete, the impact driver would shatter the material. Right-angle drills are very useful if you're working in a tight space, but are unwieldy in normal use. The same applies to other tools like die grinders and mouse-style electric sanders. →

QUICK TIP

A hot-glue gun is a power tool, and it deserves more love than it receives. Get a hot-glue gun, and thank yourself for years to come.

LEFTY LOOSEY, **RIGHTY TIGHTY**

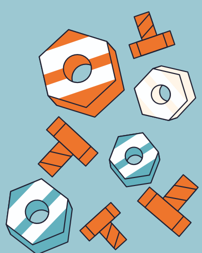

When it comes to screwdrivers, there are a mind-boggling number of choices to make. For electrical screwdrivers, it's difficult to fault Wera's offerings. Its Kraftform fixed (**hsmag.cc/KraftformFixed**) and interchangeable (**hsmag.cc/KraftformChange**) screwdrivers are precise, strong, and ideal for general workshop use.

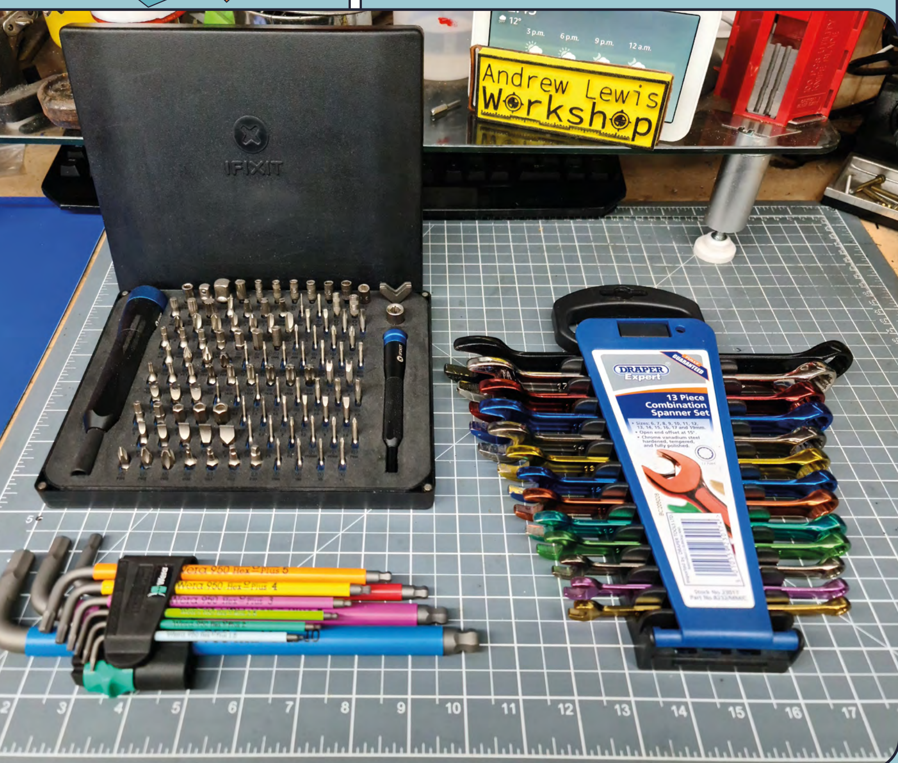

- For smaller items, Wera offers a basic Kraftform Micro Precision set (**hsmag.cc/KraftformMicro**), although for the serious maker who's looking to repair modern electronic devices, it's difficult to ignore the excellent iFixit Manta Precision set (**hsmag.cc/iFixitBitSet**).
- Spanners are another tool where there are literally thousands of options, but for general open spanners, the Draper Expert series (**hsmag.cc/DraperExpert**) offers a reliable spanner with a handy coloured finish that makes it easy to identify different sizes.
- Wera's Allen keys (**hsmag.cc/WeraAllen**), socket sets, and Torx wrenches have a similar colour-coding scheme that makes it very easy to pick the right tool when you have several on the desk.

Above
It's often said that you should buy the cheaper version of a tool first to see whether it suits your workflow, then get the higher-quality tool if you find it useful. This might be true for unfamiliar tools, but for ubiquitous items like screwdrivers and spanners, go for the expensive tool first. These are tools you'll always have a use for, and buying quality pays off

ELECTRONIC **DREAMS**

A lot of nonsense and snobbery exists in the electronics world about which tools you should be using – especially when it comes to multimeters. If you're working in an industry where a tenth of a volt is enough to ruin your plans for global domination, then you should really look into buying a very expensive set of calibrated tools. If you're an enthusiastic tinkerer or working on general repairs, a basic multimeter is going to be absolutely fine.

Thirty years ago, the first item to land on an electronics workbench after the soldering iron and multimeter would be a CRT oscilloscope. These days, an oscilloscope is much smaller, much less expensive, and works over a USB connection. For a basic unit, something like the Hantek USB oscilloscope (**hsmag.cc/HantekOscilloscope**) is fine for everyday use.

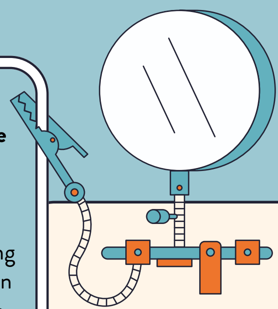

- If you're likely to be working on battery packs, a battery welder (**hsmag.cc/SeesiiWelder**) is absolutely necessary. It might appear to be a niche product, but it offers a solution that nothing else can really achieve. Trying to solder a LiPo cell with a conventional soldering iron will damage it, so mechanical contacts or welding are the only real choices you have for joining cells into a pack.
- For soldering, an SMD rework station that combines a heat gun and a soldering iron (**hsmag.cc/KatsuWelder**) is great for desktop use. The air gun is an extremely useful tool to have on the desktop, making fast work of small jobs like applying heat-shrink tubing or bending plastic pieces to shape.
- A USB microscope and borescope are useful additions to a soldering station, particularly if your eyes aren't as sharp as they used to be. Seeing clearly into confined spaces can make all the difference to a soldering job.

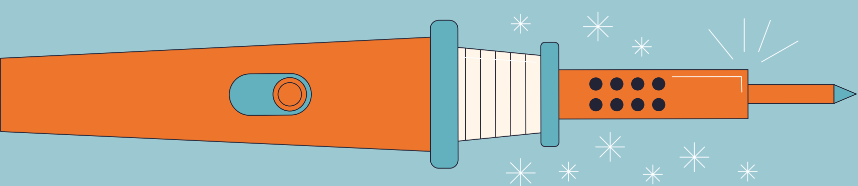

TOMORROW'S TECHNOLOGY TODAY

We are fortunate enough to have access to a few tools that simply didn't exist when our parents or grandparents would have been building their workshops. An obvious example of this is the 3D printer, which is one of the most versatile workshop tools if you know how to use it correctly. While desktop 3D printers typically produce plastic or resin parts, their strength isn't really in the production of finished parts. Their true strength lies in their ability to create jigs, frames, supports, temporary tooling, and moulds for other parts that radically simplify the processing of other materials. You can create moulds for casting, jigs and templates for accurate drilling and cutting, or even just shims and spacers to keep parts accurately positioned.

Laser cutters are similarly useful for creating short batches of repeated shapes, engraved designs, or accurate cuts on flat sheets. If you have a need to do this sort of activity on a regular basis, a laser cutter is a superb (although very expensive) addition to your workshop.

Lastly, the often overlooked workhorse of many workshops is the computer. Having reliable computers in your space to control machinery and do general work on is incredibly important. The best advice about computers is to keep multiple off-site backups of important data, and don't put all of your effort into building a single computer that does everything. Keep your computers separate so that a single failure doesn't take down the whole workshop. □

SUBSCRIBER BENEFITS:

> **Get every issue of HackSpace magazine delivered to your door**

> Learn from hackers and makers in our in-depth tutorials

> Early access to the PDF edition

> **Get a free Raspberry Pi Pico W**

hsmag.cc/subscribe

Subscribers will get a voucher giving them the chance to purchase one Raspberry Pi 5 from reserved stock at The Pi Hut (thepihut.com) for full retail price. Reserved stock means that these will be available even if they are out of stock for general purchase.

HOW I MADE

By **TOM MLADENOV**

SOFTWARE DEFINED RADIO CYBERDECK

Whether it is tracking stratospheric weather balloons, mapping out radio beacons or receiving weather images from satellites, RF (radio frequency) continues to fascinate me. What seems so far is always within reach. Many of my activities in this field require dedicated setups and separate equipment, which adds some logistical overhead, and hence I started pondering the idea of having an all-in-one tool that I can use to navigate the RF spectrum for any possible application and use it in the field. It would make use of software-defined radio (SDR) to allow using it for different purposes, ranging from listening to voice transmissions from the International Space Station and other satellites, to worldwide amateur radio traffic on the short-wave (SW) frequencies between 2 and 30MHz, for instance.

“THE UNIT NEEDED TO HAVE BASIC I/O FEATURES THAT ANY MODERN SIGNALS INTELLIGENCE SYSTEM WOULD HAVE”

Above
The Raspberry Pi SDR Cyberdeck

The unit would use a Raspberry Pi and host a suite of interesting SW applications to allow browsing the RF spectrum with it and decode various digital signals, including weather balloon beacons, ship or aircraft transponders, and Automatic Packet Reporting System (APRS) signals used by radio amateurs to bounce radio packets around nationwide repeater networks. With that initial idea, I set out to create what was to become my Swiss Army knife for the RF spectrum.

ENCLOSURE AND POWER SOURCE

As with many of my projects that I use in the field, it needs to be sturdy, resilient, and have a high degree of autonomy. The unit needed to have basic I/O features that any modern signals intelligence system would have. It needed to be waterproof, able to be taken airborne, run on internal or external power, and have both internal and external antenna connections. The latter would allow usage in a wide variety of deployments.

For the enclosure, I decided it should fit easily in a backpack and should be smaller than a shoebox. I settled on the Peli 1200 Protector Case, which measures 241 × 185 × 105 mm, offers IP67 rating, and has a pressure relief valve. The ‘Pelicase’ is fitted with a front panel frame holder from Peli, which is held in with self-tapping screws from the inside. This frame will support the user panel that will house the screen and control panel. The front panel holder frame is fitted with brass inserts to allow →

Left Front panel bus switches and indicators

> **"THE FIRST ORDER OF BUSINESS WAS TO LINE THE WALLS OF THE PELICASE WITH ALUMINIUM TAPE"**

bolting the front panel to it. Since the unit will be used for RF operations, the first order of business was to line the walls of the Pelicase with aluminium tape to shield the sensitive external antennas from noisy components inside the cyberdeck, as well as to not reveal the unit's RF signature in covert operations. For the battery pack, an Anker PowerCore 20100 power bank was used as it is a well-known and reliable power bank, ideal for this application.

Two DS18B20 digital sensors were also fitted to the battery pack with thermal silicone. Two clamps were made to fit around the battery pack, which each had a lug that attached it to threaded inserts on the front panel holder.

Below Power bank fitted in the aluminium-lined Pelicase

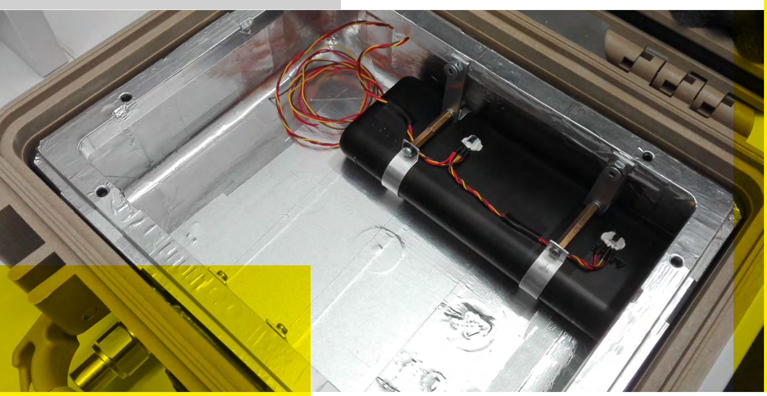

FRONT PANEL

The major components that needed to be accommodated on the front panel were the touchscreen, I/O ports, switches, and speakers. Since all electronics of the cyberdeck unit needed to be easily accessible from all angles, I decided to build the internals in such a way that the whole system stack would be attached to the back of the front panel, much like how the avionics in cockpits of aircraft are done. This way, when you pull off the front panel, the whole system comes out and can easily be worked on, rather than it being inside the box and difficult to access. This also meant that this front panel needed to be sturdy and not flex under the weight of an electronics stack that it is supporting.

I therefore settled on a 3 mm thick sheet of aluminium that my friend kindly CNC'd for me with all the required fittings. I drew the front panel in Inkscape and exported as SVG. The front panel measured 237.5 × 182.9 mm and was finished with a matte black anodization along with some white markings. The main rectangular cutout is for the 7-inch touchscreen.

The top section of the front panel shows a number of markings that depict

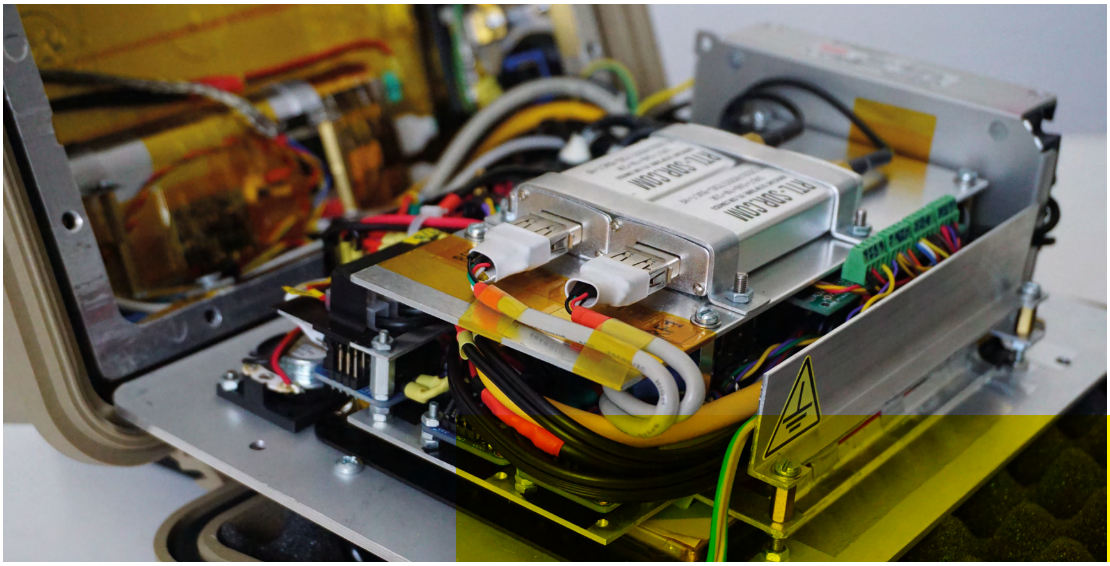

Left ←
Side view of the finished stack showing the two RTL-SDR receivers on the top

the external connections on the top side of the unit. Later on, however, the external layout would change a bit, leaving some markings inaccurate, but nevertheless, I left them this way. The bottom section of the control panel houses the speakers and main power bus controls, as well as two USB ports. The control panel looks complicated, but it's really not that bad. The RC switch controls the power source, which is set to either J1 for external power or INT for internal power. The switch marked J1 selects which external power is used, either A, which is the regulated 5 V DC input via USB or B, an unregulated 9–36 V DC input.

The switch marked CHG is to select which external power source (A or B) to charge the internal power bank from. The final two switches on the right are to power up the on-board computer, marked OBC. And finally, the switch to power on the monitor is on the far right and marked MON. Several power indicators show if the main bus has power, if the battery is charging, and which external power sources are detected. A headphone jack is available to connect a headset.

The front panel is mounted using M4 nuts in marked locations, and since this front panel has to be detached many times, I decided to use hex nuts to lessen the chance of damaging them over time. Two rack handles are mounted to the front panel to make it easy to detach from the unit.

THE STACK

As previously mentioned, the electronics stack is a big sandwich that is attached to the front panel. The stack starts with the Raspberry Pi 7-inch touchscreen mounted on the back of the panel. It has four convenient M3 threaded points to further build the stack using brass standoffs. Due to mounting the capacitive touchscreen to a big slab of aluminium, the touchscreen needed quite a bit of recalibration later on in the software configuration and setup process. On top of the screen sits the Raspberry Pi Model 3B with, attached to it, another DS18B20 digital temperature sensor, like the ones fitted to the battery pack. The Raspberry Pi has been fitted with a common heatsink. I decided to completely drill through the holes in the heatsink and then mount the Raspberry Pi to the standoffs that are attached to the back of the touchscreen. On the cover plate that tops off the stack are two RTL-SDR receivers mounted with brackets. The two SDR receivers each have a coax connector going to an external N connector for attaching antennas. On the side of the stack is a DC-to-DC converter that converts the externally supplied unregulated 9–36 V DC to 5 V DC for all subsystems. When using internal power, this unit is not used, as the power bank already provides a clean 5 V DC. The electronics of the stack are connected →

Above Right
Side panel for power and Ethernet connections

to the rest of the unit, with one big cable harness going to the unit's enclosure and battery pack.

CONNECTIONS

Since RF is at the heart of the operations of the unit, several external RF connections have been added. The top side of the unit features two N-style RF connectors which are bolted to the chassis with four screws, making them very rugged. The right side of the unit features two SMA connectors: one is currently used for attaching a WLAN antenna to the Raspberry Pi, and the other is unused and available for future use. The last RF connection is also located on the right side of the unit and provides a way to attach an external active GPS antenna. If no external GPS antenna is connected, the unit relies on its internal ceramic GPS patch antenna, which is located on the top side of the unit between the two N-style connectors. For power, I found these multi-pin military-style connectors online which are quite nice as they feature a bayonet-style connect and disconnection method. This multi-pin connector serves as external power input, but has some unused pins to allow connecting external control units or even sensors to the I2C bus of the Raspberry Pi. Finally, a chassis mount RJ45 connector allows us to connect a network cable which passes through to the Raspberry Pi's Ethernet adapter. The external connectors feature caps to protect them when not in use – they also have two rack handles mounted on either side to protect them further.

SOFTWARE

I wanted to be able to operate the unit in mainly two modes: a desktop mode, where the user controls the opened unit using the front panel, and secondly, a headless mode, where the unit's touchscreen would be off altogether, and the closed unit would be operated remotely via RJ45 hardline or a phone. Headless mode would, for instance, be used when the unit is inside a backpack or in the trunk of a car. From the software point of view, I also wanted the unit to run a control server that would expose an API over HTTPS via endpoints to allow easy integration with other systems. I created a server and client application with Python and FastAPI that runs locally on the Raspberry Pi.

Below Right
Usage in headless mode

Below
Main panel showing all subsystem data and status

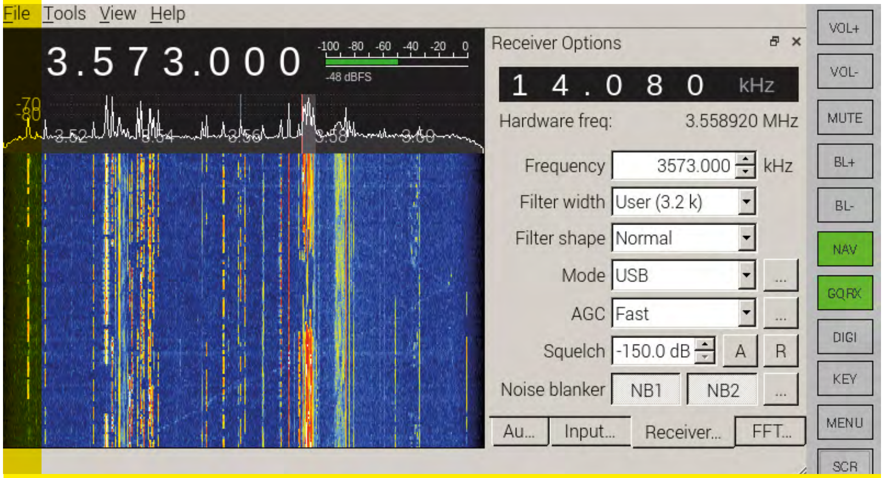

"THE FRONT PANEL NEEDED TO BE STURDY AND NOT FLEX UNDER THE WEIGHT OF AN ELECTRONICS STACK"

The GUI features all actions that the user can perform and talks to the control server using HTTP requests. For remote connectivity from a phone, the Raspberry Pi runs a VNC server to which the phone connects over Wi-Fi, allowing control over the same screen session from a phone or other computer in the network. Various applications can be started from the GUI, which all use the RTL-SDRs to browse the RF spectrum. The popular gqrx software is used as a general radio spectrum browser, allowing demodulation of various analogue and digital signals as well as viewing them on the spectrum waterfall. For tracking radio stations, the unit uses the Xastir software, which is able to take positional information from various sources and plot them on a map for situational awareness, as well as plot the GPS position of the unit. The Xastir software is primarily used for APRS information display from radio amateur stations. Since all these applications run in the same VNC session, I can connect to it from my phone and use it in the same way as I would use the front panel, which is handy when using the unit in headless mode, as mentioned earlier. Other basic functions that are directly accessible from the GUI are the backlight level, volume level, on-screen keyboard, and rapid unit shutdown. A series of coloured indicators show the status of the various subsystems. To reduce the system power consumption, I found a smart way to power off the Raspberry Pi's USB hub using a piece of software called uhubctl, which reduces the DC power consumption by around 100mA. The unit has an autonomy of between 7 and 18 hours with the Anker PowerCore 20100 power bank.

THE FIRST SIGNALS IN THE FIELD

I took the unit out on a field trip to test the various functions. The gqrx software reads the baseband I/Q stream from the dual RTL-SDR receivers and performs a fast Fourier transform (FFT) which it plots on the screen in function of time in the form of an RF waterfall. I attached a longwire high-frequency (HF) antenna to test the short-wave reception. Tuning to the amateur radio bands, we can see various signals ranging from Morse code to voice transmissions and listen to them. Far away from human-made interference, the unit shows very clean spectrograms with crisp and clear reception. Since there are two radio receivers in the cyberdeck, I can monitor two entirely different bands at the same time.

Above ↑
First signals on short-wave radio bands

Above Left ↖
Field operations

HackSpace magazine meets...

Tom Ranson

Bringing dinosaurs back to life, digitally

Tom Ranson is a 3D visualisation specialist at the Natural History Museum in London. This means that it is his job to scan anything that is or was part of the natural world, from dinosaurs to herring ear bones, and many, many things in between. Most of these scans are for researchers, but that doesn't mean they're locked away in digital ivory towers. Tom shares some of the Natural History Museum's digital treasure-trove on Sketchfab at **sketchfab.com/NHM_Imaging**. There, you can find scans of some of the museum's most iconic animals, including Dippy the diplodocus and Hope the whale.

We got in touch with Tom to find out what goes on in the vaults under the Natural History Museum, and to learn what it takes to digitally preserve some of the rarest natural specimens around.

Above
A Diplodocus skull doesn't fit easily on most 3D printers, but that doesn't mean you can't print it

HackSpace Thanks for talking with us, Tom. Can you tell us a bit about your setup?

Tom Ranson I've got three different scanners in my lab.

We've got two scanners from a company called Creaform: the Go!SCAN 20, and the Go!SCAN Spark is a brand-new one that just turned up this month. They are basically just flashing a pattern at an object and then reading that pattern something like 30 times a second, so the deformation in the pattern gives the machine the information it needs to know about what's going on on the surface. That plugs into a laptop, and the scanner plugs into a battery pack I can wear. So as long as I can get to every part of an object, I can just carry on walking. I could just scan a massive object just by walking around it. We've used these objects to scan, for example, Hope the whale, the massive blue whale skeleton that's in the Hintze Hall.

The next thing up accuracy-wise, is a laser scanner, which is attached to a granite table. It's technically mobile, because the table has wheels, but it's a bit difficult to move around. That's accurate down to .01 of a millimetre, so 10 microns. So, I scanned, for example, an Iguanodon femur which is probably just over a metre in length. And I can image that whole thing in about 15 minutes. And then the next thing down, which can scan even smaller objects with higher accuracy, is something called the Alicona Infinite Focus microscope, which is basically just a microscope with a bunch of lenses. And it takes a whole ton of images, like 2–300 images in the Z-axis, the up and down, and then stitches together, gets rid of everything that's out of focus, keeps everything that is in focus, and then it can stitch that as a 3D object. I used that microscope to scan a herring otolith, which is the ear bone. That's about 3 mm in length, and I scanned it and then printed a 15-centimetre copy.

So, I can do anything, from tiny to a massive whale. And that just encompasses everything that's in my lab.

HS That's epic! What happens next? Presumably you don't generate 3D files just for the sake of it?

TR From there, we can either send the data to researchers, or we use it in-house for our own research. My room is sort of split in two: one half is my big staging area where I teach and I scan. And then the other half is just where the 3D printers are. So I've got a Formlabs Form 3+. I've got two Prusa MK3S. I've got an Elegoo, a £200 SLA printer, and an UltiMaker.

I've got a stone slab here, with some fossilised footprints in it. The idea was that I scanned it and printed it, and then it goes to production, who can make a mount for it without ever having to touch the actual specimen. So it saves it from coming in and out of [storage] all the time. And, of course, if they drop this, I just print another one.

HS So, rather than having this priceless, 500 million-year-old piece of fossil on someone's workbench, they can make a copy and make their mistakes with that, until they're ready to use the real thing?

> "I'm trying to make the world that little bit smaller and make science a little bit bigger"

TR Exactly. Amongst our specimens in the collections, we've got stuff, for example, Charles Darwin's personal fossil collection, which is insured for millions and millions of pounds.

If we can reduce the number of times that objects like that have to come out of collections, then we can cut down on risk and do more with what we have.

HS Wow. Obviously, you don't need .01 of a millimetre when you're scanning a blue whale, right?

TR No, it completely depends on what the data is being used for. A lot of stuff that comes through my lab are teeth, to study the micro-wear on them. And it's to try and work out whatever this animal might have eaten; you can analyse the wear on the tooth, and then match that up to likely wear from greenery and whatnot, and take a guess as to what they ate.

I only started working here in August 2022, and one of the things that I quickly realised is that this is the Natural History Museum, and it's all of natural history. It's not just a bunch of dinosaurs: we've got everything. We've got an entire tower that's just skulls and skins, of every species that has existed.

We've got a lot of things that we call holotypes, which is the reference object that people would come to us for. And you go in there – it's just floor after floor of floor-to-ceiling cabinets, like rows and rows and rows, and you pull one open. And there's just the skin of the animal and a skull of the animal, for cabinet after cabinet.

We've got the holotypes for a lot of Australian species. One of the cool things that I get to do with that is that a lot of Australian researchers will get in touch with me and say 'I need to analyse this particular taxonomy, can you help?' and then I scan it to .01 of a millimetre, make the data into a lovely high-definition mesh, and then send it to them. I can do that all within a day, and it saves their 18-hour flight from Australia to come and visit our collections. I'm trying to make the world that little bit smaller and make science a little bit bigger. →

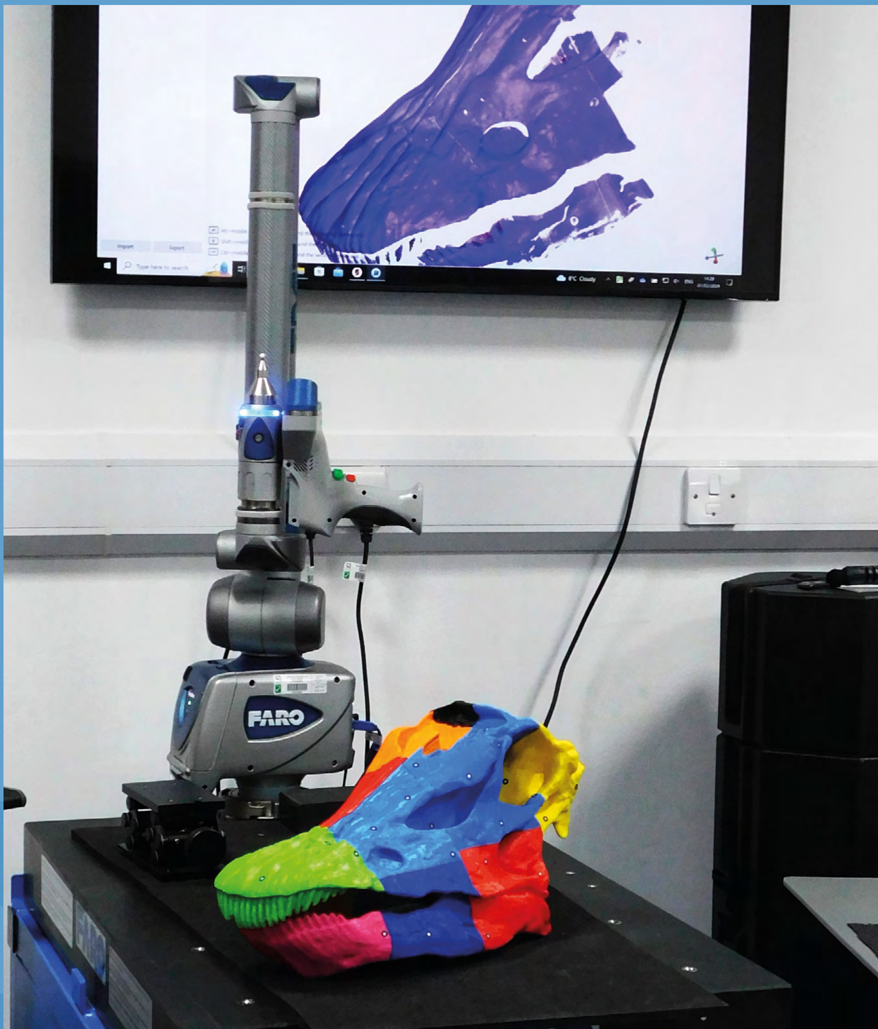

Below
When looking at the scan, Tom realised that there's a tiny Abraham Lincoln sitting on a USA one cent coin

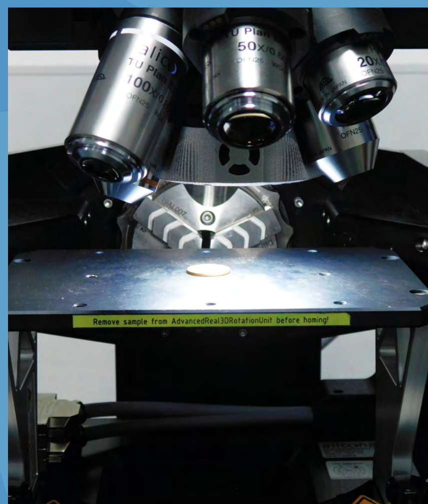

Above
The 3D model is generated as you scan, so you can see if there are any bits missing

Right
The blue laser line is scanning the 3D shape of this printed dinosaur

Above
Unlike the laser scanner, this scanner can create full-colour models of whatever it's scanning

HS That's awesome. That reminds me of the coronavirus, when the researchers working on the first mRNA vaccines hadn't actually looked at the virus itself – they just got a copy of its gene sequence in an email.

Are you trying to scan everything in the collection?

TR I would love to scan everything in the collection, but we've got hundreds of millions of specimens. I don't think it'll happen in my lifetime! I would love to do that sort of thing.

What I want to try and do is let everybody know about the science that we're doing. Millions of people come through these doors every year. And it's definitely not something that I realised, but the footprint of the museum, there is that [amount of space] again in the basement full of labs.

We do so much science here that I just never realised. I'd love to be able to get out to the public with the stuff that I'm doing. There are a couple of things that we've got coming up. For example, some of our researchers want some of the data of the Triceratops and the T. rex skulls that we have out on display. And the conservation team were going backwards and forwards for like a month, trying to solve problems like, how do we get them down? How do we get the data off of them, they're very fragile. If we move them, they might break. And then, I randomly got looped into the email chain sort of late on and I just said, 'Give me a cherry picker. And I'll go up there with my laser scanner. And I'll just scan it.' And then they're like, 'Ah, but it'll have to happen before we open because we open at 10am every day.'

I could feasibly have done it before we opened, but I wanted to rope the area off and have some people from the public side of things, talking to the public about what I'm doing. I'll answer questions, and we'll do it all out in front of the public. Because I think if someone watched me do surface scanning, I think that in itself is a really interesting process. It just looks cool.

HS You mentioned the Creaform scanner. How different is that to, say, the scanning capability you might get out of an iPhone? I don't want people to think, 'Oh this is very cool, but you need a massive budget to do it.'

TR The difference is the speed and the accuracy of the tiny area that it's scanning. So the Creaform hand scanner is a £30,000 scanner. And it's accurate to .1 of a millimetre. The FaroArm is the one that goes a bit further. And it's designed specifically to track small objects.

The wonderful thing about all of the hardware that I have is that I'm not using them in the way that the designers have intended at all. So, when we ask for help with scanning specific things, their engineers have no idea how to help us. So we have a lot of back and forth between them to help optimise their streams.

> "I would love to scan everything in the collection, but we've got hundreds of millions of specimens"

So the Creaform scanner uses these little tiny reflective dots that you put a bunch of them around the object and then the scanner is watching those and tracking where it is in 3D space, while it's being moved around. It's running at 30 or so frames per second, and taking a photo and looking at what it can see as a model and seeing what dots it can see. So, then it can use what dots it can see to understand where on the model that is, and then use the colour of the model and stitch all of those together. The iPhones have lidar scanning sensors in their camera setups now anyway. So you could do a pretty decent room-scale scan with an iPhone in a standard room. [With] the accuracy that an iPhone can afford you, the room would look quite nice. But when you start to try and scan, for example, something down to like the sort of tiny sizes, the Creaform is tuned to be able to pick something up like that. Whereas the iPhone starts to lose the detail of what you've got.

The FaroArm is a £90,000 scanner, that's like another step up. The advantage of a scanner like that is that the laser is attached to an arm with a whole bunch of encoders in it. So the software has a couple of streams of information: the software knows exactly where the scanner is in 3D space, because the scanner is bolted to the table. And then you've got all of its encoders telling you where the scanner is. So then it doesn't have to worry about positioning any of the images that it takes, because it already knows where it's looking. All it needs to worry about is a blue laser line, and there are about 2000 points across the line that it measures. →

Above
This portable scanner can be taken to scan objects wherever they are in the museum

And it's measuring those at about 25 times a second. So, you're getting hundreds of thousands of points coming into the software. It's accurate to .01 of a millimetre.

I always have little handling models in my lab. When you look at the collections we have and the hardware we have in the labs, that combination makes us probably the best in the world at what we do. Because you can come to our labs and get access to almost any species that you want, through all of mankind's history, and get an incredibly detailed look at the surface, the internal structure.

HS What are your top three models in the museum's collection for 3D printing?

TR There are a few things that I've scanned that are coming to light that haven't gone on to the Sketchfab yet that I really like. One of the ones that I've just scanned, simply because it's such a vast slab – because fossil is rock, so it's huge and it's just immovable. And the surface is actually fossilised dinosaur skin. We think that it was a very keratin-rich animal to have its skin fossilised like that, but it's the most well-preserved bit of dinosaur skin that has ever been uncovered. That will be going on to the Sketchfab at some point, and it's one of my favourites. The model is kind of nice; it's very tactile to run your finger across. The idea that it's dinosaur skin blows my mind. But I would say my absolute favourite, because it's my favourite animal, and it's one of the most complete specimens that exist in the world, is Sophie the Stegosaurus. We've got all of the data for Sophie. She's the most complete Stegosaurus fossil skeleton in the world. And that is very satisfying to print out. It's a nightmare to do it successfully, and you have to do it in tiny pieces.

HS At the moment, virtual reality, sorry, I mean spatial computing, is very much in the tech news following the release of the Apple Vision Pro. Do you expect that to change the work you're doing?

TR It's another tool that just sort of augments what we can do. It's something that I'm really passionate about. Before I worked here, I worked at the University of Suffolk and I built a Virtual University for the university staff. We had DK1 Oculus headsets, so that was right at the beginning.

> "You can come to our labs and get access to almost any species that you want, through all of mankind's history"

This data that we're scanning in the lab can be used with VR, because at the end of the day, I'm making an STL, or an OBJ file.

Before I did any of this, my actual degree was in digital film production – animation and suchlike. So I'm very familiar with building those sort of environments that you can plug into – you'll notice the HTC VIVE headset in the corner of the room. So it's something that we're looking towards because it's really easy. I've got terabytes of surface data of models and skeletons and stuff and I could just drop it into Unity.

Within 20 minutes, you could make a VR experience of picking up specimens. It would be incredible to take all the objects, 18 million specimens or something on site here, to be able to take a fraction of the locked-away specimens out into the world virtually, and give people the opportunity to play with without risking the specimen.

HS Is there much difference in the way you scan an object if it's for VR vs if it's for 3D printing vs if it's for research?

TR Not in the way that we approach the scanning. The only difference is the degree that I would go to with the post-processing of the model. With the scanning that's happening here, because the accuracy is all the way down to .01 of a millimetre at its highest resolution, Joe Public would never need that out of a mesh. And to be honest, the polygon count of a mesh that I produce would immediately crash any game if you tried to play it.

So, if it were for a researcher, they would want measurements down to fractions of a millimetre. If it were to go into a game engine, then I would have a polygon limit that I'd work to or the complexity of the geometry.

HS You release some of your models on Sketchfab. How do you decide what to release?

TR Up to this point, it's been kind of ad hoc. What I would love to do is to have a point where I'm obviously doing custom scanning jobs for people and requests that come in, but also just have a constant working with the curators for interesting models just coming through the lab. I can digitise them, and then they go off on Sketchfab. Because, unless it's owned by somebody who specifically doesn't want the data to get out or department-sensitive research, there's no reason. I have a vision that, in the future, there will be just thousands and thousands of specimens – just stick a headset on and go and play around. □

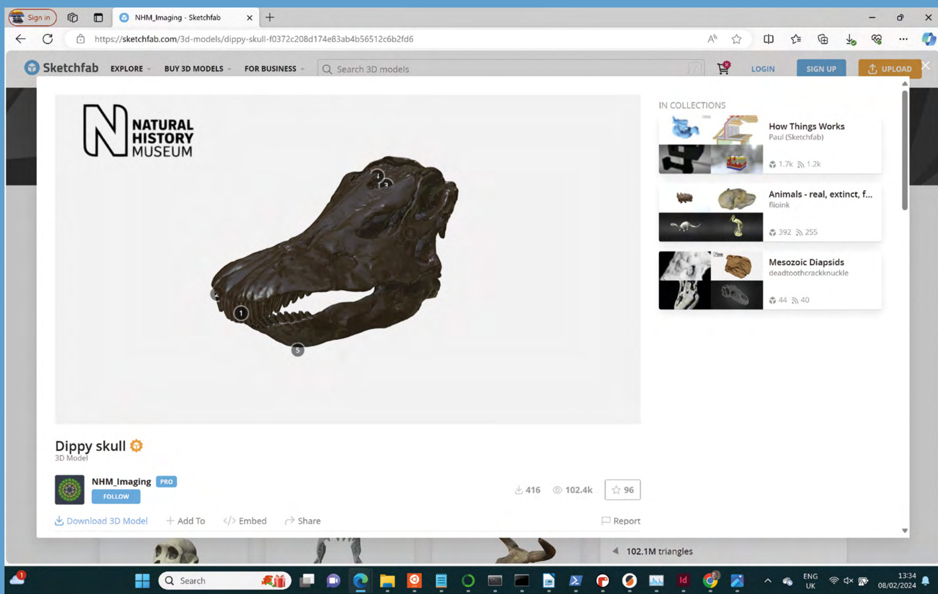

Left
Download your own virtual dinosaur from sketchfab.com/NHM_Imaging

Below
This mammoth tooth was scanned to identify wear on the teeth so researchers can work out what it ate

Above
This snake skull was CT scanned while still inside the snake. This way, researchers can handle the skull without damaging the specimen

Objet 3d'art

3D-printed artwork to bring more beauty into your life

There's nothing new about the concept of print-in-place. You print a mechanism in one go so that it moves in the way you want it to, straight off the build plate. These boxes have print-in-place latches for holding them shut, but there's also another trick – the red band around the top is a flexible filament that provides a seal when it's shut. These are joined to the PETG (for the box with orange latches) and PLA (for the box with white latches). Both of these boxes are exactly as they came off the print bed, with no post-processing or assembly. They're not water-tight, but it's a lot more sealed than most 3D-printed boxes that we have made.

We printed this on the Prusa XL. Turn to page 90 for details of our in-depth review of this new printer.

hsmag.cc/pipbox

Letters

ATTENTION ALL MAKERS!

If you have something you'd like to get off your chest (or even throw a word of praise in our direction), let us know at **hsmag.cc/hello**

AUTOMATIC MAKING

I'm just getting started on my journey as a maker and I can't decide what my first 'big' machine should be – a laser cutter or a 3D printer. I know 3D printers are much more common, but I've seen some cheap laser cutters online, and they look like they can do some impressive work. Any advice?

Anuj
Birmingham

Ben says: Laser cutters are much more serious than 3D printers. We're pretty concerned about the huge number of unenclosed laser cutters that are on the market at the moment. The laser beams that come off these can blind instantly. Yes, they usually come with goggles, but these typically leave gaps around the eye. Not only that, but anyone passing by a window is at risk. And this is before we consider the risk of fumes and smoke coming off them. It's not impossible to add these features yourself if you have the skills and experience, but it's not something you should do lightly.

A laser cutter might be cheap, but a safe laser cutter is going to cost a bit more and needs some thought into its set up (particularly with ventilation).

As to what you should get – they are very different machines with different pay-offs, so it's hard to say without more information on what you want to do. 3D printers can be slow, but they can do an awful lot. Laser cutters are much more limited, but they're much faster. There are also far more downloadable designs for 3D printers.

Our general advice, though, would be that if you can do what you want to with a 3D printer rather than a laser cutter, then a 3D printer is probably the best choice. They're safer, easier, and have a bigger community behind them.

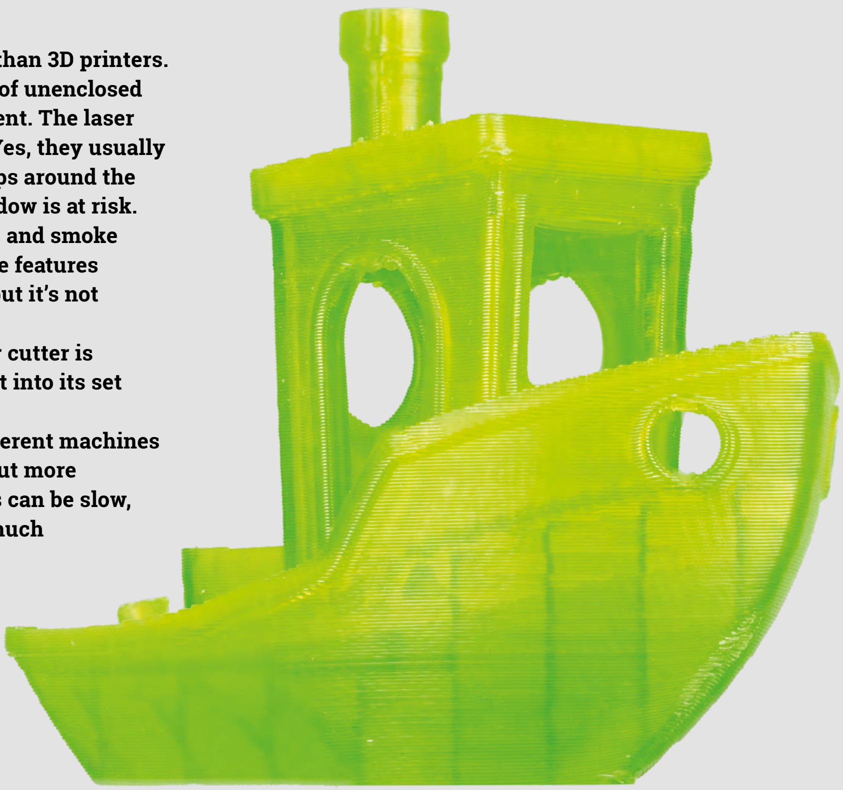

TURN UP THE VOLUME

When I was a lowly undergrad in computer science, we learned about volumetric displays. They seemed like the future, and I was pretty sure that in a few years we'd see them everywhere. Fast-forward 20 years, and I haven't seen one until flipping through the top projects section in issue 75. They still look like the future, while also looking a bit like the past.

Mary
Warwick

Ben says: Retro-futurism is the style of art – or perhaps design – of things in the past that were meant to look like an imagined future that never came. Think of space-inspired architecture, or perhaps *The Jetsons*. Here we are in 2024 – our cars don't fly, we eat food not pills, and our houses are still mostly made of bricks and mortar.

I can't tell you if volumetric displays are an idea whose time is about to come or an unsuccessful branch on the technology family tree. I can tell you that they look awesome and fully deserve their moment in the spotlight. Perhaps if we all spend the next few weeks searching for volumetric displays on our favourite search engine, some tech executive will notice the trend and commission some awesome-looking hardware.

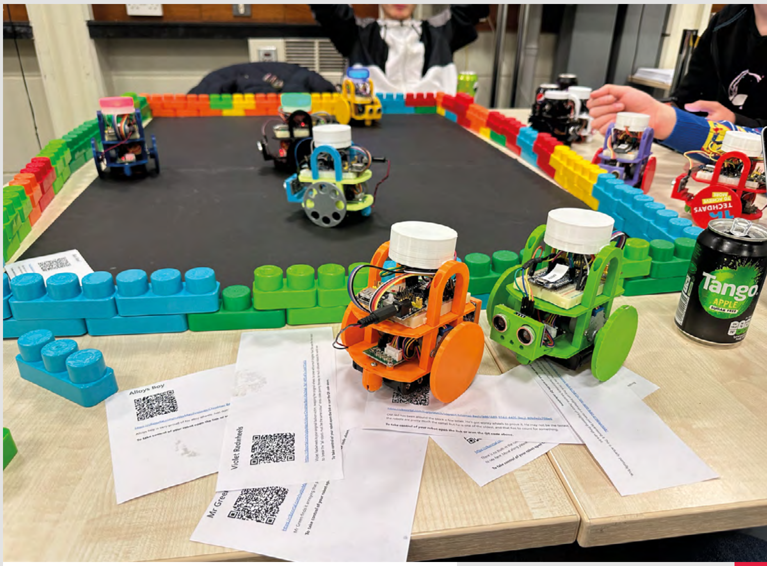

ELECTRONIC SPORT

Rob Miles talks of robot rugby. I really hope this is going to include Ellis Genge facing down a Boston Dynamics-inspired cyberman. Come to think of it, a *Doctor Who* / Six Nations crossover is probably what we need to cheer us up this spring.

Dave
Bristol

Ben says: One thing I really like about rugby is that to make a great team, you need a lot of different players. You need different skill sets and different body shapes. A team of huge, strong people can be given the runaround by a team that also has some faster (but less strong) players. But likewise, a team of fast players crumple beneath a team that also includes some strong players. In order to be a good rugby player, you have to learn to play to your strengths, whatever they are, and play alongside people who have different strengths. I think this makes it an interesting discipline for engineering.

HackSpace
TECHNOLOGY IN YOUR HANDS

FORGE

HACK | MAKE | BUILD | CREATE

Improve your skills, learn something new, or just have fun tinkering – we hope you enjoy these hand-picked projects

Leather craft – an entry-level exploration

If you have a resolution to try something new, why not consider experimenting with leather craft?

Nicola King

Nicola King is a freelance writer and sub-editor. She's looking forward to the new year, new possibilities, new opportunities… and plenty of making, creating, and hacking!

You won't be at all surprised to learn that leather-work is an ancient craft, possibly one of the oldest known to humankind. We talk a little later about its extraordinary history, but the purpose of this tutorial is to show how easy it is to take some tentative steps into the world of leather handling and manipulation. It's one of those skills that is reasonably inexpensive to experiment with too – it has its own set of tools and techniques, of course, but once these are purchased and mastered, leather craft is an engaging and unique kind of pastime.

We'll take you through some very basic leather crafting, using a reasonably thin and, therefore, easy-to-work-with leather, and we'll be making a bookmark and simple key fob. The beauty of this project is that we will be making useful accessories using leather scraps, offcuts from other people's leather projects/production if you will, which is an easy, cheap, and arguably eco-friendly way of sourcing these crafting materials. We procured a box of leather scraps for an unbelievably (for the huge amount that arrived) cheap price of £10.99 online (e.g. **hsmag.cc/LeatherScraps**), but also be aware that you can come by similar remnants and snippets from local leather goods stores, upholstery shops, or even shoe factories or handbag and belt manufacturers. We really don't want or need to be buying large pieces of expensive leather to get underway, and it's much better for leather scraps to be recycled into something useful than to end up in landfill, so do some local research and see what may be available in your area.

A (VERY) CONDENSED HISTORY OF LEATHER-WORK

The timeless tradition of leather-work will arguably be just as popular in centuries to come as it has been over thousands of years already gone by. In its earliest form, the basic needs for protection from the elements and clothing drove our ancestors to think about how animal skins and hides could be best utilised. From Egypt, to Greece to Rome, where Romans utilised leather to make sails for their boats, leather was employed for multiple purposes as people recognised just how versatile it could be. The oldest leather shoe that's been discovered thus far is reputedly 5500 years old, and was unearthed in an Armenian cave (**hsmag.cc/OldestLeatherShoe**). There's also evidence that our Stone Age forebears used bone scrapers to make leather and pelts. It's safe to say, therefore… this is a very old craft! During the Bronze Age and the Iron Age, people continued to look for new uses for leather, and jewellery and protective shields started to include this natural textile.

Through the centuries, leather started appearing in bags, belts, harnesses, saddles, and so on. In 17th-century Italy, it was used as wallpaper! Aside from just tanning leather using fats and oils and tree bark, to name but a handful of the early tanning options, craftspeople started to paint, dye, mould, carve, and stamp leather, as the material's uses evolved. During the Enlightenment, furniture, hats, and fancier shoes emerged, while the Industrial Revolution, and the consequent forms of machinery that advanced with that, transformed the leather production process in leaps and bounds. Specifically, chromium tanning was introduced in the mid-19th century, and this used a tanning agent called chromium sulphate, as opposed to the older vegetable tanning processes. This new, much cheaper, take on tanning vastly sped up the process and made the finished leather much more flexible and supple.

QUICK TIP

Be aware that too much exposure to direct sunlight can affect natural leather, potentially damaging it, as can extremes in temperature. Just try to keep your leathers at a constant, stable temperature and you're good to go.

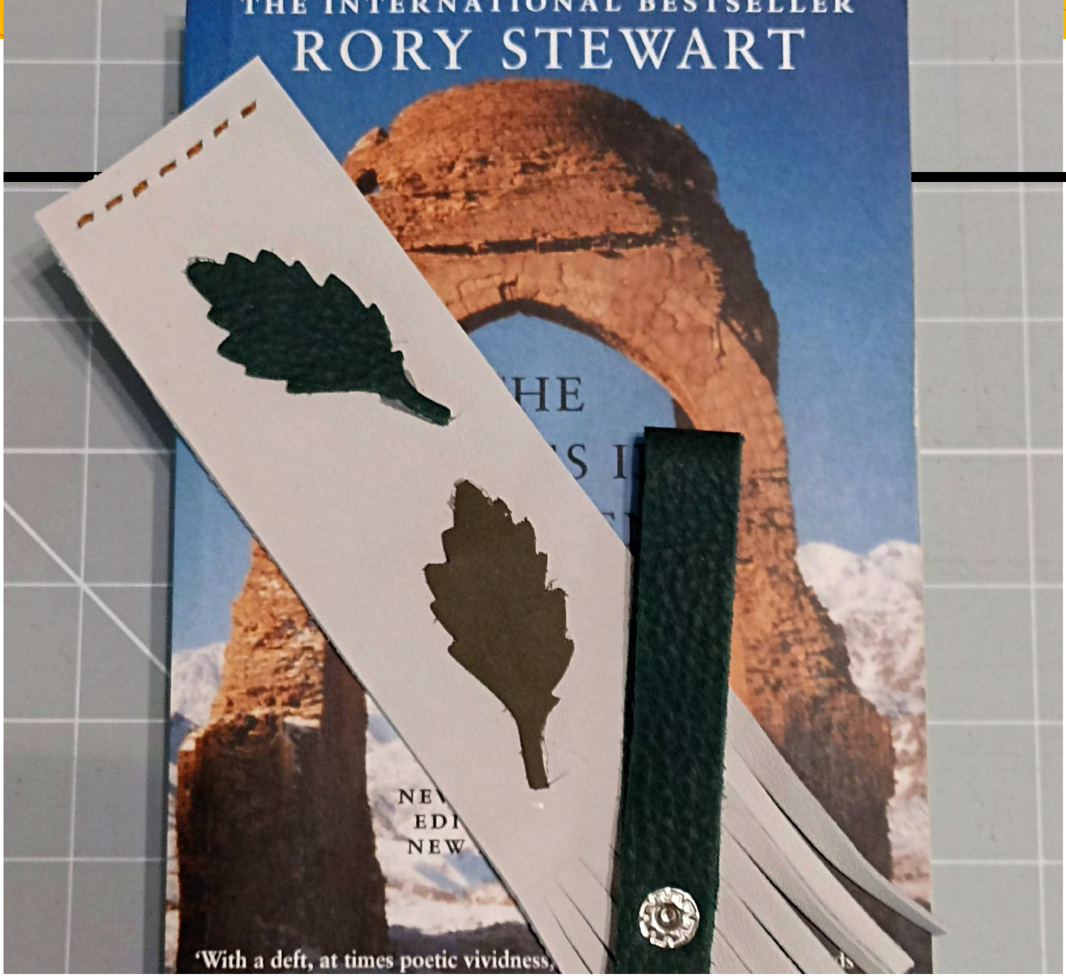

PROJECT ONE: DON'T FOB ME OFF

This key fob is an uncomplicated and almost effortless project, and would make a great gift. The thicker the piece of leather that you use, obviously the sturdier the key fob will be, but it really depends on the types of scraps you have been able to source. Our leather piece was thinner and, therefore, easy to cut with a rotary cutter. First of all, decide how long you want the fob to be. We cut a piece of leather 22 cm in length by 1.5 cm wide on our cutting mat, as we'll be folding the piece in half eventually, and the width needs to fit neatly through our keyring. Use your steel ruler to measure and cut against, being careful to cut away from yourself at all times for safety reasons. If your mat has a grid, use that to help you keep your lines straight.

Once you have your piece of leather, you can neaten the sides using some sandpaper, and just keep sanding until the edge is smooth. Once you become a little more advanced, you can use something called an edge shave / edger to gently cut shave away the edge to help round it off. You can also purchase some tools to gently burnish the edges, which adds a very professional look on thicker pieces of leather.

This helpful video gives more insight into the technique: **hsmag.cc/EdgeBurnishing**.

When you're ready, fold your piece in half, turn over one end so that you can place the keyring on, and then you can glue the ends on top of each other. Hold in place while the glue sets, and you're ready to attach some keys. Alternatively, make some holes and use a rivet to tightly secure the pieces together. →

Above
Our finished bookmark and key fob – definitely only a dip of the toe in the world of leather craft, but an enjoyable one

Below
We used a rivet to fasten the end of the key fob, and it worked a treat. For a first effort, we were pretty pleased

YOU'LL NEED

- A self-healing cutting mat
- Scrap pieces of leather
- A craft knife / rotary cutter
- Awl, leather needle, and waxed thread (hsmag.cc/LeatherKit)
- A steel ruler
- Cutting templates (optional)
- Scissors
- Strong glue e.g. a superglue, or one specified for leather
- Rivets / hole punch / hammer (optional)
- Keyring
- Sandpaper
- Edger (optional)
- Leather punch pricking tool (optional)
- Lacing/stitching punch and hammer (optional)
- Die-cuts, as a trimming (optional)

TYPES OF LEATHER YOU CAN WORK WITH

There is a wide spectrum of leathers to choose from, so let's examine some options:

- Genuine leather (from an animal hide) can come in various different 'grains', and this refers to the composition and appearance of the hide, the layers of hide included in a piece of leather. Generally, full-grain leather from the top layer of the hide is the best-quality leather. Depending on the quality of the leather, items made from real leather can last a lifetime. Proper care and maintenance can help to keep the attributes of leather intact, so conditioning it with an oil or conditioner wipe when it becomes dry, and keeping it clean and away from direct sunlight and salt water are important for it to stand the tests of time.
- Faux, or plastic, leather (sometimes called pleather) is a much cheaper alternative to genuine leather, and perhaps a more acceptable material for use by anyone who does not want to utilise natural leather, for maybe their own personal or ethical reasons. Faux leathers can look incredibly like the real thing, and come in a very convenient wide range of colours, which is also a plus point. This author has also found that it goes through her die-cutting machine pretty easily with her dies that use steel rule technology (such as Sizzix 'Bigz' dies), and is very easy to cut with just a regular pair of scissors. The major downsides are that PVC, which is usually what this is made from, is not as tough or durable as genuine leather, so will peel and crack much faster, it's not breathable, and it is also not eco-sustainable as it is made with plastic materials derived from the petroleum industry.
- In addition, you could also consider vegan leather, which tends to be made from cork, leaves, or other organic compound materials. Like faux leather, it has poor durability and workability, can have high production costs, and its feel and quality can fall short. However, it comes down to personal choice at the end of the day, so follow your own convictions and instincts.

Below Left
These pronged stitching punches are a useful toolbox addition. The choice of which to use will depend on the look you are trying to achieve but, for information, these have a spacing between each prong of 3 mm

Below Right
A colourful selection of leather scraps of various thicknesses, many of which are a really good size for creating small accessories with. Inexpensive, they are also great to practise your leather craft techniques on

PROJECT TWO: BOOKMARK THIS ONE

The dimensions that we chose for this piece were 5 cm by 20 cm. Cut your piece on your mat using your knife or rotary cutter, again using your steel ruler to cut against. You may find, when cutting leather, that it takes a few passes to cut through it depending on its thickness. This is fine, just take even strokes and take your time to cut out your shape, as we want a nice clean edge if possible. Once cut, you can sand the edges if you wish or burnish them.

Next, let's think about adding some decoration. One really useful tool for budding leather crafters is an edge stitching groove tool, which you can use to create a groove in your leather in which you can then make stitching holes with your awl/punching tool, or you can just use it to make grooves along edges simply as a decorative feature on your piece (**hsmag.cc/GrooveTool**). We're just actually going to make some holes in our bookmark with a stitching punch tool, and then we'll add a few stitches. If you are going to do the same, make sure you have a thick mat under your work so that your surface doesn't get damaged. Choose your punch, hold it where you want the holes to be (we chose to place a line at the top which we drew with a pencil first), and then tap the top with a hammer – one swift tap should be all you need to pierce the leather, but thicker leather may take more work. We then sewed through those holes for decorative purposes with our durable thick waxed thread, which is easy to sew with as the wax causes less friction as you sew, so it glides through the fabric very easily. Next, we die-cut some leather pieces and hot-glued them as an embellishment, but you could sew those too. Finally, we also decided to add a tasselled bottom to the bookmark, so just made some 5 cm length cuts to the end.

So, treat yourself to some basic tools and some leather scraps and see where it takes you.

QUICK TIP

Edge marking pens are a great way to apply an edge colouring or dye to the sides of your leather project to give it that finished, pro leather-worker look!

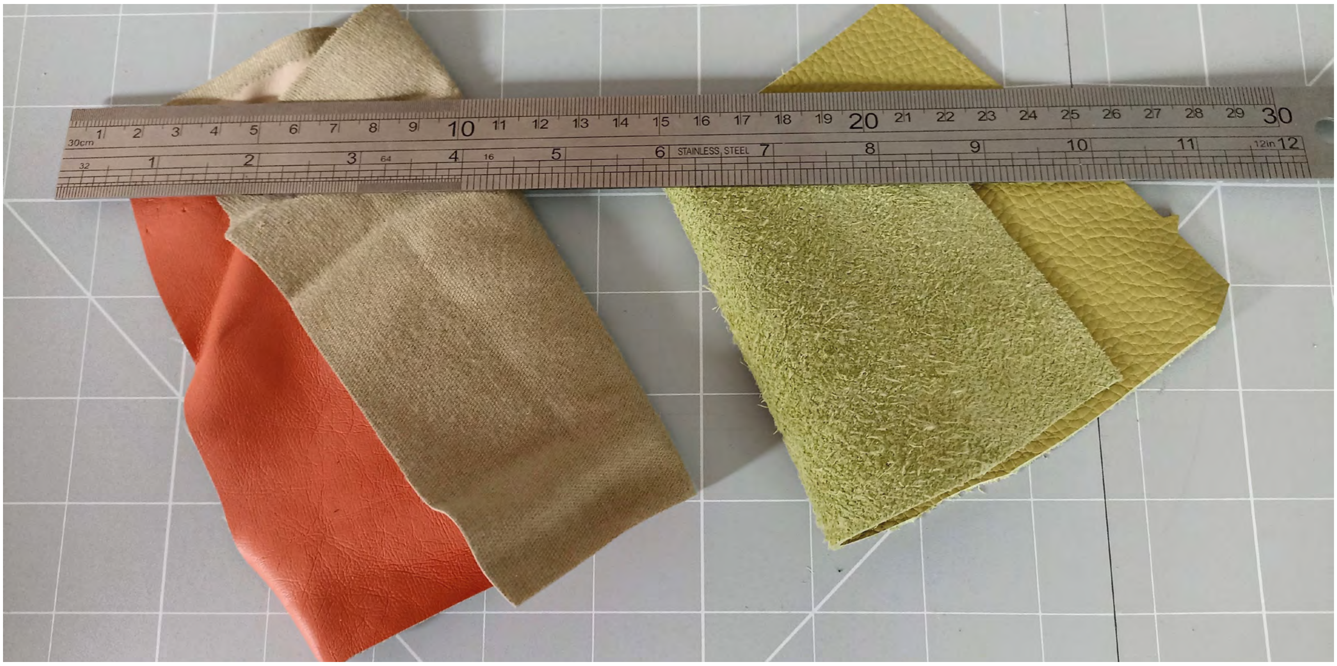

Above
Faux leather on the left and real on the right – the texture of real leather is coarser and not uniform, and it also has a much more distinct smell. If you press real leather with your finger, it will stretch like skin, unlike the faux alternative

You can even sew scraps together to make a larger piece that you could make a bag from, for example. You can hand-sew the pieces, or use a more heavy-duty sewing machine to attach them together, if the leather is soft enough and not too thick.

Making such accessories is a great way of recycling leather offcuts and, if you don't want to use real leather for whatever reason, faux leather is always an option. You can also investigate embellishing your makes once you become a little more advanced – stamping is one route. Or, you could paint your leather or apply some leather dye to it before you cut your shape out. Go ahead and make some accessories that can last a lifetime.

WHAT TO MAKE WITH LEATHER SCRAPS

Let's be honest, we have to paddle in the shallow end when we first start a new craft. There is absolutely no point in trying anything too challenging, as failure is a sure-fire way to lead to you throwing your attempted makes out of the window, while vowing never to bother trying to create anything again. So, let's consider some simple starter projects and designs which can do nothing but further inspire you to continue your leisurely leather activity. Plus, they use small scraps of leather, thus using up your precious resource in the most effectual way possible:

- Simple jewellery – for example, rustic bracelets, necklaces, and earrings
- Luggage tags (it's useful to know that you can purchase templates for all of these kind of makes, such as **hsmag.cc/LeatherTemplates** – just make a search online)
- Credit/debit/business card wallet, or small purse
- Phone case
- Double-layer leather coasters (stitching two layers together)
- Glasses case
- Scissors case

Left
A small selection of must-have leather craft tools, including waxed thread and an awl which you can use to make small holes in your fabric. Inexpensive leather sewing kits can be found online (hsmag.cc/LeatherSewing)

QUICK TIP

Join a leather-work community forum, or find a local maker space where others are having a go at leather-work. Speaking to other keen makers for useful tips is always worthwhile.

Part 02

Raspberry Pi: Mecanum robot with sensors

Expanding on our mecanum robot project, add a distance sensor to avoid crashing into obstacles. Learn how collision avoidance is used to make driving safer

MAKER

Stewart Watkiss

Also known as Penguin Tutor. A Maker and YouTuber who loves all things Raspberry Pi. Author of Learn Electronics with Raspberry Pi.

penguintutor.com

@stewartwatkiss

You'll Need

- HC-SR04 Ultrasonic distance sensor **magpi.cc/hcsr04**
- 330 ohm resistor **magpi.cc/330r**
- 470 ohm resistor **magpi.cc/470r**

The mecanum robot in the earlier tutorial (MagPi issue #135, magpi.cc/135) could be controlled with a gamepad. If you didn't react fast enough, then it was quite easy to crash the robot. Adding an ultrasonic distance sensor means that the robot can automatically stop before colliding with a wall or other object.

The sensor used is similar to those found in car parking sensors. The robot senses the distance to objects in front and automatically stops forward movement when an object is too close. With additional code, this could be used as the basis of a robot that could navigate a maze.

01 Ultrasonic distance sensors

There are various different sensors that can be used to detect distance. The one used here is based on ultrasonic distance measurement, and works by emitting an ultrasound signal and then listening for an echo response. The time taken between the signal being transmitted and the response received can be used to work out the distance to an object.

Figure 1 shows the HC-SR04 sensor used for this project. There are different models of the HC-SR04, some of which can work down to 3.3V. This circuit is based around the 5V version for maximum compatibility.

Figure 1. The HC-SR04 module has a transmitter and receiver. It sends out ultrasonic pulses and detects an echo response

02 Sensor signals

In addition to the power supply, the HC-SR04 semsor module has two pins for an input and output signal. The trigger needs to be sent a short 10μS pulse to start the measurement. The transmitter then sends eight sonic bursts and the reciever listens for an echo.

After determining how long it took to receive a response at the receiver, the echo pin is then raised high for a length of time proportional to the distance. This is shown in **Figure 2**.

The output from the echo pin is 5V, which is too high for the GPIO pins. Therefore a resistor voltage divider can be used to drop the signal to approximately 3.3V.

03 Wiring the distance sensor

The distance sensor needs a 5V power supply. As the robot can have up to 6V, this is connected after the diode used to connect to the 5V pin on your Raspberry Pi.

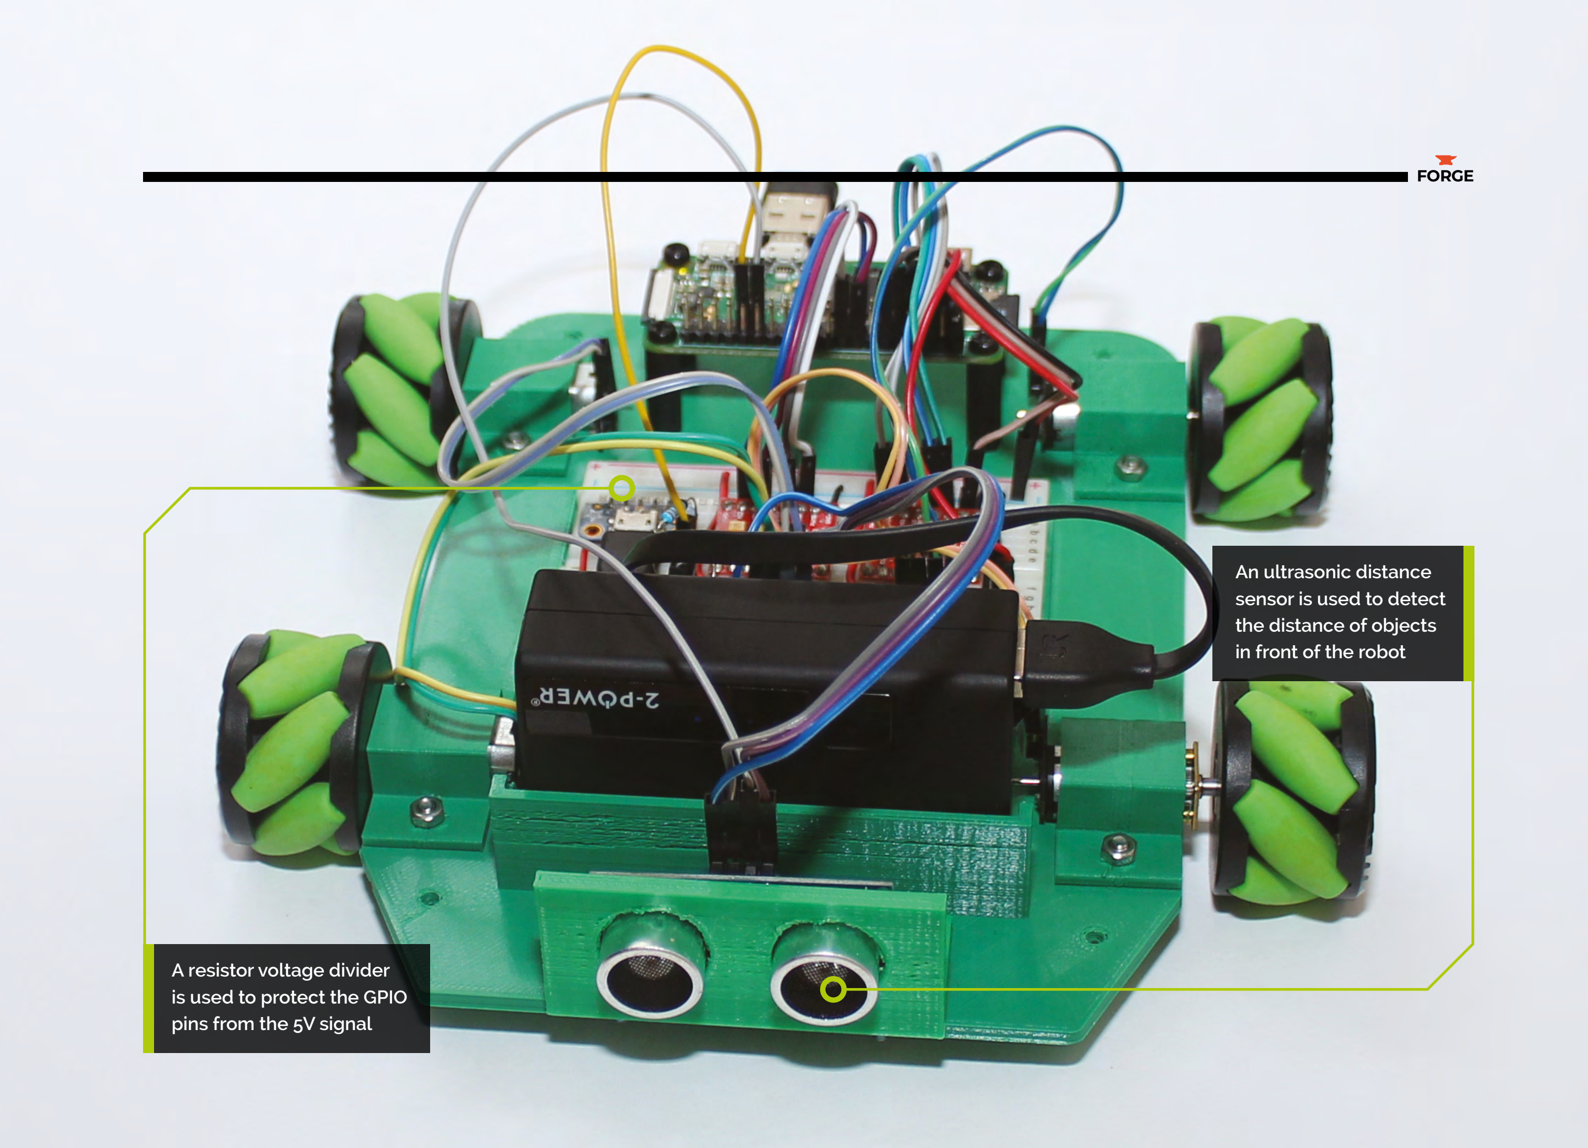

An ultrasonic distance sensor is used to detect the distance of objects in front of the robot

A resistor voltage divider is used to protect the GPIO pins from the 5V signal

The distance sensor needs to be connected to two GPIO ports. The chosen pins are GPIO 5 (physical pin 29) for the trigger, and GPIO 6 (physical pin 31) for the echo.

The trigger pin works with a 3.3V input signal so is connected directly to the GPIO pin. To drop the echo signal from 5V to 3.3V you need two resistors. If using USB for power then there is not much space for these, but they can be squeezed into the gap between the USB connector and the motor drivers as shown in **Figure 3**.

04 Resistor voltage divider

The breadboard diagram shows the resistors with long leads, which are needed to get them to fit into such a tight space. This makes it difficult to see how it is wired, so the schematic diagram of the voltage diagram is also shown in **Figure 4**.

The output is from pin 3 of the HC-SR04, which goes through a voltage divider created by R1 and R2. The voltage is shared across these resistors proportional to their resistance. The voltage across R2 will be:

```
V x (R2 / (R1+R2))
```

… which is approximately 3V. This is the signal that is passed to GPIO 6.

05 Using PiGPIO

GPIO Zero is used to control the motor driver and can also be used for the ultrasonic sensor. Using its default native pin configuration GPIO Zero can communicate with most sensors with acceptable performance. For the ultrasonic

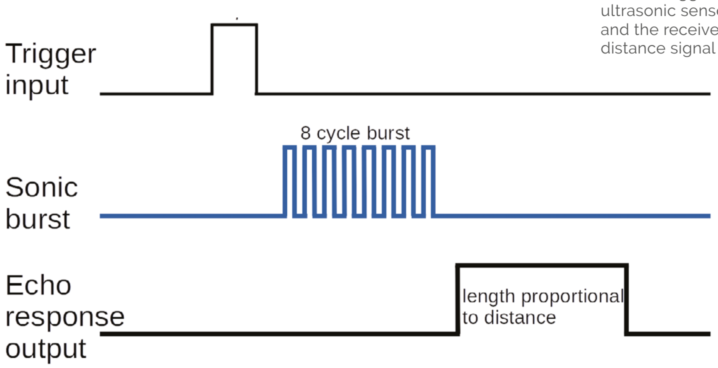

Figure 2. Waveforms showing signals used to trigger the ultrasonic sensor and the received distance signal

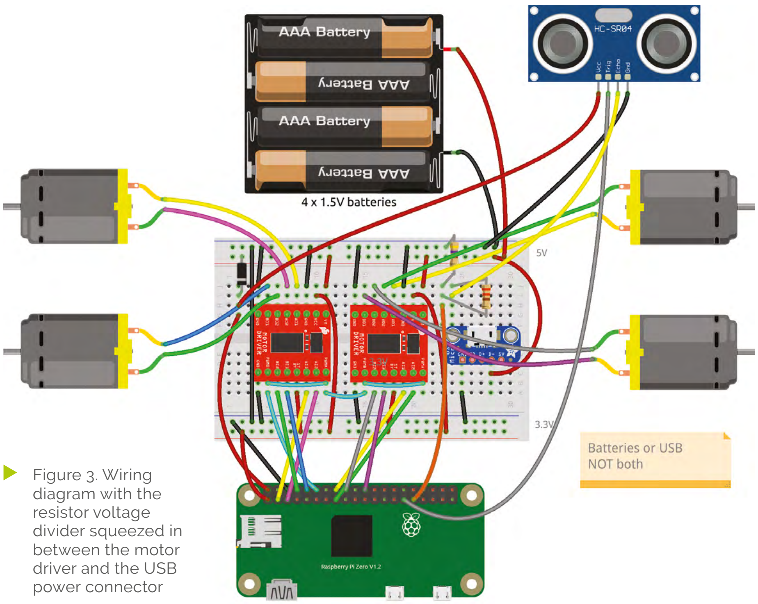

Figure 3. Wiring diagram with the resistor voltage divider squeezed in between the motor driver and the USB power connector

sensor the timing is critical, and using the native library can give incorrect results. Instead, you can use PiGPIO. This allows direct memory access (DMA), which reduces the requirement on the processor. To use PiGPIO the PiGPIO daemon first needs to be run as root, and then the appropriate `pin_factory` needs to be set in your Python program. The daemon is started using

```
sudo pigpiod
```

06 Using Python Pin Factory

With the pigpiod daemon running, then the device factory can be set within the Python code. First, you need to import the PiGPIOFactory library from `gpiozero.pins.pigpio` import `PiGPIOFactory`. Then set the pin_factory using:

```
Device.pin_factory = PiGPIOFactory()
```

This will then use PiGPIO for all devices that use the GPIO pins from GPIO Zero.

07 Measuring the distance

To measure the distance first create a DistanceSensor object:

```
dist_sensor = DistanceSensor(echo=6,
trigger=5)
```

The distance can then be queried by looking at the distance property:

```
dist_sensor.distance
```

The returned value is in meters. The maximum distance set in the code is 0.09, which works out as 9cm.

08 Using Pygame Zero

The earlier version of the robot software was written as a command-line Python program. One issue with that is that the code would stop running whilst waiting for the next key press. This means it is not possible to check for a potential collision whilst the robot is moving.

An alternative is to use Pygame Zero, which can create a GUI application which can respond to keyboard presses as well as regularly checking the value from the distance sensor. It also provides a way to create a graphical screen which can be controlled using a mouse or touchscreen. A screenshot is included in **Figure 5**.

09 How to code for Pygame Zero

The main difference with Pygame Zero is that instead of running the program in sequence the code is event-driven. The main way that code is run is through two functions `draw()` and `update()` which are each called approximately 60 times per second (depending upon computer performance). The `draw()` function is used for displaying the interface on the screen and the `update()` function normally handles any other code that needs to be run.

It is also possible to set code that runs when an event occurs which in this example is the `on_key_`

Figure 4. Schematic diagram showing the wiring between the GPIO pins and the ultrasonic sensor. The two resistors form a voltage divider reducing the voltage

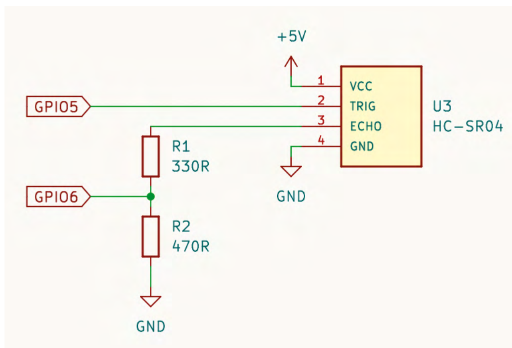

`down(key)` function which is called each time that a key is pressed on the keyboard.

10 The update function

The collision avoidance code is included in the update function. The code detects if the distance is less than the minimum distance. If it is then it also checks that the robot is going forward, otherwise you would end up with a robot that couldn't reverse either.

The code includes error handling by placing checks for the sensor inside a try block. If the code in the try block has an exception, then instead of crashing the program it will ignore the rest of that block and call the except block instead. In this case, if the sensor isn't detected, then it prints out an error message.

11 Running the code

To run the code then it needs to be run using `pgzrun`. This can be done using Pygame Zero mode in the Thonny editor, or by running on the command line using:

```
pgzrun mecanum-pgz.py
```

The program can still be used with AntiMicroX to allow the robot to be controlled using a gamepad.

If you get an error, then make sure you started pigpiod first, as mentioned earlier.

12 Mounting the sensor

Finally, you will need somewhere to mount the sensor. If you are using the 3D-printed chassis then there are mounting holes included for a printable bracket. There are small holes included on the module PCB to allow mounting, but due to their small size compared to the resolution of the 3D printer these were not used, and the board can be held with a small amount of tape or adhesive putty. The bracket is shown in **Figure 6**. You can download the code and 3D designs from **magpi.cc/robotmecanum**.

For a different chassis, it may be possible to drill holes for the sensor or make your own bracket using thick cardboard.

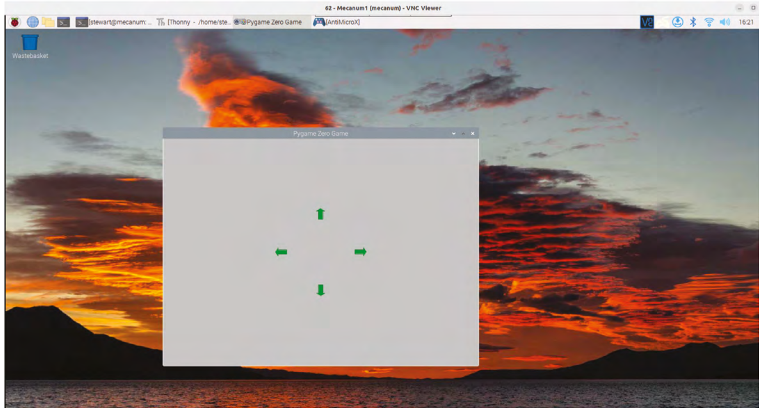

Figure 5. A very basic GUI display has been created using Pygame Zero. This allows the sensor code to run regularly to check for a possible collision

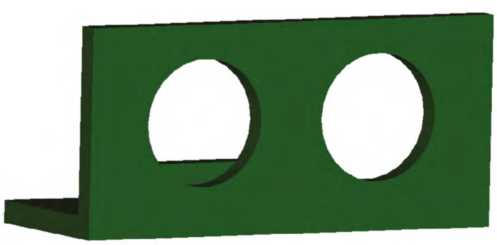

Figure 6. The bracket is designed to screw to the front of the chassis to mount the ultrasonic sensor

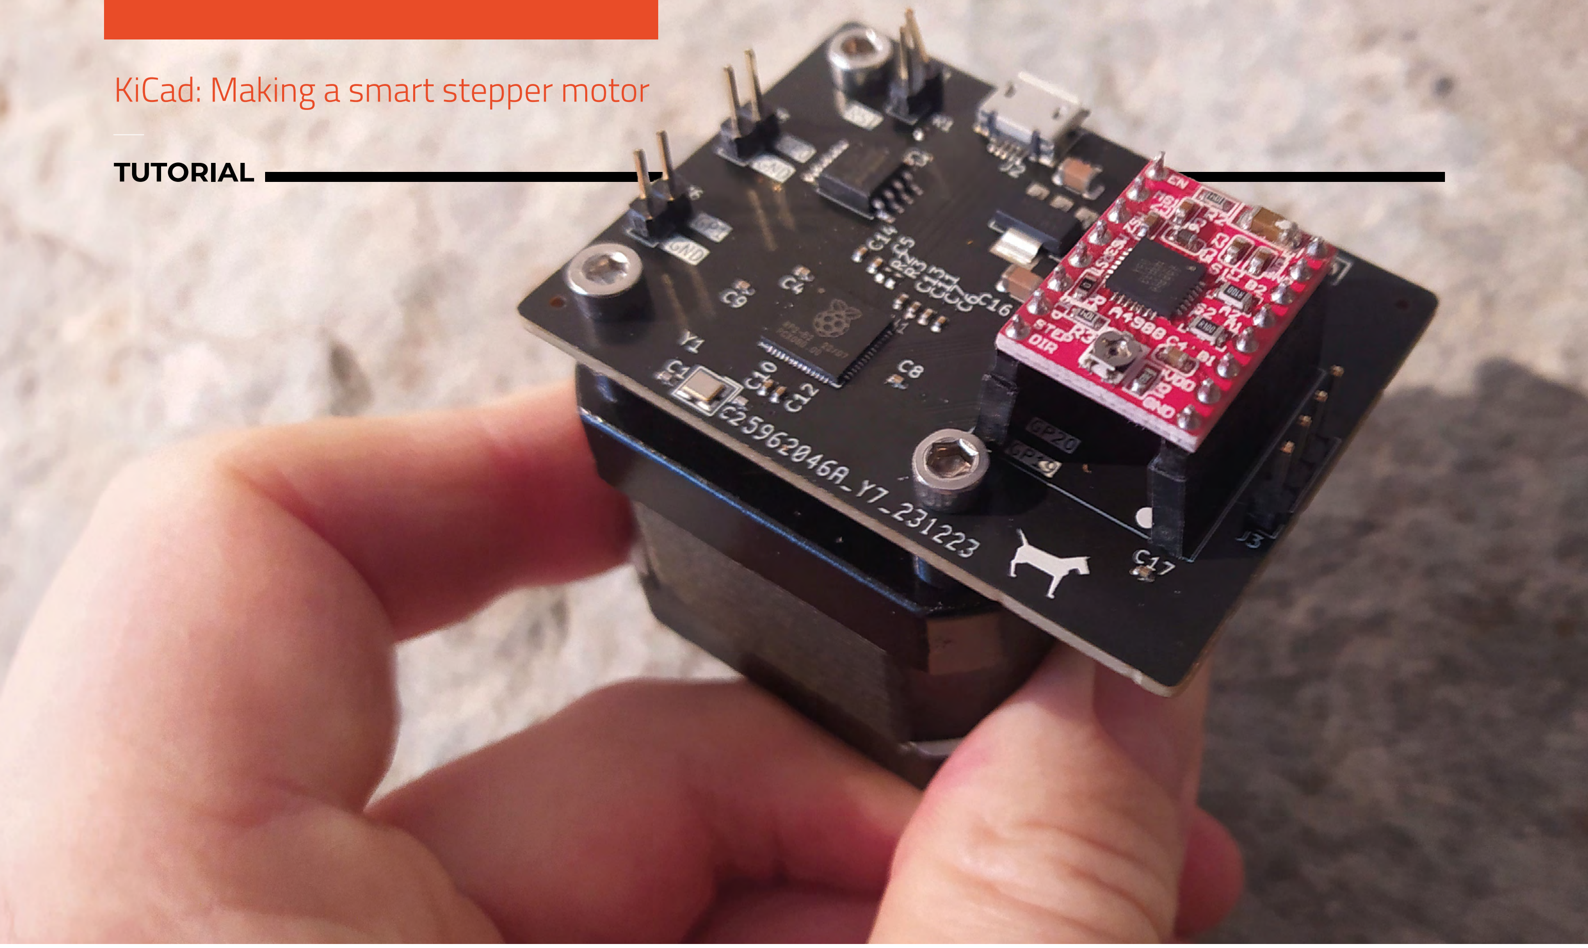

KiCad: making a smart stepper motor

In this section of the KiCad series, we adapt our earlier minimal RP2040 design to create an 'Urumbu'-style stepper motor control board

Jo Hinchcliffe

Jo Hinchliffe (AKA Concretedog) is a constant tinkerer and is passionate about all things DIY space. He loves designing and scratch-building both model and high-power rockets, and releases the designs and components as open-source. He also has a shed full of lathes and milling machines and CNC kit!

Urumbu is a mechanical concept created by Neil Gershenfeld of MIT, which concerns creating multi-axis machines simply, sidestepping some common approaches, and leveraging the development of computers. The go-to standard for many machines for many years is to feed them G-code line by line, using some kind of G-code sending application to stream the lines into a controller. This, in turn, drives the individual stepper motors to move the relevant axis. It's a solid system, but it harks from an era where parallel processing capabilities were rare and incredibly expensive.

Urumbu is interested in streamlining the making of machines, reducing both the cost and the complexity. In simplified terms, it essentially uses stepper motors (or theoretically other actuators) that have been adapted with an embedded microcontroller to run directly via USB connectivity. This means that, potentially, we can sidestep using G-code, and perhaps even avoid CAM altogether. Imagine building a machine where you perhaps parametrically define the output object and the script, or perhaps Python application, directly calculates the geometry of the form and directly controls the rapid prototyping machine connected to a convenient USB hub. This article won't cover all that, but it's well worth reading the article *Minimal Machine Building*, which you can download from **hsmag.cc/urumbu**.

You can also look around the Fab Lab depository, where you can find projects that have used the Urumbu approach, like this excellent pointing machine: **hsmag.cc/point**.

A NEW START

Let's create a separate copy of a KiCad project to work on. Open the Minimal RP2040 project we created, and then simply click File > Save As. We can then create another folder on our system and save the project into it with a new name. If you then open this folder in a file management application, you will see that all the KiCad-generated project files (the .sch and .pcb etc.) have been renamed to the new project name. We can also tidy and delete files which are specific to the old project and not relevant to the new project. For example, we won't be using the Gerbers, the CSV position and BOM files, as these will be different for the new project. Similarly, the edge-cut SVG that we created in Inkscape for the Minimal RP2040 project won't be used, so can be deleted. It's worth mentioning: make sure that you are in the right project folder before you start deleting files!

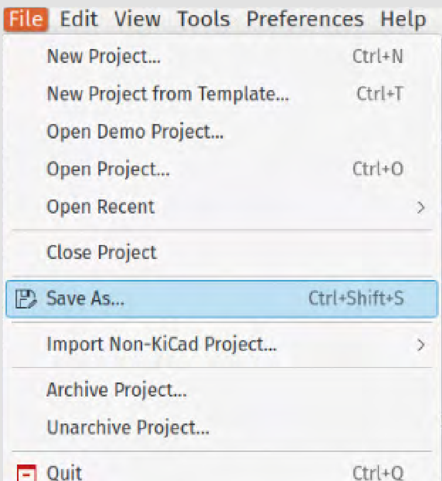

As the concept for Urumbu stepper motors is to use USB for control, the RP2040 IC is a great candidate for powering a driver board. There have been examples in the Urumbu community using NEMA 14 motors, which are convenient in the fact that they can often be controlled and powered by USB 2.0 and above. However, most small experimental rapid prototyping machines tend to use the larger NEMA 17 class of stepper motor. Check out any smaller home- or office-use 3D printer, smaller desktop CNC router, or hobby pen plotter and you'll find NEMA 17.

With NEMA 17 as our target, the first port of call is to find some mechanical dimensions and make some fundamental decisions. Looking at datasheets for NEMA 17, we find that the outer dimensions of the package are 42 by 42 mm, and that they have a set →

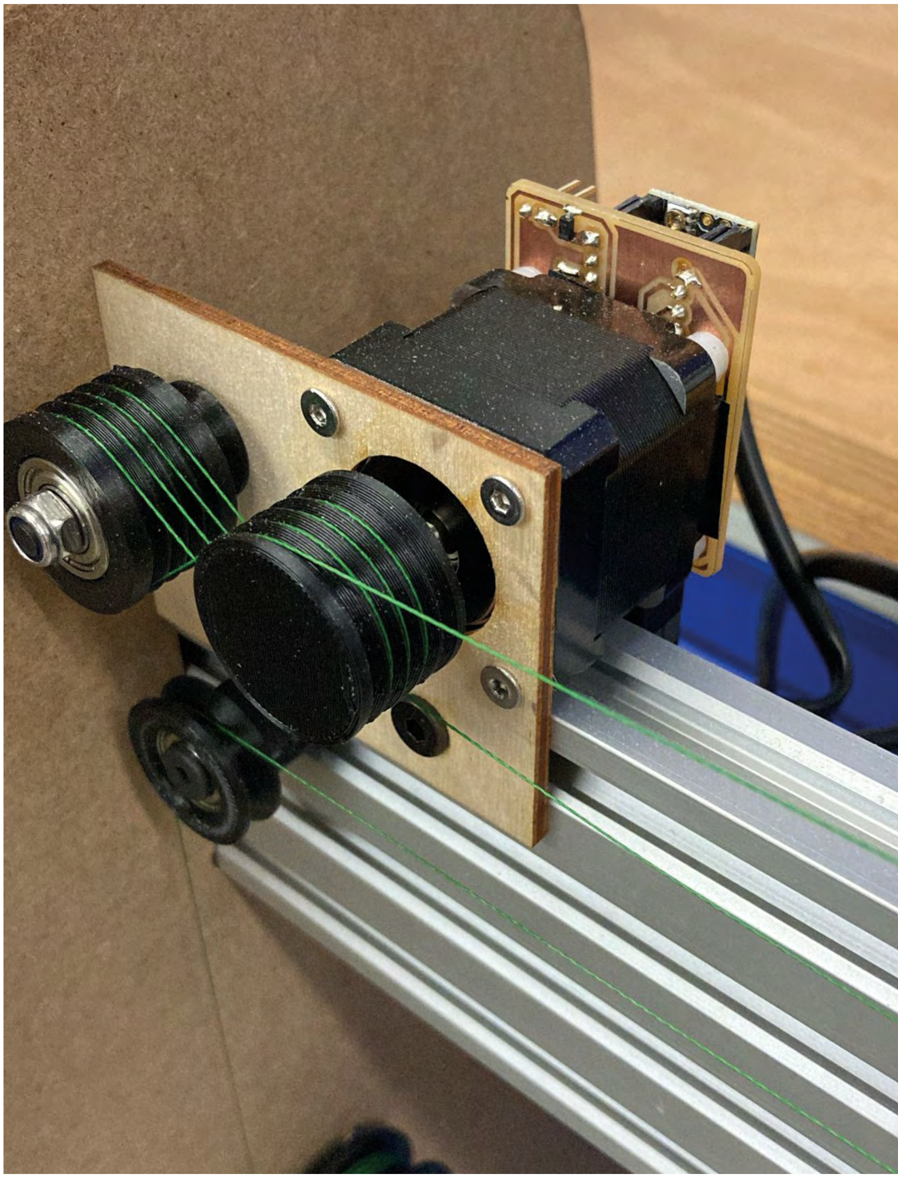

"**Potentially, we can sidestep using G-code,** and perhaps even avoid CAM altogether"

Left (opposite page) The finished PCB attached to a NEMA 17 stepper motor

Above An Urumbu-style smaller NEMA 14 motor with a CNC-milled, SAMD11-based board attached

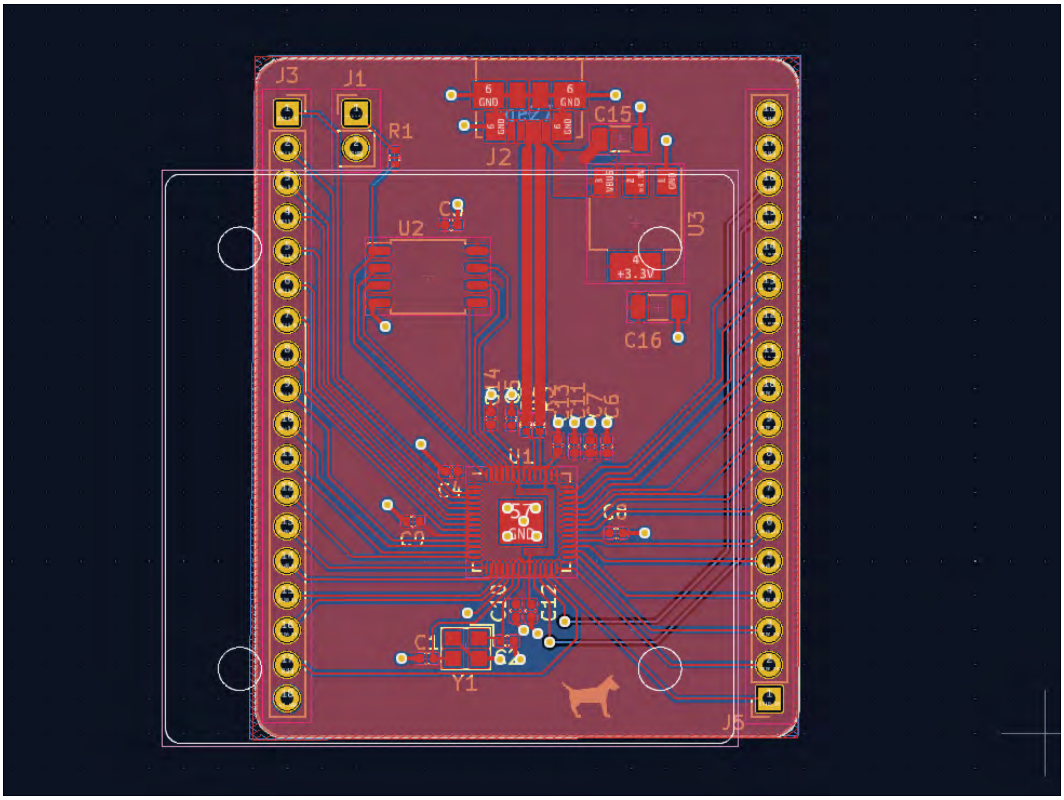

Figure 1 A NEMA 17 outline graphic laid out in Inkscape is imported as an SVG into the minimal RP2040 example to see what area we are playing with

Below Right Creating a custom symbol for any of the various similar form-factor stepper motor driver modules

of M3 bolts through the assembly in the corners of a 31 mm square. As we plan to use the minimal RP2040 circuit example as the basis of this project, it was a good starting point to draw up a NEMA 17 footprint and drop it into the minimal RP2040 design to see how things look (**Figure 1**).

We drew up a quick NEMA 17 footprint in Inkscape and imported it to the edge cuts layer in KiCad using File > Import > Graphics. It was obvious that we would have to reduce the size of some aspects of the minimal RP2040 layout, but not unreasonably. We wouldn't need all the GPIO pins broken out, so there were some easy space-savings to be had. There was probably just about enough room to also lay a motor driver IC and peripheral components on the board, but an early decision was made to actually use a module for the motor driver section as then the board would be compatible with a range of motor drivers. This meant it was necessary to expand the board dimensions in one axis. With this basic feasibility worked out, we set about editing the schematic to create our new project.

In the schematic editor, we first deleted all the GPIO breakout header sections as they weren't needed. We then used the symbol editor to create a custom symbol for our motor driver module (we covered creating symbols in the early parts of this KiCad series). With the symbol created, we then connected it to the RP2040 using labels to keep the general schematic sections easier to read.

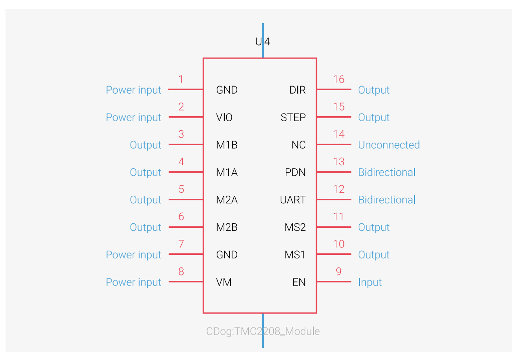

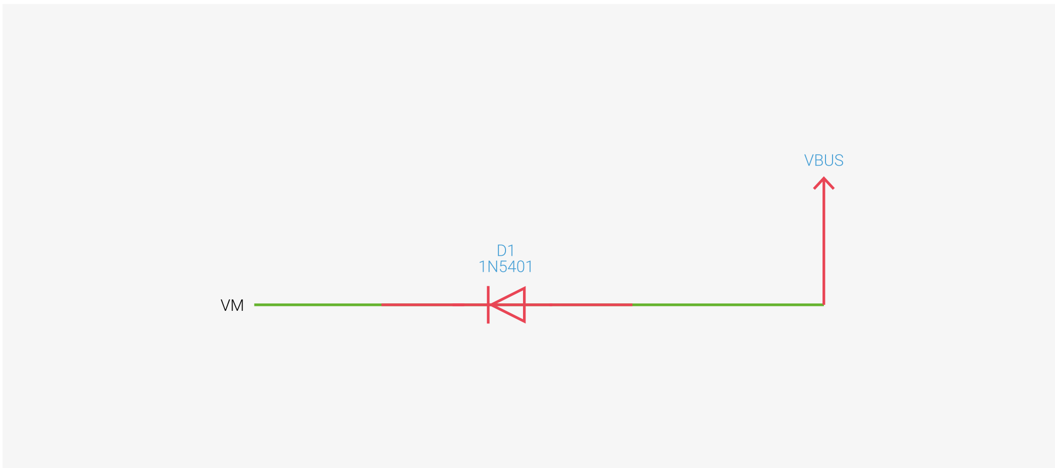

Figure 2 KiCad's workflow means that we can often simply add generic component symbols to a schematic and worry about which actual component each one will be later

One difference with running a NEMA 17 rather than a NEMA 14 is that although you could, in a slightly limited fashion, run the motor at 5 V from a beefy USB supply, it's likely you might want to run it from a larger external voltage. Most of the common stepper driver modules have the ability to do this, and have a 'VM' pin into which you can connect an external supply. We wanted to retain the ability to use either USB or VM, so we needed to add a connector for an external supply and a diode to protect the USB side of the system when the external supply was used. One of the great things about KiCad, that we may have mentioned before, is that the workflow of separate schematic symbols to which you then assign a component footprint means that we don't have to work out exactly which diode we are going to use at this part of the process. We can simply place a diode symbol and wire it into the schematic and consider the package later (**Figure 2**).

In addition to our motor driver, we wanted to add two header sockets connected to GPIO and ground to which we could attach switches to act potentially as limit switches to provide feedback and control options for any machines that we might develop with these motors. Again, we simply added these to the schematic.

With our adapted schematic largely complete, we set about making decisions on component choices and we also began to check previously used components were available. This is where things can get very tricky and time-consuming when using PCBA services. We were glad to see →

STARTING A WAREHOUSE

One way to avoid some of the component problems we have had in these projects is to pre-order components to be held ready for use in your project. This is sometimes referred to as a 'virtual warehouse', where you can buy an inventory of component stock and hold them until you are ready to place them onto a PCB assembly. This functionality is already built into your JLCPCB account and you can simply, once signed in to your account, move to the 'Parts Manager' page. On this page, you can use the 'My Parts Lib' to view and to add to your personal parts library. You can buy both Basic and Extended parts, and you can also pre-order out-of-stock extended parts for when they are hopefully restocked. For Basic parts from the JLCPCB parts warehouse, you have a minimum order requirement. However, basic parts are much less likely to go out of stock and, if they do, they are likely to have an alternate similar part available. One thing you need to know, though, is that these pre-purchased parts are only for use in assembly services – you can't suddenly have your library of parts mailed to you as a component order.

If you are creating a project and you think you are going to have a long development time where component stock might be an issue, this can be a great option for your peace of mind.

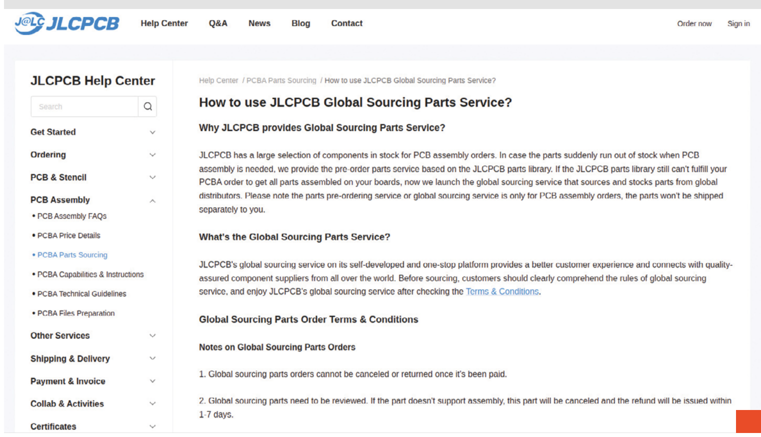

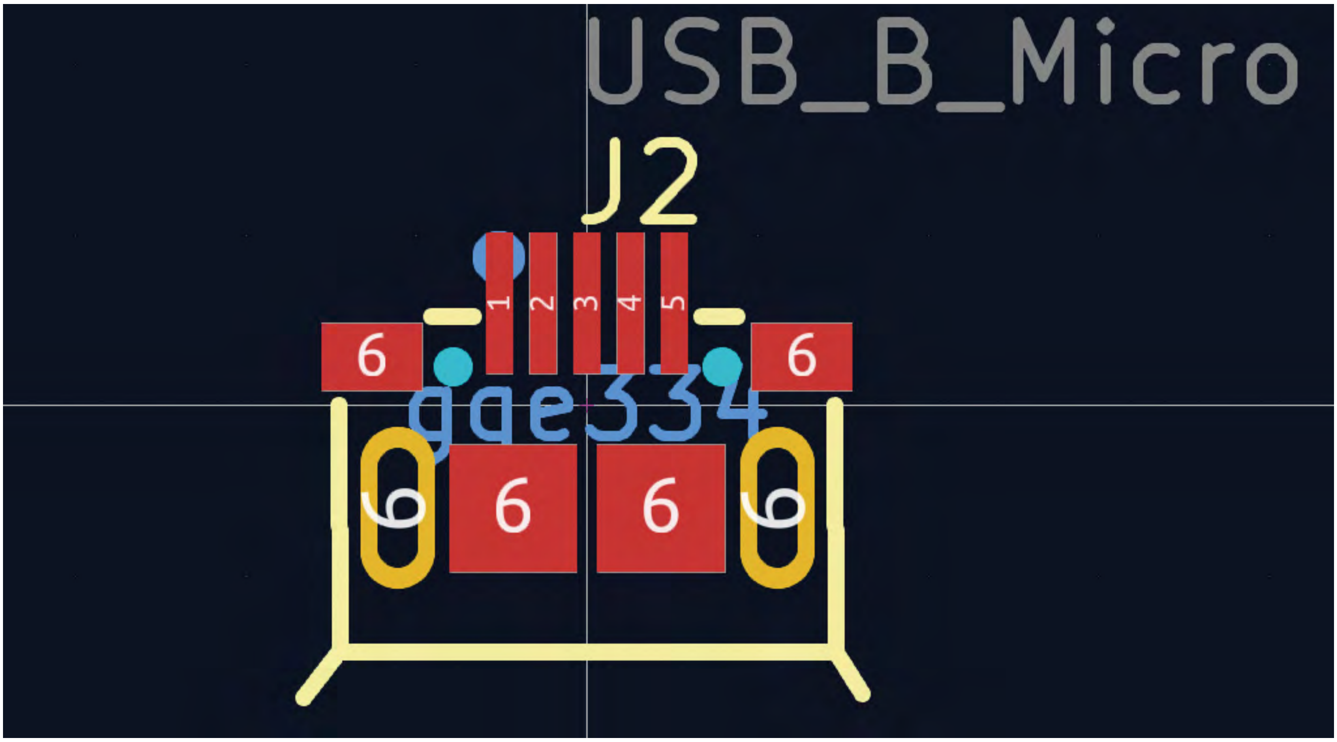

Figure 3 A new USB connector has been identified and the footprint created by converting the EasyEDA example on the product page

that the RP2040, the Winbond flash chip, and the 12MHz crystal were still in stock with JLCPCB. We also took the time to check that the smaller components, the capacitors and resistors, were all also in stock.

However, the USB socket and the 3V3 voltage regulator we had used previously were no longer in stock. Both of these are high turnover items and, at one point in this process, we couldn't identify any 3V3 voltage regulators in any package that were suitable for our project. We also found some challenges in that there would be a regulator listed, but not enough information available either in the listing or sometimes in the specific component's datasheet to actually make a decision on whether to include the part. With items like voltage regulators, it's fair to say that LCSC, the company that is the back end of the JLCPCB component warehouse, is continually changing and adding items and stock. What can appear a huge problem one day in your search results can suddenly have half a dozen more options in a couple of days' time.

We eventually found replacements for all our non-stock components. The replacement USB socket required a different footprint, and we used the approach of downloading the JLCPCB EasyEDA footprint and uploading it to the Wokwi EasyEDA to KiCad converter site, which again worked flawlessly (**Figure 3**). See earlier sections of this series for more examples of this in use.

The new USB connector had some through-hole chassis components, but was listed as an SMD. This concerned us, as the PCB assembly service charges quite a bit more for through-hole than surface-mount, but the part was attached via the single side surface-mount services without problems.

Having struggled to identify a suitable voltage regulator in any package at one point, having left the project for a couple of days we then found different stock available in the JLCPCB parts library and managed to find a drop-in replacement regulator which would sit on the same SOT23 footprint. It's definitely worth triple-checking the footprint and pinout of any swapped components to ensure that your wiring still works.

With most of the problems solved with regards to components, we set about editing the PCB layout. We needed to make the minimal RP2040

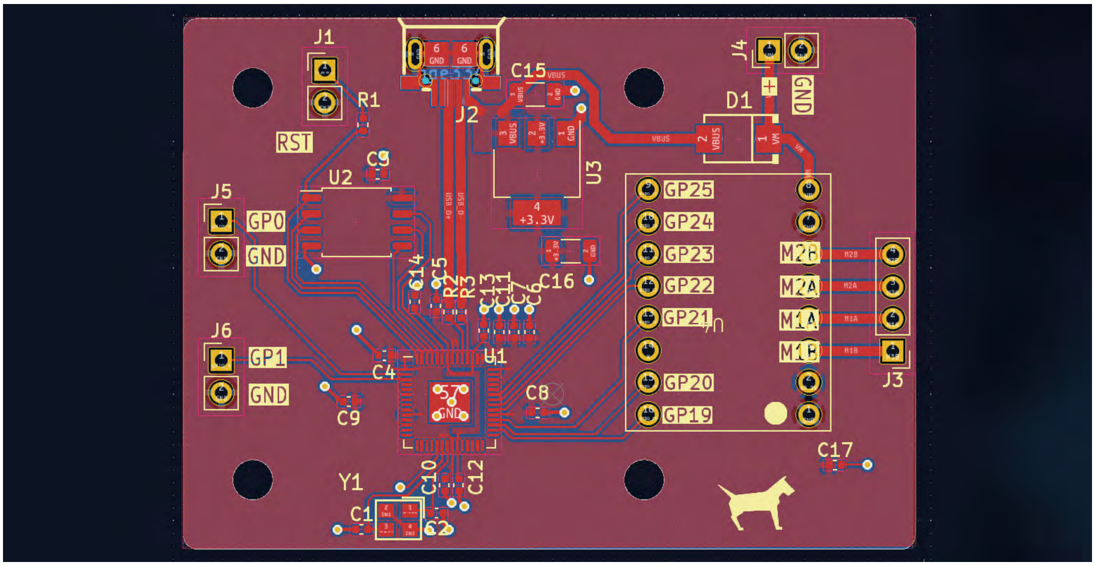

Left Below When trying to minimise redesigns and faced with component changes, you'll spend a lot of time using the search options!

Left Above The complete Urumbu RP2040 PCB layout

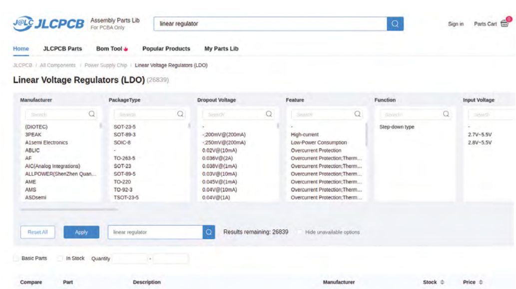

layout more compact to fit within the NEMA 17 footprint, and so we moved the actual RP2040 and crystal upwards, decreasing somewhat the distance between it and the USB socket. Sometimes, in reworking a PCB like this, the grab function is quite handy, where you can select a track, or a selection of tracks and components, and then use the 'G' hot key rather than 'M' and, instead of simply moving the objects, they are grabbed and the track connectivity remains which the tracks can move. This rarely results in a neat set of tracks in our experience, but it can be useful to create a quick new routing which you can then manually edit to neaten.

With the PCB design complete, it was the usual case of creating the Gerbers, BOM, and positional files for the project and uploading to JLCPCB. After a short production and delivery of the assembled PCBs, we created a standoff design for 3D printing using FreeCAD, and then the boards simply attach to a NEMA 17 using some longer M3 bolts. If you are interested in replicating these boards or playing with RP2040 Urumbu-style approaches, download this project from **hsmag.cc/issue76**.

WHOOPS-A-DAISY

In the spirit of failing out loud, we'd like to share a huge error we made in the production of this PCB. As you'll read in the main text, we decided to use motor driver modules rather than designing around a particular motor driver IC, as this meant we have lots of flexibility and redundancy with regard to motor drivers potentially going out of stock. The motor driver boards like the TMC2208 module, the DRV8833 module, and the A4988 module all share a common footprint with 16 pins, in two rows of eight pins in 2.54 mm spacing.

We didn't have the spacing between the rows, but a friend had a module on their desk and we messaged them for some dimensions. They sent me a collection of pictures with callipers held to the board, and also the module placed into a breadboard. We quickly counted across the breadboard to see how many columns wide the module was. It's six columns wide across the pin rows. When laying out the simple footprint for the module, we set the grid to 2.54 mm in the footprint editor and then drew one column of eight pins. We then counted across six rows and laid out the second column. Of course, that is an error: counting six rows across makes a module that would span seven columns of a breadboard, and is therefore 2.54 mm too wide. The simplest things are often the worst! The challenge with this sort of error is that they are not the kind of errors that can be detected by the DRC system, as connectivity to this simple footprint looks correct to the system. Only when the assembled PCBs arrived did we realise the error. Of course, this has been revised in the repository, so if you download this project, the footprint is correct. For the small number of boards we had manufactured, the horrid workaround is to slightly angle in the header sockets to bring them close to correct and then insert the module. Crude, but allows us to use the boards. We think it's fair to say that everyone will make mistakes; if and when it happens to you, try not to beat yourself up too much!

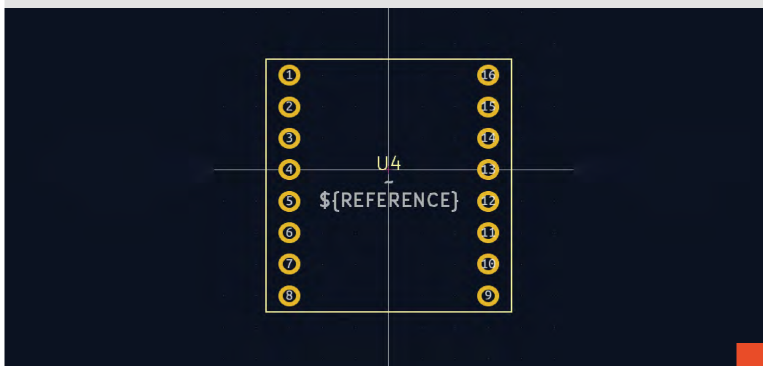

Building a modular Synth part 4: Enveloping

Getting softer edges on our notes

Ben Everard

Having given up learning the trombone in school, Ben has very little musical skill. However, he very much enjoys making noise.

Above Each potentiometer also needs a small hole for a prong that stops it rotating

In this article, we really hit the first big problem of the series. So far, we've looked at generating tones and control voltages. These are fairly straightforward problems on a microcontroller. In this issue, however, we are going to read in an audio signal, manipulate it, and then output it.

The particular module we're going to look at this month is an enveloper – this takes an audio input and a trigger. When the trigger is activated, it lets the audio signal through. However, it's not just an on/off. It attempts to mimic the volume profile of a real instrument, where the volume varies over time.

Most envelopers follow a similar four-phase profile:

1. Attack, where the sound quickly ramps up to its maximum value

2. Decay, where it drops down to the sustain level

3. Sustain, where the sound keeps playing at a constant level as long as the trigger is active

4. Release, where the trigger is removed and the sound drops away completely

In order to do this, we will have a multiplier that we apply to the input audio signal to get it to the right volume for this section of the envelope.

We need to read in the current value of the audio channel, apply the multiplier to it, and then output it. This is the same basic way any audio effect can be applied. Instead of multiplying it, we could filter it, delay it, or do something else entirely.

There are two basic ways we could try to perform this on the Pico.

1. Batch processing. We could use the direct memory access (DMA) feature of the RP2040 microcontroller to read in a set number of data points at a set timing interval, then process them all, and then output them, again using DMA to time the outputs.

2. We could put everything in a tight loop where we read in one value, process it, and then output it.

Option 1 is theoretically much higher performance. The DMA allows us to do the processing while the capture is going on, so all three (reading data in, processing, and sending data out) can all go on in parallel.

Option 2, however, is much easier, but it's only worth it if we can run the read-process-output loop fast enough.

HARDWARE

Before we can work out which option we want to use, we need to select our hardware.

For the audio input, we want the fastest analogue-to-digital (ADC) converter we can use, and this is the on-board ADC.

For audio output, we've been experimenting with two options. The 8-bit resistor ladder is really fast, but the lack of bit-depth will mean the audio is distorted at low levels. The MCP4725 I2C DAC we've used previously will give us better quality but will be slower to output. We'll see if we can output data fast enough.

We also need to add some user inputs to control the shape of the envelope. We want them to be able to define the length of the attack, decay, and release phase as well as the volume of the sustain phase. This means we need four additional analogue inputs. We could use the on-board ADCs for two of these, but this would slow down the speed at which we're able to read in the main audio input, so we'll use a separate SPI ADC – the MCP3008.

This module depends entirely on our ability to output data from the DAC quickly enough. The good news is that the MCP4725 supports the high-speed version of I2C which runs at 3.4MHz. The bad news is that the I2C controller on the RP2040 only supports the fast version which runs at 400kHz. 400kHz might sound like a lot, but it takes 30 data points to output a single value, so that meant the maximum speed we could output data was about 13kHz. And it would be slower than this because there would be additional overheads.

The RP2040 has a trick for speedy I/O: Programmed input/output (PIO). These are programmable bits of logic that you can attach to I/O pins to program your own peripherals, and are capable of implementing I2C at the full high-speed (3.4MHz) data rate. Fortunately, there's a library to help us do it: **hsmag.cc/PIOI2C**.

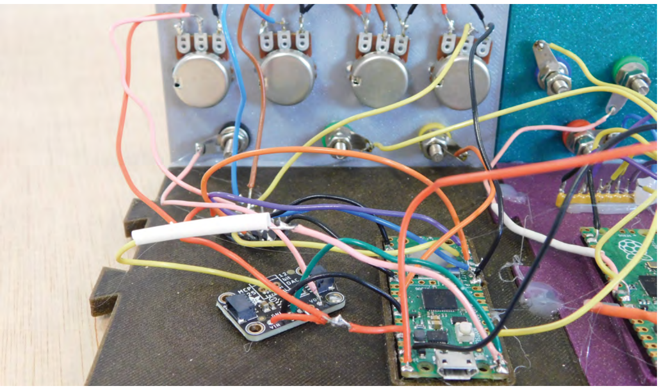

Below Left This is the messiest build of all the modules so far – it might be time to create a PCB

MCP3008

We used a stand-alone MCP3008 ADC. These don't need supporting circuitry, and we connected ours dead-bug-style by gluing it upside down and soldering wires directly onto the pins. You do have to be a bit careful to make sure you get the right pins.

We need four of the eight analogue channels. Each one needs to be connected to a potentiometer. Each potentiometer should have three pins. The pin on one end should go to 3.3V, the other end should go to ground, and the middle should go to the input pin on the MCP3008.

When looking down on the MCP3008 from the top, the left hand row of pins are the input channels 0 to 7 (with 0 being closest to the notch). On the right hand side, from the notch, there are two power pins that go to 3.3V, ground, clock, data out, data in, NOT chip select and ground. Clock, data out, data in and NOT chip select go to Pico pins 10, 12, 11 and 13.

This gives us a few ways of writing out data, but using the simplest option (a blocking write), we can get a data rate of around 24kHz. This isn't great for audio, but it's fast enough to prove the concept. Most of the time is taken up with overheads around starting and stopping the transfer, so we suspect we can get some significant speed-ups here if we need to in the future, but for now, we'll crank along at this speed.

All the code for this module and the rest of the Pico modular project is in the GitHub repository: **github.com/benevpi/PicoModular**.

At its heart, the program is really simple. We're just modifying the data and sending it out as fast as we can. The main loop that does the enveloping is:

```
        while(true) {
                adc_raw = adc_read();
                output = (uint16_t)(adc_
raw*multiplier);
                i2c_buff[1] = output >> 4;
                i2c_buff[2] = output << 4;
                err = pio_i2c_write_
blocking(pio, sm, PIN_HS_ENABLE, 0x62, i2c_buff,
3);
        }
```

This reads in the value from the audio input and puts it in the appropriate place in the three-byte buffer for the I2C write (the 0th byte is the write-to-output command). We need to change the 12-bit number into a number split across two 8-bit sections, so we first shift it back by four bits in the first byte, and then shift the remaining four bits up four places for the second byte. Since `i2c_buff` is `uint8_t`, any leftover bits get lost, and we only keep the ones we need. →

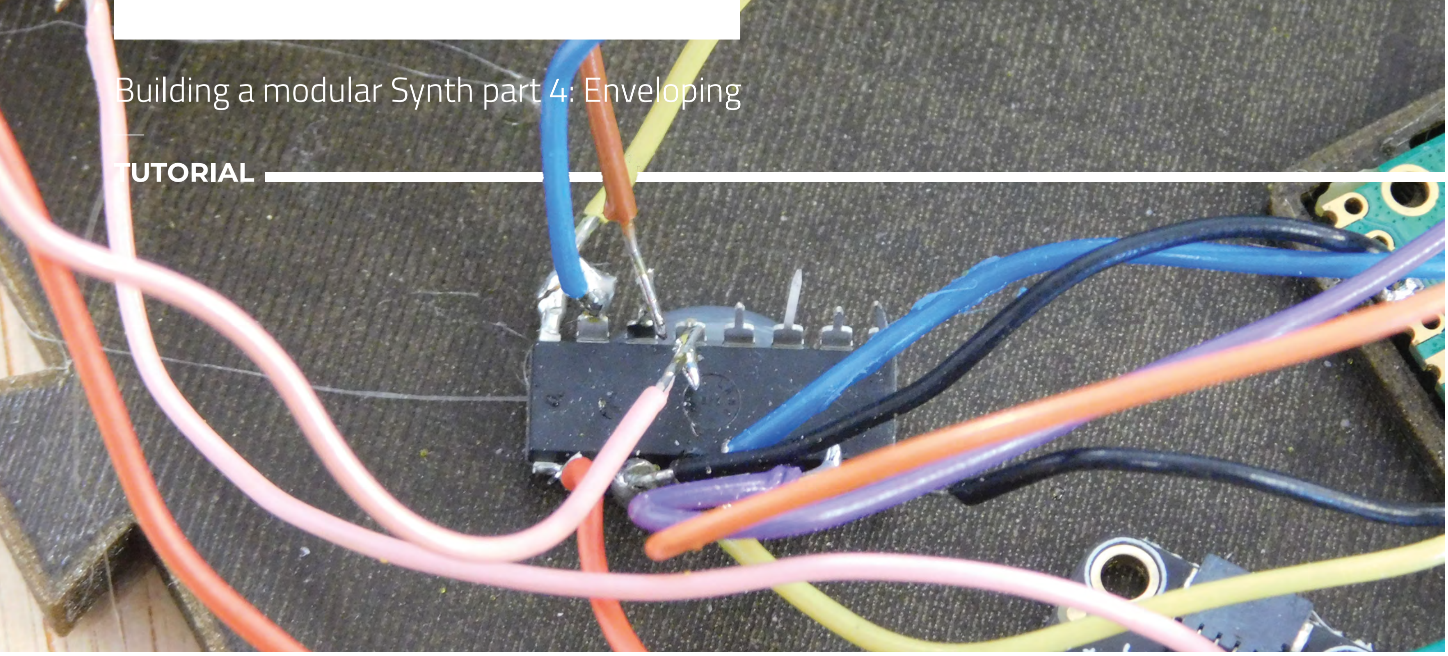

Above You can solder onto the pins of an upside-down chip, but you do have to be careful not to snap them off

MULTIPLIER

The only thing left to do is calculate the multiplier. We need to do this without impinging on the main loop at all. Fortunately, with RP2040, we can do this easily by running it on the other core.

```
void core_one_loop() {
  int attack;
  int decay;
  int sustain;
  int release;
  while(true) {
    //read trigger
    if(gpio_get(PIN_TRIGGER) ==1) {
      //really simple debounce
      sleep_ms(50);
      if (gpio_get(PIN_TRIGGER) ==1) {
        trigger_state = 1;
        //let's get the state of all the pins
        attack = read_analogue(ATTACK_CHANNEL);
        decay = read_analogue(DECAY_CHANNEL);
        sustain = read_analogue(SUSTAIN_CHANNEL);
        release = read_analogue(RELEASE_CHANNEL);

        multiplier = 0; // it should already be

        printf("multiplier: %f \n", multiplier);

        //attack stage
        for(int i=1;i<100;i++) {
          multiplier += 0.01;
          sleep_ms(attack/100);
        }

        multiplier = 1.0;

        //decay stage
        for(int i=1;i<100;i++) {
          multiplier -= (1024-sustain)/(1024*100);
          sleep_ms(decay/100);
        }

        //sustain stage
        while(true) {
          if (gpio_get(PIN_TRIGGER) == 0) {
            sleep_ms(50);
            if (gpio_get(PIN_TRIGGER) ==0) {
                 break;
            }
          }
          sleep_ms(50);

        }

        //release stage
        for(int i=1;i<100;i++) {
          multiplier =  (((float)sustain/
(float)1024)/(float)100) * (float)(100-i);
          sleep_ms(release/100);
        }
        multiplier = 0;
      }

    }

  }
}
```

This loops around continuously, waiting for the trigger pin to go high. There's a very rudimentary debouncing system where it pauses for 50 milliseconds, and if it's still high, it triggers the enveloper. While this is very rudimentary, it has the advantage of being very predictable timewise. It shouldn't need much debouncing because this won't be triggered by a mechanical switch, and it's unlikely this could be triggered by accident, but we thought we'd put it in just to be safe.

First, we have to call the `read_analogue` function which pulls the current values from the external analogue inputs. These should return a value between 0 and 1024.

The `attack, decay`, and `release` sections are all quite similar. We loop through 100 times and change the multiplier appropriately each time, and just add a suitable delay. It could be done a bit more regularly, but of all the accuracy issues with this enveloper, this isn't huge.

The `sustain` section holds the current multiplier until the trigger is released. Again, it has a simple debouncer.

We now need to assemble everything and plug it together. Once the DAC (see image caption) and the ADC (see box) are wired up, you just need to connect the audio to GPIO 27. You should now be able to connect it up with the rest of the Pico modular system.

An example setup would be to use Pico MIDI to generate sound. The CV of this should go to PicoVCOGateless, and the Gate should go to PicoEnveloper Trigger. The output of PicoVCOGateless should go to the audio input of PicoEnveloper, and the output of PicoEnveloper should go to your amplifier.

Throughout this series, we've said that we are not looking to make a high-end synthesizer. We are aiming to make an affordable modular synth to use as the basis for experiments. However, in its current form, this enveloper is dredging the bottom of that bucket. As far as we can tell, it is all asking a bit too much of the power supply and this is introducing unwanted artefacts into the audio. So, next issue, we'll see if we can smooth that out a little and work on improving the audio quality. □

I2S

We've used an I2C DAC for this project that was never really intended to stream a signal in the way we're using it. There is a different type of DAC designed for exactly this type of situation, which uses a protocol called I2S. I2S is designed to take precisely timed samples and output them at the right time.

The problem we have with this is that most hobbyist-level I2S DACs come paired with a class-D amplifier. Class D amplifiers output a PWM signal, so we can't then feed the output of them into a new module.

This might be an area that we'll look at again in the future, especially if we find a DAC that would be a good fit for this project.

Below
Each potentiometer needs power to one side, ground to the other, and the signal is the middle pin

Below Left
We used an Adafruit module for the MCP4725 DAC. To hook this up, you need to connect VIN to 3.3V, GND to ground, SCL to GPIO 1, SDA to GPIO 0, and VOUT to output

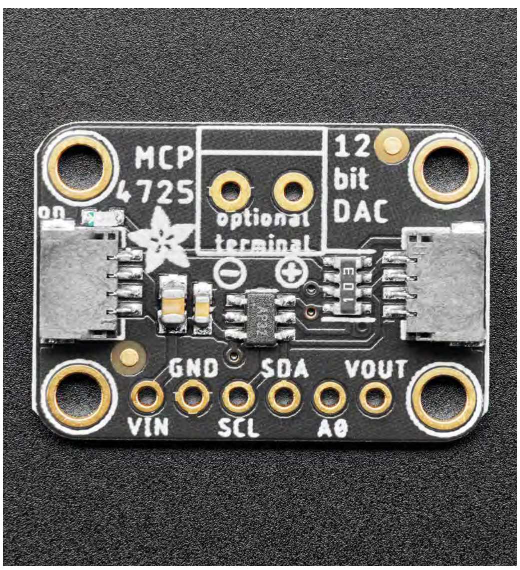

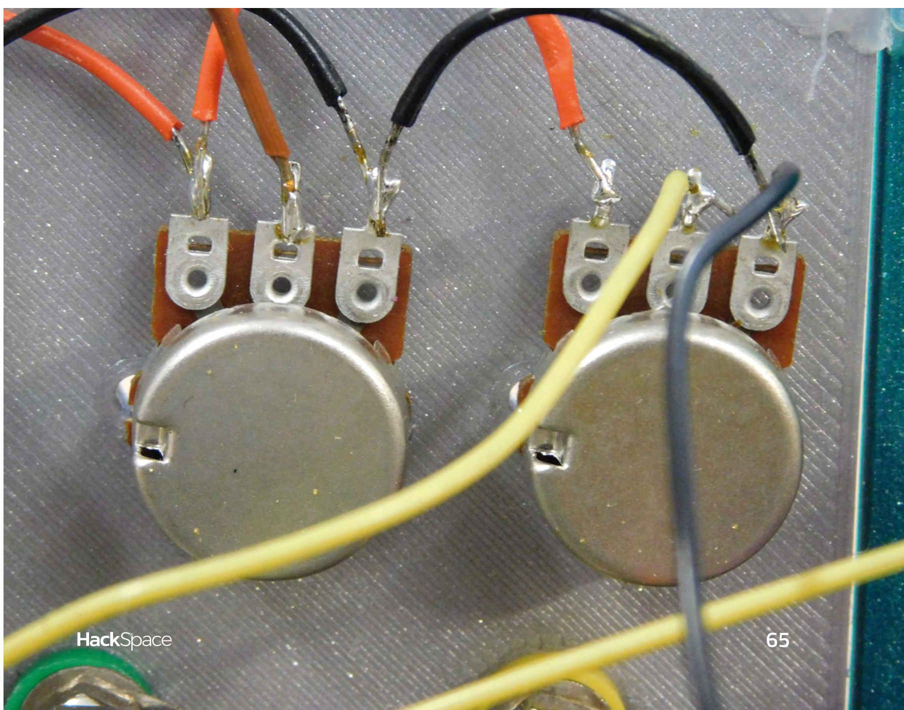

FROM THE MAKERS OF The MagPi THE OFFICIAL RASPBERRY PI MAGAZINE
NEW 2024 UPDATE
START
RETRO GAMING WITH RASPBERRY PI
3RD EDITION
180 PAGES OF VIDEO GAME PROJECTS
+ BUILD AN ARCADE MACHINE
PLAY & CODE GAMES!
File manager
Change color
Change icon
Extra blocks
Save dump
Load dump
Hexeditor
Defragment
Reformat
/ : Select
: OK
: Cancel
Display the list of the files found in the selected VMU.

Don't forget the batteries!

How to power a single-board computer or microcontroller with a battery pack

Phil King

Phil King is a freelance writer and sub-editor specialising in technology. The former books editor for Raspberry Pi, he is also the author of *Chromebooks In Easy Steps*.

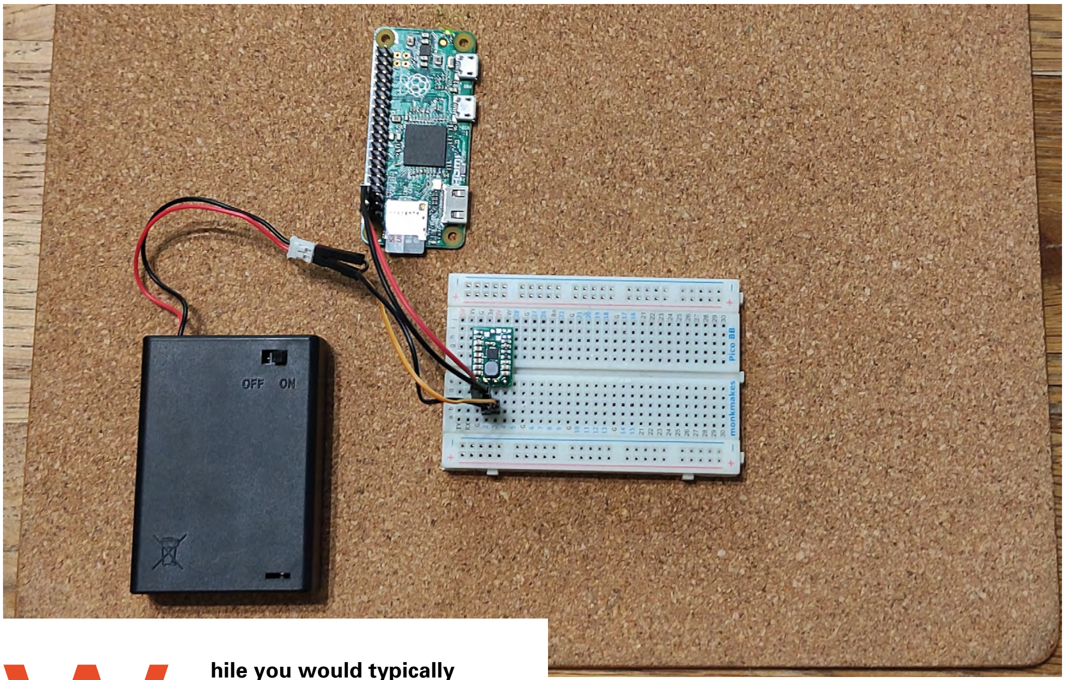

Above Right It's advisable to use a voltage regulator between the power source and board. Here, we're using a Pololu S7V8F5 to maintain a steady output voltage of 5V to a Raspberry Pi Zero

While you would typically power a single-board computer or microcontroller through its USB socket (whatever the type), it is also possible to power most boards from a source wired directly to their 5V power and GND (ground) pins. This includes all the Raspberry Pi SBC models, along with microcontrollers such as Arduino, Pico, and ESP32/ESP8266 boards. That's what we are going to explore in this tutorial, looking at the pros and cons, including the risks – you don't want to see that magic smoke!

WHY DO IT?

Depending on your project, there may be various reasons why you would want to power the main controller board directly via the pins instead of the main USB power port.

A common use case is when you want to supply power from some sort of battery source, which could be standard alkaline or rechargeable NiMH batteries (typically AA size), Li-ion cells, or a LiPo battery (see 'Battery types' box, overleaf).

For instance, you can power a Raspberry Pi SBC using a power HAT or SHIM to connect a LiPo battery. You could even use a coin cell if space is really tight.

The direct-to-pins route may also be appropriate if you want to power multiple parts of a project, such as motors and servos in a robot, from a single source – which could be higher voltage than you need for the board, in which case you'll need a buck (or step-down) converter to regulate the voltage. We'll take a look at some options for that later, but first we're going to do the 'wrong thing'!

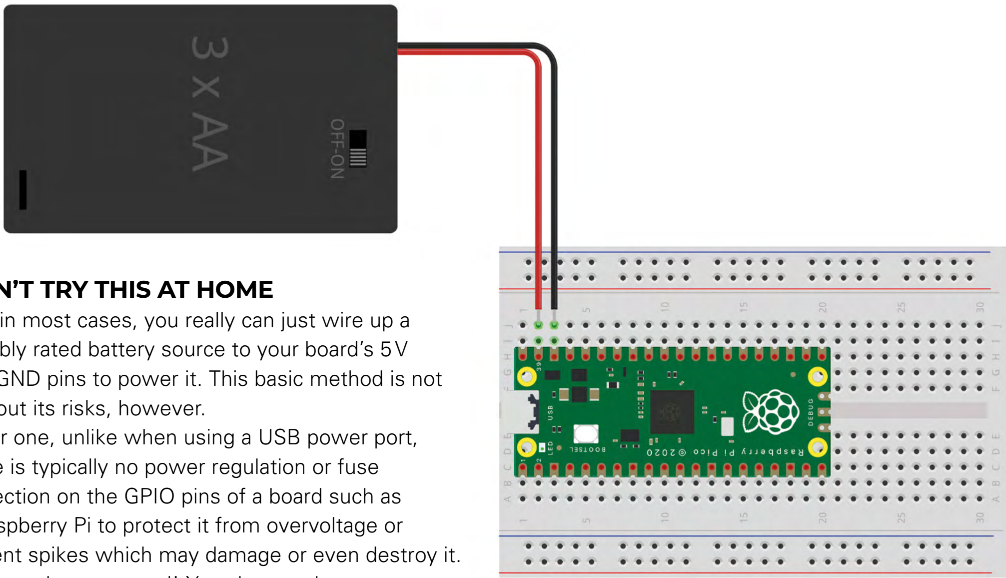

Figure 1 Wiring diagram for connecting a battery pack directly to Pico's VSYS and GND pins – not ideal, but it does work

DON'T TRY THIS AT HOME

Yes, in most cases, you really can just wire up a suitably rated battery source to your board's 5V and GND pins to power it. This basic method is not without its risks, however.

For one, unlike when using a USB power port, there is typically no power regulation or fuse protection on the GPIO pins of a board such as a Raspberry Pi to protect it from overvoltage or current spikes which may damage or even destroy it. You have been warned! You also need to ensure you connect the pins the correct way around – reverse polarity is not good.

Still, for starters, we're going to try powering a Raspberry Pi Pico microcontroller and then a Raspberry Pi Zero SBC using this method, just to check that it works.

The lower power requirements of these two boards do make it easier to supply them from a battery source, compared to larger, more power-hungry boards like the Raspberry Pi 4 and 5.

Using a standard 3 × AA battery holder (available from most electronics suppliers) with three 1.5V alkaline batteries, we used a digital multimeter to measure the output voltage at a shade under 4.5V, but this may vary according to the newness and type of batteries. If you use rechargeable NiMH batteries, you may well find the voltage is lower, around 3.6V – so if you really want 5V, you'll need to use a 4 × AA battery pack instead (don't do that with alkaline ones, unless you're using a buck converter to reduce the voltage).

We will then connect the battery pack's (red) power line to the Pico's VSYS (5V) pin, then the (black) ground line to a GND pin (**Figure 1**). We're using standard jumper wires for this basic test, but there will be a slight voltage drop depending on the thickness and length. Ideally, thicker 18AWG wires should be used.

With the wires connected to the correct pins, Pico should power up. An easy way to check is to first install a simple script to blink its on-board LED. With the Pico loaded with MicroPython and connected to a computer, save the following code as **main.py** (so it autoruns when powered up). →

VOLTAGE REGULATORS

While you could always build your own DIY regulator circuit, using resistors in series as a 'voltage divider', it's more convenient to buy an off-the-shelf voltage regulator board. There's a wide range available, suitable for almost every need.

A step-down (or buck) converter will reduce an input voltage to produce a lower output voltage; for instance, 5V to 3.3V. Conversely, a step-up (or boost) converter will increase an input voltage to produce a higher output voltage. Due to Ohm's law, this affects the relative current level of the input and output; in a step-up converter, the output current level will be reduced. Also, as no converter is 100% efficient, you should expect to lose some of the overall power level in the process.

There are versatile boards, such as the Pololu S7V8F5 regulator that we used, that can step an input voltage up or down to maintain a set output voltage. Lower-cost regulators are available, but a key factor to consider is the tolerance of their output voltage – if it's 10% on a 5V output, it could creep up to 5.5V, which is above the maximum 5.25V input voltage for a Raspberry Pi SBC (the Pico is more tolerant).

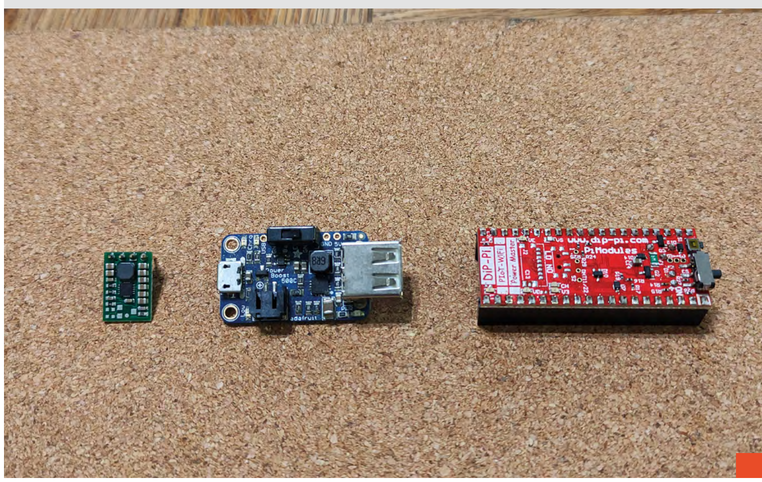

Below A small selection of regulator boards. From left to right, a Pololu S7V8F5, Adafruit PowerBoost 500C, and DiP-Pi PIoT for Raspberry Pi Pico

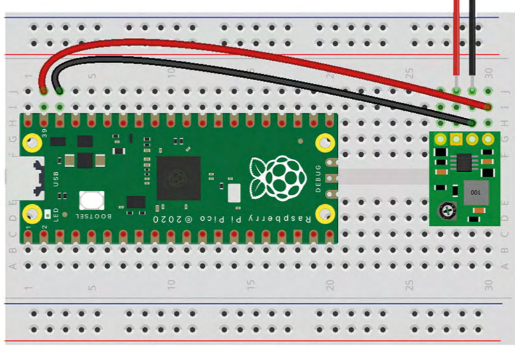

Figure 2 Wiring diagram of a battery pack connected to Pico’s pins via a voltage regulator

```
from machine import Pin
from utime import sleep

led = Pin('LED', Pin.OUT)

while True:
    led.toggle()
    sleep(1)
```

You should then see the Pico’s green LED blinking when you connect your battery power source.

Since the Pico can work from a supply as small as 1.8V, there should be no issues getting it working from even fairly low-capacity batteries. A board such as a Raspberry Pi single-board computer or Arduino is a different prospect, as it has a narrower voltage tolerance range – typically from 4.75 to 5.25V. Any lower and it may well not stay powered for long, if at all, in particular under a heavy workload or with any peripherals connected.

THE BUCK CONVERTER STOPS HERE

To obtain a more reliable battery power supply to the device, it’s advisable to use some sort of voltage regulator. A buck converter is able to take a higher DC voltage and convert it to a lower one while increasing the current. This is typically used with a higher-voltage battery source than needed, so long as it’s within the converter’s input range. This is a good option when powering the board and other parts of a project separately from a single higher-voltage source.

Since our 3 × AA battery pack is supplying less than 5V, we opted to use a Pololu S7V8F5 5V Step-Up/Step-Down Voltage Regulator (**hsmag.cc/PololuS7V8F5**). As its name implies, this tiny board can also ‘step up’ (i.e. increase) a lower voltage to 5V, with a slight loss of current – you don’t get anything for free. Based around the TI TPS63060 regulator, the S7V8F5 can accept voltage in the range 2.7V to 11.8V, making it very useful for supplying a regular 5V output to boards and other components in a project.

After soldering some header pins to the S7V8F5 and plugging it into a breadboard, we connected the battery pack wires to the appropriate inner pins (VIN and GND) and then the outer 5V pin and inner GND (again) via jumper wires to VSYS and GND pins on our Raspberry Pi Pico (**Figure 2**). You can ignore the outer SHDN (shutdown) pin, which is only used to put the S7V8F5 into low-power mode.

As you’d expect, the fairly voltage-tolerant Pico powered up without any problems. Now to try it with a Raspberry Pi Zero. With a steady 5V supply from the Pololu S7V8F5, it powered up successfully. With the generally low power draw of the Raspberry Pi Zero, you can get away with a relatively low current (typically 80 to 120mA).

With more demanding Raspberry Pi models, however, the maximum level of current supplied

A HIGHER POWER

If you need to supply a more demanding board such as a Raspberry Pi 3B/B+, 4 or 5, you’ll need a 5V step-down converter and a higher-powered source in order to maintain the current necessary to keep it running. For instance, a Raspberry Pi 4 requires at least 700mA just to boot up, but typically more (around 1.25A) when running any workload with peripherals attached.

The situation with the Raspberry Pi 5 is complicated by its use of PD (Power Delivery) to deliver more current to its USB ports. It normally requires a 27W mains power supply to be able to run at full pelt. You’re unlikely to need this for a portable, battery-powered project, but then you might be be better off instead using a Raspberry Pi Zero 2 W, which has a typical bare-board active current consumption of 350mA.

may not be enough. In which case, you can simply connect a beefier battery pack to the S7V8F5 – up to 11.8V, although you still need to beware of possible LC (inductor-capacitor) voltage spikes caused by an initial rush of current when first connecting it. If the supply is 9V or more, Pololu advises keeping connections short and/or connecting a 33µF or larger electrolytic capacitor close to the regulator between VIN and GND. See **pololu.com/docs/0J16** for more information on LC voltage spikes.

LIPO SUCTION

Next, we thought we'd try powering our boards using a LiPo (lithium polymer) battery cell, which offers a compact, more convenient alternative to regular batteries. Now, most LiPo cells output 3.7V, so unless you're using a Pico or another board that isn't too fussy about the input voltage, you'll need a buck converter to step it up to 5V.

There are two key factors to consider when choosing a LiPo battery: the capacity, in mAh (milliamp hours), which shows how much power can be stored when it's fully charged; and the 'C' rating, which signifies the maximum continuous output that the battery can maintain. To find the maximum discharge current level, multiply the capacity by the C rating: for instance, a 2000mAh battery with a 1C rating can output a maximum 2A for one hour, although this current level may be reduced if using a buck converter to boost the voltage.

> "You will need a special charger **module to recharge your LiPo cell**"

Naturally, the bigger the power requirement of your device, the quicker the battery will drain. Another thing to note is that as it drains, the output voltage will typically drop – another good reason to use a voltage regulator.

Most modern LiPo cells have a built-in protection circuit to prevent overcharge, overdischarge, and overcurrent, as well as offering short circuit and over-temperature protection. Make sure yours does. You then won't have to worry about it. You will need a special charger module to recharge your LiPo cell, however. →

BATTERY TYPES

Let's take a look at some of the most common battery technology types:

Alkaline: Standard batteries that are designed to be single-use, so don't try charging them – it can be dangerous!

NiMH: Rechargeable batteries that come in the same standard sizes as alkaline ones. The downside is that they typically have a slightly lower voltage output: 1.2V compared to 1.5V for AA batteries.

LiPo: Lithium polymer cells that can be recharged. Their slimline nature makes them ideal for tight spaces. They typically output 3.7V, so you'll need a step-up regulator to get 5V. Most LiPo cells have built-in circuitry to prevent overcharging and to shut off the output when the voltage drops too low, to avoid damage to the battery.

Li-ion: Lithium ion cells that are rechargeable and come in a range of shapes and sizes, including coin cells. Cylindrical sizes include the popular 18650, along with AA and AAA. They have a higher power density than LiPo batteries and don't suffer from the 'memory effect' that makes cells harder to charge over time. There are versatile boards, such as the Pololu S7V8F5 regulator we used, that can step an input voltage up or down to maintain a set output voltage. Lower-cost regulators are available, but a key factor to consider is the tolerance of their output voltage – if it's 10% on a 5V output, it could creep up to 5.5V, which is above the maximum 5.25V input voltage for a Raspberry Pi SBC.

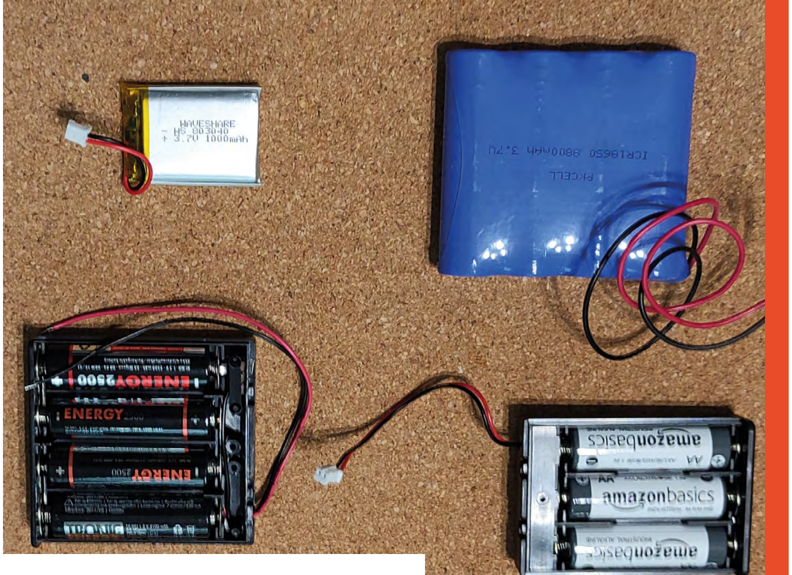

Above
Clockwise from top left: LiPo cell, 18650 Li-ion pack, alkaline batteries, NiMH rechargeables

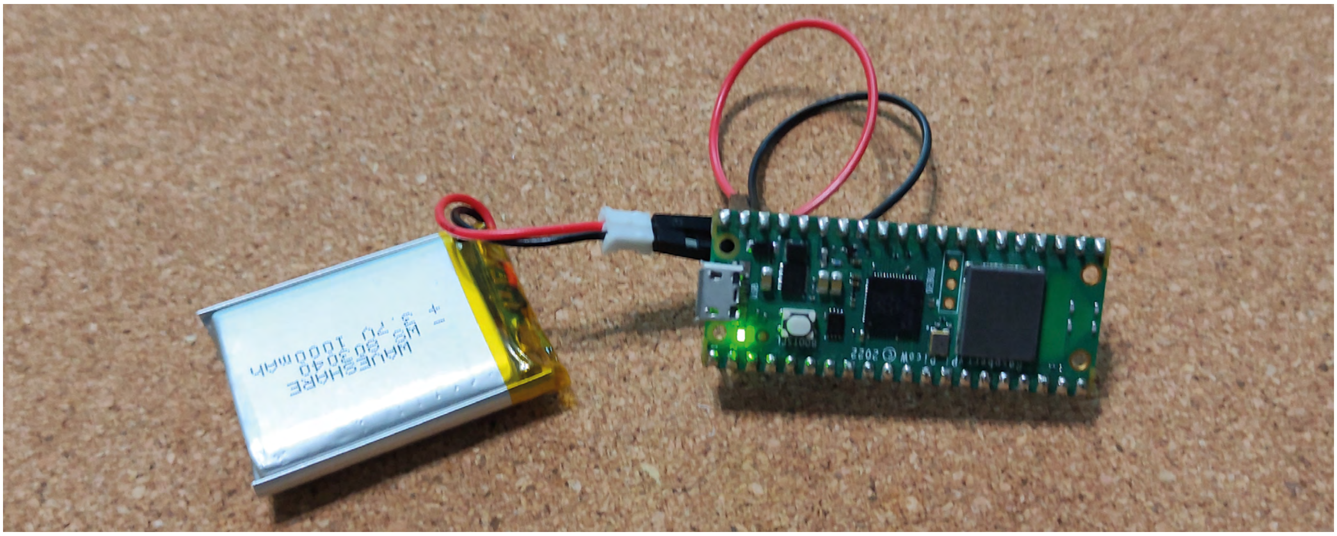

Above LiPo batteries are a good, slimline option for portable power. Here, we've connected one directly to a Pico

When it comes to connecting it up, most standard LiPo cells have a JST 2-pin or Deans connector (the latter type mainly used for RC vehicles and drones) – you can either plug male jumper wire ends into this for your connection or chop the connector off and tin the wire ends, perhaps for use in a screw terminal block on a buck converter.

With a fully charged 1000 mAh LiPo battery connected via a Pololu S7V8F5 regulator to a Pico W, we measured the output voltage at a stable 5.11 V, while the current level was around 60 mA with the on-board LED blinking. At this discharge level, the battery could in theory last for up to 16.5 hours, but in reality it will be less due to inefficiencies and the battery automatically shutting off power when it reaches a low level.

Other devices may draw more (or less) power. A Raspberry Pi Zero W typically draws between 80 and 120 mA, depending on whether it's run headless or not and whether Wi-Fi is used, so a 1000 mAh battery should in theory be able to power it for up to 12.5 hours (1000/80), although it's likely to be shorter due to inefficiencies.

UNINTERRUPTED POWER SUPPLY

As we've seen, batteries can be used to power a portable project for a limited period. While this may be fine for many projects, such as robots or a handheld games console, it's not suitable for those that need to be kept powered up for long periods, or continuously. In that case, your battery power can be used as a backup for a mains supply – or even another portable power source – to provide an uninterrupted power supply (UPS). This is ideal for projects that need to be kept online continuously to log sensor data, such as a weather station.

> **"There is typically no power regulation or fuse protection…** to protect the board from overvoltage or current spikes which may damage or even destroy it"

Even in the event of a power outage, they will stay powered up.

For this, you'll need a UPS board that can instantly switch the supply from mains to battery backup when required. Various such boards are available, including the Adafruit PowerBoost 1000C which doubles as a voltage regulator.

We tried out a Waveshare UPS HAT for Raspberry Pi Zero W/WH (**hsmag.cc/WSZeroUPS**), which is mounted on the underside of the latter, connecting with springy pogo pins – thus leaving the GPIO pins free. It also comes with a 1000 mAh LiPo battery which can be charged when the HAT is plugged into the mains. It supplies a regular 5 V supply to the Raspberry Pi Zero, on which you can run a Python script to show useful battery info such as the real-time load voltage, current (negative if battery power is being used), power, and percentage.

Other UPS boards may offer even more detailed battery data and features. A good option for Pico, the DiP-Pi range (**hsmag.cc/DiP-Pi**) offers multiple power inputs from 6 V to 18 V, battery charging and monitoring, and various extra features. For Raspberry

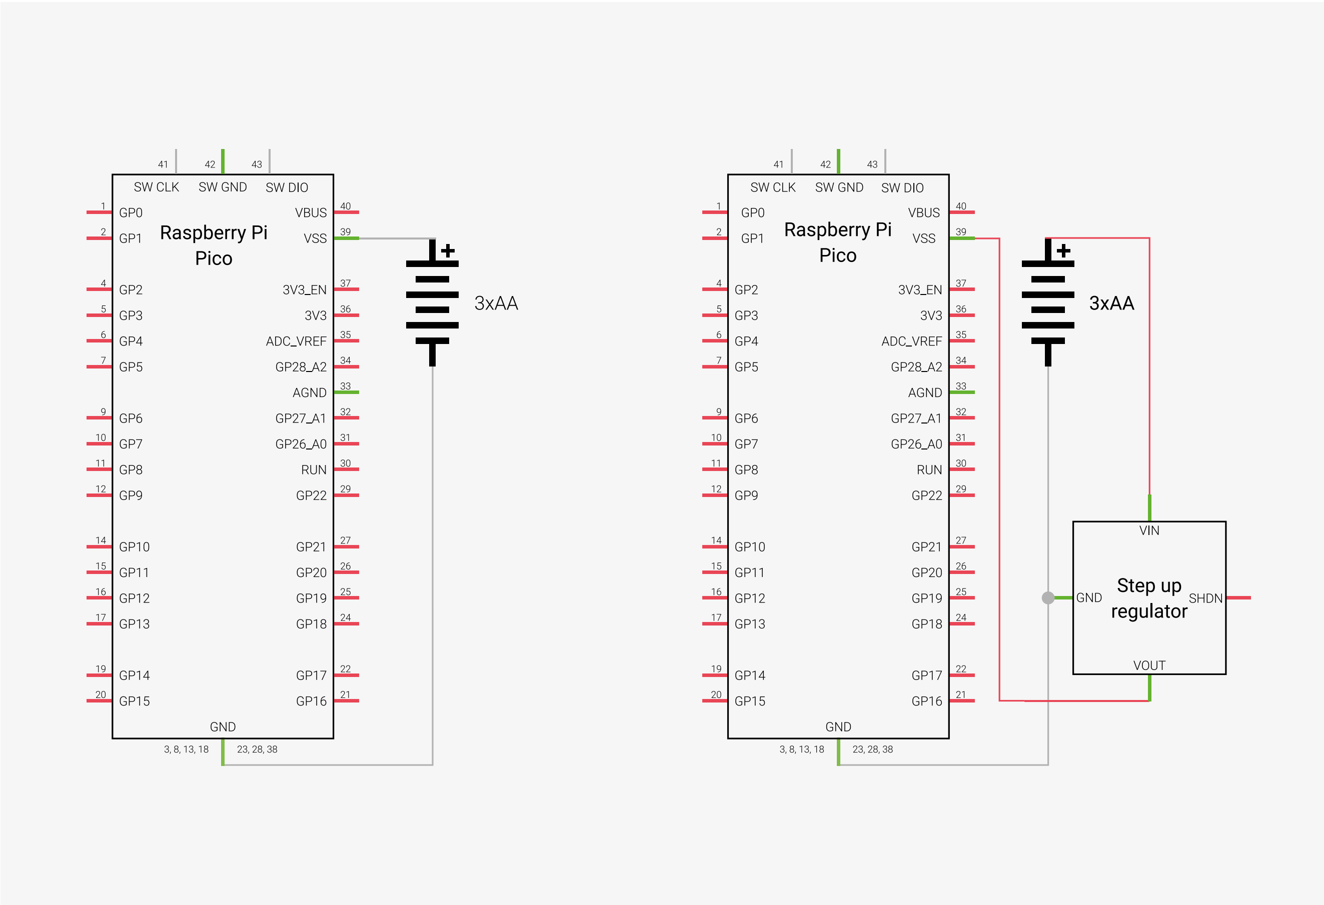

Pi SBCs, the PiJuice HAT (**hsmag.cc/PiJuice**) offers advanced power management features including a 'watchdog' timer that monitors a software 'heartbeat' and, if it's not heard for a certain period, automatically resets the Raspberry Pi.

A UPS board is also ideal for connecting alternative power sources such as solar panels, enabling you to charge a connected battery to maintain a continuous supply for remote projects away from a mains power source, such as a weather station or wildlife camera.

PORTABLE POWER OPTIONS

In conclusion, yes, you can power a Raspberry Pi or other board directly from batteries to the GPIO pins and it should work, providing suitable voltage and adequate current are supplied. However, there's nothing to protect against issues such as voltage spikes and dropouts which may damage the device or even the battery (if it lacks on-board circuit protection). For an inexpensive microcontroller, you may not be too bothered about that, and many of them can take a wide range of input voltages – from 1.8 to 5.5 V in the case of a Pico.

If you don't want to run the risk, you're much better off using some form of power regulator to ensure a constant voltage and offer protection against spikes and other issues. Some of these boards offer advanced features, too, enabling you to monitor the battery level and output, and in some cases even run custom scripts when certain events are triggered.

In addition, a UPS board will ensure an uninterrupted power supply backed up by the battery and can be used with renewable power sources such as solar panels so that you can maintain a power supply to critical projects that need to keep running continuously – for example, to gather sensor data over a long period of time. □

USEFUL RESOURCES

- **Raspberry Pi power requirements**: hsmag.cc/RPiPower
- **Raspberry Pi Pico power**: hsmag.cc/PicoPower
- **Battery University**: batteryuniversity.com
- **Circuit Journal Arduino power guide**: hsmag.cc/ArduinoPower

Above
Schematic (left) for connecting a battery pack directly to Pico, and schematic (right) for connecting it via a voltage regulator board

Adventures in multicoloured printing

Use your plastic squirter to make impressive pictures

Rob Miles

Rob has been playing with hardware and software since almost before there was hardware and software. You can find out more about his so-called life at robmiles.com.

In this article, we are going to look at how you can add colour to your 3D prints. We'll start with a look at lithophanes and them move onto tools you can use to create impressive 3D artwork. You can find the source images and the 3D printable files for this article here: **hsmag.cc/Colour3D**.

LITHOPHANES

Lithophanes are a great answer to the frequently asked question: 'Why did you get a 3D printer?' You can reply: 'To make these' and hold a thin piece of printed plastic up to the light to reveal a neat image. The next question that is asked is usually: 'Can you make me one of those?'

Figure 1 below shows a lithophane panel made by the author a few years ago of his then-new car. On the left, you can see the surface of the 3D-printed panel and, on the right, you can see how the lithophane looks when light is shone through it. If you look at the panel on the left, you can see that the dark areas of the image are thicker than the lighter ones. It's funny how such a low-quality print can look so good as an image.

MAKING YOUR OWN LITHOPHANES

If you have a 3D printer, you should have a go at printing at least one lithophane. They make nice personalised gifts. Lithophane panels are created by software that converts light levels in an image into thickness values in a 3D-printed panel. The darker the image, the thicker the panel, so that less light gets through. You can create curved or spherical lithophanes which you can make into great lights. The Lithophane Maker at **lithophanemaker.com** is a good place to start. You load an image into the web page and the software generates an STL (an abbreviation of STereoLithography) file which you then slice and print in white PLA. There are also designs available for lights into which the lithophane can be mounted. We'll now consider a few things that you need to remember.

Figure 1
This lithophane was printed on standard white PLA filament using a Creality Ender 3. It took a few hours to complete

YOU'LL NEED

- **A 3D printer** (preferably one that can switch between different filaments)
- **For coloured lithophanes** you will need some cyan, magenta, yellow, and white filament
- **For drawing plaques** you can use any filaments you happen to have around

Figure 2
The selector has been loaded with cyan, magenta, yellow, and white filaments to make a colour lithophane

PRINTER POOP

Each time the printer switches to a different filament colour, it purges the previous colour out of the print head and replaces it with the new one. The remnants of each colour change are ejected from the back of the printer so, after printing for a while, you get a little pile of 'printer poop'. These plastic fragments give you an interesting record of your printing history and are not without artistic merit, as the picture above shows.

LAYERS AND NOZZLES

A 3D print is created as a series of layers stacked on top of each other. The printer is sent a file telling it where to put the filament that makes up each layer. The printer file is produced by a program called a 'slicer' which takes an object design and converts it into a series of horizontal slices (or layers) to be printed. When a design is 'sliced', you need to tell the printer the 'layer height' which gives the height of each slice to be printed. When slicing a lithophane,

> **The darker the image, the thicker the panel** so that less light gets through

the smaller the layer height, the more levels of 'grey' you will get in the image. In **Figure 1**, you can see noticeable banding in the sky where the brightness changes from light to dark. If I had used a lower layer height for the print, this banding would have been less pronounced. Most printers regard 'normal quality' as 0.2 mm layer height (five layers per mm), but for a better-looking lithophane, you can reduce the height to 0.1 mm – although this will double the time the lithophane takes to print as the printer must now produce twice as many layers.
You can also improve the resolution of the image by reducing the size of the printer nozzle. The nozzle is the hole where the extruded plastic comes out of the printer to be added to the printed object.
Most consumer 3D printers have a nozzle which is 0.4 mm in diameter, but, with many models, you can switch the nozzle for one which is smaller – usually 0.2 mm. This lets the printer produce more detailed objects (useful for things like gears and the like) at the expense of a huge increase in printing time. The author has found that he is quite happy with the images produced by the 0.4 mm nozzle. There are some samples of each size later in this article.

DIGGING INFILL

A slicer has an 'infill' setting that determines the amount of filament extruded inside the object being printed. It is expressed as a percentage. An object printed with an infill of 0% would be completely hollow (and not very strong). An object printed with an infill of 100% would be completely solid. The normal setting for infill is 20%. This is fine for most objects unless we want them to be particularly robust. For a lithophane, we need the interior areas to be solid to stop all the light, so the infill for the slicer should be set to 100%. This also has the effect of increasing the printing time.

COLOUR LITHOPHANES

The first 3D printers printed one colour. You could only print different colours by swapping the filament during the print. However, today, you can get printers which can automatically switch between filaments during a print job. **Figure 2** above shows the AMS (automatic media selector) used with the author's Bambu Lab P1S printer (**bambulab.com**). →

QUICK TIP

The first lithophanes were produced in the 18th century. They were made from porcelain. An artist would carve a 'master' image in wax that would be used to make moulds into which liquid porcelain (a type of clay) would be poured. When the porcelain hardens and is fired in a kiln, it becomes translucent, with the thicker areas stopping more light, just like 3D-printed lithophanes.

Figure 3 You can preview the expected image on the web page, but this view is not as impressive as the final printed one will be

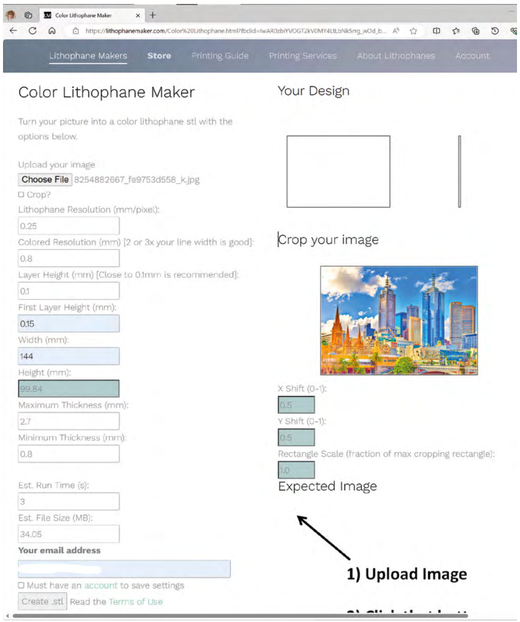

Figure 5 The magenta filament is being selected for the magenta layer of the print

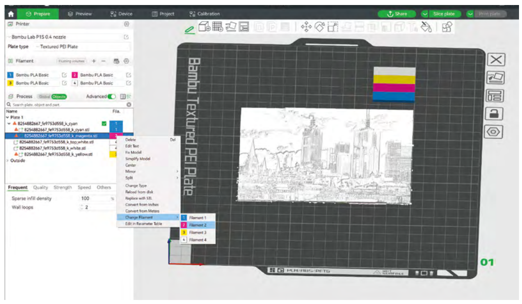

The AMS is connected to the printer which tells the AMS which filament is required during the printing process. When the printer needs to change filaments, the current filament is wound back onto the roll and the new one is loaded into the print head. The printer then pushes the new filament through the head to purge out any remaining pieces of the previous filament before resuming printing with the new colour.

If your printer can print multiple colours (or you are prepared to manually switch filaments during the print job) you can make coloured lithophanes. These work in the same way as black and white ones, but instead of a single layer representing the light and dark areas of the image, you now use four colours: cyan, magenta, yellow, and white.

Producing the lithophane files is the same as for monochrome, although the web address for the generator site is very slightly different: **hsmag.cc/ColourLithophane**.

Figure 3 shows an image being prepared for conversion into a lithophane. The image being used to create the lithophane is a slightly processed photograph of the Melbourne skyline. The output from the generator is a zip archive which contains six STL files: two for white (the bottom and the top of the lithophane) and one each for the cyan, magenta, and yellow layers.

Figure 4 shows how the layers work together to create a coloured picture. The colour layers are sandwiched between the two white layers.

Figure 4 In the actual lithophane object, these images are all directly on top of each other. In the picture, they've been separated so that we can see how they work

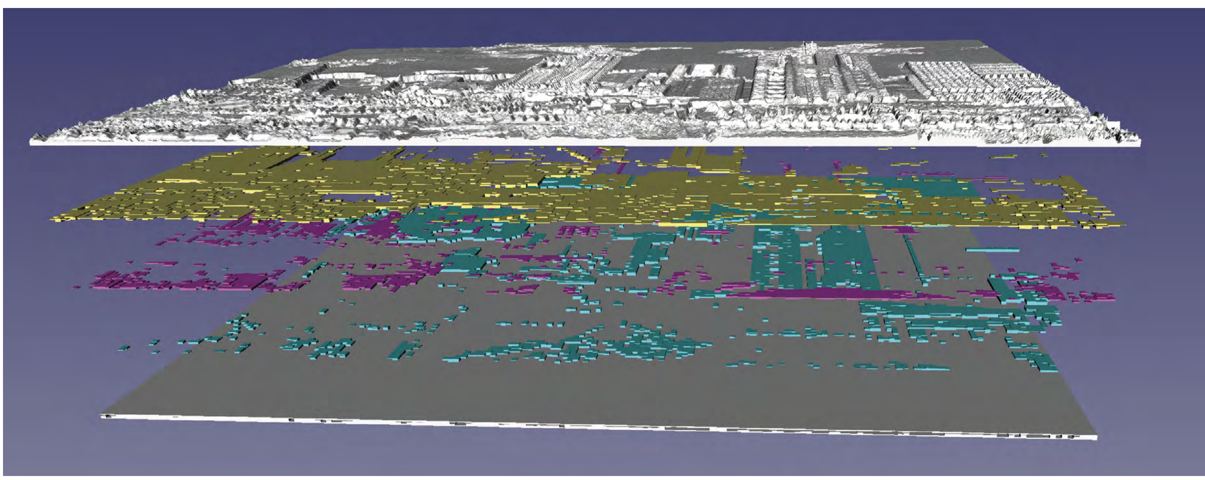

QUICK TIP

Using a media selector makes it easier to print right up to the end of a roll of filament. You can load a full roll into the selector and tell the printer to switch to that when the current roll runs out.

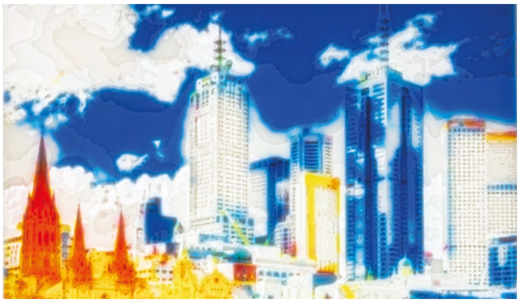

Figure 6
This print took less than two hours to complete

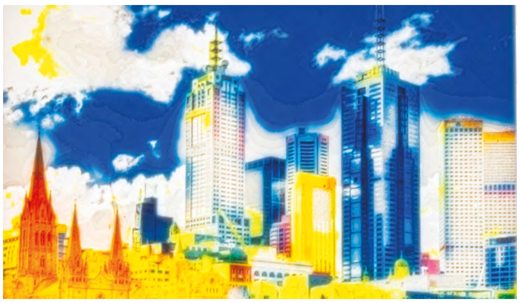

Figure 7
This print took over eight hours to complete

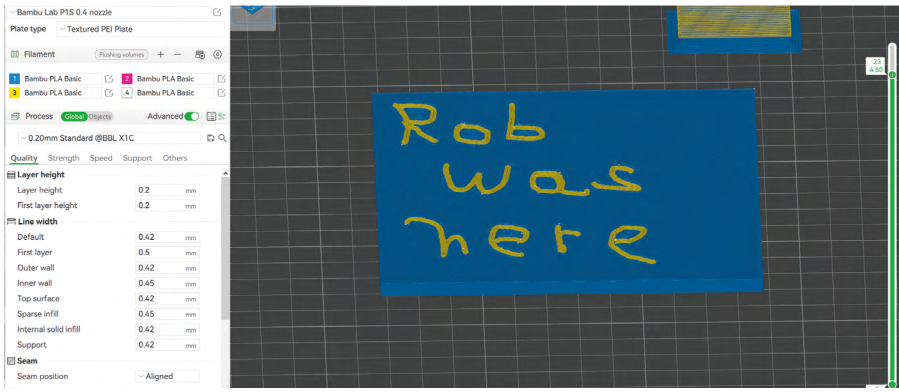

Figure 8
There are also tools in Bambu Studio that can embed text into objects and make coloured signs

Light travelling through the layers is filtered so that only the required colours are visible. The STL files are loaded into the slicer program which will produce the control file for the printer. The author uses the Bambu Studio slicer program. This can produce print files for a wide range of different 3D printers, not just ones made by Bambu Lab. You can download it from here: **hsmag.cc/BambuLab**.

> You can improve the resolution of the image by reducing the size of the printer nozzle

Figure 5 shows the Melbourne lithophane being prepared for slicing. Each colour layer STL has been imported as one element of a composite object. This means that the layer objects will be printed on top of each other, rather than being distributed around the print bed as different objects. The printer control file that is produced will contain commands to the printer to switch filaments at the appropriate times.

LITHOPHANE QUALITY

Figure 6 shows the kind of results you can expect if you use standard printer settings. There is a reasonable amount of detail, but the banding on the sky is quite pronounced.

Figure 7 shows the results achieved by reducing the layer height and using a 0.2 mm nozzle. The sky looks a lot better and there is much more detail, particularly in the buildings on the right of the picture. Whether the improvement in detail is worth an extra six hours of print time is something worth thinking about though.

COLOURED PRINTS

Lithophanes are great fun and produce very impressive results, but they are not the only way of making coloured prints. Lithophanes also need a light behind them, which restricts where you can use them. However, we can also use a 3D printer to produce coloured prints simply by taking an existing print and 'colouring it in'.

Figure 8 shows how the Bambu Studio paint tool lets you use the mouse to paint primitive messages into a 3D model. You can also use the tool to select parts of an object and colourise them – for example, →

QUICK TIP

When adding colour to a print, remember that filaments may have to be changed in many layers. The slicer will usually tell you how many filament changes are required for a given print.
You can greatly reduce print time by adjusting the colours in your design to minimise the number of colour changes.

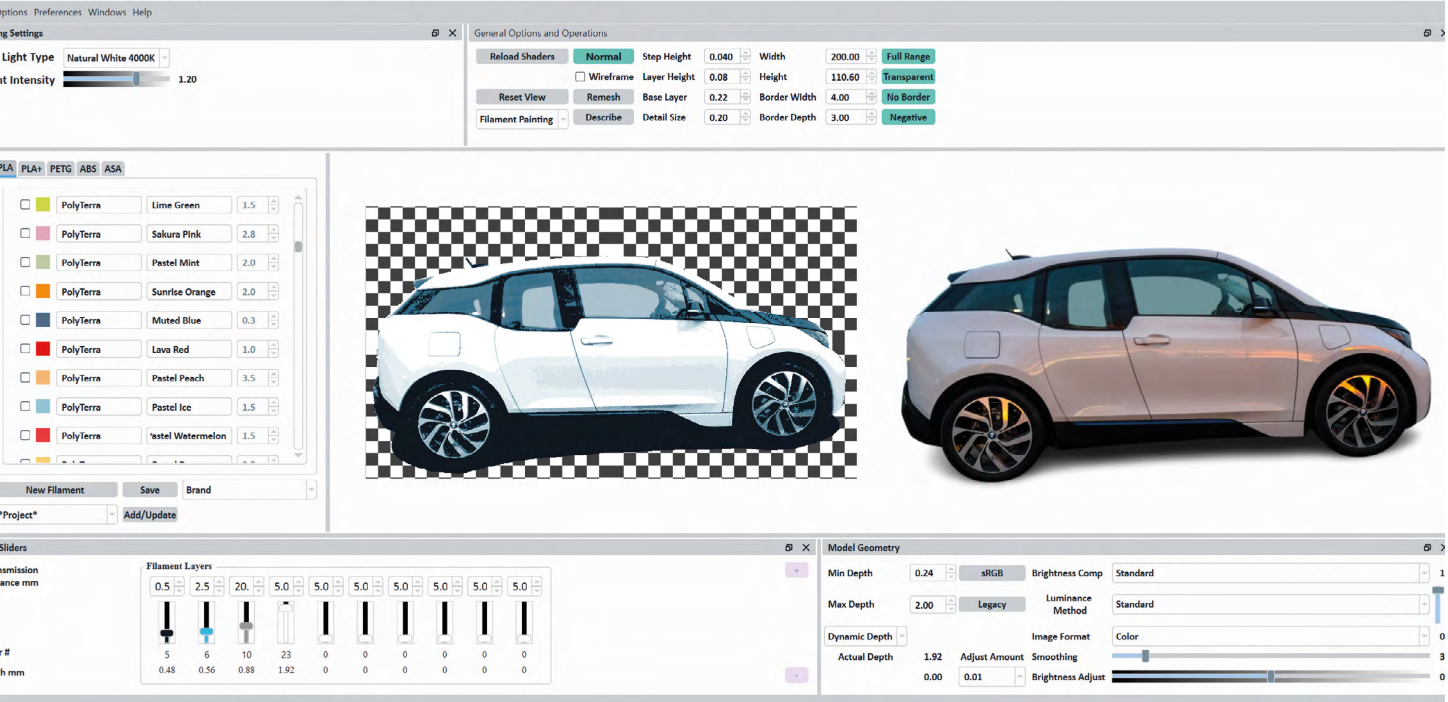

you could colour in features of an animal model. When the object is saved from the slicer, the colour changes are saved as well, and when the object is sliced, it will contain colour change commands to be used during printing.

Object painting is useful for adding colour to existing models, but it would also be nice to be able to take an image and prepare it to be 'rendered' by a 3D printer as an object. It turns out that there is a tool available to do just that, and it is called HueForge.

HUEFORGE

HueForge is a tool that takes 2D images and makes them into 3D-printable plaques. You can find it here: **shop.thehueforge.com**. **Figure 9** shows it being used to create a printable car picture. The program imports a colour image and maps colour or intensity values onto filaments that will be used to create a 3D-printable picture. You select the filament colours you want to use and drag them onto the sliders at the bottom left of the program. In **Figure 9**, the filament colours have been set to black, blue, grey, and white, which were loaded into the printer before the print. You can see these sliders at the bottom left of **Figure 9**.

You adjust the sliders for each filament to select the light intensity level in the source image that you want to map that filament colour onto.

As you can see above, the black filament has been mapped onto the darker colour, the blue onto a range of colours that match the blue parts of the image, and the grey and white filaments onto brighter parts. HueForge produces an STL file which describes the object to be printed. Before you slice the object, you need to configure the slicer to change the filament at particular layers. HueForge creates a text file which tells you what to do:

```
Swap Instructions:
        Start with Black
        At layer #5 (0.48mm) swap to Aqua Blue
        At layer #6 (0.56mm) swap to Silver Smoke
Grey
        At layer #10 (0.88mm) swap to Jade White
for the rest.
```

Above, you can see the layer colour instructions for the car picture. The first four layers are black, followed by blue, grey, and white. Bambu Studio allows you to select a layer and specify the filament colour for that layer.

You can have a lot of fun experimenting with different filament colours and adjusting the sliders to generate different 'looks' for your finished image. HueForge works even better if you spend some time preparing an image to be processed.

Figure 9
The background removal tool at 'remove.bg' was used to remove the car from the source picture

QUICK TIP

You might think that you would use red, green, and blue filaments for colour lithophanes, but, like printing on paper, the print will work by filtering out the colours we don't want rather than lighting up the colours we do.

The author of HueForge has made some useful videos exploring how to do this. You can find them here: **hsmag.cc/HueForgeFAQ**.

> **When adding colour to a print, remember that** filaments may have to be changed in many layers

Figure 10 shows the first HueForge print made by the author. It turned out very well and shows the level of detail that can be achieved. There is a very impressive 3D effect on the door handles and the panel lines on the car. With a bit more work, it is hoped that the blue elements could be made to stand out more.

It turns out that a 3D printer is not just for boxes and mechanisms but also has artistic merit too. **Figure 11** shows the kind of fun you can have. Modern printers can print large, flat objects much more quickly than previous ones, so you can turn out artwork quickly. The author hopes you have as much fun doing this as he has.

Figure 11
This print uses just three colours: black, white, and grey

Figure 10
The panel at the top right of the print is the 'purge tower'. This is used by the printer to make sure that all of a previous colour is removed from the print nozzle before it is used on the print.

QUICK TIP

By picking a source image with little detail and large areas of solid colour, you will find that you can produce good-looking colour lithophanes even with low-resolution printer settings.

HackSpace
TECHNOLOGY IN YOUR HANDS

FIELD TEST

HACK | MAKE | BUILD | CREATE

Hacker gear poked, prodded, taken apart, and investigated

BEST OF BREED

ONLY THE BEST

Machine learning and AI development boards

Artificial intelligence boards, kits, and accessories

By Marc de Vinck

Artificial Intelligence is all the rage these days. It's a very polarising subject for many people. There are those who are worried about losing their jobs to it, and then there are those people who look to the new jobs being created out of AI. To some extent, I think they are both correct. In my day job at Deep Sentinel, I work with developers every day on new AI models and technology. And although I'll never match the skills of my colleagues, it has sparked my interest in playing around with the technology and seeing if I can come up with a few fun projects for myself. And that's where this Best of Breed comes in. I'll be looking at some of the accessible technologies and accessories related to AI and edge computing. It's amazing how many products are out there that can help you make the plunge into AI. Some are expensive, and some are incredibly affordable. I've looked at AI products in the past, but since then there are many more options. So, let's look at some AI-related electronics to inspire you to join the revolution.

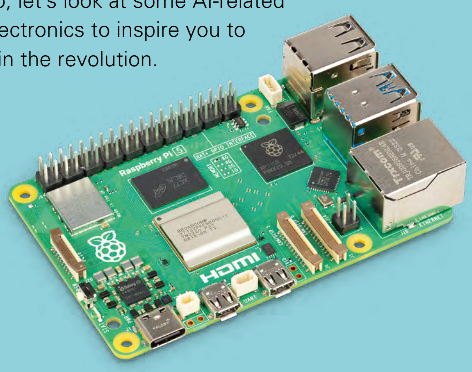

Left You don't have to have specialist hardware. A Raspberry Pi is a great starting point for a machine learning adventure

reTerminal CM4104032 – AI vs reComputer J2021-Edge AI Device

SEEEDSTUDIO ◆ $209 | seeedstudio.com

SEEEDSTUDIO ◆ $659 | seeedstudio.com

The reTerminal from Seeed Studio is a complete Raspberry Pi system in a nice enclosure just waiting to be integrated into your project. Powered by a Raspberry Pi Compute Module 4 and featuring an integrated multi-touch screen, it makes for a nice little system to get up and running quickly. The reTerminal also features Wi-Fi and Bluetooth, pre-installed Linux, 4GB RAM, and 32GB eMMC.

With all those features, and some good documentation, you can get TensorFlow Lite up and running with little effort. Now you can run object detection, image classifications, semantic segmentation, or a bunch of different facial recognition demos. This kit features a lot more, so for full details, be sure to head over to the product page to learn more about the details of what the reTerminal can do.

The reComputer J2021 is a powerful and compact intelligent edge computing box that brings up to 21 TOPS of modern AI performance. At its core is a Jetson Xavier NX 8GB Production Module featuring 384 Nvidia CUDA cores, packed into the diminutive form factor of 130 mm × 120 mm × 50 mm. In addition to the Jetson Xavier, there is an integrated carrier board featuring Gigabit Ethernet, four USB 3.1 ports, an HDMI port, and DP port. Having so much power all integrated together will help simplify your development process. If you are looking to deploy an edge AI solution around video analytics, object detection, natural language processing, medical imaging, or robotics, this is something you must consider. It's designed for more advanced users, and with all that power comes a higher price point than, say, some single-board computers, but you also get some amazing specs!

VERDICT

reTerminal CM4104032 - AI

An AI terminal ready to go.

10 /10

reComputer J2021-Edge AI Device

An all-in-one powerful AI processor for advanced users!

9 /10

BEST OF BREED

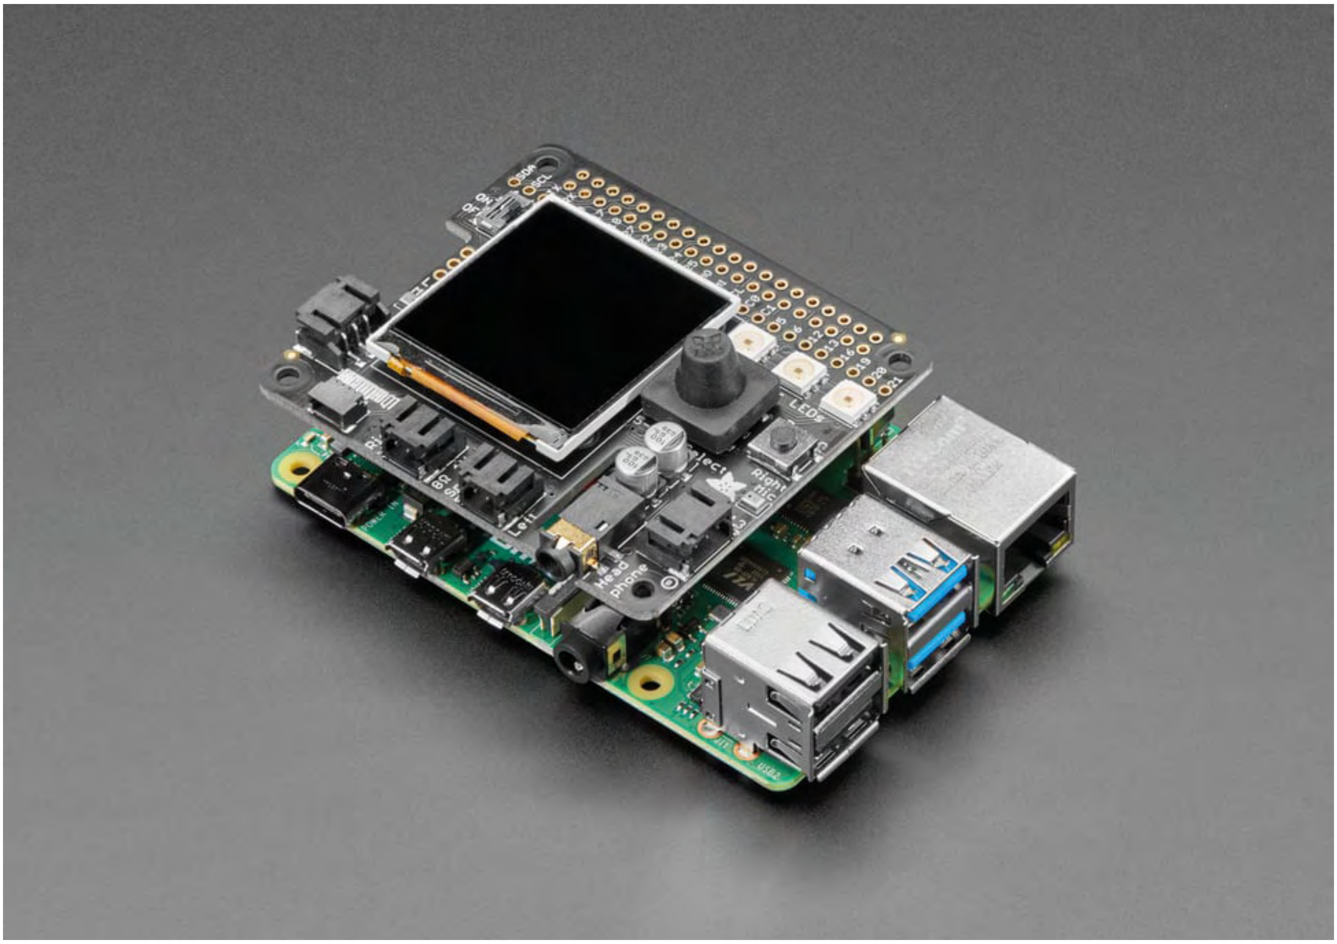

Adafruit BrainCraft HAT

ADAFRUIT $44.95 | adafruit.com

According to Adafruit, the BrainCraft HAT "allows you to "craft brains" for machine learning on the edge, with microcontrollers and microcomputers", and we really like that description. The board features a 240 × 240 TFT IPS display for inference output, a camera connection slot, joystick, button, left and right microphones, stereo headphone out, a 1-watt speaker output, and three RGB LEDs. And, to make building up your circuit simple without the need to solder, they even included two 3-pin STEMMA connectors. Just add your Raspberry Pi and you'll be off to the races. There are a few more well-thought-out features of the BrainCraft HAT, like the ability to add a fan under the HAT to keep your Raspberry Pi cool when running AI inferencing. Head over to the Adafruit website to learn more about this very capable little HAT.

VERDICT

Adafruit BrainCraft HAT

Brains for your Raspberry Pi.

9 /10

OpenMV Cam H7 R1

ADAFRUIT ◆ **$84.95** | **adafruit.com**

The OpenMV Cam, available at Adafruit, is a powerful microcontroller board coupled to a camera that is programmable via MicroPython. It allows for object recognition, face detection, and more, and all those features are embedded, meaning you don't need to fiddle with them – you can just access the data that the OpenMV Cam sends out. At its core is an STM32H743VI Arm Cortex-M7 processor running at 480MHz with 1MB of RAM and 2MB of flash. You also get a full speed connection to your computer, a microSD card slot, RGB LED, access to an SPI and I2C bus, 12-bit ADC and a 12-bit DAC, and even three I/O pins for servo control. It's a great camera with a lot of on-board power.

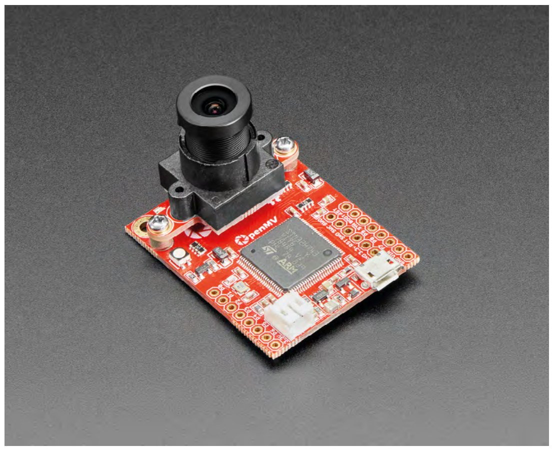

VERDICT

OpenMV Cam H7 R1

A powerful and smart camera module.

8 /10

SenseCAP A1101

SEEEDSTUDIO ◆ **$79** | **seeedstudio.com**

The SenseCAP A1101 – LoRaWAN Vision AI Sensor from Seeed Studio is a TinyML Edge AI-enabled smart image sensor. It can run a variety of different AI, allowing it to perform image recognition, people counting, target detection, and more. And for the more adventurous DIYers out there, you can also roll your own training model as it supports TensorFlow Lite.

The SenseCAP A1101 features an ultra-low-power Himax camera running at 60 fps at a resolution of 640 × 480. It can do local inferencing and transmit the data up to several miles under ideal conditions thanks to the integrated Wio-E5 LoRaWAN module.

It's an interesting device with a lot of good documentation and additional features. Be sure to head over to the website to learn more about this AI vision sensor.

VERDICT

SenseCAP A1101

Program your own smart camera.

9 /10

BEST OF BREED

SenseCAP K1100 – The Sensor Prototype Kit with LoRa and AI

SEEEDSTUDIO ◈ $49 | **seeedstudio.com**

Want to build an IoT project in only three minutes? That's what Seeed Studio claims you can do with its SenseCAP K1100 Sensor Prototype Kit. And after going over Seeed's tutorial, it does look like a simple project would be possible in about that time. This kit includes a vision sensor, temperature and humidity sensor, air sensor, and soil moisture sensor. Most importantly, it has a Wio Terminal for the brains, display, and as a central hub for connecting it all together. The Wio Terminal features an ATSAMD51-based microcontroller with both Bluetooth and Wi-Fi that is powered by a Realtek RTL8720DN module. The unit also has built-in 2.4" LCD screen, on-board IMU, microphone, buzzer, light sensor, infrared emitter, and microSD card slot. It's a robust kit that allows you to do a lot of remote monitoring and run simple object detection. It's also compatible with the 400+ modules of Seeed Studios' Grove ecosystem, which enables you to add an almost endless array of sensors and add-ons.

VERDICT

SenseCAP K1100

A great place to start with AI.

10 /10

AI & MACHINE LEARNING

RASPBERRY PI FOUNDATION ◈ FREE | **raspberrypi.org**

Now, here is something that isn't typically found in my Best of Breed roundups. A great product, and it's free! The Raspberry Pi Foundation has a nice Introduction to Machine Learning programme that is free to take. To complete the course, you should be able to dedicate 2–4 hours per week, for about four weeks. It's a self-paced programme that dives into various machine learning models and their applications. If you are looking to get started, be sure to head over to the Raspberry Pi Foundation website and check out this course.

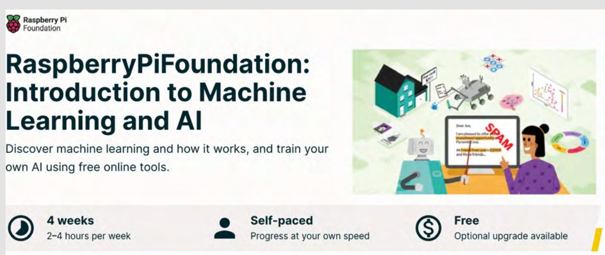

xTool S1

A little laser with a lot of features

XTOOL ◆ **£1799** | **hsmag.cc/xtools1**

By **Ben Everard**

Laser cutters are great tools for makers. They let you cut shapes out of – or engrave images into – sheet material quickly and accurately. By piecing together different shapes, you can build up 3D objects much faster than you can with a 3D printer.

If you're after a laser cutter, there are a few things to consider. Let's take a look at some of the key things you have to think about, and take a look at how the xTool S1 stacks up.

The first big decision to make is safety. Many budget machines aimed at makers are unenclosed. That means there is nothing to stop the laser beam from being reflected back towards the operator. Exposed lasers are dangerous. Very dangerous. A beam powerful enough to cut wood is easily powerful enough to blind you instantly. Yes, you can wear laser goggles, but this isn't a good solution, and when your sight (and the sight of anyone else who might happen to open a door) is at stake, then we want the best possible protection.

Below ◆ The enclosure keeps fumes and laser beams contained

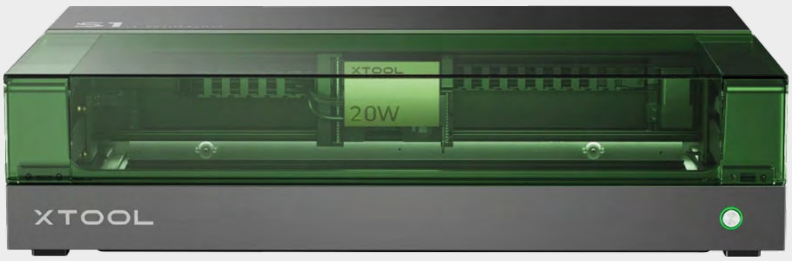

No tool is perfectly safe, but the xTool S1 comes packed with a lot of safety features. It's fully enclosed, so it should be impossible for any lasers to leak out. If you open the door while it's in use, it'll automatically stop.

FLAME OFF

Another significant safety risk when using a laser cutter is fire. The S1 has a flame detector which will stop it if what you're cutting catches fire. You can also plumb in a fire extinguishing system, but we didn't have a suitable system to test this with. Finally, you need to consider the fumes given off by the material as it burns or vaporises away. Being fully enclosed, the S1 catches the fumes but it will need to vent them. There's a hose that you can poke out of a window or plumb into an extraction system. If this isn't possible, an air filter is available separately.

Once you've looked at safety, you might want to consider power. Laser power isn't completely comparable between different systems: a diode laser will cut materials differently to a CO_2 laser of the same wattage; different wavelengths of light are absorbed by materials differently; and a laser cutter that requires mirrors will lose some power as the light is reflected.

The xTool S1 has a choice of three different power options: 2 W infrared laser, and 20 W and 40 W blue laser modules are available. We tested the 20 W version and found that it worked well with wood (both MDF and construction plywood) up to about 10 mm

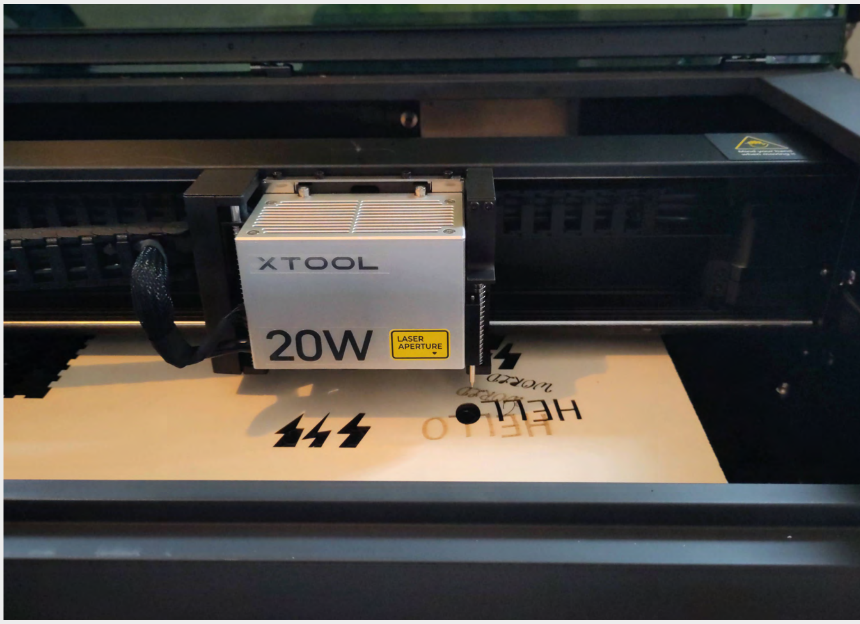

Above Left This pin is used to autofocus the laser

Above Right You can quickly and easily swap between different laser modules

thick, and worked excellently with laser ply. This cutting power is comparable with CO_2 lasers that we've used. For wood and plastics, the S1 should serve hobbyists well.

Perhaps the biggest selling point of the S1 is its ease of use, and on this, it really is excellent. There's autofocusing (with adjustable depth offset for cutting through thick materials) and a system for locating your objects on the print bed. Both of these work in slightly unusual ways. The autofocusing works with a rod attached to the side of the laser. This then probes the surface. Once it's probed, the print head whizzes off to the side to push it up and out of the way so it doesn't interfere with any of the engraving. Should the probe hit something, it's magnetically attached, so it can pop off without damage.

LASER LOCATOR

The system for locating your items uses a visible laser that's in the same place as the laser cutter. You physically place the laser head in one point, then press the button to mark the point, then move it to another point, press the button, and so on until you've marked out the necessary points. The software then converts these points to a rectangle, circle, polygon, or other shape that lets you place your work in the correct position relative to this shape.

All of this needs software to support it, and in the case of the S1, that's xTool Creative Space.

Our only gripe with the software is that it needs an internet connection to work. This is the case even when connected to the laser cutter over USB (you can connect over Wi-Fi). We can't see why this is the case, and it does make us concerned about what would happen if xTool went out of business, or shut down this particular service.

xTool Creative Space has a few features, both useful and odd. It can automatically generate QR codes for you. It also has an AI system for generating artwork that we're still getting our heads around. On this, you have to pay for each set of generated images using credits. You get a set of credits when you create your account, but after that, you have to generate more by sharing projects on the xTool website.

The S1 is also compatible with LightBurn, but not all of the features are available in this software.

Overall, we've been really impressed with the xTool S1. However, there are two areas that are potential pitfalls. The first is size. It can only cut 498 × 319 mm. You can fit material up to 600 × 400 mm, which might be useful if you don't have a good saw for cutting up sheet material, but you won't be able to cut against the edge of the sheet which, depending on what you're cutting, might mean material waste. The other thing we can't ignore is the price. The setup we tested cost £1799. You can get laser cutters cheaper than this, and you can get bigger, more powerful laser cutters for around the same price. However, given the full range of what's on offer with this machine – the safety features, ease of use, and the power – we think it represents excellent value.

VERDICT

An easy-to-use, fully enclosed laser cutter at a reasonable price.

9 /10

Prusa XL

A big printer with big potential

PRUSA RESEARCH £1798.80 | prusa3d.com

By **Ben Everard**

Below This print is based on a CT scan of a wrist. The bones are printed in PLA and they're joined by sections of flexible filament. The result is a print that moves like a real wrist

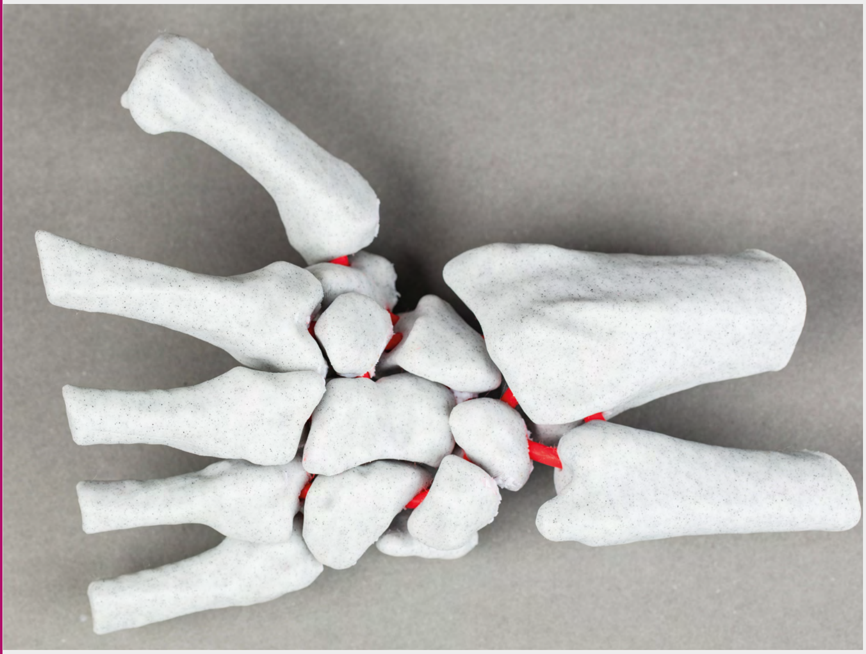

The XL is Prusa's first large format core XY printer, and this alone would be enough to make this a special printer. However, the XL has a much bigger feature tucked away at the back of the bed: a five-tool tool changer.

Let's take a step back and look at what a tool changer is in general, before looking at this printer in more detail.

The vast majority of multicolour 3D printers work by loading and unloading filament into one nozzle. This works to an extent, but it has a few limitations: it's slow, wastes a lot of filament, and only works well if all the filament is the same type of plastic. You can also get printers with more than one extruder that moves about the bed on the gantry. Sometimes these extruders are fixed together and sometimes they can move independently, in which case they're known as independent dual extruder or IDEX. These can print different filaments but, again, they have limitations. In this case, it's that you're mostly limited to two extruders because otherwise it gets very crowded. The final option for multi-material printing is tool changers. With these, there are multiple extruders, but only one extruder is engaged by the moving gantry at a time – the others are parked out of the way. The gantry should be able to quickly and automatically change to a different extruder. On paper, these are the most powerful option because they can print more different filaments faster than the alternatives. However, historically, they've been the least common because they are significantly more complex than the other options.

Let's get back to the Prusa XL, then. You can load it with up to five extruders (also known as tool heads).

One problem testing out this printer is that it has capabilities that no other printer on the market has, so you can get caught between two problems. Firstly, it's very easy to get caught up in hype about things that later turn out to be just gimmicks. And

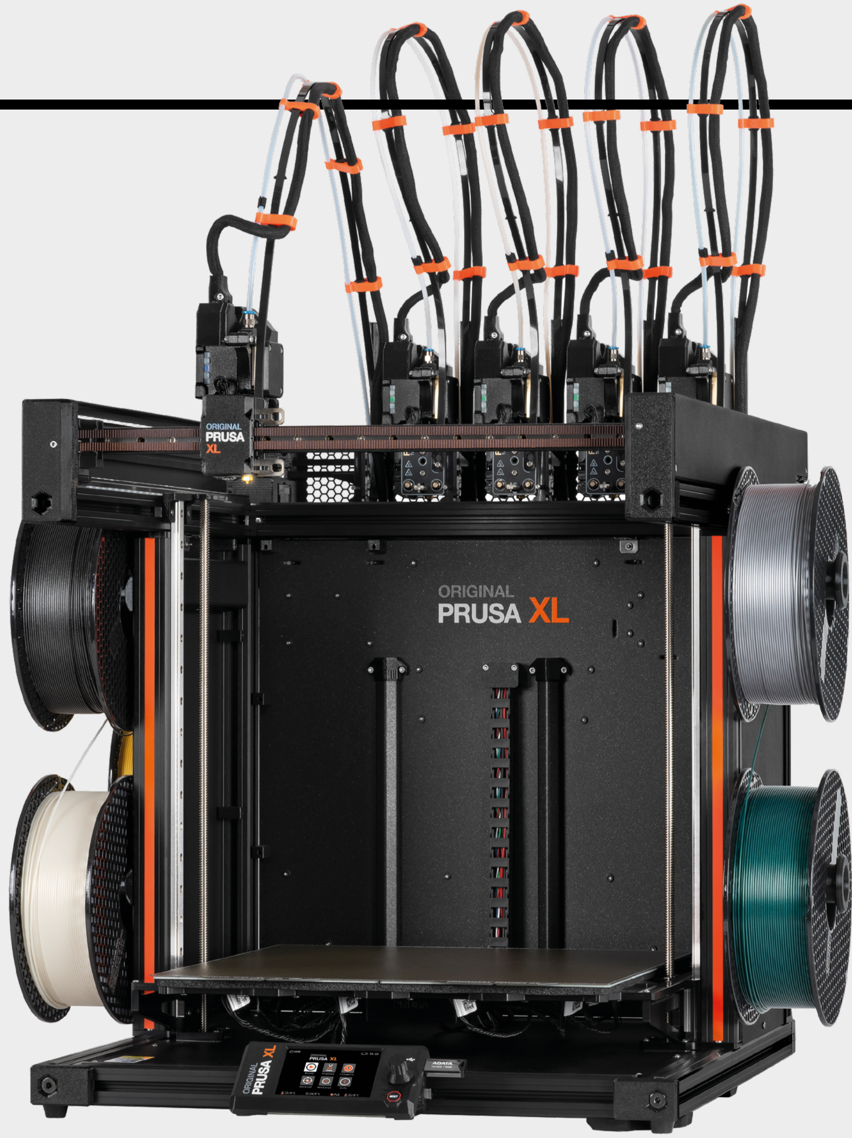

secondly, it can be hard to see the full value in features until a wide range of people have had the chance to flex their creative muscles and really find great uses for them. We'll do our best to steer a path between these two issues.

The five extruders can each be loaded up with a different filament. These can be different colours of the same filament, in which case they can print multicolour objects in a very similar way to other multicolour printers. However, it can do more. Because each extruder is entirely separate, you can use different types of plastic. There are some limitations – for example, the print bed will be the same temperature, which might cause a problem for some mixtures of filament.

The XL can mix rigid plastics, such as PETG, with softer plastics like TPU. We got a couple of test models from Prusa with which to try this out. One was a print-in-place box that included a flexible, water-resistant seal. The other was a CT scan of someone's wrist with the bones printed in PLA and the bones joined with flexible tendon-like rods printed in TPE. Both of these prints printed excellently, and have a wow factor that, honestly, we've not seen in 3D prints for a long time.

> **It is probably the most expensive printer** we'd still consider fitting in the 'prosumer' hobbyist category

While the test prints designed by Prusa are undeniably impressive, we wanted to see how hard it was to conceive and design something ourselves that made use of this mixture of materials. The two things that we came up with are a keyboard support with integrated rubber feet, and a phone case. Both turned out well, though we would like some more powerful tools in PrusaSlicer for placing filament. →

Above You need quite a bit of desk space for the XL – the spools mount on the sides, and you need to be able to reach the top of the printer to access the tool heads

> **The Prusa XL can print our multicolour cube about** eight times faster than the Bambu Lab X1-Carbon

The keyboard support could, we'll be honest, have been a print entirely in rigid plastic with rubber feet glued on. This would have worked perfectly well, but we wouldn't have been able to get feet in the shape of a lizard this way.

The phone case is another matter. We've tried 3D-printed phone cases a few times on other printers and never had much success. We've found that any filament that is flexible enough to provide some impact resistance is also stretchy enough to fall off the phone. We combined rigid and flexible material so that there was a sturdy frame around which we put soft flexible filament (40D on the Shore hardness scale). The result is a case that fits securely and snugly, while still providing a lot of impact resistance. We think there's still a little way to go before we've nailed this case design, but it's already by far the best 3D-printed phone case that we have had.

Above
Multicolour prints use far less material than traditional colour-changers

Another option is to mix different materials for supports and the main print. This can be as simple as using a cheaper material for supports – or more advanced like using two plastics that don't stick, to create supports that are easier to remove (such as PETG and PLA). There's even the option of using a soluble material such as PVA for supports that just wash away. This isn't currently well-supported by the slicer, as it causes a problem we don't fully understand with the wipe tower. Hopefully, a future version of PrusaSlicer can help with this.

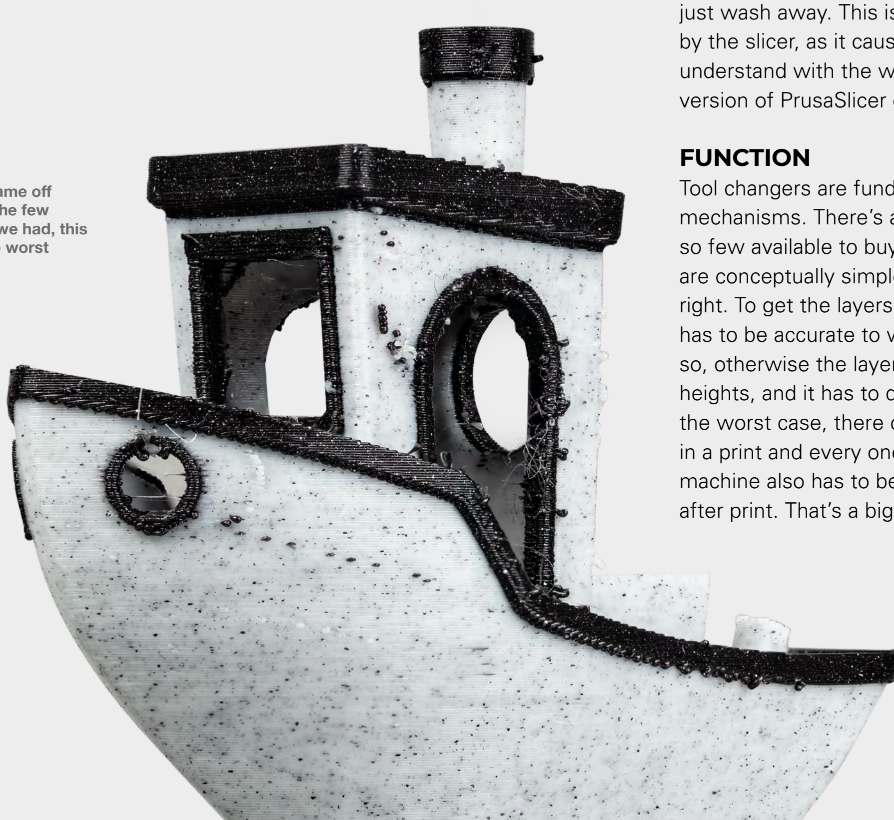

Right
Most prints came off clean, but of the few messy prints we had, this was by far the worst

FUNCTION

Tool changers are fundamentally complex mechanisms. There's a reason that there have been so few available to buy over the years. While they are conceptually simple, there is a lot of detail to get right. To get the layers consistent, each tool change has to be accurate to within a tenth of a millimetre or so, otherwise the layers will be at slightly different heights, and it has to do this over and over again. In the worst case, there could be 20,000 tool changes in a print and every one has to be spot on. The machine also has to be able to keep doing it print after print. That's a big ask.

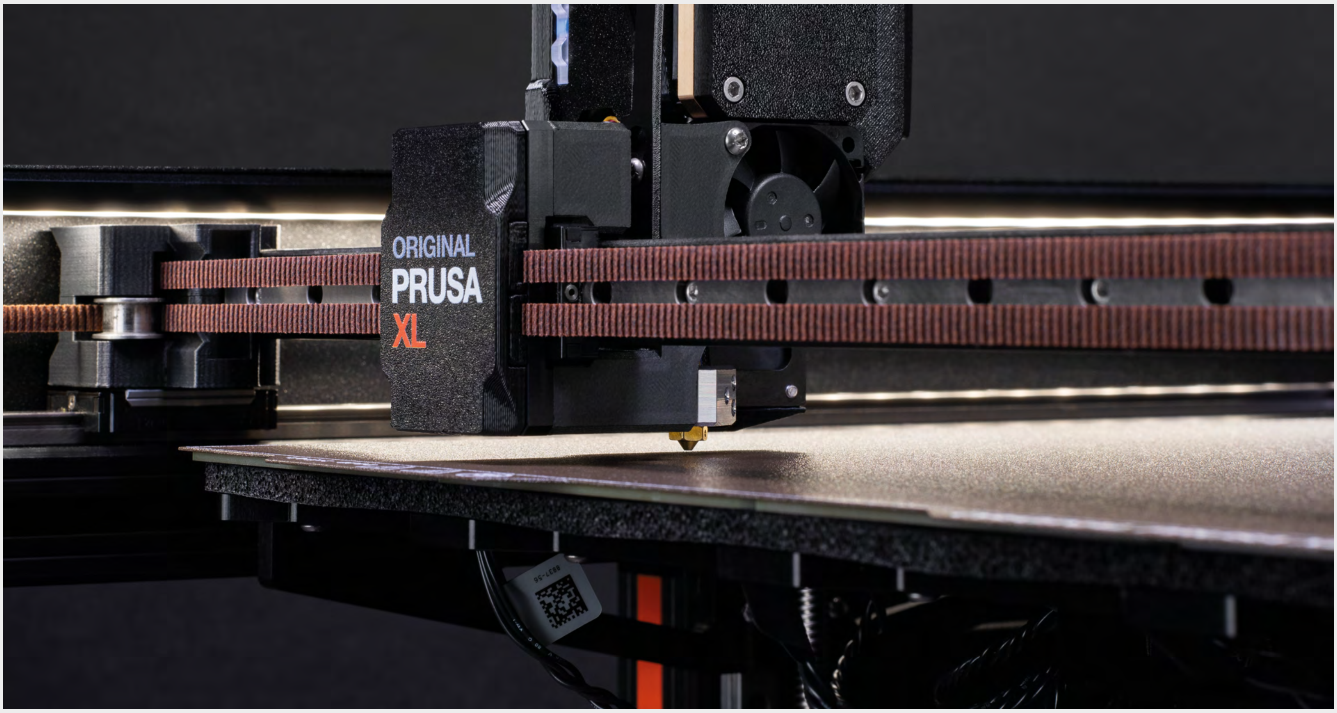

Above The Core XY motion system has two drive belts, both of which are needed to move the print head in either the X or Y direction

Initially, we did have a little difficulty with the different extruders having slightly mismatched Z-heights, which resulted in some filaments being too squashed into the print bed on the first layer, and others being too far away. We followed the calibration guide and still had the problem. We solved it by both lubricating the tool changer, and unloading the filament before doing the tool head calibration.

The only other issue we've had is getting some large blobs of filament protruding above the top of the wipe tower. This seems to happen when the extruder primes but doesn't move, so squirts a blob in place. The next time it uses the wipe tower, it then has to bump over this blob. It makes a bit of a clatter as it does so. At first, we were so paranoid about this that we kept checking on the printer and snipping off any big blobs with a pair of side cutters. However, after a while, we just got used to them being there. It looks pretty horrible, but it's on the wipe tower and doesn't seem to cause any problems.

SPEED

It felt like 2023 was the year that 3D printing became fast. Almost every 3D printer was judged on its ability to create a Benchy in minimal time. The XL is actually a really hard printer to judge on speed. When it comes to straight one-filament speed, it's reasonable, but not spectacular. There's a lot working against it. Frames with larger print volumes inevitably have less rigidity, and the tool changer both increases the mass of the print head and places it further away from the axis. The end result is a single-colour speed that's slightly slower than the MK4.

A Benchy is probably about the worst case for speed on the Prusa XL, as it contains a lot of small moves that benefit smaller machines that won't flex as much under high accelerations. The best case is something that has a lot of colours on every layer because the tool changer can switch filaments far quicker than most other printers.

For an example of this, we took a one-inch cube and coloured each face differently (other than the top and bottom). The Prusa MK4 isn't yet compatible with the colour changer (though it is in the works), so we'll use the Bambu Lab X1-Carbon as a comparison. This printer is no slouch and is probably the fastest off-the-shelf hobbyist printer. However, the XL can print our multicolour cube over eight times faster than the X1 with the same layer height. This is a phenomenal speed-up, but even this is understating it – we limited it to four colours because that's as many as the X1 can do. The XL can do an additional colour.

The only honest answer we can give about speed on the XL, compared to other fast high-end →

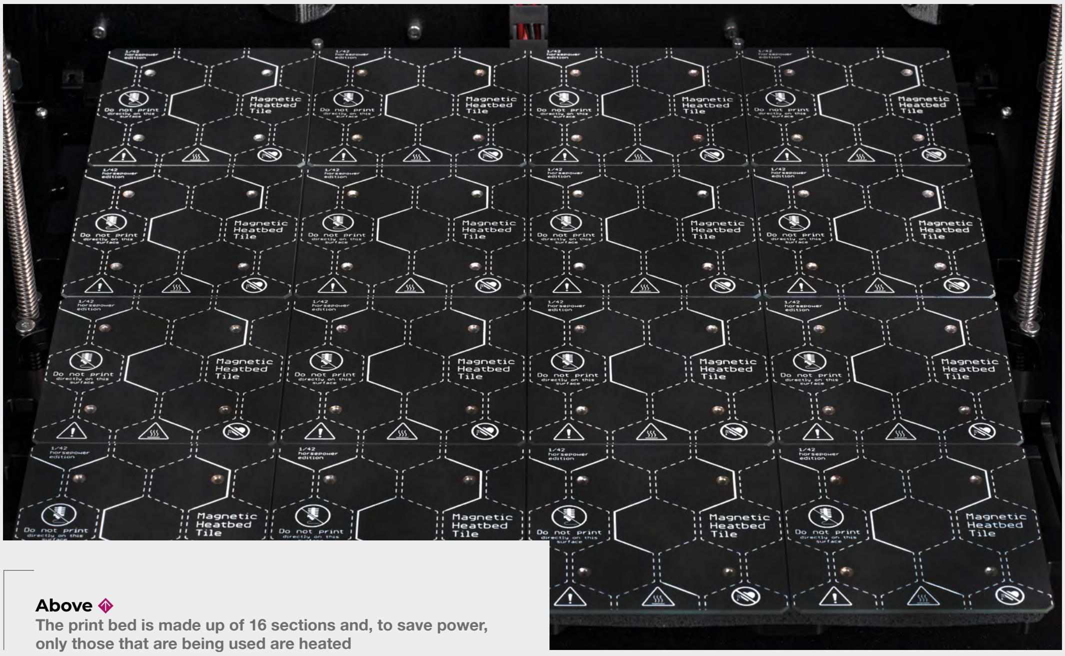

Above
The print bed is made up of 16 sections and, to save power, only those that are being used are heated

hobbyist machines, is that it's somewhere between slightly slower and eight or more times faster, depending on what you want to create. Therefore, if you want a printer for speed alone, it's really important to look at the actual models you intend to print, not just isolated benchmarks cooked up by reviewers. However, we can say that if you do plan on printing with multiple colours, you are likely to get a significant speed-up, particularly if there are a lot of colours on each layer.

Another significant speed-related issue is the ability to print lots of things at the same time. While the time to print one Benchy on the Prusa XL is about the same as printing it on the MK4, you can fit nearly three times as many on the print plate. This might seem a bit surprising, because a 36 by 36 cm print bed doesn't seem that much larger than a 25 by 21 cm print bed, but it really is. The speed-up here isn't from the printer taking less time per Benchy, but from you, the user, not having to come and take things off the print bed as often.

Obviously, this does come with the risk of one failed object on the print bed damaging all the others, but you can cancel an individual item within the print without having to cancel the entire lot.

QUALITY

The quality of prints we've had from the Prusa XL is good, but not perfect. It's not uncommon for multicolour prints to get a bit of the wrong colour filament in the wrong place. This is due to a bit of stringing, or a blob from the wipe tower, being picked up by the nozzle. It doesn't happen much, and most of the time it's easy to clean up.

SOFTWARE SUPPORT

Overall, we'd say that the Prusa XL is as well-supported by the slicing software as most existing printers. However, the tool changer opens up some whole new possibilities that we're only just getting our heads around. For example, embedding text or a logo in one colour onto a surface of another colour looks great, and you can do this in PrusaSlicer, but it's a bit clunky as it involves creating a new part and merging them. It'd be far easier if you could just do this in the same way you can emboss text or SVGs onto a part.

We also had an issue in one print where we used the colour paint tool to apply colour to the surface of an object but, on a curved surface, the colour didn't go deep enough and the underneath colour shone through a light-coloured upper layer.

Another minor problem encountered was when an object is made of multiple different parts in CAD and then merged, it can become impossible to properly select some of the parts to paint them.

In all these cases, the problem can be solved by going back to the original CAD files and manipulating them there before going into the slicer. However, we've gradually moved away from CAD to the

PrusaSlicer as it has been upgraded with more and more powerful tools for manipulating 3D shapes, and we'd like to keep moving in this direction.

These are all relatively minor issues, but we'd hope that the colour and filament selection options in PrusaSlicer get some attention soon, as it feels like there are some improvements that could really help people get to grips with this printer.

> **We did have a little difficulty with the** different extruders having slightly mismatched Z-heights

WASTE

We can't ignore the fact that 3D printers use plastic, and plastic pollution is a major source of pollution for the planet. Despite some claims to the contrary, PLA is not biodegradable in any realistic way. It is possible to recycle it (and other 3D printer filaments), but it's not particularly easy and few recycling centres will accept it. That doesn't mean we can't enjoy 3D printing as a hobby, but we need to think about our impact on the world.

The Prusa XL uses dramatically less plastic when printing in multiple colours than traditional multicolour printers. The exact amount will depend on the model, but it can easily end up using less than a quarter of the filament, especially for small models and those with lots of colours. Potentially the savings are even greater if you don't use the wipe tower, (though this option relies on having very dry filament). Not only does this save plastic, but it also saves you money. A heavy user of multi-material prints could easily find that the XL is actually a cheaper option compared to a classic multicolour printer, when considering the saving in filament over the life of the printer.

The Prusa XL is a hard printer to review, because there's so little to compare it against, especially now the E3D tool changer has been discontinued. The large print volume and weight of the tool changer slow it down, so it's not the best for large, single-colour prints, and it is probably the most expensive printer we'd still consider fitting in the 'prosumer' hobbyist category. However, it can create prints that no other printer can. This will probably prove true in ways we can't yet predict as people get their hands on the machine and start experimenting with creative ways to use the capabilities. On top of this, for multicolour prints, it's massively faster than the competition, and the lack of waste is both an environmental win and a money saver. This isn't a printer for everyone, but for some people, it's absolutely the perfect printer. □

Above The Prusa XL has the same control interface as the Prusa MK4. We found it reliable and easy to use

VERDICT

Pushing the boundaries of what 3D printers can do.

10 /10

REGULAR

CROWDFUNDING NOW

BullseyeBore Core

Perpendicular drilling every time

From $129 | **hsmag.cc/bebore** | **Delivery:** July 2024

The problem this is attempting to solve is certainly a real one – to ensure that when you're drilling a hole with a handheld power drill, you'll drill it straight. Although we've not been able to test it out, it looks like it could solve it well – it projects a set of circles onto the workpiece and, as long as those circles are concentric, you're going straight.

Our concern with this is in one area only – the price. The cheapest option at the time of writing was $129 (not including shipping). Perhaps other people have more difficulty than we do getting straight enough holes, but that seems so far beyond what we'd pay to solve this problem.

It would, perhaps, be justifiable if it were the only solution to the problem, but it's not. Search for 'Drill Guide' online, and you'll see a raft of potential solutions, from blocks with holes that you can slot the drill through to frames that allow a drill to slide through at 90 degrees. Most of these are far cheaper and also help you place and hold the bit accurately to minimise wandering. There are even some 3D-printable designs that you can run off for just a few grams of filament.

Perhaps we're just too tight, but this looks like an interesting product that is made irrelevant by its price. □

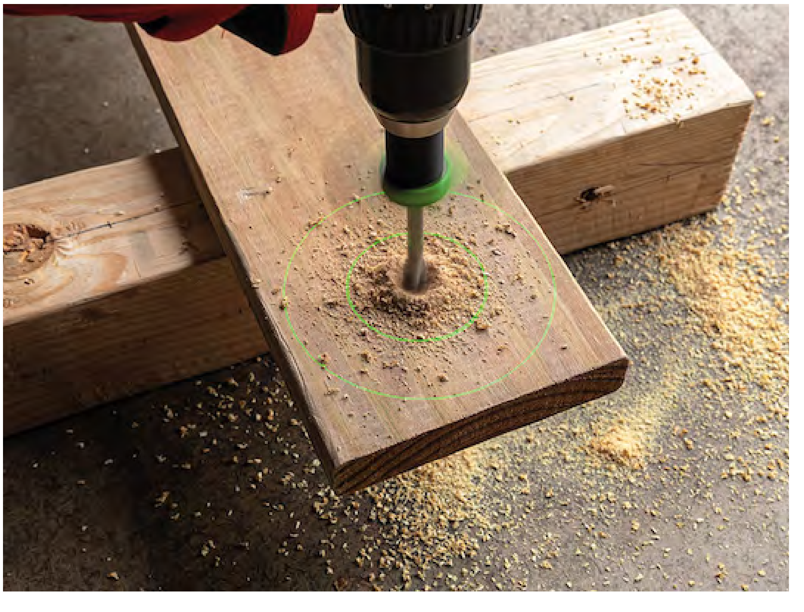

BUYER BEWARE

When backing a crowdfunding campaign, you are not purchasing a finished product, but supporting a project working on something new. There is a very real chance that the product will never ship and you'll lose your money. It's a great way to support projects you like and get some cheap hardware in the process, but if you use it purely as a chance to snag cheap stuff, you may find that you get burned.

Print-a-Kit

3D-print your own lessons

From $74 | hsmag.cc/pkit | **Delivery:** Mar 2024

Do you have a 3D printer? Do you want to learn about engineering? Then this could be for you. For about $74, this Kickstarter will provide five sets of engineering lessons. Each lesson includes 3D-printable models (and a 3D-printable box to store them in) and videos to help you understand what's going on.

While it's hard for us to properly evaluate this as we've not seen the actual models or the videos, we really like the idea behind it. 3D-printable objects can be excellent learning tools. Provided you already have a 3D printer, they're cheap and easy to create. This writer has two young children and knows first-hand how physical objects can capture people's imagination more than videos and pictures.

With 3D printers becoming more popular every year, hopefully we'll see more initiatives like this that help people gain more value from their 3D printers.

next month

issue #77

ON SALE **28 MARCH**

DIY

Keyboards and controllers

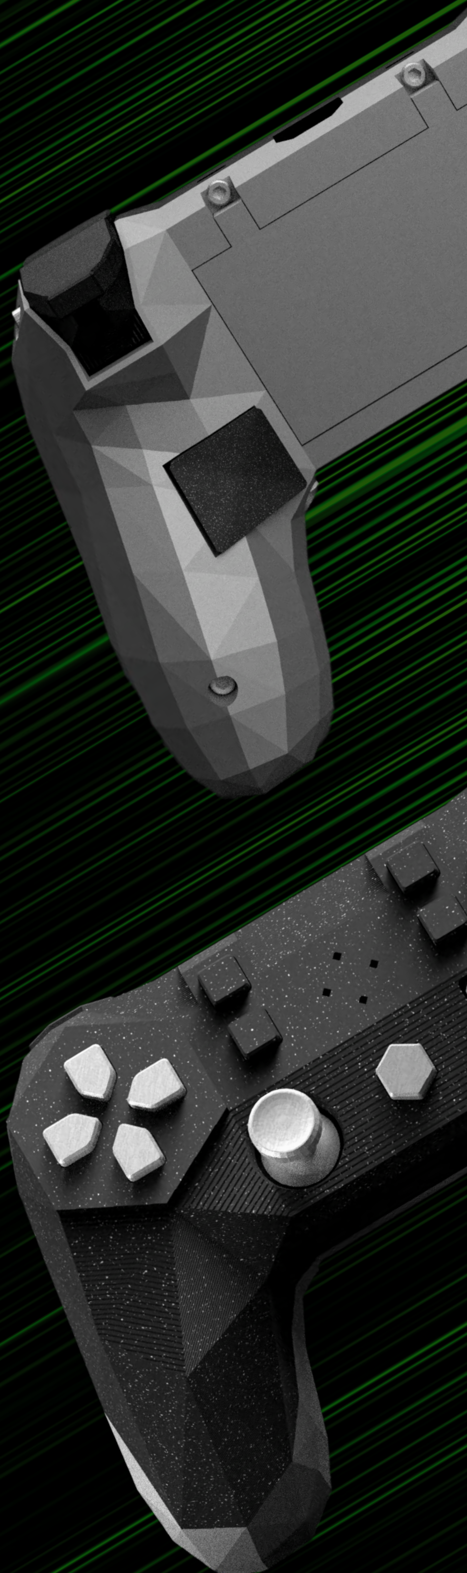

ALSO

- 3D PRINTING
- MUSIC
- SCREEN PRINTING
- TINY COMPUTERS

AND MUCH MORE

DON'T MISS OUT

hsmag.cc/subscribe

Lithophanes

We might be living in the 21st century, but our brains are still basically the same as they were when we chased prey across the savannah. Our squelchy bundle of neurons reacts differently to things that it can poke and prod.

3D printing makes physical objects: things you can hold and touch, and these give us a different reaction to the same image on a screen or on paper. We understand them in a different way, and feel differently about them. This issue, we take a look at making pictures with a 3D printer. In some ways, these pictures are much lower quality than even a basic home printer can manage – there are few colours and the resolution's terrible. However, in another way, they're far better because we can feel them and experience them.

If you've got access to a 3D printer, give it a go and see how it feels to touch an image, not just see it.

PiKVM

Remote control redefined

Manage your servers or PCs remotely!

POWERED BY Raspberry Pi

PiKVM V4 Mini

Small, cost-effective, and powerful!

- Power consumption in idle mode: just 2.67 Watts!
- Transfer your mouse and keyboard actions
- Access to all configuration settings like UEFI/BIOS
- Capture video signal up to 1920x1200@60 Hz
- Take full control of a remote PC's power

PiKVM V4 Plus

The most feature-rich edition

- More connectivity
- Extra storage via internal USB 3.0
- Upgraded powering options
- More physical security features
- Extra HDMI output
- Advanced cooling solution

A cost-effective solution for data-centers, IT departments or remote machines!

Available at the main Raspberry Pi resellers

core ELECTRONICS · PiShop.ca · RPishop.cz · Welectron. · elektorstore · PIMORONI
Pi-Shop.ch · The PiHut · PiShop.us · CloudFree · KUBII · Cytron Marketplace
ROBU.IN Your ideas, Our Parts · NETTOP COMPUTERS & ELECTRONICS · netsket mall · DIREKTRONIK · RaspberryPi.dk

HiPi.io

No reseller in your country?
Check shop.hipi.io (import fees might apply).

List of official resellers by country: