THE WORLD'S BEST RECORDING TECHNOLOGY MAGAZINE — 1985 – 2024

SOUND ON SOUND™

MUSIC PRODUCTION TECHNIQUES / INDEPENDENT IN-DEPTH PRODUCT TESTS / ENGINEER & PRODUCER INTERVIEWS / LIVE SOUND

ABLETON MOVE

LIVE — TO GO!

Juan Atkins

Inventing techno

Charlie Handsome

From rap producer to Nashville's hottest songwriter

Exclusive: SSL 2 & 2+ MkII

Next-gen desktop interfaces

Heritage Audio i73 Pro Edge

USB interface with vintage-style preamps

WIN LEWITT AUDIO BUNDLES

WORTH $1380

www.soundonsound.com

11

USA $11.99 / Canada C$12.75

ON TEST: CHORD / SPL / ARTURIA / TANTRUM / SONY / ELEKTRON / ADAM / LYNX / ERICA / KALI / BABY AUDIO / ART PRO

TECHNIQUE: SOURCE-DESTINATION EDITING IN REAPER / DAW WORKSHOPS

LIMITING FACTORS

Browse any music production forum, and you'll encounter posts along the following lines. "I dug out some old recordings from 30 years ago. All I had back then was a cheap drum machine / four-track cassette recorder / some Sellotape with iron filings stuck to it. But you know what? They sounded really good! I wish I could get that sound today!"

Limitations have also helped to shape entire genres, as Juan Atkins explains in this month's Classic Tracks. Would early house and techno have been so minimalist if its producers had had endless polyphony on tap?

It's tempting to conclude that embracing limitations is always a good thing, but I think the truth is more complicated. The warm feeling we get when looking back to work we did on limited rigs is the knowledge that we did the best we could with the resources available. And we can forgive ourselves if the toms didn't quite come through properly, or the samples were too crunchy — because it wasn't our fault.

Now that software gives us practically infinite resources, this confidence has been taken from us. If there's even the slightest flaw in what we produce, it's on us, because we surely did have the means to correct it. We can't blame the equipment anymore. Nor, when infinite good advice is only a Google search away, can we plead ignorance.

When we smile at the recordings we made 30 years ago, that's partly because we only need to take responsibility for their good points. When we wince at the mixes we delivered last month, it's because we know that any issues could have been avoided, if only we'd done yet another round of 'mastering', or added another few nodes to the automation on the lead vocal.

So what's the answer? Well, there's nothing to stop us chucking everything out of our studios and working only with a Boss DR-55 and a Portastudio today. But self-imposed limitations are contrived. The decision as to what to chuck out is just another creative option, not something forced upon us by circumstances, so it doesn't absolve us of responsibility for the consequences in the same way.

But the key point is not that limitations are good in and of themselves. It's that they force us to get the best from the kit we have. And no matter how much gear we own, there's always more we can do to get to know it. All synths are better when we learn how to program them. All mics become more versatile once we try them in multiple positions and polar patterns. And the thing that makes all of us better producers is knowing our own limitations.

Sam Inglis
Editor In Chief

"The key point is not that limitations are good in and of themselves. It's that they force us to get the best from the kit we do have."

SOUND ON SOUND

SOUND ON SOUND LTD (HEAD OFFICE)
ALLIA BUSINESS CENTRE
KING'S HEDGES ROAD
CAMBRIDGE, CB4 2HY, UK
T +44 (0)1223 851658
sos@soundonsound.com
www.soundonsound.com

WWW.SOUNDONSOUND.COM/SUBSCRIBE

WORLDWIDE EDITIONS

NORTH AMERICA

UK/WORLD

EDITORIAL
sos.feedback@soundonsound.com
Editorial Director Dave Lockwood
Executive Editor Paul White
Editor In Chief Sam Inglis
Technical Editor Hugh Robjohns
Reviews Editor David Glasper
Reviews Editor Matt Houghton
Reviews Editor Chris Korff
Production Editor Chris Korff
News Editor Luke Wood

SUBSCRIPTIONS
subscribe@soundonsound.com
www.soundonsound.com/subscribe
Circulation Manager Luci Harper
Administrator Nathalie Balzano

PRODUCTION
graphics@soundonsound.com
Production Manager Michael Groves
Designer Mick Reilly
Designer Andy Baldwin

ADMINISTRATION
admin@soundonsound.com
Managing Director/Chairman Ian Gilby
Editorial Director Dave Lockwood
Marketing Director Paul Gilby
Finance Manager Keith Werthmann

ONLINE
support@soundonsound.com
Digital Media Director Paul Gilby
Design Andy Baldwin
Web Content Editor Callum Hall
Web Editor Adam Bull
Podcast Production Manager Atheen Spencer
www.soundonsound.com
twitter.com/soundonsoundmag
facebook.com/soundonsoundmag
instagram.com/soundonsoundmag

DISTRIBUTION
distribution@soundonsound.com
International Distribution
Magazine Heaven Direct
www.magazineheavendirect.com

ADVERTISING
usadsales@soundonsound.com
North America Sales Manager Dan Brown
Regional Sales Manager David Carson
UK Media Sales Manager Guy Meredith

MARKETING
marketing@soundonsound.com
International Business Development
Nick Humbert

Printed in the USA
Not for re-sale outside North America

ISSN 1473-5326
A Member of the SOS Publications Group


3RD WAVE
ADVANCED WAVETABLE SYNTH
"Best Flagship Synth 2023"
- Alex Ball/Bo Beats
"Best Hardware Synth 2023"
- Attack Magazine
"Best New Synth 2023"
- Gear News
"The 3rd Wave might be my favorite synth released in the last 10 years."
- Sound on Sound
"One of the most powerful synths you could imagine."
- Music Radar
GROOVE
SYNTHESIS

BOUNDLESS INSPIRATION | VINTAGE ROOTS

24 voices · 3 osc per voice · 4 multi-parts · analog & digital filters

keyboard and desktop

g r o o v e s y n t h e s i s . c o m

IN THIS ISSUE

www.soundonsound.com

November 2024 / issue 1 / **volume 40**

FEATURES

82 POSITIVE GRID SPARK LIVE

ON TEST

WORKSHOPS

SSL 2 & 2+ MkII

USB Audio Interfaces

The buttons on each input channel now have associated LEDs, and the third of them now introduces a high-pass filter.

SSL's desktop interfaces get a well-deserved overhaul.

SAM INGLIS

When Solid State Logic launched their 2 and 2+, some four and a half years ago, they seemed like quite a big departure for a manufacturer best known for boat-sized mixing consoles. For a fraction of a percent of the cost of the flagship Duality desk, these wedge-shaped boxes promised the "SSL sound" in a desktop format accessible to anyone with a spare USB port. Unsurprisingly, they took the market by storm, and have proved enduringly successful.

I reviewed the 2 and 2+ back in April 2020 (www.soundonsound.com/reviews/ssl-2-plus), and although they were SSL's first foray into the world of affordable audio interfaces, found them to be well thought-out and rounded products that offered "a compelling balance of form, features and audio quality at highly competitive prices". Indeed, I've been using a 2+ as a portable laptop-based recording device ever since, and can add that it's proved reliable as well as capable.

No product is flawless, of course, but inasmuch as the 2 and 2+ had limitations, they mostly represented understandable design choices. For example, the second output pair was on unbalanced RCA phonos rather than balanced jacks: not my own preferred option, but undoubtedly convenient for some applications. And the 62dB gain range on the mic preamps was weighted towards the upper end of the scale, offering plenty of clean gain for spoken-word recording and the like, but limited headroom for very loud sources. Again, this makes good sense, given that an interface with only two inputs is unlikely to get used for close-miked drums.

Two By Two (Plus)

The original 2 and 2+ have now strolled off into the sunset, and in their place we have MkII versions of both. Given the popularity of the originals, it's perhaps no surprise that SSL haven't made any radical changes, but there are numerous small improvements both inside and out.

The new units occupy wedge-shaped cases that are very similar to those of their forebears, but marginally taller. At first glance, the top panels look almost unchanged, with the same complement of buttons and knobs. You get a large main volume control, plus separate volume controls for the headphone output(s) and a Mix knob that sets the direct monitor balance between inputs and playback. As before, this has an associated Stereo button that selects whether you'd like to hear the direct feeds from the inputs hard-panned or centrally panned.

There have, however, been some subtle but worthwhile changes to the inputs themselves. On the original 2 and 2+, each input was accessed through a single combi XLR/jack socket. High-impedance mode for DI'ing instruments was engaged by plugging your guitar into this socket and pressing the Hi-Z and Line buttons together. The new versions still feature rear-panel combi sockets for each input, but these are now used only for mic and line-level sources. Guitars, basses, electric banjos and the like plug into a dedicated pair of quarter-inch jacks, conveniently located on the front panel. These sense when a plug is inserted and take over automatically from the rear-panel combi sockets; consequently, the Hi-Z button on the MkI interfaces is no longer necessary.

Rather than remove this button altogether, though, SSL have repurposed it to add an extra feature in the shape of a switchable high-pass filter. This complements the existing 48V phantom »

power, mic/line selector and 4K Legacy buttons, the last of which still engages a combination of saturation and high-frequency boost that is said to emulate the sound of classic SSL 4000-series consoles. The three smaller buttons have also been changed from latching to momentary action, with LEDs added to visualise the status of the various settings. This is definitely a valuable change, as the position of the latching buttons wasn't always obvious at a glance.

It seems I wasn't the only 2+ user who didn't love the RCA phonos, because the rear panel now sports four balanced jack sockets carrying outputs 1-4. What's more, both interfaces now have two headphone outputs; as before, the second headphone output on the 2+ can optionally be switched to pick up the signal routed to outputs 3+4 instead of the first pair, whilst the new headphone output on the 2 is a duplicate of the first and shares its volume control. In another small but positive development, the headphone sockets have been moved to the front panel.

Elsewhere, the new 2s inherit many features unchanged from the earlier versions. They are bus-powered over USB, with no option to use a mains PSU, and connect through a Type-C socket on the rear panel. The 2+ still boasts MIDI in and out on five-pin DIN sockets. And, with low-latency monitoring built in and switched in hardware, there's no need to install any control panel software. Mac users won't even need to install drivers, as the 2 and 2+ are still class compliant.

Powers Of Two

There are, though, some fairly significant internal improvements, one of which is hinted at by the bold legending that reads "32-bit 192kHz". This signposts that the MkII units employ the latest generation of converter chips, which offer better specs than the originals. Output dynamic range is upped from 112 to 120dB, while there's now 116.5dB dynamic range on the mic inputs, a 6dB improvement. As with SSL's PureDrive preamps, however, there's no practical benefit to recording at 32-bit. Some USB devices from the likes of Zoom, Rode and Sound Devices use 32-bit float operation to make their input stages effectively 'unclippable'; as long as you're recording at 32-bit float in a compatible DAW, you can adjust the gain of a clip after the fact to recover from any apparent overloads.

Both of the MkII interfaces now feature two headphone sockets: on the 2, as shown, the second is a duplicate of the first.

That's not the case with the 2 and 2+, and 24-bit recording offers more than sufficient dynamic range to capture the full output of the A-D converters.

If anything makes clipping less likely, it's the design changes that SSL have made to the mic preamps. These now offer a 64dB gain range and can accommodate signals peaking at up to +9.7dBu, compared with +5.5dBu on the earlier iteration. The superb A-weighted equivalent input noise figure of -130.5dB is retained, meaning these are still some of the cleanest mic preamps around, and perfectly suited to the tasks such as tracking vocals and acoustic instruments that they're likely to be put to. Should you wish them to be less pristine, the 4K Legacy button still adds a quite obvious treble enhancement, which can be just the ticket on some sources, and a bit much on others. The new high-pass filter is an 18dB/octave affair turning over at 75Hz, and does its job very effectively, eliminating unwanted thumps and rumbles and thinning out overly heavy close-miked recordings.

Finally, although the 2 and 2+ still only have one pair of physical inputs, your DAW now sees an additional pair of loopback inputs. This is now a pretty ubiquitous feature in small audio interfaces, and while it's perhaps not all that useful in most music recording scenarios, it can prove invaluable for streaming, podcasting, vlogging and the like, where you need to feed the output from one software application into another. Some other small desktop interfaces go much further to cater for live streaming than simply providing loopback inputs, offering novelties such as OTG inputs for phone and tablet integration, or Bluetooth connectivity. SSL haven't followed this path, and that is fine with me: such features are of limited value in music recording, and if that's your main interest, you probably don't want to be paying for functionality you won't use.

If this review makes it sound as though SSL haven't fixed anything that wasn't broken, then, that is indeed very much the case. I already really liked the 2 and 2+, and I can't fault any of the changes that have been made in the MkII versions. I'm not sure that the improvements to the audio specs will make a noticeable difference in the vast majority of cases — if you're actually recording sources that have a dynamic range of over 100dB, you probably won't be doing it with a laptop and a stereo bus-powered interface — but it's confidence-inspiring to know that you're not compromising your sound between mic and DAW. Whilst individually small, changes such as the separate guitar inputs, the LEDs indicating button status and the 2+'s extra pair of jacks offer a real improvement to the user experience, and make for a very worthwhile refresh.

summary

This well thought-out revamp will help SSL's 2 and 2+ retain their status as popular favourites for portable recording and small studio applications.

$ 2 $229.99; 2+ $299.99
W www.solidstatelogic.com


Sound. Design. Vision.
Whether you simply need new monitors for your bedroom studio or want to design and build a commercial studio from the ground up, Sweetwater has the expertise to make sure your vision comes to life. Call our experts to see what we can do for you.
Sweetwater
Sweetwater.com
(800) 222-4700

Toontrack EZmix 3

Processing & Effects Plug-in

Version 3 of Toontrack's do-everything mixing plug-in brings AI assistance along to the party...

JOHN WALDEN

The name has always been a bit of a give-away: from the ground up, Toontrack's EZmix was designed to make recording, mixing and mastering as easy as possible. We reviewed the previous version, EZmix 2, back in *SOS* December 2013. Impressive as it seemed back then, computer technology and software design has since undergone some seismic changes, so you'd expect the new EZmix 3 to be a significant update. It is. While the underlying EZmix concept remains intact — it's a single plug-in that attempts to cover all your routine processing needs for recording, mixing and mastering, and whose approach is based around a huge suite of processing presets, each with a compact selection of macro-style controls — there's also plenty new to explore, and naturally in 2023 this includes some AI-based features.

Taking It EZ

EZmix 3's GUI has been significantly revamped. The new default layout is divided into a number of main panels: a Preset Rack (far left); a Filters/Similarity Map display (upper centre, you can toggle this between the two states); a Preset List (far right); and, once a preset is loaded, an Effects Controls panel (bottom centre). There's also now a helpful Quick Start screen, shown when you first launch a new instance of the plug-in, that focuses your attention on your initial preset choices.

The Filters tab provides a tag-based means of refining the preset selection that appears in the Preset List display. In contrast, the Similarity Map displays the presets as colour-coded points (Mix Utilities, Saturation, Amps & Distortion and Reverbs & Delays), with points that appear close together on the map offering similar styles of processing. Again, you can use filters to narrow the search, but this view is very useful if you've already selected a preset and want quickly to find alternatives that might be a better fit for the specific task in hand.

Many of EZmix's presets are, under the hood, based on multiple processes or effects. But EZmix 3's Preset Rack lets you chain up to five presets in a single instance — if you want to add compression to that otherwise perfect guitar amp preset, or saturation to a vocal preset, this means you can do so without having to load multiple instances of EZmix. You also now get four 'state' slots (A to D, located top right) that allow you to audition alternative settings with ease, and to experiment without risking losing your initial settings.

Core Strength

EZmix 3 includes an extensive collection of new core presets, and these span a huge range of potential applications. They include substantial options for mainstream instrument groups, such as guitar, drums, bass, keys and vocals, of course. But there are also smaller, more specialised selections for other categories, such as percussion, strings, brass and woodwinds. There's instant gratification to be found here. For example, you can apply any number of the impressive selection of vocal presets to your raw vocal recording, and many will add a good dollop of 'mix-ready' to the sound. There's also an excellent collection of presets specifically aimed at the stereo mix bus, or for final mastering.

Whichever preset groups you dip into, individual presets might offer something as straightforward as a single effect (for example, compression, EQ or modulation) or, in contrast, more complex signal chains (some of the mix bus presets, for example) that might combine multiple effects into a signal chain. Incidentally, for those upgrading from earlier versions, the preset libraries from both EZmix 1 and EZmix 2 are included alongside the new ones, and any preset expansion packs that you already own are not only compatible but will be fully integrated into EZmix 3.

AI Assistance

It's 2024, and so-called 'AI' is everywhere now, so it's hardly surprising that Toontrack have included some machine-learning-based features into EZmix 3, and this is evident in some of the preset groups. There are a couple of AI-based tags available in the Filters panel if you want to find these options quickly. In the main, the AI elements appear in three different types of presets: those aimed at bus or mastering tasks; in a number of 'guitar pre-EQ' presets; and in a large number of guitar and bass amp presets. They're all interesting in their own right, but the AI does something a little different in each.

For the bus and/or mastering presets, a new Analysing Track button appears in the Effects Controls »

sections and, much like the way AI is used in some plug-ins by iZotope or Sonible (for example), EZmix 3 can 'listen' to a section of the incoming signal and make informed suggestions as to the initial parameter settings. Having experimented with a range of these presets, I think on the whole that EZmix 3's AI assistance does a good job. Those new to mixing and mastering their own music or those simply needing to arrive at a workable result very quickly should find these presets very useful. You can, of course, then use your own ears and judgement to refine things further, tweaking the AI-based suggestions so that they work appropriately alongside other elements in the mix.

A second set of presets also includes the Analyse Track feature: the guitar pre-EQ presets. These come in different flavours to suit both acoustic and electric guitars, but essentially they're designed to 'pre-EQ' a DI guitar signal. The aim is to prepare the recorded sound to be at its best before you then start the actual mix work. In use, I found them very helpful, particularly the Acoustic Guitar and Rock And Metal Guitar versions. The former automatically provided some useful adjustments for acoustic guitar recordings that had been made in less-than-ideal spaces, while the latter seemed to rebalance the DI signal in a way that improved the tones I could achieve using subsequent amp modelling plug-ins. Similar in intent to Sonible's excellent pure:EQ plug-in, this 'pre-mixing' role for AI assistance is something many users may find helpful and is a welcome addition.

In the third group, a collection of new guitar amp/cab/effects modelling presets, AI technology was apparently used during the amp capture process. These span an impressive range of guitar tones and genre-based styles, from super-cleans through to thrashy metal, with plenty of stops along the way.

> "It's really not difficult for me to imagine tackling a full mix using only instances of EZmix 3."

No, the macro controls don't offer you the degree of tone-shaping provided by some of the leading dedicated amp sim plug-ins, but you can nonetheless arrive at some solid, very usable tones in very short order. EZmix 3 adds very little latency to your signal too, so real-time playing through these guitar and bass presets is a perfectly comfortable experience.

No Parameter Paralysis

There remains no danger of parameter paralysis here and everything's incredibly easy to use: you pick your preset, quickly tweak the carefully chosen parameters with the control set and then move on to the next task. But I suspect even the most ardent fan of keep-it-simple control panels will welcome the additional options that the expanded Effects Control panel brings. For any given preset, this panel now offers up to 10 control options, which is a marked increase over EZmix 2. As before, some are 'macro' style controls that, behind the scenes, adjust multiple underlying parameters, but others are more obvious single-parameter adjustments. Additional input and output controls, a wet/dry mix control, essential metering (including gain reduction for appropriate presets) and, on the mastering presets, a display of both short-term and average LUFS values are also provided. I think Toontrack have moved in exactly the right direction to get this tricky design balancing act right.

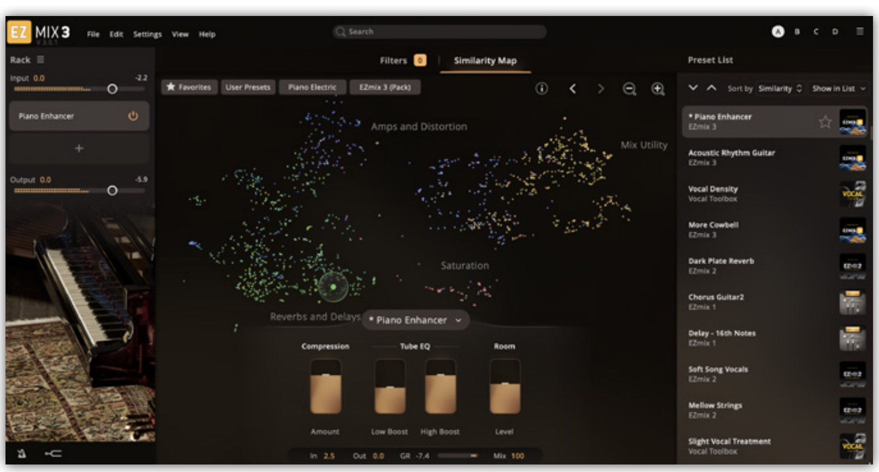

The Similarity Map display provides a useful alternative means of finding just the right starting point preset for the job.

EZmix Or Manual Mix?

EZmix 3 doesn't quite cover every eventuality — the most obvious exceptions are pitch-correction and super-creative sound design — but it could well cater for all your conventional processing needs for the huge number of routine choices you make in any typical mix project, and it's really not difficult for me to imagine tackling a full mix using only instances of EZmix 3. You shouldn't doubt the results of which it's capable either and, just as importantly, it can deliver genuinely workable solutions very quickly. For many potential users, whatever their experience level, that could be a very big deal.

It should be obvious why EZmix 3 would appeal to relatively inexperienced mix engineers or self-recording musicians. Naturally, it won't be to everyone's taste — if you prefer full control over every parameter of your processing signal-chain, EZmix's simplified approach might not be for you — but the streamlined workflow could well appeal to those with more experience who just want to achieve effective results as quickly as possible so they can focus on more important jobs, and are less concerned about their plug-ins mimicking the exact sound and look of highly regarded studio hardware. Incidentally, I should point out for those looking to get a little more out of EZmix 3 that the main controls of all presets are made available for automation in your DAW.

Over a decade on from the previous version, Toontrack have modernised the EZmix concept in some very positive ways, including a stylish new GUI, a greater degree of user control, and the integration of AI-based assistance. And none of this compromises its ease of use, which is quite an achievement.

summary

Toontrack's EZmix boasts lots of improvements, including AI-basd features. It's arguably easier to use than ever and is capable of very usable results.

$ $149 (discounted to $119 when going to press).
W www.toontrack.com

Apogee Symphony MK II Special Edition

"Apogee's Symphony MK II offers a direct pathway to the soul of the artists' music. Its unmatched clarity and fidelity redefine studio excellence."

- Bob Clearmountain

Watch

Dive deep into the art of mixing, as Bob shares his experiences and techniques with Rick Beato.

SIGN UP

Clearmountain will be hosting intimate Immersive Audio Mixing Sessions. Sign up for information on these events.

Sony MDR-M1

Closed-back Headphones

Sony's latest closed-back headphones complement the open-backed MDR-MV1.

SAM INGLIS

Japanese mega-corporation Sony have a presence in practically every aspect of the music industry. Whether you're tracking with Sony mics, releasing your songs on one of the innumerable Sony labels or playing them back on a Sony hi-fi, they really do offer something for everyone. And that something naturally includes headphones.

A vast range of over-ear, on-ear, in-ear and wireless models is available from Sony's consumer electronics arm, but studio and broadcast users are catered to by the separate Sony Professional division. This has long been home to the evergreen MDR-7506 closed-back headphones, and last year saw the launch of the MDR-MV1, which, as far as I'm aware, are Sony's first open-backed design, intended for mixing and in particular for working with immersive audio. And 2024 sees the introduction of a companion closed-back model called the MDR-M1.

Sony's marketing team are taking a slightly different tack with the new phones. The focus on spatial audio is gone, and instead, the emphasis is on what Sony call "ultra-wideband playback". The MDR-M1s proudly bear a Hi-Res Audio certification from the Japan Audio Society, meaning that they are capable of reproducing frequencies of up to 40kHz. According to the specifications, in fact, high-frequency extension is a humongous 80kHz, while at the low end, they are said to go as low as 5Hz. As ever, though, no tolerances are given.

In other respects, the specifications of the MDR-M1 are fairly standard for modern headphones. They use a 40mm moving-coil driver which presents a nominal impedance of 50Ω, with sensitivity given as 102dB/mW. Low distortion is also claimed as a positive, though no measurements are given.

Box Fresh

The MDR-M1s are sold in commendably eco-friendly paper-based packaging, and seem designed to offer a long and fruitful life, with replaceable earpads and detachable cables. Two of the latter are supplied, in different lengths; they plug into a mini-jack socket in the base of the left earcup, and can be locked in place by screwing in a threaded barrel. Sadly, though, you don't get a case, or even a soft bag.

Weighing in at a mere 216g without the cable, the new Sony cans are among the most lightweight models around. This is thanks in part to their compactness, and in part to construction which seems mostly to use rigid plastic. The earcups are attached to yokes that allow them to swivel both vertically and in the 'fore and aft' plane, and the effective length of the headband can be adjusted easily and precisely. I found them very comfortable for extended use.

The High Life

Although the Sony MDR-7506s remain very popular, and are used by numerous big-name engineers, I've never liked them without corrective EQ; they are, for want of a better word, somewhat shrill. By contrast, I'm still a fan of the long-discontinued MDR-7509s, which have a completely different, more mid-focused tonality, and I have good memories too of the 7510 and 7520. The recent MDR-MV1 open-backed headphones sound different again, with a modern, 'hi-fi' sound that foregrounds the upper bass and the high-frequency presence region.

Unsurprisingly, these new closed-back Sony cans take most of their sonic cues from the MV1, rather than any of the older models. My middle-aged ears can't tell you how well they put across 40kHz, but I can well believe Sony's claims about their high-frequency extension: even within the audible range, there's a noticeable sense or airy brightness about the top end. This is neither unpleasant nor unnatural, and there's not a hint of the grittiness or sharpness you'll find in many headphones. The low end also feels subtly emphasised, but whereas some rivals aim to create the impression of epic bottom by pushing the low midrange, bass on the MDR-M1 seems to have genuine depth. Sony's claim of low distortion is also very plausible; the MDR-M1s sound very clean, with an almost exaggerated sense of clarity.

Like the MDR-MV1s, then, the MDR-M1s are perhaps not the most neutral-sounding headphones in their class, but inasmuch as they deviate from the Platonic ideal of a perfect flat response, they do so in ways that are straightforward to get your ears around, and indeed useful in many scenarios. Both for general-purpose studio use and for mixing, especially of music with significant bass content, the MDR-M1s have a lot to offer.

summary

Clean, clear and comfortable to wear, the MDR-M1s are very nice all-purpose headphones.

$ $249.99
W https://pro.sony/ue_US/products/headphones


BAE
OUR ONLY COMPETITOR IS VINTAGE.
500 SERIES
“They were magic; every time I put my hands on these [things], it’s like every day is Christmas!”
Steve Vai (Guitarist, Songwriter, Producer)
73EQL
73MPL
IT’S NOT JUST ABOUT THE KNOBS

SPL MixDream XP Mk2 Summing Mixer

With 18 inputs and bags of headroom, this new MixDream aims to keep high-quality analogue summing simple.

MATT HOUGHTON

SPL's MixDream XP Mk2 is a Class-A 1U rackmount mixer with 18 line inputs, fed to it through a pair of (Tascam standard) DB25 D-sub connectors and a pair of XLR connectors. Like SPL's other products, it's manufactured in Germany and, as with other SPL gear I've used, the construction quality of the review unit was immaculate, inside and out.

Features

The 16 DB25 inputs are organised on the front panel as eight pairs, the left-hand eight channels having four latching buttons to switch each pair between dual mono (both signals are routed to the left and right busses, putting them in the centre) and stereo (the odd-numbered input goes to the left bus, the even-numbered one to the right). The eight channels on the right always operate as stereo sources (odd channels to the left bus, evens to the right again), and each pair has a latching -18dB pad button; SPL suggest this is to allow you to send hotter signals from digital devices/plug-ins, such as reverbs, for "higher digital resolution". There are no individual level or pan controls for these channels; the focus of this device is very much on simple, high-quality summing. The electronically balanced XLR Expansion input, activated/bypassed using a dedicated button, accepts the stereo output of another device, typically a mixer for more inputs at mixdown; for example, you could daisy-chain two MixDreams. But you could alternatively feed it any line-level signal.

Another central button, labelled Variable Output, enables an output gain control (-10dB to +5dB; when not engaged, the mixer sums at unity), so you can optimise the level for the next device in the chain. There are two separate stereo line-level outputs on rear-panel XLRs, a main mix out and a monitor output. Wired in parallel, these electronically balanced inputs deliver the signal post the output gain control. An IEC power inlet, its associated rear-panel on/off rocker switch and mains voltage slide switch, and a ground-lift button complete the I/O and user controls.

Inside the device, the audio rails operate at ±30V, so there's bags of headroom, and although the specifications in the manual could be more informative (eg. we're told there's 10Hz to 200kHz frequency response, but with no reference) they do confirm a high technical quality, with low noise, distortion and crosstalk and decent common-mode rejection.

Sum Great Reward?

I hooked up the MixDream XP Mk2 to my patchbay and fed it various signals from some DAW projects, both 'as is' and processed by outboard gear. I don't really buy into the practice of breaking a DAW mix out of the box just for analogue summing but it makes sense to me to capture a stereo mix when lots of signals are already coming out of the box; it means fewer stages of A-D conversion.

Subjectively, it sounded just as good as I'd anticipated: musical, clean and effortless, with no discernible noise or distortion, just as analogue summing should be! Though reassuring to know hot signals could be accommodated I didn't have need of the pads. The expansion facility worked well: I plumbed in my Dangerous D-Box to give me 24 inputs in total, and noticed no degradation. The output level adjustment was handy, with just enough range to set how hard the mix drove a stereo bus compressor. I occasionally wished for more facilities, such as LCR panning for the mono inputs and a pre-output insert point so stereo bus processing could be monitored but, to be fair, SPL do offer a more feature-rich summing mixer, and I generally felt that I had what I needed. That the status of all the controls was really easy to see at a glance is another positive that shouldn't be overlooked.

In short, then, this is a high-quality summing mixer that's easy to use. If you're in need of a high-quality summer without lots of 'extras' it warrants a slot on your audition list.

As well as the pair of DB25 sockets, which cater for 16 input channels between them, there's a switchable stereo expansion input on XLRs, taking the total input count to 18.

summary

A good-sounding, no-nonsense summing mixer with plenty of channels on board, as well as scope for expansion.

$ $1499.
W https://spl.audio

Erica Synths Nightverb

Stereo Reverb Processor

With a strong focus on lush, ethereal reverbs, this desktop device strikes a decent balance between simplicity and versatility.

PAUL WHITE

Erica Synths' Nightverb is a desktop-format, stereo-in, stereo-out algorithmic reverb effect processor that has enough input level adjustment to allow it to be used with the hotter signal levels of modular synths as well as things like guitars and keyboard synths. For those unfamiliar with Erica Synths, the company are based in Latvia, where they also do their manufacturing. Their products have tended to focus on Eurorack modules, desktop instruments and desktop processors — some of which are quite esoteric!

Overview

Nightverb has more in common with what were traditionally rackmount reverb processors than with pedals (there's no footswitch unless you add an external one), and it's designed with ease of use in mind, which means it has perhaps fewer user-adjustable parameters than some of the more feature-heavy software reverbs. Despite that, it's capable of generating a surprisingly wide range of reverb effects. While perfectly capable of creating small space emulations and typical studio reverbs, Nightverb's main focus seems to be on lush, ethereal reverb textures, and though the longest don't reach infinity, they can probably see it on a clear day! All of which makes it a particularly good fit for electronica and ambient musical styles.

We're informed that the stereo reverb algorithm used in Nightverb was developed by 112dB, who develop plug-ins and, more recently, Eurorack modules in Utrecht in the Netherlands, and it runs on Erica Synths' own custom DSP engine. The styling of Nightverb is clearly related to that of 112dB's Black Stereo Reverb and Black Stereo Delay Eurorack modules, which are also available from Erica Synths.

Sporting a dark-as-night theme, with a moth graphic picked out in gold, the Nightverb is visually striking, and it's housed in a rugged folded metal case with large, vintage-style knobs, clear graphics and a gently sloping profile that sits perfectly on a desk. The overall size, including knobs, is 230 x 145 x 70mm, and it weighs 833g. On the rear of the case are two balanced TRS jack inputs, two balanced TRS jack outputs, a footswitch jack, a USB socket for patch backup and firmware updates (there's no software editor, which could have been a handy touch) plus a PSU input for the included 12V, centre-positive power supply. If you plan on using a pedalboard PSU, then, you'll need one that has fully isolated outputs with a polarity swapping cable for the Nightverb in order to accommodate the centre-positive connection. Mono operation is catered for by plugging into the designated jack and the unit is also happy working with unbalanced connections. The footswitch jack is configurable for bypass, freeze or patch navigation and requires a straightforward TS non-latching footswitch. »

THE REVOLUTION CONTINUES.

Revolution 14x16: The truly high-end recording interface with 2 Auteur preamps with Cinemag® transformers

The **Revolution 14x16** is a class-compliant, 14-in / 16-out USB-C audio interface. Its **built-in Cinemag transformer-driven Auteur preamps** let you add weight and color to your tracks, and the **headphone mixer lets you tailor a headphone mix** from both hardware inputs and software playback. When your needs grow, Revolution 14x16 grows with you — add eight channels of input and output over ADAT; add external hardware processors via **4 switchable send/returns**; and sync your hardware devices via our **patented Macro-MMC clock**; treating your music to world-class clocking and low-noise conversion.

Visit **www.blacklionaudio.com** to learn how to Revolutionize your studio.

Expand your I/O with the Revolution EXP: 8x8 AD/DA Converter

S O U N D
B E T T E R
B E T T E R
S O U N D

» MIDI is handled by a pair of standard five-pin DIN connectors, and there's the option of using the Nightverb's controls to send MIDI CC data via the MIDI thru port (which can be set to operate as either a thru or a straightforward MIDI output). The user can assign custom MIDI CC messages to Nightverb's parameters if the defaults are not what's needed. For example, by default the CC message for the Size parameter is 70, which corresponds to Sound Variation. The MIDI input can be used to select patches and to control any of the Nightverb's parameters other than the gain/volume controls, again using MIDI CC messages. Unusually, when Freeze is active its pitch can be played via MIDI following a 12-note scale, which works by controlling the Size parameter (this seems to adjust the reverb tail generator's clock speed) in semitone steps.

Controls

The front panel has 15 controls if you include the Data encoder, which incorporates a push-switch. In conjunction with the Back button, a small OLED display is used both to navigate the presets (30 factory and up to 70 user), access various system settings (such as saving and naming patches), and enable a Morph function that allows the parameter values to change at a preset rate (a maximum of 10 seconds) when a new preset is selected. Presets may be re-ordered to make access easier in a live performance situation. There's also a Dirty function that provides an alternative reverb character: by default, the Nighverb's reverb coloration is based on an emulation of magnetic tape, but when Dirty is switched on you get the lo-fi character of a charge-coupled (BBD) type of delay. Illuminated buttons are used for Bypass and for the Freeze function.

The OLED display shows patch names and other relevant data and, although there's only one core reverb algorithm, Nightverb has a few tricks up its digital sleeve. Reverb Size is adjusted using the large knob at the centre, whereas the actual decay time of the reverb tail is adjusted using the Feedback knob. That one ranges all the way from a fairly dry ambience to a reverb decay that's longer than most songs! The Feedback control function also interacts with the Size and Damp controls, and if a reverb that's longer than your lunch break still isn't long enough for you, then the Freeze button will hold the reverb tail indefinitely, allowing you to play over it as a kind of background wash.

The MIDI thru port can be used to send MIDI CC data from the Nightverb's knobs.

A separate knob called Shape adjusts the spacing of the early reflections, while Spin adds modulation within the reverb-tail algorithm — this can add texture or, at higher settings, an almost chorus-like modulation. Spin also smooths out any tendency for the reverb tails to sound metallic (a technique originally pioneered by Lexicon). There are separate High and Low Damp controls, adjustable pre-delay time, stereo-width adjustment and separate bass and treble controls. The Wet/Dry control works just as you'd expect, and there are separate gain/volume controls for the input and output. In order to save the Dry/Wet knob position in a patch, while still allowing the dry signal to remain analogue, VCAs are used to balance the dry and wet signals, but there's a user option to switch to a fully digital signal path if that's preferred. Two rows of LEDs wrapped around the sides of the Size knob act as level meters. Helpful though this is, I'd have found it useful also to have a dedicated clip LED or perhaps 'traffic light' coloured LED metering for the two inputs. Bypass cuts the reverb but retains the in and out level settings.

In Use

The reverb algorithm can go from subtle room reverbs or short metallic reflections right up to long, ambient decays that hang in the air for ages. It's also quite capable of some less natural-sounding reverb variations, some with an almost granular character while others have a low end that you'd swear involved pitch shifting. The Tone control has a marked effect on the reverb character, producing warm, dreamy textures and washes when the high end is backed off. At the bright end of the scale, the reverb comes over with a steamy clarity and presence, while maxing out the early reflection spacing using the Shape knob adds a grainy texture to the reverb attack — that can sound really effective on piano and on 'plucked' synth sounds. Setting the Feedback control near its maximum can also throw up some interesting timbres. The ability to play frozen reverbs via MIDI is novel, and if you get tired of creating your own patches, you can invoke what Erica call the Magic function via the menu, which creates random presets for you. Some of the random creations are quite unusual but often still musically valid so it is worth experimenting and then saving your successful creations in the user memory slots.

Good Night?

To put this all in simple terms, then, Erica Synths' Nightverb is a simple-to-use yet surprisingly versatile reverb processor that delivers high-quality algorithmic reverb effects that sit well with synths, just as you might expect from this company. But in reality, it works just as well with any instrument, especially when working in the ambient music genre. It also integrates well with MIDI systems, while its gain structure allows for comfortable integration with typical modular synth systems.

Realistically, when it comes to good-quality reverbs, there's strong competition in a similar price range from a range of companies including Strymon, Universal Audio, IK Multimedia, Boss, Source Audio and more — but many pedal-format reverbs will struggle to handle the higher signal levels commonly found in modular synth systems. While the asking price might seem a tad on the high side for a unit built around just one reverb algorithm, the range of reverb effects on offer here is actually impressively wide, and in terms of quality you won't be disappointed if expansive lushness is your thing. The smaller space emulations are none too shabby either. The Nightverb is an impressive addition to the Erica Synths range and well worth checking out.

summary

A stylish, well-designed desktop reverb effect that's easy to use and offers a surprising amount of versatility.

$ €490 (about $550).
W www.ericasynths.lv


Standard
Standard
Standard
BOCK 187
BOCK 167
BOCK 251
sphere dlx
sphere lx

Chord Alto Studio Power Amplifier

This diminutive, studio-oriented desktop amp can drive nearfield speakers or multiple pairs of headphones with equal aplomb.

SAM INGLIS

Based in Kent, England, Chord Electronics are perhaps best known as manufacturers of high-end hi-fi gear. But they started out in the studio market, and retain a strong presence in its upper echelons. Control rooms all over the world rely on Chord power amplifiers to drive everything from huge main monitors to Auratones and NS10s, and at this year's NAMM Show, the company announced what is probably the most *SOS*-friendly product they've ever made. Offering an innovative and, as far as I'm aware, unique combination of features, the Alto is a miniature power amplifier and headphone amp designed specifically for studio use.

Now that most nearfield monitors have amplification built in, this is a relatively underpopulated product category. So if, like me, you still cling stubbornly to the idea that NS10s are a valuable tool in the 21st Century, you'll know how hard it can be to find a suitable amp to drive them. Over the last decade or so, I've been through several second-hand amps, all of which eventually fried themselves. I've bought a budget rackmount model which delivered almost as much fan noise as it did power into 8Ω. And I've baulked at the idea of dragging a huge Bryston or similar into my fairly modest room just to power a secondary pair of monitors.

So, on paper, the Chord Alto looks like the perfect product for people in my position. It's small: 1U high with the optional feet removed, and about 21cm square. It's fanless. It delivers 50W into 4Ω, so is clearly intended for nearfield monitoring applications. Unlike most hi-fi gear, it has XLR inputs as well as RCA phonos. And, as we'll see, it's much, much more than just a power amplifier.

Altoed Images

The Alto isn't quite as exotic in appearance as some of Chord's hi-fi kit, but has clearly been designed by someone with a keen aesthetic sense. It's made in Britain, and is every bit as well built as you'd expect at the price, with an impressively tough and smart aluminium shell. The rear panel boasts stereo inputs on XLRs and RCA phonos, speaker outputs on banana sockets, and a pair of line-level outputs on male XLRs. Maximum input level is specified as 6V RMS, which equates to +18dBu. The Alto takes 12V DC from a large laptop-style PSU, connected via a four-pin locking connector, and power can be cascaded to other Chord devices courtesy of an adjacent 12V output.

The front panel is eye-catching, to say the least, with visual feedback presented on three large glowing marble-like buttons. The largest of these even rotates in a satisfying way when you roll your thumb across it, although it's actually just an on/off switch. It's flanked by input and output selector buttons which light up in different colours to tell you which ins and outs are in use. Either side of these, there's a single volume control and no fewer than four headphone sockets: two quarter-inch jacks, a mini-jack and a 4.4mm Pentaconn 'balanced' output. A fourth 'marble' is neither a control nor an LED, but a window for the included infrared remote control.

The remote too is surprisingly and impressively well built, with a sleek and rigid metal case. It has many more buttons than there are controls on the Alto, so is presumably designed to work across the entire Chord range. Chord even have the forethought to include the two AAA batteries you'll need to power it, though not the tiny crosshead screwdriver you'll need to open the battery compartment! Point the remote at the device and hit the volume up or down buttons and you'll witness the Alto's party trick: its volume control is motorised, so whether you adjust it by hand or remotely, there's never a discrepancy between the position of the control and the actual volume setting.

The Chord Ultimatum

At the core of Chord's power amplifiers is a technology called Ultima, developed by company founder and design guru John Franks. There are several elements to this, starting with their own high-frequency switch-mode power supply, but the key technical feature is something called "dual feed-forward error correction", whereby the amplified signal is continuously compared with the input signal and adjusted as necessary. The Ultima design »

As befits an amp designed for studio use, the Alto features balanced inputs on XLRs as well as its phono sockets.

uses Chord's proprietary MOSFETs, operating in Class A/B, giving the Alto a significant point of difference compared with most small desktop amps, which these days are usually Class D. It also delivers specifications that rival high-end Class-D amps. THD, for example, is quoted as 0.003% into 4Ω, which is a couple of orders of magnitude better than a typical traditional rackmounting amp, and signal-to-noise ratio is an impressive 119.6dB.

The Alto's form factor makes it look a lot like a monitor controller, and it has a little of the functionality you'd want from such a thing — but if you expect it to be a monitor controller, you're likely to wind up scratching your head. For example, on a typical monitor controller, you'd expect the headphone sockets to have their own amps with individual level controls. That's not the case here. The point of the Alto is to provide a single, really good stereo amp, with switching that allows this to be fed either to the speaker outputs, or simultaneously to all four headphone sockets. It's not possible to have both the speakers and headphones active at once, or to adjust their levels independently. By contrast, the line-level XLR outputs can optionally be made active at the same time as either the speakers or the headphones; they can also be exempted from the Alto's volume control in order to feed another amp, master recorder, meter or monitor controller further down the line.

Having everything be subject to a single volume control can certainly prove frustrating if you try to treat the Alto as your main monitor controller. The setting that gave me a comfortable level from my NS10s was much too loud with most of the headphones I tried, so it isn't really feasible to toggle rapidly back and forth between speakers and cans for checking mixes and so on. Nor is the volume control detented, so you can't easily restore a previous setting. And although there is a Dim button on the remote, this merely makes the LEDs less bright! In long-term use, you will probably end up employing the Alto either as a headphone amp or as a speaker amp, depending on what you happen to be doing that day.

Power Trip

The flip side of this is that standard monitor controllers aren't power amps, and thus can't drive your NS10s, Auratones, LS3/5as and so on. The Alto not only can do so, but does so spectacularly well. Power amplifiers are not the easiest thing to A/B, but the Alto sounded better than anything else I personally have ever connected to my NS10s, and it wasn't subtle. The few other studio-oriented power amps on the market aimed at nearfield listening typically offer at least 100W per channel into 8Ω, so the Alto's 50W into 4Ω might sound underpowered, but that wasn't at all my experience in practice. Even though it presumably develops only 25W or so per channel into an 8Ω speaker like the NS10, it delivers all the level I personally would ever want, and a fair bit more besides. Unless you're the kind of mix engineer who blows through a new tweeter every week, I think it would be more than adequate for nearfield listening in most environments.

> "The Alto sounded better than anything else I personally have ever connected to my NS10s, and it wasn't subtle."

If anything, the Alto's prowess as a headphone amp is even more impressive. With its four outputs all operating at the same level, it's not the most flexible headphone amp, but boy, does it sound good. As you'd expect, given that it can deliver a specified 2.25W into 100Ω, it has no trouble driving the most awkward of loads. I expected it to sound great with high-end open-backed headphones and it didn't disappoint. What really surprised me, though, was the difference it made with relatively modest headphones. Plug the same set of cans into the built-in headphone amp on even a high-end interface, and they sound flat and lifeless by comparison with the Alto, which delivers a lot more punch and dynamic impact.

It might, of course, seem a bit perverse to pay several grand for an amplifier to power speakers or headphones that are probably worth a tenth of that. But it's the system as a whole that matters. Most guitarists would be much happier playing a Squier through a handwired Fender Deluxe than they would playing a Custom Shop Strat through a cheap practice amp, and although I don't think the analogy always holds for studio monitoring, the Alto will certainly bring out the best in whatever transducer you choose — whether that's a humble pair of passive speakers or a set of headphones that cost more than your car.

summary

Compact, stylish, well specified and innovative, the Alto is both the ideal device to drive your nearfield monitors and a superb headphone amplifier.

$ $4325
T The Sound Organisation +1 972 234 0182
E support@soundorg.com
W https://soundorg.com
W https://chordelectronics.co.uk/professional

ARTURIA

_The sound explorers

_KeyLab mk3

Universal MIDI controller

Whether sculpting a studio masterpiece or channeling the raw energy of a live performance, KeyLab mk3 is an advanced MIDI controller that will inspire your very best composition.

From its design, to its pro-grade aftertouch keybed and seamless software control - KeyLab mk3 leaves no stone unturned.

Chambord Audio CA-250

500-series EQ Module

Inspired by an early parametric design, this module offers more control than most 500-series EQs.

MATT HOUGHTON

Chambord Audio are a small pro audio manufacturer based in Montreal. Currently, their product range comprises a four-slot 500 series chassis and a number of surprisingly affordable analogue modules, including two mic preamps, a high- and a low-pass filter, a parametric EQ, a compressor and a headphone amp. It's the CA-250 EQ that's on review here, and Chambord kindly sent a pair of them, so I could try them on mono and stereo material.

Based on the Sontec MEP-250A, this is a solid-state, three-band parametric EQ: the gain, frequency and bandwidth are continuously adjustable across their range. These three functions are given colour-coded knobs (red for gain, black for the bandwidth and grey for the frequency) and the whole ensemble looks and feels classy. That said, while the labelling around the knobs is printed clearly, the font is small, and with the shadows cast by the profusion of controls it can be tricky to discern the precise settings in low lighting. It's a pretty conventional mixing console-style layout, though, and it's easy to see the knob positions, so operation soon becomes second nature.

There's up to ±12dB of gain for every band and, as an alternative to the adjustable-bandwidth bell (the Q can be set from 0.4 to 4), the high and low bands switch to act as shelving filters when their bandwidth knobs are fully anti-clockwise. There's a global EQ in (or bypass) button, and although there are no individual bypasses the gain knobs' 0dB positions are detented, so it's pretty easy to hear the effect of a single band in isolation or turn it 'off' should you wish.

The low band's generous 15 to 815 Hz range lets you dig well into the low mids, while a very low-frequency bell or shelving boost gives you the scope to create different curves that 'reach up' into the audible range. The high band again boasts a vast range (1.8 to 29.5 kHz), so is capable of anything from dipping the 'honk' in a vocal to gently lifting the air frequencies. The mid band (200Hz to 12kHz) offers plenty of overlap with the others so, again, it could cover anything from dipping an ugly 200Hz boxiness in a drum to massaging the breathy frequencies in a vocal or cymbal sizzle.

On Test

I used the modules on mono and stereo sources in a mix I happened to working on when the CA-250s arrived. This gave me the opportunity to test them on male and female vocals, drums, bass, and various electric and acoustic guitar and keyboard parts, including a Hammond organ. And I can honestly say I'd be happy to mix on a console loaded with these EQs. It's a versatile design that does exactly what you'd want a parametric EQ to do: the curves are predictable, in a good way, and while the overall sonic character isn't overly strong it definitely appeals, and sounds different from conventional digital EQ plug-ins. That's particularly the case at the top end, where it smooths things a touch in a pleasing way, but in general it's delightfully smooth when boosting, helpfully clinical when cutting and, with the gains to unity, admirably unobtrusive.

The trade-off of the broad ranges and narrow knobs that the physical form factor necessitates is that the frequency control can sometimes feel a touch over-sensitive. It's not really an issue when treating a mono source — I found it pretty easy to dial in the sound I wanted by ear — but when using two modules on stereo material it was trickier to quickly match their settings precisely; to do that required care and a little time. The two units were very well matched, though, an opinion I formed first by ear before confirming with a frequency analyser, matching the curves precisely and then checking the position of the controls. Top marks there.

If you can get the hang of matching the settings, a pair of CA-250s could be very well suited to use on subgroups, the stereo bus or even for mastering. There's also the option of using two modules to process the Mid and Sides elements of a stereo signal individually, which I tend to do a fair bit in my own DIY mastering efforts. These modules worked admirably in that role, for example enabling me to dial in a warm 50Hz bell boost on the Mid while applying a brightening high-shelf boost only to the Sides.

Verdict

So, does this sound like a Sontec MEP-250? I can't say that — I don't have one to compare it with. At this price, there are undoubtedly more cost-effective components used than in a direct clone, and on the Sontec the shelves were separate bands. But does that matter? I don't think so. Judged on its own merits, the Chambord CA-250 is versatile for a three-band EQ, it's largely easy to use, sounds pleasingly smooth and is competitively priced. If I had a few in my rack, I could put them all to good use on most projects, so if you're looking for a 500-series EQ for tonal shaping or to nix resonances, the CA-250 is well worth consideration.

summary

A versatile and cost-effective 500-series EQ, the CA-250 offers decent value for money.

$ $499 CAD (about $367 USD) per module. Price excludes tax and shipping.
E info@chambordaudio.com
W https://chambordaudio.com

Baby Audio Humanoid

Vocal Transformation Software

Transform your boring old human vocals with Baby Audio's latest plug-in!

ROBIN BIGWOOD

Mangled, manipulated vocal sounds can be a big part of many electronic, dance and experimental musical genres. And they're exactly what Humanoid, a new plug-in from California-based Baby Audio, is all about: it aims to be a one-stop shop for a range of really synthetic, processed treatments, including hard tuning.

Plug-in formats are VST, VST3, AU and AAX, and you'll need to be running at least macOS 10.11 or Windows 10. Humanoid is an audio effect plug-in with an optional MIDI input, so most serious DAWs will be able to access all its features. The interface and style seems ideally suited to the iPad too, but there's no word on a version for that OS just yet.

Human Behaviour

Fire up the plug-in and you're met with a modern, continuously resizeable single-window design, and the five user interface panels give a strong indication of what effects are on offer.

Pitch is Humanoid's take on Auto-Tune and works in a few different ways. Scale mode is conventional, detecting the pitch of incoming monophonic audio and bending it to a quantised scale degree. Preset scales and an interactive keyboard graphic are provided to let you choose your desired notes. In Note mode pitch is constrained to a single pitch drone, or in Duo mode a pair in parallel, across a five-octave range. MIDI mode lets you route into the plug-in live or DAW track-based MIDI data to 'play' the pitch correction, and it's five-note polyphonic too, making for a vocoder/harmoniser-like experience. MIDI pitch-bend is respected, to the tune of ±2 semitones, but neither velocity nor aftertouch (let alone MPE) is observed, which is a shame. Finally a Lock option lets you maintain settings here even when exploring other presets.

For each mode there's a small handful of parameters, but it's worth noting immediately that pitch correction is always 'on' even when all controls are turned down. The Quantize knob would then seem to be a sort of macro, adding artificiality by increasing retuning speed and clamping down on pitch fluctuations. Nearby there's a familiar Formant twister, and also an intriguing 'Robotify' knob. This conforms the harmonic series of the input signal to a mathematical ideal, but also boosts upper harmonics to create an unnatural but not unpleasant plastic, zingy, almost cute quality.

Throughout Humanoid more nuanced control is available for parameters that have a 'cog' button: clicking one in the Pitch section for example accesses up to four 'calibration options', which are really additional parameters in their own right. So there's a nice balance here of immediacy and depth: you can do a lot of damage with just a few mouse drags, or spend longer to really explore the effect. The same is true of the Utility panel, that limits pitch correction range, has a useful, dry-sounding and not too invasive de-esser, plus a gate for the input signal. The Smoothing parameter keys into Humanoid's FFT-based processing engine to (I quote) 'reduce harshness', but it does so by adding yet another kind of artificial, robotic quality. The Sharpen control is more conventional, applying a presence EQ curve to enliven a drab input signal.

The destructive/creative heart of Humanoid is the Synthesize module. Two separate processes take place here. The first is sophisticated resynthesis, with the module capable of generating a new, synthetic version of your vocal in close to real time. The second, less obvious one is a pitch-shift that can add a ±1 octave doubling over and above a ±12 semitone transposition to the dry/resynthesized combo. The Pitch section's scale setting is respected when it does this, which can make for interesting and useful harmonisation possibilities.

In short, there's a heck of a lot of potential here for huge tonal manipulation, »

Baby Audio Humanoid
$129

PROS
- A range of more out-there treatments including tuning, resynthesis and filtering.
- Easy and quick to use, with a clear, uncomplicated interface.
- MIDI input allows harmonisation and chords from a mono vocal.

CONS
- Artificial-sounding even its most 'zeroed' state.
- Hyperactive pitch correction is a particularly dominant presence.
- No bypasses for individual processing modules.

SUMMARY
A colourful, sophisticated, in-your-face vocal processor that can wreak havoc ranging from gimmicky to gorgeous.

NOMINEE
SOS AWARDS 2025
Electro-Voice
BIG SOUND MEETS BATTERY POWER
EVERSE 12
WEATHERIZED BATTERY-POWERED LOUDSPEAKER
WITH BLUETOOTH® AUDIO AND CONTROL
ASHLEY NORTON
Electro-Voice user
OUR ELECTRO-VOICE COMMUNITY
KICK-BACK AND RELAX
With up to 12 hours of run time, EVERSE 12 has got your audience covered anywhere – indoors or outdoors. Premium features include a titanium dome compression driver, true stereo streaming, power for mobile devices and a wireless microphone system, and much more. Great sound to grab and go.
LOUDER LOWER LONGER
Higher output, deeper bass, longer run time*
8”
NEW SIZE!
12”
SIMPLE SETUP
It’s never been easier to optimize your sound, with user-friendly QuickSmart DSP and app control/mixing
EVERSE 12
* Compared to EVERSE 8 at 95 dB SPL. The EVERSE 12 battery pack will provide 12+ hours at 100 dB SPL average with typical program material, or 6+ hours at maximum output level before limit indication. Always use the supplied weatherized input cover, or the Electro-Voice approved accessory rain resistant cover in wet or damp environments. Information in this document is subject to change without notice. All trademarks are the copyrights of their respective owners. ©2024 Bosch Security Systems, LLC
electrovoice.com/everse | #RealProAudioUnplugged

and it's all invitingly tweakable. The degree of acoustic/synthetic hybridisation is controlled by the big Transform knob, and ranges from no synth contribution at all, through synthetic but still intelligible at 12 o'clock, to complete replacement of the original input and expungement of all consonants, sibilance and fricatives at maximum. A wet/dry mix fader tempers the effect in a different way.

Factory wavetables (64 are provided) range from familiar analogue-style waves, through metallic and distorted tones, synths and SID chips, to speech/formant, noise-based and edgy digital timbres. Two parameters, Shape and Stretch, further distort these starting points, increasing harmonic complexity. Further slots are available for user-imported audio, which can be up to 4096 samples long. Anything longer gets automatically truncated, and in my tests only WAV files (and not MP3 or FLAC, say) were accepted. There's no support for true wavetable files (in whatever non-standard format you might have them), but Shape and Stretch distortion of user waveforms is available all the same. Which makes me suspect that the factory waves aren't genuine multi-frame wavetables either, but behave as if they are through algorithmic transformation.

I mentioned some vocoder-like behaviour in the Pitch section, and the Synthesize module can certainly sound something like a vocoder. In fact though the integration of input and synth that goes on here happens at a much more sophisticated level, is distinctly more agile, and doesn't have that typical vocoder squelch. The input signal is the only way of controlling this synth, though, and it has no separate modulation sources, and (in conceptual terms) just one 'oscillator'.

Moving on, Humanoid's filter is a potentially resonant high-pass and low-pass sandwich with a parametric EQ-like middle which is rather oddly labelled 'boost'. A spectrum analyser backdrop lets you easily home in on the harmonic content passing through the plug-in, and both high- and low-pass filters are capable of dizzyingly steep cuts, paring off individual harmonics if necessary: I've rarely seen or heard a filter so 'surgical'. Bipolar resonance controls also allow for smooth slopes though, as well as more familiar emphasis peaks around the cutoff frequency. The

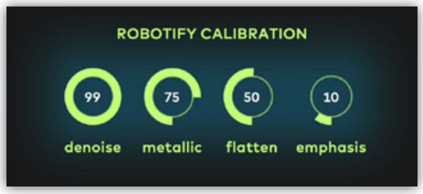

Under-the-hood 'calibration' parameters, shown here for the Robotify knob, allow for a deeper dive when required.

parametric mid band is more conventional, but still very useful for creating scoops and warmth all the way from 41Hz to 12.5kHz.

The Effects section is Humanoid's last raspberry to convention. Instead of a predicable delay/reverb you get a chorus-like stereo widener, a Warble effect that modulates the pitch of the resynthesized signal's upper partials, and a Freeze. That can be clicked or automated to turn Humanoid's output into a drone or stutter effect, muting the input. The adjacent Buffer parameter controls the length of the loop, with both arbitrary and tempo-sync'ed timing values.

> "Humanoid is both inspiring and great fun. It's quicker and more intuitive to use than any other competing vocal processor I've tried and absolutely invites experimentation."

Humanoid Boogie

Humanoid undoubtedly works best on pre-existing audio, ie. vocals already recorded in your DAW, as a fair amount of latency is inherent in the design. Buffer size is variable, and at the lowest setting it's usable in real time, just about, but with a reduction in pitch correction accuracy. Latencies of around 50, 80 and 180 ms for the three available settings were reported in Studio One, which I used for testing, at 48kHz.

In sonic terms, Humanoid's character is strong, occupying a distinct timbral

The resynthesis core of Humanoid comes with 64 varied waveforms, but you can also import your own, and then distort them in various ways.

territory, and it can never be made completely transparent. So it's unlikely you'll use it on your next acoustic folk album, but for arresting, creative vocal manipulations it offers exciting options. Alongside the aggressive pitch correction, Synthesize can completely reinvent vocals, transforming them at molecular level if necessary, and the filters can tease out anything from angelic whispers to demonic darkness. It's worth emphasising here too that Humanoid really does only work on monophonic voices: it leaves polyphonic sources or whole mixes in garbled tatters.

I have a handful of specific criticisms. The steroidal pitch correction, though characterful, is not flexible enough to encompass general problem-solving: vocals with lots of natural portamento, or that are not well performed or recorded to begin with, can result in an awful lot of adjacent-note yodelling. I would love to see a way to dial the effect right back. There doesn't seem to be a way to adjust or tune the correction's pitch centre currently either. Then generally I found the lack of module bypasses limiting. Not because I'd necessarily want to try and utilise bits of Humanoid independently (which would be wonderful, but I expect impossible because of the underlying processing scheme), but because it could be so helpful and instructive whilst experimenting to temporarily turn off individual bits, without destroying settings.

Those criticisms are outweighed by positives, though: above all, Humanoid is both inspiring and great fun. It's quicker and more intuitive to use than any other competing vocal processor I've tried and absolutely invites experimentation. Even from my guttural grumblings it was able to tease out all sorts of great new sounds, ranging from comedic and vintage-style vocal effects to silky and beautiful synths. There is a 'sound', certainly, but Humanoid can easily be different things to different people, with literally dozens of ways to use it. I enjoyed the MIDI harmonisation and filter, and for others it might be all about the tuning and extreme resynthesis. Either way, it's a plug-in no otherwordly vocal adventurer should be without.

$ $129
W www.babyaud.io

Arturia AudioFuse X8 In & X8 Out

A-D & D-A Converters

Arturia's ADAT expanders are a useful and affordable option for anyone who needs more I/O.

SAM INGLIS

Not even misplaced '90s nostalgia is likely to revive Alesis Digital Audio Tape as a recording format in 2024, but the ADAT Lightpipe digital audio transfer protocol remains ubiquitous. A single optical cable allows convenient one-way transfer of up to eight mono channels of 24-bit audio at base sample rates, or four at 88.2 and 96 kHz. It's easy to use, cheap to implement and scaleable enough to suit most situations.

One of the original use cases for ADAT expansion was to interface digital recording systems with analogue mixing consoles. Many older A-D and D-A converters with ADAT ports thus operated purely at line level on the analogue side. But it now seems that the main application for ADAT connectivity is to add extra standalone mic preamps to a computer audio interface. Consequently, there are now very few ADAT expanders with only line-level inputs, and almost none that provide only outputs. If you need such a thing, you could seek it out second-hand, but the audio performance of a 20-year-old converter won't match what's possible even in budget gear today.

Back In Black

With the AudioFuse X8 In and X8 Out, Arturia have revived the line-level ADAT expander. As the names imply, these two 1U, half-rack units respectively offer eight channels of A-D and D-A conversion. Each of them boasts two optical ports, allowing the full channel count to be maintained at higher sample rates as long as the connected device plays ball, and a BNC word-clock input with associated termination switch. Analogue signals enter the X8 In and depart the X8 Out on balanced quarter-inch jacks. As these units are standalone converters rather than audio interfaces, bus powering is not an option, so both come with wall-wart power supplies that have locking DC connectors. Both also come with optional widgets attached to the base that can be detached and used as rack ears.

One of the reasons why Arturia have decided the time is right to revisit the line-level converter is that a new use case has emerged since the heyday of the Mackie 8-Bus and Fostex VC-8. Many people are now generating control voltage and gate signals from software to control modular synths. For this to work properly, a converter or interface needs two features. It has to be able to generate or accept a voltage swing of at least ±5V; and it has to be able to work with unchanging DC voltages as well as with conventional audio signals. Both the X8 In and X8 Out are calibrated such that a full-scale digital signal corresponds to a +24dBu analogue input or output, which in turn equates to a peak voltage swing of more than ±17V, so there are no worries on the first score; and they are DC coupled, meaning that they are quite happy to engage with steady-state non-zero voltages.

These new additions to the AudioFuse line also fully exploit the power of modern converter technology. Whereas the original ADAT barely filled the theoretical dynamic range of a 16-bit signal, the X8 In and X8 Out offer 119 and 120.5 dB respectively, with equally impressive THD+Noise figures.

Rated X

The X8 Out is the simpler of the two devices. Apart from the power button at the top right of the front panel, it has only a single control: a Sync button to select the clock source from ADAT, ADAT S/MUX and word clock. The X8 In is somewhat more complex. The Sync button in this case determines whether the X8 In should be clock master or follow an incoming word clock (there's no ADAT input, so that's not an option). A second button labelled Clock sets the sample rate if the X8 In is acting as master; when it's clocked externally, sample rate changes happen automatically, but the LED for the old sample rate flashes until you hit Clock to 'accept' it. I'm not sure why this is really needed, since there's no obvious alternative course of action!

The left-hand side of the front panel, meanwhile, sports a tiny LED for each input, and four buttons. In normal use, the LEDs offer basic signal metering, turning green in the presence of a signal above -60dBFS, orange when the input level hits -6dBFS and red when clipping occurs. The manual rather daringly recommends setting levels so that the loudest peaks register orange; on anything at all dynamic, I personally would want to be a bit more cautious, and would like the option to have the LEDs turn orange at -12dFBS.

The four buttons are used to engage what Arturia call 'pads'. This is a confusing description for what is essentially the ability for each input to be switched between nominal +4dBu or -10dBV operating levels. With the 'pad' switched on, the inputs can accept peaks of up to +24dBu; disable the pad and you lose 12dB of that headroom. This is done by stepping left and right through the input list with the arrow buttons and double-tapping the Pad button. Pad settings for adjacent pairs can also be tied together using the Link button, which seems a little unnecessary.

Conclusion

After a slow start, Arturia's AudioFuse range has really come into its own in the last few years. They've developed superb mic preamps, and as befits a manufacturer of desirable synths, they've also thought about the needs of electronic musicians with products like the recent AudioFuse 16Rig. The X8 In and X8 Out are the first members of the range that aren't themselves interfaces, but they certainly maintain the standard. I have no doubt that there still are many people out there who would like to expand their computer I/O but don't want or need additional preamps, and the X8s provide a cost-effective, high-quality way to do just that.

Their arrival at my studio prompted me to finally make a change I'd been planning for a while. I ordered BNC cables and T-adaptors in order to distribute word clock to all the audio hardware that supports it, including the X8 In and X8 Out. They were equally happy to work in this configuration as they were to clock over ADAT. Everything worked first time, sound quality was exemplary, and my setup had eight ins and outs that weren't there before. If you're looking for line-level ADAT expansion and you don't need the full 16-in/16-out configuration offered by something like the Ferrofish Pulse 16, the X8s may well be your only option — and they're definitely a good option.

Things are pretty much as you'd expect round the back of the X8 In: analogue audio comes in, digital audio comes out.

The X8 Out's rear panel is pretty much identical to the X8 In, the only difference being that digital audio goes in and analogue audio comes out.

summary

The X8s are simple but effective line-level A-D and D-A converters with good specs and pricing that could be considered highly competitive — if they had any obvious competition!

$ AudioFuse X8 In $349, X8 Out $299.
W www.arturia.com

TC Electronic 2290 P

Digital Delay Pedal

The classic ’80s delay effect has been reborn in pedal form.

TED MORCALDI

TC Electronic’s 2290 Dynamic Digital Delay is a classic ’80s rackmount effects processor that went far beyond the standard delay units of the time. Thanks in large part to its powerful modulation section, it could be thought of more as a delay-based multi-effects system, with which it was possible to dial in chorus, flange, automatic double-tracking, panning, tremolo, ducking, gating, and even pitched delays (more on this later), in addition to the typical echo effects that you’d normally expect on a digital delay. The 2290 is still sought after today, and a few years ago TC created a plug-in version (with an optional hardware controller). More recently, they decided to recreate the whole thing in pedal format and the result is the 2290 P, reviewed here.

Overview

Housed in a rugged enclosure, the 2290 P offers up to 9.999 seconds of delay time and boasts the sort of retro-futuristic aesthetics that suggest it might have been plucked straight from the cockpit of a certain DeLorean. Besides the three footswitches (more on those below), the user controls take the form of many buttons and a single scroll wheel, with settings indicated on a few ’80s-style numeric displays. A sleek top cover is designed to prevent accidental parameter changes while tap-dancing around your pedalboard on stage, and it’s a thoughtful inclusion — when fitted over the top, you can still see the displays (though darkened), and the buttons are shielded to leave you with access to just the footswitches and encoder.

On the back, there’s a power inlet for the supplied centre-negative 9V DC supply (if using a pedalboard supply, you’ll need 250mA available). There are also MIDI in and out/thru five-pin DIN sockets, and a USB-C port — the 2290 P can connect to a Mac/Windows machine running a software librarian, which is mainly for preset organisation but also to configure MIDI settings, and it’s the only way to map the quarter-inch TRS expression pedal jack. There are four TS jacks for the analogue in and out. These can function as stereo in and out or, in a mono configuration, as a main in and out with an effect send/return for the feedback loop — you set the mode using a toggle switch between the two jack pairs.

Features include 128 presets in 64 banks (selected with the A and B footswitches) and, in tribute to the original, the pedal’s first 20 presets are recreations of those found in the rack unit. There are three footswitches, labelled A, B and Learn. The first two switch between different patches, while Learn is another nod to the original rack unit — that was released before ‘tap tempo’ had become the ubiquitous term! This is capable of controlling tap tempo for presets individually, both A and B at the same time, or as a global tempo control. There are many rhythmic sub-divisions programmable via the Learn switch.

Missing Manual?

It’s a shame that the manual provided is virtually non-existent. All you get is a one-page quick start guide that labels each button’s function (as opposed to explaining how the functions work in any depth) and a video tutorial online, and there are some settings that these fail to mention. The user would really benefit from an in-depth walkthrough of how each section of parameters interacts with the others and affects the overall sound. Fortunately, the manual for the original 2290 can be found at the https://tc2290.com site, which might best be described as a tribute page to the original unit, but it also provides some patch ideas, electronics history lessons... and occasional dad jokes! Because the control panels are nearly identical it’s almost directly translatable to the new pedal. When learning how to program sounds on the pedal version, then, I referred to that manual, which was very well written, and provides a clear walkthrough — almost as if the developer were in the room, talking you through every section, and how the parameters interact with other sections of the unit.

Getting Stuck In

Precision programming is this pedal’s speciality, with a wide array of parameters available that can easily be recalled in »

EVERY DETAIL OF EVERY MIX.

From award-winning hits to movie soundtracks and home demos, wherever you are on your musical journey, the new ROKIT Generation Five Active Two-Way Studio Reference Monitors from KRK Systems provide the creative tools you need. With three Mix Modes delivering three-in-one monitor flexibility and astonishing clarity, your studio and your ears deserve it.

THREE MONITORS IN ONE

- ***Mix Mode:*** *Flat frequency and phase response ideal for mixing and critical listening.*
- ***Create Mode:*** *Inspirational voicing for writing, producing, and casual listening.*
- ***Focus Mode:*** *Mid-focused voicing for critical analysis of the midrange band.*

Scan the QR code to visit
krkmusic.com and shop now.

KRK SYSTEMS®

The Sample Section?

This pedal version of the 2290 lacks the original rack unit's Sample section. That was a big part of what made the original 2290 delay a unique creative device: with 32 seconds of delay (with extra memory option) available, it was an obvious choice for many musicians who were experimenting with looping at that time. It was used by trumpeter/composer Mark Isham as his main looping device in the late '80s, and a pair of 2290s lay at the heart of Robert Fripp's soundscape looping setup in the early '90s, replacing his Revox tape machines. Fripp even stated in a 1993 *SOS* interview (free to read on the Muzines site: www.muzines.co.uk/articles/inside-frippertronics/10807) that David Sylvian managed to program his 2290 to 64 seconds, something that the TC distributor had deemed not possible! Had TC chosen to recreate this facility in their 2290 P, re-pitched, envelope-triggered samples, panning loop textures and more could have been available, and this pedal could easily have rivalled the top creative delay/looping pedals of the last few years. Not every user will need it, but I see it as a missed opportunity — as well as testament to just how far ahead of its time the original 2290 was.

» the preset banks. And after diving deep into programming the 2290 P, it became clear to me just how much the 2290 has inspired ideas in the modern, boutique pedal world. It's capable of effects ranging from beautiful, pristine modulated delays, to wildly random panning and — a pleasant surprise when I began to dig deeper into the modulation section — envelope-controlled pitch modulated delays. The precision of programming sounds on this delay feels more like a modern plug-in than a vintage device, which just goes to show how advanced the original unit was, and its ability to manipulate the stereo field is otherworldly, putting it in a different category to the average delay pedal.

The ducking delay is one of the effects the 2290 was most famous for, and this is achieved here when the WAVE parameter is set to TRIG. In this setting, the volume of the delay is lowered when an input signal is detected and then, when the input level drops, the volume of the delay rises back up to its 'natural' level. This makes the effect more present in gaps between notes/phrases and prevents the instrument signal from being muddied up by a delay signal — a very useful feature. The Reverse switch on the front panel — which I should make clear doesn't access a 'reverse delay' effect such as you might find on, say, a Boss DD-8 — reverses the effect of the dynamic control, so that it changes from ducking to gating: the effect achieved is now the opposite of what I just described, with the delay effect now only present when the input is detected. You can use the Speed knob to change the rate at which the delay fades in (ducking setting) or cuts out (gating setting with reverse switch on).

When it comes to creating beautiful ambient soundscapes, this unit surpasses most others in its field. Long echo textures, ducking to create a bed of sound underneath your playing, or panning delays with high feedback settings are easily discoverable. Some sounds I discovered in this pedal that were pleasantly surprising were the input-triggered pitch-modulation effects. In a similar sense to the ducking delay, precise programming of the delay modulation can lead to the delay's pitch changing intervals when playing stops or starts. Musical intervals can be programmed by ear with the depth setting, or random and sine modulation can be used for subtle or more abstract pitch effects.

The feedback loop when using this pedal as a mono device can be really interesting, but I much preferred using this pedal as a stereo device. Panning effects can be achieved on either the delay or the direct signal, or both or neither, as you prefer. Especially when working with it while wearing headphones, the ability to create such rich stereo effects really made this pedal stand out for me, and a lot of that character is lost in a mono setup.

When diving into the depths of the parameters, it's very easy to stumble upon happy accidents, especially once you've connected the pedal to the software. Unfortunately, except for mapping the expression pedal, the parameters cannot be directly controlled from the software. But once I began experimenting with the expression-pedal mappings — three parameters can be assigned to the expression pedal in any one patch — the effects were transported into new sonic territories. It transforms into a textural looper when set to a long delay with feedback and hints of panning and modulation, and there's a workaround that lets you use this pedal as a sound-on-sound looper: set the expression pedal toe-down position to maximum feedback, have a long delay setting of 6-8 seconds, and bypass the pedal to prevent new signal from entering the loop. Luckily, with its different bypass modes, the default bypass allows the trails to continue, and when the same switch is activated again you can add more to the loop.

The Bottom (Delay) Line

Overall, then, while I wish there were a proper manual, and some might wish that the original's Sample section (see box) had been included here, the 2290 P is a powerful and impressive pedal. It would be a great choice for any musician looking for a creative device to expand their sonic palette, whether live or in the studio, as well as a guitar player looking for a solid, top-notch delay pedal. There are cheaper delay pedals out there, of course, but this goes far beyond an ordinary delay, and it's great value for the asking price. A true classic recreated for the next generation to experience for themselves.

summary

It could use a proper manual but, with top-notch sound quality, versatile routing, precision programmability, easy preset recall, and many unique sounds, this is a powerful and creative delay pedal for a competitive price.

$ $349.
W www.tcelectronic.com

The analogue I/O can be configured either as conventional stereo in and out pairs, or as mono in/out with an effects send and return in the delay feedback path.


Truthful audio monitoring, anywhere.
GENELEC
8300A 216535
Reach the highest standards of audio production – even if you don’t have access to a professional studio.
Our Smart Active Monitors and Subwoofers (SAM™) work with Genelec Loudspeaker Manager (GLM™) software to offer the finest room adaption available, so you can cut out the guesswork and move on with confidence.
Visit genelec.com/studio-monitors
GENELEC®

Samplicity Berlin Studio

Orchestral Convolution Reverb Plug-in

The person behind 'those' Bricasti impulses now brings us Berlin's Teldex Studio...

SAM BOYDELL

Let me start by identifying a common issue: you're working on a big orchestral project that's loaded head to toe with lots of wonderful sample libraries, perhaps complemented by some drier-sounding recordings captured in your studio. Despite a decent arrangement, top-notch instruments and great recordings, because each sound was recorded in a different space your mix lacks coherency, realism and, therefore, personality.

For years, people have tried to tackle this by sending a little of everything to the same convolution reverb, to create the impression of everything being in the same 'real' space. But the essential part that this always missed is that a single IR can't account for the position in the room of different players or the corresponding time delay to the mics. If you have a trumpet section, for example, you really want to hear reverb that corresponds with where the trumpeters stood on the day (or at least where you're pretending they did!). The same goes for the strings, the woodwinds and so on.

Samplicity!

Samplicity's Berlin Studio reverb attempts to overcome this problem by providing multiple impulse responses, all captured at different positions in Berlin's Teldex Studios. This studio is famous for decades of high-quality classical recordings and film scores, and it's also where some great sample libraries have been recorded, including Orchestral Tools' Berlin range. Berlin Studio promises users of those libraries the very useful ability to place other sources in the same space.

It's not just for them, though. These are some of the best impulse responses I've heard to date, and while that is of course partly down to the room itself — it has a beautiful tone, with a natural yet characterful presence that never becomes overwhelming — it's also very much due to the skill of Peter Roos, the man behind Samplicity. I could easily see Berlin Studio replacing the classic 'hall' setting that many of us have in our default templates, and it's entirely possible that Peter is responsible for that hall because over a decade ago he created the hugely popular and free-to-download Bricasti impulses that many of us still use (and which remain available on the Samplicity site). You can also hear his IRs in some virtual instruments from the likes of Heavyocity, Scarbee and Orchestral Tools, all three of which are very highly rated for their sound. Put simply, Peter knows how to make an excellent impulse response.

Into Position

As you can see on the left of the screenshot, the plug-in's GUI includes a soundstage graphic, and you can use this to select one of 18 separate instrument/microphone placements, using the 'traditional' seating arrangement for Teldex. They're organised both by instrument and mic position, from Viola to Percussion Upward to Front Stage Room Mics, and there's lots more in between. This is where this plug-in really shines, because it allows you to account for where and how both the source and the mic would be positioned in a room. These positions are all individual impulses too: the manual states that "to provide the most exact, realistic and transparent results, Berlin Studio does not use IR synthesis, nor reconstruction from different IR sources. Switching between position options is instantaneous, without loading or calculation discontinuities."

On the right of the interface, we're given mixing flexibility similar to that of a console — it's easy on the eye and familiar in use. There are faders for the input signal, the source (used for adjusting the direct signal in the mix) and the stereo output, plus three separate mic positions: Decca Tree, A-B pair and surrounds. Each mic channel has built-in tools for adjusting the EQ, imaging, decay, tail and early reflections, as well as mute and solo buttons and a level meter. The output channel has just the level fader and EQ controls, while the input gets left- and right-channel pan pots.

The source channel has almost the same features as the mic channels, except that instead of the width and scaling controls it has left plus right channel pan pots, an Align knob and a related Auto button. These last

Samplicity Berlin Studio

$189

PROS

- Fantastic sound.
- Very light on the CPU.
- Adjustable microphone time-alignment.
- Simple to use.
- Superb value for money.

CONS

- Lacks global control linking.
- Not as adjustable as the competition.

SUMMARY

A fantastic-sounding, competitively priced orchestral reverb that can add a quality touch of realism to your productions.

Berlin Studio Professional

Just as we were going to print, Samplicity announced the release of **Berlin Studio Professional**. This extended version of Berlin Studio is a separate product that adds immersive mixing of its source and three microphone channels (Decca Tree, A-B pair and the high surround microphones). Users of Berlin Studio can upgrade their licence to Berlin Studio Professional for the price difference, while keeping access to the stereo version.

ALTERNATIVES

In terms of the quality of result, **Berlin Studio** is up there with the likes of **Spitfire Audio AIR Studio Reverb** and **Vienna MIR Pro**. It's more affordable than either of those, though its GUI offers less by way of user control.

two are used to align the source with the signal in the Decca Tree mics that typically form the main image. The knob allows you to dial this difference in manually, while the Auto button calculates it based on the selected position — and it often seemed that when I pressed this button, magic happened!

Ich Bin Berliner?

The Berlin Studio plug-in is intended for use as either an insert or a send effect. The two common approaches would be: inserting instances on your instrument group tracks, and selecting the appropriate point on the Teldex recording stage for each (eg. you'd put an instance on a cello group and select the plug-in's cello position); or setting up different instances as effects busses, to which you can send from your various tracks. I tended to do the latter.

In my experiments, I created a piece using sample libraries recorded in six different rooms and, by the end, I was totally convinced it was a one-room recording. It's worth noting that if you're to achieve the most convincing results then you'll need libraries that allow you to isolate their spot mics but, since not all libraries have that option, I also decided to find out how well Berlin Studio would work with other sources. To do that, I opened up a previous project of mine and simply added 'on top' reverb sends to Berlin Studio's Mid-Stage Wide and Front-Stage Wide microphone placements. It worked well, adding a pleasing sense of cohesion to an already dense mix, and lifting the production to another level. So this reverb is clearly usable with a wide range of material. But the best part is that I can send non-orchestral parts through this room too, just as if I had the studio next door — and it sounds legit!

Despite the superb quality of result, Samplicity position this product at the 'affordable' end of the orchestral reverb market. That's great for obvious reasons, but while the GUI is easy to use it lacks some functionality offered by more expensive competitors that could save the user time. For example, many DAWs don't currently have options for quick-linking controls inside a plug-in or an insert-level mix pot (as in Reaper; an example of what they should all do!). I often wanted to ever-so-slightly tweak the balance between mics, and I had to do this 18 times by managing somewhere from three to six faders on the individual instances. And while the manual states that "the plug-in does not provide a dedicated control for adjusting the ratio between the dry signal and the wet output, in line with its approach of mixing four separate channels", I sometimes missed simple features like that.

Having said all that, those are all minor issues, really, because once the plug-in's sweet spots are found it's unlikely that you'll want to tweak much from there. As I've written above, the sound is what it's all about — and at this price, what a difference it can make!

$ Berlin Studio €169 (about $189).
Berlin Studio Professional €229 (about $255).
W https://samplicity.com

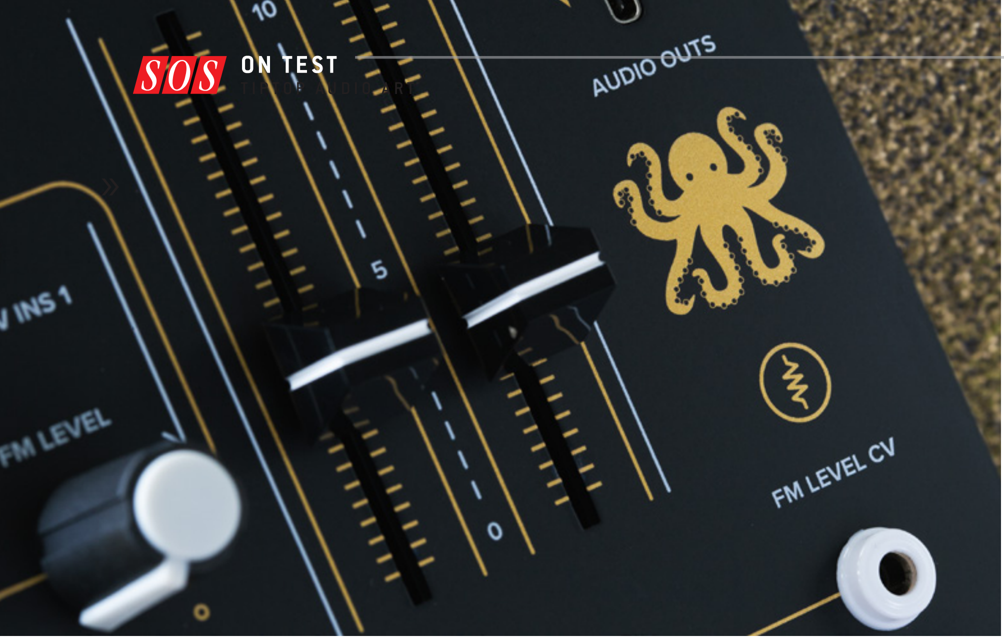

Tiptop Audio ART

Polyphonic Modular Synthesizer System

Veteran Eurorackers Tiptop Audio take on the age-old problem of polyphonic modular synthesis...

WILLIAM STOKES

To the layperson, any modular system looks futuristic. I don't think I've ever shown even a modest system to one without one or all of the terms 'spaceship', 'cockpit', 'telephone exchange' or 'complicated' cropping up in the conversation somewhere. Often it's followed by a question along the lines of 'How do you know what all those cables are doing?'.

It's not often, though, that I myself am faced with a system that strikes me as futuristic; even less so that, at least in the first instance, no, I'm not entirely sure what all those cables are doing. But such was my reaction upon sitting down with Tiptop Audio's ART range of Eurorack modules for the first time. The ART series is the company's most stridently bold product range yet, and that's saying something: it's no mean feat to confidently reissue a line of celebrated Buchla 200 series modules for Eurorack, nor is it to step out with a novel, now ubiquitous type of stackable patch cable. Many a newcomer — including me, many moons ago — will have Tiptop Audio to thank for their first rack and power supply in the Happy Ending Kit, a uniquely affordable bare-bones power supply and rack set.

So, what exactly is ART? Well, in short it's a signal. In fact, it's two types of signal. No, wait, it's just one. OK, I'll come to that. The name is an adaptation of UART (likely a familiar acronym to those versed in electronics), which stands for Universal Asynchronous Receiver-Transmitter — but that's by far the least interesting way of introducing the standard. Tiptop Audio's spiel, if poetic, hardly helps to clarify things: "ART is a new set of ideas, a paradigm," they say. "At the heart of it is the new ART control signal. Building upon the potential of the new ART signal a new generation of modules that had never existed before are now coming to life and give breath to possibilities we could only dream of prior to that."

The Octopus USB/MIDI to ART interface and voice manager.

In short, that means polyphony. Tiptop Audio have come up with a very nifty design — or two — to make it possible to send polyphonic signals around a modular system. It is mightily impressive, but there's even more to it than that. Tuning and temperature regulation for ART oscillators is a cinch — standardised, even — while signal flow is reconfigured and streamlined in choice areas. Tiptop insist they have put a lot of thought into keeping power consumption as efficient as possible and I can believe it.

So: I have in front of me a selection of ART modules. I will say right off the bat that the system looks simply beautiful with Tiptop's signature white and gold graphics on black faceplates with white-ringed jacks — only this time there are sliders, screens and little gold octopodes that remind me of the Spectre symbol in James Bond. The aesthetic is impeccable, with build quality to match. I'll also take a moment to commend Tiptop on their Mantis case, which I had not used up to now and which is so very flexible, sturdy and ergonomic.

Tiptop Audio ART

PROS

- The obvious: a brilliant new signal standard for polyphonic patching!
- The not-so-obvious: ART's other benefits and uses.
- An initial slew of beautiful, well-made modules.
- Well thought-out interfacing options for using CV and MIDI with the ART standard.

CONS

- Not being able to patch mono signals into polyphonic modules feels limiting.
- Fairly high list prices are likely.
- Two types of patch cable can also become awkward.

SUMMARY

Polyphony, standardised tuning and more: a hugely innovative development in modular, with masses of promise. Users and other developers alike may need some convincing, but with enough adopters there could be exciting things ahead — both for Tiptop Audio and for ART.

The first thing to notice are the innocent-looking jacks labelled ART In and ART Out dotted around the system; innocent-looking, perhaps, because their socket colours might been coloured a little more distinctively than light grey on the CV jacks' white, but I digress. Here's the nub: Just one digital ART signal, patched with a good old-fashioned TS cable from an ART Out to ART In jack, sends polyphony and note event messages, along with a rock-solid tuning standard and velocity information. Feel free to read that again.

The ART V/Oct Quantizer.

The second thing to notice is the presence of what look like USB-C sockets. In fact, materially this is exactly what they are; though here the connector is completely reappropriated, with the developer taking great pains to warn us not to confuse the two. These are for the other novel signal format at work in Tiptop's system, though not its titular one: the fun-to-say 'Polytip'. Polytip cables are analogue multi-channel audio cables, capable of bussing an eight-voice polyphonic signal wholesale from module to module (I'm informed it doesn't stop at eight, but all signs and tentacles point to eight being the standard here). See where we're going with this?

The ATX-1 VCO can be switched between ART, V/oct and LFO Modes.

Getting Patching

It all begins with one of two key modules: the Octopus MIDI-ART interface, or the ART V/Oct Quantizer. The former has a USB MIDI input (a USB-B port, mind you, so as not to get confused with the aforementioned sockets) as well as a five-pin DIN MIDI input, and eight separate ART Out jacks. Alongside are eight CV outs for velocity and eight drum trigger outs. Plug a DIN MIDI controller or computer into the Octopus and it'll convert its messages into ART signals; each ART output can send monophonic or polyphonic signals, depending on the destination, and it can also voice-manage polyphony across its different outputs. It can also run a clock from your DAW. Tiptop promise that the Octopus is only the first of many ART controllers in the pipeline, and I'll be very interested to see what else they come out with. After all, this is the very point of entry into the ART world.

The ART V/Oct Quantizer, instead of interfacing the ART protocol with MIDI, interfaces with CV. Patch a CV sequencer or keyboard from elsewhere in your system into one of the four available channels of 1V/oct and gate inputs to meld them into an ART signal. It can also quantise that CV to one of » a litany of scales in accordance with the ART tuning standard.

Luxuriously, I have a choice of four different oscillators for this demonstration: first there's the 8HP monophonic ATX1, another gateway of sorts into the ART protocol for reasons I'll come to. There's the Vortex, a monophonic wavetable oscillator also in 8HP, boasting two vintage D-A converters and capacity to accept an SD card for extra wavetable storage. It follows that the larger Vortex 6 — the developer's first polyphonic oscillator design — is a six-voice polyphonic wavetable oscillator, with a design inspired by the PPG Wave 2.2 hybrid synth from 1985. Finally there's an as-yet-unannounced module whose name I won't divulge at this stage, but I will say it's a magnificent eight-voice powerhouse that will surely assume flagship status for the range, providing all the sonic power and modulation possibilities I could hope for in a polyphonic analogue Eurorack module. Tantalising stuff.

I'd be lying if I said all of these didn't perform brilliantly. Also, if it wasn't already clear from the Octopus' ability to map note events sequentially over its eight ART outputs, it's possible to stack any or all of these into a discrete polyphonic instrument, harking back elegantly to the likes of Oberheim's SEMs, literal building blocks in those earliest days of polyphonic synthesis. At the time of writing, Tiptop Audio offer two Analog Bundles: one with the Octopus and one with the ART V/Oct Quantizer. Both come with three identical ATX1 oscillators, ready to be stacked.

Still with me? It's a lot to take in, I know. Whether covering the basics with ease or producing rich and resonant complex waveforms, or all-out angle-grinder-style aggression, Tiptop Audio have done astonishingly well straight out of the gate with these most crucial components of an ART system. Remember, though: these are only the first fruits of what Tiptop Audio hope will be a long-standing and bounteous area of innovation. So I'm reticent to judge the overall validity and staying power of the ART concept on the individual effectiveness of these modules, however impressive they are.

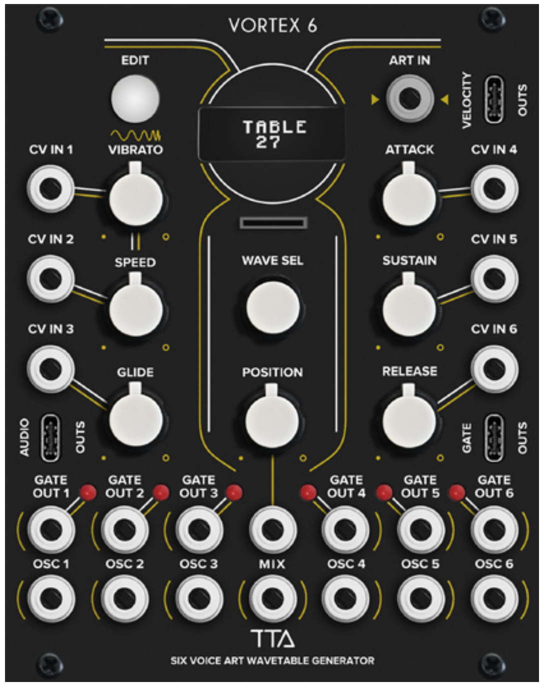

The Vortex 6 is a six-voice wavetable oscillator inspired by the PPG Wave 2.2.

> "Tiptop Audio have taken a bold step out here, and they're on to something."

The Vortex is a twin wavetable oscillator.

ART For Art's Sake

As mentioned, the ATX1 provides a good gateway into all things ART since it's the only module to offer a switch between ART and Volt-per-octave CV control (or it can be an LFO, usefully). The ATX1 — though small and monophonic — encapsulates some key ART-isms. First there's a gate output alongside the audio output. 'That's odd,' I hear you say. And you'd be right. Ordinarily we'd expect a gate to be sent from the controller in parallel to the oscillator signal itself. But since ART sends note-on and note-off messages a little like MIDI, it does make sense to have a gate output on the oscillator itself that can simply be patched onward to trigger, say, an envelope. Neatly, this also helps to keep the cabling from your controller to a minimum.

Another ART-ism on the ATX1 is a button labelled Hold To Tune. That's right: with ART, tuning is standardised, automated and — after warming up — impeccably stable. This stability extends even to migrating the module to other systems with different ambient temperatures, where it'll be able to adapt accordingly. On the ATX1, switching Volt-per-octave mode to ART mode turns the topmost Frequency knob into an octave switcher. Because, well, you now don't need to tune it. I know! Fine tuning is still offered by the knob below, with subtle detuning on one side of its travel distance and semitone increments on the other. And yes, the oscillator itself is analogue. At first I found this all very disconcerting, since it does feel a bit out of my control unlike other 'regular' oscillators. Then I realised it's just very clever indeed. The Vortex 6 has no tuning knobs at all. You just patch it in and, well, there it is. In tune. I couldn't but crack a smile.

So now to the polyphony, and the Polytip cables. They're robust, fairly stiff things — about as robust and stiff as I'd hope an eight-channel cable to be. Tiptop Audio founder Gur Milstein, speaking to *SOS* Editor Sam Inglis at this year's Superbooth event in Berlin, explained that the Polytip's design was essentially in the name of future-proofing, with Tiptop Audio wisely taking the European Commission's standardisation of the connector as its cue.

Milstein and company have done well to keep it self-explanatory with the Polytips: It's essentially business as usual, just pluralised. Patch the ART Out from the Octopus or the Quantizer into the ART In on a polyphonic oscillator. Patch the oscillator's Audio Outs Polytip output into the eight-channel Octogain VCA; then the oscillator's Gate Outs into the Trigger ins of the ADSR Octostages Poly Envelope. Patch the envelope CV Outs into the VCA CV Ins. Presto! Polyphony! And if you want to break this out into discrete mono signals, Tiptop's Octo I/O utility module offers an eight-jack breakout alongside a three-way Polytip mult, so you can retain the PolyTip path alongside if desired.

Too good to be true? Potentially. A few aspects of the Polytip format do complicate things a little bit. Having two different types of cable to think about occasionally feels somehow... awkward, even a little limiting. Perhaps it's because there has always been a wonderful freedom in Eurorack about being able to patch anything to anything, with only one type of cable in the picture (aside from MIDI). With the Polytip, now it's more like your system is divided into modules that do one thing and modules that do another — perhaps with the exception of modulation. At times, with those latter modules, I found myself wishing there was a mono input as well as a polyphonic input, just to have the option when not working with polyphony. Which is often. Modules like the OBX-inspired Octostages Poly Envelope or the (once again) PPG Wave-inspired Octopass Poly Filter do their jobs flawlessly, but it's not an option to use them with discrete mono signals unless you're willing to pay the extra for an Octo I/O on either side. On that note, while many of these modules' list prices haven't been confirmed at the time of writing, the indication is that they will hover at the premium end.

Hitting The Mono-Poly

It all leads me to wonder if the two signal types, monophony and polyphony, were ever begging to be integrated in the first place. Polyphony, while a wonderful thing, should not simply be seen as an upgrade on monophony, marketable as that idea may be. I was disappointed, for example, to see that Arturia's software emulation of the Moog Model D endowed it with full polyphony: for me that takes away a core part of its identity and significantly detracts focus from sound sculpting — presumably in the interest of placating keyboard players. I even wonder if, rather than assimilating polyphony into a conventional modular setup, Tiptop Audio might actually find themselves paving the way for a wholly separate, fully polyphonic modular ecosystem. Not better: different. That, to my mind, would make a lot of sense, and the Octo I/O would still be rendered a highly useful module for when one does want to link the two formats together.

This is really a separate discussion to that of the innovative ART signal itself, which is not just for polyphony and could therefore exist just fine within any system, Polytip or no. Sure, it can handle polyphony with aplomb, but all its aforementioned benefits have the potential to go a very long way while requiring very little user-adaptation. If there's one thing I like more than innovation, it's innovation that works with what I've already got. By this rubric, you might call the Polytip an optional extra for those wanting to dive deeper into polyphonic modular, and if that strikes a chord (geddit?) then all power to you. Gur Milstein has openly said that the Polytip came about in response to the need to manage multiple channels of audio simultaneously, and on that count it's difficult to find fault with what he and his team have come up with.

The fact that the digital ART signal flows down a conventional TS patch cable is very impressive. Of course, let's remember MIDI can comfortably flow down a TRS cable, so needless to say the path of polyphony down eighth-inch jacks is well-worn. One conspicuous

The Octogain is an eight-voice polyphonic VCA and mixer.

area in which the ART signal outdoes MIDI capability, though, is speed. Anyone with experience in sync'ing a modular system to a DAW will know that every so

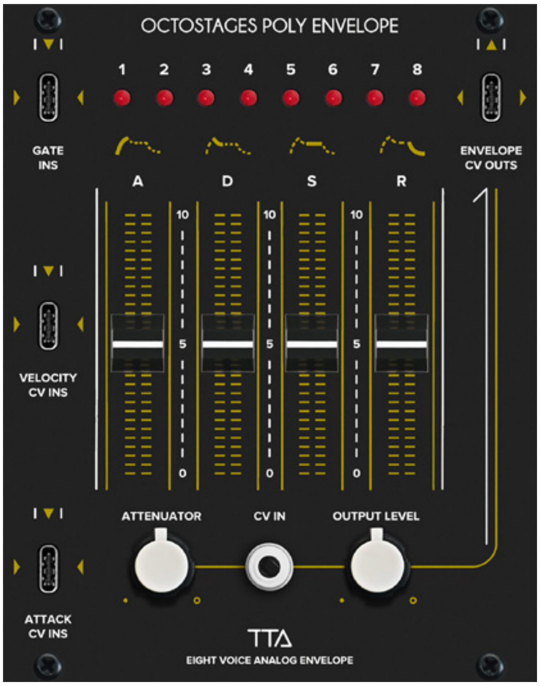

The Octostages is an eight-voice analogue envelope generator based on the Oberheim OBX.

often a little MIDI-induced latency is to be expected and mitigated. But the ART signal can keep up with CV — in fact it's 40 times faster. Which is, to understate it heavily, useful. There are also no channels to worry about. In this regard it's actually simpler than MIDI, and in practice feels similar, if not exactly the same, as simply patching with CV.

Whether that means it's ideal for those same entry-level modular synthesists or recommended only for experienced practitioners, I can't decide. This is because, on the one hand, it certainly serves to have a substantial understanding of how conventional gates and CV signals tend to behave, as well as how to work within its limitations, since that will surely better equip you to know what to do when some of those limitations are lifted. On the other hand, Tiptop Audio have torn up so much of the rulebook in this area that it feels like there would be little lost in learning the ART standard alongside the other basics. It's telling that Tiptop Audio flippantly chalk up the ATX1's ability to accept Volt-per-octave signals as being for "legacy compatibility purposes". In fact, the ART workflow throws into the mix a sufficiently different approach to that of 'ordinary' modular synthesis that I can imagine some seeing it as a different instrument entirely: just one that happens to be in 3U and one that happens to be compatible with other Eurorack modules. There are, after all, plenty of polyphonic MIDI-controlled Eurorack modules out there, as well as digital effects, interfacing modules, multi-purpose modules and so on, and the line of where these intersect with the Doepfer-ordained fundamentals is equally blurred.

This is the moment to point out that Tiptop Audio have made the ART standard universally available to developers, which is noble, granted, but also a plain means of shoring up its survival. Indeed, more than encouraging other designers to follow suit, they're depending on it. After all, an entire protocol needs far more than one developer to establish any kind of longevity. No complaints here about that; this democratic interdependence is one of the things we love about the often cottage industry of modular synthesis — and I do hope other companies get on board with this, big or small. Could we be seeing an ART-powered Roland Juno Eurorack module? Or perhaps more innovation will take place in the realm of controllers, with MPE-style ART controllers of all shapes and sizes specifically tailored to working with modular systems.

I greatly respect the innovation and confidence demonstrated by Tiptop Audio to bring this development to completion. It may not be an easy sell to developers, despite Gur Milstein's assurances that most components are widely available off the shelf, aside from the Polytip cables themselves. But aside from this, many will still argue that the wellspring that is CV, which has begotten modular synth design for 60 or so years, is in no danger of running dry and therefore in little need of this kind of development. I don't think Tiptop Audio are selling ART simply as an upgrade on monophony, though, and I daresay it's actually all its other benefits — tuning stability, near-zero latency and the rest — that will bring even the most hesitant sceptic around into accepting this as an exciting, and freeing, development. Tiptop Audio have taken a bold step out here, and they're on to something. Watch this space.

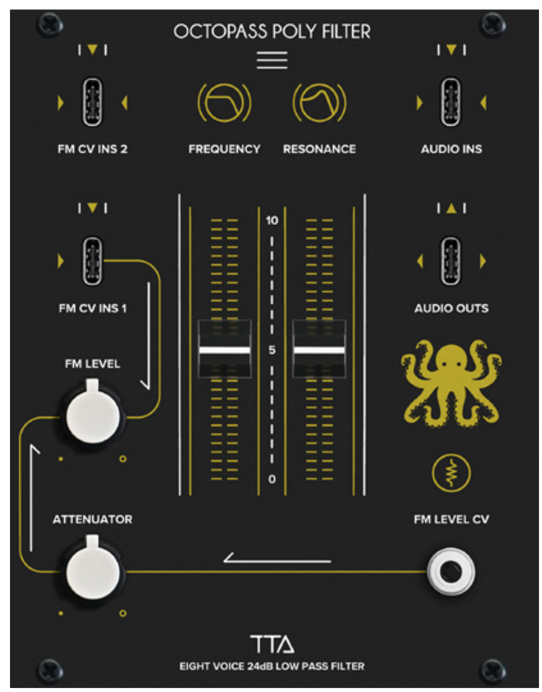

The Octopass is an eight-voice 24dB/octave low-pass filter.

$ See 'ART As An Investment' box.
W www.tiptopaudio.com

ART As An Investment

The ART system is still a work in progress, with more modules on the way and some with prices still to be confirmed. This is what we know at the time of going to press...

- **ATX1:** ART analogue VCO, $225
- **Octopus:** USB/MIDI to ART interface and voice manager, $345
- **Vortex:** ART wavetable oscillator, $245
- **Control Path:** Voice dynamics, $165
- **Quad ART Quantizer:** $255
- **Vortex 6:** Six-voice ART wavetable oscillator, estimate $350
- **Octopass:** Eight-voice 24dB/octave VCF, estimate $350
- **Octostages:** Eight-voice analogue envelope, estimate $310
- **Octo I/O:** Polytip mult and breakout, estimate $85
- **Octogain:** Eight-voice VCA, estimate $310
- **The Analog Bundle S1:** $945
- **The Analog Bundle S2:** $999
- **Polytip Patch Cable:** estimate $12

Ableton Move Groovebox

The Ableton Move puts a carefully chosen slice of Live into a beautifully portable box.

SIMON SHERBOURNE

The Move looks like half a Push, but it's more like an incarnation of Ableton's Note app. Yes, it will work as a Live controller, but it's also a standalone groovebox with four polyphonic instrument tracks assignable as drum kits, samplers or synths. It's primarily focused on a particular workflow: capturing ideas away from the computer, then opening them for further development in Live, or in Note, using local or cloud-based sharing. We'll dive into everything the Move can do, and explore how it's different from other similar format devices, and who it's for.

On The Move

The Push is a notoriously hefty bit of kit, especially the standalone version. It does not belong in a beats-on-the-tube lifestyle ad. The Move, on the other hand, is a true portable: light, slim and tough, and with an internal rechargeable battery. I've found that if a device is missing one of these factors, no matter how functional it is you'll probably not use it around the house or take it on a trip. I count a speaker as optional, but the Move does have one. In or out of the studio I love devices with this aspect ratio, comfortable in front of a laptop or between a keyboard and display.

The panel is minimalistic. Musical input and clip launching are courtesy of four rows of eight velocity-sensitive pads. Above these is a row of touch-sensitive rotary encoders. A larger push encoder is paired with a small screen for various management, browsing and navigation duties. Miscellaneous mode and function buttons grace the wings, but you'll probably have been drawn to the 16 buttons along the front: dedicated sequence gate controls, which you won't find on the Push.

Gear reviews often complain about low-light legibility. The Move's screen and every button and mode indicator are backlit, or hidden if not used in a particular view. There's nothing printed on the surface. Once you've learned the symbology, you can operate the Move in the dark. The lights and pads are very bright, aiding outdoor use. I found the second from bottom brightness setting ample indoors.

Connectivity round the back is also minimal. There's a single stereo mini-jack output for headphones or main out. An input is provided for sampling and there's a built-in mic hiding behind a tiny hole on the front. A Type-A USB port is the only MIDI conduit, allowing the Move to host a USB-based controller directly. It can do input and output, but not both at the same time.

USB-C is used for charging and computer connectivity. Like the Push 3 standalone, the Move requires a high-power USB source for external power and charging. It comes with its own 15W charger, but you can use your laptop charger, or a computer port will keep it topped up if it can push 7.5W. It would be more convenient if it could refuel from any old USB port, but there you go. When untethered, you'll get two to three hours of usage.

Orientation

When booted, the Move presents you with the Set view, where up to 32 projects can be parked. If you tap an empty slot a new Set is created with a drum kit and three instruments with random presets. Sets load almost instantaneously, though not sync'ed or fast enough to be used like scenes.

The Move uses a clips grid Session view for launching the patterns you create on the tracks. To best fit the 4x8 landscape orientation, the Move turns the familiar Session view layout on its side, so the eight clip slots available to each track run horizontally. I was initially disheartened by the lack of scene launch buttons, but quickly realised that you can simply swipe down any row for the same result.

Continuing the minimal aesthetic, the Move's back panel features just a 3.5mm headphone/audio output and input and USB-A and C ports.

Tapping one of the four track buttons to the left of the grid puts you into Note mode, where you can play and sequence the track. Drum tracks grab the left-hand side of the grid for 16 channels of sample playback. On melodic sampler and synth tracks the pads display notes in octave rows constrained to the Set's scale, or in the guitarist-friendly ascending fourths layout if you switch to chromatic.

Starting with random patches leads to instant jamming instead of the usual round of sound browsing, and is one of the main reasons I filled up the Set bank. However, if you prefer you can dive into the preset library, which allows instant auditioning as you scroll. You can undo patch loading, which is a useful luxury.

There are familiar options for recording clips in real time (preserving your timing unless you quantise after) or via the step buttons. But Capture (aka retrospective record) is where it's at across Ableton's growing ecosystem of software and devices. The Move does a good job of guessing your intentions when you hit the Capture button, and turns what you've just been playing or tweaking into recorded notes or automation. If you do this from a standing start it will guess the bpm and length of what you played. If you're playing along to an existing beat it will overdub what's there. I was initially sceptical about this when it was introduced in Live but was won over by the fantastic, visual implementation in Note, and I now rarely use the Record button.

Unusually for a groovebox, the Move allows the creation of sequences up to 16 bars long. I was able to work the same way as I do in Live: making a longer record pass then selecting a subset of the bars to loop as my working clip. It's possible to copy sections, but there's no way to delete bars — I often wished I could trim the start or crop the clip to the loop length.

The step sequencer buttons break you out of the more DAW-like approach of Note, and to some extent the Push. It's easy to drop notes and chords onto steps, adjust their length and velocity, and nudge their timing. Parameter-lock automation works intuitively by holding steps and adjusting encoders. Things are missing that you tend to find on devices in this format. There are no concepts of probability, ratchets, sequence mutations or playback modes. Lanes in drum tracks all share the same loop length so you can't have polymeters within one track's patterns, although you can between tracks. However, there are some advantages to the behind-the-scenes engine being more DAW than drum machine: you can change the step resolution to up to 64 steps per bar and still have 16 bars (so you can at least program those ratchets), and triplet timings are catered for gracefully.

Instruments & Effects

Like on the Push 3 and Note, the Move's instruments (and the effects which work alongside them) are devices taken directly from the Live desktop software. Drum tracks host a Drum Rack device with 16 channels, each powered by an instance of the new Drum Sampler module. On the Move this gives you a 16-track drum machine/sampler packed into a track, with two master effects slots.

The encoders edit whichever pad you last tapped, controlling sample start, envelope and the 'playback effect' module. Touching any encoder shows its function on the screen, so I found myself often swiping across them all to hunt for the right control. A slight frustration is that when playing pads with your right hand, the screen tends to get obscured by your left hand reaching for encoders.

A second page of controls accesses settings and filters per channel, although getting to these is slightly awkward, more of a set-and-forget procedure than a real time modulation tweak. The drum kit has two effects slots: one send and one insert. Either slot can be filled from the bank of effects: Reverb, Delay, Saturator, Chorus-Ensemble, Phaser-Flanger, Redux, EQ and Dynamics.

The melodic sampler instrument leverages Live's trusty Simpler device, again backed up by two effects. This provides pitched, polyphonic playback of a single sample. On the synth side you get Drift and Wavetable. Drift is quite a recent addition and is a two-oscillator virtual-analogue poly. Wavetable is a dual-layer wavetable synth with a huge sonic range, and multiple filter modes and routings.

Drift and Wavetable are deep, versatile, desktop-class virtual instruments. The catch is the same as it is on Novation's Circuit Tracks: user control is limited to eight pre-assigned macros. You can tweak presets, but not get into real sound design or patch creation on the hardware. On the positive side, the presets are very expressive and playable. The Move might not have the MPE capabilities of its big brother the Push 3, but the poly pressure sensitivity is nicely balanced on the pads, and the presets are set up to really respond to them.

Ableton Move

$449

PROS
- Desktop-class instruments.
- Wi-Fi and cloud file sharing.
- Controls Live.

CONS
- No audio tracks.
- Macro control only of the synths.
- Limited MIDI connectivity, and can't be used as a generic pad controller.

SUMMARY
A fun and productive sketchpad companion for Live users.

Move drum tracks are powered by Ableton's Drum Rack, with 16 channels of the Drum Sampler device: a new one-shot focused sampler featuring a multi-mode Playback Effects module.

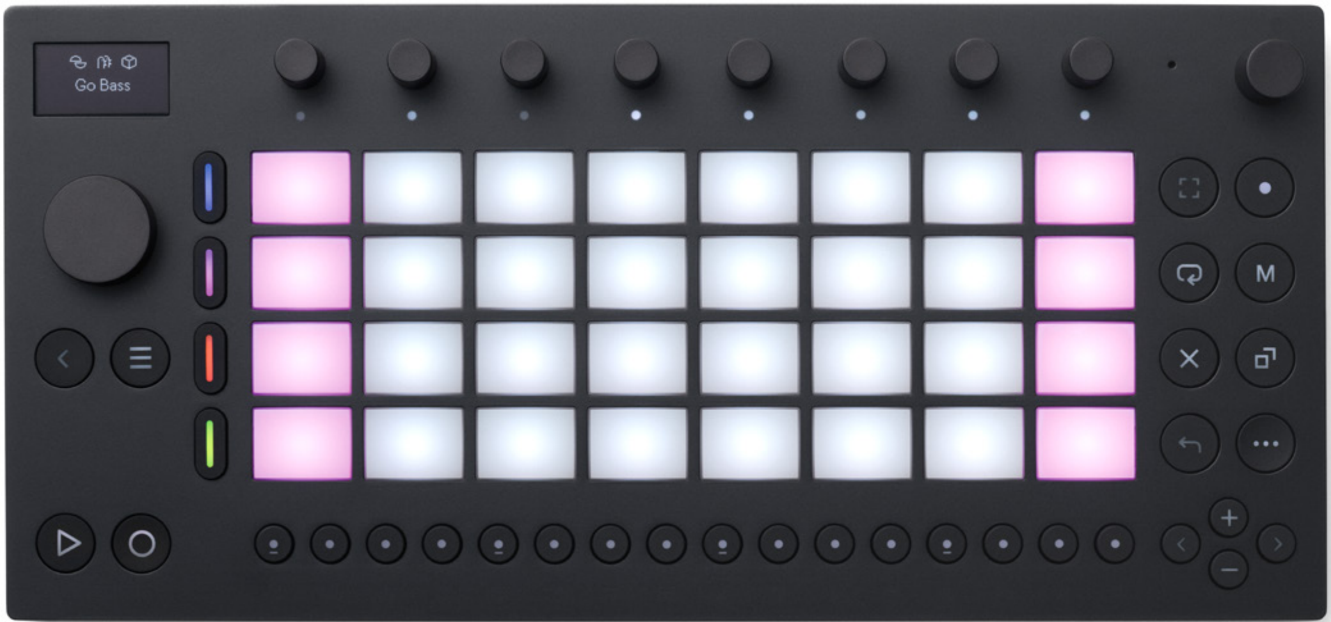

The Move is certainly more portable than the standalone Push, measuring 313.5 x 146.3 x 34mm and weighing 0.97kg.

» The other saving grace if you're squeamish about presets is that you can create your own patches in Live and export them to the Move. This includes setting your own macros, which are simply the Instrument Rack macros you know and love in Live. During testing this was only possible for drum racks and Drift patches, and required careful following of the device chain layout, but eventually this should be available for all the types, using templates built into Live's Library.

Move Manager & Ableton Cloud

Sharing between the Move and a computer can be negotiated over WiFi (yay!) or USB, and just uses your web browser. It's a basic file manager, with directories for accessing your Sets, recorded samples, sample library and instrument presets. There's no requirement to assign things to fixed slots or directory structures, you can just make or upload your own folders as you see fit. The manager shows that the Move has over 50GB of onboard space available for user files.

Move Manager is simple and effective, but it's a slight shame that there's no direct interface with Live. The Push 3 appears as a volume directly in Live's Browser. You can grab Sets directly, and drop presets onto the device, without having to go through a file export/import process as you do with the Move. Move Manager does show some signs of intelligence, though: you can export Sets as audio files directly from the browser.

Sets on the Move can be downloaded and then opened directly in Live, or imported track by track into another session. However, the Ableton Cloud system — which was introduced to support Note app workflows – provides a straight route into Live, bypassing the Manager. The cloud system lets you park up to eight compositions in the ether, which will sync any changes made in the Move or Note, and appear in the Cloud space in Live's Browser.

Once a Set is opened from the Cloud in Live it can't go back, and needs to be saved as a local copy. The workflow that Ableton have designed is pretty clear: experiment, sample, capture ideas on your mobile devices, then bring them into your full-featured production environment to develop further or inject into other projects.

MIDI, Sampling & Connectivity

More devices are starting to include a USB host port, allowing you to directly connect a USB-based controller keyboard. The Move is ahead of the game here in the compact workstation space, as the likes of the Digitakt II, Polyend Play and Novation Circuit don't offer this. A keyboard connected to the USB-A port will trigger whichever track is currently focused, with no messing about with MIDI channels.

What's missing are any other MIDI connections, and the USB port can only do input or output. Again, Ableton's vision for the device appears to be focused on the standalone sketchpad concept. There's not much scope for the Move to be the heart of a DAWless jam rig, for example. Compare this, say, to the new Roland P-6 mini sampler that I also have on the bench: this has physical MIDI and sync ins and outs as well as USB.

Having said this, working alongside a single other device is pretty cool. I connected the P-6 to Move via USB and set the Move to MIDI Out mode. Not only did the P-6 take power from the Move, it also sync'ed clock and transport without any set up. When MIDI Out mode is on, you can elect one of the four tracks to output MIDI, and use that track to sequence the external device. The Move also of course supports Ableton Link, so sync can be achieved over WiFi with DAWs and devices that support it.

Plugging the audio output of the P-6 to the input on the Move, I found audio was also passed through. This persistent monitoring is an option in the sampling settings. Talking of which... Sampling on the Move is really simple. After arming, the Move asks you to press a pad. When you do, it starts recording. In a drum track, the new sample will land on the pad you chose, where it can be trimmed and shaped with the usual drum channel controls. On a melodic sampler track, your new recording is pitched across all keys, with the pad you tapped playing the original pitch.

As well as the audio input, the Move has a built-in mic that you can sample »

» from, and you also have the option to resample. Resampling records the mix out, and starts playback as soon as you tap your destination pad. If resampling is part of your workflow, this makes it really easy to bounce or capture phrases or loops down to a pad. Unfortunately, switching to Session view stops recording, so this is not a way to capture a full song performance. There's a 'lazy chop' feature, where you can tap multiple pads while sampling to slice across several pads. Each pad gets the full recording with unique in and out points, so you can adjust after the fact.

When connected to a computer via USB-C, the Move acts as a class-compliant audio interface, so stereo audio can be passed in either direction. Currently audio output only works when the Move is in Live Controller mode, but apparently a future update will allow sampling directly from your computer or phone in standalone mode. Perhaps Ableton could also address the fact that no MIDI is passed across the USB-C connection. You couldn't for example use the Move as a sequencer or pad controller for anything other than Live.

What It's Not

We've mostly looked at what the Move is; let's look at what it's not. First, the Move is not an audio recorder or looper, and this is top of my wish list for any future development (as it was with Note). It would be great to be able to throw audio loops into a track, but mainly I'd really love to be able to use the Move to capture audio clips, well, on the move. For me this would be grabbing synth loops from hardware, or recording guitar parts. For singer-songwriters and rappers this would make the Move incredibly compelling.

You can do a kind of pseudo audio clip recording via sampling. I experimented with setting up clips with a single trigger at the beginning and trying to trigger sampling during playback, but without any sync'ed or threshold detection modes it's hit and miss. You also need to adjust clip time to match the sample length, and you probably want to switch to gate mode and adjust note length to stop your samples playing out to the end every time you stop the transport. Novation tweaked their Circuit Rhythm to make this workflow easier, but I really hope to see true audio clip functionality added in the future.

"The instruments and effects powering the Move are wicked powerful by groovebox standards."

The other thing the Move isn't is a performance machine. There's no dedicated mixer, you adjust track levels and pads individually within each track view. There are master effects which you can tweak in Session view, but they're not particularly aimed toward real-time creative use. There's no master filter or beat repeat type effects for example, or performance effects like a step repeat for variations and fills.

Likewise there's no arrangement or song mode facilities, and you're limited to eight clips per track, and therefore eight scenes. In fact it's more like seven as there's no clip-stop operation so you need to keep an empty slot as a workaround.

Ableton say that they certainly do want to deliver major feature upgrades to the Move and Note over time, and didn't rule out any of the above, as well as other sequencing and workflow enhancements. However, design is as much about what you don't do as what you do, and the vision for the Move/Note is skewed toward simplicity, fun and ideas generation, with the Push and Live providing full production and performance environments.

Like me, your initial objection might have been, 'only four tracks?' On the face of it this sounds light when most competitive devices sport eight or more. But there are tracks and then there are tracks. Tracks on a Digitakt or drum tracks on Circuit are monophonic channels. Drum tracks on the Move have 16 voices and channels, and all the synths are polyphonic, so more comparable to something like the Roland SH-4D in capacity. What's more, you can play the samples on drum tracks melodically via a 16 Notes view on the right side of the grid, effectively giving you multiple channels of monophonic melodic sequencing in addition to the dedicated tracks.

Also, the tracks are versatile: you can choose to have four tracks of drums if you like, giving you 64 sample playback channels. But yes, it would be nice to have more tracks to make it easy to add layers on different clips for control in Session view. And it's a shame that there isn't complete feature parity with the Note app, allowing you to shuffle projects in either direction. You can open a Note set on the Move if they're both connected to your Ableton Cloud but only the first four tracks will play. Thankfully at least the out-of-bounds tracks are maintained if you round-trip a Set through the Move.

Conclusion

Most of the grooveboxes and beat workstations I've lived with are full of the beginnings of ideas and half-formed tracks that I thought were interesting but didn't do anything with. I like to think I make music with hardware synths, but tracks I actually finish start and end on the computer. Even with my Maschine+ or Elektron boxes that have clear routes to getting ideas out and into a DAW there's routing to set up here, an SD card to move there, plug-ins that need to be in place. The Move removes the friction points by creating projects that are native to your DAW, and are instantly available over WiFi or cloud. So long as you use Live.

The instruments and effects powering the Move are wicked powerful by groovebox standards, and you can capture meaningful musical moments outside of the typical four-bar limit. However, if Live isn't your platform of choice, or you tend to develop and arrange music out of the box, then the Move may not be for you. It lacks the means to develop songs, and the sound design, exploration and performance depth of a standalone workstation.

Ableton have reimagined the groovebox as part of a music creation workflow that extends outside the studio, giving you a little piece of Live to go.

$ $449
W www.ableton.com

Live Control

When connected via USB, the Move has a dedicated Live Control mode, which turns it into a Live session controller. In this mode, the pad grid rotates back to the conventional clip columns layout to match Live. The eighth column controls scene launching, so you get a 7x4 view into Live's clip launcher. The cursor buttons move focus through the Session view. Most of the other buttons perform equivalent functions in Live to their standalone use and the wheel navigates through devices for control with the encoders.

YSM5-2
50 Watt, 5-inch Studio Monitor

LISTEN WITH PRECISION

YSM Series II Studio Monitors

Experience the uncompromising clarity your songs deserve with Yorkville Sound's YSM Series II Studio Monitors. With a tapered front baffle and an updated DSP system, this completely new design is unrivaled in its class.

Now built with the fiber-reinforced drivers, the YSM Series II Studio Monitors are designed to deliver a powerful low-end and clean mid-range. The tweeter and woofer are perfectly time aligned to ensure that there are no phase irregularities.

Listen with precision and create with sonic excellence. The YSM Series II Studio Monitors are available now for your project or professional studio.

YSM5-2

Also Available

YSM6-2
75 Watt / 6-inch

YSM8-2
100 Watt / 8-inch

yorkville.com

ADAM Audio H200

Closed-back Headphones

The H200 are the first headphones ADAM Audio have developed entirely in-house.

SAM INGLIS

Given that loudspeakers and headphones are the only listening options we have, it might seem surprising that few manufacturers make both. Many of the biggest names in studio monitoring have never dabbled in headphone production, whilst the market for phones is dominated by companies like AKG, Audeze, Audio-Technica, Austrian Audio, beyerdynamic, Sennheiser and Sony Professional, whose ranges are mostly devoid of loudspeakers.

This reflects the fact that headphones and speakers actually have relatively little in common from a technological point of view. Design-wise, moving-coil phones share much more with dynamic microphones than they do with loudspeakers. But there are a few manufacturers who have successfully straddled the divide, including Yamaha, Focal and Neumann. And, in 2024, that list has gained two very notable additions. Finnish loudspeaker titans Genelec announced their first ever pair of headphones — which have yet to arrive at *SOS* Towers — and ADAM Audio launched the H300.

However, readers with long memories may recall that these aren't quite the first pair of headphones to carry the ADAM Audio logo. Back in 2019, the company teamed up with Ultrasone to develop a model branded the SP-5; but although these were said to be voiced to match ADAM's loudspeakers, there was a lot more Ultrasone than ADAM in their DNA. So while the H200s might not be quite the first headphones to say ADAM on them, they're certainly the first to have been designed and made entirely in-house. And whereas the SP-5s bore a hefty price tag at launch, the new model retails at a highly competitive $150.

The 200 Club

The H200s are orthodox passive, closed-back headphones designed for music listening and general studio use. They employ a 40mm moving-coil driver with a diaphragm made from a material called polyether ether ketone (PEEK), which is said to offer an optimal balance of rigidity, light weight and temperature stability. These are housed within earcups that look conventional, but which implement a patent-pending internal airflow system that helps to control the low-frequency response.

The plastic shells of the earcups are elongated and tilted back slightly in order to fit snugly over the average human ear, and suspended from aluminium yokes that permit a small amount of fore-and-aft rotation and a greater degree of freedom in the vertical plane. The yokes also slide in

ADAM Audio H200
$150

PROS
- Very good basic sound.
- Useful and straightforward bundled corrective EQ plug-in.
- Excellent value for money.

CONS
- Physical noise transmitted through the headband and earcups can be distracting.

SUMMARY
ADAM Audio have gone back to the drawing board and designed a really good pair of headphones at an extremely tempting price!

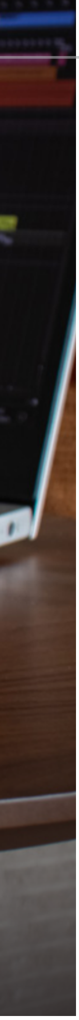

and out of the headband to accommodate heads of all sizes. The H200s ship with faux leather earpads; these are removable, and an alternative soft cloth version is available as a cost option. The headband cushion is also easily replaceable should the need arise.

Compared with other closed-back designs such as the Sony and Austrian Audio models reviewed in this issue, the H200s clamp the sides of the head pretty firmly. This helps to keep them in position and maintain a good acoustic seal, which is desirable for consistent performance, but as a consequence I'd rate them only middling for comfort. The construction also makes them distractingly prone to the transmission of physical noise, especially from the headband.

Like quite a few modern headphone models, the ADAM H200s have a socket at the base of each earcup, allowing the cable to be attached to whichever side you find most convenient. This is done using a locking 2.5mm TRS mini-jack. The supplied cable is 3m long and straight, but short and curly options can be purchased separately if you prefer.

In terms of what's actually in the box, the physical package is completed by a simple padded drawstring bag. However, there's also a virtual dimension to ADAM Audio's offering that helps to set the H200 apart from other headphones in this price bracket.

House Correction

Since the introduction of Sonarworks' Reference (now SoundID Reference) a decade or so ago, the studio world has come to accept that software equalisation can make a worthwhile improvement to the frequency balance of practically any pair of headphones. The cheaper and nastier the headphones, the bigger the improvement, and the only real down sides are practical rather than sonic: headphone correction adds complexity, can increase latency, and may require careful setting up if you want to avoid inadvertently bouncing mixes through it, or if you want to listen to other applications as well as your DAW.

SoundID Reference now supports a vast library of headphone models, and also offers loudspeaker calibration. Indeed, ADAM themselves have partnered with Sonarworks to integrate SoundID Reference calibration into their latest A-series monitors. So it's no surprise that they've decided to offer a software calibration tool for their headphones, but possibly more of a surprise that they haven't simply licensed this from Sonarworks. Instead, ADAM Audio recruited plug-in developers Sonnox — who, like ADAM, are part of the Focusrite empire — to create a new plug-in that goes by the snappy title of ADAM Audio Headphone Utility. This is available for macOS and Windows in all the usual native formats, but there is currently no 'systemwide' version, so you will need a plug-in host application to run it.

ADAM Audio Headphone Utility is, by design, about a million times simpler than the current iterations of SoundID Reference, or rivals such as dSoniq's Realphones and ToneBoosters' Morphit. Menus allow you to specify which type of earcups you're using, and whether you want the plug-in to compensate for subjective level changes »

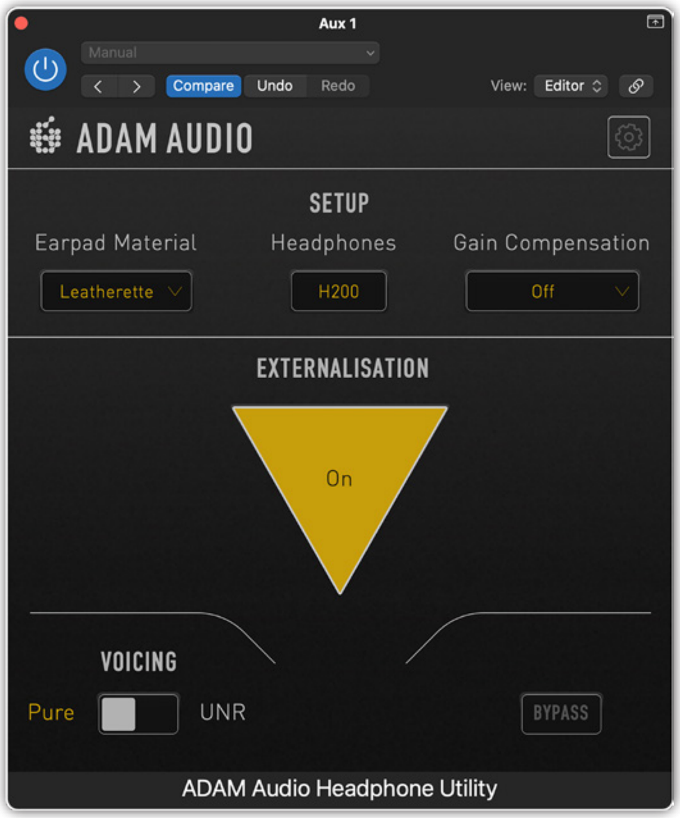

» when the correction is changed or bypassed. Another field currently does nothing but looks as though it might become a pop-up selector for multiple headphone models, suggesting that there may be more in the pipeline. But the most interesting options are the switch labelled Voicing and the large triangular Externalisation control.

The former has two options, which mirror the equivalent settings on ADAM's current A-series loudspeakers. Pure is said to offer the most accurate, neutral presentation, while Uniform Natural Response is "a dynamic, natural-sounding response curve of ADAM Audio's own design, which stems from a variety of iconic ADAM Audio legacy products, including the AX Series". Pedants might wonder why the same voicing can't be both 'neutral' and 'natural'; the two might be better described as 'clinical' and 'fun' respectively.

Externalisation is ADAM Audio's attempt to mitigate some of the other well-known issues with mixing on headphones that can affect translation to other systems. Standard headphones deliver almost complete separation between left and right channels, which is not at all what we experience with loudspeakers, where both ears hear both speakers, as well as reflections from room boundaries. Solutions to make headphone listening more speaker-like range from simple crossfeed, where some of the left channel is fed into the right and *vice versa*, to ambitious binaural processing that models room acoustics, speaker response and the human ear itself. ADAM Audio's offering is somewhere in the middle: a refined crossfeed technology that varies the amount of inter-channel bleed with frequency, in a fashion that is said to be specifically optimised for the H200.

On The Scales

Unusually, ADAM Audio offer a meaningful frequency response measurement as part of the H200s' specifications. With respect to 1kHz, the -3dB points at either end of the scale are at 2Hz and 23.5kHz. On paper this might sound limited compared with the high-frequency numbers thrown around by other manufacturers, but those other manufacturers almost never qualify them with tolerances, so full marks to ADAM for doing so.

In other respects, the specifications of the H200 are fairly typical for a modern pair of closed-back headphones. The nominal impedance is given as 32Ω, and a 1mW input at 1kHz will produce 97.5dB SPL output. This is not quite as sensitive as some rivals, but more than adequate for real-world listening levels with any source. It's also nice to see a distortion measurement given in the specifications, especially when it's a very decent 0.002% THD for a 1kHz signal at 100dB SPL.

Subjectively, I rather liked the sound of the H200 even without ADAM's frequency correction applied. In a world where most headphones seem to offer some flavour of 'scooped' tonality, they stand out by presenting the upper midrange in an agreeably forward way, and their most obvious distinguishing characteristic is a gentle but noticeable emphasis around 2.5kHz. In other respects, they struck me as being impressively close to my idea of 'neutral' right from the get-go, and definitely more so than the vast majority of closed-back headphones in this price bracket.

Some frequency analysis tests on ADAM's own correction utility showed that they, too, clearly view the basic tonality as being reasonably close to neutral. In both modes, the plug-in applies a narrow cut around 6.5kHz, which is presumably compensating for a small resonance or similar gremlin. A slightly

The H200s ship with leather-effect earpads, but a soft cloth alternative is available as a cost option.

broader cut is also apparent from 2.5-3 kHz, and suggests that a couple of dB attenuation is enough to tame the H200s' aforementioned mid-forwardness. Both modes also apply some very broad tone-shaping equalisation. In Pure mode, this consists of a gentle cut below 100Hz and a very broad shelving boost above 4kHz or so, again amounting to a couple of dB at most. Uniform Natural Response delivers more of a 'smile curve' sound, lifting both the lows and highs significantly. It's certainly less flat, but it isn't overdone, and is still closer to my idea of neutral than the default tone of many headphones. I didn't have the opportunity to compare it with any vintage ADAM speakers, so I can't tell you how closely it mimics the classic AX-series sound.

I also really liked the plug-in's Externalisation feature. With complex psychoacoustic processing such as loudspeaker and room emulation, I find you're often a bit too aware that the sound has been processed. The results are rarely completely natural, and in the worst event can sound noticeably phasey or watery. That's not the case here. There's no attempt to make it feel like you're standing in front of a huge pair of Augspurgers in a world-renowned studio, but Externalisation makes the sound stage feel much more speaker-like, and does so without obvious side-effects or down sides. I also like the simplicity of the plug-in, though it would benefit from the option to have it automatically bypass itself during bounces, and a systemwide version would be very welcome.

Right Second Time

With the benefit of hindsight, the launch of the SP-5 half a decade ago feels like a bit of a false start for ADAM Audio. They were too expensive to make an impact in a very competitive market, and they inherited enough of Ultrasone's ultra-bright family sound to make them a Marmite proposition from the sonic perspective. ADAM probably weren't too thrilled with the SP-5 review I wrote back in 2019, so it gives me all the more pleasure to report that they have absolutely knocked it out of the park with the H200s. By taking everything in-house, they've come up with a really compelling blend of practicality, sound quality and value for money. You can spend a great deal more on a pair of closed-back headphones and end up with something much less good — and the ADAM Audio Headphone Utility is a very valuable cherry on the icing.

"There's a virtual dimension to ADAM Audio's offering that helps to set the H200 apart from other headphones in this price bracket."

$ $149.99
W www.adam-audio.com

Teegarden Audio PPC-125

Capacitor Microphone

Teegarden Audio have gone back to first principles to design a versatile studio instrument mic.

SAM INGLIS

Bret Teegarden is a veteran American recording engineer with many high-profile credits in the world of Christian and Gospel music. Over the last decade he's built up a successful second career as a designer of audio equipment, with a range that includes mic preamps, valve and solid-state DI boxes and studio reference monitors, all handbuilt in Nashville. Bret has now teamed up with microphone engineer Steve Mills, who spent 25 years at Crown and played a significant role in the development of the PCC-160 boundary mic, and together they've produced a new range of microphones.

The first of them to reach these shores is the PPC-125 cardioid capacitor mic, which is the fruit of some refreshingly different thinking. PPC stands for 'Pure Path Condenser', and the designers apparently put a lot of work into coming up with an active, transformerless circuit that would offer low noise, minimal coloration and the headroom to handle high sound pressure levels. Another key design focus is consistency of manufacturing, to ensure that stereo pairs are perfectly matched.

In terms of its function and intended role, the PPC-125 is a general-purpose studio mic, but Teegarden Audio envisage it being a first-choice option for acoustic instruments. "We wanted the transient and high-frequency response of a small-diaphragm condenser," says Bret, "and the low-end warmth of a large-diaphragm microphone." To this end, rather than design their own capsule, the duo spent much time auditioning samples from different suppliers until they found one that offered the manufacturing consistency and tonal qualities they were after. They then set to work designing a custom housing and grille that would retain a clean off-axis response.

Happy Medium

Unusually, Teegarden Audio describe the PPC-125 as a 'medium diaphragm' microphone, as its 25mm-diameter electret capsule falls between the dimensions of typical large and small-diaphragm models. It's not unique in this, but whereas other medium-diaphragm mics I'm aware of tend to use a side-address form factor inherited from the large-diaphragm side

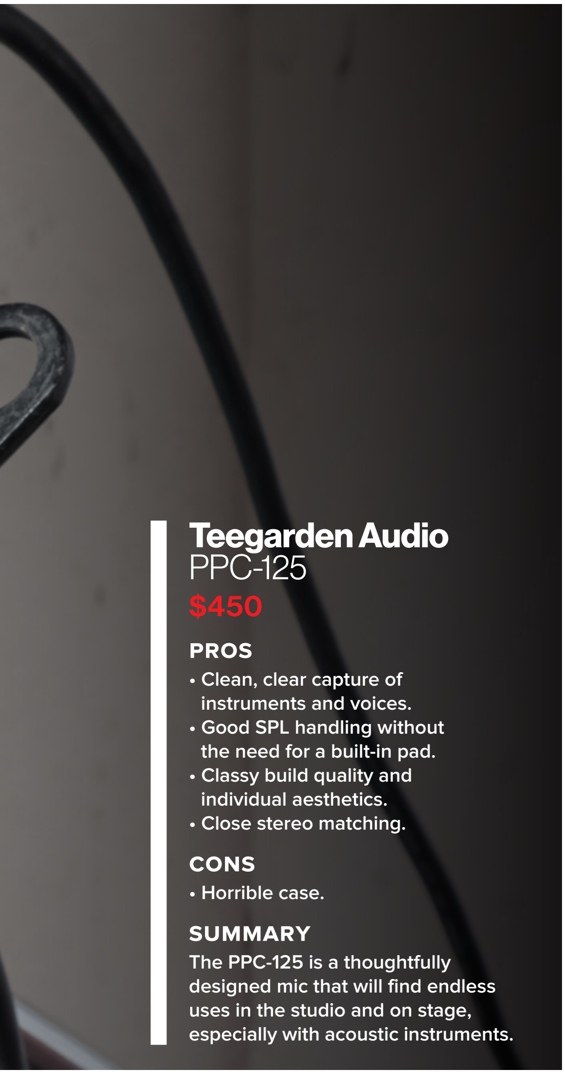

Teegarden Audio
PPC-125
$450

PROS
- Clean, clear capture of instruments and voices.
- Good SPL handling without the need for a built-in pad.
- Classy build quality and individual aesthetics.
- Close stereo matching.

CONS
- Horrible case.

SUMMARY
The PPC-125 is a thoughtfully designed mic that will find endless uses in the studio and on stage, especially with acoustic instruments.

of the divide — one thinks, for example, of the Neumann U89 or the Sony C80 — the PPC-125 is an end-address design. It is therefore a bit too stubby to be considered a 'pencil' mic, but is still much more discreet than a typical side-address model.

The PPC-125 is available individually or as a stereo pair. A pair was sent for review, supplied in a plastic case with clips and foam windshields. The build quality of the mics themselves is impeccable, but I would urge Teegarden to find a better case — the current offering looks like something you'd get with a wish.com socket set, and really undermines the sense that you're getting a quality product.

The published specifications for the PPC-125 are closer to those of a typical small-diaphragm mic than to a large-diaphragm model. No frequency response or polar pattern plots are available, but a deviation of ±2dB is specified across the entire 20Hz-20kHz frequency range. Maximum SPL is quoted as 140dB, although this is in respect of 3% distortion rather than the more usual 1%. Self-noise is a respectable 16dB (A-weighted) and the PPC-125 can operate on phantom power voltages of 24-48 V, drawing a modest 4mA or less. There are no built-in pads or filters.

The PPC-125 has unusually low sensitivity by the standards of modern capacitor mics, clocking in at a mere 6mV/Pa. For reference, the old Neumann KM84 delivers 10mV/Pa, and most newer mics are significantly hotter than that. Personally, I think Teegarden Audio's restraint in this department is to be applauded: the PPC-125 is still more sensitive than typical dynamic or ribbon mics, so you'll never struggle to get enough signal from it, but equally, you can use it on loud sources without fear of overloading preamps and converters.

Point & Shoot

The PPC-125's form factor made it feel natural to compare it mainly with other end-address mics, though I made sure to try it also in roles where I'd normally use a large-diaphragm mic, such as on vocals. The KM84 is an obvious reference point in the former context, and using the PPC-125 side by side with it made for some interesting comparisons. They're both fine mics on almost any acoustic instrument, but they do sound noticeably different. The KM84 pulls off the neat trick of being rich and detailed without being remotely hyped; there's a sort of compact, self-contained quality to its capture. The PPC-125 feels on the one hand slightly cleaner, yet on the other a little brighter and more open, with a broad presence lift running through the upper midrange and into the high frequencies.

> **"The primary market for the PPC-125 will be engineers, studios and live sound engineers looking for a high-quality instrument mic, and I think it very much hits the spot."**

In many cases this presence lift reflected what I'd want to do with EQ to the KM84 in any case, and on instruments such as acoustic guitar, I felt the PPC-125 delivered something closer to a mix-ready sound. Where the source is already on the bright side, though, you might need to be a little careful, and for example cymbals that sounded well controlled on the KM84 were a hint on the splashy side when I tried the PPC-125 as an overhead mic on drum kit. This application also suggested that the PPC-125's cardioid pattern is perhaps a little less tight than that of the KM84, as it picked up noticeably more hi-hat in the 'floor tom side' overhead mic.

With its form factor hinting at the aesthetics of typical stage mics, the PPC-125 might provoke fewer raised eyebrows from vocalists than your average small-diaphragm capacitor model. Its tonality is also well suited to the role; that subtle presence lift is not so different from what you'll encounter in many vocal-oriented large-diaphragm mics. At the low end, there's a controlled but noticeable proximity effect to work with, although it shares with most small-diaphragm mics the quality of being quite susceptible to popping.

At the end of the day, though, the primary market for the PPC-125 will be engineers, studios and live sound engineers looking for a high-quality instrument mic, and I think it very much hits the spot. Stereo matching is subjectively as good as Teegarden claim, and I didn't encounter a source that was too loud for it during the review period. It faces relatively little competition from explicitly 'medium diaphragm' mics, but there are of course numerous pencil mics in this price bracket, including the beyerdynamic MC930, Austrian Audio CC8, Sennheiser e914, Shure KSM141 and many more. Whichever way you place your money, every mic locker should contain at least a couple of mics that can be dug out and pointed at almost anything with confidence, and the PPC-125s fulfil that role admirably.

$ $450 each or $885 per pair
T Teegarden Audio +1 615 454 3981
E sales@teegardenaudio.com
W www.teegardenaudio.com

Kali Audio SM-5

Three-way Active Monitors

Kali Audio's latest three-way speakers look unassuming, but represent a big step up in performance.

PHIL WARD

Back in *SOS* June 2021 I reviewed the Kali Audio IN-5. It was the company's first active three-way monitor, and featured their dual-coincident driver technology, in which a tweeter is located at the apex of the midrange driver diaphragm. The IN-5 driver complement was completed by a relatively small 120mm-diameter bass unit, and the whole system was housed within a modest, easily accommodated enclosure for nearfield use.

If you're wondering why I'm recalling the IN-5, it's because the subject of this review, Kali Audio's new Santa Monica SM-5 active monitor, looks at first glance to be very similar. Its cabinet is of similar dimensions and the driver array, in headline terms, is the same. But despite the apparent similarities, and though it still qualifies as an 'affordable monitor', the SM-5 is a fair bit more expensive than the IN-5. So what differentiates them? The short answer, say Kali Audio, is that there's considerably more expensive electronic and acoustic engineering integrated within the SM-5, and that it consequently aspires to a significantly higher level of performance. The longer answer is... the rest of this review.

Business At The Back

I'm going to start around the back of the SM-5's black-painted and substantially tree-derived enclosure, where there's a balanced XLR analogue input, an AES3 digital input and output (unusually, on BNC sockets) and an RJ45 network socket that enables connection to Kali Audio's new Kali Control Panel application. This comes in macOS and Windows flavours and I'll describe it a little further down.

Downstream of the SM-5's inputs there's analogue-to-digital conversion on the analogue inputs, followed by a DSP module that provides both the EQ functionality accessed through Kali Control Panel and the crossover filtering required to integrate the drivers. And following the crossover, the SM-5's amplification comprises Class-D modules that provide a total of 225 Watts shared among the drivers. While the Kali Control Panel application offers the most comprehensive configuration options, the monitor can also be set up to a basic level using a rear-panel rotary encoder that provides access to input sensitivity, and a DIP switch array that enables selection of some preset LF boundary and HF level trim options.

I don't usually write about a monitor's top and bottom panels, but those of the SM-5 feature arrays of four M5 tapped inserts intended to provide bracket attachment points. Kali Audio clearly envisage the SM-5 finding its way into Atmos monitoring systems, so they've done what they can to make ceiling, wall or frame mounting of SM-5s a reasonably practical proposition. Atmos studio installation is becoming a factor in the design of nearfield monitors, and manufacturers that don't integrate some appropriate bracket mounting points risk being left behind the curve.

Sound On Surround

On the front panel are the previously mentioned composite midrange driver and tweeter, the bass driver, and a letterbox-style reflex port that blends into the lower circumference of the bass driver. Starting with the dual-coincident mid/tweeter unit, one of the great benefits of the three-way electro-acoustic architecture is that bass and midrange drivers can be optimised for their specific roles rather than having to cover both. Bass/midrange drivers are fundamentally compromised at low frequencies because they have to reach up to the midrange, and (this will come as no surprise) fundamentally compromised at mid frequencies because they have to reach down to the bass. The relief from that compromise is immediately apparent in the SM-5 midrange driver, because it has only a very small surround.

The roles of the surround in a bass/mid driver are to enable significant diaphragm excursion at low frequencies, while simultaneously dissipating the vibrational energy that propagates outward in the diaphragm at higher frequencies. It's a tough job, and designers of bass/mid drivers probably spend a good proportion of their careers trying to find surround »

Kali Audio SM-5

$1699

PROS

- **Exceptional subjective detail and imaging.**
- **Very low distortion.**
- **Informative bass.**
- **Effective Control Panel application.**
- **Compact and versatile.**

CONS

- **Slight midrange coloration.**
- **Tonal balance needs a little EQ massaging.**

SUMMARY

Despite its passing similarity to Kali's IN-5, the SM-5 is an entirely different, and extremely capable proposition. Perhaps the most impressive Kali monitor so far.

materials and geometries that can handle both roles without compromise. One of the most fruitful ways of optimising a midrange driver surround is to all but dispense with it. There's typically no need for a midrange diaphragm to move very far, so the only role for the surround is the energy-dissipation one, and that's more about material characteristics than it is about generous geometry. So the surround on the SM-5 midrange driver is no more than a few millimetres wide and actually, at first glance, it appears as if there's no surround at all, because it's so tightly integrated with the blend between the diaphragm and its waveguide.

Guiding Light

Yes, you read that correctly, the midrange driver has a waveguide. In fact, the waveguide works for both the midrange driver and its concentrically mounted tweeter and, again, it's the tiny midrange surround that makes that possible. One of the issues you see with tweeters concentrically mounted within bass/midrange drivers is that the necessarily large roll surround causes significant diffraction of the tweeter radiation. The midrange diaphragm works as an effective waveguide for the tweeter right up until the point where the surround happens, and that undoes all the good work. Take the surround away, though, and you can have a continuous waveguide for both drivers. That's what the SM-5 does. It's a really neat example of electro-acoustic engineering.

Before I move on, I'll note that the midrange driver diaphragm, rather than being fashioned from a high-tech wonder material, is made from paper. Even in the modern world of graphene

The mid driver's roll surround is very thin, which minimises diffraction.

composites, paper is still very hard to beat as a speaker diaphragm material — its particular combination of light weight, stiffness and ease of manufacture is genuinely hard to find elsewhere.

Highs & Lows

The SM-5 tweeter is a 25mm aluminium-dome device that, Kali say, employs a geometry designed to reduce the ultrasonic high-Q resonance that's often a characteristic of such tweeters. The measurement data I took showed no evidence of any high-Q resonance below the 20kHz cutoff; there does appear to be an unexpected sharp dip at around 18kHz, but it looks benign as there's no evidence that it's a resonance feature (and there's really not much music going on at 18kHz anyway). I'll describe my measurements more fully further down the page.

The bass driver, like the midrange driver, sports a paper diaphragm and, as it's optimised for low frequencies, incorporates a generous rubber roll surround. The driver appears to be of conventional manufacture, other than its dust cap being glued to the end of the voice-coil former rather than to the diaphragm. This assembly quirk has become a bit of a Kali signature and it's one I quite like for its logical engineering common sense. Kali say that the bass driver motor system incorporates a Faraday ring and non-magnetic pole-piece cap to reduce magnetic flux and inductance modulation. These are good things to have, and I'd expect the SM-5 to demonstrate good distortion numbers as a result.

Kali Control Panel

Once installed (which is slightly finicky because it requires specific installation permission on macOS), the Kali Control Panel application quickly found the pair of SM-5s I'd connected via their RJ45 sockets to my studio network switch. Kali Control Panel looks and feels slightly less sophisticated than other similar applications such as PMC's SoundAlign, but the necessary functionality is all present.

Like the monitors themselves, Kali Control Panel is designed with multi-channel monitoring in mind, and to that end it can accommodate Atmos systems up to 9.1.6, or custom formats with even larger numbers of monitors. Each monitor that the application discovers on the network is assigned a location by dragging it from the

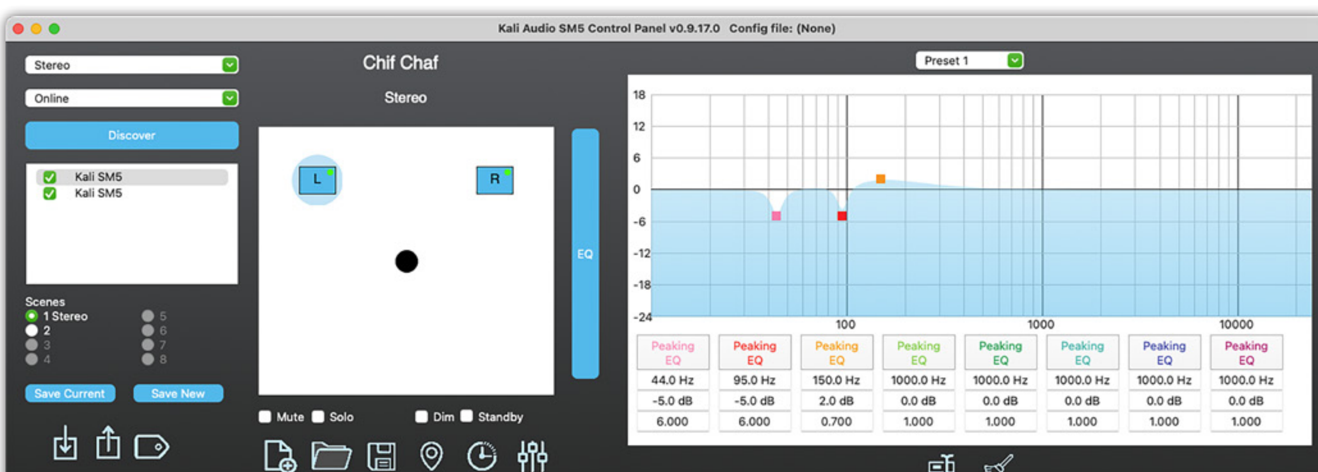

Screen 1: The Kali Control Panel software lets you apply EQ according to personal preference and to correct for room acoustics.


88M
LEGENDARY SOUND ON YOUR DESKTOP
"I love it! It's the best mobile mic preamp & audio interface available by far. I'll be taking this everywhere with me."
Paul Epworth, Record Producer (Adele, Rihanna, Florence and the Machine).
Learn More by scanning the QR code or visit neve.com
Neve
DESIGNED AND CRAFTED IN THE UK BY NEVE® ENGINEERS

Discover list to the listening space alongside. Speakers can also be renamed and identified courtesy of a flashing front-panel indicator, and multiple monitoring systems can be stored and recalled as presets. Each monitor can also be assigned an EQ, delay and level trim.

Once a monitor is assigned, the Control Panel's eight-band EQ functionality becomes available. Each band offers peak, first- and second-order high-pass and low-pass, and high- and low-shelf filtering, with frequency, level and Q controls. EQ curves can be created by grabbing and sliding the frequency 'handles' or defined through numerical entry of frequency, level and Q. It all works very smoothly. If a speaker remains connected to Control Panel over the network, EQ changes happen on the fly, but EQs can also be saved and uploaded to the speaker either over the network or by storing them on a USB memory stick. The SM-5's front panel incorporates a USB socket for EQ uploads.

The EQ is useful both as a means to mould the sound of the SM-5 towards personal preferences and for room optimisation. Room frequency response data captured with, for example, FuzzMeasure (www.rodetest.com) or Room EQ Wizard (www.roomeqwizard.com) can inform the Control Panel EQ settings to compensate for unhelpful trends and specific studio acoustic issues. Without wishing to spoil my subjective feeling about the sound of the SM-5, I found when listening that a little boost around 150Hz was useful, and at the same time I used the EQ to compensate for a couple of known low-frequency modes in my studio room. Screen 1 illustrates the curve I came up with, and Diagram 1 shows its effect on the in-room frequency response of one channel at the listening position. The red curve is pre-EQ and the blue curve is post-EQ. Given more time, I'm sure I would have refined the EQ further, and maybe added a few more filter instances.

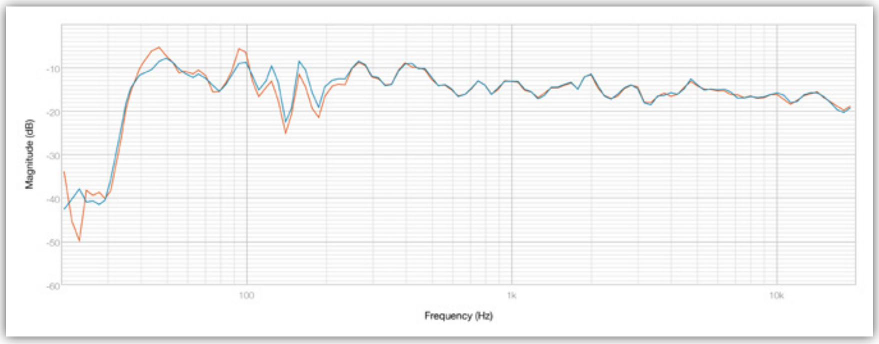

Diagram 1: The in-room response of the SM-5, with EQ bypassed (red trace), and with the EQ settings from Screen 1 applied (blue).

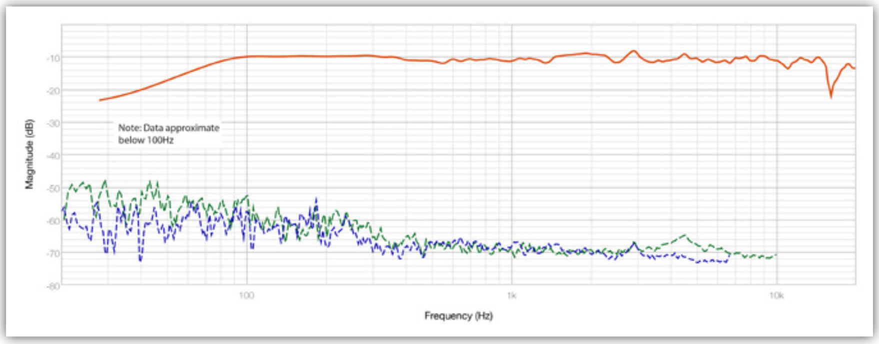

Diagram 2: A quasi-anechoic measurement of the SM-5's on-axis frequency response (red trace) and distortion, measured at 1m at 80dB SPL. Second-harmonic distortion is shown in green, third-harmonic distortion in blue.

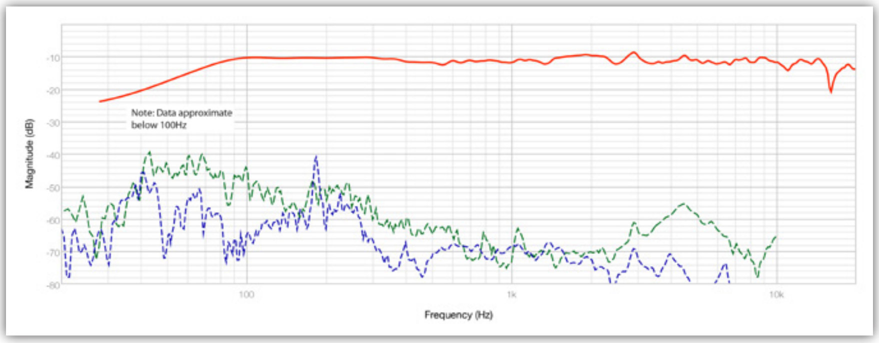

Diagram 3: As Diagram 2, but at 90dB SPL. The distortion rises accordingly, but is still well controlled, and impressively low for a ported monitor.

Measurements

As usual, I took some quasi-anechoic (using a ground plane technique with a 2m mic distance) measurements of the SM-5. Diagram 2 shows the axial frequency response and distortion, measured at 80dB SPL at 1m. The response is nicely linear with no major flaws (apart from the previously mentioned and probably not that significant suck-out at 18kHz) and the distortion is generally very low, pretty much bouncing along the measurement noise floor. Diagram 3 is a repeat of Diagram 2 but with the SM-5 working much harder at 90dB SPL. The distortion rises as expected but is still very well controlled. The distortion performance of the SM-5 is generally very impressive.

Diagram 4 shows how the frequency response varies with mic position: 20 degrees off-axis horizontally and ±10 degrees vertically (with the speaker in portrait orientation). The result, a very tight bunch of curves, is excellent and illustrates one of the huge advantages of dual-coincident drivers: consistency of dispersion.

Finally, Diagram 5 illustrates the results of placing a measuring mic very close to the bass driver and the port exit respectively, and then driving the monitor both gently (equivalent to 80dB SPL at 1m) and quite hard (90dB SPL at 1m). Looking at the driver first (the red and green curves, which are normalised), there's very little difference between 80dB and 90dB. This is as it should be, apart from a little thermal compression between about 90Hz and 250Hz, and a reduction in the depth of the dip at the port tuning frequency on the 90dB curve. The reason for that dip difference can be seen in the port close mic curves, again for 80dB (blue) and 90dB (orange). At 90dB the port airflow has begun to become turbulent, so the port output drops relatively by a couple of dB: the port Q reduces as the monitor works harder, so the port makes less of a contribution to the output. Now, I should emphasise that this kind of port compression effect is generally a characteristic of reflex loading, and all such monitors will ultimately display it to some degree (all ports run out of laminar airflow capability at some point). So this is all more educational observation than criticism. Furthermore, the port response curves illustrate an element of the SM-5 port performance that demands great praise: its lack of any significant midrange organ-pipe resonance above the 45Hz tuning frequency. That's a definite win.

Listening In

I initially found listening to the SM-5 to be a slightly mixed experience. Genuinely exceptional (striking, even) detail, clarity and stereo image focus and resolution was accompanied by a tonal balance that seemed somewhat light on lower midrange. Voices lacked an element of warmth and body, and, with my bass player hat on, I wanted to hear a bit more neck pickup. Electric basses sounded 'all bridge pickup'. In addition, the SM-5 to my ears possesses a hint of 'cuppy' coloration in the midrange that I particularly noticed on naturally recorded voices. It's not a deal breaker, and on much material perhaps wouldn't really be significant at all, but once you hear such things you can't unhear them. This was the point at which I began to use the Kali Control Panel EQ, and adding a little energy around 150Hz not only nudged the overall balance to somewhere much nearer my liking (and I think more usefully neutral in translation terms), it also masked significantly the coloration I'd heard. Once the balance was more to my liking, I was much less conscious of the coloration and could appreciate the SM-5's exceptional reproduction of mix detail. It has the resolving power of significantly more expensive monitors — and a slightly NS10-like way of making midrange details explicitly audible. There's also a highly satisfying coherence to the entire SM-5 midrange and high-frequency experience, and speaking of high frequencies, I wasn't particularly aware of anything missing at 18kHz. It's tempting, of course, to ascribe the subjective coherence of the SM-5 to the integrated

"It has the resolving power of significantly more expensive monitors — and a slightly NS10-like way of making midrange details explicitly audible."

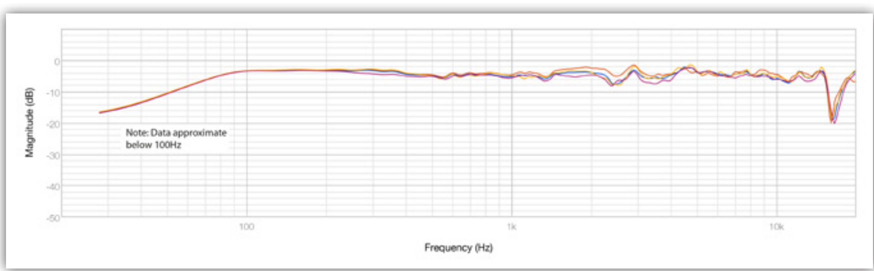

Diagram 4: Comparing the on-axis frequency response (red trace) with measurements taken 20 degrees to the side (blue), 10 degrees above (yellow), and 10 degrees below axis (purple).

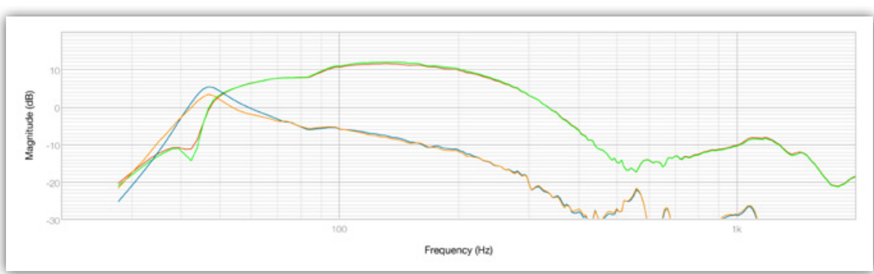

Diagram 5: Close-mic measurements of the bass driver, with the speaker playing at 80dB SPL (green trace) and 90dB SPL (red trace) at 1m. Close-mic measurements of the port are shown in blue (80dB SPL) and green (90dB SPL). The results are normalised, and show impressively consistent port behaviour at different SPLs, with minimal organ-pipe resonance.

ALTERNATIVES

This price band is not short of very capable monitors. Alternatives worth considering might be the **Neumann KH120 II**, **Focal Shape Twin**, **ADAM Audio A7V** and **Genelec 8330A**.

nature of its dual-coincident driver. So tempting, in fact, that that is exactly what I'm going to do. My feeling is that the consistent dispersion and synchronised timing of midrange and high-frequency output — fundamental characteristics of a well-designed dual-coincident driver — pay great dividends.

The more I listened, the more I came to appreciate the SM-5 in the midrange and high end, but I liked its efforts at low frequencies right from the start. It sounds as if the designers made a conscious choice not to be too greedy with low-frequency bandwidth extension (possibly because the bass driver is relatively modest of diaphragm area), and that's a laudable decision, I think. I'd much rather have accurate and informative bass with slightly limited bandwidth than massively extended bass that's light on actual subjective insight. The SM-5 is notably lacking in the pitch, timing and dynamic uncertainty that can sometimes trouble reflex-loaded monitors. It will probably never shake the walls with bass, but it will most likely enable a decent low-frequency mix to be built — and that's way more important.

I was unsure of what to expect when I first unpacked the SM-5s. Conceptually, they seemed rather too close to the IN-5s. How much better could they be? Well, acoustic memory is not always entirely reliable, but I don't remember the IN-5s leaving as great an impression as the SM-5s have. The SM-5 is not entirely free of quirks (what monitor is?) but it is hugely likeable and could easily earn its keep in a wide variety of studio and mix contexts — all the way up to professional Atmos mix studios. It is a great example of relatively cost-effective contemporary electro-acoustic design and engineering, and Kali Audio ought to be very proud.

$ $1699 per pair.
E sales@kaliaudio.com
W www.kaliaudio.com

Win!

LEWITT Recording Bundles

Worth $1380

To enter, please visit:
https://sosm.ag/lewitt-comp-1124

Hailing from Vienna, Austria, LEWITT are one of the most innovative microphone manufacturers around. From their low-noise circuitry to their striking looks, LEWITT's products combine engineering expertise with a fresh approach to design and aesthetics.

For this month's exclusive *SOS* competition, we're giving you the chance to win one of three fabulous LEWITT recording bundles, based around the company's unique Connect 2 interface. A desktop USB device, the Connect 2 is compatible with both computers and iOS devices, providing a microphone and instrument input plus stereo outputs for studio monitors and headphones. With its intuitive touch-friendly top panel, the Connect 2 aims to make capturing pristine recordings easier than ever. An Autosetup feature for the inputs will automatically set your recording levels to avoid noise or clipping, while a choice of three preamp modes gives you Clean, Warm and Vivid sonic flavours. There's DSP built in too, including a compressor and a denoiser module.

First prize in this competition pairs the Connect 2 with LEWITT's innovative new RAY microphone. Uniquely among microphones, the RAY senses the distance between it and whatever you're recording, and then adjusts its tonal balance and gain to counteract the proximity effect — all using the same phantom power supply that biases the mic's one-inch, true capacitor capsule. The RAY also ships with a magnetically attaching pop shield to stop plosives compromising your vocal takes.

Second and third prizes pair the Connect 2 with the LCT 440 Pure or LCT 240 Pro, respectively. These large-diaphragm cardioid capacitor microphones promise flattering capture of a wide range of sources, and are suitable for everything from studio recording to podcasting and live streaming.

To be in with a chance of winning one of these fantastic bundles, simply follow the URL shown and answer all the questions there, by Friday 6th December 2024. Good luck!

Prizes kindly donated by LEWITT
W www.lewitt-audio.com

ART Pro Audio Solo MPA & Solo VLA

Microphone Preamplifier & Compressor

Both of these devices are part of ART's new 'multivoice' series, which aims to provide a wide range of tonal options.

BOB THOMAS

Over the years, ART (Applied Research and Technology) Pro Audio have been one of the more prolific manufacturers of valve-based equipment for home and project studios. The arrival of their new 'multivoice' products, the Solo MPA (a hybrid microphone and instrument preamplifier that offers solid-state and valve amplification) and Solo VLA (which offers a choice of optical compressor with a valve stage or a FET compressor), takes ART's lifetime total of valve-equipped products to a very commendable 24! They've always offered decent bang for buck, and I hope there'll be many more to come. Below, I'll cover the functionality of each of these two new products in turn, before explaining how they performed for me during the few weeks I had them in for review.

MPA Preamp

Arguably, it's quite hard for anything to make the front panel of a mono mic/instrument preamp in a black 1U enclosure intriguing! However, the legends above and around the Solo MPA's control knobs somehow manage that trick very well. Balanced, rear-panel XLR and TRS jack inputs accept mic and line-level signals, while an unbalanced front-panel TS jack provides the high-impedance (>800kΩ) instrument input. The first thing to catch my eye was the gain control: scaled from 0 to 40 dB, this sets the gain of the separate FET-based preamplifier stages and it's an unusual gain range on the face of it. But it can be augmented with two other stages: first, there's a +20dB gain switch; and, second, you can apply up to 15dB of gain at the master output. So you have a generous total available gain of up to +75dB. This unusual arrangement appears to me to be designed to give the Solo MPA the ability to accept either mic- or line-level signals.

Next to the gain knob sits a variable input impedance control for the microphone input's FET preamp. This ranges from 600Ω to a whopping 80kΩ — that's by far the highest input impedance I've seen on a general-purpose mic preamp. As Hugh Robjohns explained back in *SOS* September 2023 (https://sosm.ag/qa-mic-preamp-impedance), lowering the input impedance of a mic preamp will significantly reduce a microphone's output signal. This requires higher preamp gain to obtain a usable output level, which in turn leads to an increase in noise, and maybe also in distortion. The tonal effects of lower input impedances can include: high-frequency roll off; changes in the mic's response to fast transients; a reduction in dynamic range; and an increased prominence in the low to low-mid frequencies. All of which can lead to a subjectively 'thicker' overall sound. Increasing the input impedance, on the other hand, will not only lead to an increase in the level of

the mic's output signal, but will also tend to improve the mic's dynamic range and linearise its response at low, low-mid and high frequencies.

Whether or not these effects are desirable — or even audible — will depend both on the individual microphone concerned, and on the source being recorded. Dynamic (including ribbon) mics and older transformer-equipped capacitor mics are usually more responsive to changes in preamp input impedance, because their output impedance varies with frequency. The balanced mics that we use these days typically have a low output impedance, so that they can drive very long cable lengths with no loss of signal level or quality (unlike high-impedance electric guitar pickups, for which a long cable run can have a deleterious, though sometimes desirable, effect on the sound). This makes them less sensitive to these changes.

Next comes a conventional low-cut (aka high-pass) 6dB/octave filter, with corner frequencies from 10-200 Hz. This is followed by a tilt EQ, whose shelving response acts either side of a pivot frequency (typically in the region of 900Hz), so that when the treble is boosted, the bass is cut and *vice versa*. In contrast to the older tilt controls, such as the ±3dB one in the Quad 34 preamp I owned back in the day, the Solo MPA's tilt EQ has a fairly radical ±12dB gain range.

The Voice control that follows is the source of the unit's 'multivoice' moniker, and sets the balance between the Solo MPA's solid-state and valve preamplifier stages. The idea is that, by varying the blend of these two stages, you can 'voice' the output to suit your preference — be that for the pristine cleanliness of solid-state, or the warmer, more characterful tonality of the Solo MPA's single Class-A 12AX7 valve stage.

A line of three illuminated (when active) buttons bridge the gap between the Voice control and the illuminated analogue VU meter. The first is a +20dB gain switch, and that's followed by buttons for +48V phantom power and 'phase' (polarity) reverse. After the VU meter, there's the aforementioned master output level control, ranging from full attenuation to +15dB. The Solo MPA's maximum output level is a whopping +27dBu, with maximum input levels of +19dBu (XLR) and +17dBu (instrument input).

VLA Compressor

The Solo VLA could be thought of as the Solo MPA's more conventional cousin — it's pretty much a standard, soft-knee, mono compressor. Having said that, though, this is again a 'multivoice' device: it has two parallel, switchable gain cells so as to give you a choice between the speed, punch and transient control typical of a FET compressor, and the slower attack and release of an optical compressor, the latter with added colour courtesy of a 12AT7-based Class-A valve stage.

The Solo VLA's rotary controls comprise threshold (-30dB to +20dB), ratio (1.25:1 to 20:1), attack (0.25ms to 50ms), release (150ms to 3.0s), blend (dry to 100 percent wet) and master output level (full attenuation to +15dB) knobs. Its three selection switches (optical or FET, compression in/out, and the meter display of gain reduction or output level) illuminate yellow when active, with the meter source selector switch turning red to indicate when gain reduction is being displayed.

Flying Solo

In overall sonic terms, the Solo MPA preamp performed very well with a selection of dynamic, passive and active ribbon mics, and both large- and small-diaphragm capacitor microphones, reflecting its specified ruler-flat 15Hz to 50kHz frequency response, and dynamic range of over 110dB. I'm well used to microphone preamps with 70dB (or more) of gain being preceded by a -20dB pad to cope with line-level inputs, but the Solo MPA is the first that I've used with a lower initial gain followed by a 'reverse pad' of 20dB of gain and a final boost of 15dB on the way out — it's unusual to include the output gain in a preamp's specified gain figure, but this didn't cause any problems and there was always plenty on tap.

»

ART Pro Audio
Solo MPA & Solo VLA
$300 each

PROS
- Excellent audio performance from both devices.
- Fully variable preamp input impedance up to 80kΩ!
- Preamp features low-pass filter, tilt EQ and blendable solid-state and valve stages.
- Solo VLA offers both FET and optical compression.
- Optical compressor includes a valve stage.

CONS
- Compressor lacks a side-chain high-pass filter.

SUMMARY
ART's valve-equipped Solo MPA mic preamp and Solo VLA FET/optical compressor offer great performance, value for money and a range of voicing options that you won't find anywhere else.

As with the Solo VLA, the Solo MPA's I/O are paralleled on XLRs and quarter-inch jacks, and can be set to operate at +4dBu or -10dBV.

The Solo VLA's two different compressors (one an optical type with valve stage, the other based around a FET) share the same front-panel control set and rear-panel I/O.

The input impedance control was very interesting to experiment with. It doesn't reach below 600Ω, so I couldn't access exceptionally low impedances, but I could definitely hear changes in the sound of some mics as I slowly raised or lowered the impedance. Lowering the impedance tended to make a mic sound fuller and heavier in the lower frequencies and duller in the highs, which I sometimes found a useful option. The low-pass and tilt controls then proved very useful in optimising this change for a given source, and I found myself spending more time than I'd initially thought that I would in this area! With all my mics, increasing the input impedance to 15kΩ evened up their overall frequency responses to the point that I had to rethink my approach to each mic. I couldn't discern any additional changes occurring above 15kΩ, which means that half of the impedance control's travel was wasted on me, but your own microphones and perceptions may well lead to a different conclusion.

Once I had the input impedance and tilt EQ dialled in for a given mic and source, I started experimenting with the solid-state/valve blended voicings. As you'd expect, the solid-state stage sounded clearer and cleaner compared with the warmer, weightier and (to my ears) more interesting tonality of the Class-A valve stage, and blending these two preamp voices together offers smooth and often beguiling characters to work with. Since this palette of 'blends' interacts with your chosen impedance, low-pass and tilt settings, a need for experimentation arises that can easily lead to several happy hours passing by (perhaps something to do outside of a busy tracking session, then!).

The Solo MPA's wide-ranging impedance control can coax different characters from some mics.

Setting up the Solo VLA compressor was a much less complex task. Plug it in, plug in a source (a single cable from the MPA to the VLA — I used a TRS jack patch cable; it did the job) and hook its outputs up to the desired destination and you're good to go. Other than giving you the choice between FET and optical compression, the Solo VLA can be thought of as a stereotypical soft-knee compressor. The full suite of front-panel controls plus the wet/dry blend control that gives you instant access to parallel compression makes it a really easy-to-operate, flexible compressor.

"Blending these two preamp voices together offers smooth and often beguiling characters to work with."

For a wide range of sources, its FET-based compression delivered just the clarity, transparency, speed, precision and effective transient control I'd expect of a FET gain cell, underscoring the versatility and controllability of this type of compression. By way of contrast, switching over to the optical gain cell, with its Class-A valve output stage, without changing any of the control settings, delivered a more languorous, vintage style of compression, combined with a degree of valve coloration and weight that I found very attractive. As with the Solo MPA, there's plenty to experiment with in the Solo VLA. Indeed, my only criticism really is the absence of any side-chain high-pass filtering, which I'd have found helpful.

But Is It ART?

There's a lot to like in these two units, both of which offer excellent value for money for the performances on offer. Despite the unusual internal gain-structure arrangement, the Solo MPA is a very impressive microphone preamplifier for close-miking applications. The variable impedance, low-pass filter, tilt control and ability to blend between solid-state and valve stages gives you a ton of voicing options. Up until this review, I'd rarely used a mic preamp with more than a 3kΩ input impedance, and I don't own a tilt EQ: and on the back of this experience, those are two situations that I need to change!

It's perhaps overshadowed somewhat in the voicing department by its companion preamps, but don't let that put you off the Solo VLA: it's a very good compressor, and one that I'd be happy to own. Its FET compression delivers exactly what you'd expect it to and its optical compressor with added valve-derived colour does likewise. Having the two compression styles in one box is really useful, in terms of flexibility of use and experimentation, and also in terms of value for money.

If you're in the market for a versatile mono microphone preamp and/or a mono compressor, or perhaps want a couple of options to add colour and versatility to your collection, you really should consider the Solo MPA and Solo VLA. There's a lot to be gained from having them to hand.

$ $299 each including VAT.
T Yorkville Sound +1 716 297 2920.
E sales@artproaudio.com
W https://yorkville.com
W www.artproaudio.com

i73® PRO
The first ever USB-C Audio Interfaces with
73 style Class A preamps
Heritage Audio
The sound of yesterday
built for tomorrow.
RAD
GLOBAL DISTRIBUTION
NEW YORK CHICAGO BATON ROUGE SHENZHEN PHUKET MADRID
Heritage Audio Full Analog Experience
i73 PRO EDGE
i73 PRO ONE
i73 PRO 2
Try our native plugins for free TODAY!
BRITSTRIP
HERITAGE TAPEOPLEX
HA15 PRO
HA1200 TAPESAT
SMALL RECORDING AMP SERIAL # C 17744
HA 240 GOLD FOIL VERB
14-day demo trial

Lynx Hilo 2

A-D/D-A Converter

This updated version of Lynx's mastering converter offers a much improved touchscreen, a host of new filtering options — and the best D-A converter performance we've ever measured!

HUGH ROBJOHNS

When launched in July 2012, Lynx's Hilo converter established a new benchmark level of performance that has endured for over a decade. Now, the company have raised that bar once again with a new, substantially updated version: the Hilo 2. At first glance, aside from a fashionable dark-grey colour scheme replacing the original silver, and the 2 suffix being added to the name, the new model is visually indistinguishable from the old. But much has changed beneath the skin, all aimed at eking out even higher technical performance and delivering enhanced usability and functionality.

Overview

Users of the original Hilo will immediately feel at home, as the operation and connectivity are largely unchanged. This remains a two-in, six-out, mastering-grade A-D/D-A converter, combined with an internal 16x16 digital routing matrix. The front-panel touchscreen provides direct control and display functions, but a Mac/Win computer app can be used for remote control.

The front panel remains very minimalist, with just a standby on/off button, assignable rotary volume control, and quarter-inch headphone socket. As before, every aspect of the machine's configuration and operation is controlled through a large, colour touchscreen, but this 480x272-pixel display is a major upgrade: now an IPS (in-plane switching) type with capacitive touch-sensing, it has been totally redesigned to be more responsive to touch, brighter, and with improved off-axis viewing.

In my review of the original Hilo (https://sosm.ag/lynx-hilo) I noted that "my only difficulty was with the touchscreen interface, where my podgy fingers only needed to be slightly off centre of the relatively small virtual buttons on the routing pages to either get no response or activate the wrong source". I had no such difficulties with the Hilo 2, which responded quickly, precisely and reliably to my every prod without my even thinking about finger placements. So, this new screen is a very worthwhile upgrade!

As far as I can tell, the Hilo 2's current firmware (v1.3, released August 2024) is derived from the most recent version for the previous model (v8.13), and the core functionality and graphics appear almost identical. One useful key difference that I spotted early on, though, is that the new converter hardware can be configured with a wide range of different linear- or minimum-phase filtering modes. The excellent, 70-page manual explains these different modes well but, as a quick overview, the minimum-phase options mimic ideal analogue filters, complete with their potential for 'time smearing' and altered waveshapes due to different phase-shifts versus frequency.

The linear-phase options maintain the phase relationships between different frequencies (and thus preserve waveshapes), but introduce pre-ringing to the impulse response (ie. the filter creates an output before the wanted signal arrives, which isn't something that can happen in nature!). Some people argue that pre-ringing is audible and accounts for the difference in perceived sound quality between analogue and digital systems. For those who subscribe to this view, the availability of minimum-phase alternatives will be appealing.

Both 'fast' and 'slow' options are available for both types of filter, referring to the steepness of the filter slope. The fast filter modes essentially provide a slightly wider and flatter signal bandwidth, while the slow options start to roll off a little earlier but much more gently, and the impulse response ringing is of shorter duration and reduced amplitude as a result. Interestingly, there's also a linear-phase mode with an 'apodising filter' setting (see the 'Apodising' box).

Inputs & Outputs

Other than the front-panel headphone socket, all I/O connections are made on the rear panel and are identical to the previous model. Stereo line-level electronically balanced analogue connections are on XLRs, with monitor outputs on a pair of quarter-inch TRS sockets. As with the original version, the analogue line input and output reference levels are adjustable, reaching 0dBFS for signals between +18 and +24 dBu to match professional equipment, or 0 to +6 dBV for consumer gear. The monitor outputs are also configurable for a maximum level of either +24 or +10 dBu.

Digital audio in and out is either via XLRs for AES3 or RCA phonos for S/PDIF. Both formats are active simultaneously for the outputs, but only one can be selected at a time for the digital input. A pair of Toslink optical ports is also provided and can be configured either for ADAT or optical S/PDIF formats. All the usual sample rates are supported up to 192kHz, although the ADAT interface only supports single-rate (eight-channel) and double-rate (four-channel) modes. Similarly, the USB interface provides 16 channels at single and double rates, but only eight channels at quad rates. Word-clock in and out is available via a pair of BNC connectors, and power is through the ubiquitous IEC mains inlet (100-240 V AC), or a four-pin XLR for external battery power (9-18 V DC).

As with the original model, an optional 'LSlot' interface card can be installed providing up to 16 channels in (usually configured as eight stereo sets) and 16 more out, via a USB, Thunderbolt or Dante connection to a computer. Drivers are available for the Thunderbolt and USB options, for both macOS and Windows. The Hilo 2 operates as a Core Audio device using USB on a Mac, and supports both WDM and ASIO on Windows. The Thunderbolt interface requires a dedicated Lynx driver on both platforms.

Screens

Most of the time, the display screen shows audio signal levels in various ways, and there are four options. The All I/O view does as it says, using two rows of vertical bar graphs to show the current signal levels on each and every input and output. The Horizontal view presents two larger horizontal stereo bar graphs, configured to show selected input and output levels. Perhaps the most attractive option is the Analog VU, which is a virtual pair of classic VU meters whose sensitivity is adjustable between -3 and -30 dBFS for 0VU. Again, any input or output source can be assigned to feed the meters. Lastly, there's an RTA (real-time analyser) option, with two rows of vertical bar-graphs indicating the stereo level across 30 frequency bands for any selected input or output. The sensitivity is also adjustable in 3dB steps.

The Audio screen provides access to functions like the desired sample rate and sync source, analogue line input and output reference levels, Toslink optical interface mode (S/PDIF or ADAT), selected digital input source (AES or S/PDIF), and SRC on/off for the digital input (synchronising an external digital source to the Hilo's current clock rate).

A System page accesses functions like the display language, backlighting, USB modes (eight or 16 channels), digital filtering modes, and rotary control assignments (normally it controls the monitor and/or headphone outputs, but it can be used to control other outputs or groups of outputs). Information is also displayed here

Lynx Hilo 2
From $3599

PROS
- **Stunning audio performance.**
- **Redefines the benchmark standard of converter performance.**
- **Greatly improved touchscreen.**
- **Immensely versatile and impressive converter/router system.**

CONS
- **This sort of quality costs!**
- **Signal processing and metering functions haven't yet expanded to meet the open-ended promise of the original design.**

SUMMARY
The Hilo 2 redefines the state of the art in converter performance, while retaining all the much-loved versatility and capability of the original.

The expansion slot accepts USB, Thunderbolt or Dante cards, with USB being the most affordable option.

The new touchscreen is a significant improvement over the one used in the original Hilo.

» on serial and firmware numbers, and factory defaults can be restored.

Particular system configurations and routings can be stored and recalled as Scenes on another menu page with eight options altogether; four are factory settings, and four are user configurable. The last menu page allows signal routing from any combination of inputs to any number of outputs, allowing multiple sources to be mixed and fed to multiple destinations, if desired, with adjustable signal levels. The setup process involves highlighting an output and then selecting (or muting) the required input(s). It's easier than that might sound, and routing is performed as stereo pairs of channels. Helpfully, selected sources can also be summed to mono, or either channel can be routed to feed both output channels, if needed.

Bench Tests

The best modern converters really are now performing on the very edges of human perceptibility, meaning that incremental technical improvements are more easily identified through test-bench measurements than through conventional listening tests. Regular readers of my converter reviews will know that I value the AES17 dynamic range test as a good overall indicator of converter performance. I have found it to give consistent and reliable quantitative assessments of converter performance that tally closely with my personal subjective impressions. This is probably because, in order to achieve a good AES17 score, every single aspect of the converter's design and manufacture has to be optimised, including the analogue circuitry, the grounding arrangements, the power supplies (particularly the analogue/digital supply isolation), the clocking arrangements, digital stream recovery, and more. The smallest technical weakness in any area very quickly degrades the AES17 measurement.

When I reviewed the original Hilo, it scored 121.3dB (A-weighted) for the A-D converter and 120.5dBA for the D-A conversion. Both are excellent figures which placed the Hilo comfortably within the top 10 of all the converters I'd measured at that time. I still rate anything above 118dBA as being of genuine professional mastering quality, so to have achieved this standard 12 years ago is highly impressive. Nevertheless, converter chip manufacturers have continued to improve the performance of their devices, and product manufacturers continue to build converter systems that are capable of extracting all the potential performance of these newer chips.

This is all clearly evident in my bench tests of the new Hilo 2. In my tests, this newer model gained 2dB of performance on the A-D side, measuring an exceptional 123.1dBA. That's second on my list only to the incredible RME ADI-2 Pro (124dBA). The D-A side exhibited an even greater improvement, reaching

Apodising

Apodising filters have been used in spectral analysis functions for decades as a way of reducing the effect of the abrupt start and end of analysis time windows. In digital audio applications, an apodising filter is employed to reduce or remove linear-phase filter pre- and post-ringing that is caused by the D-A reconstruction filter and that is already embedded in the digital signal by the A-D's linear-phase anti-alias filter.

In simplistic terms, an 'apodising filter' is essentially a specific type of minimum-phase low-pass filter, which is critically designed to place its first deep null at the Nyquist frequency, but there's rather more to it than that.

As we know, digital audio signal paths inherently involve both anti-alias and reconstruction filters in the A-D and D-A converters, respectively. In any system containing a series of filter stages, the system's overall frequency response is the product of the individual responses of each filter. So, for example, if the chain contains two first-order (6dB/oct) filters in series with the same turnover frequency, the overall response will be a second-order (12dB/oct) filter.

However, in the time domain, the overall impulse response of the whole chain is defined by the convolution of the individual filter responses, the significance of which is that it's possible to design a filter which, through convolution, exactly counteracts the time-domain characteristics of standard linear-phase filters. This is the 'apodising' filter and, when introduced into the signal path, the overall system's impulse response becomes shorter and smaller than that of any of the individual filters. In this way, it's possible to almost completely remove both pre- and post-ringing associated with conventional linear-phase filters.

Naturally, there are trade-offs; the most significant is that the apodising filter has to start rolling off slightly lower in order to achieve that deep first null at the Nyquist frequency. But we're really only talking from roughly 19.5kHz instead of 21.5kHz, in a 44.1kHz sampled system, and I think that's a very small price to pay for a fully corrected time-domain performance — and in systems with double or quad sample rates there's no compromise involved at all! Secondly, as the apodising filter is still a minimum-phase low-pass filter, there will still be some post-ringing, but this is natural and, as it generally falls within the human temporal masking range, inaudible.

Apodising is a very clever and proven practical solution, first described for digital audio applications by Peter Craven back in 2004, and I remain baffled, 20 years on, that so few manufacturers of digital products employ it — it's good to see it as an option in the Hilo 2.

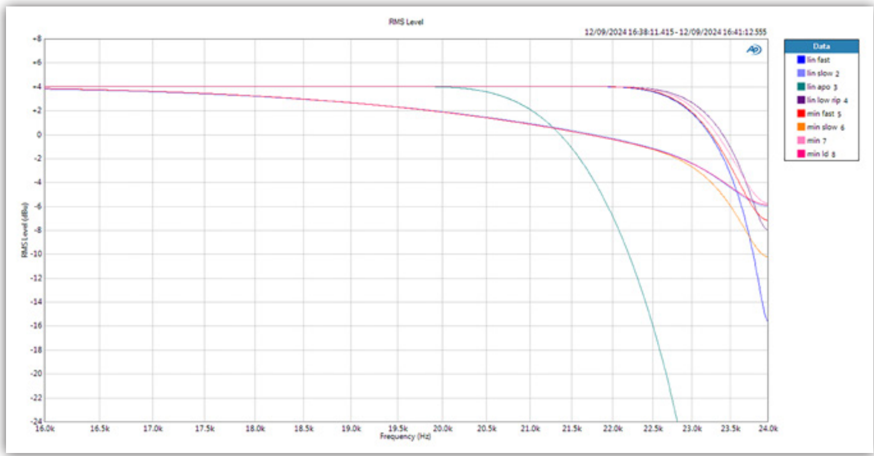

This Audio Precision chart illustrates some of the available D-A filter mode options. The red, orange and pink curves are minimum-phase types while the blue and purple curves are linear-phase. The fast modes start to roll-off steeply from around 22kHz, while the slow options start attenuating very gently from about 16kHz. The default option of Linear Fast gives greater attenuation at the Nyquist frequency (48kHz) than all other options apart from the apodising filter (green).

a whopping 128.1dBA, making it my new leaderboard winner — 2dB better than the Merging Hapi interface, which was the previous leader. These are exceptional figures, genuinely pushing against the boundaries of what is possible in a practical commercial product, and Lynx are to be congratulated on their striking achievements.

Similar improvements are apparent across many other bench test measurements too. For example, the THD+N figure for the analogue line inputs is 3dB better (now -117dB), while that for the line output is 11dB better (-120dB). Crosstalk is also improved for both line input and output, but especially the latter, which is 8dB better at -143dB (at 1kHz).

I couldn't find any significant changes in the performance of the monitor or headphone output sections, which both achieved THD+N of -107dB and crosstalk around the -135dB (at 1kHz) mark. Perhaps not quite as impressive as the main line outputs, but still a very high standard.

> **"These are exceptional figures, genuinely pushing against the boundaries of what is possible in a practical commercial product, and Lynx are to be congratulated on their striking achievements."**

ALTERNATIVES

The only direct alternative, in terms of both core functionality and class-defining audio quality, is **RME**'s **ADI-2 Pro**.

Impressions

The new Hilo 2 undoubtedly delivers an impressive step up in its technical performance compared with the original model, regaining the top spot in my own tested converter league table and highlighting the current state of the art. This is to be highly commended.

Many of the 'essential' items I described in a 'wish list' of functions in my review of the original model have also been incorporated through various firmware updates over the last decade, and have been carried over to the Hilo 2. For example, things like adjustable VU meter sensitivity and ADAT channel selection. Sadly, most of my other suggestions haven't (yet?) been implemented and, given the open-ended development promise of the original concept and 10 years of development time, I found this disappointing. I'd still like to see, for example, basic monitor controller functions, stereo width manipulation, L-R channel swapping, True-Peak and Loudness metering, and a cross-point matrix display option for viewing and setting signal routing paths.

Nevertheless, the Hilo 2 is a phenomenal, benchmark-setting professional converter that offers stunning sound quality, superb features, great versatility, and even easier operation than its forebear. The Hilo 2 is ideally suited as the heart of any professional mastering suite, transfer room, or high-end audiophile's system, delivering standard-defining audio quality in a wonderfully easy to use and configure format. The best just got better.

The filters on offer include an apodising option, intended to reduce the pre-ringing inherent in linear-phase filters.

$ Hilo 2 including: USB card $3599; Thunderbolt 3 card $4199; Dante card $4299.
W www.lynxstudio.com

MODULAR

Blaknblu Foxtrot

Eurorack Module

You may have heard it said that you can never have enough VCAs, but Blaknblu would suggest you can never have enough filters. The Foxtrot quad-multi-mode filter has 12 of them. At 38HP, it's very wide, but the front panel accommodates four individual and complete multi-mode filter modules, each with a choice of three topologies. The idea is that you've got four independent and identical filter channels. Each filter can morph from band- to low- to high-pass. You have two mono inputs, dual or stereo outputs, two CV inputs for the cutoff and an assignable 'Aux CV' that can control the filter mode morphing, resonance or panning. Lastly, there is a Boost switch for a little bit of overdrive. Behind the scenes is some clever routing to link the filters into a parallel bank controlled by a single cutoff, if you so wish.

Each filter can be one of three types. First, we have the friendly bounce of the Moog-style ladder filter, followed by the unruly nature of a Sallen-Key (SK) filter from the likes of the Korg MS20 and, finally, a buttery state-variable (SV) filter similar to the Oberheim SEM filter. The Boost switch is tailored to the specific filter design to replicate the behaviour of the original circuits when pushed to distortion. Throughout the Foxtrot, the topology is digitally modelled virtual analogue, as opposed to actual analogue circuits, but there's nothing in what you hear that would suggest these are anything other than classic analogue filters.

One thing you don't see very often is CV control over the filter mode. The modes themselves are helpful in broadening the appeal, but being able to sweep between them with an LFO or source of rhythmic voltage gives it a whole other personality.

All of this is great in a single-filter module, but this has four. Using them independently is naturally useful in any rack. If you were looking into polyphonic modular, then a slab of four identical filters would be perfect for a four-voice system. In fact, Blaknblu also do a similarly featured quad VCA called Alpha Pro that would complement that endeavour beautifully. However, it's the linked operation that gives the Foxtrot its superpower.

Flick two or more of the filter channels to On, and they will become linked. The audio outputs are summed to every output, and the filter cutoff is added to the leftmost filter meaning that you can sweep all of the linked filters from the left filter. With a single input, you find that you can add filter after filter to your signal chain. You can tune in four band-pass filters and then shift them together, or you can link a cascade of different filters, all modulated in different ways, creating complex spectral shifting. Even running two filters side by side with an oscillator in each gives a dramatically different tone when linked and unlinked.

It all hangs together heroically in the face of so much filtering, especially when pushed against the limits. If I had to find a fault it would be that the patch socket layout is a little unintuitive, and I keep grabbing the mode knob, thinking it's the cutoff because it stands out more. I'm also not convinced I'd be making full use of its features very often and in my overloaded rack every module needs to earn its space. However, it does sound pretty fantastic.

If you're drawn by the idea of these three filter models but can't quite cope with the enormity of what the Foxtrot has to offer, you'll be pleased to hear that Blaknblu have just released a single-channel stereo version called the Foxtrot Duo. This feels like a solid move, as for all the Foxtrot's versatility I can't help wondering whether the scale and scope is a bit much for most people's racks.

Robin Vincent

$	$573.99
W	www.blaknblu.com

Weston Precision Audio Phase Animated Oscillator

Eurorack Module

Weston Precision Audio's new module is called the Phase Animated Oscillator, aka the PAO, but just what exactly is a phase animated oscillator? There is a bit of physics to consider when thinking about the functions of this module, since we don't have oscilloscopes built into our heads. Basically, when two sound waves are perfectly aligned, they are said to be in phase; they sound the same and there is no audible difference. The PAO has two oscillator output banks that allow the phase of a waveform to be shifted and modulated with the other, or animated, essentially, adding harmonics to create a more complex sound.

At its heart the PAO is a triangle-core VCO with two sets of waveform outputs. The oscillator bank on the left is normal waveform outputs, and the bank on the right is the phase-animated waveform outputs. Both waveform banks contain sine, saw, triangle and square wave outputs, but Weston have added two other unique outputs on each side. On the bottom of the module the left set of six waveforms

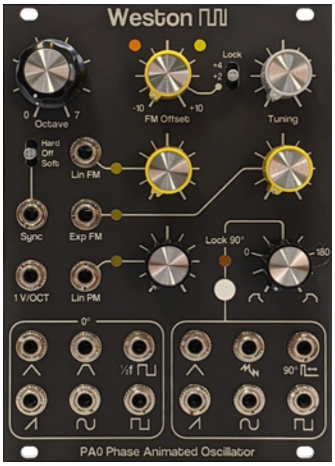

is fixed at zero degrees phase, and the right set can be adjusted between zero to 180 degrees with a push-button lock option at 90 degrees, which can be helpful for jumping to a phase.

Both sets of waveform outputs can work independently, but if both outputs are patched on each side it can create a waveform with more texture or even a stereo effect, depending on the way it is patched. If the output of the waveforms on the right is set to zero degrees, the waveforms will be the same and overlap. But at 180 degrees, they will be opposite and create a panning sensation. Though this is a single-oscillator module, it can sound like a dual oscillator by mixing two waveforms from the left and right banks and slowly phase-modulating the waveform on the right side.

The analogue sound is warm and clean, though the phase animation has the potential to create grit and harshness. It creates full and textured sound waves and therefore lends itself very well to panning and spatialisation. We hear two waveforms in phase when at zero degrees, and then the phase animation becomes audible as they are adjusted, slightly or drastically, to become out of phase. Then, there is the great CV experimentation point through the modulation inputs to be explored.

It's when you introduce CV to the Linear FM, Exponential FM and Linear PM inputs that things start to get really interesting, creating all kinds of complex results. For example, if you send three different speeds of CV pulses to the three different FM inputs, the PAO can even create interesting rhythmic sounds. The second set of phase outputs can be modulated through-zero with the Linear PM and Linear FM. To elaborate on the function of through-zero, the phase of the waveform is reversed, the frequency is increased, and the pitch remains stable.

For the discerning ear, which as synth players I think it is fair to say we all have, it is helpful to try different combinations of outputs from the two oscillator banks and compare the sounds, 6 x 6 giving you 36 different pairings in total. For example, if you pair the triangle wave output on the left with the unusual MW output on the right, the knob that controls the phase modulation produces a filtered, harmonic sweep. It sounds pristine and absolutely wonderful. Though it does not go too deep, the manual is extremely helpful in explaining the functions of the PAO, as it contains figures of the oscilloscope to explain and help us understand the waveforms and their sums in various phases.

For those who want to get into creative waveshaping with modulation options and a full, textured sound within a single oscillator, the PAO more than justifies its 18HP of rack space.

Dr Chelsea Bruno

$	$325
W	www.westonaudio.com

XAOC Devices Batumi II & Poti II

Eurorack Modules

It's been 10 years since the original Batumi quadruple low-frequency oscillator appeared. It was a solid, four-channel LFO with sine, sawtooth and square outputs with the option of switching out the saw for a triangle or trapezoid shape. It had four different modes: free running, fixed phase, variable phase and time division, and had CV control over reset, sync and frequency, which could push the range all the way down to 53 minutes. The Batumi has been a reliable purveyor of Eurorack modulation.

Prompted somewhat by the unexpected release of a clone, XAOC decided to reassert themselves with a new version that pulls in the expansions of the past while pushing it into the realms of audio as a four-channel VCO.

The Mode is now reflected in the four colours of a single LED on the left of the module, and our choices are Free, Variable Phase, Divide and Multiply. On the right is a Wave selector button which, through a six-colour LED, sets the assignable output waveform. On the original you could only set this, very inconveniently, on the back of the module via a jumper. XAOC introduced

the optional Poti 1974 breakout expansion to bring the switch to the front, which certainly helped. This is now built directly into the front panel. You can choose from triangle, saw in both directions, trapezoid and stepped or smooth randomisation.

With those small adjustments, the Batumi II feels like it's been perfected. Any criticism that could be levelled at the original has been smoothed away and it remains a completely solid four-channel source of slightly more modulation. However, two new aspects elevate this no-nonsense LFO into a surprisingly creative device.

As a four-channel VCO it gets a completely new lease of life. Each channel can be fed 1V/oct sequences and act completely independently as a melodic oscillator with sine, square or one of the assigned waveforms. While the range of the oscillator on the slider only goes up to 100Hz, it can be pushed all the way up to 5kHz with an input voltage, giving it a useful range of octaves. Tuning can be a bit tricky as you've got very little room to play with on the slider at audio frequencies. I found putting the sliders all the way to the top gave a slightly sharp G2, but all four were consistent. Outside of Free mode, you could run them as an out-of-phase four-oscillator cluster with channel A dictating the melody, or, in Divide mode, you could find octave-down sub-oscillators.

The other remarkable aspect can be found in the new and completely different Poti II expander. The Poti II adds a button so you can select each channel individually and set its own assignable waveform. Once selected, you can attenuate the incoming CV, the sine output and the assignable output level for that channel. It also adds a modulation input for each channel »

» with a three-way switch that lets you set either wavefolding on the sine wave, wave selection on the assignable or pulse-width modulation on the square wave.

So, in playing with the Batumi II, I found myself doing all the usual modulations and nodding my head in response to the predictable, solid outcomes. I very much enjoyed using the four VCOs because it becomes a great tool for chords, intervals and harmonies. All great stuff. But then I started to work it back into itself and realised that it has a phenomenal amount of power as a self-driven generator of waveforms. With a couple of patch cables, I could have a tuneful sine wave oscillator on channel A, being played by stepped voltages from channel B, while being folded by the sine wave LFO from channel C and reset by a pulse from channel D.

The ability of the Batumi II to interweave with itself is very satisfying. The Poti II is an essential companion that gives it so much tonal potential and more nuanced modulation through attenuation and shaping. It can be all sorts of things like a pair of VCOs with self-modulating wavefolding or PWM, it can be an FM sound source, a fat pulse wave machine, or a source of sound, randomness and modulation all coming apart in different directions. The Batumi II is overwhelmingly useful in a small system but superbly multi-functional in any setup. *Robin Vincent*

$ Batumi II $369.99, Poti II $99.99.
W www.xaocdevices.com

Modular Profile: Gur Milstein

Gur Milstein is the founder of Tiptop Audio, a developer with a litany of celebrated designs to their name; not least their Eurorack versions of legendary 200 series Buchla modules and the recent, groundbreaking ART system.

On his entry into modular
I probably came out of my mum's womb with two SL1200MkII Technics turntables because I just don't remember any time in my life without them. The music touched my soul, and the gear touched my technical curiosity. In the late '80s and early '90s, I was involved in the beginning of the global techno movement. Then came the moment I wanted to create this type of music myself, but getting analogue synthesizers back then turned out to be very expensive and difficult. That is when I decided to try to build one myself. This made me fall in love with physics, which led me to study electronic engineering so I could really know what I was doing. When I finished my studies, I worked for an engineering company for two years and completed the practical side of electronic engineering. In 2008, I opened Tiptop Audio from a small flat in the Hollywood Hills in Los Angeles.

On his go-to modules (aside from his own!)
I really like the Cwejman [*RES-4*] resonator; I think it stands out from the crowd with a beautiful, lush sound. I like the Harvestman Piston Honda V1, which was very harsh sounding but in a good way. The Cylonix Cyclebox is another wonderful oscillator for some digital sound madness. For control, I use the Bananalogue VCS as my go-to Serge-style slope generator.

On the story of Tiptop Audio
Tiptop Audio was born when I realised that there were a lot more people like me who were craving analogue sound and real knobs. In the '90s, we were horrified by the destruction MP3s and computers inflicted on the music industry and the collapse of nearly all synthesizer manufacturers, studio gear makers and top analogue studios. I started Tiptop Audio in 2008. A few years later, Chris Clepper and Piero Fragola joined. As we grew and the demand for our products increased, I decided that Tiptop should stand for its name. To do that, I needed people who would look into every part of a product we make to make it great. So, I built a synthesizer factory in Thailand and started producing all our products under one roof.

On the ART system
As Volt-per-octave is so sensitive to electric noise, it really is a poor choice for pitch control of oscillators. All oscillators, digital or analogue, might miss the correct note when receiving CV notes. In analogue, the effect will be a note out of tune, and in digital, it will just be the wrong note playing. The ART signal solves all these problems by sending the oscillators note value data at high speed. This means the oscillators will always play the right note you wanted, and not by chance. Once that problem was solved, we could also start making polyphonic modules that house stacks of oscillators that are all tuned properly and controlled from a single Eurorack patch cable. This is a huge step forward. As a new format, ART is challenged by the ability of the modular community of makers to embrace change. While ego, marketing strategy and backward compatibility all play a role in that, I have no doubt that ART is the right thing to take our clumsy old oscillator designs to a new level, making them more reliable, easy to work with, and even cheaper to make. Every oscillator or sequencer released these days without ART is a missed opportunity for a better Eurorack. ART is completely open and royalty-free for anyone to use.

On the culture of modular
Modular is the highest level of synthesis possible today, and as such, it does not fit a mainstream audience and results in being a sub-underground culture. It really puts the elements of sound, music and physics at your fingertips, and the community is full of people coming from all these angles: sound design, music, DIY and electronics. Therefore, this culture is composed of geeks, touring musicians, film composers, sound designers, DJs and lots of others who find this the hobby of their life and put everything into it. It is a culture with an addictive aspect to it, as the pleasure coming from using these fantastic instruments is truly immense. *William Stokes*


PROJECT SANTA MONICA
Introducing the SM-5, Kali's finest speaker yet.
Sound On Sound's Phil Ward called it "...a great example of relatively cost-effective contemporary electro-acoustic engineering, and Kali Audio ought to be very proud."
See page 62 for full review.
Starting at $1699 in the US.
Available now.
More information at
kaliaudio.com/santa-monica
NOMINEE
SOS AWARDS 2025
KALI
AUDIO

Positive Grid Spark Live Portable PA System

Having mastered guitar amps, Positive Grid's Spark system takes aim at the rest of your backline!

PAUL WHITE

The Spark Live puts Positive Grid's guitar amp modelling technology into a 150W, Class-D-amplified four-channel gigging PA that features a full-spec guitar amplifier modelling channel plus a mixer that can blend in one further mic/instrument source. The mixer also has a stereo line input, as well as allowing music playback over Bluetooth. Bluetooth is also used to connect an Apple/Android app for in-depth control over the DSP, and for pairing with the optional Spark Control X footswitch unit. For portable applications, an optional 6Ah rechargeable lithium-ion battery will give up to eight hours of performance time at a moderate volume, with a recharge time of three hours. The Spark Live is designed to be loud enough to hold its own against drum kits in smaller venues, and in larger venues, the amplifier can feed a house PA.

Number Crunching

Given its 150W power rating (450W peak), the Spark Live is pretty compact at just 464 x 251 x 306mm, making it very portable, though it feels reassuringly substantial and weighs a hefty 12kg. Its cabinet is covered with vinyl, and has an attractive grille cloth as well as protective plastic corner pieces and carry handles.

The speaker setup comprises a pair of one-inch compression tweeters feeding a custom angled horn, with the lows and mids handled by a pair of reflex-loaded, 6.5-inch custom woofers. SPLs of 118dB (121dB peak) are specified, with a frequency range of 45Hz to 20kHz (no tolerances are given). The Spark Live can be used either horizontally or vertically, and there's a fold-out leg on the rear in case you want it angled. For stand-mounting, there's a standard pole socket for horizontal use.

In all, there are three mixer channels: one on the front panel for guitar, and two more on the rear panel. Two rotary controls on the front panel adjust the overall level of channel 1 (guitar) and music (Bluetooth streamed audio or USB DAW playback). Channel 2, which has a volume control and a preset selector, is fed from a combi XLR/jack socket and has access to its own menu of amp modelling and effects, while channel 3+4 has only a pair of straightforward jack inputs and a volume control. There's also a 3.5mm mini-jack output for headphones.

A pair of five-pin DIN sockets provides MIDI In and Out. MIDI can be used to bypass the effects, to change presets, to control the levels of the individual channels and to sync a second Spark Live amp. There are also two USB-C sockets, one for use as an audio interface or for firmware updates, the other for charging mobile devices. The Bluetooth pairing button is located just above the USB ports. The recording feed produces a 24-bit/48kHz stereo signal, and is compatible with macOS and Windows.

Two line output jacks facilitate running the amp's signal to a PA or further amps or speakers. A master control section offers a three-band EQ and a volume knob.

The Engine Room

At the heart of the system is Positive Grid's proprietary "Sonic IQ audio engine", featuring signal processing that has been designed to take account of the acoustic engineering aspects of the amplifier so as to optimise the final sound. The amplifier can even adapt dynamically using its G-Sensor, automatically tuning the system depending on whether it is being used horizontally or vertically. In vertical mode the amp works in mono, while in horizontal mode it switches to wide stereo. Other features include auto loudness and clarity adjustments, compression, bass enhancement and vocal enhancement.

How can all this be achieved without making the amplifier too complicated? The answer is to put a practical number of physical controls on the amp itself, whilst leaving detailed editing and preset creation to the app. The app allows amp and effects modules to be swapped out and adjusted before being saved back into the amplifier, from where they can be further controlled using the familiar amp-style controls on the front panel. Channel 2 presets are accessed in the same way, though all adjustments have to be done from the app.

A channel 1 preset can have one amp model and up to three pre- and three post-amp effects, drawn from the full complement of Spark processing modules (33 amps and 43 effects). The front panel offers physical controls for Drive, three-band EQ and volume, while up to eight presets can be saved/recalled using a push-turn encoder accompanied by four LEDs that show either red or green. A long press of

the encoder stores the current settings as a preset, while a double-press puts the amp into tuner mode and mutes the guitar output. The LEDs show whether the note is sharp or flat relative to the closest note, but they don't tell you what note it is.

Channel 2 offers three amp models tailored for vocals, bass or acoustic guitar, this time with 50 effects covering instrument and vocal treatments. Here the user can set up one effect before the amplifier model and two after it, and there's a preset encoder dial exactly like the one on channel 1. A typical use for channel 2 might be for a vocal mic or acoustic guitar, with the stereo channel available for sources such as keyboards or the outputs from a small mixer.

Appy Now?

We've covered the amp and effects emulations in Positive Grid's Spark range before, but suffice to say there's a useful range of amp types and effects from which to choose, allowing for guitar tones that run from country clean to 'veins in your teeth' heavy rock. There's also an AI Jam Along feature: tap in a tempo on your phone, play a few lines, and the app will create a suitable backing in the same key. A number of ready-made backing tracks are also included, with chords displayed on your phone, and there's a library of presets created by other users available from Positive Grid's Tone Cloud.

Navigating the Spark Live app is straightforward: a mixer icon takes you to a very simple faders page that lets you balance the levels of the various inputs, while tapping the Channel button toggles between channels 1 and 2. The app is clearly set out, with graphic depictions of the various amps and effects that offer strong visual clues as to which 'real life' amps and pedals they are inspired by. You can turn effects on or off, switch them out, swap amplifiers, adjust their controls and so on. Communication works both ways, so adjusting a dial on the app changes the value on the amp, and *vice versa*.

The Spark Live can accommodate an electric guitar, a mic or other instrument, a stereo line-level source, plus stereo Bluetooth or DAW playback simultaneously.

In Use

Though the amplifier comes with its preset slots filled with a variety of amp tones, I didn't find any that were quite right for my own rock covers band, so I invoked the app and found that setting up my own amp/effect chains was very easy. Additional FX can be purchased, but for most applications the user is well served with a practical selection of each type of effect.

If there's a function missing, it is a master preset level adjustment on the amp itself, to provide an easy way to balance the levels of the presets you wish to use live. As it is, you need to use the app and adjust the master volumes on the amp models, which is not quite so straightforward, especially if you want to make tweaks during a soundcheck at a gig.

Overall, the app is really well thought out, and the AI features are useful if you want to explore variations of tones that fit a particular genre, or jam along with an automatically generated backing. The more generic jam tracks arranged by genre are also useful, and I like the way they display the chords on a moving grid as they play. Checking out what other users have created via Tone Cloud is also a good way of homing in on tones you like, as you can see exactly what settings they have used.

Of course, if the amp didn't actually sound good, then all of the above technology would be a waste of time. Fortunately, Positive Grid have been working on amp modelling for years and they've got it down to a fine art. As long as you give yourself time to explore the different amp models and make the necessary adjustments, you can get pretty much any guitar sound you like. The compression and reverb add a professional polish to clean sounds, while combining compression with light overdrive or a gently pushed amp will deliver those edge-of-breakup sounds that work well for blues or country-rock. There's a decent amount of touch responsiveness too, so you can control proceedings via the guitar volume control and/or playing dynamics.

When it comes to high-gain sounds, there's everything you could wish for, from classic rock to edgy metal. Add a bit of delay and the corner of a small pub becomes a stadium. You don't quite get that sense of moving air that you get from an open-backed combo — to my ears, what you hear is more like a studio recorded guitar sound, but at gigging volumes. For many players, that will be just what they want, especially if a lot of different sounds are needed during a gig. The noise gate works well for cleaning up high-gain sounds but the default setting sometimes cuts off sustained notes a little short, so I'd advise adjusting the threshold as low as you can get it.

In a home or studio setting the amp feels right at home and it can deliver recording-ready tones very easily, either over its two-in/two-out USB audio interface, by taking a line feed from the back of the amp or by sticking a mic in front of it. At pub gig levels you should be fine if just using the amp for one guitar, but if mixing multiple sources and playing with a drummer, hooking up a powered extension speaker or two would be a good plan. Lastly, the optional battery would be a great option for some seriously good-sounding busking.

Positive Grid Spark Live

$549

PROS

- **A huge range of very viable guitar tones.**
- **Separate presets for channel 2.**
- **Simple mixer built in.**
- **Spark Live app for tone editing and more.**
- **Optional battery operation.**

CONS

- **A way to balance the different preset levels on the amp itself would have been nice.**

SUMMARY

A flexible, great-sounding portable PA that offers full-fat guitar DSP, as well as options for bass, vocals, acoustic guitar and more.

$ $549, battery $79.
W www.positivegrid.com

Moog Spectravox

Semi-modular Spectral Processor

This new addition to the Moog semi-modular 60HP range is a little hard to pin down, so what exactly is the Spectravox?

ROBIN BIGWOOD

Along with the Labyrinth (reviewed in the August 2024 issue), the Spectravox is Moog's newest semi-modular standalone and Eurorack-compatible device, joining the well-known Mother-32, DFAM and Subharmonicon, with similar compact form factor, styling, and high-quality construction.

It would be easy to assume the Spectravox is just an analogue vocoder — not that there are very many of those out there — but in fact it's a multi-function unit that should have much wider appeal. Under the hood it has a true analogue vocoder's requisite pair of filter banks and accompanying facilities, and I note with some kind of miserly delight that it undercuts the currently available, hefty Moog 16-channel vocoder reissue by about £4500.

However, it's also possible to use the Spectravox only as a filter bank, routing into it mono audio from elsewhere and enjoying tone-shaping possibilities that lie somewhere between those of synth filters and EQ. In this role there's a direct lineage to the Moog 907 Fixed Filter Bank, used in various full-size modulars. In fact, a Spectravox can do things a 907 can't. The first is to introduce resonance in each filter band, and the second to move or 'shift' the filter frequencies, jointly and relatively, making the Spectravox more of an 'unfixed' bank, if you like. Spectral Shift (as it's called) can be done manually, or an onboard, hardwired LFO can do it for you, with scope for liquid, phaser-like sweeps or resonant sirens.

Finally, for distinctly good measure, the Spectravox is a simple synth voice as well, which can act as the harmonically rich carrier component for the vocoder, or can be utilised independently. It has a single oscillator with sawtooth and variable-width pulse waves alongside a white-noise source. The provided envelope generator is a basic thing, with an instant attack, variable decay, and no sustain phase, but it can help to create percussive sounds. Pulse-width modulation, vibrato and more can also be achieved using the shift LFO, via the patchbay.

Say 'Uhh'?

The way these elements interact when you approach a Spectravox for the first time may not be glaringly obvious. The only non-modular connections (other than for the 12V 2.0A mains adaptor) are a quarter-inch rear-panel mono line-level output that also doubles as a headphone out, and a front-panel XLR/jack combi 'Program Input'. In contrast with many digital vocoder alternatives on the market there's no MIDI, so the onboard synth cannot be played that way. The included quick start guide isn't a great help either, being extremely light on practical, useful information. The downloadable full PDF manual is much more what we've come to expect from Moog, includes some fascinating historical context, and has a vital 'Exploring Spectravox' section with some setup walkthroughs. A clutch of patch overlay cards are useful too, but you'll still get the most out of those after you've read the PDF. In the end, and after some experimentation, it all makes sense, and there's a lot to enjoy.

Filter Fabric

Much as I'd love to (yet again) pull on my full flower-power get-up and dive into the

Spectravox in its role as a vocoder, it makes more sense to consider it as a filter bank first. So the lab coat it is...

Getting into this mode, and others, is a matter of appropriately juggling the three switches at the lower right. Then, with no external patching, you'll hear the internal oscillator routed through the filters. External signals go into the 3.5mm Carrier patch point on the front panel and break the internal synth connection. Actually, you can also use the XLR/quarter-inch combi socket: on its own it won't get routed to the right place, but patching Program output to Carrier input solves that. The mic input does not supply phantom power, incidentally, but its gain knob was able to tease out healthy signals from various dynamic mics I tested it with. An accompanying LED indicates signal level with varying brightness, but is a blunt tool as meters go.

In no time it emerges that the Spectravox's filter bank — the one used to sculpt and shape external signals — is flexible, can pull off some interesting and remarkable tone transformations, and has plenty of character. It's distinctly unlike the vast majority of single-band synth filters we're all used to and is full of practical and creative potential.

At the same time, it is not a Moog 907 emulation. Where the rare handmade module has filters tuned to two harmonically-related stacks (250, 500, 1000, 2000 Hz, alongside 350, 700, 1400 and 2800 Hz, plus low- and high-pass filters), the Spectravox offers the funkily inharmonic line-up of 230, 320, 560, 830, 1200, 1700, 2300, 3200, 4500 and 5400 Hz when the Shift knob is at its default central position. You can then swing frequency centres sufficiently that the lowest band will go almost subsonic at one extreme, and the highest band ultrasonic at the other. A 907 can only dream of that.

Where they are similar though is in the way that the pots provided for each band are all attenuators. You can only cut the level for a band, never boost it, but attenuation can be complete in that if you turn down all bands at once you'll get silence. On the Spectravox attenuator pots are of the skinny love-'em-or-hate-'em kind, but I understand an official knurled pot kit is available to purchase to help those of the latter persuasion.

Filters 2-9 are band-pass types, with overlapping 6dB/octave slopes (as far as I can make out, squinting at my spectrum analyser). The slope flattens and broadens a touch in the upper 30 degrees or so of pot travel. Moog call filter

Moog Spectravox

$599

PROS

- Unusual tone-shaping potential for external signals and internal tones.
- Patching extends creative potential exponentially, blurring traditional oscillator/filter boundaries.
- The simple synth oscillator and white noise source punch above their weight.

CONS

- The vocoder offers poor vocal intelligibility, despite undoubted character.
- External gear required for some basic CV tasks such as level matching.

SUMMARY

A historically resonant analogue filter bank, vocoder and synth rolled into one. Satisfyingly adaptable, it can also be delightfully weird.

1 a low-pass, and 10 a high-pass, and on the control panel that's reinforced by little response curve graphics. In practice they behave much like additional band-passes in their relative positions, albeit with filter 1 exhibiting more of a shelf-type response over low-end frequency content. You could even argue that they're named the wrong way round, according to the way they often behave. Certainly filter 10 removes high frequencies, pure and simple, which is not most people's definition of a high-pass filter... It's a matter of semantics more than anything, so work with the band controls intuitively and all will be well.

Turning up the Resonance knob (yep, eat your heart out again, 907!) steepens the filter slopes and tends to make the frequency centres ring. Not to the extent they'll self-oscillate, but enough that they can and will contribute harmonic content, almost like a resonator.

How you employ band attenuation, in conjunction with the Shift knob to 'tune' the bands, is up to you. Remarkable vocal tract-like and nasal sounds are available when you turn down all bands except perhaps say 2, 5 and 10, increase resonance and set Shift quite high. Try it with your Moog Subharmonicon for a Trautonium-inspired nerd-out. Or you could sculpt a ramp going from a fully open band 1 to a closed 10, to create something more like a conventional low-pass response, with Shift opening and closing the array like it was a cutoff control. Raise band 10 and high frequencies can start to bleed through, somewhat reminiscent of the smooth sparkle of an Oberheim SEM or other continuously variable filter »

The Spectravox's back panel, as with all the 60HP range, is a sparse affair, with just a quarter-inch audio output and a 12V external power input.

between its LPF and notch states. Yet another approach could be to turn down all odd-numbered bands (for example), raise resonance, and start sweeping your Shift. The resulting moving peaks and phase interactions are reminiscent of some phaser or flanger sounds.

Which brings us nicely to the onboard Shift LFO. Hardwired to the Shift function, it has a fixed triangle wave and rates from approx 0.05Hz (one cycle every 20 seconds) to 500Hz. Even with the Amt (Amount) knob at maximum the LFO can't quite visit the extremes of Shift that can be explored manually, but the range is still more than sufficient for classic LFO-driven phaser effects amongst others. It's also great for adding subtle movement and minimising any sense of sterility associated with fixed filtering.

Rounding out this first section, I'll take this early opportunity to point out that the flip side to Spectravox's uncommon filtering chops is a certain inherent, strong, baked-in colour that's impossible to escape. The carrier filter bank skews every signal that goes through it, no matter what. Routing in a test tone or a mix, for example, and then twisting the Shift knob — even when all bands are fully open and resonance is down — will still result in obvious, radical filtration zooming across the audible spectrum. This is rather different to most low-pass filters of the modular synth world, which when they're fully open (and assuming they're not distorting) should spit out pretty much what you put in. The Spectravox's ever-present harmonic zing is certainly not a failing, it's definitely a feature, and is very much part of its character. Often it imbues everything with a scooped quality, warm and buzzy.

Voco Loco

Turning to Vocoder mode now (please feel free to turn on your lava lamp...), this way of working with the Spectravox brings the additional Program (or analysis) filter bank into play. It receives signal from the front-panel combi jack or patchbay Program input, and in keeping with its analysis role its filters are fixed in frequency and unaffected by the Shift knob.

Used conventionally, with vocal program and synth carrier signals, things can certainly sound rich and smooth, recalling those earliest Wendy Carlos vocoder tracks. However, it has to be said, intelligibility of speech or singing is very poor indeed. Now, this arguably makes for an accurate emulation of that first, iconic Moog-Carlos 10-band vocoder, which is reputed to have required an almost absurdly over-enunciated vocal style. It's not that there's no clarity at all, but there's frequently nowhere near enough for words to be made out, regardless of how extreme your delivery. Consequently it may be safer to think of this vocoder as more of a 'singing synth' than something that can get near an 'Autobahn', 'Mr Blue Sky' or 'O Superman'.

Still, in an effort to improve speech intelligibility, Moog have employed the historically-rooted scheme whereby white noise replaces the Carrier signal fed to filter bands 9 and 10 when they're triggered by fricatives, sibilance and other 'unvoiced' elements in the Program signal. The Hiss position of the Hiss/Buzz switch enables this, and it works well in principle, albeit with a rather hard-edged on/off quality. However, the use of band attenuators can tame these noises and integrate them much better into the overall effect.

Of course, the Carrier filter attenuators and Shift knob are in play here for vocoding just as they are for filtering, and that allows

Levelling Up

Audio inputs and outputs are vital for an external filter bank and vocoder, and the Spectravox has everything it needs on the face of it. But I was surprised to find that Moog have designed Program and Carrier patchbay inputs for line level (specifically 1.7V) and not Eurorack (typically 5V) signals. Of course, you can attenuate your favourite Eurorack oscillator's output before it hits them, but that'll require an external attenuator or two (there are none built in here) and it means gain-staging can become a bit hit and miss. Certainly with complex vocoding using external carrier and program signals I frequently ran into distortion, especially with high levels of resonance. But since backing off levels sometimes made no difference I suspect it's just part and parcel of the Spectravox's internal summing, and I'm minded of the various Moog synth mixer stages that are revered because they are not clean or linear.

The other gotcha is that sounds using heavy attenuation of many program filter bands simply end up very quiet. This is easily dealt with if you're able to raise gain to compensate on a mixer or audio interface, and you might even choose to make the 12 o'clock position of the Spectravox's volume control represent a normal working level, allowing for compensatory boosts above it when required. It's again potentially more of a bother in Eurorack, where you'll need to commit something with a gain or boost stage to compensate. Signal-to-noise ratio is eroded in these circumstances, inevitably, though not unworkably so. Oscillator and noise bleeds can be audible sometimes, though on the whole signal-to-noise performance is respectable, and plenty good enough for real-world use.

for a kind of easy, in-the-box tone shaping and articulation of the vocoded signal that I've certainly never seen before in hardware, and is almost as rare in software.

Additional niceties include a Hold switch (with an accompanying patchbay trigger) that will freeze control voltages arriving at the Carrier filters: it's a sort of spectral 'shutter'. Engage it while singing a vowel sound, say, and there's no requirement to keep on singing. It can also make for some fascinating effects when applied rhythmically to an incoming signal. Then there's an option to have the onboard decay envelope trigger from a certain threshold of Program signal level. Finally, huge flexibility is offered via the patch points on every band for both the envelope follower outputs and the VCA inputs. Mind-boggling possibilities open up here, not always easy to predict. One is simply to cross-patch bands to scramble the normal spectral response: bass sounds could open the highest filters. But you can also attempt tricks like patching band 1 and 10 envelope followers to external Eurorack envelope or sample triggers, letting you beatbox into the Spectravox and 'play' kick and snare elsewhere. Or with a suitable sequencer, with multiple CV gate outputs, you could rhythmically open and close individual frequency bands. This aspect of the CV connectivity is lavish, and in fact was completely missing from the prototype, engineer-built Spectravoxes that emerged from the 2019 Moogfest. So thank you, Moog, for pushing that out to us general punters.

Synth Forever

The inclusion of a simple oscillator and noise source on the Spectravox is a nice bonus, and adds surprising flexibility. In its most basic role it'll support static-pitch robot-voice vocoding, acting as an internal carrier. But it's also a genuinely playable, useful modular sound source in its own right, which in my tests had accurate 1V/oct pitch tracking over a wide range and decent tuning stability too. The VCO Frequency knob sets static tuning from a clicky 10Hz to dead-on 4kHz, and with such a span it makes accurate tuning for musical use with external CV controllers quite tricky to dial in, though it's possible with care.

The switchable sawtooth and variable pulse oscillator sound is classic, lively Moog, and supported by both a dedicated pulse-width knob and patchbay PWM input can offer movement and girth. Just as fruitful though is the continuously variable VCO/Noise mix knob, rather like the one offered on the Mother-32.

Now, for use in a modular rig you could source the oscillator and noise signals via their dedicated patchbay outputs, where they emerge pristine and energetic. But »

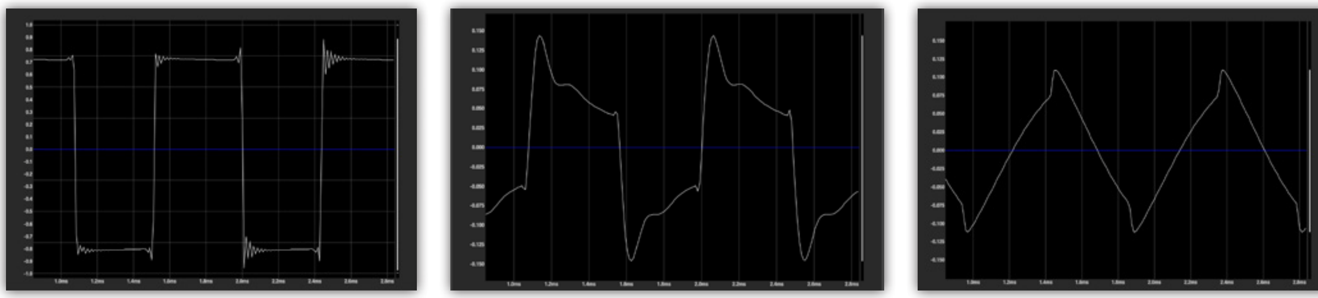

Because of the inherent coloration associated with the Carrier filter bank, a huge timbral range can be teased out of a simple single oscillator. To illustrate the point, the three waveforms here are all Spectravox pulse waves. The first is sourced direct from the VCO output, and is the raw, 'accurate' pulse straight from the oscillator. The second is from the main output with Shift in its central position. The third — basically a triangle wave — has Shift turned 90 degrees left. The crucial thing here is that all filter bands were wide open at all times: a further range of wave shapes emerge when you start to use them.

Like the other models in the range, the Spectravox's front panel is 60HP wide and can be removed from the case and fitted into a Eurorack system.

part of the charm of the Spectravox is that baked-in filter bank coloration I mentioned earlier, and using the post-filter VCA output (or indeed the main rear-panel output) makes for a very different sound. It's the same sawtooth, pulse and noise, but even with all filter bands open they'll now be far from pristine. The scooped quality is evident once more, wave shapes are nowhere near their starting points, and the noise profile becomes unusual but attractive, with somehow quite hollow, throaty-sounding low-mid frequencies and smoothed-out highs. With all those band attenuators in play I can't think of another modular-world oscillator/filter combo that is quite like this one. It'll produce uncannily human vocal timbres on the one hand, and beautiful synthetic snares and hi-hats on the other, for which the little decay envelope and front-panel (or patchbay voltage) trigger come into their own. Alternatively, max out noise level and resonance, enable one or two filter bands only, and patch a 1V/octave musical controller to the Spectravox's Shift input, and you can tease out all sorts of pitched noise timbres that track perfectly and can be played 'in tune', in their way.

Vox Aeterna

As with some other Moog semi-modulars, what you might first assume the Spectravox is all about is not quite how things work out in reality. I've described its application as a filter, vocoder and synth voice, and there is quirkiness associated with all three. As a filter it excels at tonal sculpting bordering on resonator territory, but isn't perhaps one most people will choose for bread and

Vocoders Revisited

Vocoders have always piqued interest in the general listening public, but their roots are in the purely functional, as devices that made the best of severely limited bandwidth in early long-distance voice communications, as far back as the 1920s. One of the very first musical uses came through a collaboration between Wendy Carlos and Bob Moog in 1968: two modified 907 Fixed Filter Bank modules were used alongside a clutch of envelope followers and VCAs. The rest, by way of *Switched-on Bach* and the soundtrack for *A Clockwork Orange*, kicks off a particularly colourful part of synth history that saw many competing vocoder products released through the 1970s.

Vocoders work by pairing two complementary banks of (mostly) band-pass filters. In the first, the program or analysis bank, there's an envelope follower on each band, which together form an approximation of the harmonic make-up of an input signal. The voltages generated from the followers open VCAs for a set of corresponding filters in the second bank, called the carrier or synthesis bank. A harmonically rich signal like a pulse-wave tone is usually used as the audio input for that, and for musical use it can be melodic or chordal. Voila! The energy signature of the first signal is imposed on the second, and if you're using your voice as a program input it won't really matter at what pitch you sing or speak: the pitch that's heard is that of the carrier.

Crudely, the more bands a vocoder has the more intelligibly speech and other articulate signals will be rendered when used as a program input: some digital vocoders offers hundreds of bands. But it's interesting to note that another vocoder classic, the Roland VP-330, also had only 10 bands and manages to be highly intelligible. Other factors such as band tuning, filter slopes and envelope follower response characteristics come into play too. Of course, there's a question as to whether speech intelligibility should always be the end goal, or if other characteristics and strengths could or should be prioritised: the Spectravox is one answer amongst many, and perhaps there are no wrong answers...

butter synth patches. As vocoders go it has a strong flavour and certain quite magical qualities, but vocal intelligibility is not where it excels. As a synth it's fundamentally basic but can sound distinctive and unique.

What emerged for me is that the whole of this unusual box of tricks is greater — and weirder — than the sum of these parts. It's when you look beyond obvious mainstream uses and get creative that the Spectravox seems to bloom into something really special. An example is the potential as a noise source: it's a wonderful generator of all kinds of interesting beds, sweeps, special effects and indeed synthetic drums. Then there are the sounds that emerge through feedback, such as when you patch the VCA out to the Carrier in: unpredictable outcomes emerge at the turn of every band pot. Not to mention the modulation possibilities offered by frequency-constrained envelope follower outputs, or the broadband Program level follower, which exists only in the patchbay. To think out of the box with the Spectravox is to discover its real essence I think. The more experimental I got with it the more it gave back.

Could it be better? Yes. It's unfortunate the vocoder doesn't have the scope to be more intelligible than it is, because for mainstream use (and who isn't up for that once in a while?) you'll need something else besides. We could wish for more CV ins too: LFO Amount is the one that is conspicuously missing. There are also no general attenuators or mults, which will have you looking elsewhere for some jobs.

However, the Spectravox is both unique (as far as I can see) and entertainingly versatile. It carves niches that could be irresistible to devotees of synth history, filter fetishists, effect aficionados and exploratory synthesists alike. It'll make for an iconic, absorbing and great-sounding enhancement in many a semi- or fully-modular setup.

ALTERNATIVES

Genuine analogue vocoders are rare, and apart from the expensive, 16-channel **Moog** reissue the currently available alternatives are **Behringer**'s **VC-340** (a largely faithful copy of the classic **Roland VP-330**) and in Eurorack format **AnalogFX**'s **VXC-2220**. Both offer vastly better vocal intelligibility than the Spectravox, but no equivalent of its Shift control or band attenuation (other than in the module, if you were to equip it with banks of external attenuators). In the digital realm the **Electro-Harmonix V256** pedal is a dependable choice, adjustable between eight and 256 filter bands, and sporting a MIDI-driven polyphonic synth. The **Boss VO-1** is good too, though you'll need both external program and carrier signals for that. Beyond that there are vocoders built into synths (**Arturia Microfreak**, **Korg MicroKORG**) and also plug-ins, often fabulously flexible. **Arturia**'s Moog-inspired **Vocoder-V** is a particularly potent one.

As for filter banks, in the modular world **Doepfer**'s **A-128** is a 15-band fixed bank, pure and simple. Somewhat more complex is **Tiptop Audio**'s Eurorack recreation of the **Buchla 296 Spectral Processor**. Also check out the **ADDAC 601**, with eight bands and lots of voltage control, and of course **Behringer**'s ultra-affordable copy of the 14-band Moog 914. In software, a standout option is the **914 MkII** plug-in by **Audio Damage**: it's a 14 band Moog 914 emulation and has some equivalents of the Spectravox's desirable resonance and Shift features.

$ $599
W www.moogmusic.com

Heritage Audio i73 Pro Edge

USB Audio & MIDI Interface

Does Heritage Audio's new range of transformer-based audio interfaces mean you no longer need separate mic preamps?

SAM INGLIS

From bedroom setups to Brad Pitt's château, practically every studio nowadays is based around a computer. That means practically every studio needs a computer audio interface. And that, in turn, means the market for computer audio interfaces is fiercely competitive. So it's strange that, as I've observed more than once in these pages, manufacturers have allowed an obvious gap in that market to remain unfilled for many years. The large majority of audio interfaces nowadays have mic preamps built in — but the preamps that get built into audio interfaces are of a type that lots of studio engineers don't use.

Many engineers want characterful mic preamps that can add pleasing thickness, coloration and saturation to their recordings. What we get in our audio interfaces are transformerless, chip-based preamps that offer impressive technical specs but add nothing to the sound. That's slowly beginning to change, courtesy of products like the Neve 88M and 1073 SPX-D, Steinberg UR-RT series and the Icon/Harrison 32Ci we reviewed in September. Other manufacturers such as Universal Audio have chosen a different route to the same destination, developing input circuits that can emulate the sound of older console preamps. But it's fair to say that if you want a good-quality audio interface with classic transformer-based preamps, your options are still very limited.

Heritage Audio's i73 Pro series thus provoked much interest when it was announced in January. It's billed as "the first ever USB-C audio interface with 73-style Class-A preamps built in", and as if to push home the point, it has square holes in the top panel through which the input transformers proudly protrude. And, although the preamps are definitely the headline feature, the i73 Pro offers plenty of interesting design choices in other departments too.

Three Times 73

The i73 Pro range contains three models, of which the i73 Pro One is the simplest. On the input side, it has a single channel of transformer-balanced preamplification that can accommodate mic, line or instrument signals, plus a mono line-level input. Output-wise, it sports

The i73 Pro One is the baby of the range, with just one transformer-based preamp input.

a pair of line-level monitor outputs on balanced jacks, and a single front-panel headphone output.

The next step up is the i73 Pro 2, which replaces the line-level input with a second transformer-based preamp; unlike the first, though, this does not have an instrument input option. In other respects, it's largely identical to the Pro One, with the same output arrangements.

Top of the range is the i73 Pro Edge, as supplied for review. This features two transformer-balanced input channels that both accept mic, line and instrument signals, plus a second pair of line-level inputs on balanced rear-panel jacks that can be switched between -10dBV and +4dBu input sensitivities. Output-wise, the Edge adds a second headphone socket and a second pair of balanced line-level jacks. And, unlike the other two models, it also has digital audio I/O. A pair of optical sockets accept and transmit up to eight channels of ADAT Lightpipe audio at base sample rates, with no option to switch to stereo S/PDIF.

The use of Class-A transformer-based preamps means that bus powering is out of the question, so all three models employ external 12V DC supplies. And in addition to their audio I/O, all three interfaces also have MIDI in and out. However, both MIDI connections are presented on a single six-pin mini-DIN socket, so you'll need an adaptor cable to use them. This is a cost option and is not included as standard. Connection to the host computer is indeed on a single USB-C socket, so Heritage Audio's billing remains technically true: the Neve 1073 SPX-D was available first, but has a Type-B socket.

Heritage Audio
i73 Pro Edge
$1499

PROS
- At last, someone's made a reasonably affordable audio interface with colourful, transformer-balanced preamps.
- Output trim control makes it easy to exploit the character of the preamps.
- Has built-in signal processing, and comes with a free suite of rather good native and DSP plug-ins.
- Well thought-out mixer utility that can host VST plug-ins on its auxiliary channels.
- Powerful headphone amps.

CONS
- Relatively narrow input gain range.
- Limited DSP resources restrict the usefulness of the plug-ins.
- MIDI cable is a cost option.
- Insert points would be nice.

SUMMARY
The i73 is an ambitious and impressive audio interface that combines DSP plug-in hosting with old-school transformer-based mic preamps.

With their chunky wedge-shaped desktop cases, the i73 models have a form factor similar to Universal Audio's Apollo Twin interfaces, but they're considerably larger. The Pro Edge version occupies about 20x23cm of desk space and stands a full 9cm high. The colour scheme and Marconi-style controls make obvious visual reference to the Neve consoles of yore, although the slightly wobbly knobs and wood-effect plastic end cheeks don't quite communicate the same sense of permanence and gravitas.

Finally, there's one more important factor that the i73 series has in common with UA's Apollos. These are not only audio interfaces: they're also DSP platforms that come with their own suite of plug-ins. Where it matters, these exist not only in native formats but also as processors that can be loaded directly into the i73's DSP chips, allowing their output to be monitored and recorded at very low latency.

73 Or 72?

All three i73 Pro models share the same 'master section' features, located on the right side of the top panel. These include mono, mute and dim buttons, along with a dual-concentric encoder. This has a push-button action that toggles whether »

The flagship i73 Pro Edge features a second pair of line-level inputs and outputs as well as ADAT I/O. A special adaptor cable is needed to use the MIDI ports.

the inner ring controls the volume of the main monitor outs or that of outs 3+4, and which of the two headphone amps is governed by the outer ring. LEDs indicate which of the two are presently under encoder control, but there's no visual indication of the actual volume level in either case. A five-segment LED ladder at the top of the 'master section' can display either levels at inputs 1+2 or at the monitor outputs. All the controls in the master section are virtual, in the sense that they don't adjust analogue potentiometers or switches but digital settings that can also be controlled in software.

Settings for the inputs, by contrast, are all analogue. They reside on the left-hand side of the front panel and begin with -20dB pad, polarity and +48V phantom power buttons, followed by switched and continuous gain controls. These look at first glance like the ones you'd find on a 'full fat' vintage preamp bearing the number 1073 (which, along with the Neve name, is a registered trademark, hence the abbreviations and euphemisms flying around), but there are some nuances here that are worth teasing out.

On a genuine Neve 1073, the main gain control has 22 switched positions that span almost a full 360-degree arc, and include two settings marked 'off'. This is because it has multiple gain stages. Settings in the middle of the gain range engage just one of the two mic preamp boards, and the second stage is brought in only once you go past the second 'off' setting. By contrast, the i73's preamps have just one gain stage, and the switched gain control has only six positions. As on the 1073, these are stepped at approximately 5dB intervals, so the total gain range available from this control is just 25dB. The pad extends this range by a further 20dB at the lower end, so for mics, the minimum gain available is +25dB, with the pad engaged, and the maximum is +70dB. The pad is not available for line inputs, so the total gain range in this case spans -20 to +5 dB. There's also no equivalent of the 1073's low-impedance option in the i73.

To put all this in language that will make sense to Neve geeks, the input stage in the i73 is actually not so much a 1073 copy, but much closer to the 1272 transformer-balanced line amp that was widely used in early '70s Neve consoles. The 1272 is now often repurposed as a mic preamp, and is seen as being sonically very similar to the 1073 within its more limited gain range.

The key to wringing colour and harmonic saturation from any mic preamp is the availability of an output trim control, because this allows you to crank up the main gain setting and then back things off on the way out to avoid overloading downstream devices. On the i73, the downstream device is always the built-in A-D converter, as there are no insert points into which EQs, compressors or attenuators could be placed, and the trim control runs from minus infinity to zero (unity gain). So, no matter how hard you're hitting the preamp, you can always avoid overloading the A-D converter by turning the output control down far enough — but this won't ameliorate any clipping or other non-linearity that takes place within the analogue stages.

Slotting In

Heritage Audio's marketing for the i73 Pro models trumpets the Full Analog Experience that they're said to offer. Slightly counter-intuitively, quite a lot of this 'analogue' experience is actually digital, and is made possible by the built-in DSP. In order to access this, and indeed to use the i73 at all, you'll need to install two separate pieces of software. An app called Heritage Sync runs in the background and handles communication with the Heritage website, authorisation, firmware updates and so on, whilst actual operation is carried out using the i73 Mixer utility. Heritage Sync was a little bit touchy on my Mac; it wouldn't let me use the i73 without updating the firmware, but the update failed several times before going through OK. Once you've got everything working and registered your unit, you can download the free plug-ins that complete the Full Analog Experience and authorise them to your iLok account.

The Mixer software is more than a little reminiscent of UA's Console, albeit with a slightly more obviously Neve-esque GUI. The i73 Edge Pro has a total of 12 physical input channels, comprising the two mic/line/instrument inputs, the additional line inputs, and the eight ADAT ins, while its 16 outputs are made up of the main and secondary line outs, the two stereo headphone outs, and the eight ADAT outs. However, your DAW 'sees' an additional four inputs and two outputs. Two of each are accounted for

There's a new EQ on the 500 block.

We can now ALL have the the amazing sound and flexibility of the UTA console EQ in a 500 series package.

The UTA console EQ and subsequent MPEQ-1 channel strip have been recognized for their exceptional vintage musicality and extraordinary flexibility.

Imagine having a 500 series EQ that can replicate the shape or curve of any other hardware EQ made.

The 500 series format makes this extraordinary EQ design more accesible for all.

FEATURES:

3 "Super Parametric" bands

Range switching mid band

Shape Control for blending between peak and shelf shapes

Q control range from .3 to 10

Notch Filter mode for extreme cuts

Unique "Vari-Phase" mode

input pad for added headroom

High quality, long life switches and capacitors

Introducing the UnderToneAudio

UTEQ500

NOMINEE
SOS AWARDS 2025

Howdy Neighbors!

by a stereo loopback channel, whilst the final two input sources are related to the i73's DSP features.

These DSP facilities are available only on the two mic/line/instrument inputs, and are presented as four insert slots. By clicking on these, you can load one of the four DSP plug-ins currently available, which are described in the boxout. Any processing you use in these insert slots is always in the monitor path through the Mixer utility, but the record path is split, so you can choose to capture either the wet signal (with processing), the dry signal, or both. You're only allowed to load one instance of any given DSP plug-in per channel, and in fact four insert slots is currently overkill, as there is no combination of four DSP plug-ins that can be loaded without maxing out the DSP resources.

All the other mixer channel features are common to every channel, and include a fader, pan pot, mute and solo buttons, and two auxiliary sends with separate pre/post-fader switches. Channels can be named, and adjacent channels linked for stereo operation.

Great Auxes

On some audio interface mixers, the aux sends are used to set up alternate balances that can feed outputs other than the main pair. That's not the case on the i73 Pro Edge. The main and alternate outputs each have their own separate mixer panels, while the two headphone outputs can separately be switched to pick up either of these mixers, or to be independent and thus have a mixer of their own each. So, in total, the i73's internal mix engine can run up to four independent mix panels, each routed to a separate physical output.

Any processing applied using the DSP plug-ins on the insert slots is common to all of the mixers, as you'd expect. However, each mixer has its own auxiliary send and return structure. The sends are hard-wired to stereo aux channels that are returned into the same mixer, so for example there's no way to use the sends from the main output mixer to feed the mixer for outputs 3+4. But what these aux channels do have is a single insert slot that can be used to host VST3 plug-ins. The obvious application is to host reverb and delay plug-ins for performers' cue mixes, and the Full Analog Experience includes emulations of the EMT 240 'gold foil' plate and Maestro Echoplex, in case you don't already have such effects available to you.

> "The i73's claims to uniqueness are mainly to do with its distinctive preamps and its ability to act as a plug-in host."

The mutual independence of the i73's mixer panels has some pluses and some minuses. On the plus side, not only can you set up four entirely different mix and effects balances for the different outputs, but they can each have different auxiliary effects loaded. You're not bound to use a single global reverb for everyone. It also means that soloing a channel in one mixer has no effect on the others, which in my experience is generally what you want. On the down side, the only DAW return channel that appears in any of these i73 mixer panes is the loopback channel. So, if you set the headphones to be independent but leave your DAW addressing outputs 1+2, you won't hear the DAW output in your headphones. In order to do so, you would either need to create a second DAW output feeding the headphone bus, or route the DAW output to the loopback channel instead of outputs 1+2.

One thing that is missing from the Mac version of i73 Mixer is any reference to ADAT clocking. I eventually tracked this setting down within Audio MIDI Setup on my Mac, but it would be nice if there was some indication elsewhere of what clock source is selected, or indeed if the topic was discussed in the PDF manual. I'm told that clock selection is handled within i73 Mixer on Windows, though on both platforms, changing clock source makes it necessary to close and reopen both i73 Mixer and your DAW project, which is a little annoying.

On balance, though, I think i73 Mixer is pretty well thought out, and feels

The i73 Mixer utility. The first two input channels have insert slots where Heritage's DSP plug-ins can be loaded, while the two auxiliary channels (right) can each host a single VST plug-in.

admirably polished and complete given that this is an entirely new product. There are some unexpected nice touches such as the ability to disable monitor control on the line outputs for situations where you're using the i73 with an external monitor controller, and although the skeuomorphic graphics are a bit gratuitous, everything is generally pretty clear. In an ideal world it might be nice if the UI could be resized, and the positions of the virtual toggle switches in the master section can be hard to discern at a glance, but those are minor gripes.

Heritage have developed a custom Windows driver, but the i73 uses the built-in Core Audio macOS driver. Its mixer has some latency of its own, so at 44.1kHz a 32-sample buffer yields a round-trip delay of 8ms, while the direct monitor path has 1ms or so. Compared with the dry path, the wet signal path through the i73's DSP has additional latency of about 85 samples with no plug-ins loaded, or 91 samples through plug-ins. At present this needs to be compensated for manually in order to align wet and dry versions of the same signal, but a future firmware version will offer the option to delay the dry signal automatically.

In To Out

From an operational point of view, if you've ever used another audio interface with its own mixer utility, I don't think you'll take long to get accustomed to the i73. The only slightly unusual factor is the aforementioned absence of DAW return channels in the mixer. On those inevitable occasions when you scratch your head thinking "Why can't I hear anything?" the answer is either that your DAW is not feeding whatever output you happen to be listening to, or that you've forgotten the i73 always powers up with its outputs muted. (Ask me how I know...)

The i73's claims to uniqueness are mainly to do with its distinctive preamps and its ability to act as a plug-in host. To take the input side of the Full Analog Experience first, it's perhaps worth pointing out that the input transformer in this sort of circuit precedes all the gain stages apart from the input pad. Any saturation or non-linearity the transformer itself imparts is strictly related to the content and amplitude of the input signal, and unaffected by the settings of the main or output gain controls. In other words, you can't "overdrive the input transformer" by turning up the gain. What you can do is overload the active gain stages themselves, then use the output level control to back things off. I suspect that most of what people think of as "transformer saturation" in this type of preamp is actually generated in the active circuitry or in the output transformer, if there is one. (In this case, there isn't, as the preamp is coupled directly to the A-D converter.)

Either way, it's perfectly possible to get clean-sounding recordings out of the i73's preamps as long as you can ensure the signal coming in isn't too hot. For example, when using a pair of beyerdynamic MC930 capacitor mics as drum overheads and with the output control at unity gain, I needed to have the pads on both the mics and the i73 active in order to attenuate the signal enough for the i73 to cope. At this point, the main difference compared with recording the same mics through a modern transformerless preamp was a fairly subtle richness in the low mids.

Switching off either of the pads made it necessary to turn the output control a fair way down in order to avoid overloading the A-D converter. The recorded signal was then visibly rounded off on peaks, and sounded quite different from the clean version. I liked the extra energy and life that this clipping delivered, though you'd need to be pretty careful to avoid overdoing » it. Turning up the gain and turning down the output control further took things into the realms of obvious crunch. Again, there's a place for this in some circumstances, and it's nice to have the option, but it would be easy to ruin a good take by going too far.

In my experience, this sort of obvious clipping also works better on drums than it does on many other sources. The i73's relatively narrow gain range means you may have to take care with vocal recording, for example, if you want the preamps to add warmth without introducing blatant distortion on peaks. But, again, I don't mean that as a criticism; if you want the sound of vintage-style preamps, you have to be prepared to put the effort into optimising the gain structure to find the sweet spot. Incidentally, audio specifications for the i73 were unavailable at the time of writing, but those aspects of its sound that are meant to be clean seemed good to me, and I'd particularly highlight the headphone amps, which have a lot more oomph than those found in most audio interfaces.

The Complete Package?

Although previous products have incorporated digital elements, Heritage Audio are known mainly as manufacturers of analogue gear. So the i73 Pro series represents an ambitious move, incorporating as it does not only USB interfacing but also a new DSP plug-in platform. Feature-wise, its obvious rival is Universal Audio's Apollo Twin range. These are slightly more affordable and offer more DSP power and an enormous catalogue of plug-ins, plus the ability to emulate a wide variety of vintage preamps — but they don't actually have transformers, or the same hands-on approach to setting gain structure for maximum impact.

The Plug-in Experience

Heritage Audio's high-concept sales pitch for the i73 revolves around something called the Full Analog Experience. This refers to two key features: the "73-style" mic preamps, and the inclusion of a suite of six plug-ins modelled on classic recording kit. Four of these are available in both native and DSP formats, whilst the remaining two are native-only, reflecting the fact that they'd usually be used as auxiliary effects in any case. The native plug-ins are authorised using the iLok system and give you the choice of installing VST3, AAX and AU (on the Mac). If you want to use them within the i73 Mixer utility you'll need to install the VST versions.

I've reviewed a couple of audio interfaces that feature clever ways of integrating DSP plug-ins into a native recording environment, notably Avid's Pro Tools | Carbon and Apogee's Symphony Desktop. The idea in both cases is that you can track and monitor at low latency using DSP versions of the plug-ins, then have them switch automatically to native versions to free up DSP resources for further tracking. Heritage Audio haven't attempted anything along these lines, but the DSP and native versions of each plug-in share a preset folder, so it's easy to replicate the same settings in both environments.

Two of the four plug-ins that can be used in the i73's DSP inserts are amp and cabinet simulators. HA 15 PRO emulates the classic Ampeg 'flip top' bass amp, plus a selection of six different cabs, whilst the snappily titled Small Recording Amp Serial #C17744 replicates an old Fender tweed amplifier, this time with a choice of three cabinets and a fairly flexible virtual dual-mic setup. I liked both of these plug-ins a lot, but the DSP versions are clearly quite resource-hungry — so much so that the i73 can't load them both at once. If you want to track bass and guitar simultaneously, someone will have to do without.

More general recording applications are catered for by Britstrip, a channel strip plug-in which combines a 1073-style EQ section with a compressor loosely based on the Neve 2254. This is pretty versatile, and incorporates some nice features such as a wet/dry control for the

The four DSP plug-ins that come with the Full Analog Experience comprise bass and guitar amp simulators, a channel strip and a tape simulator.

compressor and the ability to swap the order of the EQ and dynamics. In this case, though, the skeuomorphic GUI is a bit of an ergonomic drag, thanks to the use of dual-concentric controls that aren't much fun from a mousing point of view.

Finally, the HA 1200 Tape Saturator is a pretty versatile tape emulation plug-in, with a choice of three speeds, five tape formulations and continuously variable bias control, as well as input and output level settings. Its sound runs the gamut from very subtle to not subtle at all!

The two native-only plug-ins are both delay-based effects. One is an emulation of the EMT 240, a compact electromechanical reverb made using a sheet of gold foil rather than the more traditional suspended steel plate, and the other mimics the classic Maestro Echoplex tape delay. Both sounded very good to my ears, though I was not able to compare them with the originals.

Most audio interfaces these days come with a software bundle, but all too often, this is a ragtag collection of odds and sods harvested from familiar names and cut down to the point where they're frustratingly inflexible. Heritage Audio are to be commended for developing an entirely new collection, and even more for the fact that they are all very usable indeed. Apart from the amp simulators, I personally would be unlikely to track through the DSP plug-ins, but I'm sure I will get plenty of use for the native effects in a mixing context.

Included as native plug-ins but not in DSP format are Heritage's emulations of the EMT 240 'gold foil' plate reverb and Maestro Echoplex delay.

Neve's 88M, meanwhile, brings the transformers, the insert points and the Neve name, but doesn't offer an easy way to get a 'distressed' sound from the preamps, and has no built-in DSP.

Are there ways in which the i73 could be better? Of course: the DSP is underpowered, to the point where you can't run amp simulators on both inputs at once, and the benefits of the preamp design would be even greater if it had insert points. But I think Heritage have got the major calls right here. Their suite of plug-ins is surprisingly great, the software mixer is slick and feels mature, and the balance of features is well judged. And above all, this is the first desktop interface I've ever encountered that offers really obvious colour from its built-in mic preamps.

$ i73 Pro One $649; i73 Pro 2 $999; i73 Pro Edge $1499.
T Heritage Audio +1 800 605 3127
E support@heritageaudio.com
W www.heritageaudio.com


L — Out — R
L — In — R
MIDI In
MIDI Out
LITTLEEND
ATK DEC SUS REL
LEV FREQ RESO TYPE ENV
Main Volume
Level/Data
Preset Pool
Digitakt II
TRIG
Quantize
SRC
Machine
Setup
YES
Save
NO
Reload
FUNC
Perform
Save Proj
Samples
Tap Tempo
Copy
Clear
Paste
KB Setup
TRK
Mute Mode
PTN
Bank
SONG
Song Edit
Track/Mute

Elektron Digitakt II

Digital Drum Computer & Sampler

The original Digitakt was pretty good, so how about an updated version with more of... everything!

SIMON SHERBOURNE

Elektron describe the Digitakt as a Drum Computer and Sampler, and for a few years after its launch in 2017 it became more-or-less the default drum machine in hardware synth rigs. It's versatile, though, and despite being the darling of the DAWless, it features sophisticated computer integration with plug-in-based control and audio over USB.

For the Digitakt II, the word 'stereo' has slipped into the subtitle, but this is just one of many upgrades. This is the first outing for a new underlying hardware platform that pushes past some of the long-standing spec limitations shared by Elektron's instruments. So, does the Digitakt II do enough to stay on top in a crowded field? Should you upgrade if you have the MkI? And does it leapfrog flagship Elektron devices like the Octatrack II?

Drum Computer

The Digitakt II is a multitrack sampler built around the Elektron sequencing engine. It's structured like a drum machine or beat workstation with 16 monophonic, stereo tracks. Tracks default to playing one-shot samples, but can also be used for loop playback, grid-based slicing or MIDI sequencing.

The Elektron instruments share an architecture and workflow that can take some time to get your head around, but once you're over the hump you can work quickly, and get straight to work on any of the other devices. The Digitakt II conforms to the standard Elektron MO but, as we'll see, extends it in a number of ways.

Work on the Digitakt II happens within self-contained Patterns that store sound and mix parameters as well as note and automation sequences. You can build a beat in a single pattern using mutes, fills and conditional properties to create variations, or you can compose across a bank of patterns. Eight banks constitute a Project, whose other chief property is to store the bucket of live samples that are held in RAM.

Reboot

At a glance, the MkII Digitakt appears unchanged, with its pleasingly compact metal build evocative of high-end field-recorder hardware. There are some differences, such as the display now being white like on the Syntakt and the top tier Elektron devices. There's one additional button providing direct access to keyboard play mode, a nod to the fact that with 16 tracks and stereo capability the MkII lends itself to more complete compositions than just drums. Likewise Song mode now gets a top-level spot.

Elektron Digitakt II
$999

PROS
- More of everything.
- New Elektron sequencing features.
- Sounds massive.

CONS
- If it had a built-in battery it would be the perfect groovebox.

SUMMARY
A big upgrade to an already great sampler and drum machine that makes itself useful in almost any studio or live situation.

Finally, there are eight sequence page LEDs instead of four, reflecting the newly doubled maximum pattern length.

Internally, as well as a more capable brain, RAM has increased from 64MB to 400MB, and the internal flash storage from 1GB to 20GB. With more RAM, direct sampling time increases to 66 seconds, and longer samples can be imported. More significantly, the RAM allows a Digitakt II project to keep eight times as many samples in its live sample pool. The pool now has eight banks of 128 samples available for playback and 'sample locking': the Digitakt's ability to fire any sample or preset on any step.

The Digitakt II is a class-compliant USB audio device, with stereo audio in and out to your computer or phone. Extended Overbridge functionality is still in beta testing, and should bring multitrack audio streaming and remote control directly in your DAW, as with the original Digitakt plug-in.

Basic Beatmaking

One of my chief complaints about the Digitakt was the lack of a real concept of Kits. Loading sounds could be confusing and time consuming to a new user.

»

Round the back, things are identical to the Digitakt I, with MIDI in, out and thru ports and stereo audio I/O and headphone output all on quarter-inch sockets.

Things are more immediate and familiar now, with a clear workflow for browsing and auditioning Kits and Presets (previously called Sounds) without needing to understand and populate the sample pool first. A set of excellent factory kits and a couple of banks of inspiring example patterns show the results possible once you've mastered the sequencer.

With a kit in place, you can play sounds from all the tracks with the 16 non-velocity-sensitive 'trig' buttons. If keyboard mode is active, the selected track will be pitched across the keys for melodic playback. Digitakt II gains some additional trig modes. Velocity mode mirrors the 16-level mode on an MPC: allowing one track to be played at different velocities with the keys. Retrig mode is a note repeat function, again playing a single track but with 16 retrigger rates. Finally Preset mode maps the project Preset pool across trigs in banks of 16, making it easy to play and record multiple sounds within a single track. This is a more powerful version of the Slot trig mode on the Octatrack.

Pattern sequences can be recorded in real time using the current trig mode, or entered X0X style one track at a time on the keys. There's also a step record mode. Patterns can be up to 128 steps long — eight bars at the default speed — and tracks can be set to different lengths and speeds. Everything can be automated as part of the Pattern, either by capturing in real time, or by holding any step and dialling in parameters: a workflow innovated by Elektron and dubbed Parameter Locking.

Track effects are now accessed from a single page.

A first for Elektron sequencers: chance, fill and trigger conditions have independent slots.

Digitakt II tracks have swappable playback and filter 'machines'.

Sample Machines

When we first reviewed the Digitakt, each of its eight audio tracks worked in the same mode, playing samples in a familiar sampler-y way. Playback could be pitched per step for melodic sequencing, you could set start, end and loop points, and each track had a filter, bit-crusher and audio effects. While conventional, I noted that the sampler engine was unusually nimble, able to smoothly play and sweep very short loops to create synthesis and granular-like results.

Over time, the Digitakt gained new sample playback modes. The MkII inherits these and gains one more, along with swappable filter modes, which are collectively referred to as Machines. The original and default machine is called OneShot, and plays samples in the same linear way as before, but in stereo.

The Werp machine is a time-stretching mode, which breaks the audio into segments. It's a fairly crude stretching algorithm but playing with the segment control can produce some interesting and unexpected results. This mode is maintained on the MkII, but is joined by a smoother granular warp mode called Stretch. This is now the go-to mode if you want to play imported loops in time with your Pattern tempo with independent control of pitch.

Repitch mode automatically adjusts the playback speed of samples to match a set number of bars. For the Digitakt II, Slice has been renamed to Grid, presumably to more accurately reflect what this mode does. While it does slice your sample into sections, it's always regular divisions of a bar. This covers you for scenarios where you're chopping up a neat loop, but doesn't give you the manual chopping workflow you take for granted on an MPC, Maschine, or for that matter the Octatrack. You can lock which slice plays back on each step, or play from keyboard mode (though there's an issue that if you have a scale set some slices won't be accessible).

Filters & Effects

The filter module for each track can be chosen from Multi-mode, Lowpass 4, Equalizer, Comb and Legacy LP/HP. The default Multi starts out as a 4-pole low-pass, but can be morphed to a notch then a high-pass. This is super versatile, and you also get a second non-resonant high+low-pass filter stage before the main filter, which is a real luxury. The primary filter has a dedicated envelope, with a welcome Reset option to choose whether legato notes retrigger the envelope or not.

The Lowpass 4 mode is more 'analogue-y' sounding, with a more tuneful resonance. Equalizer is a single-band parametric EQ. Comb is a short delay-based comb filter combined with a low-pass, which is tons of fun for modulation effects and resonant metallic sounds, and becomes a slap-back echo or reflection at lower settings. The Legacy filter models the Digitakt I filter, which is good for backward compatibility and has a character of its own with a more aggressive resonance.

The effects offerings have been upgraded on the Digitakt II, and everything is now organised in one FX page per track instead of three. The

Is The Digitakt II The Octatrack III?

One question I've heard is whether the updated Digitakt essentially replaces Elektron's Octatrack. The Octatrack is an outlier in the range, running on an older platform incapable of modern niceties like even basic USB functionality let alone audio streaming or Overbridge. There is overlap between the devices: they are both multitrack samplers and MIDI sequencers running mostly the same sequencer, and they can both be used as a drum machines or as a workstation hub for a hardware performance.

The Digitakt II has more audio tracks, higher-quality effects, newer sequencing features and modern USB connectivity, but there are many things still unique to the Octatrack. It has a true slice mode for a start, although maybe the Digitakt II will get that one day. But it also has a set of performance features that are unmatched. It has multiple inputs, and incoming audio can be routed through tracks. It can be a looper and can sample and replay on the fly using sync'ed record triggers. And it has snapshot scenes that can be morphed between with the crossfader. It's unlikely this functionality will get rolled into a Digitakt II update, so I'm still hoping we'll see a next-gen Octa.

track effects (bit reduction, overdrive and sample rate reduction) and the shared reverb and delay sends have been joined by a chorus send and a master overdrive.

It's now possible to set specific tracks to bypass the master compressor, so you can set up a classic side-chain pumping scenario, where your kick is set as the side-chain trigger but passes through unaffected.

Modulation

The Digitakt has seriously upped its modulation game, with three LFOs per track in addition to the dedicated amp and filter envelopes. This puts the Digitakt II ahead of its larger peers the Analog Rytm and Analog Four for LFO count. The LFOs are versatile, can be triggered, sync'ed or free running, and have single- and half-cycle modes that turn them into extra envelope generators. LFO 3 can even modulate one of the other LFOs if you like.

I do wish that Elektron would find a more intuitive way to assign modulation on all their devices. The fact that mapping is done by selecting from a long scrollable list always means that I don't use them as much as I could.

The Elektron sequencer, with its automation capabilities, is itself a powerful mod source. There's also the facility of 'trigless trigs': you can set steps to trigger envelopes or other modulations without playing a note. These are especially useful when playing longer samples or loops which only have a single playback trigger, but might contain beats and notes that you want to coincide with filter envelope triggers.

Sequencer 2.0

The Elektron sequencer has been consistent for a number of years, but has had some recent updates which the Digitakt II inherits, and there are advances unique to this new instrument. In common with the Analog Rytm and Four you get the Euclidean sequencer mode for auto-generating track patterns. Euclidean sequences can be predictable, but here you have two separate layers which can be phase-adjusted and set to interact in different ways. A nice touch is being able to write the current pattern to the track.

Page Loop lets you choose specific pages in a multi-page sequence to play. Ostensibly this is so you can edit one section of a longer sequence without

The Digitakt II shares exactly the same dimensions as its predecessor, measuring 215 x 176 x 63mm.

having to keep waiting for playback to go around the horn. However, it's actually a handy performance tool, allowing you to park pattern variations or fills on different pages, then instantly jump between them. Disappointingly, the MkII Digitakt still lacks Direct Jump mode from the Rytm, which lets you do this with Patterns, but Page Loop can get similar results.

The first truly new feature is the doubling of the maximum Pattern length, putting 128 steps at your disposal. Then there's the ability to set multiple trigger conditions per step. The Elektron sequencer typically has a single parameter slot where you can set a percentage probability to all steps in a track and also place per-step Conditional Locks. These apply either a different probability amount or some other event condition like 'Fill' or 'Play on fourth of four bars'.

The Digitakt II has separate slots for Probability, Fill and Condition. So you could now set a snare hit to play only while Fill is held, but only on every second bar, 70 percent of the time. With liberal use of Conditional Locks you can squeeze a lot of variation into those eight bars, and even have completely different sequences baked into Fill mode. If that's still not enough, the Digitakt II can chain patterns together, and has a Song mode where you can playlist patterns, adjusting loops and lengths at each stage.

Breaking out of the single condition limit is massive, and I hope this gets implemented on the other devices. Even so I'd love to see more; conditionality is an innovation associated with Elektron sequencing, but other devices have taken the idea and run with it, including the Polyend Play and Teenage Engineering OP-Z, and sequencers like the Five12 Vector and Sequentix Cirklon. It would be great to have conditions or sequence modifiers that add generative operations beyond simply 'trigger/don't trigger', for example random or set pitch changes and retriggers.

Everything Everywhere All At Once

A highlight on the smaller Elektron instruments is Control All mode, which helps to make up for the lack of the performance macros and scenes found on the larger devices. It's dead simple: while you hold the Track button, parameter changes are applied to all tracks at once. It's a brilliant tool for creating build-ups and breakdowns with no prep work. It works hand-in-hand with the Temp Save and Reload functions, allowing you to totally scramble everything in the name of dancefloor excitement, then punch back to the last sane state.

The Digitakt II refines the Control All trick, by offering the chance to exclude specific tracks. For example, you can keep your kick drum unchanged, while maxing out delay and bit crush on everything else.

Also new is Kit Perform mode, which stops Patterns from loading their own parameters when launched. This allows you to tweak sounds and effects, and smoothly move between Patterns without affecting the sound — only the sequences get loaded. This is a big improvement for performing, or when adjusting sound and mix parameters across a multi-pattern song. However, to commit changes still requires saving to a kit and loading to each Pattern.

> **"The Digitakt II is still the Digitakt you know and love, but doubles up its core abilities: twice the audio tracks, stereo, double the sequence length, double the MIDI control scope."**

MIDI, Externals & Sampling

The Digitakt I has eight audio tracks and eight MIDI tracks arranged as two rows. The Digitakt II's 16 audio tracks occupy all the keys by default, but you can switch any track to MIDI in the Machine selector. Control of external hardware or software has had some enhancements. You now get two pages of CC assignments per track, and two MIDI LFOs.

CC values can now be learned from your external instrument by patching MIDI from it into the Digitakt II and wiggling controls. Mapped knobs can now be renamed. Connectivity wise, the Digitakt II has MIDI in, out and thru on DIN connectors, and computer connectivity over USB. I'd have loved to have seen a USB host port, and a battery option would have made it my perfect portable groovebox. Audio from one or two external devices can be routed into the Digitakt (stereo or dual mono), and monitored through the mixer page, which offers level control, panning and access to the send effects.

The sampling workflow remains much the same on the MkII except that you have twice the sampling time and can sample in stereo. Available sources are the input pair, the mix out, or any of the individual tracks. It seems an oversight that you can't choose either of the inputs individually. If you have a mono source connected it still gets sampled in stereo and ends up hard panned.

A nice feature is that you can choose a sampling time in steps, so if your source is clocked to the Digitakt you can grab, say, one bar ready to loop. Also nice is that after sampling something you're prompted to assign it to a track straight away if you want. (I don't know why the same courtesy isn't extended when you browse samples from storage: instead you have to load to the pool, then go and find it in a track.)

Conclusion

The Digitakt II is still the Digitakt you know and love, but doubles up its core abilities: twice the audio tracks, stereo, double the sequence length, double the MIDI control scope. There's more memory and sample slot capacity, more modulation, and more effects and filters. It also just sounds immense. This is a digital sampler, but as well as super crisp sonics, the built-in send effects and the drive and compressor combine to give the Digitakt a huge floor-shaking character.

The Digitakt is not trying to do everything like an MPC or Maschine, and I enjoy using it in a focused role as a drum machine. That said, it does have some excellent performance capabilities, it's a great standalone groovebox that can also stand in to orchestrate a compact DAWless rig, and it can act as a bridge to your computer. In short, it's 100 percent still a top contender for a place in any kind of hardware setup. As for upgrading, if your Digitakt I plays a peripheral role as sampler or drum machine it'll continue earning its keep, but if the Digitakt was my main workstation I'd certainly be looking to swap it out for the MkII.

$ $999
W www.elektron.se

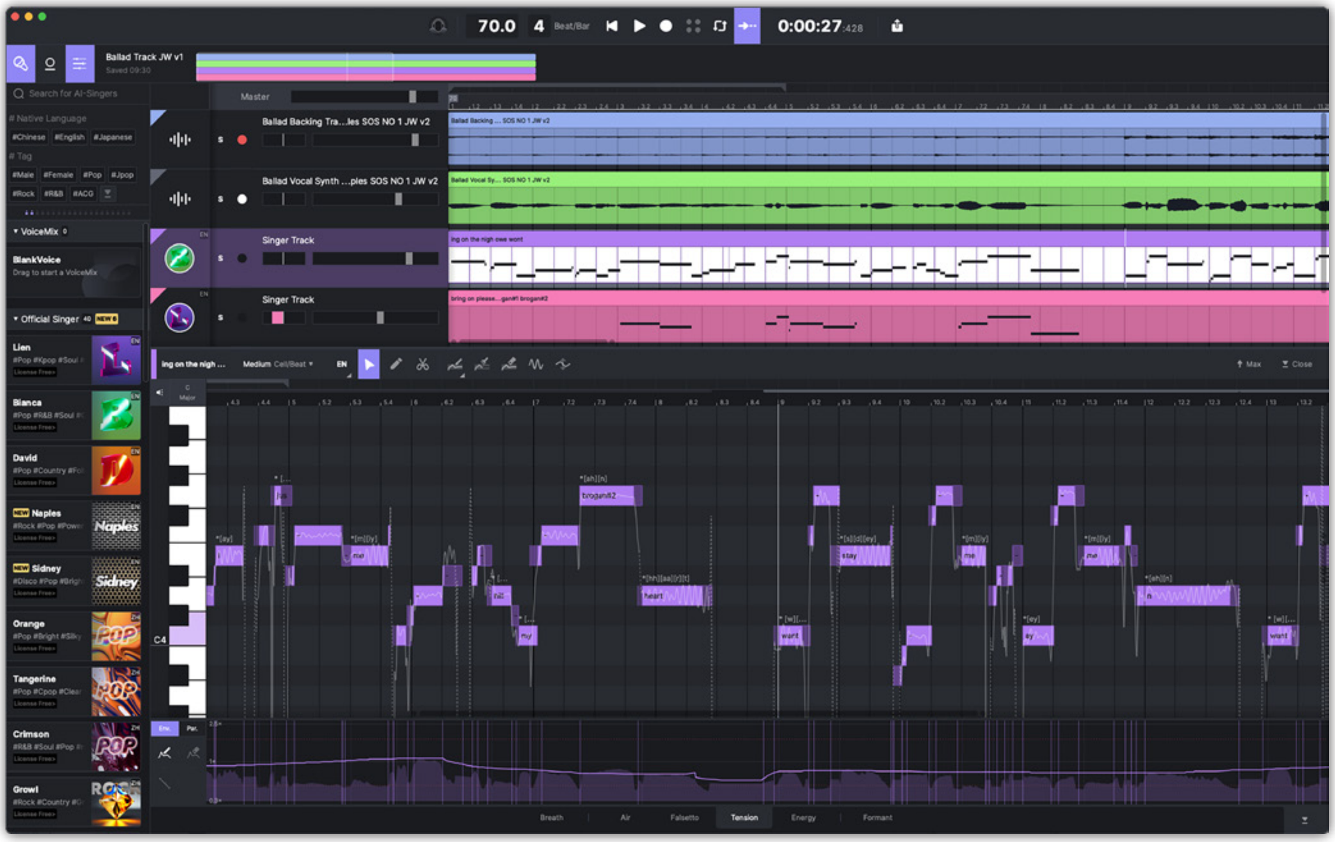

ACE Studio

Vocal Synthesizer Software

ACE Studio is an AI-powered virtual vocalist that's frighteningly close to the real thing.

JOHN WALDEN

While virtual drummers, bassists, pianists or guitarists have already passed the point where they can generate a believable impression of a human player, virtual vocals have, for understandable reasons, proved more challenging. Vocaloid has been a long-standing option for a more synthetic style. However, as I experienced when reviewing Dreamtonics' Synthesizer V in the March 2023 issue of *SOS*, on the back of advances in AI things have started to move very rapidly when it comes to more natural-sounding results. Synth V is not the only product riding this wave, though: ACE Studio is also generating a lot of interest.

So, if you could use a virtual session singer within your own music production workflow, should you be booking ACE Studio for an audition?

ACE Concepts

ACE Studio operates as a standalone application. However, the ACE Bridge plug-in (AU and VST3) can be loaded within a suitable DAW, and this provides both sync and audio routing from ACE Studio back to your DAW. The package currently offers some 40+ AI-based vocalists, each with their own singing style. The AI engine supports Chinese, Japanese and English for all vocalists but each vocal database is 'native' in one of these and there are currently five native English singers; Lien, Bianca, David, Naples and Sidney.

Transport and menu bar aside, the UI offers three main sections; the Singer Library, Arrangement View and Clip View. Located far left is the Singer Library containing a list of available singers. The other two views dominate the rest of the display. At the top, the Arrangement View provides an overview of the various Audio Tracks (for example, to house an instrumental backing track) and Singer Tracks within the current project. Singer Tracks are essentially hosts for the MIDI-based clips upon which the voice synthesis is based. The contents of individual clips — audio or MIDI — can be shown in the lower Clip View panel and it's here that the key tools for editing notes, adding lyrics, customising the pitch and automating the singer's dynamics/character (via a parameter panel at the very base of the display) are to be found.

Two additional panels, the Track Control Panel and the Mixer, can be popped open as required using the buttons located top left. The latter is self-explanatory (and shown in one of the screenshots) while I'll say more about the former below.

A couple of other practical points are worth noting. First, ACE Studio is currently only available via a subscription model. An active subscription gives you access to all of the available singers. Second, and in contrast to the difficulties that might arise in using AI-generated vocals from some online sources (where the copyright of the AI's training materials may cause issues), the

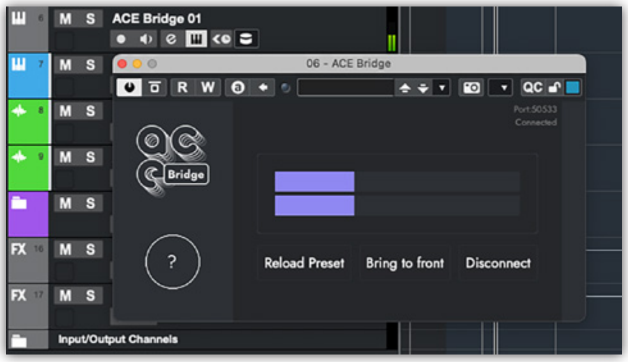

The ACE Bridge plug-in allows you to easily sync ACE Studio with your DAW.

AI vocals you create with ACE Studio are completely licence free to use within any commercial context. There are a few singer exceptions that are highlighted within the Singer Library panel and an online licensing process is available for these.

Studio Workflow

Used as a standalone application, the most obvious workflow involves importing some sort of audio backing track for your song project to use as context for any vocal parts you wish to generate. Once you have constructed whatever vocal parts you need — and these can involve multiple Singer Tracks — those vocals can be exported as audio files to use within your DAW for mixing/arranging alongside the rest of the instruments within the project.

Having selected an AI singer for a Singer Track, you can manually enter notes with the various note editing tools or import a MIDI file (for example, if you have already written an initial melody line within your DAW). At present, it doesn't seem possible to record MIDI data directly into ACE Studio. There is a third option for creating your initial 'MIDI with lyrics' data. It's included within the software but currently classed as a beta feature. It is, however, something quite special and I'll come back to it later.

Whether you have added new notes within the Clip View, or edited existing ones (pitch, lyrics or other parameter adjustments), when you initiate playback, ACE Studio has to pre-render any changes. A technical point is worth noting here; that rendering is done in the cloud via ACE Studio's own servers. That obviously requires an active Internet connection (and an active subscription). The process is fairly rapid but it does involve a short wait each time the vocal is reprocessed.

There are some really nice touches with the note editing process itself. For example, within a note, shading indicates the portion of each note's length set aside for either consonant or vowel sounds, and you can drag with the mouse to adjust this to manipulate the way a word or phoneme is sounded. You can, of course, also edit the phonemes themselves (if you zoom in close enough, these are shown above each note). You can create/edit a user pitch curve if you want to adjust that generated by the AI engine. The mouse-based graphical editing of pitch modulation and vibrato are particularly well designed and incredibly easy to use.

Regardless of which virtual singer is used, you get six parameters — Breath, Air, Falsetto, Tension, Energy and Formant — that you automate via the parameter panel at the bottom of the Clip View. Breath allows you to add breaths into the performance for added realism, while Formant lets you adjust the 'gender' of the voice. Used subtly, it can provide a useful shift in character if required. The other four parameters allow you to add various types of expression or character to the performance to squeeze as much 'human' out of the engine as you possibly can.

Three other workflow features are also worth noting. First, ACE Studio offers a very effective Vocal Double feature. This automatically creates two additional versions of the selected track on two new Singer Tracks. These will have timing and pitch variations applied and you can adjust the tightness of the double tracking before generating the new tracks. It works really well and it's impressively easy to use. After generating the doubles, you can, of course, then change the singer used or adjust the pitch to create harmony parts.

Second, the Track Control Panel I mentioned earlier provides the very interesting Voice Mix feature. Essentially, this allows you to blend multiple Singers together to create a composite voice, with control over the level of Style and Timbre characteristics used from each voice. This means you can create unique voice styles to suit your needs and then save your creations as a new voice preset.

ACE Studio
From $149 Per Annum

PROS
- Capable of genuinely useful results.
- Impressive a capella to virtual vocal conversion process.
- Very slick pitch, modulation and vibrato editing.

CONS
- Subscription model and cloud-based rendering require an Internet connection at all times.

SUMMARY
Vocal synthesis is making remarkable progress. In experienced hands, ACE Studio is capable of generating genuinely useful results.

Third — and the special beta feature I mentioned earlier — is an alternative workflow for creating the MIDI note/lyric data for a Singer Track. If you import an audio file that contains a guide vocal (cleanly recorded without reverb or delay effects or vocal harmonies), you can drag it onto a Singer Track. ACE Studio will then analyse it and convert it into a MIDI clip, add the required MIDI notes at the appropriate pitch, and attempt to add the lyrics from the audio. If you are happy to record a scratch demo of your vocal, even if your own performance is not fit for public consumption, this is a workflow that offers a massive time saver over manually entering notes and lyrics, even if some subsequent editing is still required. And, of course, your demo vocal can then be sung by any of the 40+ AI singers (male or female), and pitch, lyrics and delivery style can all be adjusted to suit the song's needs.

Get Real

So, the feature list is impressive, but how good are the actual results? Well, I'll qualify »

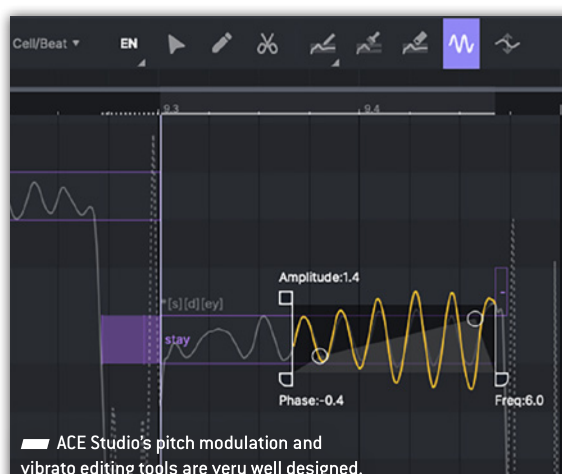

ACE Studio's pitch modulation and vibrato editing tools are very well designed.

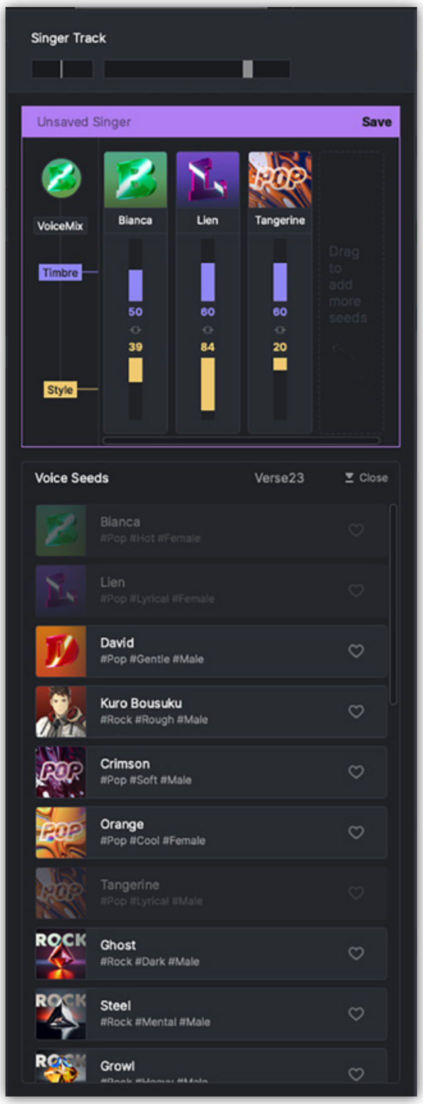

The Voice Mix panel lets you create unique blends for the included AI singers.

my comments with an acknowledgement that, despite working with ACE Studio over a couple of weeks, I'd still describe myself as a new user. Like Synth V, this is an application that can produce impressive results almost instantly, but experience will bring further incremental improvements. That qualifier made, how far might we go up the quality scale? Well, if you need some backing vocal parts or vocal hooks to place within an electronic dance track, ACE Studio can deliver. Equally, if you need to create a fully realised demo vocal for a song project (for example, so that your actual human session singer has a better idea of what they are shooting for), again, ACE Studio can do the business.

With enough experience to fully capitalise on all the features the synthesis engine offers, I think you could also pull off a lead vocal (including any harmonies) for a contemporary EDM or busy pop track and many listeners would simply not realise it was a computer-generated vocal. It might be more difficult to do in a bare-bones piano/vocal ballad, where the voice is more exposed.

In all cases, while the vocal might sound technically proficient, instilling it with all the subtle details of emotion that a really good singer brings to a performance will be a challenge. But, goodness me, it's still impressive stuff; ACE Studio's vocals will be able to pass the test in plenty of potential use-case scenarios. And, as a taster of what's possible, I've created a few short audio examples, which you can audition on the *SOS* website, at https://sosm.ag/ace-studio-audio.

Singing Competition

A comparison with Dreamtonics' Synthesizer V is inevitable. The two products attempt broadly the same task, and in terms of their UIs, the workflows have many things in common. There is a lot that could be said here so I'll confine myself to the most obvious of pros and cons. For instance, when it comes to manual pitch editing, I think the approach adopted within ACE Studio currently has the edge, particularly when it comes to manually editing vibrato. Equally, the option for 40+ singers straight out of the box, and the interesting Voice Mix feature, will undoubtedly appeal to some potential users. However, for English language vocals (I'm not qualified to comment on those of the other languages supported), I do think Synth V's voice synthesis is currently superior and provides better pronunciation, leading to a more realistic end result. I also prefer Synth V's voice-specific 'vocal modes' system to control the expression and dynamics, giving each of the virtual singers a very individual character.

There are also some technical differences between the two. ACE Studio is based upon a subscription model and uses cloud-based rendering, whereas Synth V is a one-time payment (with each AI voice also being a separate purchase) and all rendering is done locally. Incidentally, the automatic a capella-to-MIDI conversion process is also a feature that Dreamtonics have added in the current public beta of Synth V, and it works very well.

It's hard to imagine that the obvious competition between the two products is not pushing both development

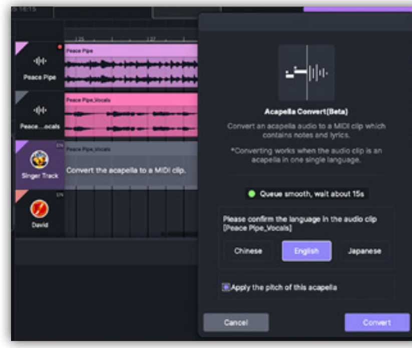

The option to automatically convert an audio vocal into an AI vocal provides an impressive workflow option, and it's quite remarkable to see how well it works.

teams forward. The resulting pace of development is an obvious up side for existing users of either product. However, as a down side, it might make a purchase decision for potential new users more difficult as some feature leapfrogging will undoubtedly occur over the next year or two through their competing release cycles. Watch this space...

Holding All The ACES

Whether you like the concept of a computer-generated lead vocal or not, products like ACE Studio and Synthesizer V represent truly groundbreaking technology. Bells and whistles of their respective workflows aside, it's worth reminding ourselves just how remarkable it is that AI vocal synthesis can even get close to sounding human. If you already have a virtual band of musicians within your studio computer, the thought of adding a virtual vocalist might be a tempting one. Personally, I think these products are now at a level where that's a realistic proposition and, in the right context, these tools have very practical applications.

Thankfully, you can find out for yourself whether you agree with me as both ACE Studio and Synth V provide options to try before you buy. In the case of ACE Studio, that's via a free trial period, and I'd encourage any potential user to take advantage of that to fully explore what's possible. Be prepared to invest a little time though; experience and experimentation are required to get the best from the engine. I know this final sentence might sound a little weird... but virtual singers are now a real thing.

$ From $149 per annum.
W www.acestudio.ai

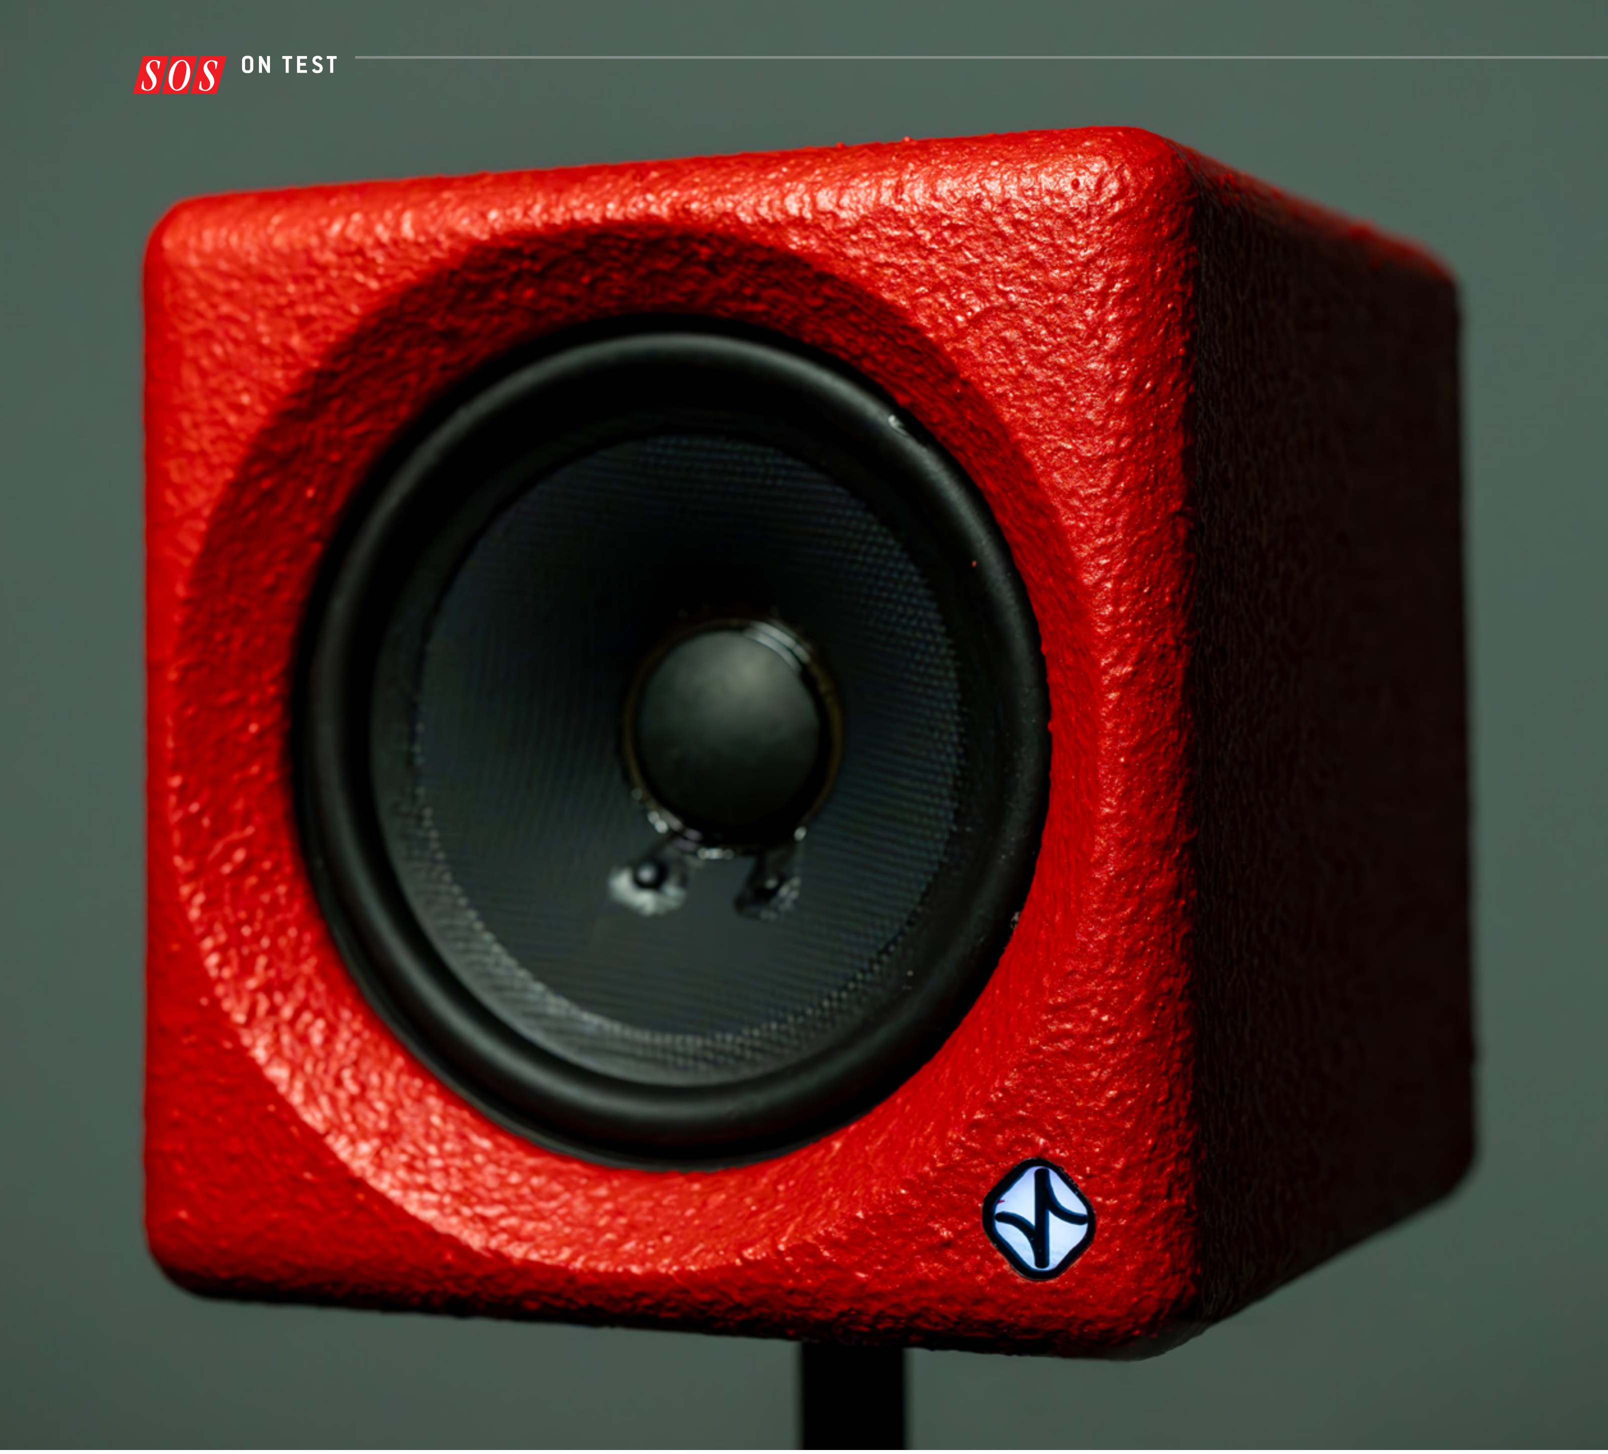

Tantrum Audio Angry Box

Single-driver Active Monitors

Sound and fury, signifying nothing... or a remarkably useful mixing tool?

MIKE SENIOR

One of your most important tasks when mixing for the mass market is to ensure that the most important elements of the music survive their unpredictable journey to the listener. Many end-user playback systems are way too small to transmit appreciable low end, for instance, and the large-scale public-address systems in public spaces rarely offer much bass either, simply because subwoofers are large and expensive to install. At the other spectral extreme, high frequencies are extremely directional and easily shadowed (by obstacles in the listening environment) or masked (by background noise), so plenty of listening situations provide no meaningful high end either.

A common response to this problem is to check your mixing decisions with a bandwidth-limited loudspeaker, thereby focusing your attention at mixdown on the midrange area of the spectrum, since that's the frequency region that's likely to reach the greatest number of real-world listeners. Probably the

Tantrum Audio
Angry Box
$999

PROS
- Genuine, professional-grade single-driver monitoring in a small, lightweight package.
- Extended-bandwidth mode offers a good deal of additional general-purpose mixing power into the bargain.
- Internal power supply.

CONS
- Less severe midrange focus than some competitors potentially delivers less robust mass-market mix translation.
- A little pricier than the competition.

SUMMARY
A powerful midrange reference monitor which also offers a great upgrade path for up-and-coming mix engineers currently mixing on headphones.

most famous example of this kind of product is the original Auratone 5C Super Sound Cube from the '70s and '80s: a small loudspeaker that not only naturally generated a midrange-heavy tone, but which also, by virtue of its single-driver, closed-box design, provided a forensically accurate presentation of the midrange balance, without the crossover artefacts and resonance problems inherent in most wider-bandwidth multi-driver studio monitors.

In recent years, a number of companies have attempted to update the original Auratone concept with varying degrees of success. Having reviewed a number of these products in *Sound On Sound* over the years (including Auratone's modern reissue), the most successful one to my mind has been Avantone's Mixcube. It's not quite as severely mid-focused as some of the original Auratones, but just as clean and fast, and at least as useful at mixdown, in my opinion. However, UK-based start-up Tantrum Audio are now offering their Angry Box as a new and slightly higher-priced competitor. Can it challenge the Mixcube's supremacy?

Hardware

The Angry Box is a cubic sealed box containing a 4.5-inch driver, a 65W Class-D amplifier and some DSP processing. The whole thing measures 14cm across and weighs only 2.5kg, which makes it significantly smaller and lighter than the Avantone Mixcube; it also has an internal PSU, rather than the Mixcube's large external line-lump. The rear panel provides a balanced combi jack/XLR input, a switched IEC connection for mains power, and a three-way volume switch, which can deliver some extra headroom should the front-panel clip-warning LED turn red. A further rear-panel Mode switch offers two different sonic flavours: a bandwidth-limited mode with negligible DSP processing for midrange monitoring purposes, and an extended-bandwidth mode for more general-purpose auditioning.

The speaker's eye-catching red textured finish feels very much in keeping with the product's disgruntled moniker, and generally feels solid and robust. Recessed into the speaker's underside are four rubber non-slip feet, although on the models I reviewed, I found that these didn't quite protrude enough to give a confident grip at all four corners when the units were placed on hard flat surfaces, so I ended up using rubber matting to ensure the speakers remained securely anchored to my stands for review purposes.

In The Mids

Given that there are two different tonal options on offer here, let's deal with them one at a time, starting with the midrange focus mode. As I'd hoped, this was extremely effective at clarifying critical questions of midrange balance and, unlike some other small speakers that I've tested in this capacity, the Angry Box did a commendable job of rendering every one of my own personal mix reference tracks with a sense of accuracy and familiarity. The transient definition and detail transmission were both excellent, too, contributing to both an impressively sharp stereo image and decent depth presentation.

Comparing the Angry Box side by side with my own trusty Mixcube, the Angry Box's midrange emphasis felt a little less severe overall, as well as being 'centred' maybe a half-octave lower, giving a slightly warmer and friendlier subjective tone — a tonal character particularly audible on dense rock textures like Linkin Park's 'Forgotten' or Royal Blood's 'Lights Out', say, as well as on tracks with full-bodied kick drums, such as Dynoro & Gigi D'Agostino's 'In My Mind' or Post Malone's 'Circles'. Initially, this left me with the impression that the Mixcube's transients were perhaps a little faster and tighter than those of the Angry Box, but this feeling all but dissipated after more extended auditioning. If you held a gun to my head, I'd still probably give the Mixcube the edge in terms of cleanliness and detail on complex arrangements like Outkast's 'Bowtie', Justin Timberlake's 'Can't Stop The Feeling', and David Guetta's 'Don't Leave Me Alone', but there honestly wasn't a heck of a lot in it.

The bottom line is that I'd happily use either the Angry Box's midrange mode or the Avantone Mixcube for any professional mixing job. Indeed, for comparison purposes I deliberately used the Angry Box exclusively to do the detailed lead-vocal fader automation for a mix I was working on (probably the most critical midrange balancing task of all), and when I subsequently checked my work on the Mixcube I really couldn't fault the results at all. Nuff said!

Now, there's an argument that the Mixcube's more severe midrange focus might give a better indication of mix translation, given that the Angry Box's presentation retains some mix elements that fall off the extremes of more restricted-bandwidth mass-market playback systems — both the high-frequency fizz in the choruses of Dr Dre's 'Housewife' and the bass low-end weight on Bonnie Raitt's 'Right Down The Line' are noticeably more present on the Angry Box than on the Mixcube. Set against that, however, is that the slightly wider frequency context provided by the Angry Box made the »

A Tale Of Two Monitors

It's important to stress that the Angry Box featured in this review is actually the second iteration of the design, incorporating a number of updates in response to feedback from early users. At the time of writing, though, some remaining stock of the speaker's first iteration still appears to be on sale from some retailers. Given that the first-iteration hardware I auditioned offered only the extended-bandwidth mode (not the more useful midrange focus option), and also had some unresolved DSP-related performance issues, I would encourage any potential purchaser to avoid that version if possible. (To clarify: the second iteration is the one with the Mode switch.)

Internal DSP allows for two listening modes: the default 'mid focus' mode, and an extended-bandwidth mode.

sonic transition from my main monitors a little less jarring, with the result that I found the process of balancing lead vocals against a backing track a touch more intuitive. The Avantone is great at highlighting relative vocal levels, but it does often push the vocal so far out front of the rest of the balance (especially if you listen in single-speaker mono, as I would usually do) that you sometimes have to fight against the urge to turn it down while automating! With all this in mind, I wouldn't personally chalk up the tonal difference between the two speakers as a clear win for either camp.

Highs & Lows

What the Angry Box does definitely have in its favour, though, is the extended bandwidth mode, which uses the onboard DSP to flatten out the speaker's frequency response such that it more closely resembles that of a traditional small nearfield. Specifically, it extends the bandwidth to 90Hz-12kHz at the -3dB points and 60Hz-20kHz at the -10dB points. Despite my concerns that this DSP processing might introduce unwelcome resonance issues at the low end, I was pleasantly surprised that this didn't seem to be much of an issue in practice — certainly not by comparison with the bass overhang problems I've often heard on small ported speakers in this price range. The sound retained a good deal of its speed and punch, such that the only substantial down side of the DSP processing at the low end was the steep ported-speaker-like frequency roll-off, a frequency-response characteristic that makes balancing kick drum and bass rather hit-and-miss depending on their spectral content (the missing bass line on Justin Bieber's 'Boyfriend' being just one case in point), and also undermines your ability to reliably compare the low-frequency balance of your mix against commercial reference tracks.

The lack of spectral extension at the high end (an inevitable by-product of the tweeterless design) will likely dissuade anyone from seriously considering the Angry Box as their main monitoring system, but if you consider the typical purchasing process for entry-level mix engineers, I can nevertheless see that the Angry Box's extended-bandwidth mode could serve as quite a useful halfway house for project-studio engineers on a budget. Imagine, for instance, that you've started out mixing on decent-quality headphones, as a lot of *SOS* readers do, but are considering upgrading to speakers. Now, for the price of the Angry Boxes, I've yet to find any pair of budget active nearfield monitors that honestly offer solid value for money in terms of increasing your ability to generate commercial-quality mixes if you've already got high-quality studio headphones. Buying the Angry Boxes under those circumstances, though, would give you genuinely professional-grade midrange monitoring, plus a 'bonus' extended-bandwidth monitoring experience that doesn't actually lose much ground to similarly priced two-way nearfields. Furthermore, the Angry Boxes might actually be a better long-term investment under these circumstances than similarly priced traditional nearfields, because they'd continue to serve a vital monitoring role even if you eventually expanded your monitoring system in the future to include higher-fidelity full-range nearfield monitoring.

"The Angry Box did a commendable job of rendering every one of my own personal mix reference tracks with a sense of accuracy and familiarity."

Sounds & Sweet Airs

Overall, then, I think Tantrum Audio have come up with an extremely attractive product here. If you're only looking for a midrange reference speaker, the choice between the Angry Box and the Avantone Mixcube may be a fairly nuanced decision. Do you prefer the smaller, lighter, and (ironically) less aggressive-sounding Angry Box, or the cheaper and more midrange-focused Mixcube? In terms of sheer midrange referencing performance, you'd be a winner whichever way you jump. However, entry-level mix engineers in particular may well be swayed in the Angry Box's favour by its 'bonus' extended-bandwidth mode, delivering as it does a good chunk of the mixing power of similarly priced traditional nearfield monitors into the bargain.

$ $999 per pair.
W angrybox.co.uk

Firesonic Firespacer

Spectral Ducking Plug-in

Most of you will be familiar with the idea of ducking, a process by which the level of one signal is controlled by that of another. For example, this is how the background music is dipped automatically in level when a DJ speaks. But ever since the inception of spectral processing, it occurred to me that a ducker that affects the levels of only conflicting frequencies should be possible. This concept has gained traction in recent years and it's what lies behind Firesonic's Firespacer, which is marketed by United Plugins and protected by a user key that allows installation on all the user's machines. All the common plug-in formats for Mac and Windows are supported, including AAX, and a trial mode is available so you can try before you buy.

Firespacer can be used in any DAW that supports side-chaining. The plug-in is inserted on the channel or bus that needs ducking and the control signal is received through the external side-chain input. There are four modes of operation: two Spectral, which introduce some noticeable latency; plus two latency-free modes that use minimum-phase filters. The Spectral modes are the most accurate, as the minimum-phase filtering modes may carve out some frequencies that don't need trimming. As to which filter or spectral mode you choose, that's down to listening to see what works best for the material being processed. A Precision control, only active for the Spectral modes, sets how tightly the spectrum of the source sound is followed, so at lower settings the effect is rather like smoothing the response curve of the spectral filters.

The resizeable GUI offers mono, L/R or Mid-Sides operation, with metering and separate gain controls for the input, output and Sides levels. The Sides also has a solo listen button. Two controls set the minimum and maximum processing frequencies so that the user can opt to exclude lows and highs from the processing, and there's also a variable release time that's best adjusted by ear to suit the material. The large Amount knob in the centre sets the depth of the ducking and there's a frequency display running across the top of the window that shows the main and side-chain signals as well as the signal level after ducking. The power button functions as bypass.

A typical application might be to duck the spectrum of pad or distorted rhythm guitar sounds in the presence of vocals, where there are likely to be significant overlapping frequencies, especially in the midrange. Similarly, you might use it to keep a lead instrument clear of whatever murk is going on in the background. There's also potential for untangling conflicts between the bass and drums. Doing these jobs with a conventional ducker often results in very obvious gain changes to the track being processed. By pulling down only the overlapping frequencies, spectral ducking sounds far more transparent and the higher latency of the spectral modes isn't an issue when mixing. Where zero added latency is important, the minimum-phase filter modes still sound more benign than straightforward ducking, though some tonal modulation of the target track may be evident at more extreme settings with some material.

In summary, Firespacer is a simple-to-use and very effective mix tool that helps bring more clarity to key parts by reducing conflict due to overlapping frequencies in other parts of the mix. *Paul White*

$ $76.
W https://unitedplugins.com

Boss SDE-3

Digital Delay Pedal

Designed to recreate the 'vintage digital' sound of the Roland SDE-3000 rack unit, the Boss SDE-3 hosts dual delays that can be used in various permutations of mono or stereo in and out, although there are some limitations if running it as a stereo-in, mono-out pedal — the inputs lock with their respective outputs when a stereo input is received. An Offset knob adjusts the relative time difference between the two delays, which run in parallel. Occupying a standard Boss compact pedal format, the case has dual jacks for the inputs and for the outputs, plus another jack for external control. There's also a MIDI mini-jack on the rear for sync'ing the delay time to MIDI Clock, as well as a Carryover on/off switch that allows for echoes to fade out naturally when the pedal is bypassed.

The pedal LED flashes at the delay rate, and the delay time can be tapped in by holding down the pedal to enter tap mode or by using an external tap switch, connected via the CTL/EXP jack input. This can accommodate two switches for tap tempo (CTL1) and hold (CTL2) or an expression pedal. Note that the SDE-3's Time knob and tap-tempo functions are bypassed when MIDI Clock is being received, though the Offset time setting is not affected by either tap-tempo or MIDI Clock. If an expression pedal is plugged into the CTL/EXP jack controls, it can control the delay time, level and feedback simultaneously, which is great for the creation of special effects. (The CTL/EXP jack function is set using a power-up routine.) Other user settings include the option for standard stereo delay or alternate panning delays, plus three output options: stereo, wet from one output and delay from the other, or direct mute for when only the delays are needed.

To make all this fit into a standard Boss pedal, three dual-concentric control are used, giving access to the wet/dry mix, delay feedback and delay time. Turning the level control clockwise raises the delay repeat volume relative to the unity dry signal. Delays of up to 800ms are available in stereo (twice that in mono), with the minimum delay being essentially 0s. The knobs' outer rings address modulation depth and rate, and high-cut filtering. The aforementioned Offset control can apply up to 100ms difference between the left and right delay lines. Turning Offset clockwise from centre offers a choice of eighth-note and dotted-eighth-note delays as well as in-between settings.

Its impressive dual-delay feature aside, the SDE-3 delivers a very musical delay sound that can be given an even more vintage flavour by dialling in a little pitch modulation and rolling off a little high end with the tone knob. Offset allows for all those tricky delays that included dotted rhythms ('Run Like Hell', anyone?) while the stereo output options are a real bonus for studio work, or live work using two amps. Boss have always made really great delay pedals but there's something about this one that just draws you in and makes you want to keep on playing. *Paul White*

$ $219.99.
W www.boss.info

Klevgrand Revolv

Reverb Plug-in

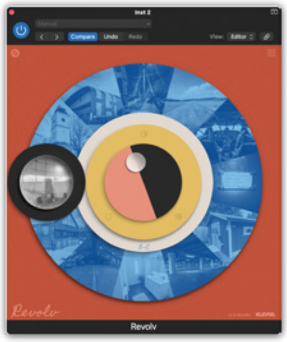

Klevgrand's Revolv is an IR-based reverb plug-in that supports all the usual macOS/Windows formats including AAX, and offers a wide range of spaces with very different characters. It's all wrapped in an interface that allows for quick auditioning of the different reverbs, and the ability to dive in deeper and adjust parameters such as mic type, mic distance and EQ if required. The IRs were recorded by sound designer Oscar Björk at various locations across Sweden, in a range of hand-picked locations.

The GUI is certainly different from the norm. It's set out as a wheel with 12 segments, each hosting an image of a reverberant space. A magnifying glass can be placed over any image to load that impulse, while a slowly moving pan shot of the selected space is shown inside the magnifying glass. A gimmick maybe, but it's a very attractive one! The large knob in the centre sets the wet/dry mix, and for many applications that may be all you need. As you move the mouse pointer over an image, the name of the space is displayed above and to the right of the wheel.

Click on the three faint horizontal lines at the top right of the GUI and you'll find six different categories: Ruins and Nature, Historic Halls, Old and Etherial, Modern Gatherings, Wooden Rooms, and Stages. Choose one and the wheel will present you with 12 spaces from that category. Here you'll find spaces as diverse as kitchens, caves, cathedrals, concert halls, parking garages, crypts, towers and stone ruins, plus some natural outdoor spaces. There's even a plane cockpit. There's also an 11-category preset browser, with many presets offering alternative mic setups for each space.

If you want to edit at a deeper level, double-clicking the magnifying glass brings up a window in which you can adjust the pre-delay and decay times, the wet/dry mix, gain, EQ, width and mic type and, where applicable, mic placement. You'll also find a brief description of the space here.

Clearly, Revolv might not be the first reverb of choice if you simply need a plate-style decay with no obvious character, but if you want to put your voice or instrument into a specific real space with a very strong sense of character, then Revolv provides lots of useful options. While it doesn't work much differently from other IR-based reverbs, the selected locations each have a unique character and the unorthodox GUI makes it quick and very easy to find a space that works. Interestingly, adding the character of some of the outdoor spaces can seem to make the sound drier than before adding the reverb, while the stone buildings add a real sense of the character of the space. If you want a reverb that tells a story, then Revolv is both approachable and affordable. *Paul White*

$ $69.99.
W https://klevgrand.com

Ghost Note Audio Conductor MkII

MIDI Controller

Within the ever-growing world of orchestral sample libraries, it's become the norm to provide a large degree of control over timbre when playing, in an attempt to garner an ever more realistic instrumental performance. At the most basic level this involves using the modulation wheel to control timbral dynamics, but extended control is often available, provided you have the MIDI faders to achieve it.

It came to pass that the Holy Grail of fader combinations was a three-fader combo, which has become something of a standard, while also being a realistically manageable scenario for most ambidextrous keyboard players. Ghost Note Audio's new Conductor MkII is exactly that: a solid three-fader MIDI box, specifically designed for the purpose of driving orchestral libraries.

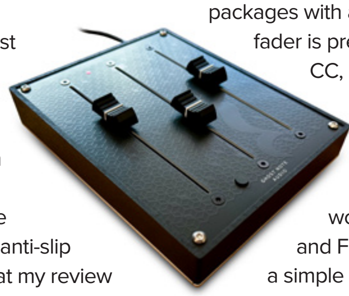

The controller itself is just the right size at 119 x 152mm. The top and bottom of the unit are assembled from aluminium panels, with a 3D-printed surrounding enclosure. The base also sports a silicone anti-slip layer, although I noticed that my review model appeared to be slightly convex in shape, meaning that it had the tendency to occasionally twist when in use. It looks pretty striking on the desktop and will nestle nicely on top of many 88-note MIDI controllers. One minor aesthetic point is that the whole unit is black, so it might have been nice to use black rivets instead of the four silver rivets that populate each corner.

On the rear is a single USB-C socket, which provides connection for both power and data. In testing I connected the Conductor MkII to my computer directly, as well as through a powered USB hub. It behaved perfectly in both instances. There's a supplied USB-A to C cable, so depending on your host machine, you may need an adaptor or a different cable.

The Conductor MkII is MIDI class compliant (macOS, Windows and Linux), so within seconds of connection I was driving Kontakt and other orchestral host packages with a minimum of fuss. Each fader is pre-configured to a MIDI CC, but can be altered easily with the browser-based editor. This requires a MIDI-compatible browser; Safari doesn't work but Chrome, Edge and Firefox do. Once engaged, a simple three-column/fader editor facilitates a number of useful settings for personalisation, before uploading instantly to the unit. Channel, values and thresholds are all assignable, along with other modes of operation, such as NRPN and pitch-bend.

As different libraries often require different MIDI CC settings, GNA have made three fader banks available, which are accessed from the momentary switch on the front panel. The power light switches colour from red to green to blue. These indicate the three banks, mirrored on the web-based editor page. Any setting alterations stay with the unit, even when powered off or disconnected. It's a very helpful inclusion and easily edited, although the black momentary switch does tend to disappear against the black fascia.

When it comes to a unit such as this, everyone will have a preference. My personal take is that I like a small degree of resistance in a fader, with a decent amount of travel, and that's exactly what we have here. Driving any number of the usual orchestral suspects is very easy, and exceptionally comfortable too. In the nicest possible way, I largely forgot about the interface, and got on with my musical task at hand, and that's really what you want with a device of this kind. *Dave Gale*

$ £149 (about $195).
W www.ghostnoteaudio.uk

Austrian Audio Hi-X20

Closed-back Headphones

If the contents of this issue are anything to go by, Autumn 2024 is delivering a rich crop of new headphones. Austrian Audio's contribution to this bountiful harvest is the Hi-X20, a closed-back, moving-coil model that sits in between the affordable Hi-X15 and the more upmarket pairs in their range.

Austrian Audio draw on the expertise of many people who worked at the old AKG headquarters in Vienna and, inasmuch as their existing models have a family sound, it's one that has something in common with AKG's, involving a readily apparent presence boost in the upper midrange and high frequencies. That tendency is particularly pronounced in the Hi-X15s, which are some of the brightest cans I've used: ideal for ensuring the drummer's click track cuts through, perhaps less so for delicate mixing decisions. The Hi-X20s are designed to present a different voicing.

In physical terms, the two models are outwardly identical, except that the 20s are finished throughout in a sober black, without the 15s' silver and red elements. Consequently, the new cans are lightweight and comfortable, and boast longevity-oriented features such as removable earpads and a detachable cable. Thankfully, they're supplied with a longer cable than the Hi-X15, along with a black drawstring bag for (somewhat) safe keeping. Like all of Austrian Audio's headphones, the Hi-X20s are low-impedance designs that use the company's proprietary "high excursion" driver, combining high sensitivity with low distortion.

The Hi-X20s do indeed have a noticeably different voicing from the Hi-X15s, but they're still recognisably part of the Austrian Audio family. In other words, they are on the bright side of neutral, but not nearly as much so as their junior siblings. In fact, they're even a bit more restrained at the top than more expensive models like the Hi-X60 — or perhaps it would be more accurate to say that the brightness of the Hi-X20 seems focused in a narrower frequency range, somewhere around 4-5 kHz. On first encounter, this might still challenge your judgement around sibilance or cymbal splashiness, but I think it would be easy enough to compensate for either with experience or EQ, and in other respects, the Hi-X20s fulfil all the requirements of a good pair of general-purpose studio headphones. They're less fatiguing, more listenable and more versatile than the Hi-X15, and given the very modest difference in price, the choice seems like a no-brainer. *Sam Inglis*

$ $149.
W https://austrian.audio

MNTRA Borealis

Modulated Reverb plug-in

It may look like yet another reverb plug-in, but Borealis ploughs its own furrow by including extensive modulation facilities — there are three modulation sources plus an envelope follower that tracks the incoming audio. It also includes a master section, with a parametric EQ and dual-mode limiter that offers modern or vintage characters. All the mainstream plug-in formats are supported across Windows and macOS platforms, including AAX.

The somewhat unconventional-looking interface, which looks like a sphere hovering over a bunch of rods, employs three-axis macro controls for parameter manipulation, allowing for the creation of unusual dynamic effects. Based around 14 reverb algorithms, Borealis employs both convolution and algorithmic methodology, from classic spaces and springs to granular, shimmer, distorted, endless and freeze with some also including echo-like effects. X, Y and Z parameters are reflected in the GUI animations.

Factory presets are arranged by character for easy searching and, once loaded, you can change the algorithms using the selector at the top left of the main Perform View screen. Clicking the logo in the top-right corner gets you back to the Perform View from any other screen. The three control axes are represented by a circular dot, a triangle and an equals sign, and these may be dragged around the GUI to effect changes via the performance macros. A control panel can be opened on the left of the GUI showing the wet/dry mix and other parameters relevant to the current algorithm. A further narrow window can be opened that shows the control mapping, which can be Modulator, Axis Position or Manual. Curves can be adjusted or reshaped by adding and dragging points, and each modulator thumbnail can be expanded to full size, where a set of preset curves is available along with a randomiser.

A '+' icon tab at the bottom of the GUI opens additional rotary controls for adjusting the basic reverb parameters, while the next tab along opens the modulation page, with settings for LFO shape, rate, intensity and so on. Shapes include custom options and sequencer-like steps. Modulators can even modulate other modulators! The Paths section is where the modulation sources are assigned to destinations, and it's a similar story for the envelope follower. Finally, the rightmost tab opens the master EQ and limiter.

By way of performance, all this would amount to nothing if the reverbs themselves didn't sound good — but they do! From realistic spaces to huge ambient washes, there's something for everyone, with the modulation capabilities producing reverbs that pulse and shimmer in a way that reacts to the input. I suspect that creators of ambient, lo-fi or chillout music are the most likely to embrace the extremes but there are plenty of treatments that would work in mainstream pop too. There's a free version of Borealis limited to a single algorithm to help you get a feel for how Borealis works, and the full version is not overly expensive. It's refreshing to see a company like MNTRA adding useful new twists to reverb, one of the oldest effects in the recordist's arsenal. *Paul White*

$ $99 (discounted to $59 when going to press).
W www.mntra.io


ADAM AUDIO
All New
H200
Studio reference headphones from ADAM Audio.
Never Leave the Sweetspot.
www.adam-audio.com

Source-destination Editing

With ReaClassical

We explore a clever set of donationware scripts that transform the popular Reaper DAW software into a source-destination editing powerhouse.

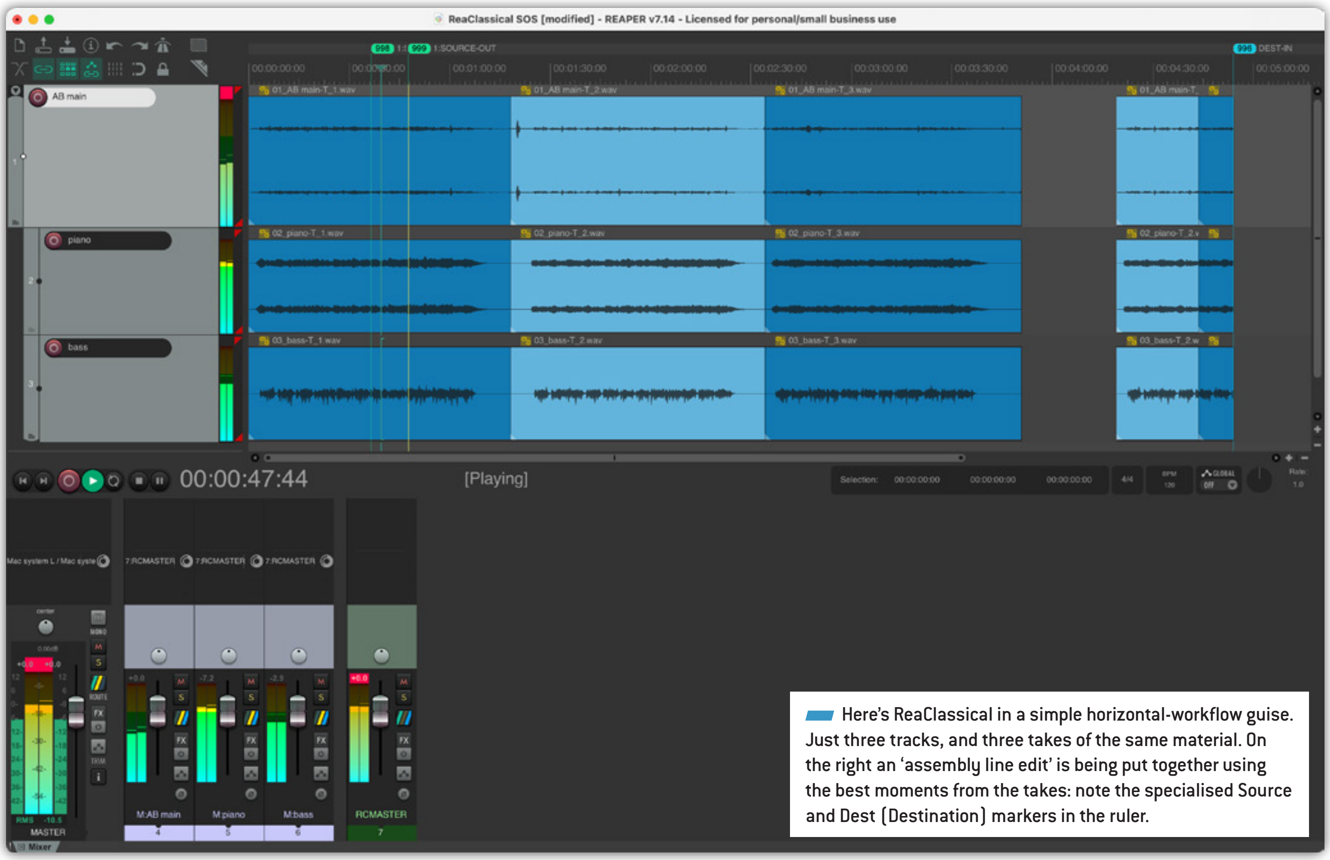

Here's ReaClassical in a simple horizontal-workflow guise. Just three tracks, and three takes of the same material. On the right an 'assembly line edit' is being put together using the best moments from the takes: note the specialised Source and Dest (Destination) markers in the ruler.

ROBIN BIGWOOD

Reaper is a DAW that's famous for its unparalleled scope for customisation, and there's a very active community of developers who write scripts to take advantage of that. A select few, though, take the script-writing further than others, to turn this DAW into something quite different. One such project is ReaClassical — a set of open-source scripts that transforms Reaper into an amazingly capable source-destination audio editor that has specialist features to support your sessions from recording right through to mixing, mastering and album authoring. I suppose I could conceivably have written this article as a review, but it would have taken way too many pages to do it justice. Instead, I thought it might be a good idea to demonstrate some of the possibilities by taking you through a few useful examples of how and why you might want to use it.

Source-destination Editing

So, what's 'source-destination editing', exactly? Still uncommon in audio circles, it's a standard approach in the video world that comes into its own when dealing with mostly acoustic musical styles, where performances are played and recorded naturally rather than performed to a click track, and overdubbing isn't typically used. Which is to say classical music, many types of folk/world music, and jazz. For such styles it's normal, in a recording session, for the recordist/producer to gather multiple takes (the 'source') of the same pieces, with a view to piecing together a 'best bits' edit (the 'destination') in post-production, cutting between takes to cover musical glitches, possibly assisted by notes made by a producer during the sessions.

It isn't impossible to do this using Reaper in its standard form, of course, or indeed most DAW software. But for this job, a source-destination (S-D) system is faster and more flexible. It supports gathering material on the recording session, with really robust track and item grouping. Editing then happens on slick and fast ripple-edit principles, so you can build edits without laborious cutting and pasting. Crossfading also plays a big part, knitting together edited material with quite different facilities from those offered by the vast majority of DAWs. In ReaClassical, keyboard shortcuts drive every aspect of this: the user is rarely required to use any modifier-key combinations, and the shortcuts are grouped tightly and logically on your computer keyboard — all of which is intended to facilitate ease and speed of use.

This isn't the only S-D-capable DAW, but virtually all of the others that are currently available are Windows-only affairs, and they're often quite expensive too: Sequoia, Pyramix and SADiE, for example. ReaClassical will run in macOS and various flavours of Linux, as well as Windows, and it's donationware. Given that Reaper is also inexpensive, this makes it a much more accessible system than most.

On The Session

To use ReaClassical on a recording session, after booting up the portable Reaper install, pressing F7 is a good first step. A Horizontal Workflow dialogue box (we'll return to that terminology) asks how many tracks are required, and the answer should equate to how many stereo or mono mic sources you're intending to use. Hit OK, and ReaClassical creates enough track lanes, with track (and item) linking already set up. Next steps are to name the tracks (double-click fields in the Track Control Panel in the usual way) and configure track inputs (right-click in the nearby level meter, or use the Routing Matrix).

You can set up tracks as you please, but it's a good practice to choose your main stereo pair of mics for the parent track, track 1: this ensures that a healthy, representative waveform is visible even when the track group is closed. When you've done this, press F7 again to 'sync' tracks and mixer channel names, and build some behind-the-scenes routings. Finally, save your new project.

Download ReaClassical

To obtain ReaClassical, visit https://reaclassical.org, where you'll find links to installers (Terminal commands for macOS/Linux, and downloadable EXE files for Windows) that do everything for you, creating a portable Reaper install with various third-party extensions that can (and should) live alongside your standard installation. A friendly PDF user guide also documents the many features I've not covered in this workshop. It's donationware (rather than freeware) and I'd very much encourage you to support its ongoing development.

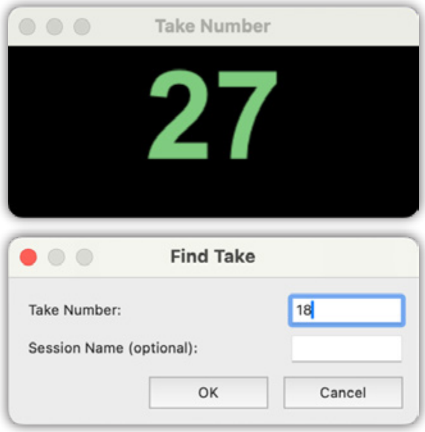

The floating Take Number window is a wonderfully clear reference during a recording session, and Find Takes is a real time-saver for quickly locating Source material, especially if you've recorded hours of it, and hundreds of takes.

To start recording, F9 is your friend. A first press record-arms all your tracks. Press again to start recording, and once more to stop. You may notice immediately the way ReaClassical names your recorded media items in the track lanes: along with the track's number and name, they're given a 'T' suffix that gives them a unique 'take' identification. This numbering is the key to tallying up recorded material with a producer's notes made on a session about what went well and what might need attention. ReaClassical keeps incrementing this number as you make more recording passes, and it keeps this harmonised between tracks even if you add new tracks (using the Shift+T shortcut) some way into a session.

Two really specialised features build on this take functionality, and they're worth their weight in gold. The first is a current take number display: Ctrl+Return or Ctrl+Enter (Command on macOS, as always) opens it. This is a simple floating window with a nice clear number: when not recording it shows in green the number of the next take, but during recording the number of the current take, in red. It brings so much clarity to a session, especially hours in! The second feature is closely related. When not recording, pressing Return or Enter opens a Find Take dialogue box. Type in a number (and hit Enter or Return), and ReaClassical relocates the play cursor to the start of that take, reading the metadata in the audio items, with not a marker in sight.

There's one more thing to mention here: playback, which is often required in a session to check balance, and for foldbacks for musicians. You can click the playback position in the time ruler or use Find Takes, and both work really well, but there's another option. Hit the keyboard shortcut A (for 'audition'), and playback will start from the mouse pointer position, when it's placed over a track lane. Point at an item in the parent track and it plays the whole mix. Point at an item in a child track, and it effectively solos that track. Its fundamental strength is its speed: you can find musical sections simply by repeatedly pressing A as you move your mouse over one or many items. Audition is used elsewhere too, as we'll see.

In The Edit

So far, we've only dealt with 'source' material: audio captured in the session. But the key to a source-destination editing system is taking that and building a separate edit from it. So let's look at that now.

For simple, shorter jobs, continuing with a 'horizontal' setup can work fine. Thinking about those first principles of source-destination editing again, what we can do now is grab the best bits from our various session/source takes, and start splicing them together to form a better composite whole. A typical first step is an 'assembly line edit', working through a piece from start to finish. Here's how it might go.

Using Find Takes and Audition, find a take that represents the best start to the piece of music you've recorded. It may go awry some way in, and that's OK — it's what this style of editing is all about, in fact. We want to mark just the 'good bit', and for this we use two specialised ReaClassical marker types: Source In and Source Out, placed with keyboard shortcuts 3 and 4, respectively. These can be placed during playback or when playback has stopped, and the markers can be dragged manually if necessary; the locations need to be roughly right, musically speaking, but needn't be laser accurate.

Next, click in the time ruler above an empty spot in the timeline, probably after the session takes run out. Press 1. This places a Dest In marker (short for 'destination in'), which marks the spot at which the edit will start to be built. (You should now have three points marked.) Then press F3, the shortcut for 'Assembly Line edit', to copy the marked source region to the destination marker location — without the several extra steps involved in a traditional select, copy, paste process. What's more, the Dest In marker »

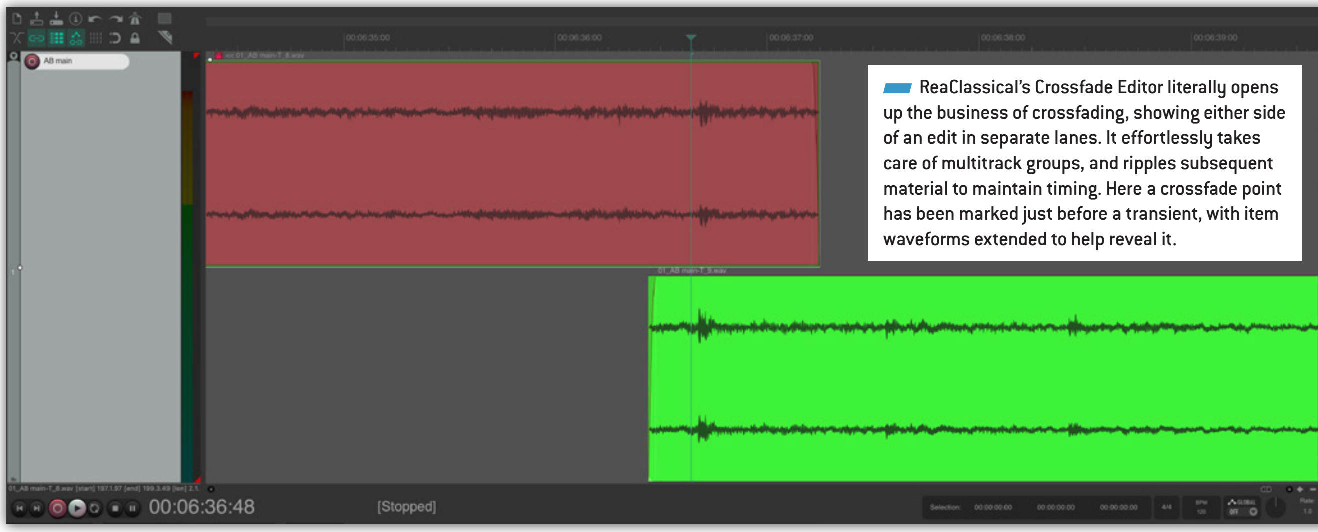

ReaClassical's Crossfade Editor literally opens up the business of crossfading, showing either side of an edit in separate lanes. It effortlessly takes care of multitrack groups, and ripples subsequent material to maintain timing. Here a crossfade point has been marked just before a transient, with item waveforms extended to help reveal it.

automatically moves along, so it's perfectly placed to receive the next source material.

It's now a case of working through the music, choosing good sections within takes and marking them with 3 and 4, each time pressing F3 to add to the rough edit. You might have to do this a few times or a few dozen, depending on the length of the music and how good the performance was, and potentially swap backwards and forwards between source takes. But in the end you should have something that represents an entire piece or movement. Put the mouse pointer over the beginning of this new edit and press A to hear it. It should be complete, but the transitions from take to take are likely to be clunky, despite being equipped by default with short crossfades. Which is why ReaClassical has its own dedicated Crossfade Editor...

Fading Up

Any good DAW, including Reaper in its standard form, can make crossfades between abutting or overlapping sections of audio. But a problem is that you're forced to work in a single track lane, with the waveforms of the audio involved overlapping and thus potentially being obscured. While that's fine for a single fade on one track, it's not very useful for multitrack work that positively relies on crossfading, and might include hundreds them, all of which need to be perfect to the point of complete inaudibility! ReaClassical's fade editor works differently, with vastly better clarity and flexibility, courtesy of a dedicated two-lane fade editor. It'll show the end of one audio region and the beginning of the next separately (ie. either side of the potential crossfade). Next up, then, I'll take you through the basics of this, which may seem unfamiliar to begin with but will, I promise, soon become second nature.

In your clunky assembly edit, scroll back to the start and select its first item. Press F to open the Crossfade Editor. The normal track group display is now replaced with two lanes: the red item is the one you're going to fade from; and the green item is the one you'll fade to. What's needed now is information: you need to hear the edit as it stands, to gauge how near or far it is from being correct. It might be close or, very likely, there may be a beat or two missing or duplicated.

The best way to check this is, once again, to use the A (audition) keystroke. Put the mouse pointer over empty space in the lane above or below the red item (I say that because the item could itself appear in either lane), and press A ReaClassical plays across the edit to a mirror point on the other side. If you'd moused over the red item directly you'd get playback from your pointer up to the edit, and if over the green item, from the edit to your pointer.

Assuming it's not already perfect (possible, but unlikely!), point again into an empty lane area and press Z. This extends the waveforms of both items so that they overlap, the idea being that this'll help you spot the matching shapes of matching musical material. Transient peaks and spikes are particularly useful for this. Assuming you can see some correlation, as in the screenshot, then the next steps are easy.

The Other Crossfade Editor

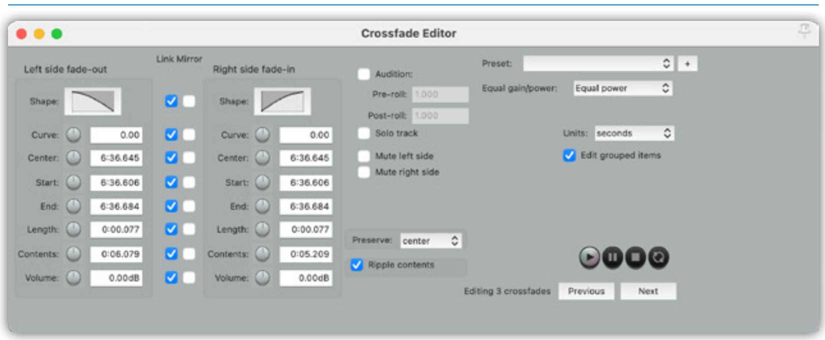

ReaClassical's Crossfade Editor is a game-changer in terms of speed and flexibility. But its scope broadens even further when used in tandem with Reaper's similarly named built-in Crossfade Editor window. That offers additional parameters for fade length (slightly longer fades are often a solution for tricky transitions) and fade 'centre' position (letting you explore alternative out/in positions without having to repeatedly do the Z/click/drag/X combo, as quick as that is). For really challenging crossfades, it'll also let you dial in different fade shapes. It's all up for grabs, but there's one caveat: never use its Previous and Next buttons! Instead use only ReaClassical's dedicated Q and W equivalents.

» Click in the red item where you want the crossfade to occur; just before a transient is always a good candidate. A cursor is placed as a clear visual guide. Click and drag the green item left or right to line up its matching transient vertically and, finally, press X to re-trim the items and write the new crossfade — which you'll immediately need to audition again to see how successful it was.

No good? Listen more, press Z again, try a different crossfade point or timing, and then X to make another new version. Like it? Good. In this case, we want to move along to the next edit whose crossfade needs polishing. Having to locate that manually would be a drag but, without leaving the Crossfade Editor, pressing W takes you straight to the next edit — then you can work on that one just as you did the last. In fact, W is actually part of a pair of shortcuts, together with Q, that go left and right between edits (and items generally), so you can fix all fades in an assembly edit in no time. And when that's done, pressing F once again leaves the Crossfade Editor.

Second Sight

In an ideal world, our assembly edit and fade work (what might be termed a 'first edit') would result in a perfect

Vertical Limit

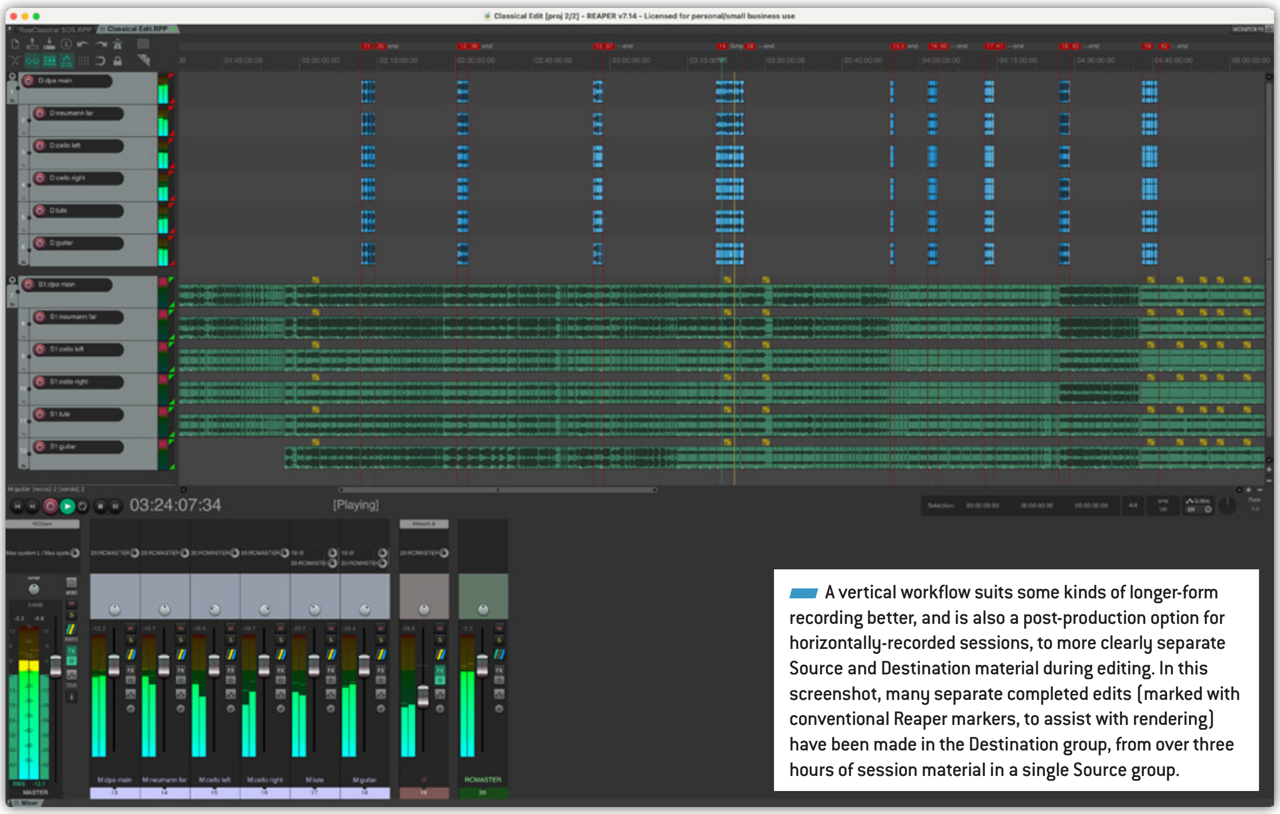

A vertical workflow suits some kinds of longer-form recording better, and is also a post-production option for horizontally-recorded sessions, to more clearly separate Source and Destination material during editing. In this screenshot, many separate completed edits (marked with conventional Reaper markers, to assist with rendering) have been made in the Destination group, from over three hours of session material in a single Source group.

Take a cross-section of source-destination DAWs and you can find different paradigms for how source and destination material is managed and displayed. In ReaClassical the simplest is the 'horizontal' workflow that you set up with F7, and which I've explored in the main text. That works great for session recording and for relatively simple editing jobs, like single pieces or tidying up live concert recordings. But there's an alternative (you guessed it!) 'vertical' workflow...

The vertical workflow is accessed by pressing F8 after creating a new project. You're asked how many tracks you want, as before, but then you'll end up with a single dedicated Destination track group and six Source groups by default (though you can add as many more of these as you need). As with the horizontal setup, after adding track names a further press of F8 syncs them across the multiple groups. All groups use the same unified mixer channels, so it's dead easy to manage.

"But what is vertical for?" you may ask. Well, it can make for a visually neater and more compact working environment when recording well-drilled ensembles that aren't going to make all that many takes. An example would be an orchestra making only a handful of takes of a long movement. Each take would end up in a separate source group, such that matching musical material between them can be closely vertically aligned. This makes auditioning between them phenomenally quick, with little or no scrolling required. On the session, the F9 keystroke is still used to arm tracks and initiate recording a take. But in this workflow, when F9 stops recording ReaClassical returns the play cursor to the previous record start location and arms the next Source group down (automatically creating an additional one if required). On the next press, you'll record into that.

A vertical workflow is also good for editing longer, perhaps multi-day recording sessions that started out in horizontal and may have run to many hours in length. It allows a separate Destination track group to run concurrently, which can be used to build edits close to the Source takes, in the same project. For this you'd open your horizontal-style project (with its single track group) as normal, but then a press of F8 'converts' it to the new workflow. As it happens, your session material stays put in the uppermost D (Destination) group to begin with, but it's a simple job to transfer it into a Source group. Press the accent (`)key (the one to the left of Z on my UK keyboard) to zoom out fully. Right-click-drag over all recorded items in the D track group, and drag them down into S1. Job done: you've now an empty, independent Destination group that'll receive any kind of S-D edit at any DEST markers you've placed. If you don't need the remaining five source tracks, you can either ignore them, or select them, right-click and choose Remove Tracks.

A four-point edit is where you specify in and out locations in both the Source material and the Destination 'best bits' composite — it lets you carve into an existing edit to make later improvements, maintaining all other edit points and crossfades around it. Here a four-point edit is about to replace an existing edited section that includes a crossfade.

musical product. In reality, it'll probably need more polish: anything from repairs of individual notes to replacing longer sections. This is the moment we utilise true 'four-point' S-D editing. Compared with our assembly edit (which was technically a 'three-point' edit), there's one difference: we mark both IN and OUT points for the source material and the destination within the existing edit.

Imagine an edit you've made that needs just one tweak: a new, alternative section to be dropped in, in one place. Use the audition action (A) over your edit to find out where the section in question starts and finishes. Mark it with Dest In and Out markers (keystrokes 1 and 2) and then go back to your source material and audition that to find a better alternative. Mark that with Source In and Out (3 and 4). Then, pressing 5 executes a 'make S-D edit': it lifts the source material, drops it to the destination, and ripples everything that follows along the timeline, including any fades you've previously worked on. ReaClassical also helpfully scrolls to the location and places the play cursor at the end. Remember, though, that this new region will likely need fade work to transition in and out of it, so select the item before it, enter the Fade Editor, fix the 'in' crossfade, and press W to move to the 'out'. Fix that, and you're free to leave the editor.

Next Steps?

As this is an introduction to ReaClassical, I've kept things fairly high level, and only run through some of its most potent core editing functionality. But while there should already be plenty here to get your teeth into, its abilities extend further — much further, in fact. You can perform S-D editing with time-stretching, for example, and there's flexible signal and effects routing in the mixer, support for reference tracks and 'room tone' (recorded silence, used to pad transitions between tracks), snapshot-style automation, some great mastering-style plug-ins, and even CD authoring and DDP file creation. The developer (chmaha) also tells me that there's much more on the horizon. So there's lots more to explore here if I've piqued your interest!

We breathe some new life into the Mai Tai synth, with some creative patches for you to explore.

Alien Squirrelscape: modulating an electric-piano-style patch beyond recognition using the Mod Matrix.

ROBIN VINCENT

It is time to revisit Mai Tai, the often-overlooked Studio One synthesizer. While it may be named after a cocktail, this 32-voice polyphonic virtual analogue synth has little to do with propping up a tiki bar on a far-flung tropical beach and more to do with sumptuous pads, fat basses and penetrating leads. It's a classic subtractive synth and has a huge amount of character and versatility. It has a traditional architecture, plenty of modulation opportunities and a bunch of effects.

So, for this workshop, we're going to initialise a patch and set to work on putting together some patches that lean into the different strengths of Mai Tai in the hope that it will encourage further exploration of your own.

Mountain Webbing: a simple yet effective string pad patch.

Alien Squirrelscape

One of the best things about creating your own patches is that you can call them what the heck you like, and no one can do anything about it. So, let's kick off with the wild chaos of Alien Squirrelscape. This is not a workshop for the squeamish.

The purpose of this little beastie is to dive head-first into the Mod Matrix and push things around in an alarming manner. But we should start with an initialised patch. So, drag in an instance of Mai Tai to create a nice new track, click on the patch list where it currently says 'default', and select '+ init'. You will be faced with a simple singular sawtooth sound. Let's get to work.

Osc 1: Set to a sine wave and enable RP (random phase).
Osc 2: Enable, set to sine wave and dial the octave down to 16'.
Noise: Turn the level down, although you can always add this back in to taste.
Character: Enable and set to Voxii for a bit of a strangled vocal vibe.
Filter: Disable Soft, put the Drive up to 9 o'clock, Res to just past noon, Vel to 2 o'clock and Key all the way around. Set Cutoff to 1 o'clock.
LFO 1: Enable, set to triangle, enable Key and Free and set the rate to about 10 o'clock.
LFO 2: Enable, leave on sine wave, disable all three buttons and set the rate to about midnight
Amp Env: Leave Attack at zero, put Decay to 11 o'clock, Sustain to 2 o'clock and Release to 3 o'clock.
Env 2: Bring Attack up to 2 o'clock, Decay to 4 o'clock, Sustain at zero and Release to 1 o'clock.
Env 3: Attack to noon, Decay to 1 o'clock, Sustain and Release to nothing.
FX A: Enable Delay and increase the time to 1 bar. Enable Reverb and increase the Size to noon and the Mix to about 10 o'clock. (If you can't see the Modulation and Effects section at the bottom, click on the MOD/FX button at the bottom left.)

Now, give it a play. You will hopefully hear an unexpectedly beautiful electric piano with a ghostlike, distant and ambient quality. We are going to mess that up entirely with the Mod Matrix.

Click on Mod A and add the following modulations to the first six slots. The upper slots (above the modulation amount sliders) are the modulation sources, and the lower ones are the destinations.

Slot 1: Source LFO 2, Destination Osc 1 Pitch. Push Amount to about 75 percent.
Slot 2: Source LFO 1, Destination Osc 2 Pitch. Amount 100.
Slot 3: Source Env 3, Destination LFO 1 Frequency. Amount 100.
Slot 4: Source Env 2, Destination LFO 2 Frequency. Amount 100.

This patch, named Abrupt, is great for dramatic build-ups.

Persistent Eyebrow: a patch that combines a guitar-like plucking sound with ethereal washes.

Slot 5: Source Env Main, Destination Filter Cutoff. Amount 25.
Slot 6: Source LFO 1, Destination Character Sound. Amount 75.

You can probably see that the basic idea is to modulate the pitch with the LFOs and use the envelopes to modulate the speed of the pitch modulation. It creates a wonderfully vibrant and strangely organic explosion of animated, long-release events. Groovy.

Mountain Webbing

This is probably the finest string pad I've ever come up with; it's just lovely. The patch is actually quite simple and is based upon slightly detuned sawtooth waves, a cinematic envelope and a delayed bit of modulation. Start with the init patch.

Osc 1: Set to a sawtooth wave and enable RP. Turn Spread up to 3 o'clock, Sub up to 9 o'clock and push the Fine-tuning knob left just a little bit to -9.00.
Osc 2: Enable, set to sawtooth wave and dial the octave down to 16'. Enable RP and Sync. Push Level up to about 1 o'clock and the Fine tuning to the right to 13.00.
Noise: Enable and set Level to 9 o'clock (adjust to taste) and Colour all the way up.
Character: Enable and set to GrandClass with Amount all the way up and Sound set to L14.
Filter: Disable Soft, put the Drive all the way up, Punch to 8 o'clock, Res to zero, Vel to 1 o'clock and Key to 2 o'clock. Set Cutoff to 2 o'clock.
LFO 1: Enable, set to sine wave, disable all three buttons and set the rate to about 10 o'clock. Set the Delay to a little past noon.
LFO 2: Leave disabled.
Amp Env: Put all four knobs to just past noon.
Env 2: Put Decay to maximum and everything else to zero.
Env 3: Unchanged.
FX A: Enable Chorus, Delay and Reverb. The default settings are pretty good, but maybe the delay mix should be reduced slightly and the reverb mix could be increased to almost 3 o'clock.
Mod A: The modulation is much simpler here, with Env 2 controlling Filter Cutoff at about 40 percent, and LFO 1 set to Osc 2 Pan at 100.

Have a go, and hold some chords. I hope you find it as delicious as I do.

Abrupt

Here's a simple but dramatic patch: a phenomenal build-up followed by an abrupt drop. Simple things can be awesome...

Osc 1: Set to a square wave, put PWM on 11 o'clock, Spread to 10 o'clock and Sub all the way round.
Osc 2: Enable, set to square wave, put PWM to noon and nudge the Fine tuning to the right a tiny bit.
Noise: Leave disabled.
Character: Enable and set to CharacterSaw, with both Amount and Sound set to noon.
Filter: Turn Soft on, put the Cutoff to zero and the Res to just past 1 o'clock.
LFO 1: Enable, set to sine wave, enable Key and Free and set the rate to about 11 o'clock.
LFO2: Leave disabled.
Amp Env: Set Attack and Release to zero, Decay to just before 9 o'clock and Sustain all the way up.
Env 2: Set Attack to 3 o'clock, Decay up a tiny bit and everything else to zero.
Env 3: Unchanged.
FX A: None
Mod A: In slot 1, set Env 2 to control Filter Cutoff at about 100 percent; in slot 2, set LFO 1 to control Osc 1 Pulse width at 75.

Persistent Eyebrow

To highlight Mai Tai's versatility, here's a sound that has electric guitar-like elements combined with a subtle ringing feedback and washing ephemeral emphasis:

Osc 1: Set to a triangle wave, Octave to 4' and enable RP. Put the Level all the way round, Spread to 10 o'clock and Sub to 9 o'clock.
Osc 2: Enable, set to sawtooth wave and dial the octave up to 2'.
Noise: Enable and set the level and Colour to 9 o'clock.
Character: Enable and set to Spherical with Amount all the way round and Sound set all the way left.
Filter: Put the Drive and Punch to noon, Res up a little bit, and Key to 10 o'clock. Set the Cutoff to 10 o'clock.
LFO 1: Enable, set to sine wave, enable Free and set the rate to just past 9 o'clock.
LFO 2: Leave disabled.
Amp Env: Put Attack to zero and set DS and R to 1 o'clock.
Env 2: Keep Attack at zero and put »

This patch, called Muffin FM, demonstrates Mai Tai's frequency-modulation abilities.

Demented Wurzel is a Multi Instrument patch comrising four instances of Mai Tai.

Decay to 10 o'clock, Sustain and Release to 11 o'clock.
Env 3: Unchanged.
FX A: Enable Chorus, Delay and Reverb. The Chorus should have its speed wound up to 2 o'clock and the Depth reduced to 11 o'clock. For Delay raise the Low/High to 10 o'clock, set time to 1/8, FB to 2 o'clock and Mix to just past noon. Set PingPong to Panned and enable the Reverb button. For the Reverb set the Damp to maximum, Size to 3 o'clock, High to 10 o'clock and Mix to 1 o'clock.
Mod A: Env 2 controlling Filter Cutoff at about 80 percent and LFO 1 is also set to Filter Cutoff at 70 percent.

Once you start playing, you'll immediately push into the pluck sound, but if you hold some notes, you'll encounter the rising tide of the slowly modulated filter.

If it's feeling a little bit light and fluffy then click on FX B, and let's give it some excitement. Enable Gater, set it to Slackjoint at 1/2 speed. Enable Distortion, set it to Soft Tube 2 and push up the Drive. Hold a wide chord and play some lead over the top. You might want to increase the voice count to 32 to prevent note stealing. Change the Distortion settings to taste.

Muffin FM

Frequency modulation in Mai Tai is not exactly nuanced. You can't modulate one oscillator from another, but the LFO goes up into audio rate, which gives us possibilities. I've fought with it quite a bit and can only offer you this patch as a place for experimentation within the context of a malfunctioning telephone exchange:

OSC 1: Set to a sine wave
OSC 2: Enable, set to sine wave.
Noise: Disable
Character: Disable
LFO 1: Enable, set to sine wave, enable Key and Free and set the rate to around 3 o'clock.
LFO 2: Enable and set to a sine wave at 9 o'clock.
Amp Env: Put A, D and R to zero and set Sustain to maximum.
Env 2: Unchanged.
Env 3: Attack to between 9 and 10 o'clock and Sustain to maximum.
FX A: No effects.
Mod A: Set Voice Pitch to LFO 1 Frequency, LFO 1 to Osc 1 Pitch and then again to LFO 2 Pitch, all at 100 percent. Set Env 3 to LFO 1 Frequency at -25 percent and LFO 2 to LFO 1 Frequency at -25 percent.

Play with the LFO 1 Rate and the Env 3 and LFO 2 modulation amounts to find your favourite clangy sounds.

Demented Wurzel

Lastly, let's look at a really simple way to combine multiple Mai Tais in a cacophony of sine waves. The idea is to use eight sine-wave oscillators in an unholy drone. Sadly, drones are not possible without holding notes down, because there's no direct access to the VCA.

Drag in a New Multi Instrument from the Instruments browser. In the Multi Instrument panel, click on '+ Instruments' and add a single Mai Tai. Load up the init patch, enable both oscillators and set them to sine waves, and increase the Release on the Amp Env to about 12 o'clock. Now we're going to duplicate the Mai Tai and make some small adjustments to give us our Demented Wurzel organ. To create multiple Mai Tais, simply hold Ctrl and click-drag the Mai Tai into space. Do it again until you have four side-by-side. Then set them as follows:

Mai Tai 1: Set Osc 1 to 16' and Osc 2 to 32'. Nudge the Semi on Osc 1 to the left two semitones. Enable all three effects under FX A and set the delay time to 1/8D.
Mai Tai 2: Set Osc 1 to 8', nudge the Semi left one semitone and increase Spread to 10 o'clock. Set Osc 2 to 16'. Enable the FX A effects and set the delay to 1/2T.
Mai Tai 3: Set both oscillators to 4'. On Osc 1 push Semi right to 7 semitones and Spread up to 10 o'clock. Enable the FX A effects with the delay set to 1/4D.
Mai Tai 4: Set both oscillators to 2'. Pull the Semi down four semitones on Osc 1. Push the Spread up to 11 o'clock. Turn on the same effects and set the delay to 1/2.

Play single notes for some nicely disturbing organ sounds. We can make this more interesting by adding an Arpeggiator. You'll find one under Note FX in the Multi Instrument window. Make sure it sits in the chain above the row of Mai Tais. Set it to Random, set it to 1/4 time and a two-octave range and you will start to appreciate all those offset delay times we put in.

I hope that's given you some ideas for your exploration of Mai Tai.

Want to get the lo-fi vibe using Cubase's stock plug-ins? Here's how.

JOHN WALDEN

Whether it's within the abundant supply of relaxation-meets-study music, embedded firmly in some types of hip-hop, or just blended subtly into a whole range of other popular genres, lo-fi is very much back in fashion. If you want to take your pristine recordings and add a touch of sonic degradation to create a warm, nostalgic sound there are some very popular third-party effects plug-ins tailor-made for the job. But if you have Cubase, it's also a style you can achieve using just the stock plug-ins — with no need additional expenditure required. So let's see how Cubase can help you embrace all those perfect imperfections!

Tools Of The Lo-Fi Trade

Typically, signal-chains for lo-fi feature a number of elements. For example, electrical noise or vinyl crackles might be applied. Tape (whether real or emulated) can be used to simulate the pitch modulation caused by varying tape speed or wow, flutter and dropouts. Distortion can be added via real or virtual analogue circuity or tube distortion. Sample degradation through bit-depth or sample-rate reduction can be used to 'downgrade' the sound. Pitch modulation or resonator components are often added to reverb or delay effects. And, finally, the bandwidth of the sound can be restricted using filters.

As with any effects chain, changing the order of the effects in the chain can lead you to different results, and you can use as few or as many of these options simultaneously as you wish, and adjust the wet/dry balance of individual effects or the whole chain to taste.

Stock Options

Given the typical processing options described above, the first screenshot shows some obvious candidates from Cubase's bundled plug-in collection that might fulfil each role. To give you a better idea of the kinds of effect this example signal chain can deliver, I've created some audio examples that you can find on the *SOS* website (https://sosm.ag/cubase-1124). These are based around two particularly common targets for lo-fi processing: a piano part and a drum loop.

The full signal chain I used for my lo-fi processing experiment including Grungelizer, BitCrusher and FX Modulator.

So, what have I got in my example effects chain and why? For 'turntable noise', Grungelizer is definitely the plug-in for the task. It lets you add noise, crackle (with a turntable speed switch) and distortion effects as well as an element of mains hum. For this workshop's experiments, I started with the suitably named LoFi 1 preset and dialled in to taste from there. Really, the only thing to note is that the noise elements added by Grungelizer are 'always on', whether or not the instrument it's applied to is playing. Now you *might* want the sound all through your track, but often you won't, particularly if you're processing multiple parts in this way, where the effect will 'stack' unhelpfully. Here, I've added an (optional) instance of the standard gate plug-in immediately after Grungelizer and set its threshold so that it mutes the noise from Grungelizer when there's no instrument signal playing.

For a tape-style detuning effect (caused when the tape's playback speed varies, resulting in variations in pitch), I've used the FX Modulator plug-in. This includes a preset called Old Tape Record, and I used this as my starting point. I then focused on just the Pitch module (I removed the Reverb, Chorus and Filter modules). The result is a long, slow, gentle modulation curve that's applied to the pitch of the incoming audio. I tweaked the default curve to make the effect a little more obvious in the audio examples, but you can push things much harder and fully customise the modulation curve to suit your needs. Cubase Elements users could replace FX Modulator with one or more of the modulation-style plug-ins such as Chorus or Flanger: add a little parameter automation and you can recreate much the same sorts of gradual pitch-drift.

For distortion duties, there are lots of options but I chose Quadrafuzz. I added a single band of tube-based distortion in the 1-5 kHz frequency range and boosted the output gain a few dB to make it more obvious. On its own, this gave a somewhat crunchy, crisp, and bandwidth-limited flavour, but you can then use Quadrafuzz's built-in mix control to dial in the effect to taste. Elements users could easily substitute in DaTube, or either the Tape or Tube Saturation modules of the MixConsole's Channel Strip to achieve an effect that's in the same ballpark.

BitCrusher is the obvious choice for downgrading your sounds. Again, there's a preset called LoFi that provides a good starting point, but the plug-in's Depth control makes it easy to dial in whatever degree of bit reduction your ears are comfortable with. The Sample Divider and Mode buttons provide plenty more variation, while the Mix knob can again be used to blend the wet/dry balance to taste.

By this stage, you may well have all the options you need to downgrade your sounds to pleasingly lo-fi status but, just for good measure, I finished off my example effects chain with instances of ModMachine and RoomWorks (choosing the Rhodes Beneath The Waves and FX LoFiverb presets, respectively). The former works well on my drum loop, adding a subtle movement, but you can raise the Mix value to increase the amount of 'seasickness' it induces. The latter works better with the piano part, injecting an additional element of lo-fi to the ambience. Again, you could easily substitute alternative plug-ins for either task or both, but these options are built into Cubase.

Insert Or Send?

The whole chain can be set up using the insert slots on the relevant instrument's track, of course, but you may well want easier control over how much of this lo-fi processing you hear versus the dry sound, and in that case it makes better sense to have the whole signal chain sitting on an FX Channel that's fed by a send from the instrument track. That's the approach I adopted here, and I chose to make the send pre-fader. That way, you can set-and-forget the actual send level, and simply use the instrument's track fader (unprocessed signal) and FX Channel fader (processed signal) to set the desired blend of clean and lo-fi.

As shown in the screenshot, I also used the High-Cut and Low-Cut filters in the FX Channel's Pre section to bandwidth-limit the lo-fi processed element between 250Hz and 5000Hz, for an even less hi-fi sound. Again, I've provided audio examples to illustrate the possibilities.

On, Off, Wet, Dry, Up & Down!

With the initial plug-in selection and settings sorted, it's then a case of dialling in exactly what combination of these effects you wish to use in any particular case. The options are pretty much endless: turn individual plug-ins on/off via their bypass buttons; finesse the unprocessed to processed balance further using the wet/dry mix controls found in many of these plug-ins; or drag individual plug-ins up or down in the FX Channel's insert slots to change the order in which they process the sound. The choice is yours, but I've included another audio example to illustrate some of the possibilities with the same drum and piano examples.

If you feel the need to expand your lo-fi processing options, there are some excellent third-party free-to-download options including iZotope's Vinyl and Caelum Audio's Tape Cassette 2.

Free Lo-Fi Plug-ins

Hopefully you can see, then, that Cubase's bundled plug-in collection can do a decent job of pushing your sounds into lo-fi territory, but that's not to say we couldn't identify a couple of areas that aren't quite so well catered for. And if you're not ready to stump up for one of the dedicated third-party lo-fi effects processors such as XLN Audio's RC-20 Retro Color (probably the most well-known), you can still usefully supplement Cubase's plug-ins with some rather good freebies.

For example, for those turntable-style noises, iZotope's free-to-download Vinyl plug-in is an excellent alternative to Grungelizer. It includes emulated turntable noise based on different decades (with more sonic compromises the further you travel back in time) and speed settings, as well as a whole range of noise types that you can sprinkle across your audio to give it a more lo-fi retro sound.

Another personal favourite is Caelum Audio's Tape Cassette 2, which is also a free download. As the name suggests, it emulates the distinctive sound of cassette tape (anyone under 25 might need to ask their parents about this), including tape saturation, noise, and the wow and flutter created by variations in the tape speed on playback. For a degraded tape sound, this perhaps does a more convincing job than just adding pitch modulation using FX Modulator, Chorus or Flanger, but a combination of both can also be very effective.

Again, I've included some audio examples to make a comparison between the stock plug-in approach and what additional lo-fi character these third-party options can deliver. Oh, and while we are on the subject of 'free', don't forget Steinberg's own and rather excellent free-to-download Lo-Fi Piano expansion for HALion/HALion Sonic. This is a cool piano library in its own right, but there are also a number of lo-fi effects options within it.

I used a single band of Quadrafuzz to add some tube distortion to my lo-fi processing, but there are other distortion options within the Cubase plug-in collection that could serve a similar role.

There's nothing more dispiriting than opening a session only to see the Missing Files error. We show you how to find your audio again!

JULIAN RODGERS

Opening a Pro Tools session and encountering the Missing Files dialogue box can be alarming, but there's no need to panic. This situation can occur if you've moved to a different computer, disconnected a hard drive or reorganised your sample library, causing Pro Tools to lose track of where certain files are located. The default option Pro Tools offers is to automatically find and relink those missing files, but in my experience this isn't the best choice. It depends on why the files are absent, but for the kind of work I do, when a file is reported missing, it's probably not on the system any more. Automatic relinking can be a very slow way of confirming what I already suspect. A more targeted approach is often better.

Missing In Action

The first thing I do when faced with the dreaded Missing Files message is to select Skip All and check out the session to see exactly what's missing. If all you're missing is an alternative take that isn't being used anyway, then you're probably good to go. If you can't see anything obviously missing on the timeline, the first place to look is in the toolbar. At the bottom of the Main Counter section, you'll see the Timeline Data Online Status and the Session Data Online Status indicators. If anything is missing then one or both of these will be red. If the Timeline status indicator is red, you have a missing clip on the timeline; if the Session indicator is red, you have a missing clip in your Clips List.

Assuming you have a missing clip on the timeline, look at the Clips List in the right-hand sidebar of the Edit window. Be aware that audio files on your drives, as opposed to sub-clips of those files, appear in bold. Any missing files will appear in italics. If you're wondering how they can have gone astray, one of my most helpful pieces of advice is to avoid doing anything to your session folders outside of Pro Tools. For example, don't tidy up your Audio Files folders using Finder or Windows Explorer; do it from within Pro Tools, or you'll be doing the equivalent of putting the washing up away in someone else's house — they won't be able to find anything!

If you click on a missing file in the Clips List, you can quickly diagnose the problem. Go to the Clips List drop-down menu and select Show / Full Path to see exactly where Pro Tools is looking for the missing media. Often this confirms to me that searching for the missing file is going to be pointless. For example, the file might be on an external drive which isn't connected to your computer, or (a favourite amongst students I have taught) it's on the desktop of another computer! However, if you have reason to think that the file is mislaid rather than unavailable, and you have a good idea where it's going to be, a manual relink will be faster than the automatic option.

Better Than Cure

Ideally, we wouldn't see the Missing Files message at all, and two top tips for avoiding it (beyond the aforementioned rule about not messing with your session folders outside of Pro Tools) are to check that you have Automatically Copy Files On Import ticked in the Pro Tools Preferences, and to always use Save a Copy In... and to tick the Audio Files in Items to Copy box. This will make absolutely sure that you don't have any stray media currently outside your session folder when it's time to archive.

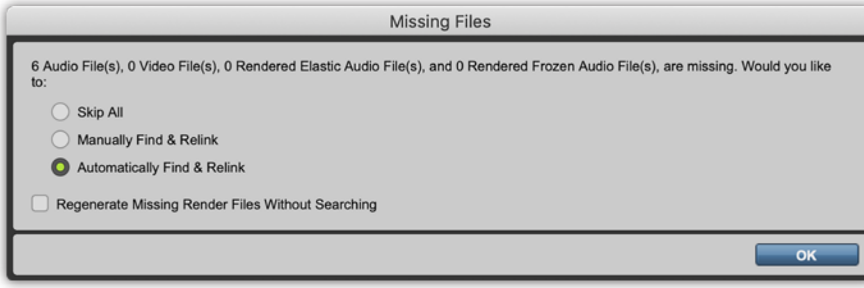

Heartbreaking. Worse, the most common reasons for files being missing won't be helped by the Automatically Find & Relink option!

A Link To The Past

When the Missing Files dialogue appears, select the option to Manually Find & Relink, then click OK to proceed. If you've already skipped the relinking process, or if you've selected automatic relinking and want to revisit the manual options, you can easily access them from within Pro Tools. Navigate to the top menu, click on Window and select New Workspace (or use the shortcut Option/Alt+;). This will bring up a Workspace Browser from which you can manage your files. »

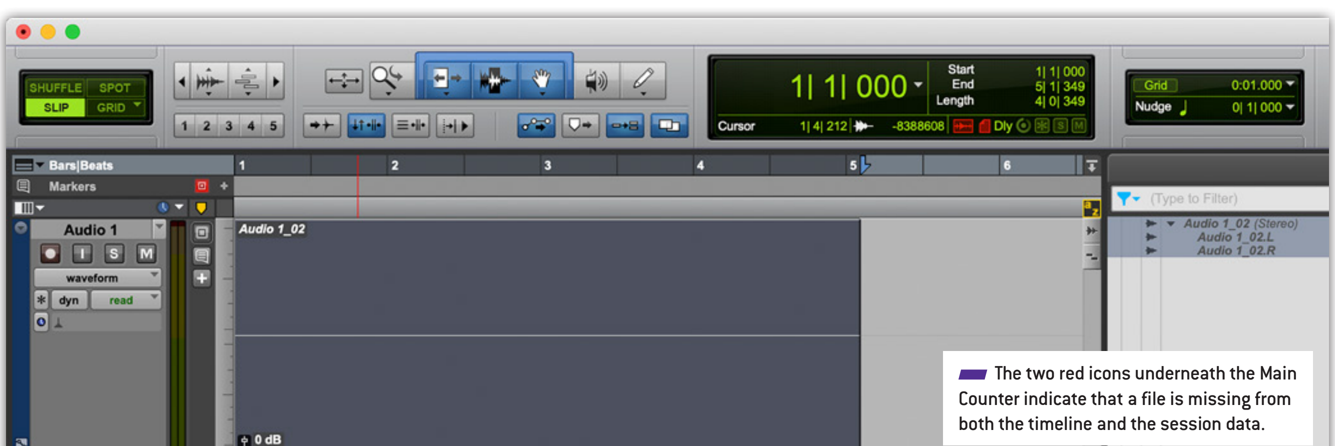

The two red icons underneath the Main Counter indicate that a file is missing from both the timeline and the session data.

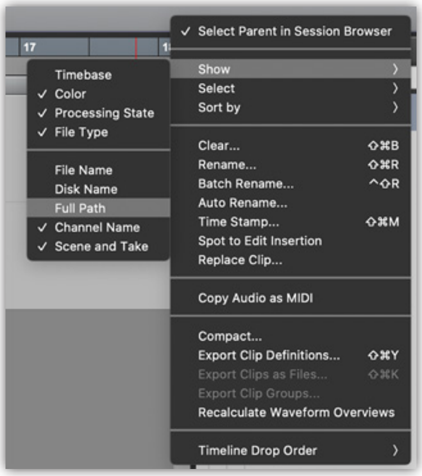

Clicking on Show / Full Path in the Clips List will tell you where Pro Tools is looking for the missing media. This can be a quick way to confirm that the drive your audio was on is no longer connected.

» Within the Workspace Browser, you'll see a list of locations on the left, such as sound libraries and hard drives. Locate your active project folder here and double-click on the session's Audio Files folder. This will display all the audio files associated with your project. If some files are missing, they'll be marked as offline.

To relink these offline files, right-click on one of them and choose Relink Offline. This will automatically select all missing files and open the Relink window. At this point, you need to point Pro Tools to the actual location of the missing files. Navigate to the folder where you know the files are located and tick the box next to it.

Next, select the missing files by clicking on the first one and then pressing Command+A (or Ctrl+A on Windows) to select them all. Click Find Links to bring up the linking options. By default, Pro Tools will search for files using both the name and File ID, which usually gives the most accurate results. If the File ID has changed, you can opt to search by name only. After setting your preferences, click OK.

If Pro Tools successfully locates the missing files, you'll see a link icon appear next to them. To finalise the process, go up to Commit Links and confirm by clicking Yes. Your files should now be properly relinked and visible within your session, allowing you to continue working without further interruptions.

If you're unfamiliar with the File ID, this is a useful feature of Pro Tools whereby it allocates a unique identifying number to every audio file. This means it can distinguish 'Kick_01.wav' in your blastbeat metal track from the 'Kick_01.wav' in your smooth jazz instrumental, even though they share the same name. This also applies to 'Audio_01.wav' if you're lazy about naming things...

"If you have reason to think that the file is mislaid rather than unavailable, and you have a good idea where it's going to be, a manual relink will be faster than the automatic option."

Riding The Waveform

While missing files means we lose audio data, it is also possible to have the right audio but incorrect waveforms in your Pro Tools sessions. Thankfully this is simpler to fix than missing files, but it can cause just as much confusion. This waveform issue is caused by a problem with caching. Pro Tools calculates waveform data in advance and stores it in cache files. The Session Wavecache file is stored in the session folder, and can travel with the session, in order that there is no need to recalculate if it's opened on another system. If the cache file is deleted or missing, Pro Tools will recalculate and build a new one. Like any file it can become corrupted or become out of sync with the files in the session. This isn't common, but it can be fixed easily.

I find I sometimes encounter this problem when exporting audio from my video editing software for editing in Pro Tools. If I have made a mistake when exporting (forgetting to unmute a track or selecting the wrong timeline length, say), I'll sometimes find that if I re-export using the same file name, Pro Tools uses the waveform overview from the previous file. Re-exporting with a different file name avoids this but there are also ways to fix this in Pro Tools.

If what you are seeing doesn't match what you are hearing, then this waveform cache issue is probably at the root of it, and it's extremely disruptive to the process of editing. Deleting the waveform cache file from your session folder will force Pro Tools to recalculate the waveforms next time you open the session, but there is a more precise way. Select the clip in question in the Clips List, click on the disclosure triangle at the top of the Clips List to open the Clips List menu, and select Recalculate Waveform Display.

Cache Rich

Missing files can sneak up on you, and one reason for this is that you can delete an audio file from your drive while it's in use, and Pro Tools will continue to play back quite happily. You'll only find out you have an issue when you next open the session. This is because of Pro Tools' Disk Playback Cache. This setting in the Playback Engine Window caches audio data into available RAM to improve read/write speeds from your drives.

Be aware that Wavecache files are also generated in the Workspace Browser and stored in a separate file. This is where the cache of the mini waveform overviews found in the Workspace Browser is stored. It's very common for lots of files that haven't yet been previewed in the browser to lack these waveforms. To calculate these waveforms, select multiple files or folders in the Workspace Browser, right-click and select Calculate Waveforms.

Missing files and missing or incorrect waveforms can be an irritating barrier between opening Pro Tools and actually getting to work, rather than scratching your head and wondering why Pro Tools is being difficult today! Hopefully these tips and ideas help minimise these issues in the future.

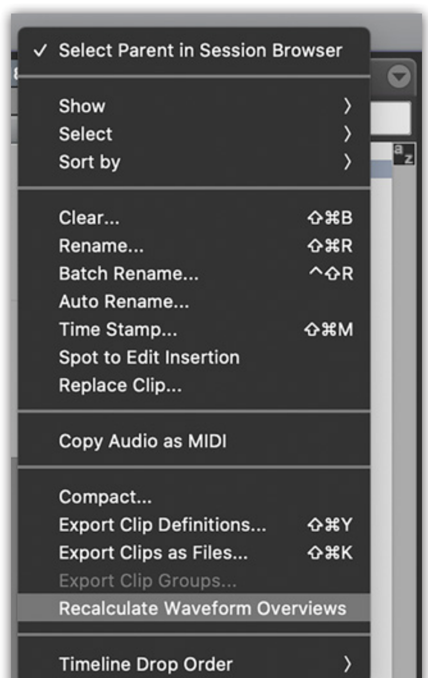

If the waveforms don't match what you're hearing, you can have Pro Tools redraw them.

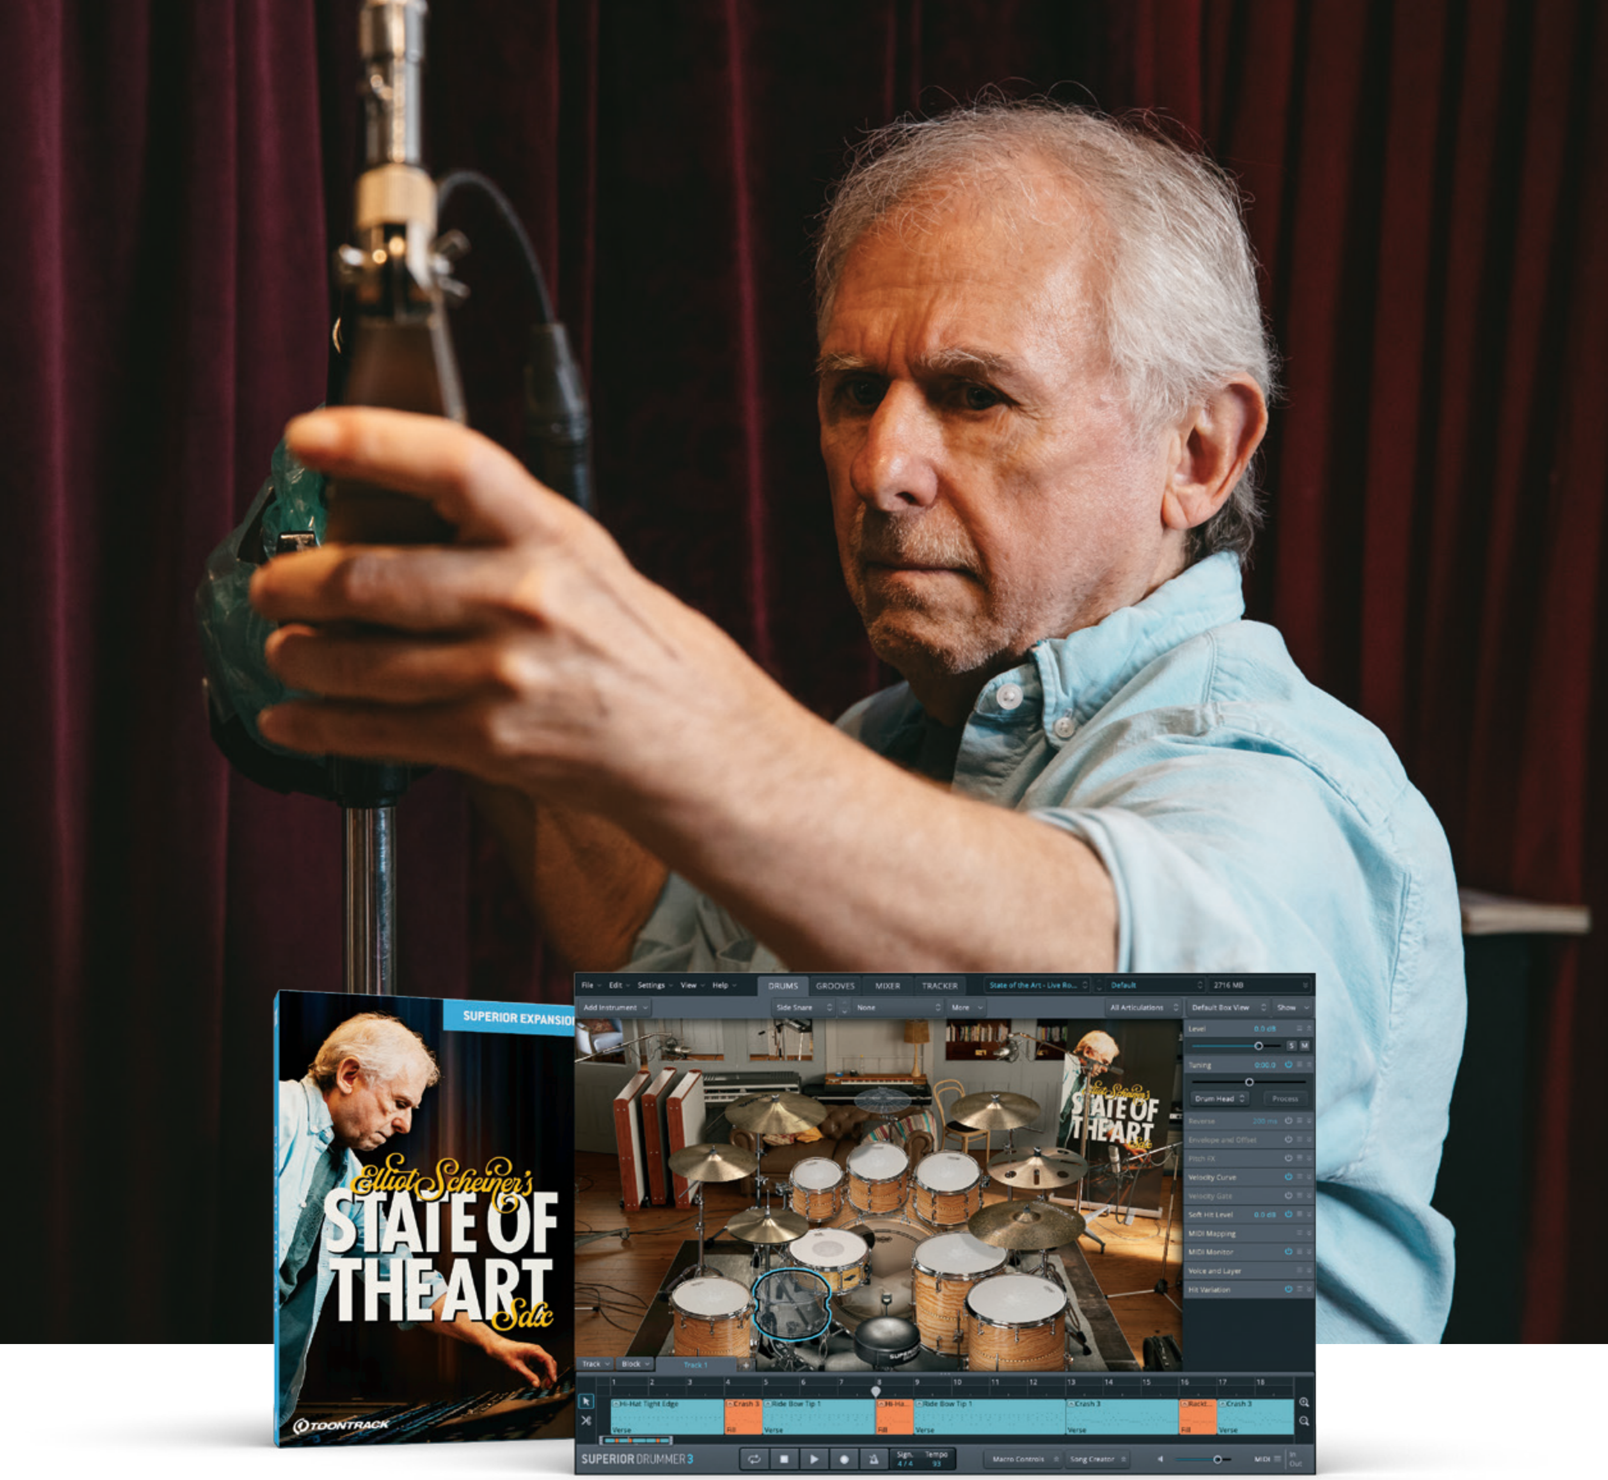

BEAUTIFULLY ACOUSTIC.

The best of drums captured with the best of gear by one of the absolute best in the history of the business: **Elliot Scheiner**. Have a listen to the **State of the Art SDX** for Superior Drummer 3 on our site today.

WWW.TOONTRACK.COM

TOONTRACK

Arturia DrumBrute Impact

Spotlight

Hardware Drum Machines

There's been a tremendous renaissance in hardware drum machines recently. We round up some of the best.

LUKE WOOD

Much like the samplers we covered recently, hardware drum machines have seen something of a resurgence in recent years, and there's now a healthy selection of modern options for those looking to craft and sequence drum sounds outside of their DAW. Some are all-analogue affairs that seek to recreate the ever-popular sounds of classic Rolands; others take a more modern approach and aim to do something a little different — while a few offer the best of both worlds. This month, we shine the *SOS* spotlight on a selection of instruments that suit a wide variety of needs and budgets.

Acidlab Drumatix & Miami

Acidlab's Drumatix recreates the sound of Roland's TR-606, but offers an increased voice count and a far more comprehensive set of controls than the original unit. As well as additional snare and bass drums, the Drumatix also adds clap and rimshot channels, and includes a set of tuning, decay and tone controls that make it possible to fine-tune the sound of the drum hits. A pattern mode makes it possible to create sequences up to 16 steps long, and multiple stored sequences can be combined to build longer phrases using a track mode. A main mono output is joined by six individual voice outputs, along with trig out, MIDI in and DIN Sync connections.

Acidlab Drumatix

If it's TR-808 sounds you're after, then it's worth taking

Acidlab Miami

a look at Acidlab's Miami, which closely replicates the features and sounds of the sought-after classic but with MIDI and DIN Sync connectivity.

$ Drumatix €590, Miami €690.
W www.soundonsound.com/reviews/acidlab-drumatix
W www.soundonsound.com/reviews/acidlab-miami
W www.acidlab.de/

Arturia DrumBrute Impact

Arturia's compact DrumBrute Impact packs in 10 analogue drum sounds. Along with the expected selection of kicks, snares, toms and so on, you also get a two-operator FM synth channel that can be used to craft an array of additional hits. The instrument is capable of storing four banks of 16 patterns, each of which can contain up to 64 steps, and it's possible to chain up to 16 patterns together to create full songs. There's a Swing parameter for introducing some timing variation, as well as a Random control that will automatically trigger additional hits — both controls can be applied globally or to individual tracks. As well as an assortment of per-drum decay, pitch and tone controls, there's a Color function that can apply some per-track or per-step saturation; if things are still left sounding too clean, the output section houses a global distortion effect. Four individual outputs provide access to the kick, snares, hi-hats and FM drum channels, and are joined by main mix and headphone outputs, clock in and out connections, MIDI in and out on DIN sockets, and USB MIDI.

$ $349.
W www.soundonsound.com/reviews/arturia-drumbrute-impact
W www.arturia.com/products/hardware-synths/drumbrute-impact/overview

Alesis SR-16

Although it was originally released back in 1990, the SR-16 is still a part of Alesis' current product line-up. It comes loaded with 233 realistic drum sounds that are available both dry and processed using the company's popular digital reverb processors, and benefit from a Dynamic Articulation feature that varies the sound of each based on playing velocity. It includes 50 patterns that have been played in by real drummers rather than programmed, and which offer A and B parts as well as a choice of two fills. In addition to its velocity-sensitive pads, the SR-16 can be triggered via MIDI, and offers a pair of footswitch connections for external control of playback and part/fill switching.

$ $159
W www.inmusicstore.com/alesis-sr16.html

Alesis SR-16

Behringer RD-6

Behringer Rhythm Designer Series

Another take on some Roland classics comes from Behringer, whose RD-6, RD-8 and RD-9 offer recreations of the TR-606, TR-808 and TR-909 respectively. The compact RD-6 comes loaded with eight drum sounds and a 16-step sequencer, offers hands-on control over each voice's level and boasts a built-in global distortion. There's plenty of connectivity on offer, with MIDI (USB and DIN) and sync I/O joined by a pair of trigger outputs, and a main mix output complemented by individual outputs for each voice. Stepping things up a gear, the RD-8 packs in 16 drum sounds, a 64-step sequencer, a Wave Designer feature that can be used to shape each voice, and a 12dB/oct dual-mode (low- and high-pass) filter. Swing, Flam and Probability parameters help to add some variation to sequenced patterns, and a total of 93 knobs, switches and buttons provide a wealth of hands-on control. As for connectivity, you

Behringer RD-8

Behringer RD-9

get a main mix and individual voice outputs, three trigger outputs, sync I/O, MIDI in, out and thru on DIN sockets, and USB MIDI. Lastly, the RD-9 offers the same sequencer, dual-mode filter, Wave Designer and array of socketry, but with 11 drum sounds inspired by the TR-909.

$ RD-6 $129, RD-8 $319, RD-9 $329.
W www.behringer.com/series.html?category=R-BEHRINGER-RHYTHMDESIGNERSERIES

Elektron Analog Rytm MKII

As its name suggests, Elektron's Analog Rytm MKII features an all-analogue signal path, but it also includes a digital sample layer that's capable of capturing and editing both internal and external sounds. There are eight drum voices equipped with an analogue sound generator and sample playback engine, all of which benefit from a two-pole multi-mode filter, analogue overdrive, filter and amp envelopes, a pair of effects sends, an assignable LFO and a dedicated LFO fade envelope. Elektron devices are renowned for their powerful sequencing capabilities, and the Analog Rytm is no exception: there are 12 drum tracks that can contain up to 64 steps and boast independent scale and length settings, as well as supporting both sample- and »

Elektron Analog Rytm MKII

sound-per-step changes and parameter locking. There's no shortage of effects, either, with send-based delay and reverb processors joined by global stereo distortion and compression. As for connectivity, you get main and individual outputs, a stereo input, MIDI in, out and thru, DIN Sync out, CV/expression pedal inputs and both audio and MIDI over USB.

$ $1799
W www.elektron.se/en/analog-rytm-mkii-explorer
W www.ikmultimedia.com/products/iloudmm

Erica Synths Perkons HD-01

Erica Synths like to do things a bit differently, and say that the Perkons HD-01 "tears down the borders between drum machine, synthesizer and drone instrument". The instrument features four hybrid voices that pair a digital sound engine with an analogue multi-mode (high-, band- and low-pass) filter and overdrive, with Tune and Decay controls allowing users to craft all manner of interesting sounds. A four-track sequencer offers a choice of four time divisions and multiplications per track, as well as per-step ratchets and probabilities and a range of shuffle and groove settings. There's an onboard compressor and BBD-style delay, along with an LFO with morphing waveforms that can be routed to up to eight destinations per voice. Every voice gets its own effects loop and audio output connections along with a trigger input, and the device is equipped with clock and MIDI I/O. The company's range also includes the LXR-02, a more compact six-voice machine that combines a powerful digital sound engine with a 64-step sequencer.

$ Perkons HD-01 $2059, LXR-02 $599.
W www.soundonsound.com/reviews/erica-synths-perkons-hd-01
W www.ericasynths.lv/shop/standalone-instruments-1/perkons-hd-01-black
W www.ericasynths.lv/shop/standalone-instruments-1/drum-synthesizer-lxr-02/

Erica Synths Perkons HD-01

Erica Synths LXR-02

IK Multimedia UNO Drum

IK Multimedia UNO Drum

Another example that merges the analogue and digital worlds is IK Multimedia's UNO Drum. The compact device features six channels that can be switched between an analogue drum voice or sample playback, and six dedicated to samples. The analogue engine takes care of kicks, snares, claps and hi-hats, while a collection of factory samples offer a range of toms, rimshots, cowbells, rides and crashes. The instrument's 64-step sequencer allows users to automate up to eight parameters per step, and it's possible to create longer grooves by chaining together up to 64 patterns. Five performance effects (Roll, Fill, Random, Swing and Humanize) make it easy to introduce variation to programmed patterns, and a global effects section houses a stutter effect, an analogue drive and a compressor. Audio and MIDI I/O are provided by four mini-jacks, and the instrument can either be powered via USB or four AA batteries.

$ $269.99.
W www.soundonsound.com/reviews/ik-multimedia-uno-drum
W www.ikmultimedia.com/products/unodrum

Korg Drumlogue

Korg Drumlogue

Featuring new analogue circuitry designed by the engineer behind Korg's ARP 2600 M, MS-20 Mini and ARP Odyssey, along with seven versatile digital voices, the Drumlogue is said to mark a paradigm shift in the world of drum machines. The analogue voices provide kick, snare and tom sounds, while their digital counterparts offer six sample-based sound sources along with the Multi Engine from the company's Minilogue XD, meaning the instrument is capable of playing fully fledged synth voices as well as drum sounds. A 64-step sequencer makes quick work of creating complex patterns and polyrhythms, and includes probability, alternate trigger pattern, micro offset and groove pattern settings that can all be applied on a per-step basis. As for effects, you get send-based delay and reverb processors and a master effects section that can be bypassed on a per-part basis. Four individual outputs complement the main stereo out, and there's also MIDI I/O (via DIN and USB), sync I/O and a USB host port for hooking up an external controller. If you're after something smaller, the Volca Drum and Volca Beats both squeeze an array of drum and percussion sounds and a 16-step sequencer into a compact, portable package.

$ Drumlogue $399.99, Volca Drum/Beats $159.99.
W www.soundonsound.com/reviews/korg-drumlogue
W www.soundonsound.com/reviews/korg-volca-beats-bass-keys
W www.korg.com/us/products

Roland TR Series

Seeing as so many drum machines draw their inspiration from Roland's legendary early instruments, it's no surprise that the company's current line-up includes a few of their own modern alternatives. The TR-06 and TR-08 use Roland's ACB (Analog Circuit Behaviour) technology to faithfully recreate the sound of the TR-606 and TR-808, while introducing


AUDEZE
Studio-trusted
Audeze Sound for Everyone
Only $399
NOMINEE
SOS AWARDS 2025
MM-100
SERIES

Roland TR-06

Roland TR-08

Roland TR-8S

» some useful modern functionality. While the much-loved drum hits remain the same, the onboard sequencers offer some advanced functions including sub-steps, step-loop and per-step probability, while their analogue connectivity is complemented by audio and MIDI over USB. If you're looking for an all-in-one solution, the larger, 11-track TR-8S packs in ACB-powered recreations of a whole host of iconic instruments including the TR-808, TR-606, TR-909, TR-707 and TR-727, as well as several modified versions and the CR-78 CompuRhythm, a predecessor of the TR-808. If that's still not enough, the instrument includes a selection of FM sounds and is capable of blending in both factory and user-imported samples. Designed specifically for hands-on performances, the TR-8S boasts per-track control over levels, tuning and decay, and offers a fully fledged effects section and a powerful sequencer. Stereo analogue inputs and outputs are joined by multi-channel audio and MIDI via USB, and there are eight analogue outputs that can either carry individual voices or function as trigger outputs. Roland also offer the TR-6S, a more compact and affordable six-track version of the TR-8S.

$ TR-06 $399.99, TR-08 $411.99, TR-8S $749.99, TR-6S $411.99.
W www.roland.com/global/products

System80 880

System80's take on the TR-808 actually comes in the form of a 60HP Eurorack module, but can be purchased with an optional desktop mounting kit with built-in power supply, so it's not exclusive to the modular crowd. Six fixed single voices and five switchable dual voices bring the total count to 16, and the on-board sequencer provides up to 32 steps and can store 12 banks of 16 patterns. The drum voices can be triggered externally via MIDI, and synchronisation with other gear is possible via MIDI Clock, DIN Sync or clock pulse, and there's also a pair of assignable trigger outputs. As for audio connectivity, the instrument boasts 11 individual voice outputs along with a master mix out.

$ $995
W www.system80.net/product/880

Vermona DRM1 MKIV

The latest version of Vermona's DRM1 introduces a range of enhancements that include a redesigned power supply and improved noise performance, as well as a more detailed set of per-instrument channel controls. There's also the option to order a unit fitted with analogue trigger inputs that respond to dynamic levels and can convert incoming gate signals to MIDI messages. Its core functionality remains the same as earlier versions, though, with eight dedicated drum channels offering a wide variety of kick, snare, tom, hi-hat and clap sounds, as well as being capable of producing a whole host of metallic-sounding percussion, and even some more extreme 'zaps' and laser sounds! Eight horizontal rows provide a dedicated set of parameters for each voice — including the likes of decay, pitch, attack, noise level, resonance and so on — along with an output/insert connection and the optional trigger input (if specified).

$ DRM1 MKIV $819, DRM1 MKIV Trigger $919.
W www.vermona.com/en/products/drums-percussion/product/drm1-mkiv-2

System80 880

Vermona DRM1 MKIV

INSIDE TRACK

Charlie Handsome

Songwriter and producer Charlie Handsome has drawn on his huge array of influences to reinvent modern country music with Morgan Wallen and Post Malone.

PAUL TINGEN

"Many people spend a lot of time trying to get the perfect sound," says Charlie Handsome. "I don't like doing any of that. When I work with people, I'll say, 'I don't know shit. I don't know how you guys do any of this stuff.' I'm the least knowledgeable about the technical side of sound of all the producers that I work with. I just do the bare minimum. It's about getting out the idea, as long as it sounds cool enough.

"One of my first major records was Travis Scott's 'Drugs You Should Try It', from his mixtape *Days Before Rodeo* [*2014*]. It was just re-released on streaming platforms a week or two ago [*reaching number two in the US*]. I played guitar on that song. My friend FKi and I made the beat, sang some shit over it, and sent that to Travis Scott. He recorded over our vocal. That's the song. There never were any stems. There was no additional work. It's out. It sounds fine. Nobody gives a fuck. People like that song.

"Artists will go, 'Oh, we should get this song mastered.' You could. I'm not saying you shouldn't. But you could also just put the shit on Spotify. They'll run it through an algorithm that'll match your sound level with other records on Spotify. It's almost like they master it. My point is this: pick one of your favourite songs. If the sound quality isn't all the way there, it would still be one of your favourite songs. You wouldn't say, 'Oh, they didn't record this through this or that.' You wouldn't care. The song is the most important thing. If something I've done sounds polished, it's probably because whoever I collab'ed with worried about that. Or it's the mix. Because we still use mixers and mastering people to finish records."

Cultured Palate

Charlie Handsome's success thus demonstrates the truth of the old adage that good songs count more than everything else. In the 10 years since 'Drugs You Should Try It', Handsome has become one of the world's leading songwriters and producers. His Instagram account lists Morgan Wallen, Post Malone, Juice WRLD, Drake, Jack Harlow, the Weeknd, Khalid, Kanye West, Young Thug, Travis Scott, Lil Uzi Vert, Kodak Black, Kid Cudi and Lil Durk as his main credits, and there are many more, including mike, Cardi B, Tate McRae, Lil Wayne, Gunna, MIA and the Kid Laroi.

Handsome has been involved in some of the biggest songs and albums of the last decade, among them Jack Harlow's 'First Class' (2022), and Morgan Wallen's *Dangerous: The Double Album* (2021) and *One Thing At A Time* (2023). Handsome has won several BMI Pop Awards, Album of the Year at the Country Music Awards for Wallen's *Dangerous*, and a Grammy nomination, for Best Country Song, for Wallen's 'Last Night' (2023). In September, Handsome was number one in MusicRow's Top Songwriter Chart, due to his work on all 27 tracks on Post Malone's number one album *F1 Trillion*, as well as on Morgan Wallen's 'Cowgirls', Dylan Scott's 'This Town's Been Too Good To Us' and Moneybagg Yo's 'Whiskey Whiskey'.

It's a hugely impressive list of credits, which is all the more striking for its variety of genres, from hip-hop and trap to pop, rock, and country. There are similarities in this respect between Handsome's journey and that of star producer Andrew Watt, who featured in last month's *SOS*. Both started out as aspiring rock guitar heroes, got into production through an obsession with hip-hop, and ended up working across multiple genres. "It probably starts with being a fan of music in general," reflects Handsome. "I was buying CDs back in the day, and after that you had Limewire, Ares and Napster, and I would download as much stuff as I could. I would listen to a wide variety of music. When I first started playing guitar, I was learning all the Nirvana songs and all the Led Zeppelin records. From those two, I tried to learn every classic rock song, all the '80s heavy metal records, all the metal stuff. I used to shred. I put on all the effing Yngwie Malmsteen and Joe Satriani records! At the same time I was listening to rap the whole time. I was playing everything I could on guitar, and in my free time I was listening to rap. I just wanted to have a wide palette."

Emo Country

Handsome was born Ryan Vojtesak in 1989, in Atlanta, Georgia. He started playing guitar when he was 13, and during his teenage years "played in bands, doing metal, emo, post hardcore, pop punk, whatever was popular at the time. But hip-hop got me into production. The whole journey of my career was because of hip-hop. I could cite a lot of references, but some of my biggest ones were Dungeon Family, OutKast, and especially Kanye. Because hip-hop was sample-heavy for so long, you were getting many different genres. Soul, rock, blues and jazz records all went into hip-hop songs. The sky was the limit in terms of what you could do with hip-hop beats. My first beats were all sample flips, essentially Kanye beats. And then I started finding my own sound and stepped out of sampling, because at the end of the day, if you want to make some money, you got to stop sampling records.

"I'll give you an example: the song 'Go Flex' on Post Malone's first album [*Stoney, 2016*]. I wrote the music for that at 18 or 19, when I was in college. It was a different song at the time. Fast forward almost 10 years later, with Post and I in the studio. It was by then a mix between a folk guitar with trap drums and claps in the hook, because at that point I was already doing trap production. So that is why that song is such a mix of genres. But it was not intentional, as in 'Oh man, I want to mix all these different genres.'

"I could keep going. If you look at the Juice WRLD records I did, a lot of them are based on my emo sound from when I was in emo bands, with those types of riffs, and then combining that with trap drums. Same with the Travis Scott record, 'Drugs You Should Try', which was essentially what I would have done on a dark emo record. So I've always mixed genres, and I think that's helped me in my career. It's also helped a lot with my transition to country music. Morgan [*Wallen*] and I ended up doing records that no-one else in country had done yet. Obviously, Tim McGraw and Nelly paved the way back in 2004 with their collab 'Over And Over', but I feel Morgan and I pushed it even further down that road."

Charlie Handsome had a big hand in the making of Morgan Wallen's hit albums *Dangerous* and *One Thing At A Time*.

Charlie Handsome: "If Morgan doesn't take a song I've written, I don't call a different artist and give it to them. The song is dead. I'm not giving his sound to somebody else."

Hey, Good Looking

Handsome took on his artist name at the start of his career, when told by a friend that Vojtesak was "too difficult to pronounce and too weird for a producer name. Purely as a joke I thought of Charles Manson, the cult leader, and Handsome, because a friend of mine was poking fun at me, saying I was ugly. I needed an artist name quickly, because my career was beginning to happen, so I combined the two." In 2014, Handsome was introduced to Che Pope, president of Kanye West's label GOOD Music, and decided to move to Los Angeles. He met up with Post Malone in that same year, and in 2016 he worked with his idol Kanye on the track 'Fade', which was released on *The Life Of Pablo*. Handsome conducted his first country session in 2017, with Ernest, and then Florida

Georgia Line, and the enormous success he achieved in the genre, plus lower taxes, contributed to his decision to move to Nashville in 2021.

Although Handsome claims to be "the least knowledgeable" about the technical side of music making, he does know his way around "the bare minimum". He recalls: "When I was in college, I was writing songs on my acoustic and got an eight-track Korg recorder that went to hard disk. This allowed me to layer guitars and make up stuff. Then I got Fruity Loops, from a roommate. I'm really competitive, and thought, 'I can make better beats!' He gave me his program, and I fell in love with that and started making beats all the time with Fruity Loops. In 2011 or 2012, I downloaded Pro Tools for the first time. I switched to Logic in 2016, and I am still using that.

"I prefer Logic because I do a lot of time-stretching and a lot of pitch stuff. It's on almost every song I do. It really is a lot. When you hear the final product, you don't realise that the song started in a different place. For example, you'll hear a song I did that is in, let's say the key of D, with a tempo of 120 bpm. But when I started, I probably played that song at 90 bpm in the key of G. I then sped it up and then I switched it. I treat a lot of the riffs and a lot of the vocal recordings I do on the mic almost as a sample. I flip them. By the time you hear the beat or whatever we end up doing, it's completely different.

"Pro Tools is very slow at doing that. You have to do five, six things to make it work and sound good. Whereas, in Logic, you can put on the varispeed function and it literally needs just one click of a button to up or down the tempo, or the pitch, and it sounds great. So that's what I do. People are always aggravated in the studio with me for my tempos, because when you're using varispeed, for the song to stay perfectly in a key, something in the algorithm means that you end up with a tempo with decimal points. My tempos are like 83.946 or 140.323 bpm. Just weird. Afterwards, a lot of times we'll time-stretch it to the nearest whole number, but sometimes we don't."

Riffing

"I have a few different formulas when it comes to creating stuff. Most of the time I make stuff very, very quickly and very mindlessly. I don't really sit around with an intention. I don't sit around and think, 'All right, I got to make a song and it needs to have this feel.' Throughout the years when working with different artists and producers, a lot of people come from the approach of 'We want to make a song like...' and then they have a reference song. I hate to make songs that way.

"What I do generally is pick a random tempo, the first tempo that comes in my head. Then I press record, not knowing what I'm going to play. I play for eight bars. I hit stop, open a new track, hit record, layer the eight bars with a lead, stop, start again, and layer the eight bars with a pad guitar, for which I use Valhalla Shimmer most of the time, with a big reverb that sustains for a long time. Once I have those three layers, which takes about two minutes or less, I'll usually sing over it and then chop up the vocal like a sample, and I'm done.

Charlie Handsome: "I try to make beats that are already a hit, and then all we have to do is write the right song to it."

And I move on. I make usually 10, 15 of those in a row. After that I take a break and then I look back at them and pick the best ones and/or put drums on a beat.

"In doing these ideas, I really homed in on a style of guitar playing that mixes my influences, and that's less about soloing and shredding, and more about good riffs. The vocal melody also is very important, but I try to make sure that there's some sort of riff or musical direction or melody in the beat itself that can carry it. I try to make beats that are already a hit, and then all we have to do is write the right song to it. And, if you have the right hook, you can have trash verses that nobody cares about. You shouldn't be doing this, but if you have a hook that's a hit song, you can get by with bad verses.

"A lot of the time I don't do the full production any more on these initial ideas, because I don't want to waste time. It's more efficient to do the full production once we have a song. So I build up a huge library of music. When I'm in sessions with people and we're figuring out what we're doing, rather than play something on a guitar, I just pull from my library. Many of the songs that I worked on started this way. Every Juice WRLD song was probably made this way. 'Last Night' by Morgan Wallen was based on a voice note on my phone, that I'd chopped up and made into a loop in my computer, adding some layers. It was just eight bars and one day in the studio I pulled that up and we wrote the song to it.

"There's not just one way to write songs. I make a lot of music, and in the hip-hop world, I used to send my music out to friends who do rap production, and *vice versa*, a lot of those guys will send me their music. I'll put on drums and will make beats out of their stuff. I have more fun putting drums to someone else's music than my own. It's a new challenge to figure out what works for their stuff. But that's predominantly not any more how I'm approaching the business these days. The idea of 'I send this loop to this guy and hope he makes a good beat and that he plays it for somebody,' is somewhat of a lottery. Now I mostly sit in the room with the artist and figure out what kind of songs we want to do. I do still collab, in the sense that I flip other people's music. As I said, for the most part I don't make full beats unless we start writing a song, at which point I'll finish it. There are a lot of songs where I'll do the music and the beat. Or, I'll hire a band. We'll get a seven-piece band and we'll track, remake the song, and just build it up."

Old School

Handsome was heavily involved in the making of Post Malone's third UK and US number one album *F-1 Trillion*, on which he applied all the above-mentioned writing and production methods. It's not the first time Malone has flirted with country, but it's his first album that contains authentic-sounding country from start to finish, with guest performances by country stars like Dolly Parton, Brad Paisley, Tim McGraw, Chris Stapleton, Luke Combs, Morgan Wallen and others. The album also features an all-star band, including pedal steel player Paul Franklin, fiddle player Larry Franklin, guitarists Bryan Sutton and Derek Wells, keyboardist Dave Cohen, drummer Aaron Sterling, and more.

"The album started with Post's vision," explains Handsome. "We had both wanted to do a country album for a long time, but if he hadn't had his own idea of how he thought his album should be and had just given me a blank slate, there's a chance that I would have taken it in a very progressive country route, which is a sound that I've been creating. I would nevertheless have figured out a distinct and unique sound for him, because that's something I do with all artists. For example, if Morgan doesn't take a song I've written, I don't call a different artist and give it to them. The song is dead. I'm not giving his sound to somebody else. I try to keep the sounds separated.


Charlie Handsome has worked with Post Malone since his first album in 2016.


F-1 Trillion has been a worldwide hit for Post Malone.

"Post had a specific vision. He told me that he did not want a lot of the elements of progressive country I do with Morgan. He said, 'That's not what I'm looking for. I'm looking for like an older sound that's more reminiscent and classic. I love the '90s. I want a lot of the influences and sounds that feel like '90s country.' Later we expanded to also include influences from the '80s or '70s, going all the way back to '50s and '60s in some cases. Obviously you can't recreate the past and you can't contend with the classics, because they're classics for a reason. They've been around for 60 years. Instead he wanted songs that could live next to those songs.

"We did do some modern stuff, but we also reached back and we wanted to give everything an authentic feeling. That's why we got people like Paul and Larry Franklin on the album. These are the guys who have been giving records that authentic country sound through the years. At the same time, what makes the record different is who Post is as a writer. Some of his lyrics, some of his phrasing, some of the melodies, definitely took a direction

F-1 Trillion was recorded in surprisingly traditional fashion at Nashville's East Iris studios.

that feels more authentic to him. You are not listening to a country song, you're listening to how Post Malone is doing a country song."

Grown From Seeds

F-1 Trillion was co-produced by Handsome and fellow star producer Louis Bell, with help on a few songs from Jonathan Hoskins. Handsome, Bell and Post Malone have co-writing credits on all songs. As is common these days, several other songwriters were involved as well on most songs. The writing and recording process took place for the most part at The Cave in East Iris Studios in Nashville.

"Some of the songs started in the room, sitting around playing acoustic guitar together, usually Post, one of the writers, and me, and we'd just kick around ideas. If someone had a good chord progression or another idea, we would build on that. But I think the bigger, more memorable songs, the singles and so on, came from me already having a riff or an idea. I'd have half of a track made, and we wrote to that. We brought in the band later. For me, that's how we got the best results.

"So I tended to come in with the musical seed idea, and then I co-wrote the lyrics and the melodies. I also figured out who I wanted to work in the room together, as in 'I want this writer with that guy.' Regardless of whether I had the concept for a song or someone else, it's also my job as a producer to say, 'You know what? I don't love that idea.' Or 'I do like that one, but I think we should switch it to this. Instead of making it about this girl, let's make it about that part of the relationship.' It was a matter of coaching the room once we got something going.

"The lead single, 'I Had Some Help', is an example. That song started as a riff I recorded as a voice note on my iPhone. I just take my phone, press record, and make something up. Then I'll either chop up what I have done, or I have another producer work on it who is signed to my company, Jonathan Hoskins. In this particular case, I sent the voice note to him, he chopped it up, added some drums and bass and different stuff, and sent it back, so I had something that I could write to. Musically and beat-wise, that song was just Jonathan and I. We then wrote the song to it, with input from co-writers Post, Morgan Wallen, ERNEST, Ashley Gorley and Chandler Walters."

Finding The Music

Regarding the many country legends who made their way to East Iris, Handsome notes: "Post and I invited the guest singers on the album, and I picked all the writers and the musicians. I hand selected

Handsome Plug-ins

Charlie Handsome tends to make his beats at home, though he admits that he does not have a studio as such. "It's crazy. I'm just really slow at doing stuff outside of making music. I've had the house I'm in for a year, and the house before that two years, and didn't build a studio. I have studios that I go to and use, and some of them I pay for, and some of them are free. I do have some stuff at my home, I just wouldn't consider it a studio. I have Barefoot and Yamaha HS8S speakers, and also some KRK Rokits somewhere. I have a lot of guitars here, and my laptop, my UAD Apollo, and my Shure SM7 to lay down quick ideas. I have a Sony C800 mic as well, but I keep it at one of my friend's studios. Most of the time when I work at home I use headphones.

"I've never been a big equipment guy. Coming up, I was using basic rap gear, like the Avalon 737, the [*Universal Audio*] 1176 compressor, stuff like that. Many of my songs are recorded direct to the Apollo, and then I just mess with a few plug-ins. Since I started, nothing's really changed. My first go-to for electric guitar has been and still is Waves CLA Guitars. I picked a preset called 'Late Night', and modified it. That's my base. It starts with a big reverb. My second go-to is the Valhalla VintageVerb. And my third is the Valhalla Shimmer. They're all stacked with that same CLA plug-in.

"I've been using the stock Logic chorus a lot. It has a lot of cool sounds. I also like Goodhertz. I've been using their plug-ins for years. The Lossy plug-in is probably my favourite. It kind of effs up the audio, which is what I like about it. It puts a bunch of artefacts into your audio, and does these weird EQ things. But that's almost it in terms of plug-ins. A lot of the rest is just pitching, using the basic Waves SoundShifter plug-in, to speed up and slow down audio to get different sounds, or time-stretching. What I use is very basic. I don't use a ton of stuff. It's about which guitar I play, and which pickup I use.

"I predominantly do clean guitar. I don't do distortion. Actually, having said that, I use Minimal Audio Rift a lot. It's really cool. It's a guitar distortion plug-in, but it comes with a lot of delays, which is what I use, and I take out some of the distortion. All these plug-ins are for electric guitar. With acoustic, my go-to situation for almost every song I do is to record it straight into the iPhone voice memos. Then I import that, chop it up so it's in time with the grid, and I put a Waves CLA76 compressor on it, and that's literally it, apart from that I may EQ a little bit of the low end and the high end out. I usually have no reverb."

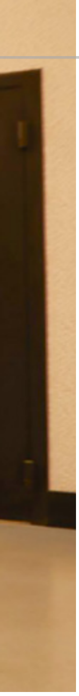

everyone as far as the personnel went for this project. When it came to track the records, we tried to stick with the same players as much as we could throughout. Once the band gets to the studio, they listen, chart the song, and play it immediately. And then we start doing live edits. The guitar player may play a riff and we'll say, 'Let's switch it to this.' Or we say, 'Let's move this note here.' Stuff like that.

"One of the things we were looking for was musicality. Instinctually, I would have no guitar solos, because the modern-day young listener doesn't want to sit through that. They just want to click a two-minute song and move on. But for this album, it was like, 'No, we're going to have fun,' so we had ripping guitar solos and did musical breaks. We were like: 'All right, all four of you are going to solo. Let's go!' We had all these great musicians involved, and they can jam and create bigger sounds. There's more layers on these songs than any of the other songs I've done in my career. Normally I don't go that far. I try to do minimal things.

"Louis [*Bell*] was involved in the writing and production, and he recorded Post's vocals. I think he's done that for nine years now, and they have their routine. And Louis is good at creating libraries. At the beginning of the project he AI-ripped the drum beats of tons of classic country songs. He had hundreds of them. So while I was doing the music, he would go, 'What tempo are we at?' I'd say '135', and he'd have 50 drum grooves that work at 135 bpm. There was this huge catalogue of ideas we could aim for, and that we could time-stretch if needed to make them fit. However, this album was not about programming. I had hired Aaron Sterling to play the drums, and while the libraries sometimes provided a guide, Aaron would play something new.

"Everything we did was part of creating an album as an entire body of work. I do hope that albums will be able to co-exist with individual songs. Doing an album meant that we were able to explore having many ups and downs and different moments. That's a big part of it, and I like that. Even the track listing was important. How do you start? Then what do you want to feel as you're listening to the songs in succession? When listing to the whole album, you want to feel all the highs and lows and emotions involved."

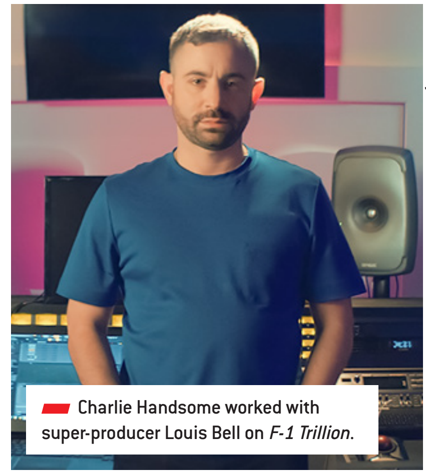

Photo courtesy of Genelec

Charlie Handsome worked with super-producer Louis Bell on *F-1 Trillion*.

And, Handsome could have added, making the songs the most important thing is far more likely to make the journey of listening to an entire album more rewarding.

TalkBack

Katie May

WILLIAM STOKES

At just 26, Katie May has already worked with some of the biggest names in the industry. She trained at Real World Studios in Wiltshire, UK, where she is still based, and has worked on records by Peter Gabriel, the 1975, Harry Styles, Roger Waters, Gabrielle Aplin and many more.

At the moment I can't stop listening to
It depends on whatever session I'm doing. When I work on a pop session, I want to listen to rock music. When I work on, like, a jazz session, I want to listen to techno stuff! When you work in a studio, you need to constantly remind yourself to listen to music for pleasure. For a long while I was just listening to BBC Radio 4 whenever I got in the car after a session. I wouldn't listen to any music, because the last thing I wanted to do after a 12-hour day was force myself to listen to an album. So, I've been going through all of my old Spotify playlists from three or four years ago and listening to what I was listening to back then. And finding a lot of comfort and nostalgia in that. There's an album called *Somewhere* by Sun June, which I listened to religiously about three years ago, and which I've been getting back into now. It's just a really nice dream pop album. But it's not even necessarily about how good the album is. It's just, as soon as I press play, I'm immediately transported to that place I was in. I love that feeling of having such a visceral connection with an album, that you can't really explain. It's as if, immediately, your brain goes back to those exact feelings that you were feeling at that stage in your life.

The artist I'd most like to record
I feel like my classic artist that I've always wanted to record is Laura Marling. But that's not impossible for the future, because she's been to Real World before. You know what, I'd love to work with Fiona Apple, just because there's so much going on in all of her stuff. I loved *Fetch The Bolt Cutters*. There are so many layers of intensity and raw energy and emotion in what she does. It would be a really interesting thing to watch happen in front of you.

I mean, my favourite part of working in a studio is the privilege of being able to watch someone perform in front of you, especially somewhere like Real World where the control room and live room are all the same. So for me, being able to turn around and watch someone who's so talented put so much effort and energy into a take, whether that's a guitar take or a vocal take — or anything, that's such an incredible thing. And no-one gets to see that, except the people in the room with the artist. So watching someone like Fiona perform in front of you would be amazing.

The first thing I look for in a studio
I look for vibe! That is paramount over everything. I try and approach it from an artist perspective: if I was an artist walking in, about to work with an engineer and a producer who I hadn't worked with before, what sort of a room would I need to feel most comfortable? And I find a lot of that is in the small things: the ambient lighting, or having a nice rug out, or a couple of plants up, or whatever. I personally put that above how great it sounds. If you can't get the artist to perform something that's true to them, where they have confidence in what they're doing, it won't be a good record, regardless of how good the mics are.

The person I would consider my mentor
When I started working at Real World, Oli Jacobs was the head engineer here, so he was like my mentor because he was my manager at the time. He very much supported me, and he still does now. [*Former Real World engineer*] Oli Middleton as well. Both of them, I love to bits. I think part of the reason I felt so excited about doing this as a career is because I've had people around me who really lift me up. I feel like there can be quite a lot of competitiveness in this industry; I see some engineers almost step over other engineers to get to their goals. And it always feels really counterintuitive to the whole music industry and the whole process, because you want to be around people who lift you up and make you feel good about what you're doing.

My go-to reference track or album
The way I tend to view references is as a really good way of resetting your ears. A lot of people bring up reference tracks when they want their mix to sound more similar to that album, or whatever. If I'm in a new room, often I'll just put up a song I know really well. So, like, 'Reckoner' by Radiohead. I've listened to

that so many times, something like that. As long as it's something that I know really well, then that's exactly what I would put up. There's one song called 'Good Guy' on the record *Crushing* by Julia Jacklin; that's just my perfect ear reset, to remind myself what my 'flat ear frequency response' is. I don't think I ever put reference tracks on as a means of trying to make something 'sound more like that', because I think that can be quite limiting for how you mix.

When we get bands down here, a lot of times they'll come and they'll say: I want this record to sound like this. And we'll say, OK, cool. We can, you know, choose mics appropriately and place them appropriately for that sort of a sound. But inevitably, it always sounds different because you've got a different room, you've got a different drummer — you're working at a different time of day, even, it's a different season in the year, it's a different year! So it's never going to sound the way that classic album sounds in your head. And that's perfectly fine! It should always be different. Otherwise, what's the point in doing it?

Katie May: "If you can't get the artist to perform something that's true to them, where they have confidence in what they're doing, it won't be a good record, regardless of how good the mics are."

My secret weapon in the studio is

This probably isn't the most technical answer of all time, but what I'd like to say is that I do have a particular level of awareness that I try and bring to every single session. That might just be a social awareness of how to make someone feel comfortable in the room, or trying to read how they might be feeling that morning so you can react to them in the best way, to get the best take out of them.

I've assisted on quite a lot of sessions where it's quite high stress, because the producer might be quite an old-school producer who likes to make the drummer feel like they're not doing a very good job or whatever, in order to get the best take out of them that way. But for me, there's nothing nicer than when a band are in a room together and they all feel comfortable with each other, they'll feel like it's a musically safe space, and they feel like they can talk about things, that everyone can listen to each other. A lot of that comes from the engineer: if the engineer is making everyone feel uncomfortable, even if it's just in their body language or the way their face looks when they react to people and their ideas, those small sorts of things can just get into people's heads. And immediately everyone's taken out of the process. It has such a trickle-down effect, I think. from the recording to the end of the day, to the end of the project. If everyone feels like they're really involved in the process, everyone has more fun doing it. And beyond that, if they're then touring that record a year later, they'll look back on those songs with fond memories, instead of being like, "Oh, that was when I had the worst week of my life, and it was so stressful, but we got a good album out of it." And that then plays into how they tour it, how happy they are playing those songs back.

I've got the luxury of having so many nice mics at Real World, so I've never had to deal with only having one nice mic or whatever in a room, because we have so many nice things. I feel like I've been so exposed to those really nice, expensive, boutique bits of gear that I've realised that, really, it's not about that. I've heard so many bad recordings made with that gear!

The studio session I wish I'd witnessed

Well, the baby part of myself just wishes I could see some of the AC/DC recording sessions back in the day! *Back In Black* was my favourite album when I was a kid. That's the reason I started playing guitar. It »

Katie May in the Wood Room at Real World.

would be interesting to see how differently people worked in the studio back then as opposed to now. The whole rock & roll thing of being in the studio doesn't really exist any more, at least from what I've experienced. It's all peppermint tea instead of cocaine now! Which I think is a good thing. But it would be interesting to look back and actually see how sessions used to be run, just purely out of curiosity.

The producer I'd most like to work with
Well, I would have said someone like Jack Antonoff, but then I've worked with him already! Obviously there are the classics, but I don't think I'd want to have met someone like Steve Albini, because I think it would take some of the magic out of the process as well. Watching Jack work was amazing, but it did also take some of the magic out of it for me. Obviously we work with quite a few big producers at the studio, and that's a really exciting part of the whole process, watching them do what they do. But I also find that there's so much to learn from every single person you work with. Even if it's like a producer who's not very well known at all, there's still so much that you can learn from that one person. They're not necessarily less valuable than a big-name producer.

The studio experience that taught me the most
There was one thing which I think definitely informed how I think about sessions now: we did one of the Breath's last albums. The Breath are a lovely folk band who are on the record label at Real World. Stuart McCallum and Ríoghnach Connolly. They're the nicest people on planet Earth, and so funny. They did an album with Thomas Bartlett, it was just Ríoghnach, Stuart and Thomas doing the session. Seeing the three of them work together... Thomas: his brain is all over the place, in the best way possible. As soon as they did a take, as soon as he had an idea, Thomas would just immediately add, like, a keyboard part. Without even telling anyone what he's doing, he's just going to do it. He does, like, a four-bar section of it. And then, just for speed's sake, he would just then loop that section as and when it needed to be in the Pro Tools session. There was something just really cool about seeing how fast it was done, and how it wasn't about necessarily capturing the most 'human' performance or anything. It was just about keeping the excitement in the room really high, just building it as quickly as you can, because if the momentum continues to grow, the artist will get more excited about what's happening as well.

That was similar to seeing Jack working with the 1975. He was just so fast at everything he did. He jumped between instruments really quickly. And as soon as there was a lot of energy in the room, he would either swap to a different song, or he'd ask people to get involved and get their opinions on things, just to keep everyone really engaged at any one time. As soon as you start to lose that engagement, people don't really care about what's happening any more. They don't care about the part they're playing, they don't really care about the song any more. But if you keep everyone's energy high all the time, have breaks as and when you need them to make sure that you are not destroying yourself; that excitement and that momentum will translate all the way to the end, when it's released.

That's something that I've tried to keep in my head. It doesn't matter if it's not recorded perfectly. It doesn't matter if something's out of phase. You can fix it later. All of that stuff doesn't matter, as long as the energy is there and everyone's excited about what's happening. You're just building these things up, bit by bit. The other approach is being really careful about everything you do, but then writing or recording that song turns into a huge slog. And as soon as it's a slog, it's not fun. No-one's having a nice time. If it's a slog, then something needs to change. Something's gone wrong. Maybe you just need to go have a nap, or maybe you need to change to a different song. But trying to force something to happen when everyone's not having a good time, that's not going to get you anywhere.

The advice I'd give myself of 10 years ago
Just trust yourself! There are a lot of people with a lot of opinions in this industry. I used to go on forums all the time and ask all these questions, try and learn from all these people online, but a lot of the time the people you're learning from are either people who don't do it full time, or they're people who have really strong opinions but don't necessarily have the experience to back it up.

There are lots of people who think that what they believe is the be-all and end-all of of every single area of audio. And I just feel like, if someone was there 10 years ago, telling me: "Katie, there are no rules to any of it. It should just be fun. It should feel like painting. It should feel like, like the most creative, exciting thing in the world." I would have told myself just to trust my gut a bit more and not look to other people for validation. And no-one really said that online. Everyone was like, "Oh, you need this mic to make something sound good. You need this interface. You need this cable." For God's sake! None of that is true.


The NAMM Show
the PROS mix here
Register Now!
namm.org/attend
The NAMM Show Music
The NAMM Show Sound
The NAMM Show Live
January 21-25, 2025
Anaheim Convention Center • Southern California
Events
January 21–25
Exhibits
January 23–25

Classic TRACKS

Model 500 'No UFOs'

Widely credited with not just inventing techno but also coining the name, Juan Atkins tells the story of his genre defining record, 'No UFOs'.

TOM DOYLE

Listening now to 'No UFOs', Juan Atkins' pioneering techno track released under the name Model 500, it's hard to believe it was actually made in 1985. With its stripped-back, driving machine beats, squelchy synth bass line, stuttering vocal effects and emphasis on strict repetition, it was a glimpse into the future and is now considered to be the first techno record.

"It still sounds fresh today," Atkins says, in a rare interview with *SOS*. "I think it's timeless and the reason why is because a lot of my productions were made at the beginning of an era. These records kind of set the tone. But, y'know, that era is still in progress."

While the Detroit producer made his name creating purely electronic music, his musical roots stretch back to his childhood and more traditional instruments. He first picked up a guitar at the age of 10, before he branched out into playing drums and bass.

"I have a brother that's 10 months apart from me," he says. "I used to talk my parents into buying him musical instruments for Christmas, so I could play them. I talked them into buying a drum set and I switched from guitar to bass guitar when I was 13 or 14 years old. I actually asked my grandma to buy this Rickenbacker copy bass.

"I learned by ear," he adds. "I wasn't really a trained musician. I would play along with a lot of records. Eddie Kendricks' 'Keep On Truckin'' on Tamla... that was a record that I learned to play to, because it was a real drum-heavy record."

The now 62-year-old Atkins believes that his initial, non-electronic grounding helped him greatly when he first turned to synths and drum machines.

"Yeah, I mean, I think any type of experience with any kind of musical instrument would definitely translate or transfer into future music endeavours. So I think, had I not had those early instruments, who knows if I would even be making music now?"

It was in 1979 that Atkins' interest in electronic music was first sparked. His grandmother owned and played a Hammond B3 organ, and he would often accompany her to the Grinnell's music store in Detroit when she went there to buy sheet music. One day, the teenager happened to wander into a back room where they kept the synths.

"Of course, these were like baby synths in a way," he points out. "Monophonic, like the Minimoog and the Korg MS-10." To his surprise, Atkins' grandmother bought one of the latter synths for him. When he took it home, a whole universe of sound opened up.

"What I was doing at the store was I was creating all of these different, weird sounds. Like what a UFO landing would sound like, and things like this, and I just put all that stuff in my recordings."

Cybotron

Home recording for the teenager was at this point a fairly primitive operation. Using two Kenwood cassette decks and a Yamaha four-channel mixer, he began making tracks involving tape-to-tape overdubs, first creating white or pink noise beats on the Korg MS-10.

Artist: Model 500
Track: 'No UFOs'
Producer: Juan Atkins
Label: Metroplex
Year: 1985

"You can make a pretty good electronic kick sound using gated pink noise," he says. "White noise you can use for like hi-hats and cymbals and snare. Pink noise you can use for kick drum and toms... y'know, the lower region sounds."

As time went on, Atkins began to experiment with EQ to make his recordings sound clearer and punchier. "I became an expert at doing this," he recalls, "because I realised that during each pass you get a generation. So, my very first recordings, by the time I got to the end, man, the drums were so buried in static and muddy sounding. I had to learn how to tweak all of the bass out. So by the time it got to that last generation, it would even out. With each pass, you'd just add a little bit more bass, and a little bit more bass, until the last pass, it was like, you can leave it sort of flat and natural."

During his days at Belleville High School, Atkins became friends with another two future Detroit techno innovators, Derrick May and Kevin Saunderson. "Belleville was in the suburbs, and it was only maybe 10 percent Black people in the whole school," he says. "So you automatically bonded. Especially at lunchtime, because all the Black kids sat at one table. Derrick and Kevin were in my brother's class, and they would come by the house to visit my brother. So, somehow me and Derrick and Kevin became friends."

Although the three didn't make music together at this time, they would study records by Prince, Yellow Magic Orchestra and, crucially, Kraftwerk. They were further inspired by DJ Charles 'The Electrifying Mojo' Johnson's shows on various Detroit radio stations, titled Midnight Funk Association but involving genre-blurring playlists featuring anyone from Parliament to Tangerine Dream, Giorgio Moroder to Devo.

But it wasn't until Juan Atkins enrolled at Washtenaw Community College in Ann Arbor, Michigan, and played his home-recorded tapes to a fellow student, Rik Davis, that he chanced upon his first key collaborator. »

Juan Atkins' Metroplex Soundworks studio was based around a Tascam eight-track tape recorder and mixing desk.

"I met Rik as a fellow synthesist," he remembers. "He was doing the same thing that I was doing, like a one-man band, so to speak. And then when he heard the demo, he invited me to come over and jam with him."

The first time he visited Davis at his home studio, Atkins was amazed by how much music tech his classmate owned. "I went over, with my little Korg MS-10 under my arm," he remembers. "I walked into his studio, which was a bedroom in a two-bedroom apartment. He had the lights basically off, shades down. So all I could see was the LEDs coming off the keyboards and all of the sequencers and drum machines. I thought I was going into a new dimension or something, man. I thought I was in the cockpit of an alien spacecraft.

"He was a lot more advanced," he adds. "This guy had a Boss DR-55 [*Dr Rhythm*]. You could create your own drum pattern. It was basically what I was creating with the pink and white noise, but it was encapsulated in actual hardware now. Then he had the [*Roland*] RS-09, which was a synth that was basically all string programs. And he was using a lot of ARP stuff... the Odyssey, the Axxe.

"I was a bit intimidated," he admits. "But y'know, I kind of held my ground. And so we started playing."

The first track the pair created together was the electronic, funky 'Cosmic Raindance' that in 1981 became the B-side of their first release as Cybotron, on their Deep Space Records label. The A-side, meanwhile, was 'Alleys Of Your Mind', a starker, Kraftwerk-styled synth track. The single went on to sell more than 10,000 copies in the Detroit area alone and many people believed it was a European record. Part of the track's appeal was that it sounded machine-tight in its beats yet had a human looseness to its synth parts. "'Alleys Of Your Mind' I played by hand," Atkins remembers. "That bom-bom-bom bassline."

Cybotron's 'Clear', released in 1983, pushed the duo further in the direction of minimalist electro, while clearly being inspired by Afrika Bambaataa's 'Planet Rock'. It featured syncopated rhythms, pitch-shifted and flanged vocals and proved to be another ground-breaking track.

"The way we did Cybotron," Atkins says, "it wasn't like we were living, sleeping, eating together or whatever, y'know. I'd go home for a while and come back. Sometimes we had time to make our own [*tracks*], like, do our own thing. 'Clear', for instance, was a song that when I came into the studio one day, I said, 'Hey, Rik, take a listen to what I've been working on.' And *vice versa*, he'd do the same thing too."

By this point, Atkins had learned to programme bass lines using the step sequencer in the Sequential Circuits Pro-One triggered from a drum machine. It was a fairly random exercise, however.

"You could trigger the notes by the accent from the drum machine into an audio input trigger," Atkins says. "So I would put the rhythm in and then I would play the notes. That bass line on 'Clear' was done after testing about three different types of note sequences.

It's only four notes actually sequenced. But when you flipped the switch on, they kind of randomly came out. So there was different note patterns that I put in, and then when I hit it and the 'Clear' bass line came up, I said, 'OK, yeah, that's it.'"

A lot of the time Atkins and Davis would rehearse their tracks together live at the latter's home, recording rough versions on cassette, before booking studio time to tape them using more pro equipment. "It was laid down in stereo," Atkins says, "with a professional two-track, on a Studer or something."

In addition to his studio activities, Atkins was DJ'ing alongside Derrick May and Kevin Saunderson as Deep Space Soundworks. During this period, they began using a Roland TR-808 programmed with their own beats to transition between records.

"When drum machines were introduced," he remembers, "you had people that would release straight-up rhythm track records, with a drum machine playing their own custom patterns. So it made sense as a DJ to take this drum machine out and play homemade patterns and rhythms to segue between certain records.

"The thing is there's an art to mixing," he stresses. "It's not about just taking two records and slamming them into each other. You gotta find a record that fits the previous record, beat-wise. So sometimes a raw rhythm pattern would be a great way to segue between different segments of your set."

Metroplex Soundworks

Juan Atkins and Rik Davis parted ways as Cybotron after one album, 1983's *Enter*, when the latter became more interested in pursuing a rockier musical direction. Before long, working alone, Atkins discovered the MSQ-700 MIDI sequencer, launched by Roland in 1984.

At this point, he built his own studio, Metroplex Soundworks, in the basement of his grandmother's house, equipped with a Tascam eight-track tape machine and 16-channel board, along with other synths including a Sequential Six-Trak. One of the first results was 'No UFOs', an Atkins solo track that was to be his first release as Model 500.

For the track's beats, using MIDI, Atkins chained together a Roland TR-909 and Sequential Circuits DrumTraks. "When I'd MIDI out of the DrumTraks and MIDI into the 909, it was a perfect sync," he says. "The only thing about MIDI in the early days was that... like if I started the DrumTraks, the 909 didn't start automatically. You'd have to make it start. And sometimes whether [*the patterns*] played in sync or not depended on where you started."

Atkins meanwhile used the MSQ-700 to trigger the Six-Trak for the 'No UFOs' bass line and added stab and drone effects played on his Pro-One, fed through his Lexicon PCM-60 digital reverb unit. He then added his part-spoken, part-sung vocal, recorded using a Sennheiser MD 421, before capturing parts of it for distinctive 'freeze' effects using a Lexicon delay. In the final stage, he mixed the track down to a Tascam 22-2 two-track, domestic, seven-inch-reel tape recorder.

"It wasn't like a professional machine," says Atkins. "It was that little one that you could buy from Tascam, and I mixed it down to that. 'No UFOs' was a total home recording."

Released in April 1985 on his Metroplex label, it proved to be a hugely inspirational record. Atkins says this didn't take him entirely by surprise. "Well, y'know, you got to remember that this was the first record I released on Metroplex, but I had been doing Cybotron since 1980. So, I sort of had a five-year head start in terms of knowing what the public liked, or what our scene liked... people, y'know, from the ghetto in Detroit, pretty much, in terms of dance music.

"So," he adds, "it was sort of a no brainer. The experimental phase was somewhat over with by that time."

Techno! The New Dance Sound Of Detroit

How 'No UFOs' exploded on the dance scene, particularly in Chicago, was almost by pure luck. Derrick May's parents had moved out west to the Illinois city and on his visits there, May began checking out the clubs and touting copies of the

Model 500 record. Before long, DJs were spinning the track and it quickly earned the distinction of being the only American dance record that influential figures in Chicago, such as Frankie Knuckles at the Power Plant, were playing. The 909 soon became the sound of Chicago house, after May gave Knuckles one of his beatboxes.

"That drum machine was the very drum machine that was used on 'Time To Jack' by Chip E and the first JM Silk record [*1985's 'Music Is The Key'*]. So they passed along that 909. Just like we used to pass around stuff in Detroit."

Some of the key gear on 'No UFOs', such as the Roland MSQ-700 sequencer and Sequential Six-Trak and DrumTraks.

By 1987, May and Saunderson were fast building their own reputations as DJs and musicians: the former producing the landmark dance track 'Strings Of Life' as Rhythim Is Rhythim; the latter hooking up with singer Paris Grey in Inner City, whose 'Big Fun' would go on to be an enormous hit the following year. Along with Atkins, the trio earned the nickname the Belleville Three, after their old high school.

Then, in 1988, May met Neil Rushton, of the British Virgin Records dance subsidiary label, 10 Records. "By Derrick's association with Chicago and Neil Rushton's fascination with Chicago," Atkins explains, "they decided to try to curate this compilation." The working title for the various artists album was however named after their home city: *The House Sound Of Detroit*. Until Atkins objected.

"I said, 'Well, no, this ain't the house sound of Detroit'," he laughs. "This is techno music." Atkins had first used the term back in Cybotron with their 1984 track, 'Techno City'. It had been inspired, like their band name, by futurist author Alvin Toffler's books, *Future Shock* (1970) and *The Third Wave* (1980).

"*Future Shock* was the reference book that we used in the course in high school called Future Studies," says Atkins. "And this course actually set me up for a lot of the lingo and things that I used, or me and Rik used. Even the term 'Metroplex' comes from the word 'Metrocomplex', which is a sort of a mesh of the words metropolitan and complex."

"I said, 'Well, no, this ain't the house sound of Detroit.' This is techno music."

Toffler also wrote about 'techno rebels', and in 1988, Atkins felt it was the perfect word to describe this new, modernist sound. He delivered a new track, called 'Techno Music', to 10 Records, for inclusion on the new compilation. "After I submitted that track," he says now, "they decided to change the name of the album."

May, Saunderson and Rushton agreed to the record being renamed *Techno! The New Dance Sound Of Detroit*, and a musical genre — for which Juan Atkins was largely responsible — was born. A theory soon emerged that Detroit being an industrial city had inspired the sound, with its machine-like rhythms.

Atkins only partly buys into this. "Yeah, I think for any musician playing any style of music, your surroundings subconsciously affect your music. But it's more subconscious than actually conscious. I mean, I wasn't looking at the skyline, saying, 'OK, this music fits', y'know. But I believe, definitely, it's just sort of a product of your surroundings on a subconscious level."

Collaborations

When techno began to boom, particularly in the UK, a plan was hatched by British record producer and then-owner of ZTT Records, Trevor Horn, for the Belleville Three to unite as a group, named Intellex. Horn apparently envisaged the trio as a "Black Pet Shop Boys". But the project quickly floundered when Derrick May told Horn he wouldn't appear on *Top Of The Pops* to promote the band's records.

making music are making music. But that's probably the only real downside. It's too easy [*laughs*]."

Perhaps that's one of the reasons why, when Juan Atkins performs live shows as Model 500 these days, he's returned to hardware synths. "We revamped the show to actually play all of the sounds directly from the machines," he says. "Instead of recording the sounds into Ableton or something. Because basically what we were doing before was recording a lot of the bare tracks. You'd hear certain tracks like [*1995's*] 'Starlight' or whatever, where we would play different sounds over the top of what was the original track. But now everything is done from top to bottom with live machines."

The Future

In 2018, both Moodymann and Luciano created new remixes of 'No UFOs', released on Metroplex. The former, similarly Detroit-based producer took the track on a journey involving the sound of waves and an original era TR-606 Drumatix. The latter Swiss/Chilean DJ reimagined it as minimalist house. "Moodymann's mix was quite interesting," says Atkins. "Luciano's mix was Luciano. I mean, it fit in his crowd or his type of set. But, y'know, my all-time favourite was the original mix."

As to whether Juan Atkins can hear the influence of 'No UFOs' in modern dance music, he remains characteristically modest.

"Music is constantly evolving," he states. "However, yeah, a lot of stuff I hear now, I can hear a direct correlation to what I was doing in the early days to where it is now. So, I guess you could say that was the pioneer, so to speak."

Finally, then, having had a vision of music's future back in the 1980s, can Juan Atkins predict where dance music might be headed?

"Don't ask me what's coming next," he laughs, "because I don't know what's coming next."

"Yeah, uh," Atkins laughs, "Derrick is a great talker. God bless him. The same way he talked us into the deal, he talked us out of the deal. I mean, I can understand in a way, but you don't go in and say to the managing director, 'We're not going to do *Top Of The Pops*.' And then they said, 'Well, we can't do the project. How are we going to push the record?'

"The unfortunate thing is that he didn't discuss this with anybody before he went in and told the guy we wasn't gonna do it. So basically, the whole deal fell apart based on that statement."

Down the years, Atkins continued to make his own singles as Model 500, before his first album under that name, *Deep Space*, arrived in 1995. Along the way, he collaborated with Doug Craig as Channel One and in 2013 hooked up with Moritz von Oswald as Borderland. Then, last year, he revived the Cybotron name for the *Maintain The Golden Ratio* EP, although without Rik Davis.

"We couldn't work it out to actually get back together to do this thing," he says. "And so, he gave me his blessing to go forth with the name. Because Cybotron is iconic, y'know, and I mean, of course, it's synonymous with the birth of techno, in a way. There was so much public interest in releasing new Cybotron music."

These days, Atkins works in the box, Ableton being his DAW of choice. He names Spectrasonics' Omnisphere as his favourite soft synth. "I use a lot of plug-ins," he says, "because I've always been an advocate of technology, of progress. So, y'know, I'm not gonna sit here and say, 'Hey, I'm an advocate of change and progress', but still be saying, 'Oh, I'm just gonna use a Pro-One for the rest of my career.'

"Now definitely, it's a lot easier. Because, man, MIDI cables that don't work... and you're like an hour or two into your project. That's definitely a change that I welcome, where everything in the box is making the MIDI connections. The only downside I would say is that now it's so easy that a lot of people that probably shouldn't be

How I Got THAT SOUND

Torsten Kinsella • God Is An Astronaut 'Falling Leaves'

JOE MATERA

Since forming in 2002, Irish instrumental outfit God Is An Astronaut have developed a unique sound that incorporates ambient electronica, post-rock guitars, cinematic soundscapes and captivating melodies. Asked about the origins of a favourite sound, God Is An Astronaut's Torsten Kinsella details how he achieved the layered rhythm guitar tones on 'Falling Leaves', from their most recent album *Embers*.

"For the rhythm guitar sound on 'Falling Leaves', I used a 1969 Marshall Super Bass with a 1971 Orange 4x12 cabinet, featuring Heritage and Vintage 30 speakers in an X-pattern. I used four microphones: a Sennheiser 421-U, which I find less harsh than newer models, and two older Shure SM57s, the Unidyne III, which sound more melodic and sweeter than modern SM57s. Additionally, I used the Royer Labs 121 ribbon microphone to add weight.

"I placed the mics close together on the Celestion Heritage speaker to minimise phase issues. For the Celestion Vintage 30 speaker, I used another Unidyne III Shure SM57 and a Royer 121 ribbon mic. I recorded the sound with a vintage 1963 Fender Jaguar. I repeated the process with a Fender Jazzmaster, but used a 1974 Orange 4x12 cabinet and a 1973 Marshall Super Tremolo, with a different Sennheiser 421-U and a Royer 122 instead of the Royer 121, while keeping the same SM57s."

Perfect Phase

"Each amp/cab recording — four channels each — had the Sound Radix Auto-Align plug-in inserted. Although the phase was good, this plug-in optimised it further, perfectly aligning the four mics, which I blended to create a cohesive guitar sound. I repeated the process with another guitar and panned both guitars hard left and right.

"When I mono'ed the full guitar mix, it lost some level compared to the stereo field. To fix this, I used a Hiwatt DR103 with a Hiwatt 4x12 cab with Fane speakers down the centre, using a Sennheiser 421-U and a Unidyne III SM57 with another Fender Jazzmaster. I mixed the Hiwatt just loud enough to prevent an obvious volume decrease in mono, though it still felt lacking.

"I used another Sound Radix plug-in called Pi to improve the phase relationships between the different amps, as they were slightly detracting from each other. I placed the Pi plug-in at the end of each amp group in Pro Tools and assigned them to the same internal group in Pi. When I mono'ed it, the phase relationship was near perfect, resulting in a more coherent and full sound.

"I also considered how it would sound on smaller devices, using a Crane Song Phoenix II plug-in with the Radiant setting to enhance the mid tone, giving it a fuller harmonic sound. The entire guitar mix went through a Chandler Curve Bender and the Chandler Zener Limiter, adding a touch of treble and mid, with the Zener Limiter adding colour and cohesion, which really finished the guitar sound."

Hear The Sound

W https://open.spotify.com/track/3zqThRe8zbDcQjLvZAbuHB

W www.youtube.com/watch?v=NoSSmzA7NGs

Sample Libraries

Sample Logic Arpology X

Kontakt Instrument

★★★★★

Sample Logic's original Arpology instrument was reviewed in the September 2014 issue of *SOS*. Ten years down the line, we now have Arpology X, a direct evolution of the concept but taking advantage of a decade's worth of progress in UI design, Kontakt's scripting capabilities, sampling technology and computing power. While aimed firmly at media composers, Arpology X could undoubtedly find its way into any contemporary electronic music context.

The substantial 28GB of sampled instruments spans an impressive breadth of organic and electronic sounds. The blend of up to four of the core sounds can controlled via the X/Y pad but, interestingly, this pad is dual-mode; you can flip it to also modulate the comprehensive effects processing options by blending four different effects presets. Unsurprisingly, arpeggios form an important part of the feature set. Sample Logic have not pulled any punches here; the step-based sequencer/arpeggiator available for each individual sound engine is very deep, including the ability to sequence properties on both of the dual X/Y pad systems. There are also ADSR, filter, tuning and other sound-shaping elements for customising the core sounds in each engine.

The instrument offers a massive selection of presets and a slick, tag-based browser for finding what you need. Frankly, if you never even peeped under the hood, you could spend endless happy (and creative) hours just exploring these. Sonically, Arpology X is top-class. Providing you get as far as applying some hands-on control of the dual X/Y pad, then any of the presets could get you rolling on a new cue; it's inspiring stuff. If you want to experiment without the deep dive, the very clever randomisation options are most certainly worth exploring. You can choose to let these create you a complete preset from scratch or, thanks to a whole palette of choices to target just specific parameters for randomisation, simply generate variations of an existing sound. Either way, there are endless possibilities. All that said, when you are ready, that deeper dive into the modulation and arpeggiation engines is hugely rewarding and will let you fully exploit the considerable potential of the instrument.

Whether it's high-octane rhythmic electronica for action, drama or modern horror, or more subtle movement for a sense of the magical, mysterious or ethereal, Arpology X has something (actually, lots of somethings) to offer. The sounds are of the highest quality and the modulation and arpeggiation makes it a hugely inspiring instrument to play and compose with. Yes, there are plenty of other options for cinematic sound blending but, if you are a media composer always keen for some fresh sonic options — and have the required budget — Sample Logic's Arpology X is most certainly worth exploring. *John Walden*

$299.99

www.samplelogic.com

Heavyocity Oblivion: Aggression Designer

Kontakt Instrument

★★★★★

Heavyocity's Oblivion: Aggression Designer maintains the 'fierce hybrid' style of its Gravity 2 predecessor, unleashing a fresh set of rhythmic pedals, leads and basses, textures, pads, drones, stings, noises, risers and dives. Oblivion (8.39GB installed) requires Kontakt or Kontakt Player 7.10.5 or later and runs on the Gravity 2 sound engine, described at length in my March 2024 *SOS* review.

For this library Heavyocity collaborated with composer and sound designer David Levy, famed for his work on the *Doom Eternal — The Ancient Gods* video game. In their quest for "ruthless power, relentless energy and serious sonic pain", both parties fired up their modular and analogue synths (Moog Modular, Metasonix, Make Noise Strega, Soma Labs Enner, Lyra 8, Pulsar 23), tube-based outboard gear and DSP plug-ins to create a total of 336 unique (and undeniably aggressive) sound sources.

The partnership has generated some excellent four-bar pedal loops. Presets such as the brutal 'Raging Road Rash', the banging, Prodigy-like 'Lazer Face' and the aptly named 'Crusher Waves' are infused with Levy's signature blasting, ripping sound, and the majestic 'State Of Emergency' (which sounds like a dramatic movie cue) and explosive 'Spiked Pinger' carry the percussive force of a rock kick.

In a somewhat less violent vein, the hustling, propulsive 'Double Vision' and 'A Demon Runner' loops provide instant rhythmic action for composers, while synth players will enjoy the big, beautiful Moog-like bass sound of 'Say Goodbye'. Heavyocity's programming virtuosity comes to the fore with the insanely syncopated 'CnC Doom Factory', and 'Factory Electric' (which reminded me of the *Dr Who* theme) makes brilliant use of the Macro Sequencer's vivid tone-shaping. As with Gravity 2, the 72 pedal loops are divided equally between straight 4/4 and triplet-based 12/8 time.

Oblivion's playable basses and leads will be nirvana for those seeking aggressive synth patches. The three-channel bass presets 'Driver Open' and 'Bad Nightmares' are fierce, snarling monstrosities which slam like a barn door, while the gigantic twang of 'Magnetized' and forceful 'Magma Lashes' are epic distorted bass timbres. In the leads department, 'Lead With The BFG' and 'Polymorpher' are both furious, filthy howling rackets designed to spread fear and alarm — a vicarage tea party it is not.

Moving into calmer waters, a set of clean, quiet pads include the breathy, gently swelling 'Fallout', the radiant 'After Mourning' and the lush hybrid orchestral tones of 'The Abandoned Lab'. Underpinning these serene textures are 24 ominous-sounding pitched drones and 36 scary atonal drones, which Heavyocity have concocted into a diabolical cacophony — that's intended as a compliment, by the way.

Rounding off this excellent library is a wildly inventive collection of stings and modular noises comprising impacts, sub thumps, clangs, wild metallic sound effects, pink noise explosions, fizzing crackles and hair-raising electronic screams, a celebration of creative cyber industrial sound design. *Dave Stewart*

$149

www.heavyocity.com

»

JASL-Audio
Striking Strings Glissandi

Kontakt Instrument

★★★★★

Striking Strings Glissandi adopts a fairly sharp focus that is likely to appeal to the professional or semi-pro composer community, who frequently reach for tools with a singular purpose in mind. This is the first unique offering from JASL-Audio, a company presided over by Jochem Weierink, a talented film composer in his own right and clearly equipped with useful insight.

SSG includes just under 6GB of sampled content, divided up into four instrument patches. The first three are similar, represented by a Full String Ensemble, Lower Strings (Bass & Cello) and Upper Strings (Violin & Viola). This is not a Kontakt instrument for virtuosic legato lines, being very much geared toward the common dilemma of creating string risers in a controlled manner, where you want strings to naturally rise to a climax.

This 'riser' element is controlled from the Glissandi fader, located centrally within the Kontakt interface. Play a single note or chord, and fader movement upward will result in a rising pitch. It's worth noting that at it's lowest position the note is centred, but as it reaches its upper limit the tonal centre becomes more ambiguous. This does mean that the resultant climax is perfect horror-fodder, but it is trickier to generate an exacting octave leap.

By default the on-screen Glissando fader is linked to the modulation wheel. It can also be linked and coupled with the Dynamics fader, which increases the climactic quotient considerably. You can also maintain independent control, should you have a box of hardware faders to hand. Highlighting this crescendo even further, the Accent control governs the degree of re-articulation at the point of note release, which beautifully places a pin on the end of any Glissandi phrase.

To the right of the instrument window, you dictate the articulation employed at both the beginning and the end of the riser. This is also available through keyswitching in the upper keyboard register, with helpful control of crossfade between the two articulations. This means that you can glide from a tremolo to a pizzicato, or a staccato to a sul pont, although in all cases, the tonal centre is chaotically obscured.

A basic mixer allows selection of signal between spot and Decca Tree miking. The Reverb fader only operates if an initial signal is present, as reverb is generated within the plug-in, rather than recorded.

The fourth and final included instrument adopts a similar ethos to the initial three, but the resultant crescendo is geared toward a tonal cluster. This is another very handy effect which extends the instrument's usefulness further.

SSG offers an incredibly ingenious solution to the age-old problem of controlling risers. The ability to dictate your own speed of rise and transition, along with control of dynamics and release, makes SSG pretty indispensable, should it be an effect that you frequently require. It lends itself perfectly to cinematic and tension scoring, but its focused detailing may exclude more mainstream use. *Dave Gale*

€178.80

www.jasl-audio.com

Vienna Symphonic Library
Synchron Solo Violin & Solo Cello

VSL Synchron Player Instruments

★★★★★

Two new libraries fill a vital gap in VSL's catalogue. After seven years of constant recording in the Synchron Stage had yielded no solo strings, users were beginning to grow impatient. Bowing to public demand, VSL began work on a large sampling project which now bears fruit in the shape of Synchron Solo Violin and Synchron Solo Cello. Both run exclusively on the free Synchron Player and respectively require 25.1GB and 22.6GB of storage space.

The new instruments feature violinist/concertmaster Marina Dimitrova and cellist Florian Eggner, principal players of the Synchron Stage Orchestra with extensive classical, film music and sample recording experience. In the interests of musical consistency, sessions (described as "heavy", "exhausting" and "insanely intense") were scheduled so that the players could execute the same articulations on the same day.

This Olympian sampling marathon has generated an encyclopaedic articulation list. Short-note styles include spiccato, staccato and détaché with a choice of bold and agile attacks, ricochet, saltando, a lighter 'performance spiccato' option and 'saltando glissandos', which combine a bouncing bow stroke with up and down slides. In addition to their normal, short, soft and no-vibrato deliveries, the détachés also feature graceful up and down tone slides and despairing semitone falls.

Long notes are played with normal, fast and soft attacks with no vibrato, poco (light) and molto (strong) vibrato; for note offs, you can choose between a normal release tail, the abovementioned comic falls and animated, ear-catching octave slides. Also included are regular and snap pizzicatos, col legnos, tremolos, fast 16th-note measured tremolos, trills, harmonics, crescendos, diminuendos and a furious, edgy sforzato artic that will easily cut through a cinematic drum battalion, all offering multiple performance variants.

The solo strings' performance legatos sound absolutely sensational. Played over a G3-E7 range, the violin legatos epitomise the instrument's expressive powers, whether playing soaring melodic themes, high-speed runs or quiet reflective passages. The cello (C2-A5) is no less emotive and adaptable, sounding equally great on passionate fast arpeggios and slow sombre melodies. I found the instruments' agile/fast attack legato to be beautifully smooth and responsive, while the evocative, slower-moving lyrical style, exuberant Bollywood-esque portamentos and Zigane semitone slides add further creative options.

Two new features are worthy of note: Sequence Control switches between a series of different articulations on each key press (great for rhythmic spiccato and staccato ostinatos), while Slot XFade lets you adjust the split point between playing styles (such as no vibrato and poco vibrato). The libraries are available in Standard and Full versions: the samples are identical, but the latter contain more microphone positions.

Immaculately played, sonically rich, stunningly realistic and superbly playable, these solo strings represent the pinnacle of VSL's sampling achievements — 20 years of experience distilled into two remarkable virtual instruments. *Dave Stewart*

Standard Library €220, Full Library €310

www.vsl.co.at

Audio examples of this month's libraries are available at www.soundonsound.com.

SOS SHOWCASE

David Carson (Regional Sales Associate): david.carson@soundonsound.com

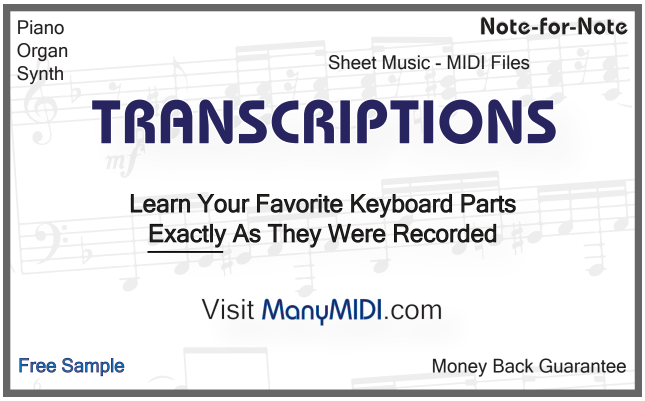

AD INDEX

To Advertise in ***Sound On Sound*** please contact: **usadsales@soundonsound.com**

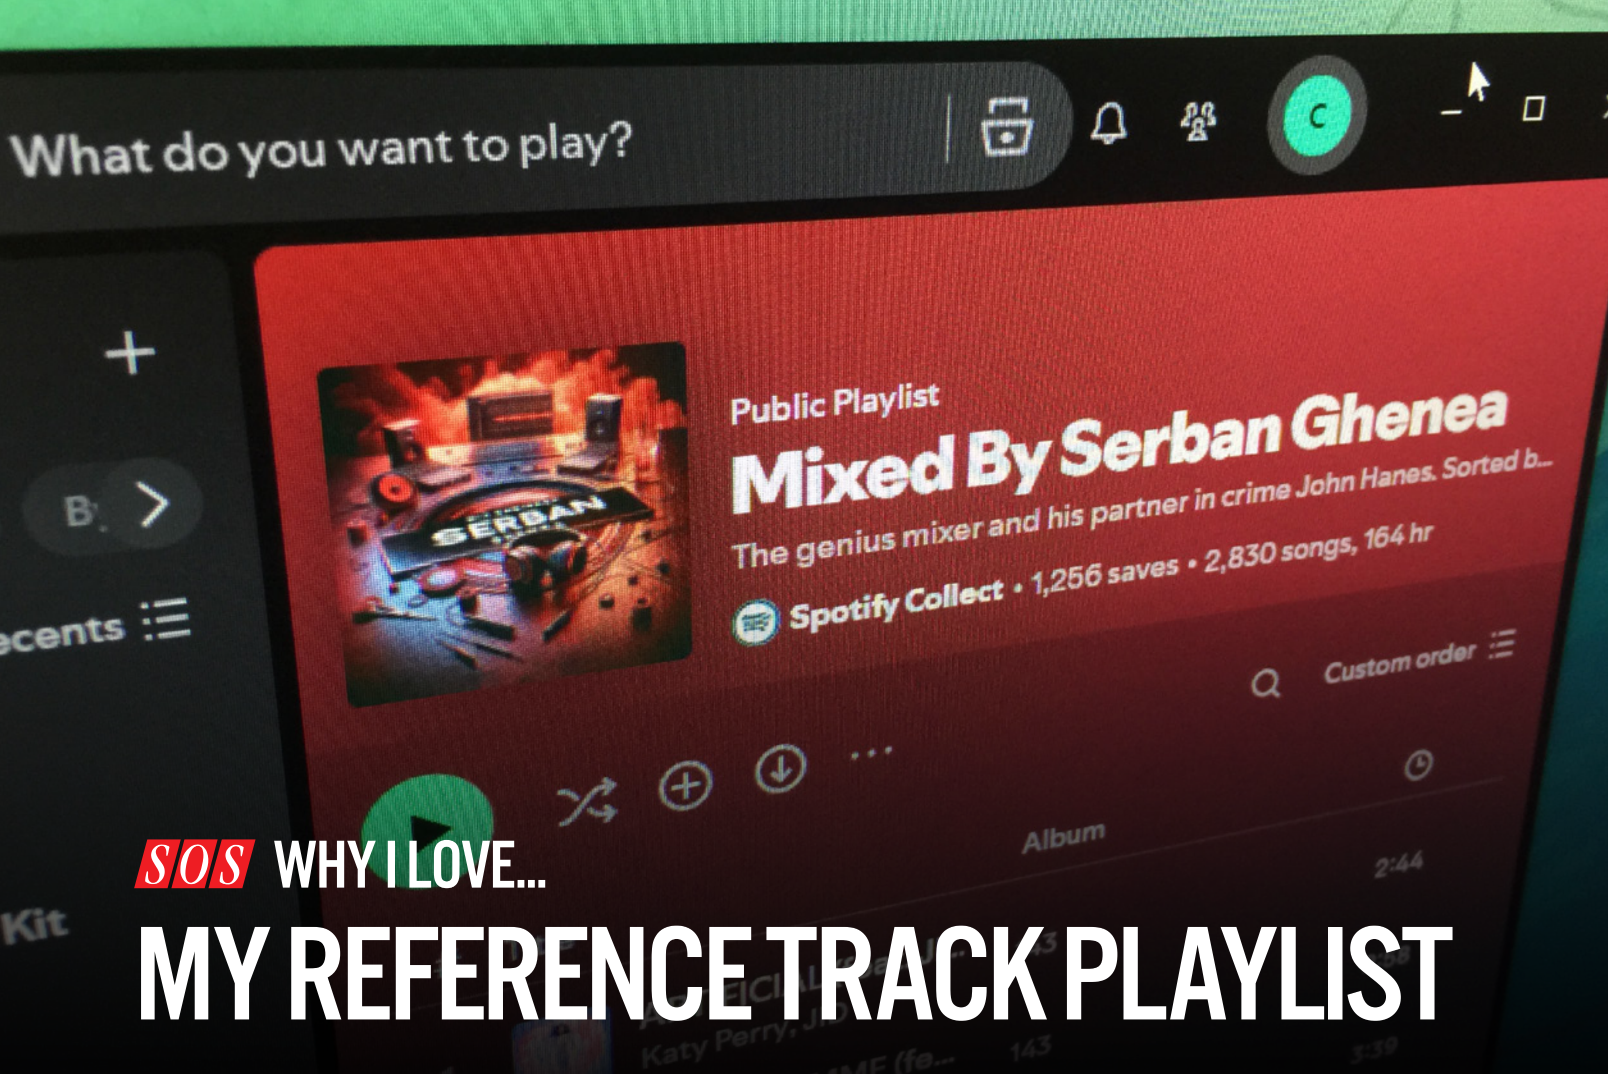

SOS WHY I LOVE...

MY REFERENCE TRACK PLAYLIST

MATT HOUGHTON

What first attracted me to the idea of recording and mixing wasn't just a love of music (which could just as easily lead you to be a songwriter or a fan), but a deep-rooted desire to know why certain records got me up on my feet and dancing, or lying in a dreamlike state, immersed in a wash of sound that made the hairs on my neck tingle. Some of the sounds that triggered such reactions seem laughably simple to me now, but at the time they invoked such a sense of awe that I could only imagine sorcery was involved. Just how did they do it — and, more importantly, how can I?

From my early teens, I learned to play synths and guitars, and acquired technology I hoped would enable me to hear, understand and replicate the sounds hitting my ears. I developed recording, editing and mixing skills that allowed me to bend sounds to my will, and blend them to taste.

I now love making and shaping music every bit as much as I used to enjoy listening to my favourite artists. But the heady mix of 'magic' and ignorance that made so many records so enticing has, of course, largely given way to science and understanding. I wouldn't turn the clock back for a moment, but there's a tinge of sadness — who wouldn't want to experience that again?

So I reckon it's hugely important to work at keeping that feeling alive, and it's one reason I love curating my reference track playlist. For years, I've kept a library of reference tracks, and while once it was all on CD, hindsight tells me I was too snobby about streaming services for a bit too long. Today, while I use full-quality references for certain things when mixing, I now use reference tracks mostly to keep my ears accustomed to my speakers, room and headphones, or to gauge the character of new gear, and increasingly I've been relying on a well-known streaming app on my phone. My monitor controller has a Bluetooth receiver (streaming and Bluetooth, I know!), so I can whip out the phone and play any number of tracks straight away. Away from the studio I can still listen on the same headphones, which is great.

A happy side-effect is that it's dead easy to spend time reappraising and updating my reference track playlist, a process to which I dedicate a half-day once every couple of months. I'll seek out new mixes that I could consider exemplars in some respect, and this sort of active listening really is a joy. Sometimes, I might also spend a while figuring out how a specific producer's sound has evolved over the last few years. On the most recent occasion I spent a good few hours travelling down a Serban Ghenea-shaped rabbithole, listening to the most wonderfully present, yet deep and spacious pop productions. Just how did he do that — and, more importantly, how can I?

"I'll seek out new mixes that I could consider exemplars in some respect, and this sort of active listening really is a joy."


AUTO-TUNE®
PRO
VERSION 11
AUTO-TUNE® PRO
THE PROFESSIONAL STANDARD FOR PITCH CORRECTION
The iconic Auto-Tune effect with 4-Part Harmony Player, ARA2 support, improved Graph Mode, and more.
AUTOTUNE.COM
NOMINEE
SOS AWARDS 2025
SCAN ME