THE WORLD'S BEST RECORDING TECHNOLOGY MAGAZINE 1985 – 2023

SOUND ON SOUND™

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

EXCLUSIVE FIRST REVIEW

RODE NT1 5TH GENERATION

THE 32-BIT MIC: NEVER MAKE CLIPPED RECORDINGS AGAIN

Groove Synthesis 3rd Wave

Wavetable wonder synth

Recording & mixing SZA

The story behind her hit comeback album

Latest CPUs benchmarked

Have Intel recaptured the crown?

WIN

McDSP EVERYTHING PACK

WORTH $2999

REVIEWS: POLYEND / ADAM / AUSTRIAN AUDIO / AUDIENT / TWO NOTES / DANGEROUS / TERRY AUDIO

TECHNIQUE: HANDS-ON SOUND DESIGN / MIX RESCUE / DAW WORKSHOPS

www.soundonsound.com
03
0 29074 05361 3
USA $11.99 / Canada C$12.75

SOFT MACHINES

Machine learning is a ticking time bomb underneath the arts. Tools such as DALL-E and Stable Diffusion create visual artwork that is sometimes uncanny and often frighteningly convincing. ChatGPT will generate short-form text that is distinguishable from real human writing only by its perfect spelling. It can even code.

The challenges are greater with audio and video content, but they're not fundamentally different. It's only a matter of time before AI can produce music that 'passes'. And it won't be created using layers of overdubs, by mucking about with MIDI, or in any other way we can observe and make sense of. It will simply emerge fully formed in response to our polite inquiries.

Does this mean the end for music as a human career? No. Music differs from writing, painting and photography in that live performance is fundamental to it. An artificial intelligence is pretty much the definitive studio artist. The parts of our job that involve playing 'Mr Brightside' to a handful of uninterested punters in the Dog & Duck on a Friday night are safe — for now.

More fundamentally, the creativity of machine learning is essentially synthetic. The scope of an AI is set by the body of work on which it's trained. An AI trained entirely on symphonic classical music would never be able to come up with rock & roll, and *vice versa*. Machine learning is fantastic for creating pastiches, and as such, is a real threat to composers of library music, for example. But it's debatable whether it could anticipate or drive the emergence of new musical genres, at least without extensive human input.

And this is something that the alarmist view of machine learning tends to overlook. What you get from an AI is highly dependent on what you put in. Crafting the prompts that will deliver what you're looking for is a different skill set from writing and recording music the old-fashioned way, but it's a skill nonetheless. Is it optimistic to think that people who've spent their lives working in music might be able to interact with musical AI more effectively than others? And that this could open up a new income stream for some? Time will tell.

Before AI music itself becomes big, though, it's a safe bet that you'll see plenty of AI-generated writing about music. And in a world where anyone can create a website stuffed full of artificially created 'reviews' or 'tutorials', the value of an established and trusted source like *SOS* will be even more apparent. Machine learning can do amazing things, but it can't get a new synthesizer or microphone out of its box, plug it in and put it through its paces. And as long as we are still making music the human way, human experience really matters!

Sam Inglis
Editor In Chief

"Machine learning is a ticking time bomb underneath the arts."

SOUND ON SOUND

SOUND ON SOUND NORTH AMERICA
1083 VINE STREET #727
HEALDSBURG, CA 95448, USA

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
Head Of Design George Nicholson Hart
Designer Alan Edwards
Designer Andy Baldwin

ADMINISTRATION
admin@soundonsound.com
Managing Director/Chairman Ian Gilby
Editorial Director Dave Lockwood
Sales Director Robert Cottee
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
paul.dacruz@soundonsound.com
Tel: 707-569-6021
Advertising Manager Paul DaCruz
Regional Sales Associate David Carson

MARKETING
marketing@soundonsound.com
Business Development Manager
Nick Humbert

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

ISSN 1473-5326
A Member of the
SOS Publications Group

CLASSIC. MODERN. EFF

MASSIVE WAVETABLE SOUND |

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

GROOVE
SYNTHESIS
groovesynthesis.com

FORTLESSLY VERSATILE.

REAL WARMTH AND CHARACTER

ers · built-in wavemaker · arp+song sequencer · knob per function

3RD WAVE
ADVANCED WAVETABLE SYNTH

IN THIS ISSUE

www.soundonsound.com

March 2023 / issue 5 / **volume 38**

FEATURES

62 GROOVE SYNTHESIS 3RD WAVE

ON TEST

WORKSHOPS

FOCAL®

HI-FI HEADPHONES I CLEAR MG PROFESSIONAL

The Clear Mg Professional open-back headphones are aimed at all music creators looking for a reliable tool in terms of quality and transfer speed. Thanks to their high-performance magnesium drivers, they offer exceptional dynamics, combined with a capacity to reproduce the tiniest details across the full sound spectrum: you totally forget you're wearing the headphones, as they offer the same listening sensation as the best monitors.

Expect the best

MADE IN FRANCE

$1 499

LISTEN PROFESSIONAL

Listen Professional circum-aural closed-back headphones are the essential work tool for music production. Sound engineers have total control over their work, whether they're in the studio or on a business trip. Besides their undisputable acoustic performance, they also provide maximum comfort.

Listen to your music, not to your headphones!

$299

Focal.com | Distribution by Focal Naim America Contact us for more information : 1 800.663.9352 / info@focal-naim.com


Depth
FILTER
LP 12dB Edgy
Cutoff
Res
Drive
Mix
COMPRESSOR
Tune
Amount
Type
Limiter
Input
FOCAL
FOCAL
FOCAL
Clear Mg Professional
apollo
PREAMP
Neve
Neve
UNIVERSAL AUDIO

Austrian Audio OC7 & OD5

SAM INGLIS

Instrument Microphones

Austrian Audio's new mics are intended to turn heads!

Microphones with the same operating principle tend to share the same form factor. Thus, whereas large-diaphragm mics are typically addressed from the side, moving-coil dynamic and small-diaphragm capacitor mics are usually end-address types. But it's not impossible to do things differently, and side-address designs like the Sennheiser e906, Shure Beta 181 and Josephson e22s have practical advantages, especially on live stages where space is an issue.

With the OC7 and OD5, Austrian Audio have gone further down this road. In essence, they've taken the small-diaphragm capacitor capsule from their CC8 pencil mic and the active moving-coil element from their OD505 stage vocal mic, tweaked the tuning for instrument applications, and adapted them to work in a side-address format. The really clever part, though, is that these side-address housings are swivel-mounted, allowing them to be rotated through 220 degrees independently of the rest of the mic.

Two Shades Of Grey

Externally, the only thing that differentiates the two mics is the colour scheme. The OD5 is black, whilst the OC7 is a sort of mid grey. They are shipped in the same padded foam cases that are used for other Austrian Audio mics, which offer good protection at the cost of taking up more space than would seem entirely necessary. Like most mics, they deliver their output through a male XLR connector mounted in the base of the 'stalk'. This output is balanced and transformerless. The stalk houses the electronics — both mics require phantom power, as the OD5 has active circuitry — and hosts recessed slide switches for a 10dB pad and a second-order high-pass filter with two turnover options. (These are 40 and 80 Hz on the OC7, and 80 or 120 Hz on the OD5.)

Austrian Audio aren't the first company to implement a swivelling headbasket. It's an idea that goes back a long way with manufacturers like Schoeps, while Blue Microphones have made several models with rotating heads, and the Electro-Voice ND44 and ND46 are direct competitors for the OD5. However, Austrian Audio's implementation is distinctive. The stalk of the microphones terminates in a curved arm, to which the capsule housing is attached at a single point. With the thumbwheel slackened off, the housing »

can rotate freely, and when it comes to tightening the thumbwheel, fingers alone easily provide sufficient torque to lock the capsule firmly in position. In short, the system works as advertised, and has the advantages you'd expect. The angle of the capsule can be freely and precisely set with no need to adjust the mic stand. Perhaps a bigger plus is that there's never any need to slacken or re-dress the cable, which can always be run directly down the boom arm of the mic stand.

The 220-degree arc of rotation is more than you'll need for any conceivable circumstance, but is limited to a single plane. It allows the capsule to 'look' up or down, but not to one side. As such, it's actually quite similar to something like the yoke mount for the Shure SM7B, albeit much more compact. However, unlike the SM7B, the OC7 and OD5 don't have an integral mic stand attachment: they're held by a conventional mic clip, which is cleverly shaped so as not to block access to the filter and pad switches in normal use. This grips the mic by friction alone, so the mic can be rotated left-right within the clip if you wish.

Spinning Top

The idea of the OD5 and OC7 is to offer studio-quality performance in a form factor that is more convenient for instrument miking than a typical end-address mic. And on the form-factor side, I think Austrian Audio have fulfilled that brief pretty well. For example, it can be a challenge to position a close mic on a snare drum so that its tail doesn't get in the way of the hi-hat or rack tom. That's much less of an issue with the OD5 and OC7 because you can keep the 'stalk' parallel to the ground and rotate the capsule to look at your preferred bit of drum head. Positioning within a bass drum through a hole on the front skin is likewise easy.

Depending on the angle of rotation, the front-to-back depth of the OD5 and OC7 is typically around two inches — about twice that of the e906, but much less than any end-address mic. In terms of their cross-section, however, these are actually not particularly compact mics, and in fact they're not a great deal narrower than an SM7. In most contexts this isn't an issue, but it can make them visually prominent, and makes them less than perfect for using alongside a second mic in situations where you want to blend the sound of coincident mics on a guitar amp or snare drum. If you want the ultimate in discreet compactness, you'd be better off looking at 'bug' mics with clip-on mounts.

Both mics sport the same 10dB pad and 80/120 Hz high-pass filter options.

Given that both models are primarily intended for close-miking, specifications such as self-noise are arguably less important than the ability to accept humungous sound pressure levels, and neither the OD5 nor the OC7 disappoints on this front. Both can handle over 150dB SLP even without the pad switched in, and more than 160dB with. Sensitivity is a sensible 10mV/Pa for the OC7 and 6.7 for the OD5, so neither will cook your preamps when confronted with a loud source. And although self-noise is, as mentioned, a rather secondary consideration in mics of this type, figures of 19dBA for the OC7 and 21 for the OD5 don't rule them out from being used in distant-miking applications at a pinch.

Five Or Seven?

So, do the OD5 and OC7 offer studio-quality performance to complement their innovative ergonomics? In a word, yes — but they sound quite different from each other. The OC7 uses the OCC7 capsule that was developed for Austrian Audio's CC8 pencil mic and which, in turn, was derived from the classic AKG CK1 design. In the OC7 it sounds much as it does in the CC8: clean, neutral and extended, if perhaps a hair darker owing to the shadowing effect of the housing. Lacking any sort of intrinsic push in the midrange, it can seem quite plain and even flat on some sources, but a capsule of this fundamental quality responds very well to EQ, making it easy to add some excitement back in if the recorded sound is too plain. It's a blank canvas, capturing the sound of the source in natural fashion or allowing a character to be sculpted by the engineer. And, as you'd expect from a small-diaphragm capacitor mic, it has a well-developed cardioid pattern and a nice clean off-axis response.

By contrast, the OD5 is related to Austrian Audio's OD505 handheld stage vocal mic, which has a 'mix ready', forward tone. That's carried over to the instrument mic, with a broad presence boost ephasising the upper midrange from perhaps 2 to 6 kHz. When used close up you also get plenty of proximity effect, so the overall sound presents a very effective balance of low-mid warmth against upper-mid bite and edge. It's a sound that often requires little additional EQ on sources like drums or amps, though equally, if the source itself is a bit fierce and toppy, it might all be too much! The OD5 is not refined or smooth in the way that the OC7 is, but the off-axis response and pattern conformity seem pretty decent for a moving-coil design, and in any case, there wouldn't be much point in having the two mics sound the same.

If you're the kind of person who likes to mic everything up with capacitor mics to achieve a pure, clear, sound, the OC7 will slot right into your collection. If you prefer things a bit more aggressive straight off the bat, the OD5 should tick all your boxes, and as it's an active mic, it shares the OC7's ability to drive long cable runs and its indifference to loading from preamps. And although the ergonomic benefits of the design aren't always relevant, you'll have cause to thank Austrian Audio next time you have to fight your way through a jungle of cymbal stands to reach the snare drum.

summary

The OD5 and OC7 bring ergonomic innovation to the world of instrument mics and offer an appealing sonic contrast, with the OC7 supplying a 'vanilla', neutral tonal quality and the OD5 introducing some rock & roll excitement.

$ OC7 $519, OD5 $299
W austrian.audio

Sweetwater

Music Instruments & Pro Audio

The Brands You Want. Advice You Can Trust.

FREE 2-YEAR WARRANTY*

WE'VE GOT YOUR BACK

At most music stores, customers have to spend an extra 10-30% of an item's cost to make sure their gear is covered beyond the manufacturer's warranty.

You don't have to pay hundreds of dollars — or even more — for extra gear protection at Sweetwater!

Our Total Confidence Coverage warranty gives you 24 months of coverage backed by our team of Sales Engineers, Support Specialists, and our factory-certified Service team at no extra cost.

*Please note: Restrictions may apply. Please visit Sweetwater.com/warranty for complete details.

DEDICATED SALES ENGINEERS

The advice you need to buy with confidence

FAST, FREE SHIPPING

No minimum order

FREE TECH SUPPORT

Got a question? We're here to help

Call (800) 222-4700 or visit Sweetwater.com today!

S-CAT Bass-Synth

Analogue Synthesizer

S-CAT's bespoke Bass-Synth does exactly what it says on the tin...

DAVE GALE

A once-fashionable dance outfit cited that you need to put 'Bass In The Place', but if your synth or bass line lacks that low-end grunt, S-CAT have a bespoke solution that could fill your frequency void.

Along with other products from the S-CAT line-up, the Bass-Synth is made to order. The less-than-oversized device sits neatly on the desktop, with an uncluttered panel that slopes toward you. It's got that Moog-style feel, with chunky knobs, pots and switches that ooze quality. Complete this picture with a chunky pair of wooden end cheeks, and you've got a beautiful piece of analogue synth design.

What You See Is What You Get

Switching it on, the Bass-Synth emitted a very low series of pulses before I clambered to plug in a keyboard, to control pitch and gate. There is no MIDI included with this device, only 1V/octave CV/gate via 3.5mm sockets on the back panel. Despite its analogue makeup, warm-up time is pretty immediate; no more warming up the room before you make music!

The Bass-Synth's sonic journey begins with three switches, which activate square/pulse, sine and saw waveforms. These are enabled or disabled, rather than offering any individual amplitude adjustment, emanating from a Curtis CEM 3340 VCO chip. This is the chip of legends, providing the sonic backbone for instruments such as the SH-101. At the top left of the unit is a Transpose knob, switchable between four octaves, and this is accompanied by a sweepable fine-tune pot, which will shift pitch up by a major third or down by a fourth. There's plenty of volume to call upon. I was operating at just 50 percent of its output, with plenty of signal in evidence.

The VCO feeds a 48dB/octave low-pass filter; it's undeniably rich and smooth, offering a depth of tonal colour which chimes with the build quality of the instrument. When fully open, it's as bright as the sun, but it's the smoothness with which the filter sweeps that creates the tonal magic. In use, I did find that the cutoff was quite lively at the lower end of the sweep, but still manageable, with lower frequencies adopting a real sense of warmth and authority.

Envelopes, Triggers & Modulation

The Bass-Synth offers a trigger switch, which immediately engages the gate signal, and hence the envelope. This is also a simple affair, with only an attack and a decay/release control to call upon, but with a full sustain available. There is also a Latch mode, handy for rumbling drones. The attack and decay phases themselves are relatively short, with just a couple of seconds to call upon, and while they can also be agile, they are far from the snappiest envelopes in the synth world. This feels slightly surprising, given the relative shortness of phase lengths.

There are no modulation options on the panel, with these elements being available via connections on the front or rear of the casing. Four 3.5mm CV connections to the rear allow pulse-width control of the square wave, hard-sync'ing, and frequency/amplitude modulation. You'll need an external source to utilise these, such as a Eurorack LFO. The front panel provides a single quarter-inch jack input, which exploits the tip/ring trick to either control the filter cutoff or input an audio signal, which in turn feeds the filter. This feels a tad cumbersome, not least because two mini-jack connectors for use with Eurorack would have been more practical.

Conclusion

In use, this is a powerful and beautiful sounding device, which sits wonderfully alongside other synths without sub-oscillator credentials. The filter sounds truly beautiful too, but with a lack of resonance control or ability to route the internal envelope to the filter's cutoff it exhibits limitations for anyone wanting more scope for bass-line creation.

summary

The Bass-Synth is a beautifully made, quality product which will happily coexist with many single-oscillator synthesizers or Eurorack systems. Its lack of onboard MIDI may require additional expense for anyone without CV/gate-based equipment.

$ £324.95
W www.spacecataudiotechnologies.com

Two notes ReVolt Guitar & ReVolt Bass

Amp Simulation Pedals

These new all-analogue simulations run a real tube on a high-voltage supply, and cover a useful range of tonalities.

PAUL WHITE

Two notes are well known for their digital amp and cab simulators, but several years ago they released the analogue LE preamp series, and in many ways the new ReVolt amp and speaker emulation pedals can be regarded as their successors. The ReVolt pedals not only feature analogue drive and tone circuits, but also all-analogue speaker emulation. It's a path that's already been trodden successfully by Tech 21's SansAmp range, and a refreshing change from the myriad digital emulations now available. There are two separate models, one intended for electric guitar, the other for bass guitar.

ReVolt Guitar

The Revolt Guitar is a three-channel pedal that incorporates a 12AX7/ECC83 valve in the signal path, and this runs on high-voltage power rails, with a view to ensuring that the pedal delivers an authentic valve tone and responsiveness. The pedal includes send and return jacks plus outputs on both an unbalanced jack and a balanced XLR DI feed. There's also a 3.5mm 'phones output and a 3.5mm aux in jack for use with 'play-along' tracks and the like. Note that the cab sim is applied only to the headphone and XLR outputs, which means that the always 'unsimulated' jack output of the pedal can be plugged into a guitar amplifier — obviously a handy feature for on-stage performances, where you can plug into your amp while sending the speaker-emulated XLR feed to the PA.

Measuring just 19.2 x 11.7 x 7cm this pedal weighs a reassuringly hefty 830g, and its three footswitches are used to select the three amp channels (bypass is achieved by pressing the first two pedals at the same time). Each switch has a different-colour status LED, while the valve, which is visible through a small window, also has a backlight. The pedal requires a 12V DC power source capable of 600mA or above, and a suitable laptop-style PSU with a short mains lead is included. An internal voltage multiplier provides the high voltage (200V) for the valve circuitry. The balanced XLR DI output is protected against the application of phantom power.

Of course, even those of us who use amps and pedals work with DAW software these days, so Two notes have thoughtfully included a lifetime licence for their Torpedo Wall Of Sound plug-in, along with a selection of 10 of their DynIR cabinets, which means you can experiment with different speaker types and virtual mic setups when recording as an alternative to using the built-in analogue simulator. (Further DynIR cabinets are available to purchase.)

The three channels are all based on the characteristics of real amps. Channel 1 is designated American Clean, and it wouldn't be over-reaching to suggest that this is inspired by early Fender amplifiers; it ranges from clean to a warm saturation, but stops short of full-on distortion. Channel 2 is British Crunch so, again, no prizes for guessing what this is designed to emulate — especially when we see the word 'Plexi' in the description. Channel 3, called Modern Lead, is for those who like loads of gain. With its tight lows and harmonically rich distortion, this »

Professional recording gets personal—say hello to **MBOX Studio**™, powered by Pro Tools Studio™. Get everything you need to create, record, edit, and mix music and audio in pristine sound quality. Easily reamp guitars. Work with other DAWs. Connect Bluetooth® devices. Produce podcasts. And more—your creative possibilities are endless.

Ready to sound incredible? Get MBOX Studio now at these retailers:

AM Alto Music · AMS · B&H · Guitar Center · Long & McQuade Musical Instruments · Sam Ash Music Stores · zZounds

DISTRIBUTED BY HAL LEONARD

» channel lends itself well to high-gain rock and metal.

Clearly there isn't space to give each amp its own three-band EQ, but the top-panel control layout for each amp section comprises four knobs. Gain and Volume are provided for all three. American Clean has Bass and Treble EQ controls, while the variable Boost, with its white knob, is located in the British Crunch channel, which just has a Bass control for EQ. Modern Lead has EQ for Mid and for Treble. The Boost, which is adjustable up to 20dB and centred at 1.1kHz, can be used with any channel by pressing the active channel footswitch to toggle it on or off; when active the backlight behind the valve changes from orange to red.

The front panel hosts the two mini-jacks used for the headphones output and aux input, as well as small toggle switches for cab sim on/off and for '4-Cable' pedal connection mode, the latter activating the insert sockets. All the main power and signal connections are located on the rear panel. These include MIDI in and out on mini-jacks (defaulting to channel, bypass and effects loop switching) and effects loop jacks, as well as the main input and output jacks and the DI balanced XLR, which has a ground-lift button. When using the balanced XLR output the maximum signal level available is +11dBu. There are also options for using MIDI to recall channels with a user-defined default effects loop state. An online link for the instructions to achieve this sent me to a 404 error page, but I'm told this page is now live (the pedals weren't yet available for purchase during the review period).

If you plan to use the ReVolt Guitar for headphone practice you'll find the insert points very useful, since any guitar amp models heard dry can tend to sound rather sterile when there's a complete lack of environmental acoustics; add a little ambience reverb or a simulated spring reverb from another pedal and the sound really comes to life. This isn't a problem when playing live, of course, as the venue will have its own acoustic properties, and you may well also be using a pedalboard to add delay or reverb effects. In the studio, there are so many reverb and delay plug-ins available that adding life to the sound is just not an issue.

All three 'amps' have their distinct sound and respond well to playing dynamics, which in turn lends them a comfortable playing feel. Being analogue, there's absolutely no latency to worry about and although the analogue speaker sim doesn't have quite the detail of an IR-based approach, it certainly sounds very believable — and it's sure to keep your live-sound engineer happy too! In the studio, I'd be temped to take a clean feed into the computer and to use the Wall If Sound IR-based plug-in, so as to give me more tonal options. Channel 1 delivers well-balanced and very usable clean tones, with the boost adding some useful 'hair', but I particularly liked the second channel with its 'Milton Keynes' UK vibe, as it can go from a sweet almost-clean tone to a nice bluesy crunch, or you can kick in the boost and go up to classic rock. Channel 3 gets you into Black Sabbath territory and edges towards that 'veins in your teeth' Scandi metal sound.

ReVolt Bass

The ReVolt Bass comes in the same physical format, and again includes the valve stage and offers three channels. The first, Classic Clean, is voiced to deliver the big, solid sound of an Ampeg SVT '76. I found this channel to have an exceptional tonal range, covering everything from a soft and deep Motown thump to twangy funk, and even with the gain right up, it remains relatively clean.

Channel 2 is called Vintage Dirt and is inspired by Marshall's JMP Super Bass MkII, built in the early 1990s. Again, this is a powerful sound but it has the ability to add in a little dirt. It can replicate the sound Jack Bruce obtained from earlier Marshall amps when playing with Cream, or any of the classic rock bands that followed, come to that!

Channel 3 is an original Two notes amp, designed to offer a lot of gain while remaining 'tight'. With the treble and gain up full this produces a very authentic fuzz bass tone, and this benefits from just a little clean sound being blended in, to keep the sound solid. Back off the drive and pull back the treble for a still dirty but less fuzzy sound. (You might want to try this channel if you're in a Stranglers covers band!)

Both pedals feature effects-loop jacks on the rear, to allow connection of other pedals between the amp and speaker emulations.

All three channels have Gain and Volume controls, but again the tone controls are different. The first channel has Bass and Treble; channel 2 has Bass and Mid; and channel 3 just a Treble control, with the remaining knob being used as a Dry/Wet blend for that channel only.

Final Thoughts

As amp replacements, then, both these pedals sound good and clearly have a lot to offer, and you can also use these pedals with an amp, to expand the range of tonal options. To do that you could plug directly into an existing guitar/bass amp that's set to clean, though this approach means you also get the preamp of your 'real' amp colouring the sound, so to make this work you may find that you have to make some top-cut and bass-boost EQ changes. If your amp has one, going into the loop return jack might produce more accurate results — this way, you're just using the amp's power stage and speaker.

The circuitry of both these pedals is creditably quiet, and if you do hear unwanted noise the chances are that it will be coming in via your guitar cable rather than being generated by the pedals' circuitry and, of course, unwanted pickup noise does tend to get louder when you turn up the channel gain.

With sensible pricing, solid build quality, simplicity of operation and appealing sonic performance, both of these ReVolt series pedals do a fine job with the absolute minimum of complication.

summary

These two all-analogue devices run a real preamp tube at high voltage, and do a great job of replicating the sound of a useful handful of guitar and bass amps, whether for use in combination with a real amp or to replace one.

$ $399 each.
W www.two-notes.com

Sonic rEvolution

The Next Generation Summing Console
Featuring two new Analog Custom Color Circuits

2-BUS-XT

SACRIFICE NOTHING™

www.dangerousmusic.com

Isuzi I-M1S

Dynamic Microphone

It may not be esoteric, but sometimes a good-quality mic at a very reasonable price is all you need!

PAUL WHITE

You may not be familiar with the brand name Isuzi, as their main output appears to be stringed instruments and their accessories, but they've recently branched into the pro audio sector with the I-M1S stage microphone. Built in China to reduce costs, it may not bring anything particularly new to the table but it does offer a very decent level of performance at a low price point.

Constructionally there's nothing unusual here — we have a gently tapered cast-alloy body with a thread machined into it at the upper end to receive the screw-on capsule basket. Both the body and basket are finished in satin black, and the foam insert inside the basket can easily be removed for washing. The cardioid-pattern dynamic capsule uses a neodymium magnet, and sits in a fairly rigid rubber insert. It provides a nominal frequency response reaching up to 16kHz (no tolerances are given). Sensitivity is quoted as -51dB, ±3dB (referenced to 0dB =1V/Pa) at 1kHz. A standard three-pin XLR output connector is secured to the bottom end of the mic by a single grub screw in the usual way. A mic cable is included along with a mic clip, thread adaptor and zip-up vinyl storage pouch.

On the side of the mic is an on/off switch. From a sound engineer's perspective, I'm not a fan of switches on live vocal mics as vocalists often turn off their mics, leaving you wondering where all the signal has gone. Fortunately, on the I-M1S, there is a grub screw set into the slider that can be used to lock it in the on position.

All In Hand

Turning to the performer's point of view, the mic feels well balanced and it sits comfortably in the hand. I also like the fact that the manufacturers have resisted the idea of applying one of those non-slip rubber coatings, and instead stuck with a paint finish. All too often those rubber coatings turn into a sticky goo after a couple of years — it feels like you're performing with your hand full of partly chewed wine gums!

My first vocal comparison was with a Prodipe dynamic microphone that is similar in both styling and cost. The sensitivity and tonal balance were very similar, and both benefitted from a low-cut filter on the desk or preamp when miking vocals, as popping could become an issue when used up close. Overall, the vocal sound came across as clear and solid, with a useful but not overdone presence lift. In comparison with a couple of my other dynamic mics, the Isuzi I-M1S had slightly more output and seemingly more low end, which becomes most apparent when close-working, where the proximity effect kicks in. However, that beefy low end comes into its own when miking toms and suchlike, and I also got very usable results using the mic to record my guitar amplifier. In summary, then, the Isuzi I-M1S is a very capable all-rounder for anyone on a tight budget.

"Overall, the vocal sound came across as clear and solid, with a useful but not overdone presence lift."

summary

A more than serviceable microphone for stage vocals, and plenty of other uses too, at a price that's hard to argue with.

$ $79.99
W www.isuzi.co.uk

MusicLab Real Guitar Series 6

Software Instrument

MusicLab's virtual guitars get a major overhaul.

NICK MAGNUS

It's been a while since any MusicLab virtual guitar instrument has been featured in *SOS* — March 2018 to be precise. That review of RealGuitar (the acoustic models) explored its promotion to version 5, a substantial upgrade with many new features. A new Steel String model was added as a separate plug-in to complement the original Standard model. The introduction of Multi Performance Mode was also particularly welcome, allowing the performer much more flexibility when playing chords; notably, melodic lines and internal movement within chords (voice leading) could now be injected in the midst of a strummed chordal part. Since then, all other guitar models in the range (RealStrat, RealRick, RealLPC and RealEight) have received the version 5 treatment, RealEight and RealRick having languished at v1 for some time prior to a brief incarnation as v4. RealStrat also received a shot in the arm — the addition of RealStrat Elite, a separate plug-in featuring a completely new sound set sampled individually from all three pickups, with additional patches simulating double-tracking and 12-string, with options for standard or baritone tuning. As of v5.1 the GUIs were all made resizeable. For those unfamiliar with these remarkably flexible, playable instruments I recommend referring to the 2018 v5 review for a detailed overview of their features (now common to all models): www.soundonsound.com/reviews/musiclab-realguitar-5.

Version 6 remains much the same in terms of core functionality, one notable improvement being found in the Output section: two faders to adjust the relative volume of upstrokes and downstrokes. This may seem trivial, but the ability to adjust these (particularly when reducing the volume of upstrokes) greatly improves the dynamics and realism of strummed chords. So what is significantly new in v6? The answer lies not so much in the instruments themselves, but in what MusicLab have incorporated to make them an all-in-one solution for virtual guitarists.

The Magic Potion

Guitarix may sound like a headbanging Gaulois from Goscinny and Uderzo's *Asterix* comics, but is in fact a Linux-based guitar amp simulator. As a PC user, I had unsurprisingly never come across it, so a little snooping was in order. The program is free, and being open source is under constant development. I scoured the Internet for a PC version, but no, it's resolutely Linux. Nevertheless, MusicLab have cunningly ported it over to work on PC and Mac within their instruments. During the guitars' installation/upgrading process a message pops up asking if you want to install Guitarix — the sensible thing to do is say yes. Guitarix is simultaneously installed as a separate VST3 plug-in, more of which later, to quote Kirsty Wark on *Newsnight*. As seen in the RealEight screenshot, the Guitar Amplifier section occupies the top left of the GUI. By default it's inactive; clicking the 'on' button for the first time invokes a request for permission to access Guitarix — again, no point in saying no.

Guitarix comes with over 200 presets, divided into Acoustic, Electric and Bass categories; the Electric presets are sub-categorised into Clean, Crunch, Distortion and Lead. The three main categories also include a selection of Effects presets. All these presets are available to any MusicLab instrument, so if you fancy distorted acoustic steel string guitar, go for it. Clicking the Edit button opens Guitarix's control panel, from where you can tweak the presets or »

ARTURIA
_The sound explorers

_PolyBrute Noir

Discover the analog abyss

A sleek limited-edition of the Brute flagship emerges.

A morphing analog polysynth, concealed within a stealthy black exterior. A striking new finish; luxurious dark wooden appointments; signature ionized parts.

Strictly limited availability.

Guitarix's modules can be resized to reveal or hide their controls by dragging their headers up or down. Click a module's '+' icon to insert a new one below it, or the '-' icon to delete. To bypass a module, click the 'tick' at the far left of its header. The Amp Stack's Master Volume is always active. If that module is bypassed, only the four pre-amp controls on the left are affected; the Bass Boost, Presence and Reverb remain accessible.

create custom amps from scratch. Your own tweaks and creations can be saved, appearing as an additional 'Guitar Amp Presets' category, and are available to Guitarix in all subsequently instantiated instruments. You can also save a complete setup, including all current guitar and amp settings as a single file. These, amazingly, can be loaded into any other 'Real' instrument, even the acoustics — great if you simply want to switch guitars but keep every other aspect the same. Brilliant.

Go With The Flow

Clicking the Guitar Amplifier's Edit button on an instrument's GUI opens the Guitarix interface in a separate window. This GUI is divided into two halves; the left column always contains an Input module and an Amp Stack — certain parameters of both can be disabled, but neither module can be removed. Whilst you can insert modules between these, in general you'll add modules below them, in signal flow order. The choice of modules is nothing if not comprehensive, with multiple types and in some cases numerous variations under categories including Tone Control, Distortion, Fuzz, Reverb, Echo/Delay and Modulation, many sporting recognisable brand names. The identity of some is rather less obvious, but interesting to investigate nonetheless. Given so many modules and their variations, attempting to describe them would be futile — and as MusicLab don't provide a separate manual for Guitarix, it's up to the user to explore and discover for themselves. The right column hosts stereo effects modules to further process the signal arriving from the left column. Clicking on the '+' icon of any module adds a new 'blank' below it — just select the desired module type from its drop-down list. Note that once in place, modules cannot be re-ordered, so think carefully about their position in the signal flow before spending time making the perfect setting, as the only way to reposition a module is to delete it, re-insert it elsewhere, and recreate any settings you'd previously made.

Caveats

As the Linux version of Guitarix is constantly evolving, we must assume MusicLab's ported version is also a work in progress. As such, some aspects appear underdeveloped: there is no preset system for the individual modules, and no MIDI automation for any aspect of Guitarix, although MusicLab say this is under consideration. And whilst the preset load/save dialogue includes the option to integrate other external amp sim plug-ins you may own (instead of Guitarix), this is not currently operational. As mentioned earlier, the separate VST3 plug-in can be inserted into any audio track, but also has its limitations. It defaults to having only the Input and Amp Stack modules, it has no presets and no means of saving what you create, nor loading the ones stored within the MusicLab instruments, so you're starting from scratch every time. However, if your DAW supports the saving of track templates, that serves pretty well as a substitute preset filing system.

The Linux Original

Anyone interested in the Linux version that spawned MusicLab's take on the subject might like to check out the Guitarix homepage, found here: guitarix.org. It's interesting to note the differences between the two — the Linux is graphically richer, there are minor differences to some modules, and it also offers MIDI automation. Whilst MusicLab don't provide a user manual for their Guitarix, there is a Wiki manual for the Linux version, found here: sourceforge.net/p/guitarix/wiki/Main_Page. It explains the purpose of many of the modules, as well as offering helpful tips on how to get a good sound — all equally applicable to MusicLab's version of Guitarix.

Conclusion

With its plentiful selection of tube amp models, amp and cabinet convolution impulses, a dizzying selection of distortions, compressors, EQs, delays, reverbs and more, Guitarix delivers an impressive range of tones. It enlivens the sound of MusicLab's acoustic guitars greatly, in particular bringing out the true character of the Steel String. Most impressive though are the overdriven electric guitar presets — tones that I've struggled (and often failed) to obtain with other 'mainstream' amp sims just seem to pop effortlessly out of Guitarix.

summary

Existing owners of MusicLab's guitars would do well to upgrade just to get access to Guitarix — you may well find yourself putting your preferred amp sims to one side! If you're a newcomer to the virtual guitar world, MusicLab's v6 instruments offer a complete all-in-one solution and a diverse choice of models, all with great playability and realism.

$ RealGuitar 6 $199, RealStrat $139, RealLPC $139, RealRick $139, RealEight $139.
W www.musiclab.com

Audient iD44 MkII

USB Audio Interface

Audient's iD44 was pretty good to start with... and it now it's even better.

SAM INGLIS

Audient's iD series of desktop interfaces seem to have nailed a sweet spot in the market, offering excellent audio quality and functionality whilst remaining affordable. The flagship of the range is the iD44, which offers four analogue inputs featuring Audient's high-quality mic preamp design, two sets of optical I/O for expansion, and a console-style master section with talkback and hands-on monitor control.

Dark Powers

In the four years since its launch, the iD44 has proved popular and capable, but Audient have identified a few opportunities for improvement, and they've implemented them in the new MkII version. The new unit seems to occupy the same robust metal chassis as the original, which was reviewed in February 2019 (www.soundonsound.com/reviews/audient-id44), but has gained a new logo and a very smart charcoal-grey colour scheme. It has the same complement of controls and physical I/O, with one small exception: the first headphone socket is now duplicated on both 6mm and 3.5mm jacks. These carry the same signal at the same level, but they can be used simultaneously if needed, and plugging a second set of phones in doesn't seem to compromise the first in any way.

Behind the scenes, however, the audio specifications have been substantially improved. These are listed in unusual detail on Audient's website, and make for impressive reading, especially the output dynamic range of 126dB A-weighted. I can't imagine that the audio performance of the original was exactly holding anyone back, but it inspires confidence in the manufacturer when the numbers are good and, just as importantly, when they're explicit about the conditions under which those numbers were measured. Unlike many desktop interfaces, the iD44 can also accommodate very hot mic inputs: with the input pad engaged, the maximum input level is quoted as +28dBu, which should mean no risk of clipping with loud drummers or guitar amps. Given the emphasis on audio quality and the I/O quotient, it's not surprising that the iD44 MkII still can't be bus-powered, though I could wish the supplied 12V PSU had a longer cable.

The iD control software will also be very familiar to anyone who's used the original iD44, which is no bad thing. The main improvement here is that the physical I/O has now been augmented by a stereo loopback input. This can be sourced from any of the five stereo DAW return channels, the iD44 Master Mix output or any of the four Cue mixes. In these days of Zoom meetings and live streaming, it will enable a lot of setups that might otherwise require complex workarounds.

Coming Of Age

I think Audient themselves would acknowledge that the MkII represents a refreshment rather than a reinvention of the iD44, and some would argue that they could perhaps have gone slightly further. They could, for example, have added a Bluetooth input, or incorporated a built-in talkback mic rather than expecting the user to sacrifice a recording input to use talkback. For the most part, though, the truth is that Audient just got it right the first time around. There are not many interfaces of this size and form factor that have the potential, when expanded, to handle a really serious recording session with 20-odd inputs and multiple cue mixes, but the iD44 is certainly one of them. At the same time, it's equally at home in a humble desktop environment catering to the odd guitar DI and vocal overdub. And as with other Audient interfaces, what's perhaps most impressive is the number of features carried over from large-format console design. From level-compensated speaker switching to a mono button that can be configured to feed left, centre or right channels, the iD44 MkII is full of nice touches that reveal its pro audio heritage.

summary

Audient have updated their flagship desktop interface to keep it at the head of the pack in terms of audio quality, without losing any of its professional features.

$ $699
W www.audient.com


EXPRESSIVE E
osmose
is coming

Icon Martian

Cardioid Capacitor Microphone

If you're looking for a mic to lend weight to lean sources, the characterful Martian could be just the thing.

PAUL WHITE

Icon mics are manufactured in Latvia, which is already well-known for other high-quality microphone brands. The Martian mic reviewed here is part of Icon's Space series of microphones. While the two alien 'eyes' mounted on this mic's spherical basket might lead you to believe that it is some type of stereo microphone, it is in fact a fixed cardioid-pattern, mono mic. Like those employed by compatriots JZ Microphones, the handbuilt capsule uses Golden Drop technology, which aims to allow the diaphragm to move more rapidly in response to transients and thus deliver a precise HF response. The capsule is shockmounted internally, and the bespoke cradle provides further shock isolation as well as support for the included metal pop filter. An elegant two-part aluminium storage case for the microphone itself is included, though this doesn't hold the shockmount and accessories.

Taking a closer look at the hardware, the standard of engineering is impressive, though I can't help wondering if those two circular 'eyes' were added purely for cosmetic purposes rather than as some kind of acoustic adjustment. It is possible to tilt Martian's head up and down over a very wide range: a mechanical stop kicks in at just under one complete revolution. The mesh basket comprises two layers, the inner one much finer than the outer. There are also two layers making up the pop filter, which is a very elegant curved metal structure that clips onto the cradle at two points, and can be swivelled to match the angle of the mic's basket.

The cradle mount has quite an unusual mechanism for securing the mic. A curved arm swings open to the allow the mic to be put in place, then when the arm is pushed back to secure the mic, there's a neat locking mechanism that requires the supporting pillar at the end of the arm to be twisted to release it. At the rear of the cradle is a US-threaded standmount that uses a ball-and-socket mechanism to allow it to be both tilted and rotated with a thumbscrew to lock it in place.

Internally we find Class-A, transformerless circuitry, again handbuilt in Latvia, using good-quality components. Standard phantom power in the range 36-52 Volts is required. A frequency response of 20Hz-20 kHz is specified, and the included plot shows a nominally flat response with no specific presence peaks and just the odd wrinkle, especially at the low end. The sensitivity at 1 kHz into a 1kΩ load is 33mV/Pa and the maximum SPL (for 0.5% distortion) is 134dB. There are no pad or roll-off switches. An equivalent noise level (DIN/IEC A-weighted) of 10.5dB is specified, which is fairly typical for this type of microphone. The overall weight is 0.77kg and the dimensions 118 x 78 x 149mm. A helpful manual is provided, and this includes a few practical application tips.

In Use

I tested the mic with voice and acoustic guitar and found it to have a very warm sound that seemed to smooth out any high-end transients. If working with somebody who has an unusually bright voice that lacks weight, or with an acoustic guitar that exhibits brittle-sounding highs, this mic would be a good choice to redress the balance, though with my own acoustic guitar, which sounds well-balanced in the room, this mic seemed to pull back the high-end zing from picked notes. I tried several different miking strategies but always ended up hearing a smoothed top end. I found pop rejection to be better than average, even without the pop shield fitted, and of course this being a cardioid microphone, you can adjust the mic distance to fine-tune the amount of proximity effect. Percussion tests showed the mic to be perfectly capable on transients too, where the warm low end helps add weight. The smooth high end can also be helpful when miking electric guitar speakers.

Despite its nominally flat response, I feel that the Martian could be regarded as a 'character' mic, good for warming up the lows and smoothing over-prominent transients. Indeed, it sounds not unlike some tube mics I've tried. Sometimes this is exactly what you need, whereas at other times you might require something more detailed. The build quality is impressive, the styling distinctive and the price very reasonable for a microphone of this quality, but if you're looking for a mic to record your own vocals, make sure that it suits your voice before you decide — which is sound advice no matter what mic you are considering.

summary

The Icon Martian is a well-engineered microphone with a distinct personality, both visually and tonally!

$ $699.99
W www.mixware.net
W www.iconproaudio.com

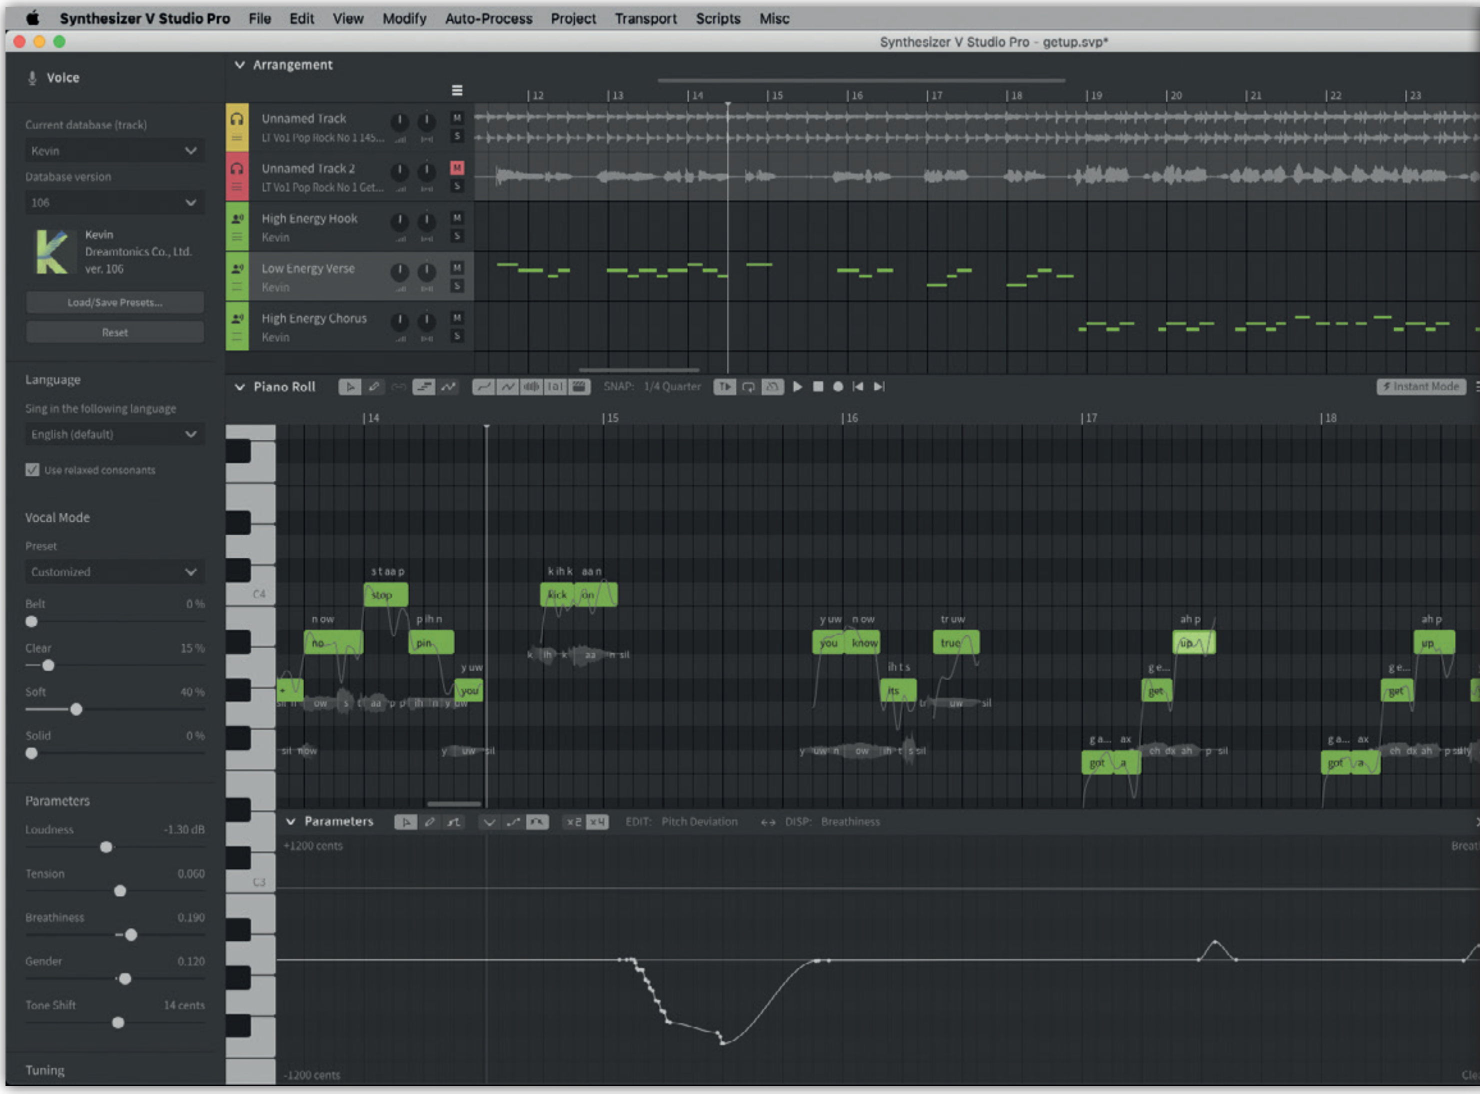

Dreamtonics Synthesizer V

Software Vocal Synthesizer

Can Dreamtonics really deliver a session singer in software?

JOHN WALDEN

Virtual musicians, in the sense of virtual instruments with powerful performance elements, are now a fixture in the workflows of countless composers, producers and recording musicians. However, while the drummer, bassist, keyboardist and guitarist of your virtual session band can now undoubtedly produce the goods, what about the vocals? Well, Dreamtonics might well suggest that their flagship product — Synthesizer V — can do just that. Session singer in a software-shaped box, anyone?

Voice In The Machine

The human voice — spoken or sung — is a hugely complex instrument, making the technical challenge of synthesizing it a very considerable one. A number of brave developers have tried, though, and for solo voices, perhaps the most widely known product is Yamaha's Vocaloid. While the potential of the technology was clear to see when *SOS* reviewed the Sonika voice database for Vocaloid 2 in March 2010, the workflow was somewhat laborious. Backing vocals or obviously processed EDM vocal styles were possible (and became a thing in their own right) but creating a lead that might fool the listener into believing it was a 'real' voice remained out of reach.

Of course, in music software terms, 2010 is a long time ago. While Yamaha have continued to move Vocaloid forwards, over that same time span, competition has also appeared. One of these newcomers is currently gaining a lot of interest; Dreamtonics' Synthesizer V.

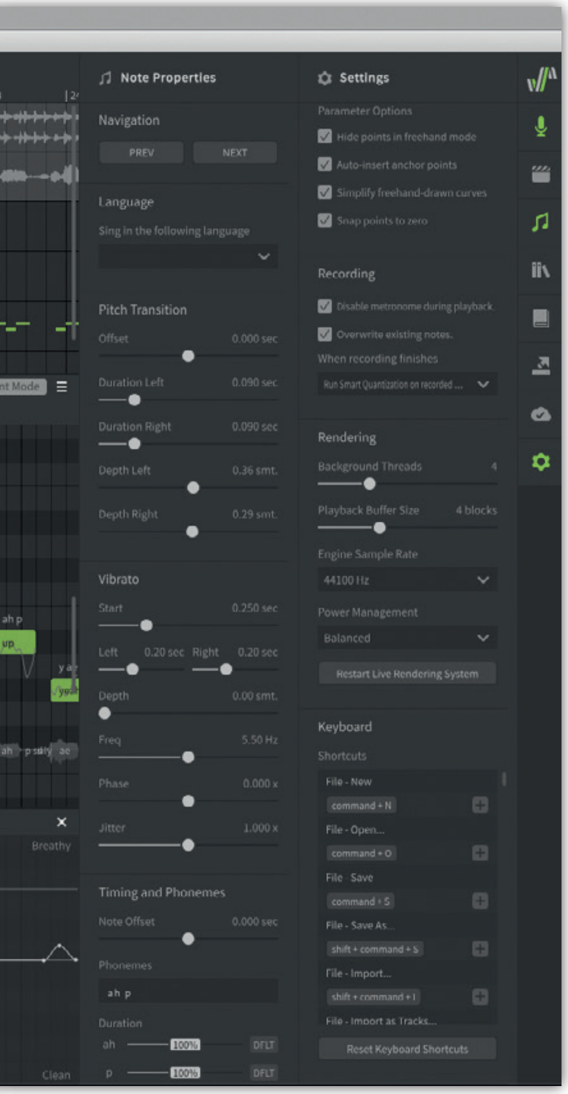

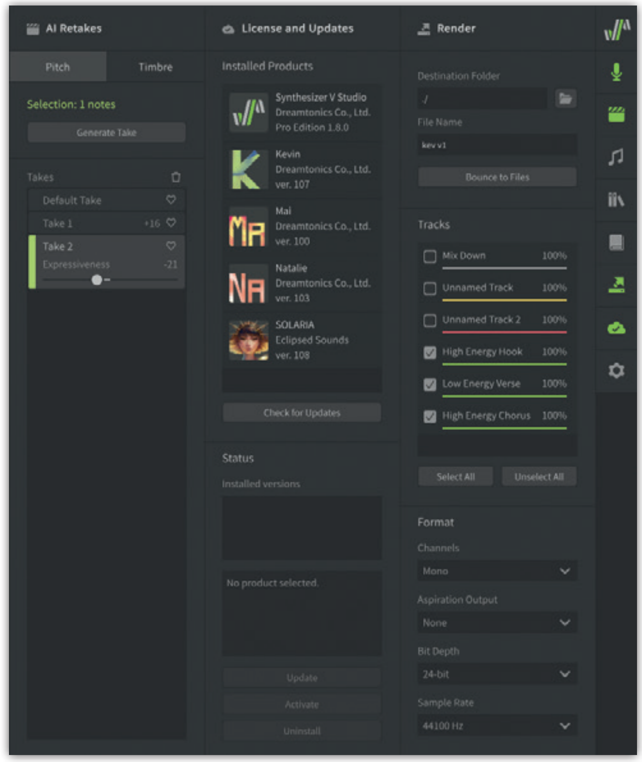

Many of Synth V's features are neatly organised into themed sub-panels that can be popped open as required by the user.

Dreamtonics are also based in Japan, and Synth V is very similar in concept to Vocaloid. Running either standalone or as a plug-in, the software has two main components; the synthesis engine and a selection of voice databases for individual 'singers'. The most recent iterations of the engine include AI elements with machine-based learning to improve the realism of the end result. The current selection of voice databases (built from recordings of real singers and available individually as separate purchases) include native Japanese, Mandarin Chinese and English singers. I had access to a number of the native English voice databases for this review but, for newer voices, the engine does enable them to switch between these three languages.

Two versions of Synth V are available. The Basic version is free-to-try with some function/feature limitations, but does at least allow you to experience the synth engine in action. The paid-for Pro edition obviously removes any limitations and, as described more fully below, provides an extensive list of editing and style options that can be applied to the synthesized voice.

It's also worth mentioning that — at the time of writing at least — the available documentation is lagging somewhat behind the development of the software itself. Dreamtonics are actively working on improving matters on this front, but it did leave me unsure whether I fully understood all the features during the course of the review. Watch this space...

Sing Something Simple

Synth V's UI contains three key elements. First, a vertical strip of buttons (far right) allows you to toggle open/closed a series of sub-panels, each focusing on a specific set of command options and that can be placed (by dragging) either on the left or right sides of the overall UI. Second, the Arrangement panel provides a DAW-like 'project window' containing a vertical arrangement of the tracks within your current project and a bar-based timeline display. Mini note displays along this timeline provide useful visual feedback for the overall arrangement.

A project can contain multiple synth voice tracks based upon one or more of the voice databases you have installed. Very usefully, you can also add audio tracks (termed Instrumental Tracks) into the arrangement. In the standalone application these might most obviously be used for an instrumental mix as musical context for your synth voice creation. All the tracks — voice synth or audio — have volume, pan, mute and solo options. There are no effects options but it's perfectly adequate for the core task of creating the synthesized vocal(s).

Third, for the selected voice track in the Arrangement panel the Piano Roll panel shows the MIDI-like note 'blobs' that represent the melody and timing of the sung performance. The display also shows the engine's AI-generated pitch curve, with features such as pitch »

Dreamtonics Synthesizer V

$89

PROS

- Truly remarkable technology.
- Can create genuinely useful results.
- Good selection of voice database options, and Solaria is capable of excellent results.
- Given just what's possible, the price is also remarkable.

CONS

- Results equivalent to an accomplished singer rather than an exceptional one.
- Occasional quirks in an otherwise remarkably smooth workflow.
- Documentation currently a bit thin on the ground.

SUMMARY

Synthesizer V is a truly remarkable product. Dreamtonics are bringing the concept of a software singer right to the very edge of believable.

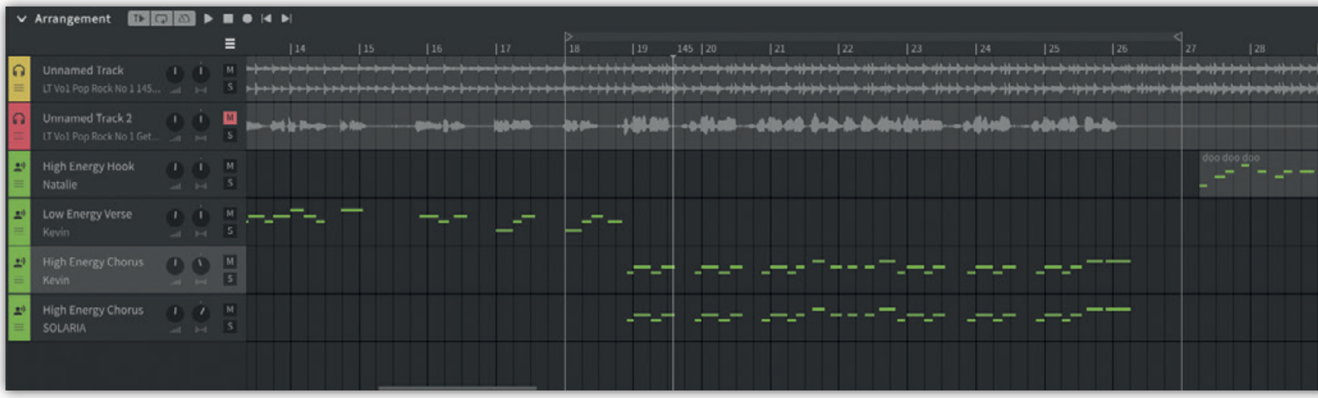

The Arrange panel lets you organise your various voice and audio tracks, see an overview of the current project and configure time signature and tempo information.

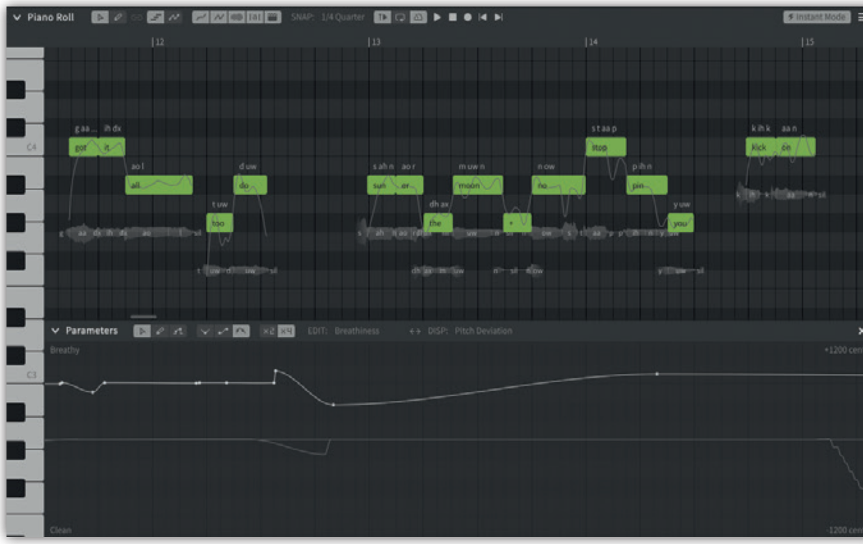

The Piano Roll panel provides a similar experience to working with a MIDI editor, including options for parameter modulation.

slides between legato notes and vibrato on sustained notes. However, these properties are fully editable, and the Note Properties panel can be used to specify the pitch transition and vibrato settings for individual (or selected groups of) notes. The Piano Roll also shows any lyrics you have added for each note blob (just double-click on a blob and type the word you want) and the phonemes the engine has assigned to these words to generate the appropriate pronunciation.

At the base of the Piano Roll panel, you can pop open further sub-panels that allow you to create modulation curves for key properties of the voice. These include Pitch Deviation, Loudness and Tension, amongst a few others, and, as with any virtual instrument, modulation can add considerable expression and realism to the performance.

You can add notes into the Piano Roll editor manually, via a MIDI file or by recording them from a MIDI keyboard. New notes are given a default 'la' lyric and, if you enable the Instant Mode button (top right of the Piano Roll panel), Synth V automatically revises the vocal waveform as you add or edit notes. The note editing toolset is very much like a MIDI editor, allowing you to change pitch, position and length very easily. Keyboard shortcuts are also supported, and you can check the default configuration for these in the Settings panel.

Having added a voice track and assigned the required voice database, generating a sung vocal requires you to enter/create a pattern of MIDI-like notes and type in the lyrics for each note. As you enter or edit notes/lyrics, Synth V works away in the background generating the resulting vocal, and you can see the calculated waveform within the Piano Roll panel. You then simply hit the Play button to audition the result.

The first few times you use Synthesizer V, you may find yourself staring at your studio monitors with your jaw smacking off the floor. Yes, the odds are that there will be some further editing work required — words where the pronunciation or phrasing doesn't flow in a truly 'human' fashion, for example — but, equally, there will be other phrases (and, quite possibly, a lot of them) where the degree of realism is, frankly, quite staggering. This is way beyond the 'good enough for backing vocals or robotic EDM vocal hooks' stage; it's right into prime-time lead vocals territory.

Talent Pool

While Dreamtonics have their own collection of voice databases available for Synth V, third-party developers are also involved in this element of the process. Other companies developing voices for Synth V include Audiologie, AHS Co, Quadimension and Animen.

Solaria — probably the most impressive of the English-based voices I looked at here — was developed in collaboration with Dreamtonics by US-based Eclipsed Sounds. However, they have a second singer — Asterian — that, at the time of writing, was about to be launched. This voice database is being built from the vocals of US singer David Hollaway. I got to try a pre-release version of Asterian and I've also heard David sing/speak via YouTube; if you can imagine the voice of Barry White in software, then David's virtual singing incarnation will have you in the right ballpark. It's deep, rich and very impressive. It's also very different stylistically from any of the other existing voice databases, suiting soul and perhaps some hip-hop styles. It would also work well on song styles related to musical theatre or film. I'll be very interested to see (hear?) what Eclipsed Sounds have to offer next.

Neve

8801

THE COMPLETE CHANNEL STRIP

THE 8801 DELIVERS NEVE® CONSOLE SOUND IN A CONVENIENT 1U RACKMOUNT DESIGN WITH POWERFUL TOTAL RECALL™ AND DIGITAL CONVERSION.

DESIGNED AND CRAFTED IN THE UK BY NEVE ENGINEERS

NEVE.COM

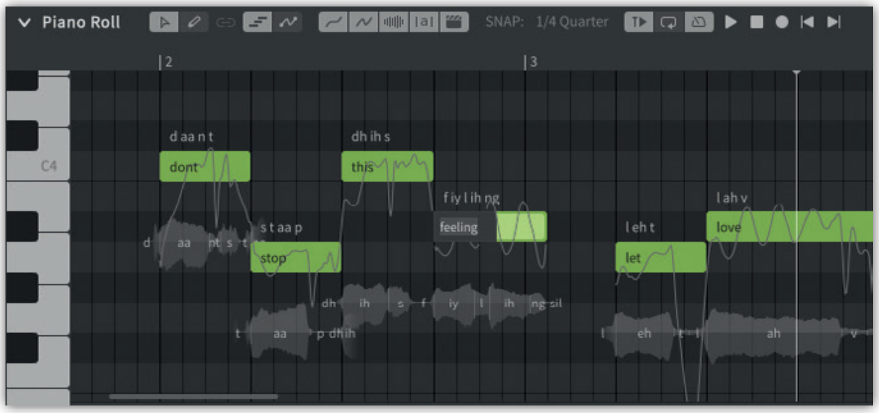

Clicking on a note within the Piano Roll allows you to enter lyrics for your selected vocalist (voice database) to sing.

Whatever Dreamtonics are doing with their AI-based algorithms under the hood, it is very, very clever indeed.

Get Real

Once Synth V has created its initial vocal, it's then up to you to decide just how much additional finessing is required. At the Track level, if you are using one of the newer AI-based voice databases (as are all the ones I had access to), the Voice panel offers a set of 'Voice Modes'. These vary for each voice but might have labels such as Soft, Airy, Passionate or Light. You can use the Vocal Mode sliders to create blends of these various Modes and, to a large extent, these then dictate the style of vocal delivery created.

These seem to operate at the Track level so, if you want your singer to sing a verse in a combination of Soft and Light modes, but a chorus that blends the Power and Passionate modes, then you simply create those vocal parts on two separate tracks within the Arrange panel. You can set the Voice panel's Parameters sliders — Loudness, Tension, Breathiness, for example — at track level, but these parameters can also be modulated via the Piano Roll's Parameters sub-panels. The Piano Roll's Parameters panel also includes a Pitch Deviation pane. This modulates the pitch relative to the pitch curve displayed within the main Piano Roll display (that is, it 'deviates' the pitch). If Synth V's pitch variations are not always exactly what you want, this modulator allows you to customise pitch in a very precise fashion.

Within the Piano Roll, if you select individual notes, the Note Properties panel lets you apply specific settings for pitch transitions (from one note to the next) and vibrato (its onset, depth and frequency, amongst other options). Controlling vibrato did leave me scratching my head a little and I'm still not sure I fully understand how these various controls interact. By default, Synth V's AI does a good job with pronunciation but the Note Panel's Phonemes settings — which allow you adjust the relative duration and strength of each phoneme within a word — let you finesse the pronunciation and customise the delivery of a word if required.

If this isn't already enough, there are other commands to assist you in the performance creation process. It's here that some comprehensive documentation of the full feature set would be really beneficial as I'm sure there are features — for example, the intriguing AI Retakes panel, the Ornament Selected Notes command, or working with note Groups — that I'm not yet fully exploiting and that I suspect would improve either workflow or the quality of the end result.

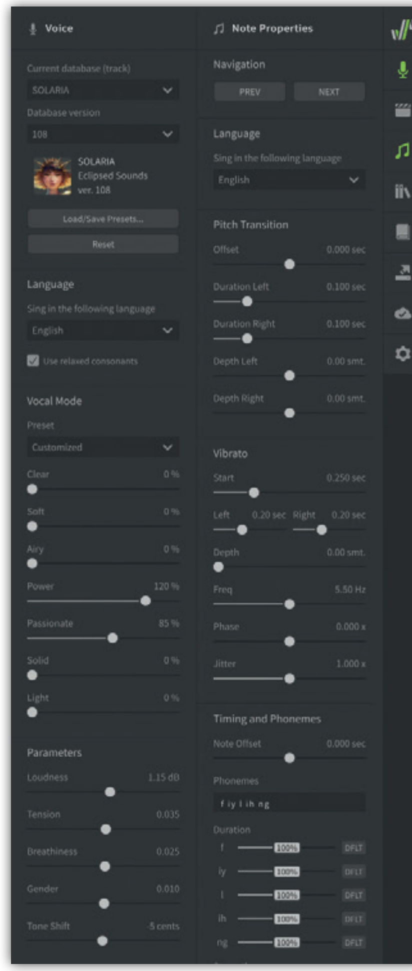

For adding character to the performance, the Voice and Note Properties panels contain many of the key controls.

Cover Me

Once down the vocal synthesizer rabbit hole, you may encounter the numerous YouTube cover versions of popular songs, where a Synth V (or Vocaloid) generated vocal has been used to replace the original. Yes, there is the good, the bad and the downright ugly. However, the best of these are a really impressive showcase of just how far this technology can be pushed.

There are also some interesting and educational elements to the workflow used to create many of these cover songs, from stem extraction from the original audio, pitch curve analysis via software tools such as Praat (developed by academics from the University of Amsterdam) and the use of Synth V's own scripting language (a user-created script set called 'Real Vocal' seem to be the most popular of these) to import the original vocal's pitch information into Synth V, giving your synthesized vocal elements of the pitch inflection contained within the original.

Some YouTubers also make their Synth V files for these vocals available to download so you can then import them into your own system, and these can be very instructional. If you want a head start, search YouTube for Synth V covers of tracks by Adele, Paramore, Linkin Park and Kate Bush; prepare to be both amazed and amused.

Singer Auditions

I was able to try four voice databases during the review. Mai is a bundled with Synthesizer V Studio and, while built from a native Japanese female voice, it sings very effectively in English. The sound is youthful and would suit pop or dance styles. The other three — Natalie,

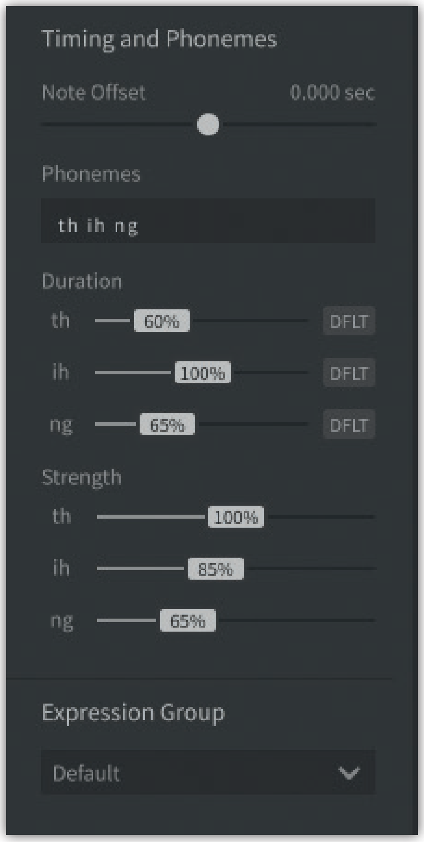

The Note Properties panel's Timing and Phonemes section gives you a fine degree of control over pronunciation.

Kevin and Solaria — are all native English voices. Natalie's voice could carry a range of styles (most obviously in the dance/electronic genres but also ballads and indie pop-rock), while Kevin also appears quite flexible and might go from EDM through to pop-rock styles.

However, to my ears at least, the cream of the vocal crop is Solaria (see the 'Talent Pool' box for more details). Based upon the singing talents of Broadway-trained vocalist Emma Rowley, this voice can cover a very wide range of styles from intimate pop, powerful EDM, pop-rock and off into symphonic rock. It's also the voice that seemed to produce realistic results with the least manual editing. The results can be very impressive indeed.

With each voice offering its own character, and each also having its own set of Vocal Mode options for customisation, there are a lot of musical, song or performance styles that you could coax out of just these four voices. If I was to have a wishlist item, though, it would be for both male and female voices with a 'rasp' Vocal Mode. You can hint at this with some of the above voices but, as yet, it doesn't really get you into Chris Cornell or Lizzy Hale territory.

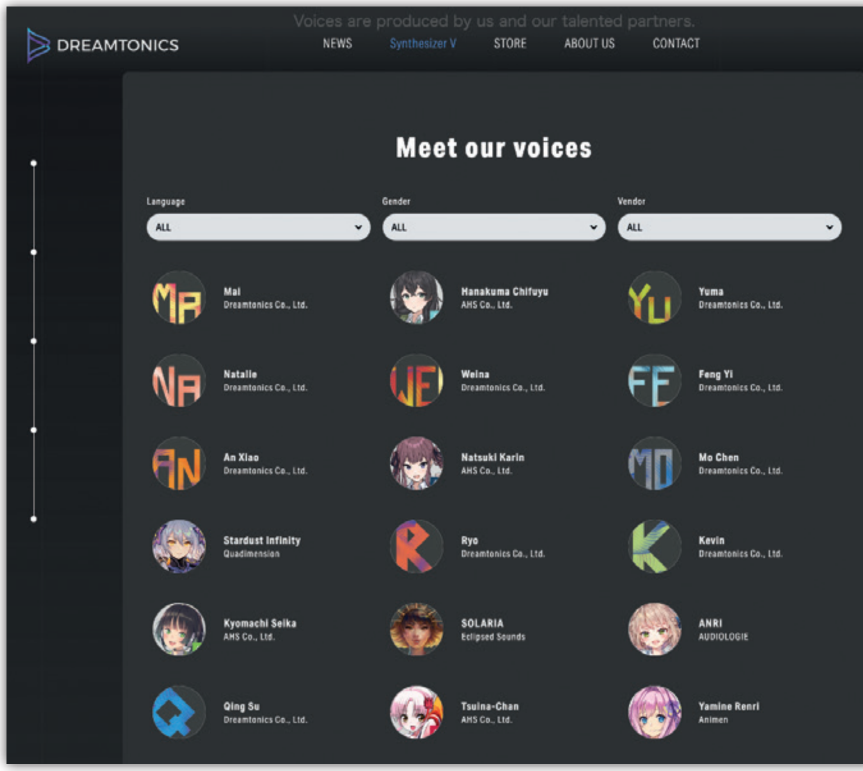

Dreamtonics' website features an ever-expanding catalogue of virtual singers that can be purchased to add to your Synth V projects.

DAW Tour

Inserted as an Instrument VST within Cubase 12, Synth V's plug-in format provides exactly the same (resizeable) working environment. Your Synth V creations and settings are also saved within your overall DAW project.

The plug-in seems to support multiple audio outputs (I was able to activate these within Cubase) but I couldn't find an obvious means to then route individual vocal tracks from the Synth V plug-in to different channels within the Cubase MixConsole. I might simply have missed something here but, if not, it would be a great addition. I did experience some other occasional less-than-fluid workflow wrinkles (for example, pop-up dialogue boxes that hid themselves behind other DAW windows) but nothing of any major consequence.

Are We There Yet?

So, the feature set and workflow are impressive, but can Synthesizer V actually produce a believable natural lead vocal performance? Well, a qualifier or three aside, amazingly, I think it can. Over my years of doing product reviews for *SOS*, there have been a few occasions when my jaw has seriously hit the floor (and one product — Melodyne — that did that twice!). Synthesizer V is one of those moments. Given that synthesis of realistic singing is such an ambitious aim, that Synthesizer V can even get close feels like music software as science fiction. Putting aside the artistic desirability of a virtual singer for a moment, the underlying technology is seriously impressive.

Those qualifiers? Well, first, it's still perfectly possible for a Synth V vocal to sound very obviously unnatural. Dreamtonics have done a tremendous job of getting the maximum level of 'natural' out of the minimum user effort, but you may still find yourself needing to dig in and finesse the pronunciation, tuning or vocal character. However, the tools are there to do just that.

Second, your definition of 'usable lead vocal' will undoubtedly depend upon the musical situation that you are working within. If you simply want to add a few vocal hooks to an EDM track to play to your mates, the required quality bar is going to be somewhat lower than if you are creating a theme song for a Hollywood blockbuster film. Context is everything but, once you get to grips with Synth V, I think it's surprising just how far you might go up that quality spectrum.

Third, and perhaps obviously, while Synth V might let you add vocals to any song when you can't actually sing them yourself, it can't make a bad song into a great one. Whether your vocals come from a human or a piece of software, writing a great song that's going to connect with a human audience is still down to you.

Deep Fakes?

The impressive nature of the technology aside, does the world need computer generated vocals in its music? I get that this concept would be the total antithesis of musicality to some musicians, producers and music consumers. Indeed, in part, it's a view I can share. Of all the elements of a great song, it's generally the vocal — its words, its phrasing, its emotions — that makes the most intimate connection with a listener. Can Synth V make that connection in the same way that, for example, singers such as Adele or Joe Cocker or Whitney Houston or Robert Plant (and many others) can? While you can certainly get it to simulate emotion in various ways, impressive though the technology is, the results are not in that superstar vocalist territory.

If you want to audition an example that I think captures the essence of where Synth V is currently at, do a YouTube search for 'Synth V Adele Easy On Me cover'. This is perhaps an unrealistic comparison (Adele has one of the most iconic voices of our generation; lots of human singers can't get close) but, even so, the Synth V generated vocal is a fabulous demonstration of what can currently be achieved, and is undeniably impressive... even if it doesn't deliver that 'something special' emotional intensity of Adele's original.

Yet. Because the obvious question is where this AI technology might go? Deep fake video can already fool our eyes into believing we are seeing actual video footage of people (for example, movie stars, or other celebrities, alive or dead) doing stuff they have never actually done. How long before we see that deep fake concept applied to the voices of famous singers? Where the AI 'learns' their sound and the varied characteristics of their voices and can recreate it to a level that our ears are fooled? My time with Synth V suggests that's going to come, and it would be a combination of fascinating (from the technological perspective), exciting, and frightening all at the same time.

Singer In A Box?

Deep fake might be for the future; so who might buy Synthesizer V Studio right now? Well, it's undoubtedly a niche product but, if you are a songwriter/producer who doesn't sing, whether to simply add vocal hooks or backing vocals, for demo projects with your vocal ideas as guides for session singers, or to pitch songs to artists, then Synth V — with suitable voice databases — now provides a remarkable solution. For many of these kinds of applications, in practised hands, Synth V will easily exceed the required quality bar.

What about a lead vocal on a commercial release? Well, those musicality issues aside, I'd not find this too difficult to imagine in a pop, dance or EDM context as vocal styles often diverge from 'natural' anyway. For singer-songwriter ballads backed by just an acoustic guitar or piano? Well, you could do it, but you would have to be very thorough with your Synth V refinements to completely fool a listener paying very close attention. You might get them second guessing though... and that's an achievement in itself. That said, maybe someone has maybe already pulled this trick off and the rest of us don't know about it yet. If it hasn't happened already, then I suspect it's not far away.

Synthesizer V Studio is a technical marvel. Synth V's vocals might not have all the nuances or character of a truly exceptional vocalist, but it's still a remarkable piece of software. It's also remarkably affordable, the price of entry for the engine and one or two vocal packs being accessible to almost anyone.

Whether you approve of the concept of computer-generated lead vocals or not, Dreamtonics appear to have taken actual magic and turned it into code. Synthesizer V is groundbreaking music technology.

Audio Examples

To provide an impression of what is possible with Dreamtonics' Synth V, I've created four short audio examples which you can find on the web version of this article. Each is based around a different musical style and, in all cases, consist of a simple backing track (to provide the musical context) and a single lead vocal produced within Synth V (no double-tracks or harmony parts).

These examples were created after I'd been using Synth V off and on over about a two-week period within which I was working on the review itself. I guess that makes me still a 'new' but perhaps not a 'novice' user. Even within that period of time, however, the workflow most certainly became faster as I gained familiarity with the various controls and options.

To provide a (hopefully) realistic impression of what's possible based upon this level of experience, I deliberately limited the amount of time I spent on each of these examples. The backing tracks were quickly roughed out and (a touch of compression and/or reverb aside) are just intended to provide a bed within which to hear the vocal part. I then spent around 60-90 minutes creating each of the vocals you hear from scratch. This involved a number of steps. First, I created a suitable melody on a piano, and this was then imported into the Synth V plug-in running within my Cubase project. Second, I then selected a suitable voice database for the specific project and wrote (no prizes for originality!) and entered the required lyrics for each of the melody notes. Finally — and where the most time was spent — I worked my way through the performance, adjusting note timings, pitch data and pronunciation/emphasis of the individual phonemes as required. I also experimented with the various Vocal Modes that each voice offers to provide contrast within the different parts of each example.

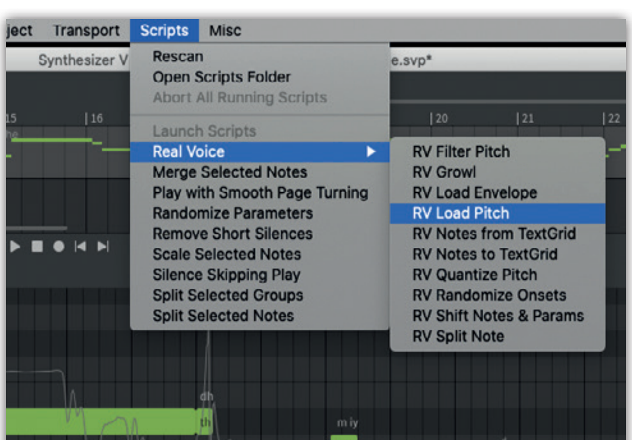

Synth V includes its own scripting options and there are some interesting user-developed scripts available online.

$ Dreamtonics Synthesizer V Studio Pro $89; voice databases from $79.

W www.dreamtonics.com

Rode NT1 5th Gen

XLR & USB Capacitor Microphone

Has Rode's new USB-capable NT1 made clipping a thing of the past?

SAM INGLIS

Innovation in music technology can take many forms. Sometimes it means implementing features and creating products that have never been seen before. But it can also involve bringing existing developments to new markets, by finding ways to manufacture them more affordably. This latter kind of innovation has always been at the heart of Rode's business model, but in recent years, they've also scored some impressive design firsts. The NTR, for example, goes where no ribbon microphone has gone before, with its laser-cut ribbon, insanely high build quality and extended frequency response.

The new NT1 5th Generation, arguably, innovates on both fronts.

One Up

Rode have had an NT1 in their line-up for more than 25 years now. It's always been a strong option for those seeking a no-frills capacitor mic, and I'd hazard that it's featured on more well-known recordings than people are letting on! It's a fixed-cardioid, large-diaphragm mic, and the no-frills aspect of the design means you don't get pad or filter switches. Unlike many fixed-cardioid mics, its one-inch capsule is truly single-sided, rather than being a Braunmühl-Weber design with the rear diaphragm disconnected. The current iteration of this capsule is called the HF6, and is unique to the NT1. In fact, this capsule is one aspect of the design that hasn't changed in the new 5th Generation model. The cosmetics are also practically identical to those of the previous 4th Generation model, with a smart dual-layer headbasket and functional matte black body.

So what has changed? Well, the key to the 5th Gen's newfound powers is found in the base of the mic. The conventional XLR connector from the 4th Gen model has been replaced by a new, patent-pending socket that can accept either a female XLR or a USB Type-C plug. When connected the old-fashioned way, the NT1 5th Gen behaves much like its predecessors: it's a conventional capacitor mic, which requires phantom power and delivers an analogue signal through your preamp of choice. And in this role, it has some eye-catching specifications, most notably an incredibly low self-noise of just 4dBA. But on that front, nothing much is new, because the 4th

Gen already offered the same analogue performance. The focus of the next generation is digital.

All In One

Unlike many USB mics, the NT1 5th Gen is an input-only device, and doesn't have a headphone output of its own. Nor does it have a physical gain control. As we'll see, this isn't necessarily an issue, and for spur-of-the-moment recordings where you don't need to monitor anything, you could just plug in and go. For most use cases, though, you'll want to install the Rode Connect utility. This allows you to aggregate multiple Rode USB devices and perform low-latency cue mixing. It also reveals that there's a lot more going on inside the NT1 than mere A-D conversion.

Clicking on the NT1 icon within Rode Connect brings up an editing window that exposes multiple parameters, none of which is accessible from the mic itself or available when it's used as an analogue source. First up is a gain control that runs from 0 to +60 dB in 1dB steps. Below this you'll find radio buttons for a high-pass filter which can be engaged at 75 or 150 Hz, but it doesn't stop there.

During their massive growth over the last 30 years, Rode have absorbed other companies including Aphex, makers of the original Aural Exciter enhancer and Compellor compressor. Their designers' know-how has been put to good use in the NT1 5th Gen, and four icons to the right of the gain control engage simple 'one-button' noise gate, compressor, Aural Exciter and Big Bottom processing. These are implemented digitally and have no controls other than the on/off button, but since they come after the gain control in the signal path, it is possible to manipulate the amount of compression that is applied by increasing or decreasing the gain.

The obvious limitation here is that if you increase the gain too far, you risk making the signal too hot to handle and causing clipping. Except you don't — and that's where the unique aspect of the NT1 5th Gen design comes in.

Window Shopping

Every device designed to capture or process audio has a dynamic range: the ratio between the greatest amplitude that can be represented within the system, and the noise floor of the system. The NT1 has an incredibly low noise floor and can accept sound pressure levels of up to 132dB SPL, so considered purely as an analogue device, it has a massive dynamic range. It would take a very, very loud acoustic event to cause the mic's analogue circuitry to clip; and at the other end of the scale, you'd have to be recording something vanishingly quiet in a completely isolated acoustic environment before its noise floor became apparent.

Making a digital recording from the NT1 or any other mic involves passing the signal through several other devices: a mic preamp, an A-D converter and a digital recorder such as a DAW program. Each of these has its own dynamic range, and we can think of them as a series of 'windows' through which the signal must pass. In a modern digital recording system, the dynamic range of each of these individual elements is much greater than that of any acoustic signal we're likely to be capturing at the mic. However, in order for the signal to make it from the mic to the end of the chain unmolested, it has to pass through all of the windows in turn, and if these windows are out of alignment with each other, clipping or noise problems can result.

The classic situation in which this problem arises is when we're overdubbing to a previously recorded source. Struggling to hear the input signal clearly over the playback, we turn the input gain up when we should be turning the playback down. In doing so, we push the dynamic windows of our mic preamp and our A-D converter apart. Go too far, and peaks in the analogue signal become too hot for the converter to handle. This trap is entirely avoidable, though it has to be said that the ergonomic design of software mixers doesn't always do enough to prevent us falling into it.

Range Finding

When connected over USB, the NT1 5th Gen internalises several of these windows. The signal isn't just being captured by the capsule and delivered through the impedance converter. It's also being preamplified by Rode's 'Revolution' preamp circuitry and converted to digital. So, on paper, there is potential for exactly the same trap to open up. If you push the preamp gain too high in order to hear yourself better, you risk overloading the digital stages and causing clipping.

The NT1's built-in Aphex processing can be accessed only through Rode Connect.

And that's exactly what happens if you select the NT1 as a 24-bit source within your DAW. Set the gain too high, and the recorded signal will be clipped. However, if you've been monitoring it via Rode Connect, the chances are you won't have heard any ugliness during the recording. What's going on?

The answer is that the signal delivered by the NT1 to Rode Connect is effectively impossible to clip, because the NT1 5th Gen doesn't just have one A-D converter: it has two in parallel, and together these have a dynamic range that exceeds that of the mic and its internal preamp combined. To abuse the window analogy further, it's as though Rode have placed a second window directly above the first. One of the A-D converters is aligned such that it's impossible to make the analogue circuitry in the NT1 generate a signal hot enough to clip it, even at maximum gain; the other is aligned such that its noise floor is lower »

Rode NT1 5th Gen

$249

PROS
- A very affordable studio-quality mic that stands up against some much more expensive rivals.
- Incredibly low noise floor.
- Comprehensive accessories including shockmount, pop shield, USB and XLR cables.
- USB operation introduces optional Aphex processing and 'unclippable' 32-bit floating-point mode.

CONS
- No storage case or box included.
- USB operation doesn't support 44.1kHz sample rate.

SUMMARY
The NT1 5th Gen is an ideal starter mic for podcasters and home studio owners. You might want to add to it later on, but it'll be a long, long time before you outgrow it.

On the upper track in this Reaper project, I've recorded myself speaking into the NT1 at maximum gain (first clip, left) and minimum gain (second clip, right). When these clips are normalised, there is no discernible difference between them, as shown on the lower screen. By contrast, the grey clip on the lower track was recorded with maximum gain at 24-bit and remains irrevocably clipped even after I adjust its clip volume.

» than that of the NT1's capsule and electronics at minimum gain. Adding the gain range of the preamp to the dynamic range of the mic itself suggests that this dual converter would need to have a total dynamic range of nearly 200dB to completely eliminate clipping. That is way more than can be represented in a 24-bit fixed-point signal, and so the native digital output of the NT1 is a 32-bit floating-point signal. This is what Rode Connect works with.

As was previously mentioned, this means that clipping is never audible within Rode Connect. And if your DAW supports 32-bit floating-point signals (not all do), the same applies here. Your recordings may show waveforms that look clipped, but by applying clip gain or normalising, you can always restore a sensible level with adequate headroom.

Stay Connected

I tested the NT1 with two 32-bit capable DAWs — Reaper and Pro Tools — and provided you remember to set up the session for 32-bit files, 'unclippable' recording works exactly as expected. In Reaper, you can use the Item Properties to adjust the volume of, or normalise, recorded clips. In Pro Tools, you can do the same using Clip Gain. In both cases it's completely non-destructive, and no matter how hot the gain setting, it's always possible to recover a non-clipped version of the recording (although in Pro Tools, the recording will continue to look clipped even after you apply Recalculate Waveform Overviews from the Clips window). The only fly in the ointment is that the NT1 is limited to 48kHz sample rates and multiples thereof, and can't be used at 44.1 or 88.2 kHz. So if, for instance, you're a remote session singer or voice artist and need to overdub to existing sessions, you may need to retain the ability to record analogue.

The experience of using the NT1 with Rode Connect on a recent MacBook was very slick, though I didn't have access to other USB Rode gear, so can't comment on how well the aggregation works. The preset Aphex processing sensibly errs on the side of subtlety, with the Aural Exciter adding a bit of sparkle to dull voices or acoustic guitars, and the Big Bottom doing just enough to give speech that Hollywood trailer effect. It's perfect for podcasting, live streaming and voiceover recording, and the noise floor really is eerily low.

When it comes to music recording, the 3m USB cable and the option of aggregating further NT1s through Rode Connect mean it's a more practical option than most USB mics, but for anything that involves multiple instruments or pairing with other types of mic, you'll need to revert to old-fashioned XLR connection. You won't get the Aphex processing or high-pass filter, and you'll need a mic preamp, but the good news is that the NT1 is not only super-quiet, but also sounds really good. It's clean but not wholly characterless; compared with a U87, for example, you don't quite get the same richness in the midrange, but there's a gentle boost above 10kHz which adds an appealing and relatively subtle air to most sources.

If you disregard the NT1's USB capabilities, its obvious competitors include the Austrian Audio OC16 and Sony C80. These are both excellent mics, and a competent engineer should be able to make good recordings with any of them. The NT1 doesn't have the presence lift of the OC16, or the slightly dry sound of the Sony with its smaller capsule, but it too is a mic you can point at practically anything and expect good results. And not only is it the cheapest of the three, it's the only one supplied with cables and a pop filter.

Not that I think music-oriented purchasers should disregard the NT1's USB capabilities. A USB mic might not have much of a role to play when recording a drum kit, but there are times where we can all use a simple but effective plug-and-play option for Zoom calls, impromptu overdubs, or to throw in the laptop bag in case inspiration strikes. This is a mic that punches a long way above its weight.

It's All Here

The NT1 5th Gen comes in a smart cubic cardboard box which is absolutely stuffed with accessories. These include the matching SM6 shockmount with removable pop shield, a good-quality mic cable, and a smart and tough 3m USB cable terminating in Type-C connectors at both ends. The only thing obviously missing is any sort of case or storage box for the mic itself, but at this price, it's hard to complain when you are getting absolutely everything you need to actually use the mic. As ever with Rode, build quality is excellent for the price, and in fact the only minor niggle I have arises precisely because of the high quality of the USB cable. This has a toughened exterior which makes it relatively rigid and spring-like, so the price you pay for durability is that bumps and scrapes are easily transferred to the mic.

$ $249
T Rode +1 562 364 7400
E info@rode.com
W www.rode.com

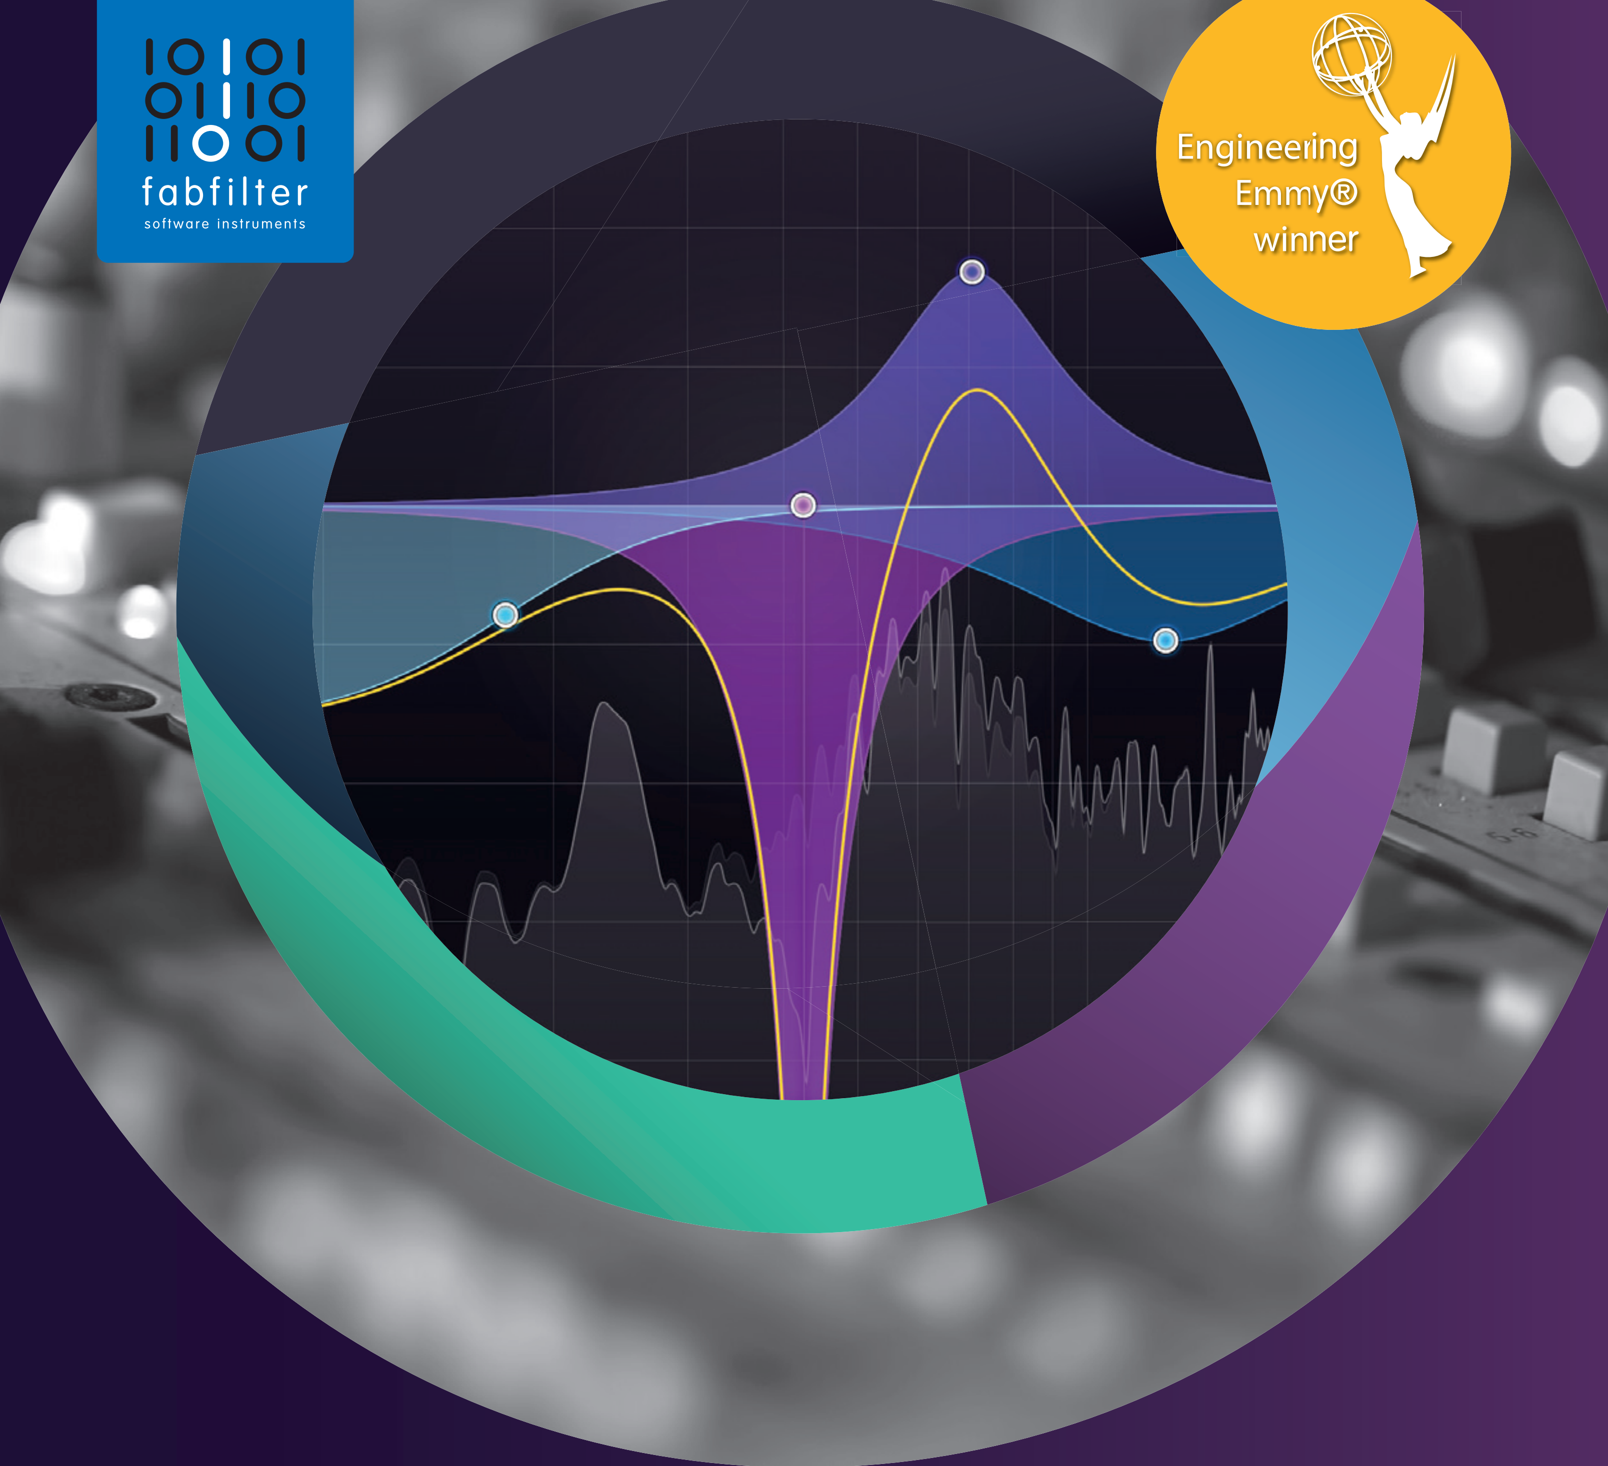
fabfilter
software instruments
Engineering
Emmy®
winner

Antares Auto-Tune Pro-X

Pro-X lets you dig very deep with your pitch correction, but its Auto mode helps keep things incredibly simple.

Pitch Correction Software

We check out Antares' latest fine-tuning of their flagship product.

JOHN WALDEN

It's hard to underestimate the impact that Antares' Auto-Tune had on the world of music production when it was launched in 1996, when it broke the link between pitch and time changes, made pitch correction quick and convenient — and created the warbly 'Cher effect' that's been used by countless other artists since. Its name quickly made the transition from proper noun to popular verb! Of course, it's now far from the only game in town, and competition has increased to the point that both real-time and offline pitch-correction processors are bundled with most serious DAW software. Nonetheless, Antares have continued to refine Auto-Tune and expand its capabilities, as well as embedding the algorithms in various hardware products, and recently the company released Auto-Tune Pro-X, the latest iteration of the software/plug-in version. Does it offer anything your DAW doesn't already? Let's find out...

On The Shoulders Of Giants

We've reviewed several iterations of Auto-Tune over the years, including the first to be given the 'Pro' suffix in *SOS* October 2018 (https://sosm.ag/1018-antares-auto-tune-pro). And while there are certainly improvements, this latest release takes a similar approach to what's gone before, and can best be viewed as an evolution of the already impressive feature set. So there remain two main modes of operation, called Auto and Graph.

Auto provides the familiar 'auto-tune' insert process/effect, its carefully chosen control set enabling you to apply just as much automatic pitch correction as you desire. That control set has been expanded over the years, and now offers considerable flexibility, especially when you dip into the Advanced view, as well as a low-latency configuration to allow for live pitch-correction effects (obviously the latency of your audio interface must also be taken into account).

Graph mode lets you dig considerably deeper, and offers far more precise control. It provides a sophisticated graphical interface with very useful visual feedback and all the tools needed to fully finesse your vocal's pitch (and timing, if required) in the finest of detail. It's already both flexible and impressive, and the quality of the underlying pitch-shifting can be very transparent if 'invisible correction' is your aim, though there's still plenty of scope for more obviously creative effects.

Fresh Tune

Antares have made a number of changes to the GUI. For example, the plug-in window is now fully resizeable and its high-resolution vector-based graphics scale perfectly. Assuming you have the screen real estate available, squinting to see Auto-Tune's finer details is a thing of the past. The GUI can now also be switched between Light and Dark modes, to match your personal preference or the current lighting conditions. Both the Quick Settings and Preferences menus have undergone a redesign, which make for easier operation, and new users will benefit from the more comprehensive advice offered by the revamped Tooltips system (this can, of course, be toggled off if it's not required).

Some new global controls have been added. These include undo/redo buttons, a global bypass button and, interestingly, a Mix knob, and all of these are found in the top-right region of the main window. I found the first of these particularly useful: the ability to step backwards/forwards through your history of control adjustments means you're much freer to experiment, since you know you can return to the safety of an earlier configuration should your adventures lead you astray.

Apple Silicon hardware is now supported, of course. Also new is the Input Type Learn button: press this and Auto-Tune will 'listen' to your input signal, after which a machine-learning process attempts to identify the best settings (and underlying Auto-Tune algorithm) to use. In addition, Antares have included a collection of presets created by artists who use Auto-Tune, and these provide

Antares Auto-Tune Pro-X

$459

PROS

- Pro-X brings a number of useful UI and workflow refinements.
- Auto mode is impressively flexible for natural, transparent pitch correction.
- Yes, you can still create 'that' vocal effect and you even get artist presets to help you do it.

CONS

- The temptation to create 'that' vocal effect!
- Other than the considerable investment for a perpetual licence, nothing else — this is powerful stuff.

SUMMARY

Pro-X is a worthwhile evolution of the top end of the Auto-Tune line. The most flexible all-round automatic pitch-correction available combined with powerful graphical editing when you need to use it.

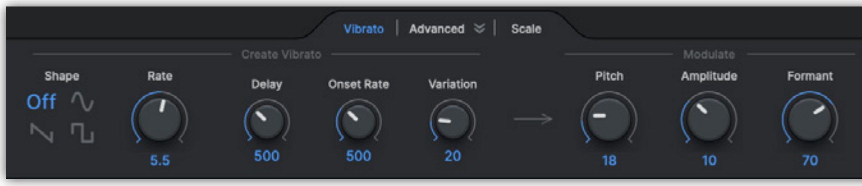

Auto mode's Advanced panel has been redesigned for Pro-X and provides further options for easy adjustment of Vibrato and Scale settings.

some useful starting points for both corrective and creative (including 'that' effect) applications.

In terms of general workflow improvements, one of the most interesting and useful new options to me is Multiview, which allows you to access any instance of Auto-Tune in a project from the GUI of any other, courtesy of a drop-down menu in the top left of the GUI of every instance. This means there's no need to close/open each individual instance or navigate between tracks. So, for example, you could have it inserted on a whole stack of backing vocals and work on the tuning of all from a single GUI. It's a very clever idea and a potentially a great time-saver.

Auto Mode

The Auto mode offers levels of subtlety well beyond the clichéd Cher/T-Pain-style hard tuning. In Pro-X, Auto mode retains the choice between Classic (based on the Auto-Tune 5 processing algorithms, upon which many producers originally built their creative vocal processing sound) and Modern styles of processing. But features such as the Flex Tune and Humanize controls mean that, even when letting Auto mode do the heavy lifting, you can finesse the processing to obtain a more 'natural' outcome.

In Auto mode you can, as before, confine yourself to the Basic view, with its intuitive and minimalist control set, or you can pop open an Advanced tab in the lower half of the window. The contents of this tab have been revamped for Pro-X, and it can now be toggled between views of either the Vibrato or Scale controls. The available controls and options are not necessarily new, but they do seem to me to be better organised. In short, they provide very flexible scale settings (including user-customised scales), and both corrective and creative options for vibrato control.

The Auto mode might be an 'automatic' mode, then, but it's beautifully conceived, delivering a scalable amount of control that provides something suitable for almost any level of user. And, while you can just set-and-forget the Auto mode controls so they apply to your whole track, do bear in mind that the settings can also be automated in your DAW. So if you do need to apply Auto mode's processing in a more selective fashion at different parts of your vocal performance, it's perfectly possible to do that.

Hard Graph

Graph mode is an offline process that's intended for those who require a more surgical approach to pitch manipulation. For this level of editing, Auto-Tune again has plenty of competition, most notably from Celemony's Melodyne but there's plenty more from the likes of Synchro Arts, Waves and, of course, similar functionality is deeply embedded within many DAWs (for example, Cubase's VariAudio and Logic's Flex Pitch processes). All of these platforms are now mature enough that the underlying algorithms are of a quality that make them capable of transparent pitch-shifting over a several-semitone range, and what most marks them out as different is their tool set and what they can offer in terms of workflow improvements.

To my mind, Auto-Tune undoubtedly has some very appealing elements to its Graph mode workflow, and these have been refined further in the Pro-X release. This includes something of a graphical overhaul of the UI — largely improvements in the way the tools are organised but, thankfully for existing Auto-Tune users, all the familiar options »

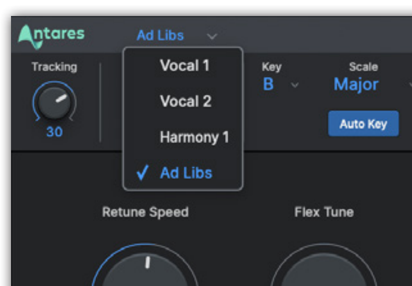

The new Multiview feature allows you to switch between multiple instances of Auto-Tune within your project from a single window.

When it comes to more detailed editing, Graph mode's different Correction Objects provide you with different options suitable for different tasks.

are still present, correct and easy to find. One of Graph mode's most appealing features is the ability to work with different types of Correction Objects: Notes, Curves or Lines. These provide targets for Auto-Tune's pitch correction and are applied to the original pitch of the audio (shown as a red curve on the display). These different types can be mixed and matched within a single editing session, such that each section of your audio can be a different object type, and collectively they provide considerable choice as to how you'll approach your manual edits. Lines are perhaps best reserved for special effects, while Notes provide a means of rapidly re-pitching sections of audio, whether for correction or to attempt something more creative, like re-writing a melodic phrase. Curves let you dive deep into the pitch processing, in a very precise fashion; if you need detailed control over a specific note or phrase, this is the way to go.

Long-time users will often develop a preference for working mostly with Notes, Curves or Lines, so it's great that Pro-X now has the option to automatically create your preferred Object type after your audio is tracked; this is configured in the Preferences panel. Meanwhile, the new zooming options provide for improved navigation in Graph mode, in which editing often requires you to get up close and personal. There are also

Subscription: Going Ultimate

The subscription model for music software has pros and cons, but that's a debate for elsewhere. If you're happy with this approach, Antares offer a couple of different plug-in bundles and Auto-Tune Pro-X is included in the larger of them. Called Auto-Tune Ultimate, this starts at $14.98 per month for an annual subscription. You get more than Auto-Tune Pro-X, though: the bundle includes a number of other versions of Auto-Tune which are streamlined for specific tasks, as well as Antares' full-suite of vocal plug-ins, such as Articulator, Aspire, Choir, Mic Mod and Harmony Engine, and the Vocal EQ, Slice and Vocodist plug-ins.

Paul White was impressed with Vocodist back in *SOS* November 2021 (https://sosm.ag/antares-vocodist). Vocal EQ is a multiband dynamic EQ, in much the same vein as FabFilter's Pro-Q 3 or Cubase's Frequency 2 plug-ins, but with an important twist: Vocal EQ includes a pitch-tracking feature rather like that of Sound Radix's Surfer EQ. This allows you to lock a specific frequency band to the fundamental pitch of the notes within the performance, or a selected harmonic of that fundamental pitch. It then tracks the frequency of the fundamental (or harmonic) as the input note changes. Like any novel processing approach, it can take some time to work out how to use it, but the potential is obvious, particularly if you want to tame a specific harmonic (or harmonics; different bands can be assigned to different harmonics) within a vocal.

As part of the Auto-Tune Ultimate subscription, Vocal EQ's Tracking feature brings a novel twist to the concept of dynamic EQ.

Slice is a virtual instrument designed for slicing and playback of vocal samples. You can, of course, load your own sample content, but the user also gets access to a library of vocal samples in a range of musical styles, some of which have been provided by name artists (for example, Bon Iver and Junior Sanchez) and with new content being added regularly. While lots of DAWs now include some element of slice-based sample playback, with a feature set obviously designed for vocal processing, a set of 14 in-built effects, and some powerful playback features/options to explore, Slice is a powerful means of revamping existing vocal recordings to create new performances. And just as importantly, it's a heck of a lot of fun to use!

a number of new helpful customisation options in terms of the Graph mode display (for example, the ability to toggle on/off the red input pitch curve or to show the waveform within a Note Object).

Best Of Both Worlds?

Pitch correction/manipulation tools come in all sorts of shapes and sizes. Some focus on easy-to-use 'auto' processing, while others place greater emphasis on more detailed editing capabilities. And some, including Auto-Tune Pro, offer the flexibility of both approaches. So, is Pro-X the best of both worlds in a single plug-in?

When it comes to the automated pitch-correction process, Auto-tune is undoubtedly the original but I think it's also pretty easy to make a case for it being the best. If your vocal just needs a little gentle nudge, it really is a process anyone could accomplish. And even if you prefer not to explore the depth and power of Graph mode, you can still accomplish a great deal in Auto mode, whether you're after corrective processing or creative effects. This was probably the case before Pro-X, but this latest version enhances Auto mode's workflow in ways that will benefit both existing and new users. And while it's very much a case of evolution rather than revolution, personally I'd see that as a positive rather than a negative; there's no point breaking what already works very well, if you can simply refine and improve it.

I think it's also true to say that when you do need a more surgical approach, Graph mode has all the tools you could require. While Pro-X delivers some useful UI and workflow refinements to this mode too, the case for 'best in class' is perhaps less clear here, given the seriously strong competition and, of course, Melodyne's polyphonic capability. Nonetheless, if detailed graphical editing is something you turn to for those particularly difficult phrases that automatic correction can't quite resolve, Pro-X is very capable and will certainly get the job done. Having the option of both Auto mode and Graph mode within a single plug-in is also very appealing.

If you're looking to buy your first professional-level pitch-correction plug-in, Auto-Tune Pro-X is undeniably a seriously good option. The only real downside is a fairly hefty price tag for the perpetual licence, though you can try before you buy (Antares offer a free 14-day trial of their full product range) and there are more affordable subscription options too, which include other goodies (see box). If you're serious about your pitch-correction, Auto-Tune Pro-X is certainly worth exploring.

Test Spec

Cubase Pro 12.0.52; Auto-Tune Pro-X (v.10.0.0 x965)
iMac running Mac OS 10.15.4, 3.5GHz Intel Quad Core i7, 32GB RAM.

$ Antares Auto-Tune Pro-X perpetual license $459 (includes one-year subscription to Auto-Tune Unlimited). Upgrades from earlier versions also available. Subscriptions which include Pro-X start at $14.58 per month.

W www.antarestech.com

Vintage Vibe Deluxe 73

Electro-mechanical Piano

Vintage Vibe's Rhodes recreation just keeps getting better.

GORDON REID

From the late 1960s to the early 1980s, you had just one choice of manufacturer if you wanted the sound of a Rhodes piano. Rhodes. Other electro-mechanical pianos existed in the form of the Wurlitzer EP200, rarities such as the Hohner Electra-Piano, the Weltmeister Claviset, and even various Pianets if you didn't mind the lack of a sustain pedal, but only a Rhodes sounded like a Rhodes, and it ruled. Then Yamaha launched the DX7, Kurzweil released the K250, and these, plus the numerous digital pianos that appeared soon after, propelled the bulky, heavy and sometimes troublesome Rhodes toward the fringes of popular music. But some players' love for it never waned and in 2007, more than 20 years after the demise of the Mark V, the Rhodes Mark 7 appeared. (I wonder what happened to the Mark 6?) If you never saw or heard one of these, I'm not surprised; it came and went without leaving a ripple on the surface of the music industry. But its spiritual successor, the Rhodes Mark 8, has now entered production, guaranteeing that the Rhodes name lives on, albeit outside of the mainstream that it dominated 50 years ago.

Meanwhile, former rental company and repair shop Vintage Vibe also released a range of 44-, 64- and 73-note Rhodes-inspired pianos. I reviewed the 64-note sparkle-top version for *Sound On Sound* in April 2013 and was impressed. Curvy and sexy like a Wurlitzer EP200 on the outside but with the soul of a Rhodes on the inside, it was beautifully built, its keyboard was more even than that of my

vintage Rhodes, and it sounded superb. Since then, Vintage Vibe have deleted the 44-note model while introducing two improvements to both the 64- and 73-note models: new hammer tips distributed across multiple zones to create an even better response, and an updated preamp with a revised EQ and tremolo circuit. But today, there's a much more significant update called Variable Voice Control, and it's this that makes it worth revisiting the Vintage Vibe piano in 2023.

Variable Voice Control

To appreciate Variable Voice Control, you have to understand how the Rhodes mechanism and the sound it produces can be adjusted to suit the player. You've probably noticed that some Rhodes pianos have a warm, rounded sound, while others of the same model and age can sound much brighter with more 'bark'. This isn't an accident, it's a consequence of the spatial relationships between the piano's tines and pickups; if you're willing to attack these with a selection of screwdrivers and spanners, you can revoice it by adjusting the height of each tine and its associated tone bar. It's not really a job for an untrained amateur because obtaining consistency across the width of the keyboard can be tricky, but it can be done. Unfortunately, this operation also changes the 'feel' of the piano slightly because the distance travelled by the hammer to strike each tine is affected, as is each tine's relationship with the damping mechanism. In the late 1970s, a company named Dyno-My-Piano simplified the revoicing procedure by devising a mechanism that shifted the whole harp assembly in relation to the pickups but, now that I've seen what Vintage Vibe have invented, I'm gobsmacked that no-one previously thought of it because it's clearly a better solution. It works like this...

The controller for the Variable Voice Control mechanism is nothing more than a slider mounted next to the preamp panel. Behind the scenes, this moves a physical ramp that adjusts the height of a short metal rod. This rod then pushes against a bar that adjusts the angle of the frame on which all the pickups are mounted. When the slider is pushed fully to the left, all the pickups are rotated so that their poles are at their lowest position, somewhat distant from the tines, and this generates a warm and rounded tone. As you push the slider to the right, the pickups are angled upward so that the poles move closer to the tines, increasing the brightness and 'cut' of the sound. This mechanism has two very significant advantages over manual adjustment. The first is that you can revoice the whole instrument in a fraction of a second — even while playing! The second is that, because the pickup assembly is being rotated and the height of the tines is unaffected, the action is unchanged. It's simple, brilliant and bloody obvious once someone else has thought of it.

In Use

To test this, I accessed the sockets on the underside of the instrument to hook up the Vintage Vibe Deluxe 73 production prototype supplied for this review. (Despite necessitating a bit of grovelling, this is a much more sensible arrangement than that of vintage Rhodeses because these have their outputs immediately behind the keyboard, requiring the use of a right-angled quarter-inch plug and some tape to ensure that the cable doesn't become entangled in your left-hand pinkie.) Having done so, I discovered that something was amiss. Everything worked as expected, except that the Variable Voice Control appeared to do nothing. After a few moments trying to work out what I was doing wrong, I decided that the right approach was to blame the instrument, find a suitable screwdriver, and open the thing up. It took just moments to remove the lid, and this revealed the problem; the rod that moves up and down was no longer underneath the bar that rotates the pickup frame, but had somehow become displaced and was wedged beside it. I have no idea how this could have happened — the force needed to do this would have been considerable. Nonetheless, it took just seconds to fix, whereupon everything worked as it should. I reported this to FX Rentals — who kindly supplied the review instrument — and later discussed it with the chaps at Vintage Vibe who told me that they have now modified the production units to prevent this from happening.

Having removed the lid, I also took the opportunity to inspect the piano mechanism. This revealed that the new model is as much a work of art today as it was in 2013. I then tested the action and was pleased to find that the keyboard was level, responsive and clatter-free. However, it's still likely to be hard work for players who have never experienced the joys of muscle development through playing a Rhodes. But if there's one area in which the Vintage Vibe piano falls a little short, it's that of the sustain pedal, which is a push-rod design (like a Rhodes) rather than the more sensible cable design (like an EP200) that allows you a degree of flexibility to position the pedal to taste. Furthermore, the pedal itself is much lighter than that of an original Rhodes and, on a polished floor, it's inclined to go walkabout. To avoid problems, keep a roll of gaffa tape to hand or do your best Grateful Dead impersonation and refuse to perform on anything less than an expensive Persian rug.

Next, I checked the revised preamp. This was reassuringly quiet, but I'm not sure that the EQ is entirely to my taste because even moderate bass boost can make the piano sound a tad woolly. Had it been up to me, I would have lowered the shelving frequency a little. I also have slight misgivings about the stereo (panning) tremolo, which is generated by a square-wave modulator so that the sound jumps from side to side rather than sliding from one extreme to the other. If one is being a Rhodes purist, this is as it should be, but I feel that a triangle »

Vintage Vibe Deluxe 73

$8799

PROS
- It looks gorgeous.
- It sounds gorgeous.
- The action is first class.
- Variable Voice Control is a brilliant innovation.
- It will outlast me, and possibly you too.
- All in all, it's a better e-piano.

CONS
- The high-frequency response and extended sustain can accentuate the natural dissonances of a Rhodes.
- The domed lid makes it unsuitable as a platform for other keyboards.
- A heavier, cable-style sustain pedal would be an improvement.
- It can wobble a little if you get energetic when playing it.
- It's expensive.

SUMMARY

It's far from cheap but, having had the opportunity to play a Vintage Vibe Deluxe 73 with Variable Voice Control, I'm not sure that I'll ever be fully satisfied with a vintage Rhodes again.

The Vintage Vibe mechanism is a labour of love and a work of art, as well as a way to make the instrument go 'doinnggg' at any of 73 pitches.

wave would have produced a more musical effect.

From the outset, I liked the Vintage Vibe's lower and mid ranges very much, but a few notes toward the upper end of the keyboard were significantly louder and more aggressive than the rest, jumping out and biting me on the ears. Since I still had the lid off, I adjusted the pickup positions of the offending notes, sliding them back a millimetre or two until each sat comfortably alongside the rest. I then checked that the Variable Voice Control was unaffected by my dabbling, which it was. Then I played again. Ah... that was much nicer.

Now we need to consider another issue with Rhodes pianos; if you play chords in the upper octaves of an equal tempered 73 or 88 you obtain a harsh, discordant initial tone. This doesn't last very long because the high-frequency components generated by an original Rhodes decay quickly, but the Vintage Vibe is designed to have a wider frequency response and a longer decay, so the effect is more pronounced. Indeed, it seemed even more pronounced than I had expected, so I grabbed a tuner and inspected each of the notes in the top octave-and-a-half. Several were a few cents out of tune, so I tweaked them and the discord was reduced. To be honest, I would have loved to have taken the time to fine-tune the whole instrument, perhaps even stretching it a little to sweeten the sound. But I doubt that FX Rentals would have thanked me, so I left well alone.

Once I had finished tweaking, it was time to stop analysing and just enjoy the instrument. Expletive deleted, it's nice! After a while, I found that my favourite setup was with the Variable Voice Control almost fully to the right and the treble EQ almost fully anticlockwise, much like a guitarist selecting the bridge pickup and turning down the treble. Adding a little tremolo then completed the sound and, in my view, the results would have graced any recording. But if I wanted a darker, smokier tone, all I had to do was whip the slider to the left, saving hours of revoicing. Happily, Variable Voice Control is also available as an upgrade for earlier 64-note and 73-note Vintage Vibe pianos, which I imagine will be of interest to many existing owners.

Inevitably, there are a few shortcomings. For example, the Vintage Vibe's unbraced Wurlitzer-style legs mean that's it a little less steady than a Rhodes, so a bit of over-energetic pounding can cause it to wobble under your fingers. Secondly, the top is domed, so it's not suitable (for both aesthetic and geometric reasons) for use as a platform for other keyboards. Furthermore, it comes without a hard lid, so you'll need a flightcase unless it's going to sit in your studio (or living room) for its entire life. It's also worth noting that there's no 88-note version, although this doesn't bother me as much as it did 10 years ago; I suspect that the 73 is wide enough for most players, especially since the very highest and very lowest notes on an 88 can be strange beasties at the best of times.

Conclusions

The Vintage Vibe piano is a beautiful instrument to look at, to play, and to hear. Sure, the e-piano emulations in modern workstations and modelled soft synths can be amazing, and I find that they can sit in a track as naturally as the original instruments. But there are some players for whom the experience contributes to the performance, and there's no doubt that sitting behind a Vintage Vibe 73 is quite different from sitting behind a workstation or a MIDI controller hooked up to your laptop. It's also quite different from sitting behind a clattery old Rhodes. So, while it may not sound or feel identical to your favourite Rhodes, I think that there's a good chance that you'll prefer it. It's not hard to understand why. In no particular order... its build quality is superb, it's lighter than the original, its action is more consistent, its active electronics are quiet, its sound is excellent, and the Variable Voice Control is a stroke of genius.

On the other hand, Vintage Vibe pianos are far from cheap and, at two to three times the cost of a restored Rhodes, even the smaller models will be beyond the reach of most players. But I don't think that anyone with the means will be disappointed. What's more, I suspect that the Vintage Vibe sitting beside me will last just as long as the 50-year-old beast alongside it. I'm not sure that I'm quite ready to view an electro-mechanical piano as an heirloom, but this one comes close.

$ Deluxe 73 $8799, Vintage Voice Control $1399.
W www.vintagevibe.com

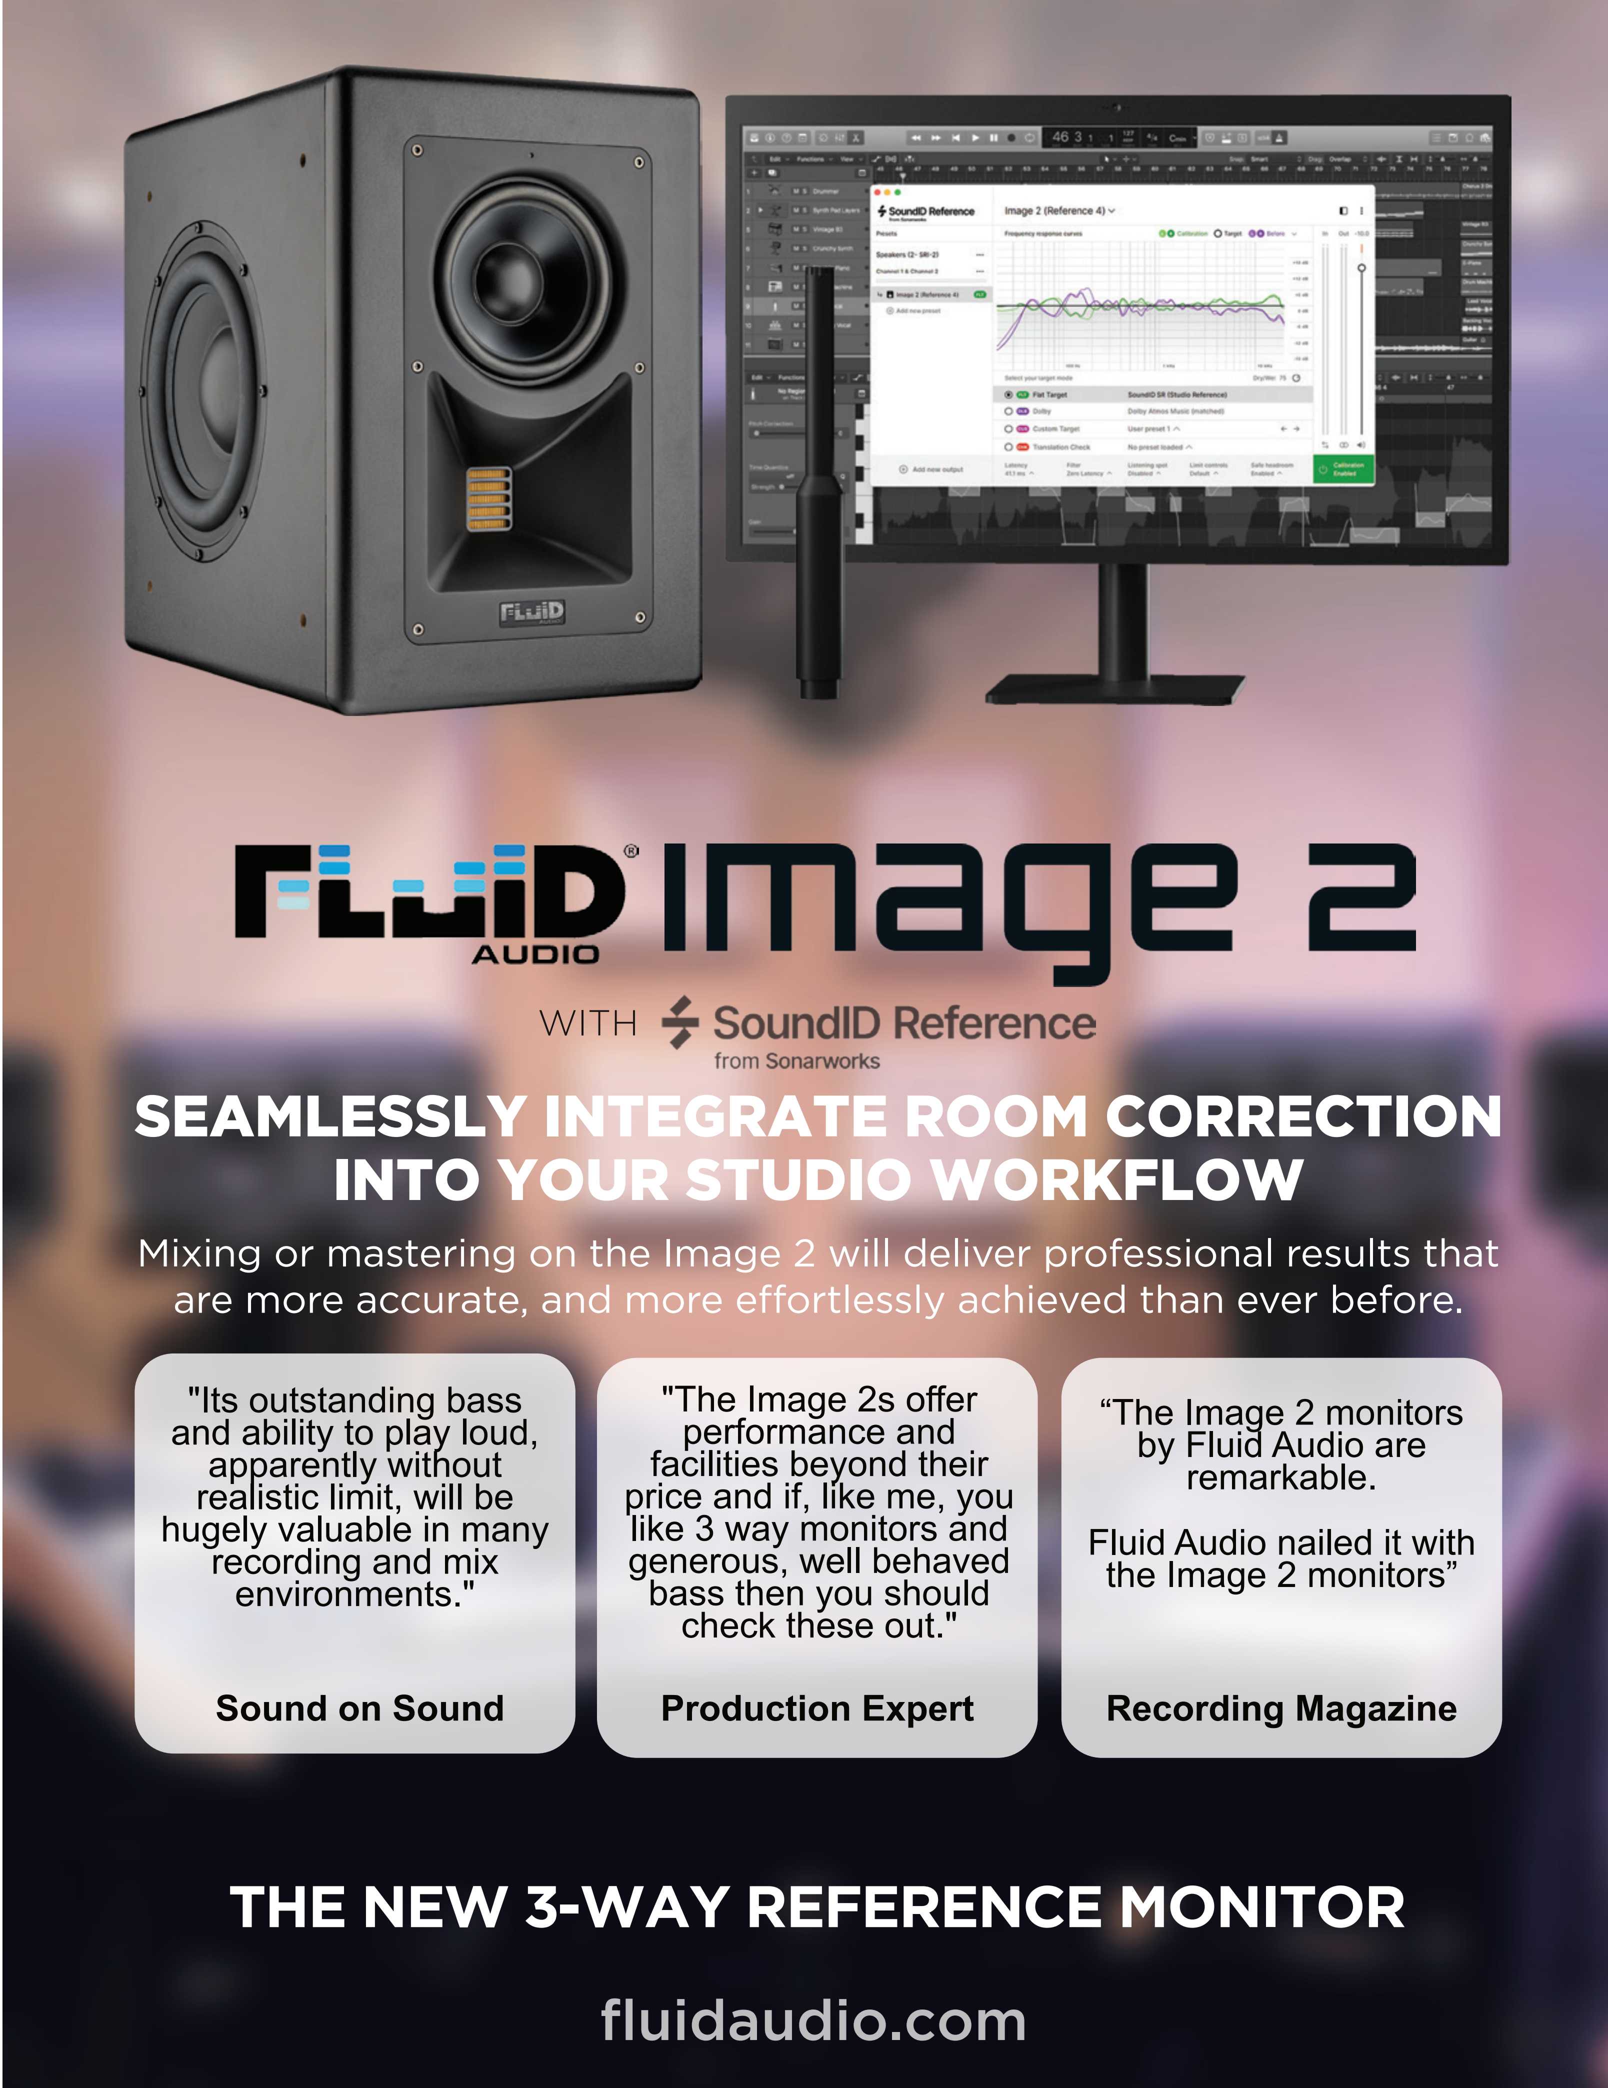
FLUID AUDIO Image 2
WITH SoundID Reference
from Sonarworks
SEAMLESSLY INTEGRATE ROOM CORRECTION INTO YOUR STUDIO WORKFLOW
Mixing or mastering on the Image 2 will deliver professional results that are more accurate, and more effortlessly achieved than ever before.
"Its outstanding bass and ability to play loud, apparently without realistic limit, will be hugely valuable in many recording and mix environments."
Sound on Sound
"The Image 2s offer performance and facilities beyond their price and if, like me, you like 3 way monitors and generous, well behaved bass then you should check these out."
Production Expert
“The Image 2 monitors by Fluid Audio are remarkable.
Fluid Audio nailed it with the Image 2 monitors”
Recording Magazine
THE NEW 3-WAY REFERENCE MONITOR
fluidaudio.com

MODULAR

ALM/Busy Circuits Pamela's Pro Workout

Eurorack Module

If the name Pamela's Workout is unfamiliar to you, you must either be new to Eurorack (welcome), or have such niche taste in modules as to purposefully avoid any recommendation of them at all. One of the most popular Eurorack modules anywhere, ALM/Busy Circuits' Pamela's Workout smashed into the modular world a decade ago, setting the gold standard for clock division and rhythmic trigger generation in the process. The follow-up, Pamela's New Workout, set that gold standard all over again, and the proof is in the pudding: it's currently the second most ubiquitous module in Eurorack according to the website Modular Grid, and ranked eighth in its list of the top modules of all time. Now comes a third iteration, Pamela's Pro Workout: a module that manages to celebrate all the functionality of its predecessors while simultaneously making them look comparatively feeble.

Since Pamela's Pro Workout builds upon such a firm foundation, it's hard to appreciate it, let alone review it, without at least a cursory look back at the Pamela family tree. Laudably, throughout this series there is next to no backtracking in design, model to model, as is often seen in similar series. The core features of previous Workouts remain true in the Pro, with new features and functions simply, joyously, slathered on top.

With the original Pamela's Workout, Matthew Allum's company hadn't just produced an efficient digital clocking solution with their very first module. Suddenly, Pam's — as it affectionately came to be known — could be found at the very heart of systems everywhere. Eight beautifully editable clocked trigger outputs of rhythmic goodness were on the bill, along with a four-digit, seven-segment display. Far more than simply a clock division and multiplication module, triplets, pulse-width editing, random skipping and much more were all achievable with ease.

"Pamela's Workout grew out of a need to flexibly synchronise various clocked modules and external devices," reads the introduction to the original's manual, "but do so in a flexible and feature-full manner as to also enable more creative and playful rhythmic usage." It's that very maxim that defines the Pamela's Workout dynasty, blurring the line between utility and creativity; the primary takeaway from this review if you're considering buying one, at the end of the day.

Pamela's New Workout then pushed all of its predecessor's functionality to the next level, retaining the core principle of editable clock divisors but adding a multitude of new tools. Most conspicuously, instead of simple pulses it offered waveform selection, meaning it could suddenly be used as a clocked modulation source or cycling envelope generator, complete with editable cycle width, level and offset. On top of this, it expanded on the original's rhythmic capabilities considerably, offering programmable Euclidean rhythms, beat-based looping and more. Twin CV inputs could be assigned to control any of the output parameters, allowing for integration with a wider system or clever self-patching, with all of this complexity saveable to any one of 200 banks. Most conspicuously, Pamela's New Workout now offered a detailed LED screen, making navigation far easier than on the original.

Once again, the most obvious defining characteristic on Pamela's Pro Workout is its display, which

is around twice the size of that of its predecessor. Crisp and full-colour, all the Pro's parameters are now presented in a spacious and clear environment, unlike the occasionally cramped display of the New. It's even possible to customise the UI colour scheme. Beyond simply making things easier, ALM/Busy Circuits have made sure to use the screen to the maximum of its ability, including things like an oscilloscope with adjustable resolution to visually track the movement of any output. It also makes things like Euclidean step programming far easier, visualising the number of steps and subsequent triggers instead of prompting the occasional quick mental maths. There's enough space here to prompt

> **"It may share the 8HP form factor of the original Pamela's Workout, but the power of the Pro means it may as well be 10 times that size."**

imaginings of even more ambitious uses for the screen: a dashboard-style display for each output, perhaps, with a scaled down oscilloscope, time division and modulation all displayed together. For now, though, it's a case of more or less one piece of information at a time — big, bright and sharp.

Across the board, Pam Pro exponentially improves and adds to the functionality of the Workout series. It has a faster CPU with dual-core processing, outputting at a 12-bit resolution compared with the New's 10. While the New could go from clock divisors of 512 to multiplications of 48, the Pro can go from a staggering /16384 all the way to x192. There are now many more waveforms to choose from, including trapezoid, hump and a choice of exponential or logarithmic envelopes. The Run and Clk inputs can now be repurposed to make a total of four CV inputs. There's slewing, loop 'nap' and 'wake', and cross-modulation between any outputs. It also offers wave inversion and two different types of ratcheting pulse waves — both examples of things you could achieve with the New in various ways, but here presented on a platter to make things quicker, easier and ultimately more fun.

The later firmware updates that pushed Pamela's New Workout further are of course present in the first instance with Pam Pro, only significantly rebooted. One major example of this is the quantise function, introduced relatively late on in the New's tenure to allow any waveform to be quantised to one of a host of scales: this could amount to anything from a random stepped voltage jumping around a major scale to a saw wave cascading through a pentatonic scale, like an 8-bit games console. Once again, this function is given a significant boost on the Pro, which not only uses its screen to visualise the notes as they would appear on a chromatic keyboard but allows for user-preset scales by cycling through an on-screen keyboard with the encoder and selecting the desired notes. The smooth random wave that appeared in a later firmware update on the New is present and correct, here displayed in all its 'Mario hills' glory. I must say, before Pamela's New Workout I hadn't heard the phrase 'Mario hills' used to describe a waveform before, and very much hope it's an ALM-coined phrase. In any case, you may consider it formally inducted into the accepted Eurorack lexicon from here on out.

It may share the 8HP form factor of the original Pamela's Workout, but the power of the Pro means it may as well be 10 times that size. It's now so far beyond the realm of triggers and clock divisions as to encroach across the border into other territories entirely, no least that of sequencer and digitally controlled oscillator. Where can the series possibly go from here? Right now, it seems, almost anywhere. Perhaps into the molecular multiverse inhabited by the likes of Expert Sleepers' Disting. Here on Earth, though, one thing is certain: when it comes to clocking, modulating and a whole lot more, Pamela's reign is far from over. *WIlliam Stokes*

$ $339
W www.busycircuits.com

2HP Slice

Eurorack Module

Founded on the principle that good things come in small packages, 2HP specialise in making modules to usefully fill those awkward spaces in your rack. And yes, they're all 2HP wide, which is just 11.16mm. Some jack plugs are wider than that...

So what does this particular 11.6mm do? Well, it records audio and plays it back as a loop. This loop is triggered by the Trig button or a trigger signal, either in a momentary or latched manner. The size of the loop is controlled by the Size knob or incoming CV, and is variable between a leisurely two bars and a high-pitched 1/256. And those divisions are relative to the incoming clock because, unlike 2HP's other audio repeater, Freez, Slice is clockable. That means that in-time stutter, ratchet and glitch effects are yours to command. It also means that you can S&H the bejesus out of the loop size, feed whatever nonsense you like into the clock input and rejoice in the time-based noise chaos that ensues. All good fun. 2HP describe the Slice as "click-less", the result, they explain of "short, envelope-based windowing on the slices", which is useful, but likely to be of greater interest to the subtle stutter crowd than the noise chaos brigade.

As with all 2HP modules the size is a curse as well as a blessing. With all five sockets occupied it can be difficult to reach the Size knob, Trigger button and Triplet switch (there's a Triplet switch!), but this is more of a 'control with CV' module than a freestyling knob-manipulation performance thing and so this is less of a problem than it could be. As ever, though, the price of entry is low and the bang for HP buck is unbeatable. If you want flashcore mayhem or just the occasional tasteful stutter effect, 2HP's Slice will get you there. *David Glasper*

$ $149
W www.twohp.com

Modular Profile: Matthew Allum

As the man behind the phenomenally successful British developer ALM/Busy Circuits, Matthew Allum's impact on the Eurorack landscape at large has been, to say the least, considerable. His company recently celebrated its 10th anniversary — no mean feat in what has in recent years proven to be an unforgiving marketplace — marking the occasion with a third iteration of the powerful, vastly popular and perfectly named Pamela's Workout, also known as the module that started it all for Busy Circuits.

On his entry into modular
Around 2010 I got a Doepfer Eurorack system after being disappointed with the then-new OP-1. Since I was a teenager I'd dreamt of owning a modular system: Doepfer really made this accessible. A frustration trying to sync that with an MFB sequencer, as well as some older DIN Sync-based equipment, motivated me to create a clocking module to solve this problem. I based it on an [*open-source prototyping platform*] Arduino that had been gathering dust on the shelf, together with some basic electronics knowledge. I have a technical background, but back then it was more software based. The surprising interest from demoing it at an early Brighton modular event encouraged me to do a small run of 50 of these 'Pamela's Workouts'. I needed to sell approximately half to break even, but I figured if nothing else it was fun to learn about manufacturing and creating a physical product. Fortunately the 50 sold out within a week and the Busy Circuits journey began!

On his go-go modules
My go-tos include the Doepfer A-111-6, a fun, great-sounding bass synth in 12HP; the Feedback Modules CR, KM and BX Pre-mixers, which are essential for magic drum coloration; the Super Synthesis 2OPFM, a simple, cheap FM VCO — perfect clonks! Also, the ADDAC System ADDAC112 VC Looper & Granular Processor. It's the ultimate granular monster.

On 10 years of ALM/Busy Circuits
This past November we celebrated 10 years of ALM/Busy Circuits. Over that time it's grown from a part-time hobby to a business with employees and offices in London and Chicago — though, small ones! I'm extremely proud of the modules we've put out over that time. Not just Pam and more bread-and-butter utility modules, but also the more ambitious re-imaginations of classic hardware, brought into the Eurorack environment: digital FM, like Akemie's Castle, classic rackmount sampling like the Squid Salmple and digital effects processing like the MFX. All done with emphasis on functional, good design, usability and fun. We've also done our first complete system, the System Coupe, and our new case range is doing well.

We've made many new friends and have got to travel to shows around the world, meeting new people. We've done collaborations with artists I admire, as well as other manufactures like Worng Electronics. We've also built up demo and education content across social media platforms, resulting in a well-respected Instagram account where users can get immediate inspiration and ideas to get the most out of our modules, plus longer-form content on our YouTube channel.

On Pamela's Pro Workout
The Pro is the third Workout iteration; each one being a complete overhaul of the last whilst keeping the form factor and core interaction model. Where the New Workout introduced modulation features on top of clocking, the Pro upscales everything with big improvements to usability and performance as well as new pattern creation features like cross-output modulation and 'off-grid' flex timings. Firmware updating is now trivial too — you just drag and drop via a USB-C connection — so we have plenty of scope to add new features over time, as we've done with other modules like the Squid [*Salmple*] and the previous Pam. That is, until we ran out of space, which was one of the motivations for the upgrade!

On the culture of modular
Eurorack used to be an obscure, geek-only oddity, but it's now commonplace in studios, with users approaching and using it in so many different ways. I feel this is parallel to the culture, which was once centered around a single Internet forum but is now much, much broader and wider-reaching. Events like Superbooth are hugely popular and also mirror this. I always thought Eurorack would maybe spawn its own musical genre, but it seems that instead its just seeped into everything! And rightly so. It's such a fun, creative technology — both audibly and visually. *William Stokes*

Dangerous Music 2-BUS-XT

16-channel Analogue Summing Mixer

This classy device boasts some impressive specs and has applications that extend way beyond summing.

MATT HOUGHTON

Launched in 1999, the 2-BUS was, according to Dangerous Music, the world's first standalone summing mixer. Back then, the music production world's migration from consoles and tape machines/digital multitrackers to DAWs and audio interfaces was nascent, and in the ensuing decades things have changed considerably. We have access to computers that can run many plug-ins, and those plug-ins can do more and sound better. What's more, the digital summing in our floating-point DAWs is pretty darned good. If you wish to mix inside the box, it's perfectly possible.

But lots of us still see (and hear) plenty of value in a hybrid setup, recording to the computer but coming out of the box to use high-quality analogue processors. I tend to mix with a lot of subgroups and often run them and the master bus through outboard compressors, EQs and saturation boxes. When working with outboard, there's potentially a role for a good analogue summing mixer. With everything on a patchbay it's easier to create different signal paths than when you set things up as 'external plug-ins', and this approach bypasses unnecessary converter stages too. A high-quality summing box with plenty of headroom also means you needn't think too hard about the levels at each stage: just tweak controls until things sounds right, keep a loose eye on your gear's meters, and you only have the level at the final A-D stage to think about. I've used the eight-channel summing of my Dangerous Music D-BOX for a number of years, but have often wondered about a setup with more channels, so I was keen to check out the company's new 2-BUS-XT.

Dangerous Conversation

If you wish to perform summing in the analogue domain without a console, you have a broad choice of active or passive summing boxes. The latter can be tempting, as they offer reasonable performance for very little outlay, but you need amplification and if you demand the best performance they cannot match a well-designed active summing amp. For example, there's no common-mode rejection and no isolation of the D-A converter's ground from the audio ground of the summing mixer, two factors which will limit how low the noise and crosstalk can be made.

Dangerous Music
2-BUS-XT
$1999

PROS
- Superb build quality.
- Much more than a summing mixer!
- Impressive on-paper specifications.
- Equally impressive subjective sound quality.

CONS
- Master bus insert might be nice.

SUMMARY
Dangerous Music have updated the final summing amp in their range to a very impressive 'second-gen' spec, and some very appealing tone-shaping facilities have been added in the process.

Dangerous Music have long offered some of the best active summing devices, in terms both of specs and build quality. They also have a wonderful habit of moving with the times, and were quick to see that engineers wanted more from such devices than the summing alone. My first-generation Dangerous D-BOX, for example, now 14 years old, combined an eight-channel summing amp with a range of other facilities (source selection, great D-A conversion and monitor-control facilities). Even today, it holds its own: it presents a lovely clean and detailed sound stage, with distortion well below audible levels, crosstalk below the noise floor and generous headroom.

Good as the D-BOX and other early products were, Dangerous have gradually been refreshing their range, developing new facilities and refining existing ones, as well as further improving the technical performance. Alongside a range of other high-quality gear, including monitor controllers, mastering routers, A-D and D-A converters, a compressor and an EQ, there are now three second-generation summing devices. In *SOS* April 2016, Frederick Norén reviewed the 2-BUS+ (www.soundonsound.com/reviews/dangerous-music-2-bus), and in August 2019 Hugh Robjohns looked at the D-BOX+ (www.soundonsound.com/reviews/dangerous-music-d-box-plus). The noise floor, distortion and crosstalk of these devices is even lower than on their predecessors. As Frederick explained, one factor which made this possible was the availability of smaller components, which allowed designer Chris Muth more freedom in laying out the circuit: think shorter, straighter lines and six-layer PCBs.

Ins & Outs

The 2-BUS-XT appears nominally to have replaced the 2-BUS-LT, yet draws just as heavily on the 2-BUS+. As you'd expect, it boasts class-leading technical specs, and the build quality is nothing short of exemplary. The signal path is entirely analogue, but digital control is involved and, as with my trusty D-BOX, some buttons perform different functions in response to short and long presses. Unlike the D-BOX, there's an internal power supply which receives mains AC through a rear-panel IEC inlet. This caters for US and European

The X-Former and Coherence effects, the latter a parallel processor with its own level knob, are available for the stereo bus or channels 15+16.

voltages, though you must fit the appropriate fuse if moving from one region to the other.

At the heart of the 2-BUS-XT is a 16-channel active summing mixer, whose balanced inputs are presented on two rear-panel AES/Tascam standard DB25 D-sub connectors. On the front, every channel has a green Signal Present light, and they're demarcated as stereo pairs. For each DB25's first pair (1+2 and 9+10) a dedicated Mono button routes both channels to the left and right mix busses, placing both signals in the centre (there's no pan facility). The 2-BUS-LT had this facility for every pair and some might miss that, but it's not often you'll need more than four mono channels right down the centre. One thing to watch out for is that if you send a stereo signal to the channel pair and hit mono, the balance between centre and hard-panned sounds will change; I noticed this when trying to 'collapse' a stereo drum and bass bus to mono.

On the rear, two balanced XLR pairs provide stereo outputs Main Out and Mon Out, which are identical. Their level is determined by the Sum Level Trim pot on the front panel, and this can be set anywhere from full attenuation (anticlockwise), through unity (a shade past 1 o'clock) to a considerable amount of gain when fully clockwise. This allows you to set an appropriate level for your recording device or, potentially, to 'drive' a master bus processor that lacks its own input gain control. There's no facility to monitor the effect of any external stereo-bus processing, though to be fair Dangerous do point out that the 2-BUS-XT is really designed to be used with other gear (eg. routers, monitor controllers) that make that possible, and other summing devices in their range possess master-bus inserts.

A third XLR pair, Exp In, is intended to receive the output from another Dangerous Music summing mixer, to increase the number of inputs. During my tests, I fed the output of my D-BOX's eight-channel summer to this Exp In, giving me 24 channels of summing in total, and then routed the 2-BUS-XT's Mon(itor) output to the D-BOX's Analog In XLRs, while routing the Main Out to »

Talking Crosstalk

The published specifications are impressive. The whopping +27dBu maximum input level means the 2-BUS-XT can accommodate the signal from pretty much any D-A converter. The frequency response is within 0.1dB of flat, way beyond the audible band (10Hz to 50kHz). And total harmonic distortion (THD), intermodulation distortion (IMD) and noise are commendably low. Dangerous seem particularly proud, though, of the vanishingly low levels of crosstalk: rejection is given as >109dB at 1kHz. To put that figure in context, at the same frequency SPL's MixDream achieves 97dB and Rupert Neve Designs' Orbit 103dB. But crosstalk usually rises considerably as you move higher up the frequency spectrum. I didn't have the opportunity to measure this myself, but Dangerous kindly shared their own Audio Precision plots with me, which compared this device with a top-performing competitor, and showed the 2-BUS-XT achieving ≥100dB right through the audible band. Could this be a reason it sounds so good? Perhaps. It's hard to be sure, but it's the best performance I've seen and such figures aren't achievable without a huge amount of care and effort going into every aspect of the design and build.

The Exp input is designed to accept the output from another Dangerous summing mixer to create a larger analogue summing system — but you could use it simply as another stereo input, taking the total at mixdown to 18 channels.

my A-D converters. In practice, you could patch in the main out from a mixing console or a feed from your DAW/converters, or just use them as two extra inputs.

A Little Sum Thing Extra

The 2-BUS-XT also incorporates two switchable 'analogue coloration' circuits, which can be applied to channels 15+16 or to the stereo mix bus (or neither). These were not present on the 2-BUS-LT and they're most welcome additions!

The first, labelled X-former, runs the signal pair through output transformers (a model they've not used previously, and which I'm told was made after yet more listening tests). Unlike on the higher-priced 2-BUS+, this effect is simply switched on or off: there's no gain control to 'drive' the transformers harder. But that can still be achieved by running hotter levels into the 2-BUS-XT. Essentially, what this processor does is add a pleasing blend of odd- and even-order harmonics, which means it has a bigger subjective impact on some sounds than others. The effect is always pretty subtle, as befits a processor intended for master-bus or subgroup processing, but it's noticeable nonetheless, often manifesting itself as a pleasing thickening of the sound that adds a little 'weight'. Interestingly, I found myself wanting to use it not only for processing mixes but also for printing the effect on individual sounds in the mix; vocals and electric guitars, for example. It's a nice bonus that a 'summing mixer' can do other jobs in a studio!

The contribution of the other effect, mysteriously labelled Coherence, is much easier to discern. It's a parallel processor, with a button engaging the effect for the desired channels and a pot in the Mix Bus section determining how much of the effect is mixed in. Dangerous say it "injects dimensional asymmetry" and can "subtly tighten up the source by pulling down spiky transients, gently widening and tilting the track forward without any phase skulduggery". Make of that what you will! Again, it seems to introduce harmonics, though a different blend than the transformer effect. Subjectively, it serves to pull things into focus, brightening them in a fairly relaxed, natural sort of way. I loved using Coherence on an electric guitar bus in a thickly layered rock mix; it was perfect for injecting a little life and presence into the wall of rhythm and lead guitar parts, without treading on the toes of the lead and backing vocals. And on a picked electric bass it was easy to shift the emphasis toward the pick and 'twang'. It's nicely judged, though if applying it to individual sources (again, a lovely effect to print) you should be careful when making comparisons, since it inherently adds level. That's less of an issue when processing subgroups or the mix bus at final mixdown, where it's much more a case of "if it sounds good..."

> "I loved using Coherence on an electric guitar bus in a thickly layered rock mix; it was perfect for injecting a little life and presence into the wall of rhythm and lead guitar parts."

Switching of these effects is achieved with two pairs of buttons, one for each processor on channels 15+16 and another for the stereo bus. With a simple short press of the processor's button, the effect is engaged for that channel pair (when not engaged, it's a true hard-wired bypass). A longer press will apply the effect only while the button is pressed; when you let go it reverts to the previous state. So you could long-press to audition the difference between applying the transformer to the full mix or, say, only the drum bus. Again, a nice touch that makes life easy.

Adding It All Up

I've loved working with the 2-BUS-XT. On a technical level, there may not be much wrong with in-the-box summing today, but there's something effortless and forgiving about working out of the box, and I still hear a benefit when using good outboard. Once I'm there, it's great to have the option of ultra-clean, quiet summing, and it doesn't get cleaner and quieter than this. With 16 channels available I was able to do much more than I've been doing with my eight-input D-BOX, and the two worked really well in tandem. But this is so much more than a summer: it's a very capable sonic sweetener too, and I love the dual-function buttons.

Is there anything I'd change? Not much. In fact, the only thing I really longed for was a master-bus insert send/return, so I could audition the effect of any master bus processing via the Mon Out from the front panel. But if that bothers you, Dangerous offer precisely this facility in the 2-BUS+ already! In short, then, if you're in the market for a simple yet very high-quality 16-channel summing mixer, and the tonal bells and whistles appeal, this one's very hard to beat.

$ $1999.
W www.dangerousmusic.com

ALTERNATIVES

High-quality active summing mixers with 16 or more channels are also available from **A-Designs**, **Neve**, **Rupert Neve Designs**, **Heritage Audio**, **Speck Electronics**, **Tonelux**, **Looptrotter Audio** and **SPL**, amongst others. But none outperform those in **Dangerous Music**'s portfolio.

2400 Audio Imperium NG

Passive Monitor Console

This digitally controlled monitor controller boasts an impressive array of features.

HUGH ROBJOHNS

Danish company 2400 Audio have been in business since 2013, producing passive monitoring controller systems of one sort or another, all with the underpinning premise of passing pure, unadulterated audio using an entirely passive, balanced signal path, but with digital control for versatility and the possibility of remote control. The Imperium product name has pervaded the company's entire range, first with the Legend Imperium models, then with Legacy Imperium units, and now the latest 'Next Generation' Imperium NG versions. Pleasingly, the two previous Legacy Imperium models can, if desired, be upgraded to the latest Imperium NG standard, which speaks well for long-term customer support. All units are hand-built in Denmark and the company offer a variety of customisation options.

Two physical versions of the Imperium NG are available: a model with a fixed I/O capacity housed in a 1U rackmount case, and a larger 2U rackmount version which has more and expandable I/O. The core features of both systems are very similar though and, as mentioned above, there is a wealth of customisation and configuration options available for both models, which I'll cover later. For this review I was supplied with the NG 1U unit fitted with the Mastering, Wi-Fi, and USB options.

Overview

The core Imperium NG 1U is a rackmounting unit extending around 205mm behind the rack ears. Most of the rear-panel audio I/O connectivity is via 10 XLRs which provide two stereo inputs and three stereo outputs — all analogue, balanced and line-level. In addition, two pairs of quarter-inch TRS sockets provide a stereo, balanced send and return loop for external signal processing, while a further pair of TRS sockets labelled as Hybrid I/O serve as a user-configurable stereo input (default) or output. This last facility allows the unit to be configured with three stereo inputs (the standard condition) and three monitor speaker outputs, or two stereo inputs and four stereo outputs — whatever is required to suit a specific installation, and many will find that setup flexibility very useful.

Supplementary rear-panel connections include a B-type USB 2.0 (to host) socket, MIDI in and out on a pair of 5-pin DINs (for remote control), an RJ45 socket, and a 12V DC coaxial power input (a 2A/12V wall-wart PSU is included). However, all is not quite as it may appear since the RJ45 port is not compatible with standard Ethernet; it's actually a customised 'link port' intended for the connection of future 2400 Audio products! Moreover, the USB socket only works if the optional USB card is installed. The wall-wart power supply is only there for the digital control electronics — there is no active circuitry in the audio path at all — but, as the signal path is created entirely via relays, without the digital control you can't connect an input to an output, let alone adjust the volume!

Since I've already mentioned configuration options, let's look at those now. The NG 1U model can be enhanced with options named: Wi-Fi, USB, Mastering,

2400 Audio Imperium NG
$930

PROS
- Very high-quality digitally controlled passive signal path with precision relay-switched attenuator.
- Unique multi-platform remote-control capabilities.
- Configurable with a range of useful options.
- Mastering option adds application-specific extra features.

CONS
- None of any practical relevance to most users.

SUMMARY
This very high-quality passive monitor controller has a relay-switched volume attenuator and, making it unique, is remote-controllable from computers or mobile devices.

Trinnov, Barefoot MEME, and Custom Caps. The NG 2U model expands slightly on that list with additional stereo inputs (up to four) and outputs (up to eight).

The easiest option to explain is Custom Caps: a selection of individual button caps with alternative labelling to replace the standard front-panel fitment — and owners can request their own personalised designs, too. By way of example, the review model was supplied with several alternative button caps including “Shhh...”, “REAL LOUD”, and even a ‘skull and crossbones’ graphic!

Predictably, the Wi-Fi and USB options permit remote control of the Imperium NG from a computer or compatible mobile device like an iPad, using the multi-platform TouchOSC app as the user interface. 2400 Audio provide convenient start-up configurations for TouchOSC, but the system is totally user-configurable if you want a custom setup. Of particular note, the computer USB option opens another useful possibility: the review unit was supplied with an Elgato StreamDeck USB control panel, which can be configured to provide one-button access to any Imperium NG monitor-control features and functions, extending those available via the unit’s own front-panel controls.

As you might anticipate, the Trinnov and Barefoot MEME options allow direct interfacing with those particular manufacturers’ systems, controlling Trinnov’s ST2 room correction system’s preferences and Barefoot’s MEME monitor emulation modes, all via the same TouchOSC remote-control system that controls the Imperium.

Hardware Controls

Looking at the physical hardware unit, a power on/off switch sits at the left of the front panel, although this only controls the internal DC supply; the wall-wart PSU obviously remains powered while plugged into a live mains socket. Moving right, the lower row of three buttons selects one of the three available inputs (there’s no mixing facility), while the upper row selects various monitoring conditions. In the standard unit these buttons access a Parallel Output function (assigning the monitored signal to two output destinations simultaneously, for example to feed a subwoofer along with the main stereo monitors), activate the insert loop, and bypass the volume control attenuation (with suitable protective interlocks!)

In the review unit, which was fitted with the Mastering option, these default button functions were repurposed to provide direct access to features more commonly required in mastering. However, none of the standard unit’s facilities are lost in this rearrangement as they can all still be accessed via the TouchOSC interface. So, on the review unit the Parallel Output and Insert functions were replaced with dedicated Cut-Left and Cut-Right buttons, while the Level Bypass mode was replaced with a Swap L-R option. (The original button caps were included with the unit.)

Monitor controllers typically use a rotary volume control, of course, but the Imperium NG doesn’t. Instead, spread across the centre of the front panel is an array of eight numbered, programmable buttons, which allow instant access to various user-defined and instantly repeatable listening levels. The company argue that this approach provides a much better workflow than the traditional rotary volume controls: since most of us use only a few reference listening levels it makes good, practical sense in offering those as guaranteed, easily repeatable button presets. Also, in most monitor controllers the volume control is performed by a dual potentiometer, which inevitably suffers stereo mis-tracking, particularly at the quiet end of the range. The Imperium NG implements volume control using combinations of relays switching closely matched resistors, giving 256 discrete attenuation levels with perfect stereo imaging at all levels. One rather nice feature for multi-user installations is that the Imperium NG allows up to eight different ‘Level Profiles’ to be stored, so different engineers can recall their own personalised volume button preferences, if required.

Over on the right-hand side of the front panel is another group of six buttons, with the lower bank selecting the output destinations (again, only one at a time unless the Parallel Output option is activated). If an extra output »

The rear panel’s XLR connectors cater for two stereo sources and three stereo destinations — but these are only part of the Imperium’s story!

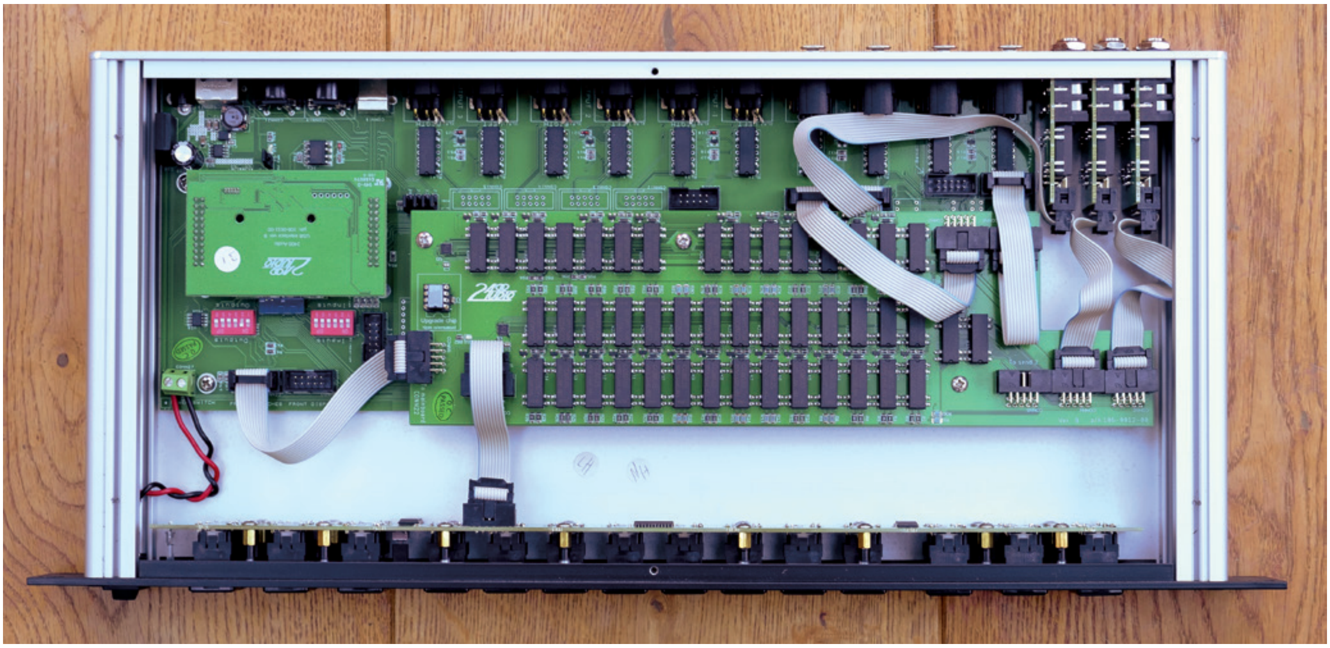

Though it's digitally controlled, the signal path is analogue and entirely passive.

destination is required the Hybrid I/O facility can be configured accordingly, in which case the default Input 3 becomes the new Output 4 (and a different button cap can be fitted for identification, obviously).

The upper row of buttons normally activate a left-channel polarity inversion, and select mono and cut (mute) modes. Sensibly, pressing the Polarity and Mono buttons together auditions the stereo difference or Sides signal (left minus right), but that task is made easier with the Mastering option installed, as the Polarity button is replaced with a dedicated Stereo-Difference (DIFF) function, giving a more convenient single-button press. Usefully, the Mono button can be programmed to activate automatically when a specific output destination is selected, if required, so feeding a mono-check speaker is easy to set up.

TouchOSC

Although the core functions of the Imperium NG can be controlled directly from the front panel, 2400 Audio actually designed the unit with remote control very much in mind, primarily using the TouchOSC app.

ALTERNATIVES

The **Coleman Audio M3PH MkIII** is another passive monitor controller, broadly comparable in features and price. An active alternative at a similar price is the **Dangerous Music Source**.

(This must be purchased separately, but a voucher is included for a 15% discount on the Mac OS/Widows version.)

Hexler's TouchOSC app provides full cross-platform compatibility (Mac OS, Windows, Linux, iOS and Android) and is completely user-configurable — although a standard TouchOSC control layout is downloadable to make it easy to get up and running with the Imperium NG. More importantly, TouchOSC provides a very simple, effective and low-cost way to access and control considerably more features than would ever be practical to control from 1U rackmount hardware.

As already mentioned, the link between the Imperium and whatever is hosting the TouchOSC app can be via the built-in MIDI I/O using a MIDI-USB interface (although control information is passed as SysEx data and apparently not all USB MIDI interfaces accommodate it correctly). Alternatively, either the USB or Wi-fi connections can be used if one (or both) of those options is/are installed. Controlling the NG directly from an iPad or smartphone is a unique feature which might appeal to users who move around the studio a lot.

The default TouchOSC layout presents a 'Monitor Console' tab displaying all of the Imperium's functions in a straightforward way as visual buttons, and clicking or pressing on them activates the corresponding function. Input sources are arrayed on the left, output destinations on the right, and the various auditioning modes are spread across the centre, duplicating and extending the functions of the hardware unit. The eight preset volume buttons are presented at the bottom of the GUI along with a virtual slider acting as a variable volume control (while also displaying the current attenuation level in decibels).

A second tab accesses the eight programmable Level Profiles mentioned earlier, and with the Mastering option installed different profiles can be assigned to the different outputs, too, which is helpful in level-matching different speaker systems. A third tab opens three sub-pages: Scenes, Assign Table, and Advanced. The Scenes page allows preset combinations of all Imperium settings to be stored, facilitating fast A/B comparisons of different setups, while the Assign Table page allows additional functions to be associated with the Input and Output buttons, such as automatically selecting mono or a specific level profile to particular outputs.

The Advanced tab contains several more sub-pages to configure various set-once-and-forget features to suit specific installations. These include how the insert connections work: conventionally, or bypassing the attenuation to allow external volume control, or as a permanent live send — perhaps to feed an external meter display, for example. There's also an option to decide whether the listening level jumps or ramps up/down to the new volume (ramping is only available with the Mastering option installed). The reallocation of front-panel button functions is also determined here. With the Trinnov or Barefoot options installed, additional sub-pages configure the remote-control capabilities for those devices and, finally, a full factory reset can also be initiated from the Advanced tab.

Impressions

The Imperium NG is a well-designed and impressively versatile passive monitor controller, and made all the more flexible thanks to its elegant digital control paradigm and the use of third-party controllers like TouchOSC and StreamDeck.

Arguments rage on social media as to the pros and cons of passive and active monitor controllers, but the truth is that it's all about the design and implementation, and there are good and bad examples of both formats. In the case of the Imperium NG, the sound quality is as good as it gets; there's literally nothing between the source and destination apart from a few sealed relays and some precision resistors. And for anyone who dislikes chattering relays, the ones used here just chirp quietly rather than click noisily! With an entirely passive signal path there is no underlying mains hum, no headroom limitation, and no amplification artefacts of any kind, while the precision-matched relay attenuator ensures perfect channel matching at all listening levels and precisely repeatable volume settings.

However, it also means that frills and extra functions are absent. There's no digital input D-A, and no headphone monitoring or built-in talkback functions, for example. There's also no output buffering, which means that routing a signal to multiple outputs inevitably results in a slight level loss to all destinations. There is no level correction when summing to mono, either... but do these things matter to most users? Probably not.

As a straightforward studio monitor controller, the Imperium NG has a lot going for it, especially in terms of its innate sound quality and flexible controllability. The Mastering option adds a couple of features which improve its suitability in that environment, too, although one function I consider essential in a mastering application, but which is currently missing, is the ability to trim input levels to allow level-matched A-B comparisons. There is a workaround using the Scenes facility, but it's a rather slow and clumsy workflow.

However, 2400 Audio could address this shortfall quite easily in a future firmware update simply by allowing the user to store (temporarily) a level offset value for a selected input, such that switching inputs also adjusts the attenuator setting appropriately. The Mastering option already implements exactly that capability to level match different outputs (through the level profiles function), so it's not much of a stretch — and having raised this with 2400 Audio I'm told it's already under consideration for a future update. Indeed, the company have a history of introducing useful updates to expand functionality quite regularly.

Price-wise, the Imperium NG offers an attractive solution in comparison to other high-end monitor controllers, but with very versatile remote-control options which give it a unique advantage over more traditional designs. I was impressed with the quality, capability and versatility of the Imperium NG and it's certainly worthy of adding to any shortlist if seeking a high-end monitor controller.

$ 2400 Imperium 1U NG base unit $930. With Mastering, USB and Wi-fi options $1520. TouchOSC under $20.
E sales@2400audio.com
W www.2400audio.com

Groove Synthesis 3rd Wave

Wavetable Synthesizer

Californian start-up Groove Synthesis haven't just recreated the PPG Wave, they've made it even better.

RORY DOW

If the picture of this synthesizer is causing you *déjà vu*, it's for a good reason. The 3rd Wave is a spiritual successor to the PPG Wave 2. Even down to the beautiful blue colour, Groove Synthesis have gone to great lengths to emulate the iconic wavetable synthesizer's sound, feel and good looks. But why stop at a straightforward recreation when you could improve on the old one?

Groove Synthesis are a new company of highly experienced synth builders. Bob Coover, a name you may recognise from his work with Dave Smith on everything from the Prophet-12 to the OB-6, founded the company. The PPG Wave is a dream synth for Bob, and with his new company, his assembled team of synthesis experts and his experience doing DSP for Sequential, he set about making the ultimate reboot.

Groove Synthesis 3rd Wave

$4995

PROS

- It sounds awesome.
- A generous 24 voices.
- Old and new-school wavetable technologies.
- Four-part multitimbral with effects for each part.
- Four stereo outputs.
- Superb build quality.

CONS

- If you plan to add your own wavetables, the 16 empty slots may fill up quickly.

SUMMARY

Groove Synthesis bring some Californian surf mojo to the synth world with the PPG Wave-inspired 3rd Wave. It pays excellent homage to that 8-bit wavetable sound whilst elevating it with 24 voices, enhanced oscillators, alternative modern 96kHz wavetables, extra filters, effects, and a sequencer.

Give Us A Wave

The PPG Wave was the first commercially successful wavetable synthesizer, released in the early 1980s. It is revered for its harmonically rich, crunchy, 8-bit sound, and it made a refreshing change from traditional analogue synths. In 2023, if you want to pick one up on the second-hand market, it will be costly. As I type this review, there is one on eBay for £15,000. And, in their old age, they are notoriously unreliable.

The PPG Wave 2.3 (its last incarnation before PPG closed shop and Wolfgang Palm, the designer, went to work with Waldorf) featured eight-voice polyphony, two oscillators per voice, 24dB/octave low-pass SSM filters, 30 wavetables, and even eight-part multitimbrality. The 3rd Wave ups the game considerably with 24 voices. It retains the analogue low-pass filters (using the SSI2140, a modern SSM2040 replacement) but adds a digital state-variable filter. On top of that, there's four-part multitimbrality, effects (two per part), 32 8-bit wavetables, 64 96kHz user wavetable locations, four stereo outputs, oscillator sync, linear FM, an arpeggiator, a sequencer and more modulation than you can shake a mod wheel at.

The striking blue design is, of course, an homage to the PPG, but so are the knobs and the lozenge-shaped buttons that cover the front panel. The build quality is superb throughout, with a five-octave, synth-action keyboard that feels lovely to play. It supports aftertouch, but not poly-pressure, which is a shame, although you can hook up an external poly-pressure keyboard and take full advantage that way. Standard pitch and modulation wheels are joined by over 70 knobs (a mixture of potentiometers and encoders), nearly 40 buttons, and a lovely bright monochrome display. Overall, I cannot overstate the premium feel.

The basic operation of the synth is done using the many dedicated controls. For any parameters that don't have a button or a knob, you use the screen plus four soft encoders and four buttons that appear directly above and below. Bob Coover's time at Sequential is evident in the design, so if you've used a Sequential synth manufactured in the last 10 years, you'll feel right at home.

Wave Mechanics

A 3rd Wave preset is made with up to four parts. Each part is a sound comprising three oscillators, filter settings, modulation, effects, etc. Parts have volume and panning and can be layered or key-split. They can even be assigned separate MIDI channels if you want to sequence them individually. With 24 voices, you can sequence four completely different parts and still have six voices per part (voice allocation is dynamic depending on how many parts are active). And, with four sets of stereo outputs, each part can be processed differently using your outboard equipment, in addition to two onboard effects. It's like having four PPG Waves in one box.

A part is made up of three oscillators. Each one can load a different wavetable, unlike the PPG, which had two oscillators which shared a common wavetable. Wavetables can be one of two types — either classic legacy wavetables or one of 64 'User' wavetables. The 33 legacy wavetables are modelled on the original 8-bit tables from the PPG Wave 2.2 and 2.3. In contrast, the User wavetables are all-new, high-resolution, anti-aliased, and sampled at 96kHz. There are 64 User wavetables, 48 of which come pre-filled, and 16 are empty for your own creations, although all 64 slots can be overwritten if desired.

In addition to the two wavetable types, there are seven virtual-analogue waveforms. They include sine, triangle, square, sawtooth, supersaw and two variable noise options. Each oscillator can frequency modulate or hard sync its neighbouring oscillator, although, for some reason, sync doesn't work when using User wavetables. Most of the waveforms can do some waveshaping via the Pulse Width control. These waveforms sound full and convincingly analogue. Even without wavetables, the 3rd Wave makes a powerful polysynth. I spent a fun afternoon making sounds using only these waveforms, and the results were »

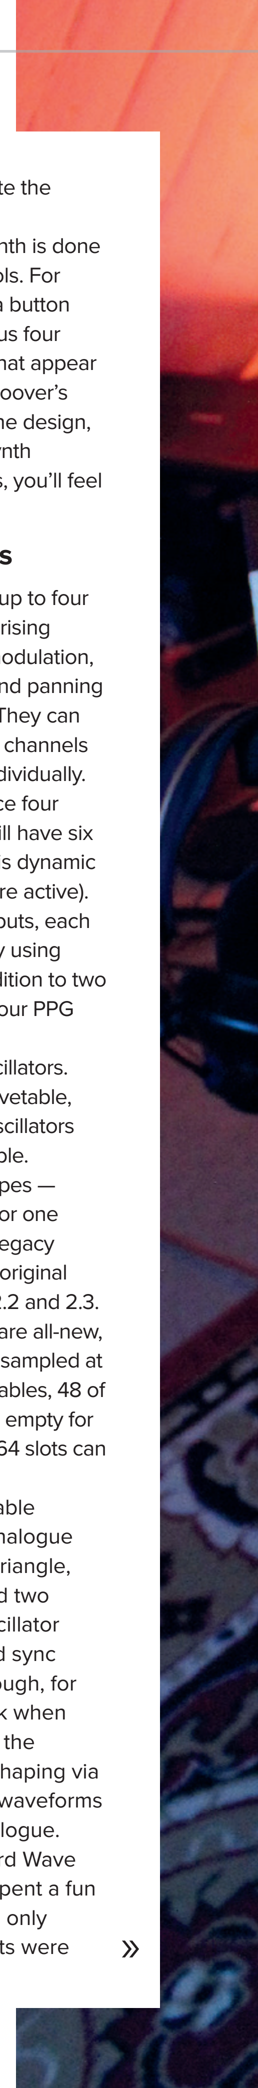


3RD WAVE
ADVANCED WAVETABLE SYNTH
ARPEGGIATOR
TEMPO
MODE
RANGE
MULTI-PART
EFFECT 1
ON
CLOCK SYNC
TYPE
MIX
PARAM 1
PARAM 2
EFFECT 2
PAN SPREAD
MONO
STEREO
ATTACK
DECAY
SUSTAIN
RELEASE
DESTINATION
AMOUNT
VELOCITY

Rejoice, as the power supply is built in and requires nothing but an IEC cable. The USB Type-B socket handles USB MIDI and access to the internal flash memory for wavetable audio transfer, OS updates and patch backup. MIDI in, out and thru are on 5-pin DIN plugs. There are one sustain and two expression pedal inputs, an audio input (unbalanced) for recording wavetables, four pairs of stereo outputs (unbalanced) and a headphone socket, all on quarter-inch jacks.

as good as from any of my dedicated analogue polysynths.

A standard option on modern wavetable synthesizers is waveform interpolation, which smooths the transition from one wave to another. The PPG had no interpolation, but in the 3rd Wave, it is called 'Wave Flow' and has its own dedicated front-panel button. It works on all three oscillators simultaneously. Wave Flow is the key to smooth evolving sounds, which, as it turns out, is something the 3rd Wave is phenomenally good at.

If accurate PPG emulation is your aim, you will have to disable Wave Flow, but that's not the end. There was a PPG feature called the 'upper wavetable'. The last four waveforms in the original PPG wavetables were always triangle, pulse, square and sawtooth. The 3rd Wave hides these last four waveforms by default since they can sound jarring compared to the first 60 waveform positions. When you turn on the upper wavetable, these last four waveforms are exposed, as well as wavetable 30, which has another 64 waveform positions that can be used if you modulate past wavetable position 64. This was how the PPG dealt with modulating beyond the end of a table. So in the quest for perfection, Groove Synthesis have recreated this idiosyncratic feature.

Another option in pursuing an accurate PPG sound is 'Waveform Smoothing', not to be confused with the wavetable smoothing, or Wave Flow mentioned above. The PPG had minor pitch errors from note to note on the keyboard caused by its lack of interpolation. Enabling Waveform Smoothing will correct these pitches, but anyone looking for the exact sound, warts and all, will want to keep this option disabled. Finally, there is a function called voice drift. Despite being digital, the PPG suffered from oscillator pitch and other parameters drifting over time, and this is now a parameter you can apply per patch.

It should be evident by now that Groove Synthesis have gone to great lengths to get the PPG sound perfect. Multiple elements must be in place if you want an authentic PPG sound: use an 8-bit wavetable, disable Wave Flow, disable Waveform Smoothing, switch the envelopes to PPG mode, increase voice drift, and enable the upper wavetable. You'll also need to load the same wavetable into two oscillators and disable the third oscillator. Thankfully, plenty of presets in the 500-strong factory library recreate famous PPG sounds. And they sound impressively close to the original.

The fact that you can switch these elements of PPG emulation on or off makes everything more flexible. Do you want a PPG sound with exponential envelope shapes instead of more linear PPG ones? No problem. A PPG sound but with all pitch instabilities fixed? No problem. Or even a modern 96kHz wavetable with all the pitch instabilities of a 40-year-old instrument. I love this approach to emulating old gear because it allows you to cherry-pick the bits you like best.

Modulate Me

Of course, wavetables are nothing without modulation, and the 3rd Wave certainly delivers. The Wave Envelope is a six-stage loopable envelope generator primarily designed to sweep through the waveforms in a wavetable. There are three per part (one for each oscillator), but you can assign them to other destinations in the 16-slot modulation matrix.

Another source dedicated to traversing wavetables is the Wave Surfer, a solid Californian name if ever I heard one. Wave Surfer is the largest encoder on the front panel. It allows you to sweep through the wavetables of all three oscillators at once. The effect is like a tone control for the preset, simultaneously changing the starting point of all three wavetables. It's a brilliant idea, and I found myself constantly using it to explore alternative timbres of any patch I had loaded.

You can assign plenty of traditional modulation sources to wavetable scanning or anything else via the modulation matrix. These include four DADSR envelopes, switchable between exponential and PPG (more linear) modes. There are also four LFOs with 10 waveforms, delay, reset and optional tempo sync. With four-part layering, that's a potential of 16 LFOs, and 19 envelopes per patch. And there are the usual MIDI sources, velocity, pressure and expression pedals. Mind-blowingly complex patches, anyone?

Two interesting modulation sources are barely mentioned in the manual, 'Audio In' and 'Audio Out'. Audio In allows you to use the audio from the input jack on the rear of the unit. Audio Out, similarly, is the final mixed audio generated by the synthesis engine. You can even use the oscillator signals as a source. Interestingly, all these audio-rate modulation options are really happening at proper audio rates. Many synths downsample their modulation, running it at a slow sample rate in order to save CPU, but the 3rd Wave runs its matrix at audio rates.

I also like the ability to modulate all the Wave Envelope segments' positions and times. Imagine a complex envelope with the height of each point moving with the help of an LFO or two. There are enough modulation options in this synth to keep the most adventurous sound designer busy for a very long time.

Making Waves

Creating your own wavetables has traditionally been a mind-numbing affair capable of reducing any grown adult to tears. To save your mind and your tears, Groove Synthesis have included the Wave Maker, which analyses audio and creates a wavetable. Instead of being a separate piece of software, Wave Maker is built directly into the synth. You can record audio from the rear-panel audio input jack or copy a WAV file to the internal flash memory via USB. Wave Maker will analyse the file and attempt to extract 64 waveforms.

Getting a good wavetable from a piece of audio can be hit-and-miss. Groove Synthesis recommend tuning audio to 93.75Hz (MIDI note F#1). This is equivalent to a single-cycle wavelength near 1024 samples, the period length that the 3rd Wave uses. If your audio isn't tuned to this frequency, you can have Wave Maker try to pitch-shift it first. Then you adjust the sensitivity level, which will alter the level of timbral changes that Wave Maker looks for in the file. The push of a button is all that's left to create your custom shiny new wavetable.

Although Wave Maker is fun to play with, you may wish to import wavetables you've downloaded from the Internet or made yourself using software like Xfer Records Serum or Kilohearts Phase Plant. The 3rd Wave requires a WAV file at 96kHz, 1024 sample wavelength, with 64 waves per wavetable. Serum wavetables use a 2048-sample waveform, but I converted some that I have created over the years, and I can confirm they sound fantastic played through the 3rd Wave. Groove Synthesis tell me they are working on an update to Wave Maker that will allow direct import of Serum format wavetables without the need to convert. That should be a real timesaver.

My only complaint is that there are only 16 spare wavetable slots. With so many free wavetables available online, it would be easy to add hundreds. You can overwrite the 48 factory-supplied User wavetables, allowing for 64 total writeable locations, but that would change the character of many factory presets, so it's not an ideal solution. One potential workaround is to export the user wavetables and import a different set, which is easy to do via the USB flash drive, but I still wish they'd included more slots.

The Other Stuff

It is proper to get carried away talking about wavetables, but the 3rd Wave has plenty more to offer. The 24 analogue filters (one for each voice) sound superb. The filter saturation sounds excellent, and there is switchable resonance compensation and dedicated envelope and velocity front-panel controls. The analogue filters self-resonate too, which is a nice bonus.

A digital state-variable filter opens the sonic palette further. Mode is variable between low-pass, notch/band-pass and high-pass. There is no way to change the

The 3rd Wave is a substantial instrument, measuring 976 x 362 x 136mm, and makes a strong case for blue becoming the official colour of wavetable synthesis.

routing, but I didn't find that a problem. The digital state-variable filter is always first, followed by the analogue low-pass. Unlike the analogue filter, the digital one does not self-resonate.

If, like me, you were raised on '90s dance music, you'll appreciate the Unison mode, which also offers a Chord Memory function for all those Detroit stabs and rave pads. Unison can be enabled per part, allowing you to detune up to 24 voices. I was almost scared to try the full 24, but, in reality, you don't hear much difference after about six voices. Using Chord Memory is simple: hold a chord and press the Unison button. Then any key you press will play the chord transposed. The Chord Memory is saved with the patch, which is a nice touch.

Each part can have two active effects. There are three types of delay (BBD, Tape and Stereo), three reverbs (Room, Hall and Superplate), chorus, phaser, flanger, ring modulator, Leslie speaker and distortion. With four parts and two effects per part, there is potential for eight active effects in any patch. Some effects use more processing power and, as a result, cannot be used in specific configurations. In practice, I found this wasn't an issue. Only the most insane patch designer would want to run eight hall reverbs simultaneously.

There are only two parameters for each effect on the front panel, but some effects have more options once you open their on-screen editors. An optional limiter can also protect the signal from clipping going into the effects. I didn't notice any of the internal effects clipping during my testing time, so I didn't feel the need to use it. But it's a nice option to have.

The quality of the effects is good. They don't match the main synth engine's quality and obsessiveness, so most patches sounded better when treated by external, high-quality effects. Of course, that is true of almost every synth with onboard effects, but it seems more noticeable here because the dry synth sounds so damn good.

Like the PPG Wave, the 3rd Wave comes with an arpeggiator and sequencer. The arpeggiator is fairly standard with its modes (up, down, random, etc.) and octave range. It can be sync'ed to tempo, and there's a hold function if you like to have both hands free for expert knob twiddling. It gets interesting when layering parts. Each part has its own arpeggiator, which means you can layer up to four different sounds, each with different settings, panning, effects... The results can send notes cascading into the darkness for days.

The sequencer is perfect for those who like to get ideas down fast. A sequence comprises 24 patterns with up to 32 measures (or bars) per pattern. You can even record parameter automation. Patterns can then be chained in a playlist, or 'Song', with repeats.

The sequencer works like a MIDI looper. You set the length of a pattern, then record and overdub (with the aid of a metronome if required) to your heart's content. You can record parts separately in different 'tracks' within the pattern. Any knob twiddling will be recorded into the pattern as well. If you make a mistake, you have to record again, although an erase button will work whilst you hold the button, so if you time it well, you can rub out the occasional duff note.

One nice touch is that the quantise function works 'live' on playback. You can record unquantised and then experiment with quantise values afterwards to see how they sound. If you don't like them, spin the knob back to 'off'. Once you have a good pattern, you can duplicate it to overdub additional notes or automation and chain the results into a song. I like the 3rd Wave sequencer. It's unfussy and very usable, and you'll never lose a killer chord progression waiting for the DAW to boot up.

Conclusion

I'll cut to the chase. The 3rd Wave might be my favourite synth released in the last 10 years. Allow me to explain why.

There are four main areas where a synth can fly or fall. The first, and arguably most important, is the sound. In this regard, the 3rd Wave does not disappoint. The sound is classy, expansive and detailed. At the risk

ALTERNATIVES

The closest synth to the 3rd Wave currently in production is the **Waldorf M**, which I reviewed in April 2022 (www.soundonsound.com/reviews/waldorf-m). The M is a recreation of the Waldorf Microwave, the synth Wolfgang Palm worked on after the PPG Wave when PPG closed their doors. Like the 3rd Wave, the M offers 8-bit wavetables alongside analogue filters. It is only eight voices, but it's a desktop unit, which might suit anyone who doesn't need another keyboard.

Otherwise, the **UDO Super-6** might be worth a look (www.soundonsound.com/reviews/udo-audio-super-6). It is modelled very loosely on the Jupiter-6. Although it doesn't use traditional wavetables, it does offer digital oscillators with many waveforms, 12 voices, analogue filters, and a similar only-the-best-will-do approach to the design and manufacture.

of sounding like a hi-fi salesperson, the highs are intricate with no harshness, and the low end is well-balanced and capable of shaking the windows if required. The PPG emulation side is impressive too. Groove Synthesis' dedication to nailing the exact sound is remarkable, right down to the tuning peculiarities and low-end imaging problems found in the original.

The PPG obsession is only half of the story. Using the high-resolution 96kHz wavetables, triple oscillators, extra filters, FM, hard sync, modulation capabilities and four-part layering, you can go way beyond anything the PPG could do. Whether complex, ever-evolving textures, expansive pads, classy polysynths, layered arpeggios, simple analogue recreations, digital basses, FM pianos or CS-80 brass emulations — the 3rd Wave does them all and remains elegant and convincing throughout. In the sound department, the 3rd Wave is a resounding success.

The second area that can let a synthesizer down is the user interface. Once again, we're in good hands. The 3rd Wave is so easy to use that you barely have to read the manual. There are controls for nearly everything, and no dual functions, shift-clicking or menu-diving (well, maybe a bit, but only for rarely used settings). The 3rd Wave is a proper synth with proper controls.

The third potential downfall for any synth is power, or lack of it. You'll know the struggle if you have used a four- or six-voice synthesizer. It doesn't matter how great the synth sounds; if it suffers from voice-stealing, it can certainly spoil the mood. With a massive 24 voices, four parts, a stereo output for each part, and dedicated effects per part, the 3rd Wave can effectively work as four 6-voice synthesizers, two 12-voice synths, or a single 24-voice monster. Groove Synthesis could have stopped at eight or 12 voices, but they didn't. Once again, the 3rd Wave knocks it out of the park.

The fourth and final area where a synth can fail is build quality, and, once again, the 3rd Wave lands in the top tier. It looks fantastic, the materials are premium, the aluminium casing curves in all the right places, and everything feels rugged and well-built.

The 3rd Wave scores top marks on every single aspect by which I can judge it. The only elephant in the room is the high price, but considering the no-compromise approach that Groove Synthesis have taken, I can't fault it for being expensive. And it's still considerably cheaper than eBay's £15,000 PPG Wave. I know it's only February, but my money is on the 3rd Wave winning best synth of 2023.

$ $4995
W www.groovesynthesis.com

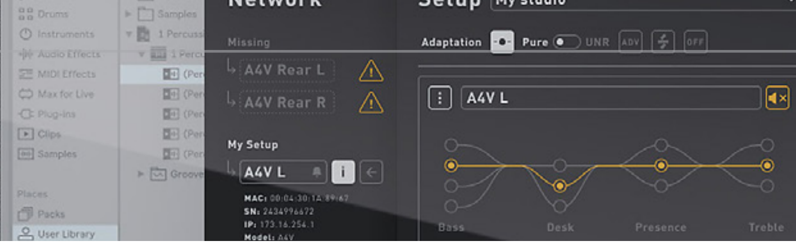

ADAM Audio A7V

Active Monitors

With newly designed drivers and Sonarworks integration, ADAM's new A series promises a lot of monitoring for your money.

BOB THOMAS

It doesn't feel like 10 years since I reviewed (in *SOS* June 2012), and subsequently bought, ADAM Audio's A77X, a three-way version of the company's two-way A7X nearfield monitor. Much has happened with ADAM Audio in that time, including their purchase by UK-based Focusrite PLC in July 2019. Focusrite CEO Tim Carroll said at the time: "With so much expertise between us in acoustics, sound reproduction, DSP, Audio-over-IP, and control, the opportunities are abundant." With the launch of the new ADAM Audio A series, which includes the A7V reviewed here, it would appear that at least two of those opportunities have now been very definitely seized.

ART Of Noise

The A7V's vertically oriented black cabinet features ADAM Audio's signature thick, deep-bevelled baffle, which strengthens and stiffens the cabinet, helps minimise cabinet diffractions, and provides a stable mount for the monitor's high- and low-frequency drivers. The baffle itself has an attractive aesthetic, with the front-facing bass reflex port's large, flared twin apertures sitting beneath and on either side of the 7-inch bass driver. It almost creates the appearance of a pair of cartoon eye sockets! The port and the internal profile of the flares are designed to optimise airflow and thereby to minimise port noise which, in practice, I found that they did very successfully.

The A7V features a new bass driver, the cone of which is made of MLM (Multi-Layer Mineral), a unique blend of mineral fibres baked together to create a lightweight, highly stable composite material. In combination with a new magnet assembly, this new material not only enables the A7V bass driver to deliver high output levels with low distortion, but also makes for a great-looking cone.

The A7V crosses over at 2.8kHz, at which point ADAM Audio's signature X-ART (Extended Accelerated Ribbon Technology) high-frequency driver takes over. Hand-built in the company's Berlin factory, the X-ART driver's pleated ribbon construction delivers precise transient response and highly detailed resolution up to 50kHz. In the A7V, a new 120x70° HPS (High-frequency Propagation System) resin and glass-fibre waveguide has been designed specifically to match the X-ART's dispersion angle in the crossover region to its accompanying bass driver. In the horizontal plane, this matching, coupled with the waveguide's wide dispersion angle, results in a smoother off-axis response and creates a larger usable 'sweet spot'. In the vertical, the HPS waveguide's 70° dispersion is designed to help reduce unwanted desk and console reflections.

Another feature of the HPS waveguide is that it can be rotated through 180°, which allows you to position the A7V horizontally or even upside-down, thanks to the M8 inserts on the underside of its cabinet and using a series of mounts that will soon be released by ADAM Audio. Personally, I prefer to position two-way loudspeakers horizontally whenever possible, as I feel that this improves the subjective integration of the high- and low-frequency drivers but, as the saying goes, your mileage may vary.

Voice Activated

In the A7V, amplification is supplied in a hybrid fashion, with a Class-D, PWM amplifier providing 90W RMS to the MLM bass driver, and a 15W RMS Class-A/B amplifier handling the X-ART high-frequency driver. Although these RMS figures may seem low to some, these amplifiers endow the A7V with a frequency response of 44Hz-41kHz (-3dB) at up to 108dBSPL (RMS) at 1m, which is going to be plenty loud enough for anyone using the A7V in the nearfield.

I work in a compact space in which my monitors are backed up quite close to a wall, so I was a little sorry to

ADAM Audio A7V
$1599

PROS
- Delivers superb sonic performance at an attractive price point.
- A Control remote-control application provides an easy to use, intuitive user interface.
- Sonarworks SoundID room calibration profiles can run on its embedded DSP platform.
- Easy to work on for long periods.

CONS
- None

SUMMARY
A superb nearfield monitor that delivers an assured and superbly detailed audio performance, and which can run Sonarworks SoundID profiles on its embedded DSP.

see that the A7V's on/off switch and input sensitivity controls have been relegated to the cabinet rear panel. This panel also houses the monitor's Room Adaptation and Voicing settings, and its balanced XLR and unbalanced RCA phono analogue inputs.

Although there are physical switches present for room compensation and voicing, in a first for ADAM, these two functions — plus EQ and crossover duties — are handled by a DSP system that has been designed not only to provide greater tuning precision and predictability than can be found in all-analogue monitors, but also to offer the possibility of future enhancements and upgrades.

Room Adaptation is controlled from the rear panel by four push buttons that cycle through the available cuts and boosts at frequencies that ADAM Audio label as Bass, Desk, Presence and Treble.

The Voicing button cycles through the three available options: the flat-response Pure; UNR (Universal Natural Response), which, according to ADAM, is a dynamic, natural-sounding response curve based on a variety of previous ADAM loudspeakers including the A7X; and Ext, which selects a Sonarworks room calibration profile that has previously been loaded into the onboard DSP from the A Control software application.

The A Team

An Ethernet port on the A7V's rear panel allows you to connect a computer running the A Control application to it, giving you real-time control not only of the Room Adaptation and Voicing functions outlined above, but also access to a six-band fully parametric equaliser in the ADV (advanced) mode. If you are running a stereo or multi-channel setup you'll need to connect your A7V via a router so that your computer can address all your monitors. A Control runs only on computers running Windows 10 and above and Mac OS 10.15 and above.

Installing the A Control software went very smoothly on my iMac running 10.15, and the two connected A7Vs were quickly recognised, with full control being instantly established. Once you've checked that the virtual L-R positioning matches the real world, operating A Control is a simple and intuitive process.

Having remote control of each monitor's rear-panel functions and internal six-band parametric equalisers is a real convenience in itself, but if you're already using, or plan to use, Sonarworks SoundID calibration software (registration of your A7V gives you access to a 60-day fully functional trial licence), A Control's ability to load a SoundID-generated calibration profile into the A7V's DSP platform, and that platform's ability to run that profile independently of the computer, will no doubt prove extremely attractive to you. If you've no idea what I'm talking about, read John Walden's review of SoundID in *SOS* July 2021: www.soundonsound.com/reviews/sonarworks-soundid-reference.

In Use

I spent a good few hours playing music through the A7V and experimenting with the rear-panel settings so that I could get used to their performance. In the Pure mode, my initial impression was of a monitor with an effortlessly clear and detailed treble, a midrange that displayed real clarity, detail and presence and a fast, well-controlled and detailed bottom end with great transient response. To me, the A7V's ability to resolve fine details is one of its great strengths, as is its ability to create a wide stereo soundfield in which instruments and vocals are precisely delineated and positioned.

For a great many users, the room adaptation switching will be sufficient to match the frequency response of the A7V to their musical environment. With its intentionally limited range of EQ adjustment I found this facility to be quick and intuitive to set up, again using the Pure voicing as my starting point. Obviously, if you don't have a computer that is capable of running the A Control application, this will be the only voicing available to you.

Listening to two of the most detailed recordings that I know of, and know well, (*Spes* by the Finnish choir Cantus with Frode Fjellheim, and *Via Crucis* by L'Arpeggiata), I heard small details with great clarity and definition: the dying reverberations in the ceiling of the church where *Spes* was recorded, and the tiny details of the renaissance instrumentation of L'Arpeggiata.

Going to the lower end of the scale, Deadmau5's Grammy Award-winning *4x4=12* features a synthesized low end that not only takes some handling, but is also extremely detailed and which poses a real challenge to any loudspeaker, small or large. Admittedly, the A7V is never going to be able to move enough air to equal the low end of a much larger monitor, but given the way that it handled and resolved the detail in the opening track of the Deadmau5 album, I'd be pretty confident that making low-end mixing decisions on the A7V would work out well. However, if I was working regularly on tracks with a lot going on in the extreme low end, I'd be inclined to add ADAM Audio's Sub 8 or Sub 10 into the mix.

I really liked the overall sense of detail and clarity in this voicing and, as with my A77X, the X-Art driver's effortless delivery and lack of harshness doesn't tire you out and makes working with the A7V for extended periods a pleasure rather than a chore.

Next, I tried the UNR mode. With its pushed low end and high end giving a slightly »

The rear panel houses controls for accessing the Room Adaptation and Voicing presets, as well as an Ethernet port for connecting to ADAM's A Control software.

scooped feel across the midrange, this voicing has a lot going for it. Personally speaking, I wouldn't use it for tracking, mixing or mastering when I was being paid to do a job, but I felt that it had a lot going for it when I was messing around with my synths, doing overdubs and listening to streaming services for pleasure when doing studio maintenance and the like.

The ADV voicing's powerful six-band, fully parametric equaliser is accessible only from the A Control application. The primary use of this EQ will either be to tailor the response of the A7V to your own personal taste, or to try to persuade it to sound like another monitor of your acquaintance. If you want to try to optimise the A7V for your room, I feel that using it successfully would require either a lot of patience and care or the use a standalone frequency response analysis program, which seems a bit of a hard way of achieving a result that Sonarworks SoundID will turn out with very little effort on your part.

This final voicing slot, EXT, is where you can store and access the calibration profile that you have generated using Sonarworks SoundID. All you have to do is to download and install the included trial version of the program and, with a Sonarworks-calibrated measurement microphone (or any other flat-response capacitor microphone) and an audio interface in hand, in 20 minutes and with very little effort you'll have produced the calibration profile for your A7Vs in your room and exported it in '.adam' format. Loading that profile into A Control and applying it to your A7Vs takes seconds, and ADAM have provided a tutorial on how to do all this at https://bit.ly/400W4m3.

With my self-generated calibration profile loaded into the review A7Vs, switching between it and the Pure setting was interesting and revealing. With the profile active, the A7Vs felt slightly more present, somewhat tauter in the high frequencies and more even in the bass. It was a subtle difference that's difficult to describe in words, but I did like it and I'd definitely use that profile when checking my own and other mixes. I think that for tracking, overdubbing and mixing, I'd be tempted to stick with the Pure and UNR voicings, as I do like the sound of my room and I know how to get the best out of it. However, if I didn't particularly like the sound of a room or I was taking the A7Vs out on location, I'd definitely use SoundID (as I do already) to optimise the A7V for such spaces.

ALTERNATIVES

The ADAM Audio A7V is parachuting into a hotly contested area of the nearfield monitor market, where you'll find competitive offerings from manufacturers such as **Eve**, **Dynaudio**, **Focal**, **Genelec** and **Neumann**. The A7V's remote control and Sonarworks integration are big pluses in this context, though some of these rivals offer proprietary alternatives such as Genelec's GLM system.

Final Thoughts

The new A7V is, in my opinion, a significant step forward in the evolution of ADAM Audio's nearfield monitoring. The X-ART treble/midrange driver and MLM bass driver combine to produce a highly detailed presentation that I found very attractive to listen to and easy to work with. The effective onboard room adaptation switching, and the remote control facilities and Sonarworks SoundID integration offered by the A Control software, provide an intuitive user interface that allows for on-the-fly switching between ADAM's voicings and a Sonarworks room-correction profile.

Both technically, sonically and operationally, the A7V is one of the best nearfield monitors that I have heard at its price. If you are looking for an affordable, high-quality two-way active monitor speaker, you really should have a listen to the A7V.

$ $1598.98 per pair.
W www.adam-audio.com

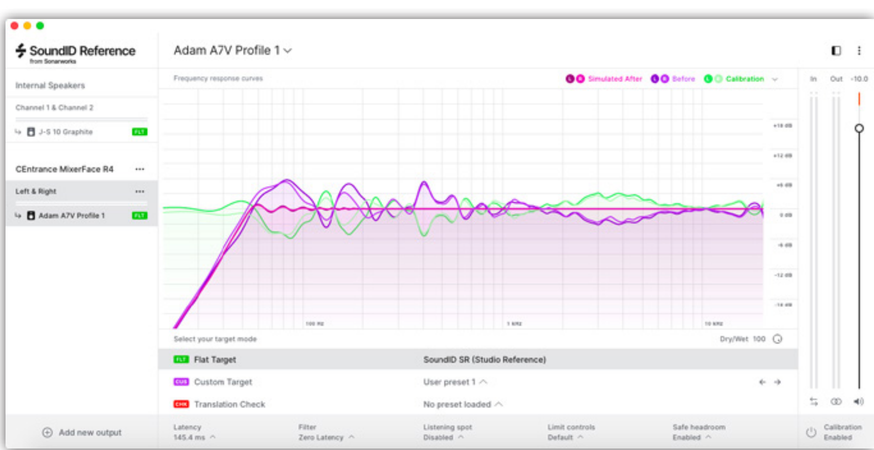

Sonarworks SoundID correction curves can be uploaded directly to the A7V.

LEWITT

WINNER
SOS AWARDS 2023

2023 Sound On Sound
winner for best microphone

Sometimes, patience pays off.

We spent 10 years in research and development to create the LCT 1040 microphone system, with feedback from our fellow enthusiasts all over the globe. We're so happy with the results, and glad to see you are too.
Thank you for voting!

- 4 all-analog tube characteristics - Clear, Warm, Dark, and Saturated
- Developed in close cooperation with Grammy-winning studio professionals
- Detachable remote control for unrivaled workflow
- Blendable FET circuit and seamless polar pattern control

LCT 1040 Ultimate microphone system.

lewitt.link/lct1040-sos

NEIL ROGERS

Terry Audio CEQ

Six-band Stereo Equaliser

Some love a Pultec. Others prefer a Neumann or a Lang. Marshall Terry decided to combine his favourite aspects of these much-loved equalisers in one box!

One thing you'll notice over time if you get to try out lots high-end outboard gear (which, happily, as a writer for *SOS*, I do!) is that once you discount the logos, knob choices, colour schemes and so on, lots of products are actually pretty similar from an aesthetic point of view. I don't mean that to sound jaded or critical in any way, and in many ways it's a positive thing: it means these tools tend to feel familiar, and that enables us to manipulate audio quickly and intuitively. But as you can see from the main photo, the product I'm evaluating here bucks that trend in quite spectacular fashion! Terry Audio's CEQ is, without a doubt, the boldest, most original-looking piece of outboard I've had the pleasure of reviewing to date.

Overview

Spread over 5U of rack space, the CEQ is a six-band stereo EQ with a combination of active and passive bands (it's mostly passive, as I'll explain). Its front panel displays no obvious settings or values, with just a handful of hieroglyph-esque markings giving the most basic indication of what control is doing what, and some coloured dots and lines hinting at a scale around the knobs. Boutique, artisan, call it what you like: it takes a very confident audio equipment designer to develop and release a high-end product that looks like this!

That designer is Marshall Terry, who some readers may know best as the chief technician at Shadow Hills Industries. There, he's played a major role in developing and building the various products for which that well-respected brand is known, and through both that role and his own endeavours, Terry has become involved in a number of bespoke builds for high-end clients. That perhaps gives a little insight into how Terry Audio and this distinctive EQ came into being.

Developed over five years, the big idea behind the CEQ was to combine the best elements of Terry's favourite vintage inductor-based EQs, from the likes of Pultec, Lang and Neumann, into a single hand-wired device. A modern unit that collected together the 'greatest hits' of all these designs would offer an engineer access to their magic without any of the burden of maintaining old, rare equipment.

While the CEQ may not come 'cheap', it's certainly much less expensive than acquiring (and maintaining) several

expensive vintage EQs and cascading them to achieve a similar thing. But Terry also makes the point that even those who can afford to do that can quickly run into problems. For example, using several such units inherently means the presence of multiple line amplifier and transformer stages in the signal path, and that can easily lead to excessive coloration. Balancing all these considerations to create a device that mastering engineers are happy to use is a technically challenging and expensive endeavour, and I think that puts the price into perspective somewhat.

Listen To The Bands

When you mention the words 'vintage passive EQ', most people's thoughts will turn first to the classic Pultec equalisers, and at least two of the CEQ's six sections look to these devices for their inspiration.

Starting with the bottom end, the CEQ provides 'cut' and 'boost' resonant shelves which are designed to offer the low-end magic of a Pultec EQP-1A, but with some additional flexibility. The bass boost section has frequency options that range from 40 to 200 Hz, and when used in combination with the bass cut, provide access to more options for shaping the bottom end of a mix than you might expect with this style of EQ. This section was comfortably my favourite during the review period; thanks to the ability of the low-cut control to carve out 'low end' up as far as 480Hz, I found dialling in pretty generous amounts of extra weight across a whole mix perfectly achievable. It's tricky to convey this sort of thing in words but, as with a Pultec, most of the magic can be found in the way these two bass bands interact; if you experiment with boosts and cuts in tandem, it's possible to bring the bottom end of a mix forward in a surprisingly focused way.

The mid boost section is based on the Neumann PEV equaliser, which was found on mastering transfer consoles in the 1960s and is celebrated for its broad and gentle bell-style EQ curves. On the CEQ, this band covers an extremely wide frequency range, from 310Hz right up to 8.2kHz. I work primarily as a tracking and mixing engineer (I'll come back to this point later) and upon first using this band I was struck immediately by just how liberal I could be with the knobs — to the extent that I often found it helpful to exaggerate an EQ boost to get a better sense of how it was affecting things.

Because of this and the large frequency range, it took me a while to learn and, crucially, to appreciate how this band interacts with the high-frequency shelf. But once I had become a bit more familiar with it, I found I was able to get just the results I was looking for. More often than not, that meant increasing the sense of presence, but without adding harshness around the cymbals or the 6-8 kHz area in vocals.

As I mentioned earlier, the CEQ can *almost* be classed as a passive EQ: five of the six bands are passive. The lone active band is the mid cut section, and Terry describes this as "the most transformative and creative control". The Q value for this midrange circuit is described as 'dynamic', which is a shorthand way of saying that its bell-shaped curve becomes narrower the more you cut. Helpfully, this means you can use this section to achieve very different things. You can, for instance, pull out a little general harshness or low-mid build-up with a pretty broad curve. Alternatively, you can home in on a specific problem pretty precisely. As with the midrange boost, I initially struggled to hear and appreciate the subtle changes this section can often deliver, and I found myself wanting to exaggerate my EQ moves just to get a better sense of what I was doing. During the review period, I often got good results using this band to remove low mids in a way that helped me shape any low-end boosts I'd dialled in. But I could imagine this mid cut band would become more and more useful the more you learned and developed your understanding of how this EQ's stages interact.

The treble boost section is another one that's inspired by classic Pultec designs, and Marshall Terry was keen to mention his careful selection of MPP inductors: a choice guided by his desire that the CEQ should be able to add subtle coloration in a particularly smooth-sounding, pleasing way. You're given only one knob to control this section, but you can do an awful lot with it because over part of its range it selects bell EQ boosts, while for the remainder it engages a shelving filter.

It's always a pleasure hearing a high-quality passive EQ add a sort of 'neutral' sense of brightness and clarity to a mix, and the CEQ really didn't disappoint in this respect. I often liked the effect of adding a little high-shelf »

Terry Audio CEQ

$5950

PROS

- **Offers the sonics of some revered vintage EQs.**
- **Excellent-sounding Pultec-style low end.**
- **Surprisingly versatile midrange cut and boost options.**
- **High-end boost filter sounds just that: high-end.**
- **Distinctive looks.**
- **Encourages purely ear-based judgements.**
- **Switched and continuous control options.**

CONS

- **Takes time to properly learn and appreciate.**
- **Some will find the lack of front-panel markings frustrating.**
- **The extra circuit control features can be very subtle.**

SUMMARY

This distinctive-looking, high-end six-band inductor EQ takes inspiration from several classic equalisers. The result is instantly impressive, and the layout encourages a focus on purely listening-based decision making.

» boost, set quite far down the frequency spectrum (it ranges from 1 -18 kHz); this tended to elevate even dense mixes in a very broad, natural-sounding way. Used in combination with the treble shelving cut control, you have plenty of high-end 'sculpting' options. Sadly, I could only keep hold of the review unit for so long, and time pressure limited the nature of my experiments with this band to some extent. But I was left with a clear sense that there's much more to be discovered here, and curious to know how I'd end up using the HF bands after more sustained use.

The Elephant In The Room

The aesthetics of the CEQ will no doubt split the room, but every single client who visited the studio asked me what it was — and more than one remarked how 'cool' it looked. From a user point of view, of course, the lack of values and descriptions will be a talking point, and I'm genuinely torn on the pros and cons.

Marshall's reasoning is that he didn't want the controls to feel overly cluttered (I've reviewed a few devices like that!), and that the CEQ is intended as a tool whose operation should be based on 'feel' and 'musicality': something that encourages you to be guided more by your ears than your eyes or assumptions. That worked: without question, the CEQ led me to choose settings that I probably wouldn't have selected had I always been conscious of the values, and using it gave me some positive food for thought. Still, I did find it frustrating at times, particularly when I was in 'workmanlike' mode, wanting to achieve a certain sound quickly. So it's perhaps worth pointing out that the recall sheet shows the values, meaning you can easily find out what you're doing. Also, I did find that I grew accustomed to what was where pretty quickly; I imagine that, over a few months' use, this sense would develop to the point where you just know!

In Use

The CEQ may be billed as a mix-bus and mastering EQ, but I'm sure Marshall Terry wouldn't mind me describing it as being more akin to an instrument than a tool. In part, that's down to the lack of EQ values on the front — a deliberate attempt to encourage you to listen carefully — but it's also down to how the different controls interact, and I developed a real sense during the review period that the CEQ's full potential will only really be properly appreciated in a critical mastering environment, once you've spent sufficient time getting to know it.

One of the main roles in which I used the CEQ was as a stereo mix-bus EQ, and in that context, I tended to dial in broad changes that I liked, and stick with those same settings when working on a group of tracks together, to maintain sonic consistency throughout a project. It's an approach that works brilliantly, even if it doesn't fully exploit the flexibility and range of options described in the manual.

The CEQ can sound beautiful, and the unusual interface can be refreshing or even liberating. However, in real-world use I also found that its action could be very subtle — to the point where my tracking and mixing-focused ears sometimes struggled to discern quite what difference some of the controls were making when processing the more dense mixes I was working on. That applies especially to the Shift and Buffer settings (discussed in the box), though I could often hear a subtle low-end 'bump' when switching the output transformer in and out of the signal path. Generally, I found the CEQ to be most satisfying when working on sparser acoustic styles of music: tracks that had more space just made it easier to appraise and appreciate the enhancements.

The Joy Of Six?

This lovely, high-quality EQ makes as big a visual statement as it does a sonic one. It's also one that requires a significant investment, both financial and in terms of the time you'll need to spend learning how to get the best out of it. But it is intended to be a distinctive 'centrepiece' for serious mixing and mastering engineers and, as I mentioned above, you really need to judge its price in the context of what this hand-built product aims to replace and improve upon. This is high-level, almost obsessive audio electronics, and although the changes that some of the more esoteric features provide can be extremely subtle in practice, I can only admire the dedication that goes into developing and making equipment like this available to discerning engineers.

$ $5950
E marshall@terryaudio.com
W www.terryaudio.com

Additional Controls

I've covered the CEQ's EQ controls in the main text, but there's more to this device than EQ alone, and a few switches provide the engineer with control over the sonic character. First, the Shift control is used to change the capacitors used in the low-boost section of the EQ, the idea being that you can choose between a softer, more 'vintage' vibe and a more modern sound that's a bit more 'precise'. Next, we have the option to switch the custom-wound output transformer in and out of the signal path, again allowing you to choose between a cleaner sound and a more coloured one. Finally, we have the option to employ or change a buffer at the input stage — this is another pretty subtle component-level change, which can make the upper midrange a hint more/less present.

Despite the CEQ's size and heft, the rear panel is perhaps best described as 'minimal'!

Win! McDSP Everything Pack

Worth $2999

McDSP's user list reads like a *Who's Who* of world-class mix engineers: Dave Pensado, Tony Maserati, Mike Dean and Michael Brauer, to name just a few, all swear by McDSP plug-ins, while almost every one of our Inside Track interviewees has a McDSP processor somewhere in their signal chain.

Founded in 1998 by industry veteran Colin McDowell, McDSP started out making plug-ins for Pro Tools TDM systems, but have since branched out into native processors, and more recently released the groundbreaking APB, a rack unit that combines the sound of analogue outboard with the convenience of plug-ins.

For this month's *SOS* competition, we've teamed up with McDSP to offer you the chance to win the company's flagship software bundle, the Everything Pack Native. This comprehensive Mac OS and Windows suite comprises some 23 plug-ins covering pretty much every processing option you'll ever need, from their analogue-modelled Channel G and 6000-series channel strips to no fewer than nine dedicated dynamics processors. EQ options include the 4020 Retro EQ, the AE400 and AE600 Active EQs, the company's pioneering FilterBank plug-in, and the cutting-edge 6020 Ultimate EQ. Also included are the Revolver convolution reverb, EC-300 Echo Collection tape-delay emulation, noise reduction tools, the DE555 Advanced De-esser, the tape-emulating Analog Channel, the SA-2 Dialog Processor, and the lo-fi distortion Futzbox plug-in.

To be in with a chance of winning this fantastic collection, simply visit the URL shown, and answer the questions there, by Friday 7th April 2023. Good luck!

To enter, please visit:
https://sosm.ag/mcdsp-comp-0323

Prizes kindly donated by McDSP
W mcdsp.com

Polyend Play

Sample Sequencer

Polyend's Play takes a unique — and very hands-on — approach to sequencing samples.

SIMON SHERBOURNE

Polyend's Play shares the same form as the company's Tracker, but offers a very different workflow. By comparison, Play appears to be a more conventional sample-based groovebox, but its freeform 'pick and place' grid sequencing approach challenges traditional assumptions about tracks and mixer channels.

On paper, Play overlaps with peers like Elektron's Digitakt and Novation's Circuit Rhythm: eight voices of mono sample playback, MIDI sequencing lanes, pattern-based sequencing and effects. But Play's angle is a focus on fast creation, facilitated by a generous 8x20 pad grid, a set of generative and probability-based tools, and a palette of real-time effects and beat manglers. Despite the resemblance to the Synthstrom Deluge, Play is not a 'workstation' like the

Polyend Play
$799

PROS
- Fast, creative.
- Real-time Perform effects.
- Interactive, inspiring pattern generation and randomness.
- It's a MIDI sequencer too.

CONS
- No audio inputs.
- Accidental encoder touches jump modes.
- Limited sample manipulation and modulation.
- Some frustrations while learning its quirks.

SUMMARY
A refreshing take on the drum-machine/groovebox format with step-sequencing and randomness features that can break you out of predictable patterns.

Deluge, MPC or Maschine — it doesn't, for example, have internal synth engines or record or sample audio.

First Impressions

I had an instant crush on the Play when it was revealed. It's a gorgeous device, smaller and slimmer than it looks in photos. Power is via the USB-C connection so it's portable, although doesn't have the cable-fee convenience of an internal battery. There's a generous full-colour screen supported by chunky buttons and a dial, and 15 touch-sensitive encoders. The rest of the panel is a grid of tiny RGB pads. The pads have a low profile and slippy surface so you can swipe through rows or columns when entering triggers or sliding through performance effects.

Connectivity is minimal. The USB connection provides MIDI comms as well as power, then there's MIDI in and out via mini-jacks. A single stereo audio output means you're either running to headphones or direct outputs. A quarter-inch breakout adaptor is provided in the box, along with one MIDI jack-to-DIN adaptor. There's no audio over USB; when pairing the Play with things like an Elektron unit, OP-1 or laptop it would have been convenient to run with a single cable (even better if there were multitrack outputs).

Sample banks (and the system) reside on a removable micro-SD card. Play comes with a 32GB card containing factory and artist sound collections. Similar to the Circuits, the Play uses the concept of sample packs, which are loaded into working memory for use in your Project. These are easy to manage without any special software: a pack is simply a folder of folders containing samples. There's a suggested convention for organising and naming folders which dovetails with Play's Randomize and Fill features, but you're free to go your own way.

You can import individual samples from the SD card or load a whole pack by selecting a folder. There's a limit of 6 minutes of 44.1kHz mono audio, although my own pack of individual hits reached the 255 sample limit before running out of memory. Importing a sample pack or opening a Project containing one can take a minute or more, so probably not something you'd want to do mid-gig.

Pick & Place

Once you're up and running you're presented with the default view: the eight-lane step sequencer for creating sample-based patterns. Patterns start as 16 steps long and fit into the visible grid. They can be extended up to 64 steps, and each lane can have an independent length, playback speed and order (with various sequence scrambling modes). Tapping the Pattern buttons zooms out to show the 128 Pattern slots in the Project.

Years of drum-machine indoctrination suggest that the next step is to load sounds into tracks, or load a kit. Play doesn't work like this. Although the rows on the grid have track-like properties (length, speed, mute, etc.) the grid is a freeform canvas. Sample selection, sound, playback and mixer settings are all stored on an entirely per-step basis, with no default sound or effects settings owned by any particular track.

The basic workflow for building a pattern manually is dubbed 'pick and place' by Polyend. You pick a sample from the loaded pool, tweak sound settings, then place the sound onto steps in the grid. You could scatter a sound all over the grid; of course it still makes sense to follow some kind of track layout, but it's basically up to you.

The four right-hand pad columns can be used as a trigger grid or keyboard for quick note selection as you drop steps, and these can also be used for real-time recording of the 'working step'. However, none of the pads are velocity sensitive, and you can only trigger one sound at a time — Play is not for finger drummers. If you plug a MIDI keyboard in you can use it to play and record chromatically with velocity.

This way of working is fast and voice efficient, and can help you to avoid predictable routines. However, it did keep tripping me up in various ways until I got used to it. It's hard to get out of the habit of editing a pattern by just tapping a step in a track and expecting it to be the same »

Minimal in terms of design and connections, the Play's back panel features just an audio out and MIDI I/O on 3.5mm sockets, a micro-SD slot, a USB-C port and a power button.

» sound as its neighbours. If you want to, say, add some snare hits to your nominal snare track, you need to hold down one of the existing snare hits first to copy it. And It might not be clear what sounds are on each step, although you can preview steps by holding the encoder button. You have to train yourself to be aware of when there is or isn't a selection. I found the initial period while discovering and getting used to these things frustrating, but I did get past it.

Touch & Tweak

Sounds already in the grid can be tweaked by holding pads and adjusting encoders. Whole tracks can be selected via the last pad on each row, or you can hold the Shift button to make a ranged selection, which can be a sub-selection of a track, or any rectangular area across the grid. Selected ranges can be copied and pasted, nudged around with the Move control, or have any other settings adjusted. When parameters have different values within the range they show as a tilde, and any changes are applied as relative offsets.

The encoder system is another unique design feature that has a learning bump to get over. Each encoder controls two parameters. Touching an encoder promotes it to the screen, where you can toggle which of the two parameters is focused using the left-hand buttons. It's absolutely clear looking at the layout how to use it, but my brain kept messing it up in practice — regularly grabbing the wrong encoders. This improved with practice but the issue that persisted was that it's easy to accidentally switch modes by the slightest touch of any other encoder. I switched off the mode that lets you double-tap knobs to switch to the Shift layer after I kept doing it unintentionally.

There's a modest amount of sound tweaking possibilities. You can adjust level and panning, and apply a high- or low-pass filter, bit-crushing, overdrive and reverb/delay. Sample editing starts and ends with start and end times: there's no looping, slicing or time-stretching. The closest thing to an envelope is start/end fade: there's no other conventional modulation. All the focus is on sequencing and per-step parameter changes, assisted by features such as Chance and Randomize. Fair enough, these can imprint a lot of movement, but there's no way to write automation outside of sound steps ('trigless trigs' in Elektron speak) so no way to modulate notes as they play.

Fill & Randomize

Play's Fill button holds a wealth of clever ways to generate rhythms in selected areas of space. For example, you can select a whole track, then generate a random pattern with variable density using the current sound. Multiple taps of Fill each generate a new pattern. There's also a deterministic Euclidean pattern option if you prefer.

Three specific drum modes — Kicks, Snares and Hats — generate patterns that make sense for those sound types. Each grabs a random sound from the appropriately named folders within the loaded sound pack. If you select two or three lanes at once you can do a 'Beats' mode Fill, where the Play layers all three drum types. If you don't like the sound selection you can hit Fill again and it will pick anew.

Fill is only the first step in Play's generative journey. Randomize offers a range of ways to mess up stuff already in your Pattern. Like everything on the Play, this applies to the tracks or range that you select. Modes include Notes (with one- or two-octave range), Volume and Sample Length, which can add lots of interesting and natural variation. More fun are Sample In Folder or Sample In Pool, which dial in random sample flips per step. Then there's Texture and Space, which modulate settings like the filter and overdrive, or the delay and reverb. Finally Nuke and Duke Nuke (nice) vary a whole bunch of things at once. The key is that you control the amount of randomness, from subtle variation to full mayhem.

You can reset the randomisation to re-roll the dice as many times as you like, then if you hit on something you want to keep you can commit it. This prints the random settings into the steps and resets the Randomize amount to zero, so any further randomisation will be applied to this altered state. Switching to any other encoder mode cancels any randomisation being auditioned.

Take A Chance

Fill and Randomize generate sequences and step settings that are written into your Patterns. Chance sets up generative processes that run as your pattern plays. This provides real-time randomness, and ways to build in variations and cycles that are not tied to the 64-step pattern limit. Users familiar with Elektron's devices will feel at home here as it echoes their Trig Conditions' functionality. The Chance encoder lets you select an action and a chance of that action occurring. The most obvious thing is to set a condition for when the step plays. This can be a percentage probability or a frequency, for example play 50 percent of the time, or only play on the fourth loop.

There's more to Chance than trigger conditions: it can apply randomness to other settings, such as note (in scale), timing, filter cutoff, sample start, etc. It's brilliant, but I quickly wanted to be able to do more with it. For example, you can only have one Chance action running on each step. I found myself wanting to commit Chance actions in the same way as Randomizations, and add more. I also wanted to be able to

Outside The Box

Sequencing of external MIDI devices is handled in a separate view which gives you eight separate lanes and MIDI-specific encoder functionality. Like on the audio side, the Play is less track-centric than most other sequencers, and any step can trigger a different MIDI output, channel and program. You can create sequences on the MIDI tracks in similar ways to audio, although only Random and Euclidean Fills are available, and Randomize has a smaller set of options.

The MIDI tracks are billed as polyphonic, although only in the sense that you can assign preset chords to steps. For true polyphonic recording you need to select multiple tracks and your incoming notes get shared into separate voices. (It would be quite nice if you could sacrifice some tracks and designate a single track as poly for ease of handling). With a MIDI keyboard attached, the Play works as a little hub, redirecting incoming notes out to the last selected MIDI port/channel.

Fun stuff like track play modes work with MIDI, but sadly the Perform effects do not. This is understandable for audio effects like the filters and overdrive, but I hope Polyend will make banks like Transpose, Rearranger and Repeat available to MIDI sequencing. Of course it would have been great if the Play had an audio input like the Digitakit or Circuits, so you could integrate external elements and run them through the Perform and master effects.

apply Chance to other functions such as the Step Repeat or Performance Effects. (I'm reminded of the Step Components on Teenage Engineering's OP-Z). Talking with Polyend it sounds like they are already looking at ways to re-think and build on the Chance features.

Patterns & Variations

The main working block on the Play is a Pattern: eight lanes of up to 64 steps, with all the associated settings. As there is no separate mixer stage, a Pattern is self contained. Well, almost — Solo and Mute states persist across Pattern changes, as does Tempo, which could be a pain if you want to set up a live set in one Project. Patterns can be stored and recalled from a Pattern Grid much like on the Circuit. Tapping a Pattern during playback cues it, or you can launch instantly in step by holding Shift.

Adjacent Patterns on the grid can play as a chain, so you can lay out multi-Pattern sections or whole songs on the generous grid. Tapping a pad during Chain play will move the position to that point, and you can also flip in and out of looping a single Pattern at any time. All very nice.

The Play has an additional concept called Variations, which can be used alongside or even instead of Patterns. Variations are track-based alternate sequences within a Pattern. Each track has its own Variations lane (displayed via the pad column next to Solo) with 16 slots. You could make fill variations on a snare track, melody variations on a tone track, etc. But Variations within a track can be entirely different, with different samples, mixer settings and so on.

I think the main idea of Variations is that they are a way to add smooth progressions without saving a new Pattern for every small change. However, you can also treat them as your primary framework for arranging and performing. You can choose to show/hide all track Variations as a block instead of individually, and you can trigger a whole column of Variations at once, emulating the track and scene way of working from Ableton Live or Roland's MC-707.

I love the idea of Variations but didn't gel with it in practice, perhaps because it's another view to switch to, accessed in a different way to Patterns. Also the Variations view overrides everything else and can get in the way of Pattern selection, for example. I actually ended up fairly happy just using Patterns, taking advantage of the spacious grid to organise smaller variations and fills underneath my main Patterns.

Perform

My frustration with trying to make Variations fit into a workflow faded as I discovered more immediate and fun ways to create non-repetitive beats. The Play is set up for performance. Mutes and solos are always present, and Pattern changing is fast, but the real jewel in the crown is Perform mode, which transforms the entire grid into a palette of real-time transformers, split into a rainbow of effect banks.

There's a transpose section, low- and high-pass filters, overdrive and bit-crushing, all applied progressively up the pad columns like virtual faders. Then there are sequence re-arrangers (my favourite), beat repeaters with pitching (think jungle snares), and banks of delays and reverbs. All of these effects are applied to selected tracks, but the final bank is a sample-grabbing looper active on the combined output.

As a side note, a Project also has customisable master effects: EQ, limiter, Space (spread) as well as the delay and reverb.

There's loads of goodness in the way Perform mode works. You can combine effects from each bank; effects are momentary when held, or latched with a tap. Effects are only applied while you're in Perform view, and latched effects are remembered, so you can set up a cool transition or fill combo and punch it in whenever you want. Encoder parameters like Reverb, Filter and Sample Length can be tweaked in Perform mode, but are then reset when you exit out. This facilitates awesome breakdown/build-up madness such as flipping the samples on all tracks, maxing out reverb and filters, then kicking back into the main beat.

Finishing Thoughts

The Play offers a different take on the idea of a sample groovebox: it doesn't sample, chop or loop, you can't finger drum on it, and there's limited scope for sound design. Instead it's all about sequencing and real-time performance. And using it does feel like play. It clicked for me when I pulled back from the headspace of 'programming' a beat and relaxed into interacting with the generative and probabilistic features. Just randomising samples proved an inspiring source of happy accidents. I like that performing a dynamic live arrangement is relatively straightforward and relaxed, with the Perform effects and safe ways to go off the rails then get back on track. Comparable devices often require a lot of project setup and practice to get similar results.

There are some things that are clunky, or could work better. Polyend have already pushed two significant updates and promise more, with meaty new functionality. As with many drum machines and grooveboxes, my main outstanding question is how to capture what I've made. The first update added the ability to export Patterns and Chains as mixes or stems, which is great. I'd love to capture more of a performance, and with no multichannel output this leaves you recording a stereo mix. And maybe that's fine, but it would help if you could imprint more by automating Perform effect triggers or stacking Chance actions.

The Play takes a little getting used to; it's not going to be a fit for everyone's way of working. But I've found it a welcome change and an ideal companion in my synth setup, allowing me to get less hung up on beat creation and enjoy the bigger musical picture.

$ $799
W www.polyend.com

Wes Audio ngLeveler

Digitally Controlled Analogue Level Automation System

MATT HOUGHTON

The idea of digitally controlled VCAs is far from new — but this implementation feels like the future!

There have been several computer-controlled analogue level automation systems released over the years, so to some extent Wes Audio's ngLeveler reinvents the wheel. But what a reinvention it is! Essentially, you have 16 channels of VCA-based level control in a black 1U 19-inch rackmount box, with a bunch of LEDs that tell you what's going on inside it. The specs are good, with low distortion and crosstalk, plenty of headroom, and hardware digital control resolution of 2500 steps per channel. This translates to a very clean-sounding device, which offers very smooth analogue level adjustment. From an end-user point of view, once you've hooked up your gear to the ngLeveler's analogue I/O, it's really all about the cleverness of the software control app (which usually runs in the background) and the DAW plug-in (which runs in Mac OS and Windows hosts that support AU, VST2.4, VST3 or AAX). That said, there's also the possibility of integrating third-party control surfaces.

Getting Started

A good chunk of the 63-page manual is dedicated to setting things up, not because the ngLeveler is complicated but rather because you have the option of USB or Ethernet connectivity, and the manual necessarily goes into some detail about the different scenarios. For the purposes of this review, all you really need know about that is that the ngLeveler supports direct connection by USB or Ethernet, or Ethernet via a LAN, and that I had the review unit working successfully in all those setups.

The software side of things (for my Mac, at least) is a hefty 472MB download, and this unpacks to install nearly 1GB of data on your computer. The reason for this is that it includes the control software and plug-ins for all the 'ng' units in Wes Audio's portfolio: the installer checks all connected units for their firmware version and updates any for which there's a newer version. So while the download could be lighter, it's great that owners of lots of Wes gear have such a low-hassle way of update everything. With only the ngLeveler connected to an M1 MacBook Pro, this firmware check and update took under a minute.

Once installed, you have the two pieces of software I mentioned above. The GCon app opens automatically on boot-up and sits on your OS's menu bar. This is where you take care of all the admin: updating firmware, configuring MIDI ports for hardware controllers and so forth. The DAW plug-in communicates with the hardware (and any configured control surfaces) via the GCon app. All versions that are compatible with your OS appear to be installed automatically, so in my Reaper-based Mac system I could see AU, VST2 and VST3 versions; it might have been nice to have the option to install only those I required. The plug-in can be inserted on any track in your DAW project, since it doesn't receive or process any audio — it's just for control and metering — but in practice, to make it easy to locate, I found it was best to give it its own dedicated, labelled track.

I've used plenty of Wes gear connected by USB over the years and it's always worked smoothly right off the bat. This time, I found I had to iron out a wrinkle. Everything had appeared to install perfectly, and audio signals were clearly flowing through the hardware as they should, but when I instantiated the control plug-in it identified no attached hardware, despite the hardware being connected via USB. Restarting both my DAW and the hardware didn't solve this, so next I tried rebooting the Mac and, finally, it worked as I'd expected. A more detailed reading of the manual suggested that the installer should have prompted me to reboot my machine, which is a necessary step. It didn't, so that's a potential 'gotcha' to watch out for, though thankfully it's easy to overcome.

The hardware connections are all on the rear panel. These comprise four AES/Tascam standard DB25 D-sub connectors for the audio I/O, an RJ45 port and USB-B port for the control connection and firmware updates, and an IEC inlet for the internal universal switch-mode power supply. So that side of things really couldn't be simpler.

Wes Audio ngLeveler

$1999

PROS

- Superbly conceived plug-in control.
- Decent build quality.
- Good, clean sound, with optional THD coloration settings.
- Multiple units can be used together.
- MCU and HUI control surface support.

CONS

- None.

SUMMARY

This may not be the first level automation box to hit the market, but it's by far the slickest implementation I've seen to date, and could be useful in a number of different scenarios.

Around the front there's an on/off rocker switch, along with three LEDs to indicate when the unit is powered up, when it's connected to a computer, and when data is being received. Thirty-two more LEDs, one pair per channel, indicate input and output signal present. It couldn't be less confusing!

Plug-in GUI

A tour of the plug-in GUI quickly reveals what functions this mysterious 1U black box can perform, as well as the immense amount of control you have over it from your DAW. When you first open the plug-in, you'll be in Fader view, with 16 separate silver-capped virtual faders, one for each channel. These set the level over the full gain range of the VCA (full attenuation to +15dB). A button at the bottom allows you to switch to Trim view, which is differentiated on the GUI by its red fader caps. In this view, you have much finer control over a smaller range (-10 to +5dB). Importantly, both settings are active simultaneously, so you could, for example, use your DAW's automation system to move the regular faders, and then use the Trim faders to apply an offset without overriding that automation. When in Fader view, a number below each fader displays the corresponding Trim fader's setting, and in Trim mode it shows the main Fader level. In both views, input and output level meters appear to the right of each fader by default, though either or both of these can be turned off. A final level control is the input pad, which applies 6dB attenuation.

Each channel can, individually, be soloed (muting all other channels), muted, or set to solo safe (ie. not muted when another channel is soloed), and has the option of switching in a THD processor. The THD button cycles each channel through three distortion settings: off, medium and high. We've seen this facility on other Wes processors so I won't dwell on it too long here, other than to say that having 16 channels of this subtle coloration available is really handy. For a while, I used the ngLeveler along with the Dangerous 2-BUS-XT (reviewed elsewhere in this issue) to process a number of subgroup busses, and the flexibility this combination gave me to massage the tonality without using other processors was a pleasant surprise.

Above each pair and each quad of channels, there are a number of useful grouping controls. Adjacent odd/even channel pairs can be linked, as two »

Control Surfaces

Physical control surfaces are integrated using the GCon app — this acts as a hub, communicating with your control surfaces, with the DAW plug-in and the ngLeveler hardware.

The currently supported control surface protocols are HUI and Mackie Control, and these can be set up in a couple of ways: In Server mode, an attached control surface manages the ngLeveler exclusively, while in Server Mediation mode the control surface manages both the ngLeveler and your DAW, with any controls not being used by the ngLeveler (eg. transport controls) working with your DAW in the usual way. In either mode, it's possible to have any number of control surfaces managing any number of ngLevelers.

The manual also details how EuCon devices can be used, as well as indicating which knobs/faders/buttons on typical control surfaces are mapped to which functions. I tested the control surface support with my old Behringer BCR2000, connected over an M-Audio MIDISport 8x8s to a Windows 10 machine, and it worked as expected. Whether you need physical faders is an interesting question: my Windows system has a touchscreen, and because you can make the plug-in fill your screen it's pretty easy to control with touch.

The plug-in presents by default in Fader view. Channels can be linked as stereo or dual mono and, interestingly, can also be reverse-linked in mono or stereo pairs.

The Trim view, with red fader caps, offers finer adjustment. At the bottom of both views is a channel grouping feature, which can apply to any of four parameters.

mono channels (whereby they retain separate controls) or as a single stereo track. In the latter case, the tracks become one, with a single control set, and when used with an external control surface they present as a single channel too, so you could, for example, use one eight-fader control surface to control eight stereo pairs without changing banks. A more novel channel-linking mode, Link Opposite, can be applied to an odd/even channel pair, or to two channel pairs. As the name implies, when channels are linked in this way and you move the fader of one up or down, that of the other travels the same distance in the opposite direction. The idea is that you can send a signal from one channel to a processor, and bring the return to the other channel: as you 'drive' the processor by moving the 'send' fader up, the 'return' fader comes down to compensate for the level change. But you can also use it to crossfade between two mono or stereo signals.

A final linking facility is provided at the bottom of the GUI: each channel can be assigned to one (or none) of four channel Groups, and you can specify which of the following Group parameters you wish those channels to follow: Level, Trim, THD and Pad. Change this on any channel withing the Group and you change them all. It could be handy for all sorts of reasons, for example to treat multi-channel mic arrays as one or give you separate control over the OH/room and close mics on a drum kit recording.

There are a number of other nice, convenient touches. The GUI can be scaled to 100, 125, 150, 175 or 200 %; as a reference, 150% nearly filled the screen of my 16-inch M1 MacBook Pro. Alongside the expected preset save/recall system, an A/B/C compare facility allows you to experiment with and audition different balances or, say, THD settings across the board, with a single click. You can copy and paste settings between scenes too.

Settings are, of course, all saved and recalled with your DAW project, and pretty much any aspect of the plug-in can be automated. Interestingly, if you link two faders and then move either to record automation to your DAW, that automation is written only for the fader that you move. When reading the automation on playback, both faders will move but only for as long as they remain linked in the plug-in. This arrangement is necessary, since otherwise the onboard linking would fight against the incoming automation data. I mention it here as it has the potential to confuse if you fail to read the manual!

On The Level

For all the slick brilliance of the plug-in control, it's worth returning to where we started: the hardware. In use, it is every bit as clean and pristine as the specs suggest and, THD 'icing' aside, you wouldn't know it was there. It had no problem coping with high input levels, or passing on sufficient levels to drive the next device in the chain.

I used the ngLeveler in a number of different scenarios during my review tests. I've already mentioned using it alongside an analogue summing system for processing subgroups sent from my DAW, and in that role the stereo grouping was really handy. Such a 'virtual faders' role, effectively replacing the faders of an analogue console, is probably what this device will be used for most. But you needn't let your imagination rest there, since it's possible to have multiple devices connected to your computer, and you can use them to control and automate levels pretty much anywhere in your signal chain. I tried, for example, using the ngLeveler to automate the aux send levels from my analogue console's channels, riding the levels to outboard reverbs and delays, and that worked really well. I really valued the Link Opposite facility too, as a means of driving outboard gear into saturation or greater gain reduction without me needing to tweak the hardware.

There's genuine potential for one or more ngLevelers to breathe new life into vintage analogue consoles, and the possibilities if you were to integrate them into a wider hardware setup based around Wes Audio's other plug-in controlled processors are almost jaw-dropping Hmmm... I wonder if Wes will ever deliver a console with this tech embedded!

$ $1999
T MusicMax Distribution +1 614 897 0007.
E sales@musicmaxdistribution.com
W https://musicmaxdistribution.com
W https://wesaudio.com

Balanced audio I/O are provided on DB25 connectors, and the ngLeveler can communicate with a computer over USB or Ethernet.

LAUTEN AUDIO®

WHAT'S INSIDE

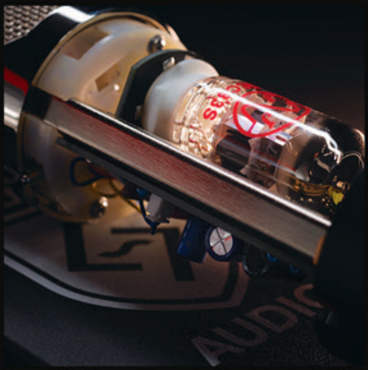

High-resolution capacitors & resistors, a JJ ECC83 dual triode vacuum tube, and a vintage-inspired output transformer.

CLASSIC OR CLEAN

Choose the clarity of a modern tube microphone or the color of a classic tube microphone at the flip of a switch.

A BETTER CAPSULE

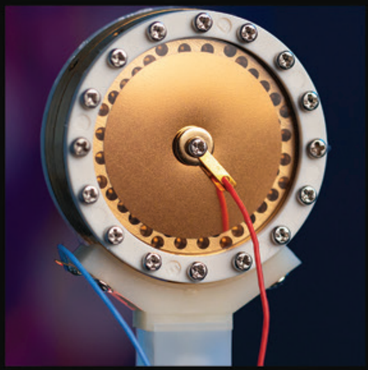

Gentle compression and a uniquely musical balancing of sounds starts with our new 32mm hand-tuned transducer.

Audient EVO SP8

Mic Preamp & ADAT Expander

Audient's EVO ecosystem grows further with an unusually intelligent ADAT expander!

SAM INGLIS

Audient's EVO 16 made an immediate splash on its launch last year. Designed to compete against products like the Focusrite Scarlett 18i20 and PreSonus Studio 1824c, it offers features not usually found in its price bracket, such as digital control over preamp gain, and a colour display featuring Audient's Motion UI graphics. With eight mic preamps built in, it also proved the worth of their Smartgain feature: the EVO 16 can 'listen' to your sources and set the mic gain automatically, which is a boon for self-recording musicians in particular.

The EVO 16 also features two sets of optical I/O. Depending on the sample rate, these allow up to 16 additional inputs and outputs to be added, making it possible to record really quite large multitrack projects. However, Audient's own ASP 800 and 880 expansion units don't quite fall into the same price bracket, so initial EVO 16 owners might have been driven to explore competitors' products such as the Focusrite Scarlett Octopre or Behringer ADA8200. Now, with the advent of the EVO SP8, that gap in Audient's own product line has been well and truly filled.

Spot The Difference

The EVO SP8 is almost identical to the EVO 16 in I/O complement and appearance. It even has a USB Type C socket, but this is employed only for firmware updates: the SP8 cannot act as a computer audio interface. Consequently, it lacks the EVO 16's headphone outputs, and its speaker and function buttons. As on the EVO 16, there are eight analogue inputs on combi sockets, the first two of which offer a 500kΩ high-impedance option for connecting electric guitars and so on. On the output side, there are eight line-level (+12dBu) outputs on balanced quarter-inch jacks. It features two optical inputs and outputs, to allow the full eight channels to be transferred at sample rates up to 96kHz, and a word-clock input.

As on the EVO 16, control is entirely digital. Each input has a button that selects it for editing, whereupon you can manually change the gain, phantom power, mute status and so on. Alternatively, you can engage Smartgain, which works exactly as it does on the EVO 16 (see our review from August 2022 for more details).

The EVO SP8 is compatible with any audio interface that has an ADAT optical input, and if you've used this method of expansion before, you'll be aware that it is fundamentally 'dumb'. The ADAT Lightpipe protocol allows the transfer of digital audio and clock information, and that's it. There's no provision for control data, and in fact the original Alesis ADAT machines had an entirely separate system for remote control and transport sync. So it may come as a surprise to find that Audient have belatedly managed to piggyback a control system on top of Lightpipe.

Fully Connected

For this to work properly, they recommend that the EVO 16 and SP8 be connected both in-to-out and out-to-in, even if you're not using the SP8's outputs. When you open the EVO mixer on your computer, you'll see the SP8's input channels

Audient EVO SP8
$499

PROS
- Offers eight channels of very decent preamplification at a keen price.
- Includes Audient's Smartgain and Motion UI technologies.
- Integrates with the EVO 16 interface to form a complete system.

CONS
- Confusing mute behaviour when paired with EVO 16.

SUMMARY
Designed to be paired with Audient's EVO 16 but also a very capable partner for other audio interfaces, the EVO SP8 is a well-priced and sophisticated ADAT expander.

labelled, by default, Digi 1-8 or 9-16. These can be incorporated into EVO 16 cue mixes just as you would with any other ADAT expander, but unlike generic expander inputs, these mixer channels also display gain controls, phantom power and instrument buttons. In other words, Audient have somehow implemented two-way communication between the SP8 and the EVO 16 such that settings made on the SP8 front panel are reflected in the EVO mixer, and *vice versa*. Even better is that this communication encompasses Smartgain. Press the green button on either EVO unit and you'll have the option to enable Smartgain on any selection of inputs from both. The benefits of Smartgain get greater the more channels you have, so the ability to do this in a fully expanded EVO rig with two SP8s attached is a pretty persuasive selling point.

At present I have only one negative to report from this, which relates to the way muting is handled. The SP8 and the EVO 16 allow individual inputs to be muted from the front panel, and the EVO mixer also displays mute buttons for each channel. For the SP8, these two mute options are independent: muting the input on the unit itself prevents any sound from reaching the computer, whilst muting the channel in the EVO mixer merely stops it appearing in any cue mixes. On the EVO 16, however, the two are linked, so muting the EVO mixer channel also mutes the input and *vice versa*. For my money, the SP8 arrangement is preferable, but either way, it's confusing, especially as there are no visual clues that reveal what's going on, and no way to see at a glance which inputs are muted on the front panels.

Once you're aware of this discrepancy, though, it's easily worked around, and I'm sure there will be scope for firmware updates to make things more consistent in future. It certainly doesn't undermine the benefits of systemwide Smartgain and two-way control over the preamp settings in your expander, which are obvious. The limitations of the ADAT protocol often mean that preamps added this way feel like the poor relation, but in the EVO system, they're all part of one happy family. And although you won't get the same integration with other audio interfaces, the SP8 is very aggressively priced and makes a highly competitive alternative to the ADAT expanders mentioned earlier.

$ $499
T Audient +44 (0)1256 381944
W www.audient.com

Kali Audio IN-UNF

The satellite speakers are coaxial designs, and use the same drivers as Kali's other IN-series monitors.

PAUL WHITE

Nearfield Monitors

Kali's new three-way system brings serious monitoring to your desktop.

Designed as part of Kali's Project Independence series of loudspeakers, the IN-UNF might at first appear to be a desktop 'multimedia' speaker system, but appearances can be deceptive. Yes, there are two small satellite speakers and a compact bass speaker, but those satellites uses the same four-inch midrange driver and one-inch soft-dome tweeter as Kali's respected IN-5 and IN-8 monitors, in this case passively crossing over at 2.8kHz (the bass/mid crossover, on the other hand, is active). The tweeter is arranged coaxially with the mid driver cone and is also phase-aligned so as to produce a coherent point source. Note that the mid driver cone is also shaped to act as a waveguide for the tweeter. In order to make the satellites adjustable in both the horizontal and vertical planes, they have a spherical form factor, and sit in specially designed silicone 'pucks' or cradles. Developed specifically for desktop monitoring, the design takes into account desktop reflections and also incorporates adjustable EQ to compensate for the boundary effect of working close to walls.

The system can easily generate a monitoring reference level of 85dBA at just 0.8 metres (roughly arm's length), and has a maximum SPL (peak) at 1m of 103dB. The Bass Unit measures 29.4 x 13.3 x 49.4cm and the satellite speakers are 15.8cm in diameter, so you don't need a lot of desk real estate. With an overall system weight of 11.2kg, the IN-UNF system is also portable. The amplifier ratings are 60 Watts (continuous) per channel for the mids and highs, and 100 Watts for the bass unit. A system frequency response of 47Hz-21kHz is quoted at the -3dB points (-10dB figures are 39Hz-25kHz). EQ options allow for adjustment of the LF, mid and HF levels.

Kali Audio IN-UNF
$599

PROS
- Ideal for smaller desk-based studios.
- Clean, well-balanced sound, especially in the mids and highs.
- Solidly built.

CONS
- Some bass overhang, but nothing excessive given the size of the system and its bass response.

SUMMARY
A compact yet surprisingly effective monitoring system that would be ideal for smaller desktop setups.

Kali describe the bass unit as the heart of the system. Frequencies below 280Hz are handled by a pair of horizontally opposed 4.5-inch low-frequency drivers located at either side of the case. The plastic parts of the case feel solid and well-damped, while the drivers are protected by perforated steel grilles. An advantage of them being horizontally opposed is that vibrations that could be passed through to the desk below should largely cancel out. The low end is summed to mono, as the vast majority of stereo information is carried by the mid and HF drivers. A benefit of this is that if there are any mix issues caused by opposite-panned low-frequency sounds that might cancel out when heard in mono, these issues become immediately obvious. In addition to the drivers, the bass unit also houses the system electronics, including the amplifiers and the physical TRS jack input connections, as well as banana-plug connections to the satellite speakers. These are arranged on the two end panels.

Other connections include a USB-C input supporting 24-bit audio at up to a 48kHz sample rates, as well as unbalanced 3.5mm jacks and optical inputs, extending the connectivity options to include gaming consoles and televisions. USB-C allows direct connection to a laptop or an iOS device via a Lightning connector and Apple's optional Camera Connection Kit. Power is delivered to the bass unit via a detachable mains cable, not an external adaptor.

Also sharing the panel space is a bank of DIP switches relating to the position and orientation of the bass unit. There's a volume control with a centre detent position and a Sleep switch to put the system into standby manually if the digital input is in use. If the digital input is not in use, standby is entered automatically after a period of inactivity, whereupon the power light turns from blue to red. Once audio starts playing again the speakers come back on line within a second or two.

Note that the bass unit needs to be located on a desktop, as it is not designed to be floor-standing and will most likely sound wrong if placed on the floor. It can, however, be set up either horizontally or vertically, so you can choose the arrangement that suits your desk layout.

Operation

For my tests I set up in my office, which also doubles as my 'studio B', positioning the bass unit vertically so that I could place it behind my laptop with a satellite speaker on either side. The way the satellites rest in their cradles makes it easy to aim the tweeters towards your ears, and there are detailed instructions on the Kali website if you need to adjust the DIP switches to account for room placement. At a comfortable playback level, there was very little in the way of unwanted vibration finding its way onto the surface of the desk.

While my appraisal of the sound quality is purely subjective, I have other speakers and headphones I regularly use for reference, and given the compact size of this system, I have to say that the IN-UNF is pretty impressive. Those coaxial satellite speakers provide a solid stereo image with great clarity and smoothness, while the bass unit does an excellent job of handling low bass notes. In comparison with my main studio monitors, which produce an exceptionally dry and tight low end, the IN-UNF does exhibit a degree of bass overhang, which is something we've come to expect from small ported monitors, but I'd still be happy to mix on these speakers in a small studio or desktop environment, given that you can always check how tight the bass really sounds using headphones.

Not having to worry about the placement of a floor-standing sub takes a lot of guesswork out of setting up reliable monitoring, and because of the physical format of the system, which feels reassuringly solid, it is very easy to fit into a desktop setup. If you have to work in a tight space but don't want to work exclusively on headphones, then I recommend you take a closer look at the Kali IN-UNF package.

$ $599
E sales@kaliaudio.com
W www.kaliaudio.com

The bass unit houses all the system's electronics, including DSP and amplification. Its two bass drivers are mounted on opposite sides, alongside the various I/O connections.

Caelum Audio Flux Pro

Modulation Effects Plug-in

Drawing some parallels with Cableguys' ShaperBox series of plug-ins, and perhaps overlapping with Output's Movement and Logic Pro's Step FX, Caelum's Flux Pro is an affordable, easy-to-use, graph-driven effects plug-in which, despite its apparent simplicity, can produce some addictive rhythmic effects that completely transform the source material. It runs on Mac OS, Windows and iOS hosts that support the VST3, AU, AAX or AUv3 plug-in formats.

Flux Pro's operation is based on three graphs that can be sync'ed to a time grid and then used to control four effects slots, into each of which can be loaded one of 10 modulation effects, including basic level modulation. Any or all of the effect parameters can be controlled by the graphs, which can be constructed with straight lines and Bézier curves, and sync'ed to a tempo grid, the outcome of which is to add a rhythmic effect or perhaps sync'ed effects sweeps to whatever's fed in, be it a drum loop or a continuous drone or pad.

Ready-made shapes (there are 48 of these) can be dragged directly onto the graphs, after which they can be further edited and saved in the User Bank if you don't want to build your own shapes from scratch. It is also possible to use MIDI triggering, both single-shot and retriggering, to initiate the modulations, and the graph shapes can be output as MIDI CCs to control other plug-ins or devices. If you have a DC-coupled audio interface, you can also output the graphs as CV information to control modular synths. Logic Pro X users will find the plug-in both in the Audio FX and MIDI Controlled Effects listings.

To get started, there are 115 presets, arranged by instrument and style, and the in-built effects are Filter 1, Filter 2, Delay, Delay (Fixed), Width, Phaser, Chorus/Flanger, Rate Modulation, Ring Modulation and Utility. The filters include formant modes. It's possible to create reverse-like effects, pitched up or down by as much as an octave, by controlling the Fixed Delay effect from a graph. And there's also an external sidechain input, so that the processing can be made to respond to a signal from another track. Offset buttons allow the graphs to be offset by a choice of musical measures.

The top section of the screen is given over to the graphs, and coloured selection buttons at the bottom left of the GUI are used to select graph A, B or C. A pulsing neon light follows the graph's path, with red nodes for position adjustment and blue nodes for creating curves. In the mid section of the screen, two of the four effects slots can be viewed at any one time, each having a sensibly modest number of controls. The percentage bar below each dial is used to set the amount of modulation from the current graph, with a yellow line appearing round the controls to show the range of modulation.

At the bottom right are the preset modulation shapes, which can be reversed or flipped vertically. In fact, the only trick I felt was missing is that you don't see the effect of parameter modulation in the lines around the controls, as you do in some plug-ins.

If you need more sophistication, then there's no arguing that you'll find more scope using ShaperBox 3, but I love the considered approach of Flux Pro, which strikes a good balance between simplicity of operation and sonic flexibility. The range of effects is impressively wide and it's easy to come up with results that sound musical. The cost is very modest too, so it's definitely worth a closer look. You can also check out Caelum's free 'taster' version, Flux Mini 2, which offers similar control over the level and a high-, low- or band-pass filter. *Paul White*

$ $59.99.
W www.caelumaudio.com

Caveman Audio BP1 & BP1 Compact

Bass Preamplifiers

The BP1 and BP1 Compact are superb bass guitar preamplifiers: they deliver slightly different audio performances but both are of the highest quality and suitable for live and/or studio use. Like all Caveman Audio products, they follow Steen Skrydstrup's principle that the signal leaving a guitar should reach the input of its amplifier without any additional noise being added, and I discussed this philosophy in detail in my AP1 review in *SOS* September 2022: www.soundonsound.com/reviews/caveman-audio-ap1.

Other than the Compact's steel chassis being 41 percent narrower than that of the BP1, there aren't massive differences in appearance between the two. Both front panels feature a Mute footswitch and a pair of detented, chickenhead knobs for gain and level. Although both pedals possess a serial effects loop, only the BP1 has a footswitch to activate it.

On both units' rear panels, you'll find a high-impedance instrument input (labelled Cable I/P on the BP1, and merely Input on the Compact); a transformer-balanced XLR DI out; and quarter-inch jack sockets for the effects loop send and return, and for the amp and tuner outputs. The BP1 possesses additional rear-panel space for the optional (€80 extra) transformer-balanced XLR Radio (wireless system) input; a 9V DC output to power an attached tuner; indicator lights that display the health of attached cables; a ground-lift switch; and an IEC mains connector for the onboard switch-mode power supply. Power to the BP1 Compact comes courtesy of an external, owner-supplied 12V DC centre-negative PSU. There is one invisible internal difference between the two units in that the tuner output of the BP1 is transformer-balanced, whereas that of the Compact is not.

Unsurprisingly, the BP1's interior, despite its greater volume, looks rather more congested than that of the BP1 Compact. After all, its input and outputs (save those of its loop and amp out) are transformer balanced, it carries an additional footswitch, it has to leave room for an optional Radio input, and it runs on an internal power supply! However, as you'd expect from Caveman Audio, the circuit boards, all-discrete componentry, wiring and construction in both products are all of a very high quality indeed.

Electronics-wise, the BP1 features the 5MΩ input impedance buffer found in the company's Custom Shop systems, together with the 1073-inspired preamplifier and output circuitry of the company's rackmount flagship BASC1 Bass System. Although the Compact's input buffer is a slightly redesigned version of the BP1's, it retains the 5MΩ input impedance and components. The Compact also features the 1073-inspired preamp, though its output stage is an original Skydstrup design. The circuit alterations in the Compact have resulted in a very slight difference in sonic delivery between it and the BP1.

> **"Playing bass through the BP1 for the first time was a revelation."**

Both units' latching footswitches feature 'silent switching', which eliminates switching transients when activating the Mute function or, on the BP1, the effects loop (which is always active on the Compact). Both pedals' effects loops sit after their respective input buffers, with both the loop send and loop return presenting output and input impedances optimised to interface with effects pedals. With this arrangement, your bass guitar's pickups are correctly loaded by the buffer, you get the best out of your pedals, and the result is a considerable improvement over an unbuffered, pre-input pedal chain. The amp out and tuner out in both pedals work perfectly as advertised, and also offer additional options in studio settings.

The BP1 offers more facilities than the BP1 Compact, and has an internal PSU.

Powered up and connected, via their DI outputs, to the microphone inputs of my RME audio interface, both the BP1 and the Compact were essentially silent in normal operation at my normal 85-95 dBA monitoring level, illustrating perfectly the efficacy of Steen Skrydstrup's approach to maintaining signal integrity. Playing bass through the BP1 for the first time was a revelation, the pedal delivering a crisp, clear definition to every note that I played, most notably in the bass and low midrange, without ever sounding harsh. Turning up the gain on the BP-1 brought an increasing level of saturation and a sense of weight that became quite addictive. The Compact, despite its electronic differences, delivered the »

The smaller, more affordable BP1 Compact.

same clear definition and gain-driven saturation in what I felt was a slightly warmer presentation which, to me, was very attractive in its own right.

The Caveman Audio BP1 and BP1 Compact are both superb bass guitar preamplifiers that deliver sonic performances which completely justify their price. Returning either of them, but especially the BP1, is going to be a wrench. If you're in the market for a pedal-format bass preamplifier of the very highest quality, then I highly recommend that you check them out.
Bob Thomas

$ BP1 $899, or $999 with Wireless input. BP1 Compact $575.
W https://caveman-audio.com

Pope Audio BAX 2020R Dual-channel Baxandall EQ

A staple of hi-fi amps since its invention in the 1950s, the humble Baxandall EQ comprises two gentle shelving filters. Put it in a high-quality professional product and it can be really useful, particularly in mastering and for stereo bus processing. I was impressed by just such a product, Pope Audio's BAX 2020 500-series EQ module, when I reviewed it a couple of years ago (www.soundonsound.com/reviews/pope-audio-bax2020). In particular, I found that its Hi control offered me a simple, forgiving way to add top end to individual tracks or a whole mix. Building on its success, Pope have now developed a dual-channel 19-inch rackmount version, the BAX 2020R, though there's a little more to it than popping two of the 500-series modules and a power supply into a rack.

Not only does the rackmount version boast 6dB more headroom, but it adds another EQ option too: the Hi band of the module was fixed at 20kHz, and in the 2020R this is switchable to 40kHz. This might strike you as being unnecessarily high, but the filter's gentle slope means that its effect can reach much lower down the spectrum, and you can use it to lift the 'air'. The Lo shelf is, as on the original module, fixed at 20Hz. Again, while that might sound extreme, the filter's curve reaches well up the spectrum, and it allows you to add weight or reduce low-frequency build-up.

When using the 2020R as a mix-bus EQ, I could hear a subtle but pleasing sense of saturation, even with no equalisation being applied. The 500-series version offers some custom op-amp options, but in this rackmount version there's a fixed choice: essentially, it's what was considered the 'full-fat' model for the module, with Avedis 1122 op-amps at the input and driving the output transformer. It's perhaps due to the bass-heavy nature of some of the material I was working on during this

review, but I found that I appreciated this coloration more than I did when writing my original review of the modules. In fact, I often found myself leaving the 2020R in my mix-bus chain solely for this purpose.

"The new 40kHz Hi option is a great addition, and one that I think will (pardon the pun!) really open up this unit to mastering engineers."

When I did engage the EQ, though, it felt instantly usable and familiar; a great option for gently shifting the overall tonal vibe of a mix.

This EQ is not just for bus processing, though: it's also a surprisingly useful tracking tool. In a busy tracking session, you often don't have the luxury of time to allow you to dial in parametric EQ settings precisely, so a tool that you can reach for to quickly nudge the top and bottom into the right tonal territory is often just the ticket.

The new 40kHz Hi option is a great addition, and one that I think will (pardon the pun!) really open up this unit to mastering engineers. As forgiving as the shelving EQ curves are on a Baxandall EQ, the 20kHz option can still be a little heavy-handed when you wish to make really delicate adjustments to sources that contain shrill-sounding cymbals, or any midrange unpleasantness around 2kHz. I quickly found myself preferring the 40kHz frequency option on a whole mix too, when dealing with finer mastering-style judgements. I loved that I could gently shift the overall tonal balance with a high-end lift, without emphasising any harshness lower down.

If you're interested in using this EQ for tracking and mixing, I think you'd be well covered by either this rackmount model or a pair of the 500-series modules. But if you're approaching this from more of a mastering perspective, the 40kHz option and additional headroom should make the 2020R the more compelling model. *Neil Rogers*

$ $1535.
W www.popeaudio.co.uk

HANDS-ON

MIDI Mapping For Intuitive Sound Design

Getting tactile with the MIDI controls can breathe new life into stale music.

OLI FREKE

A lot of sound-design tutorials focus on the technical aspects of the tools at our disposal — how they work and what they do. Whilst it's always an advantage to have a theoretical understanding, this article will take a less analytical approach. Instead of designing a sound from scratch by working through the basics, we'll discuss a way to design sound by working with an existing phrase or loop directly. The advantage of this technique is that it ignores the screen and relies on your ears and brain to feel the right settings.

We'll be mapping hardware knobs, faders and buttons to various sound control parameters in order to manipulate them in real time. Which sounds pretty straightforward and basically what analogue synths have always done... but there are some subtleties that bring this to life and (I've found) make the whole sound-design process far more fun — and productive.

This technique also solves the age-old issue of stale loops. We've all been there; a burst of creativity generates a great-sounding loop, but after a 'mere' two or three hours of repeating the same eight bars, it gets boring for some inexplicable reason! Mapping controllers to multiple parameters and recording a jam session with them can solve this, creating new and surprising sounds and ideas for arrangements and key moments in a piece of music. And the beauty is that once the principles are understood, this can work with any recorded audio, MIDI part or musical element in a mix by adapting it as needed.

The Setup

The best way to get a feel for this technique is to try it. What you'll need:

- A musical phrase in your DAW or equivalent. This could be a MIDI part triggering a soft synth or a piece of pre-recorded audio. MIDI is best as you can get into the real parameters of the synth, but an audio phrase can be used by inserting effects plug-ins and manipulating those instead.
- A control surface with any combination of knobs, faders or buttons that can send MIDI control data to your DAW.

The key thing with this simple setup is to make sure you can control various parameters of the sound in question from the control knobs and faders, and to make sure that there's one knob per parameter and no menu-diving. So for example, knob one could control filter cutoff, knob two controls resonance, knob three controls the decay time of the volume envelope, and a fader is controlling an effects return level.

If you're working with pre-existing audio, then you can apply modulation to plug-ins placed after that audio, such as filters, EQ, delay or something more exotic.

Take some time to map a decent number of parameters to controllers. There are obvious ones like filter, resonance, LFOs and envelopes, but it's definitely worth trying some less obvious parameters, just to see what happens. If you don't have a lot of handy control knobs, mapping mod and pitch wheels from a synth to two parameters at a time is still going to give some results and is worth trying. Just overdub a few passes using different parameters.

Jamming With The Knobs

So far so good, you might say. Sounds like you've recreated a basic synthesizer, what's the big deal? Well, the magic is in seeing what the setup can do. Instead of just playing the keyboard and giving a knob a quick tweak every so often, what we want to do here is to loop a part and find out what that phrase is really capable of, from a sonic perspective.

Because there are (hopefully!) going to be at least three different parameters interacting, by experimenting with them as the loop plays, one can cover a lot of ground quickly to find a sound that sympathises with the musical part really naturally. And it might be from surprising settings that you wouldn't have thought to 'mathematically' design. Remember, you're twiddling knobs and listening to the results rather than drawing visual automation on a screen. This is very definitely a mouse-free activity!

Arrangement-driven Experiments

Not only this; don't be content once you've found something that works. Especially in loop-based music such as techno, continually tweaking these parameters over the duration is what keeps such music (when it works) alive and interesting. The trick is to record a good few minutes of »

DAW Control

Most DAWs make it pretty easy to map MIDI controllers to parameters nowadays, but here are a handful of examples for some of the most popular.

Ableton Live

Ableton Live has long been able to map external MIDI controllers to virtually any parameter with a couple of swift moves:

1. Click 'MIDI' in top tool bar.
2. Click the synth or plug-in parameter you want to control.
3. Move the external control knob, fader or button you wish to map that control to.
4. Click MIDI again, and that's it.

Apple Logic

1. Click the synth or plug-in parameter you wish to control.
2. Go to Preferences / Control Surfaces / Learn Assignment (or hit Apple+L).
3. Move the hardware knob you wish to control the parameter.
4. Click 'Learn Mode' to complete.

Steinberg Cubase

1. Prerequisite: Your hardware device has been mapped, as shown in the MIDI Remote panel (either from the presets list or manually).
2. Right-click the synth or plug-in parameter you want to control, and select 'Pick for Remote Mapping: parameter name' from the drop down menu.
3. This opens the MIDI Remote Mapping Assistant, and moving the knob or fader will connect that parameter.
4. Click 'Apply Mapping' and it's mapped.

Once a MIDI controller is mapped and the phrase looped there are two choices: Record the controller data as MIDI data as the phrase loops, or set up an audio channel to record the output as audio. There are advantages and disadvantages to both, but the end result of either allows you to listen back to your manipulations and select the best sections to use. MIDI can be fiddly and sometimes unpredictable, but does allow for later changes if required; audio can't be altered after the recording, but I personally prefer this as it forces you to work with what you have — and there's usually more than enough great material to choose from that it hardly feels like a limitation!

MIDI mapping in Ableton Live couldn't be simpler...

these tweaks and then select the best bits. The recording might comprise the controller data itself or simply the resulting audio on a new audio track. They both have their advantages: MIDI can be tweaked if something was nearly perfect, but audio forces the decision to be made and you to move swiftly and decisively on.

And don't be afraid to go extremes — high resonance, fastest LFOs, lowest filter with highest envelope — it doesn't matter if it doesn't all sound great all the time, you are almost guaranteed to stumble across something unexpected and wonderful in the process.

By the way, I'm quite convinced that all those articles about 'how artist X made sound Y' are misguided, in the sense that it's better to just create new sounds using techniques like this, and finding the unexpected, than trying to reverse-engineer someone else's happy accident that they themselves are very unlikely to be able to recreate!

Feel The Music

The reason this works so well, I believe, is because one's hands are directly connected to the sound in real time. We're back to 'feeling' the effect we have on the music and we can be guided by our taste as to what sounds good and what doesn't. We're freed from the tyranny of looking at the screen and thinking, for example, that the EQ graph looks a bit 'non-standard', worrying about it and — worst of all — changing it back to something that 'looks right'. And the mere fact of not letting that happen is half the battle. Using these techniques one can:

- Fine-tune the best settings for a part in a song or track.
- Create an arrangement and discover unexpected 'key moments'.
- Refine a whole groove. Why limit yourself to one sound's parameters when a kick, bass line, lead and drums can all be adjusted alongside each other? (I've found this one to be the most fun in terms of discovering the groove.)

Do Try This At Home

A great exercise to prove why this is so powerful: create a random one-bar 16th note pattern. And I do mean totally random — atonal, multi-octave, the works! Then choose a starting sound of a classic analogue waveform — something rich like a sawtooth. With a good collection of knobs mapped to parameters challenge yourself to manipulate the sound until it sounds good. There will generally always be a setting that makes it work — or at the very least suggests a usable musical fragment — with built-in excellent sound design. And I'll bet it won't be a sound-design setting that you would have deliberately created first and which then would have inspired you to compose that random pattern... Food for thought!

And let me finish by reminding you that looping a random bar is the exact method Phuture used for their ground-breaking and genre-defining 'Acid Tracks' back in 1987. What is that track but a 12-minute exploration of that eight-note loop and all the sonic variation that could be extracted from just cutoff, resonance and envelope? And inadvertently kick-starting decades of acid house and techno. The power of the intuitive live tweak! So, free your eyes and mouse and get back to feeling the groove!

A Novation Launch Control XL is perfect for hands-on MIDI manipulation. Note the annotations on tape — it's easy to forget what's controlling what when you come back in the next day...

Control Surfaces

There are an overwhelming number of MIDI keyboards and controllers available these days. A few are listed below, but knobs and faders have been sending continuous controller (CC) data for decades so there should be a device to suit any need and budget, new or second-hand.

Almost all modern keyboard MIDI controllers come with some, more, or many knobs and faders. Key manufacturers include Akai, M-Audio, Arturia and Novation, who have models to suit almost any budget. On the other hand, If you already have a MIDI keyboard of some sort (and if you're reading this I shall be highly surprised if you do not!), it can be augmented by a dedicated box of control knobs, buttons and faders in various combinations. I prefer knobs as faders always feel more 'volume' based to me, but this is an entirely personal preference! Buttons are only useful for on/off (which has its place), but for me, something like a Novation Launch Control XL has the perfect blend — 32 knobs, eight faders and a few buttons — as does the Akai MIDIMix. Alternatives include the budget-friendly Korg nanoKontrol2, with eight knobs and faders, or the discontinued Behringer BCR2000 if you can pick one up second-hand.

Lost In Translation

Understanding Client Feedback

As a mix engineer, it's your job to give clients what they want — but what they say and what they mean aren't always quite the same...

NEIL ROGERS

If you have even a small amount of experience mixing music for other people, there's a good chance you've found yourself on the receiving end of vague, contradictory or slightly abstract mix feedback. Exactly what does a 'glassy' guitar sound like? How do you add 'more excitement' to a mix that you've already tried to make sound as exciting as you possibly can? In this article, I'm going to dive into the top five 'client speak' comments that I encounter as a mixing engineer, and explore not only what they might mean, but also how we can respond to them and tweak our mixes to get everyone on side.

"Can the vocal be a bit louder/quieter?"

Feedback around the theme of vocal levels can often be an easy fix. Frequently, all that's needed is a small adjustment in overall volume, or a little riding of levels at different points in the track. Sometimes it can simply be about reassuring an inexperienced artist that it's perfectly normal for a vocal to sit 'loud and proud' in many styles of music.

If you're struggling beyond simple level changes, one likely problem is that other elements in the mix, such as guitars or keyboards, are fighting for space with the vocal. Although this is likely to be an arrangement issue, sometimes it just needs to be dealt with in the mix. The choice of vocal effects is a big factor, but we'll look at this point later in the article. Initially, I would recommend keeping things simple, and it's worth having a quick review of your mix to make sure you haven't overlooked basic

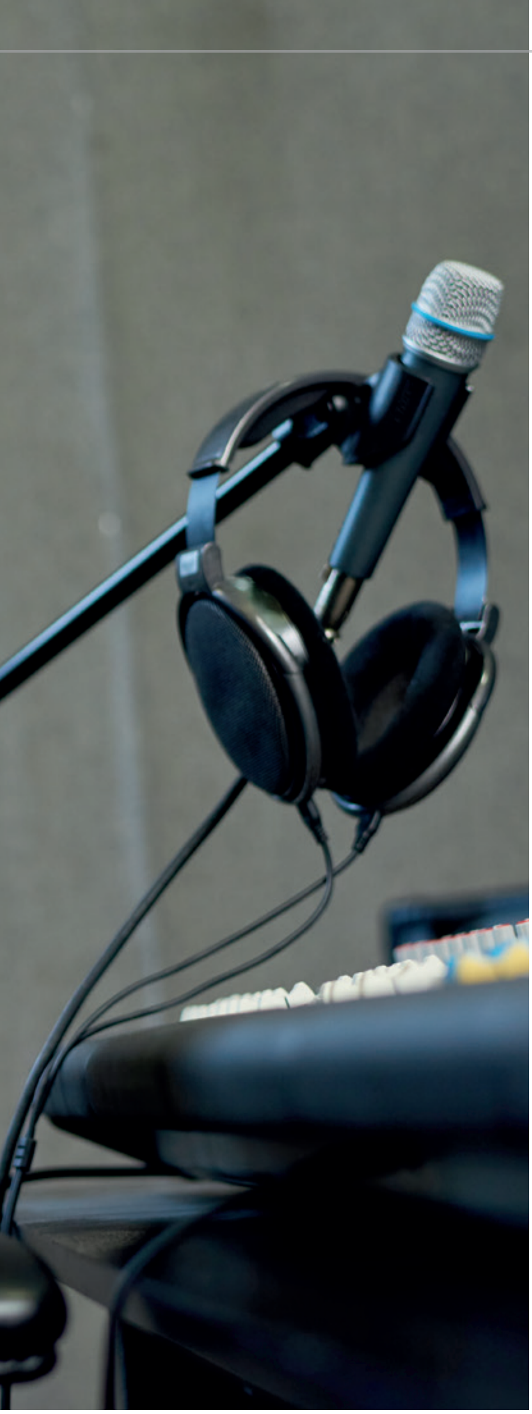

mix options such as panning things out of the centre of the stereo field, or using enough level automation to bring other instruments up and down in level around the vocal.

If none of the above is enough, we can look at exploring a few more ‘advanced’ mixing techniques. One technique that you might often hear mentioned but which can be quite hard to master is using a plain old EQ to remove competing frequency ranges from other instruments that are sitting around the lead vocal. A common mistake I often see is excessive boosting of the top end ‘air’ frequencies of a vocal in search of a quick fix. I would encourage you instead to focus more on the midrange of your mix, and a useful technique to experiment with is to identify a frequency range where it feels like the intelligibility and ‘excitement’ of the vocal are located. Depending on the voice, this is usually somewhere between 1.5 and 5 kHz. Try to then ‘carve out’ a little space in that region on a few other instruments and see if the vocal seems to sit more naturally. If it feels like the other instruments are being compromised too much, you can then explore more ‘dynamic’ techniques such as side-chaining an EQ or multiband compressor, so that the key frequencies you identified are ducked out of competing instruments only when the vocal is present. This is relatively easy to do now in a DAW and there are a couple of plug-ins, like Trackspacer by Wavesfactory, that make this very easy indeed. A little should go a long way, however, and I’d definitely recommend trying to develop your hearing around some of the more old-school ‘static’ EQ techniques — it’s a way of working that can lead to leaner and more focused-sounding mixes.

“More bass please!”

One mix session that always sticks in my mind involved a brief stand-off with a band who had got completely lost in a ‘committee-style’ mixing process. It reached the point where I had to politely »

If a client asks for 'more bass', try and establish whether they want the bass guitar louder, or the mixer to be bassier as a whole.

draw a line as we started going through the bass player's partner's mix notes — which primarily involved turning everyone else down and the bass guitar up.

Anecdotes aside, mixing bass can be tricky, and on top of the well-documented problems of mixing bass in small, bedroom-style studios, we also have to contend with modern playback systems, which often sell themselves on their ability to produce exaggerated bass frequencies from modest-sized devices. The first thing we need to figure out when we get client feedback around the low end is to decipher whether the artist (and often, especially, the bass player) wants to be able to hear the bass guitar or synth more clearly, or whether they want the overall mix to have more or less bottom-end 'weight'. If it's an intelligibility issue, we can follow a similar process to the one I suggested for vocals, creating space for a bass part by removing clashing frequencies from other instruments. This can be more straightforward than it is with vocals, as there's often plenty of superfluous low end that can be filtered out on guitars and synth parts. It's also less risky to add frequencies on a bass part, and the area around 800Hz up to 2kHz on a bass guitar can often be quite dramatically boosted before it starts to sound unpleasant or interfere too much with other parts.

If the issue is more about the overall 'weight' of a mix, one skill relating to bass that is harder to develop is learning to appreciate the difference between the real low frequencies — anything below around 70Hz — and the more audible range just above this but below the lower midrange. I had a brief period mixing on Yamaha NS10s a few years ago, and was shocked at just how much of a 'bump' those classic studio speakers have around 110-120 Hz. As a result, the mixes I was doing at that time were quite lean around that area. I quickly noticed, however, that those mixes seemed to 'travel well' outside of the studio — especially on AirPods and the like. I find that If you're careful with not overloading that 100-150 Hz frequency range, you can then have a more generous lower end, with more sub-type frequencies that will be appreciated when listening on better speakers or headphones. This can keep your bass-hungry clients happy but also ensure that the mix isn't swamped when heard on modern playback systems.

Compression is the other big factor in achieving a tight, professional, client-pleasing low end, and it's easy to underestimate just how much compression is sometimes used on bass in modern mixing. Without going too deep here, it's worth trying to separate bass compression into a couple of different jobs. The first priority is often to 'pin down' the really low end, so that individual notes don't dramatically go up or down in volume. This might typically involve quite heavy-handed use of high-ratio compression, or even limiting, to create a solid low end that can be nice and prominent in the mix. I like to focus just on the sub-200Hz range, with a multiband compressor or by multing the bass track. Once that range is more controlled, you can start to think more about the presence and character of the bass, and shape it with another compressor that won't just be triggered by the low end. It can be hard to get to the point where you feel like you have real control over the bottom end, but once you develop this, it can make a huge difference in giving clients what they want, whilst preserving the integrity of the mix that you've worked so hard on.

"I'm not sure about that reverb..."

Another common topic of mix feedback is the choice of effects. And with effects

In your quest for a 'shinier' mix, it can be tempting to simply boost the high frequencies across the whole mix — but it often pays to be a bit more selective.


Classic warm, silky character
Precision-engineered 1” gold-sputtered capsule
Ultra-high-resolution A/D (192kHz)
Ultra-low-noise (4dBa)
Advanced digital signal processing
High SPL handling
THE ULTIMATE STUDIO MICROPHONE
Dual Connect output with XLR and USB connectivity
32-bit float digital output
MADE IN AUSTRALIA
NT1 5TH GENERATION
STUDIO CONDENSER MICROPHONE
RØDE
The Choice of Today’s Creative Generation.™
rode.com

Although mixing remotely is now the norm, getting the clients involved can help give them a sense of 'ownership' of the mix.

Communication Is Key

The way music is mixed has changed a great deal in recent decades, and it's noticeable just how comfortable bands and artists now are with remote mixing. The artist not being present for the mixing stage has many advantages, but it also creates problems if not managed correctly. I'm old enough to remember all-analogue mixing sessions where everyone in the band might be riding a fader during the later stages of a mix. The thought of this makes me shudder now as an engineer, but I can clearly remember feeling, as an artist, that we were all responsible for the mix and that the engineer was the skilled person helping everyone get what they wanted.

You don't have to go completely old-school, but I do think it's important to try to give artists a sense of 'ownership' of their mix. This can be as simple as having a proper chat before you start, and identifying some reference tracks. As a mix develops, you might then find using screen-sharing tools such as Audiomovers (or even real humans in the same room!) helps get a mix finished quicker. It might seem like extra work, or an annoyance, but when you come across issues like I've discussed in this article, it becomes an easier and more collaborative process to unpick any problems, or even to avoid them in the first place.

being one of the most subjective aspects of mixing, that's perhaps not surprising. I have some clients who egg me on to add more and more reverb, and others I have to negotiate with to allow me to add even a hint of ambience around their voice. There's no right and wrong, of course, but artists can easily overthink and get insecure about this common and often liberally used element of music production. Often they just need a little reassurance that the listener won't think it's weird or unusual! If I'm recording a band, and I know I'm going to be mixing the material we're working on, I make a point of trying to get to know what they like effects-wise during the tracking stage, but if you're just mixing for someone, it's well worth trying to feel out what they might like — see the box on mix communication.

Aside from taste, it's important that effects feel like they're part of the music and not just 'stuck on'. There are a few easy ways of making effects feel more natural and cohesive, with the most obvious being the use of EQ on the effect return channels. Removing high frequencies — from reverb and delay returns especially — can help make an effect 'sit well' and can also help with other mix issues such as sibilance. Cutting low frequencies from the effects returns can help avoid low-mid build-up.

Automation is also your friend here, and riding the levels of effects throughout a mix is an easy way of keeping them under control while also adding interest. Dynamics processing can also be great for getting that balance between presence and depth, and a common technique is using a side-chained compressor on the effects return to duck the level of the effect when a vocal is present. Every time the vocal stops, the compressor is released and the effects swell up and catch the ear. One more little tip that I've found helps, with stereo spatial effects especially, is using a good pair of headphones to fine-tune and set my final effect levels — I normally always end up turning them down a touch! Not only can you hear the finer details to make these judgments, but there's also a very good chance nowadays that your client will be judging your work whilst listening on a pair of AirPods or similar. Lastly, effects, more than anything else in this article, are something to not be precious about as an engineer — it's a simple matter of taste. Be bold, but if the client doesn't like it, save that awesome effect you've created as a template for a future project and move on!

"Make it brighter/ sparklier/glassier/ shinier!"

It would be easy to think that when an artist uses these adjectives, they always want the mix to be rebalanced tone-wise to sound 'brighter'. Often that is the case, and if you feel confident with your monitoring, it could just be a case of you not quite judging a genre-specific norm right (some styles are really bright!), in which case you simply need to dial in a bit more high-frequency EQ. But when people use terms like 'sparkly' or 'shiny' (I still have no idea what 'glassy' means!) it's often worth taking a moment to try to figure out exactly what they mean. By just boosting the very high frequencies across a whole mix with a shelving EQ, we can end up chasing our tails as things like cymbals or vocals become too shrill or harsh sounding. Often, brightening up just one or two elements, such as a guitar or snare drum, will mean you can get the thumbs up without trashing the rest of the mix.

Another cause of a mix appearing too dull out in the real world is that it has too much low end or low-mid information, which our ears can perceive as a lack of brightness. Mastering engineers, for example, often get that nice 'shiny' sound not by just adding loads of high frequencies, but by first carefully removing low mids (somewhere around 250Hz is a great place to start) and controlling the very bottom end with compression and a little EQ. Hopefully, some of my tips in my 'More Bass Please' section can help with this.

"Can it be more exciting?"

Being asked to make a mix more exciting can be deflating, because it's not like you haven't already tried your best to create a big, impactful-sounding mix! However, you can't take this stuff personally: it's about figuring out what you have under your control and what the client actually means. It could be a simple(ish) issue of you not pushing things enough with saturation and parallel compression, say, or you may have misjudged certain genre-specific stylistic things like distorted vocals or the use of drum samples.

It could also, frustratingly, be a case of your mix not being loud enough. Thanks to the loudness normalisation features on streaming services, the argument for slamming mixes has, thankfully, become much less compelling. What hasn't changed, however, is that many artists still associate that 'crispy' and clipped sound you get from heavy limiting with a feeling of excitement. What I tend to do if I realise a client likes that 'noughties' kind of sound is to try to build it into the mix itself, rather than just slap it on the end with a limiter. If you use saturation and limiting on individual tracks, or groups of instruments, you then have the freedom to automate some dynamics back into the mix. Hopefully, you can get that 'exciting' sound but also ensure that the mix doesn't become boring after the first 20 seconds.

A great bit of advice I was given on the topic of excitement by a well-known producer was about not being timid when mixing music. A neat and tidy mix, with everything politely coming in at just the right level, might please us technically, but can often sound boring. If a new part, chorus or effect is coming in during a track then don't be afraid for it to really come in! A great exercise if you don't do so already is to listen to music you like on your studio monitors or headphones. I'm regularly shocked at just how 'bold' certain production choices — especially around effects or the low end — are when heard in the more detailed clinical environment of the studio. We don't want gimmicky mix choices getting in the way of a great song, of course, but the mixes that catch my interest out in the real world sound bold, exciting and creative. Achieving this is easier said than done, but it's often necessary to remind ourselves that this is, ultimately, what we're trying to achieve.

MIX RESCUE

Adding Energy

Ever thought you were nearing the end of your mix, only to realise the track needs a whole lot more impact?

MIKE SENIOR

How do you add energy to a mix? I mean, I imagine that every DAW user has at one time or another been faced with a project where they've gone through all the usual mixdown motions, but ended up with a result that feels somehow a bit lacklustre, plodding, and generally uninspiring. In other words, it lacks energy. And in my experience, it's this malaise that most commonly leads *SOS* readers to approach the Mix Rescue column. With that in mind, I'd like to focus this article on a mix overhaul I recently did for singer-songwriter Jeff Hirata, where I applied a variety of different techniques I've learnt over the years for enhancing this rather intangible mix quality.

Energetic EQ

When Jeff sent me an MP3 of his song 'Sunshine', he was already well aware that his mix didn't feel as exciting as he'd hoped, given the cheery musical content. But he was stumped about how to improve matters. When he added more low end, or more compression, or extra guitar layers, it just made things muddy and sacrificed clarity. When he pushed the drums up in the mix, the sound became too aggressive and lacked cohesion. When he tried boosting brightness it just made the overall sonics abrasive and fatiguing. More insidiously, though, he just felt that the mix was making the music seem almost boring, despite what I agreed was solid songwriting and a decent set of performances.

Now, Jeff wasn't wrong by any means in reaching for the mix tools he did, but in each case the problems he was addressing ran a little deeper than he'd identified. Take the idea of boosting the low end, for instance. In principle, this is a great way to make a mix sound more exciting, but there were two good reasons why I was able to do this in my remix where he'd previously struggled. Firstly, I cleared space for low-end boosts on the kick drum and bass guitar by filtering out low end from more than a dozen other tracks: drum overheads, drum room, low tom, congas, steel pans, acoustic/electric guitars, and backing vocals. And, secondly, I realised that the fundamental frequencies of Jeff's bass-guitar part seldom strayed below 50Hz, which left plenty of free mix real-estate for me to add a sneaky programmed sub synth layer. So, yes, more low end on your kick and bass will usually translate into a more exciting mix, but only if you plan enough space in the spectrum for that to fit.

Likewise, Jeff was right that brightening a mix often makes it feel more urgent and immediate, but you need to work in a targeted manner. So while I certainly brightened my remix with high-frequency EQ boosts and saturation effects, I was able to avoid the harshness problems Jeff had originally encountered, by:

- Filtering high end out of less foreground parts, such as the hi-hat and hand percussion.
- Softening fatiguing upper-spectrum transients with low-pass filtering (for the kick), clipping (for the snare), limiting (for the tambourine), and multiband limiting (acoustic guitars).
- Rebalancing overprominent noise consonants and sporadic high-frequency resonances on the (numerous!) vocal parts with de-essing, multiband compression, specialist spectral compression (from ProAudioDSP's DSM plug-in), and region-specific EQ processing.

Lively Dynamics

Jeff's impulse to lean on his compressors had something going for it too, because compression can indeed add liveliness and movement to the musical balance. But the catch is that compressing the wrong things or dialling in the wrong settings can just as easily kill a mix stone dead! My main advice here is not to get too carried away with per-track compression settings (where you run greater risk of compromising the musicality of each individual part), and focus more on compressing 'ensemble' signals. So in my remix, for example, I had no compression at all on any of the software drummer's individual instrument channels, but instead compressed the drum room mics, the drum kit submix, the main mix bus, and a drums parallel channel. These compressors introduced subtle music-related level interactions between each drum and the rest of the arrangement, thereby providing a more energetic-sounding mix without a loss of musicality — as well as an increased »

Structural Cuts

In addition to the arrangement adjustments discussed in the main text, I also made a couple of structural changes. Like a lot of project-studio song productions, the song opened with an instrumental intro that was basically just treading water on the verse chord progression and, from a singer-singwriter's point of view, that's rarely the best use of the audience's first 10 seconds of attention! So I decided to replace the preamble with just a little guitar fill, and bring Jeff in right from the outset to begin telling the song's story and building that all-important personal connection with the listener.

Likewise, I saw no reason to keep a very similar instrumental section between the end of the first chorus and the beginning of the second verse. All it seemed to be doing was smoothing out arrangement differences between those two song sections, which seemed counterproductive given how effectively an abrupt textural change at this point in a song can grab the ear for the second round of verse lyrics. Again, I just excised this section, as well as further thinning out the beginning of the second-verse arrangement to make even more of a virtue of the section contrast.

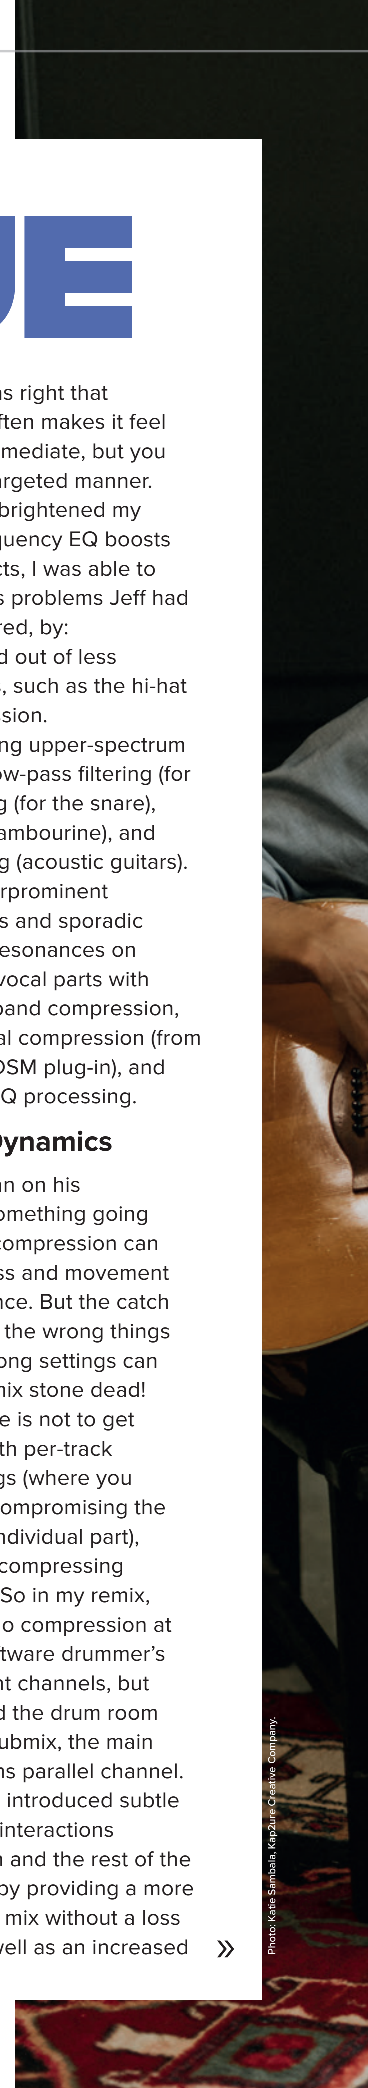

Photo: Katie Sambala, Kap2ure Creative Company.

This month's featured artist is Jeff Hirata (https://linktr.ee/jeffhirata) with his song 'Sunshine'.

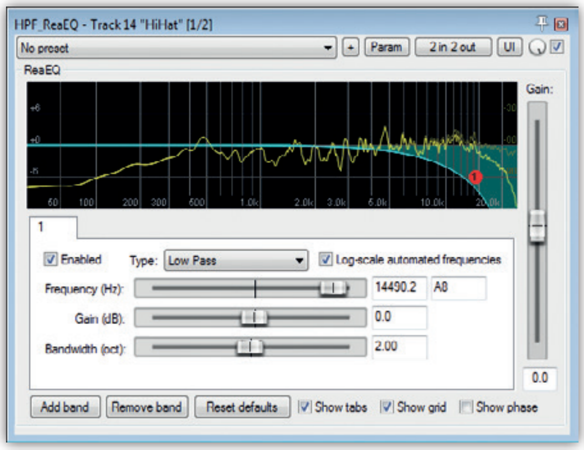

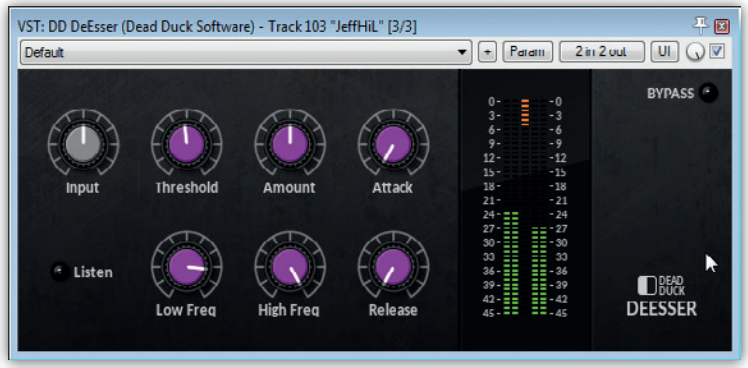

Here you can see various different processes Mike used to add brightness to Jeff's mix without making it sound harsh at the same time: cutting high frequencies from the hi-hat with Cockos' freeware ReaEQ; smoothing the tambourine transients with Dead Duck's freeware Limiter; de-essing backing multiple vocals with Dead Duck's freeware De-esser; and controlling the lead vocal's upper-spectrum consonants and resonances with Pro Audio DSP's DSM spectral processor.

illusion of ensemble cohesion into the bargain!

But however you decide to use compressors for your mix, it's vital to understand their limitations, so you don't expect them to deal with mix-balance problems they can't reasonably be expected to handle. Nowhere is this more important than with lead vocals, where no compressor is intelligent enough to consider lyric intelligibility or melodic phrasing — all it sees are signal levels. In my 'Sunshine' remix, for instance, I had plenty of compression on Jeff's vocal (a chain involving a fast limiter, a slower compressor, and a fast parallel compressor), but while that certainly helped give the performance a subjectively assertive attitude, it was actually my detailed fader automation that consolidated its place in the mix balance and ensured that every word cut through clearly.

Online Resources: Audio Clips, Videos, Mix Project & Multitracks

As usual, I've provided a selection of audio examples to accompany this article (including the full-length 'before' and 'after' mix versions) on the *Sound On Sound* website at https://sosm.ag/mix-rescue-0323. In addition, though, I've created a special resources page at www.cambridge-mt.com/jeffhirata where I've posted several demonstration videos including:

- Track-by-track deconstructions of my remix's vocal and guitar arrangements.
- A detailed explanation of my synth sub-bass setup.
- A full mix playthrough highlighting my use of fader automation on the different tracks.

You'll also find download links for my complete Reaper remix project (with a folder of screenshots so you can scrutinise my settings in detail) and for the song's raw multitrack files in case you fancy trying a remix of your own!

Automation is also crucial for breathing life and humanity into the mix balance as a whole. You see, one of the things that makes a production seem engaging is if listeners are constantly alerted to new and interesting facets of the musical material. By simply going through your arrangement and turning up the most interesting moments, you'll actually make the music itself seem better, as well as generating the illusion that all the musical performers are communicating with each other organically — even if everything was actually programmed and overdubbed one track at a time! As such, it should be no surprise that this kind of automation played an enormous role in my remix of 'Sunshine', where I did detailed fader rides across at least 20 tracks.

Widescreen Imaging

Although Jeff highlighted EQ and dynamics processing as potential

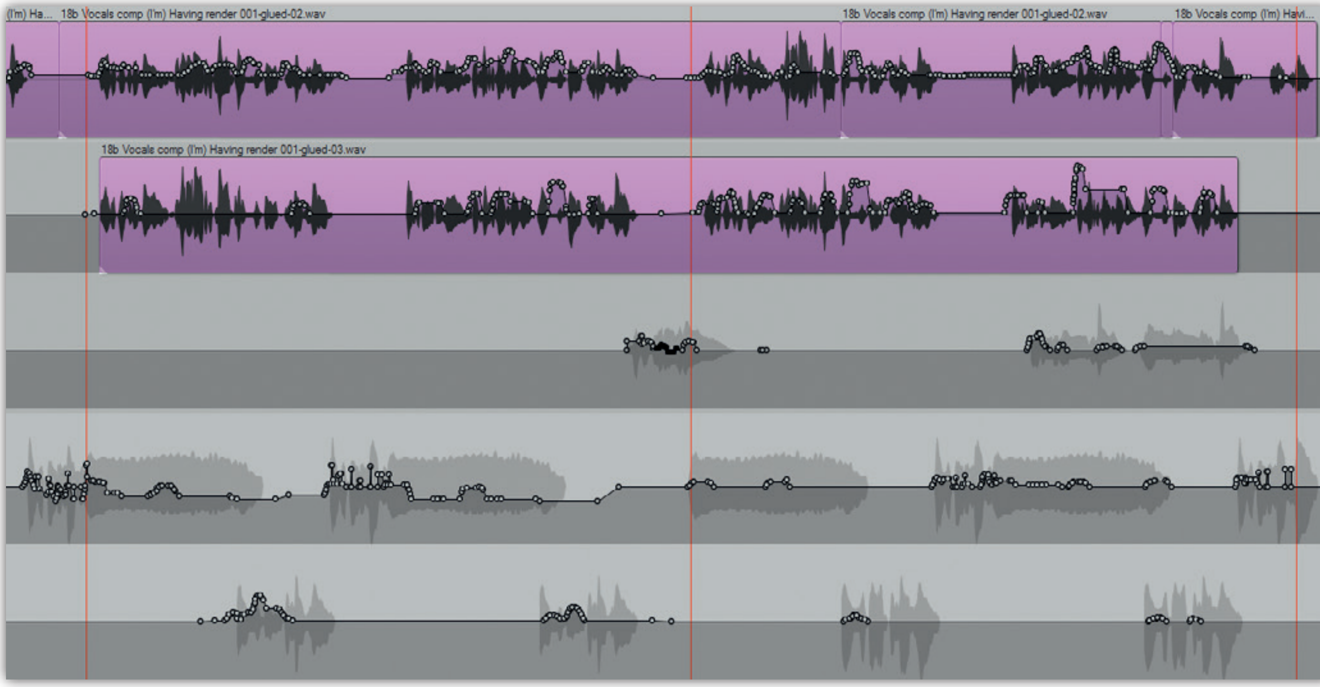

Although a good deal of compression was used to add energy to this remix, it was detailed automation that was responsible for solidifying the final mix balance.

sources of mix hype, the concept of stereo widening didn't really make it onto the radar as much, beyond double-tracking the guitar and vocal parts to allow for opposition panning. Yet, a broad panorama is one of the most reliable ways to make almost any mix feel more vibrant and impressive. Not that I'd recommend just slapping some kind of stereo doohickey over the master bus, because widening treatments tend to work best, in my opinion, when you carefully select the specific tracks you're going to widen and use different widening tactics for different instruments.

So in my remix, for example, I avoided extreme widening of any of the production's most important musical elements (kick, snare, bass, rhythm guitars, lead vocals), so as to avoid compromising their mono compatibility. For percussive parts, I decided to use simple Mid-Sides processing or EQ-based widening to avoid any of the undesirable rhythmic flamming associated with delay-based patches, whereas on vocals I was happy to benefit from the gentle thickening side-effects of pitch-shifted micro-delays.

Laws Of Layering

Another method Jeff had used to power up his mix was layering, and by the time he sent his multitracks to me they were swimming in rhythm-guitar and backing-vocal tracks: more than 50 tracks in total! There's no question that layering's a widespread technique for mainstream pop-rock songs like this so his heart was in the right place, but you can easily have too much of a good thing in this respect.

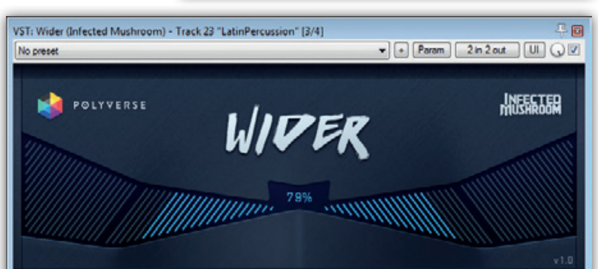

For stereo widening, Mike used a combination of traditional M-S-based processing from Voxengo's freeware MSED and EQ-based widening from Infected Mushroom's freeware Polyverse Wider.

With electric guitars, an ever-present concern is that layering multiple recordings of the same guitar sound can quickly smooth out the crunch and character of each individual part into a chorusey 'background blancmange'. This is why I deliberately avoided layering more than two tracks of any given guitar sound together in my remix, and even in those sections where two or three pairs of electric guitars were playing together, I made a point of differentiating their timbres (in one case reamping one of Jeff's supplied DI parts with my own amp simulator plug-in) in order to retain some clarity and definition. I followed the same tack with the backing vocals, jettisoning anything beyond a double-track for each harmony line. I then also supplemented Jeff's recordings (all of which featured his own voice) with new layers that my daughter and I recorded, thereby expanding the subjective size of the ensemble more convincingly.

No matter how epic your layered texture, however, there's a limit to how long it'll impress the listener. For better or worse, we humans lose interest in that kind of thing quite quickly! This is one of the reasons, in fact, why less experienced recording musicians tend to record too many layers, because the texture always seems to sound best just after each new part has been added, but that psychological buzz naturally fades with time. So, from a mixing perspective, it's wise to vary the layering for any repeated song sections (eg. the choruses) in order to give each iteration enough novelty to maintain the listener's interest. And this process offers another golden opportunity to ratchet up the mix's energy levels by muting/rebalancing the layers to enhance the feeling of build-up. Here's how I did this in my remix, for instance, by modifying the guitar layering in each successive chorus:

- **Chorus 1:** A mono muted rhythm part on acoustic guitar, and four-layer stereo electric guitar accents.
- **Chorus 2:** The acoustic guitar part changes to two-layer stereo, and the electric guitars are EQ'ed for more warmth.
- **Chorus 3:** Two-layer acoustic guitars and four-layer electric guitars move to rhythmic strumming, with support from two-layer electric-guitar accents.
- **Chorus 4** (breakdown following the climactic middle section)**:** Four-layer acoustic guitars and six-layer electric guitars all play one strummed chord per bar.
- **Chorus 5:** The bass is muted, the muted rhythm part returns on stereo two-layer acoustic guitar, and supporting accents are provided by two layers of acoustic guitar and four layers of electric guitar.
- **Chorus 6:** The bass returns, with four-layer acoustic guitars and four-layer electric guitars moving to consistent strumming while a further two layers of electric guitar provide accents.

Not only does this arrangement avoid the listener ever having to experience the same chorus sound twice, but it also helps create a powerful momentum towards the middle section in the first instance, and then towards the final choruses after that. And if you look at the screenshot of my vocal arrangement, you can see a similar trajectory there, with the chorus texture progressively thickening in two 'intensity ramps' (just like the guitar arrangement does), as well as providing an extra lift for chorus seven, given that the guitar texture has already maxed out by then. Now, I realise it can be difficult to imagine how this all sounds from just the text and pictures here, so I've created some special demonstration videos at www.cambridge-mt.com/jeffhirata in which I deconstruct my remix's guitar and vocal arrangements in detail to show how each layer sounds, and how all the layers fit together in conjunction with the lead vocal and its changing delay/reverb effects.

There are many other ways of adjusting the arrangement to energise a mix, however. The easiest is just to flex your mute buttons so that the listener gets a better opportunity to appreciate the qualities of each track. So, in my remix, I started with just acoustic guitar and lead vocal, whereas Jeff's original mix already featured snare drum and steel pans right at the outset. Your mute buttons can also add drama by suddenly refocusing the listener's attention on something new, for example where I dropped the drums out abruptly at the end of the first chorus to highlight the re-entry of the chorus vocal, acoustic guitar, and shaker.

The 'Wow' Factor

Taken together, these tricks were already doing a pretty good job of presenting Jeff's recorded material in a much more engaging light, but it was apparent from our initial discussions that he'd heard some of my more extensive Mix Rescue transformations, and was keen for me to push into more creative territory if I felt it might elevate things further. With this in mind I found myself pondering what to »

Timing Edits

With any upbeat song like this, tightening the timing and tuning can have a powerful impact on how cheerful and bouncy the mix sounds. For this project, it was the timing that demanded the most work, and for two reasons. Firstly, although Jeff's playing wasn't particularly wayward, he was playing against programmed drums, whose rhythmic implacability tends to throw even small groove discrepancies into higher relief. And, secondly, many of the rhythm parts were layered, so even minor rhythmic disagreements between layers were blurring the layered note onsets and reducing the punchiness of the combined texture. As usual, my first process involved conforming the bass with the drums while all the guitars were muted, and then reintroducing the guitars layer by layer to refine their timing against the bass/drum foundation by ear.

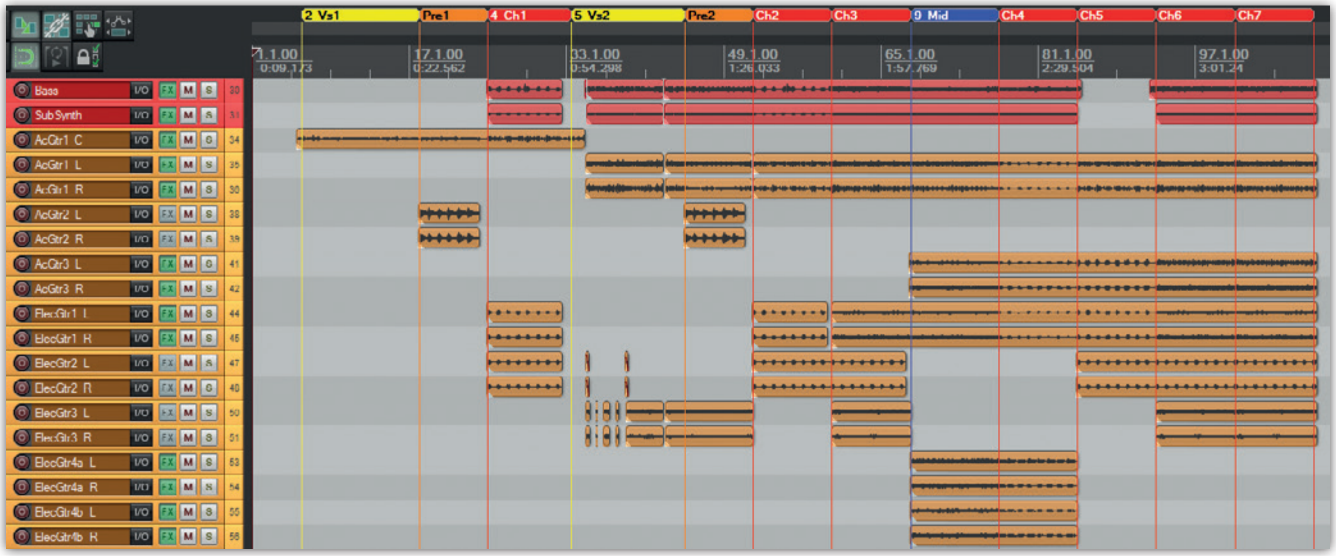

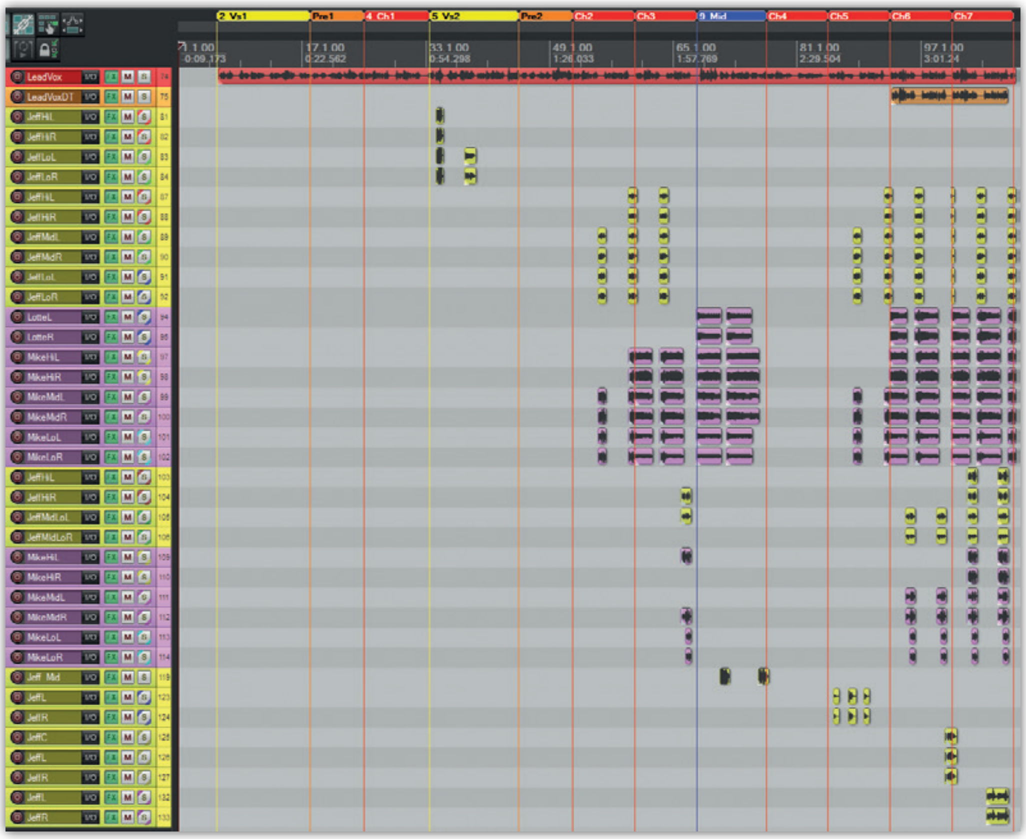

These two screenshots show the final acoustic/electric guitar and backing-vocal arrangements from Mike's remix of 'Sunshine'. As you can see, the arrangement is designed to build up in two 'ramps', the first culminating in the middle section, while the second culminates in the final choruses. Notice also that, as part of this, no verse or chorus has exactly the same arrangement as another, which helps maintain the listener's interest throughout the production's timeline.

» do about the middle section. The problem was that the music seemed to demand some kind of 'step up' in the arrangement energy from the preceding chorus section, but there were no new musical parts at that point from which to conjure this extra impetus.

So my thoughts turned towards the idea of recording something new myself. This is where you can quickly get into trouble as a third-party mix engineer if you've not adequately established ground rules for such experimentation — I'd suggest the maxim "If you hate anything I've added, I'll take it out without question" as a good place to start! But even in a situation like this, where the artist has actively solicited creative input, it's important to have a clear rationale for any additions you make. In this case, my thought process went something like this:

- I want a part that sits in the upper pitch registers, because the low end and low midrange are already saturated with the bass and guitars in that section of the song.
- I should probably avoid anything that's too noisy or rhythmic, otherwise it might interfere with the lead vocal's intelligibility at a point in the song where Jeff's singing fresh lyrics.
- The guitar parts become a little more sustained during the middle section compared with the preceding chorus, so I'd like to add something that emphasises that contrast.
- It'd also be nice to add something with a bit of 'wow' factor to help Jeff fall back in love with his own song after so much frustration during his own initial mix process.

If I'd only had to address the first three of these considerations, I'd probably have fallen back on that old stalwart, the Hammond organ pad — just tweak the

Remix Reaction

Jeff Hirata: “Wow, Mike, you did an absolutely incredible job! When I first heard the mix I was honestly shocked to hear how amazingly happy, bright, and exciting the song had become. It’s a dream come true — I can finally turn the song all the way up in my car to rock out, and also turn it all the way down and still hear everything clearly. The main vocals are the most important part of the mix and are done to perfection, but my mind was absolutely blown when I heard how you’d used the background vocals to uplift the song to a whole new level! They add a richness and complexity that a professional mix should have, while also introducing the sense of lightness and fun that the song was so desperately missing. The bridge used to be my least favourite part of the song, but with the backing vocals it’s now something I look forward to! You did hundreds of other little things that also added up, like removing the snare drum from the beginning, and more slowly building up the drums. It’s one of those things you hear and wonder why you never thought of it yourself, since it made so much sense after hearing it!”

drawbars to funnel the frequencies into your mix’s spectral pockets, add some tasty swirl by automating the Leslie speaker’s speed switch, and Bob’s your uncle! But that didn’t feel like it would be a dramatic enough move to really blow the lid off this particular section, so I resolved to take a riskier route.

Taking my inspiration from one of the songs Jeff and I had referenced during our initial discussions (the Rembrandts’ ‘I’ll Be There For You’), I returned to the vocal microphone and layered up some new sustained four-part backing-vocal harmonies based on the word “Sunshine”. (Never a bad idea to reiterate the title hook!) Then, realising that this same sustained harmony-vocal sound also seemed to work quite well for the second and third groups of choruses, I added some similar layers there too, while being mindful of the need to weed out a few layers earlier in the song, so that the arrival of the middle section still made enough of an impact.

Harness The Energy

There are a lot of different tools that you can use to boost the sense of energy in your production: EQ, compression, automation, stereo widening, editing, layering... but each of those things can also work against you if you’re not careful! I hope this real-world case-study has provided some pointers for getting the best out of them in practice.

We show you how to get started recording audio in Studio One.

ROBIN VINCENT

In recent workshops, we've got down to the basics of recording, editing and working with MIDI. Now it's time to turn our attention to audio. Recording sound is a primary function of any DAW, and I think it's good to go back to the original concepts. You might find that you've been doing it wrong all this time, or perhaps can find more efficient ways of achieving your aims. Studio One continues to evolve, so the tools and processes we use all the time can change. You might even find something you didn't know existed that will improve your workflow.

From the Studio One welcome page, you can do a little bit of configuration on your audio interface. Well, I say configuration, but all you can really do is choose the interface and set the buffer size (or Device Block Size as Studio One likes to call it). Sample rate settings and input/output configurations are managed elsewhere, and you can't get to them at this point. So, first make sure your audio interface is selected.

For audio recording, it's likely that we're going to want to monitor our audio through the Studio One mixer, and perhaps through its plug-in effects. This requires a low input and output latency, which means a low Block Size setting. You are aiming for a round-trip latency of less than 10ms or so. The 'round trip' latency is the input latency plus the output latency. When you plug in your guitar, the audio interface buffers the signal coming in; that causes a short delay, which is the input latency. It then goes through Studio One and maybe through the Ampire plug-in and back to the audio interface, where it's buffered again before going to your speakers; that's output latency. Those two things (the 'round trip') contribute to the total amount of latency you experience. Anything more than 10ms and you'll begin to hear the delay to the audio when monitoring on headphones.

Lowering the latency will make monitoring feel snappier, but it also puts pressure on your computer's processor. A Device Block Size of 128 or 64 samples should be a good compromise, allowing you to hear what you're playing without distracting delay, and without hammering the CPU. Alternatively, we can use Direct Monitoring to bypass the Studio One mixer and avoid the buffering latency, and we'll talk about that later.

Templates

Now it's time to create a song and get recording. Click on New at the top of the welcome page, and you'll be presented with a list of helpful templates. I used to ignore these things, but for version 6 PreSonus have tidied them up, and I'm starting to appreciate the help. For this workshop, we're going to select the Record Now template for a new audio recording. After you've clicked on it, more options will appear on the right side of the window. We can set it up as a single track, two tracks for vocal and guitar, or a full band.

The Guitar/Vocal and Full Band templates are very interesting. They set up multiple channels for recording, and add a selection of effects and dynamics processing. There's also an Impact drum machine loaded for the Guitar template, and for the full band you get seven channels set up for drums, all squirrelled away in a folder track. We'll deal with those in more detail in another workshop, as there's much to go through and probably lots to learn about setting up tracks, inserts and routing. For now, let's keep things simple and opt for a Single Track.

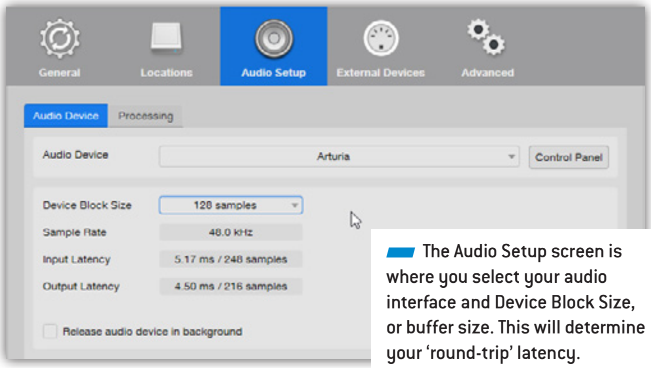

The Audio Setup screen is where you select your audio interface and Device Block Size, or buffer size. This will determine your 'round-trip' latency.

I should point out that the template doesn't default to the sample rate set by your audio interface, so you should check and ensure it's set correctly. You may have other systems in your computer that are using the audio engine, and swapping sample rates unexpectedly is rarely good.

Once we're in Studio One we get treated to the new internal tutorial system. It introduces the template and then gives you a tour over the course of five or six screens. While it doesn't let you do anything while it's running, it highlights exactly what you need to do to start recording and is spookily similar to what we're talking about.

Audio Tracks

The template gives us our one track ready for recording. It also loads an Input Channel and a channel strip for the track into the mixer console. The inspector isn't open and the layout is nice and clean, so we can focus on dealing with that audio track.

We first need to check that the input and output options in Studio One correspond correctly to the actual inputs and outputs on our interface. Click on the down arrow next to where it says Input 1 on the audio track header, and select Audio I/O Setup. In this window, we can see all the physical inputs that our audio interface presents, listed on the horizontal axis of a grid from left to right. The vertical axis displays all the software inputs that Studio One has created, and the intersections within the grid show where the two match up.

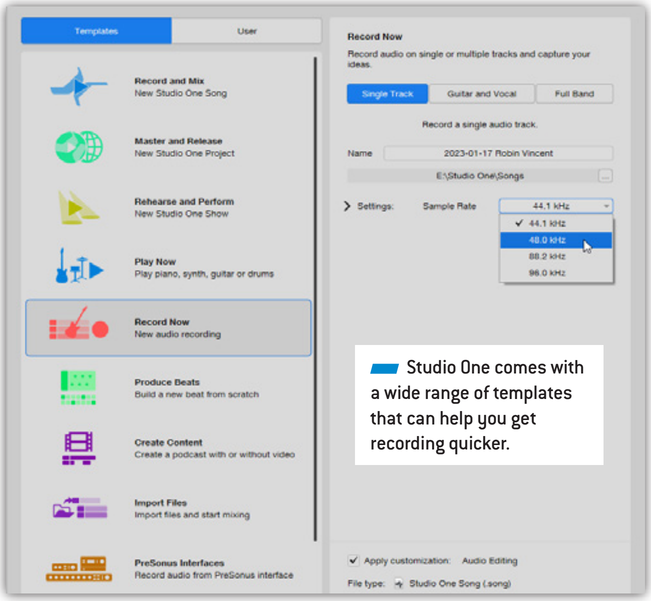

Studio One comes with a wide range of templates that can help you get recording quicker.

The Audio I/O Setup window lets you assign your audio interface's physical inputs and outputs for use in Studio One.

By default, Studio One will create a couple of inputs and assign them to the first couple of physical inputs. These may not be the ones you intend to use, but you can reassign them, rename them or create additional Studio One inputs. Simply click in the grid to connect a Studio One input to a physical input. There's also an Output tab where you can do precisely the same thing with the output configuration. This means that when you're assigning inputs to tracks in Studio One, you don't have to browse through all the physical inputs of your interface directly. Instead, you select 'virtual' inputs you created, named and assigned to those physical inputs.

In the screenshot, I've called the first Studio One input Vocal, and I've assigned it to the first physical input on my Arturia Audiofuse, where I have a mic plugged in. I've plugged my guitar into the second input, and named the corresponding Studio One input Guitar. The output from my synth rig goes into inputs 3 and 4 on the interface, and so I've created a stereo input for those.

Monitoring

Back to our track in the arrange window, and next to where it says Track 1 are four buttons. M (Mute) and S (Solo) should be familiar; we're not interested in those right now. The other two are Record Enable and Monitor. If we want to record onto that track through the selected input, Record Enable needs to be turned on. That's easy, right? The Monitor button determines how we hear our input signal. We can either monitor what we're playing or singing through Studio One, or do it outside of Studio One.

Let me explain. Normally speaking, when recording an instrument or your voice, you want to be able to hear what you're doing. When using microphones, you'll need to listen using headphones so you don't get feedback and so that none of the playback music gets picked up by the microphone. With your headphones on, it's difficult to hear your own singing or the twanging of your guitar unless that signal is returned back to your headphones. This is called foldback, or live input monitoring.

If we monitor through Studio One, by enabling the Monitor button, we'll be able to use effects that we can hear as we sing or play. You could, for example, add reverb to a vocal, or use Ampire on your guitar. There's nothing to stop you from adding or changing these effects later, but hearing them while recording can be very useful. The downside is that you will encounter the full round-trip latency of your system. If that's too large, your monitoring will sound weird and laggy, which won't help your performance.

The alternative is to use Direct Monitoring. This routes the input directly to the output within the audio interface, thus bypassing buffer latency. But as what you're hearing isn't passing through Studio One, it isn't going through any plug-in effects you want to use either. Direct Monitoring is probably the best choice if you are familiar with running external gear or mixers and generally handling your audio outside the computer prior to recording.

Input Level

Once you've sorted out how you will hear yourself, you need to get the right level into Studio One for recording. Conventional wisdom says that you should aim to have the highest peaks reach roughly -12dBFS. FS stands for 'full scale', so this is 12dB below the largest value the system can accommodate. This gives enough headroom for incidental bursts of performance energy, while keeping the signal far enough above the noise floor to ensure that hiss won't be a problem.

Your audio interface may provide some metering, but it's usually quite rudimentary. Fortunately, Studio One has some lovely fat Input Channel meters that will help you to hit the right level. In the template, the Input Channels are already open in the console. Regardless of your monitoring settings, these will show the incoming level.

While singing or playing, adjust the gain controls on your audio interface so that the loudest parts reach the -12dB mark, which is the first line in the upper quarter of the meter. The Input Channel also has a gain control, which ranges from -24dB or +24dB. This can be useful for adding further gain if your mic preamp can't deliver enough, but it can't negate the effects of recording at too high a level and clipping the A-D converter.

Here, Track 1 (the vocal track) is in Record Enable mode (the red button below Mute), and has input monitoring enabled (blue button). The large meter on the Input Channel at the left shows that the level is somewhere around the -12dB mark.

Input Effects

You may have noticed that the Input Channel also has insert points. These give you access to all the same insert presets, dynamics processing and VST effects that are available on other channel types, but the key difference is that any effects inserted here will be printed onto the recorded track. So if you load up Ampire on the guitar Input Channel, the recorded guitar track will capture the output of Ampire. It would be like recording through an external effects unit before the signal reaches the interface.

Whether this is useful depends on how you like to work. One of the key things DAWs did was provide us with a safety net by making effects non-destructive, as they are when used on other channels, so that we could change them without having to re-record. But sometimes committing to a sound is helpful!

When you're ready to go, make sure the Record Enable button is lit and hit Record on the transport bar. We'll talk about what happens after that in the next workshop...

JULIAN RODGERS

What does 2023 have in store for Pro Tools users?

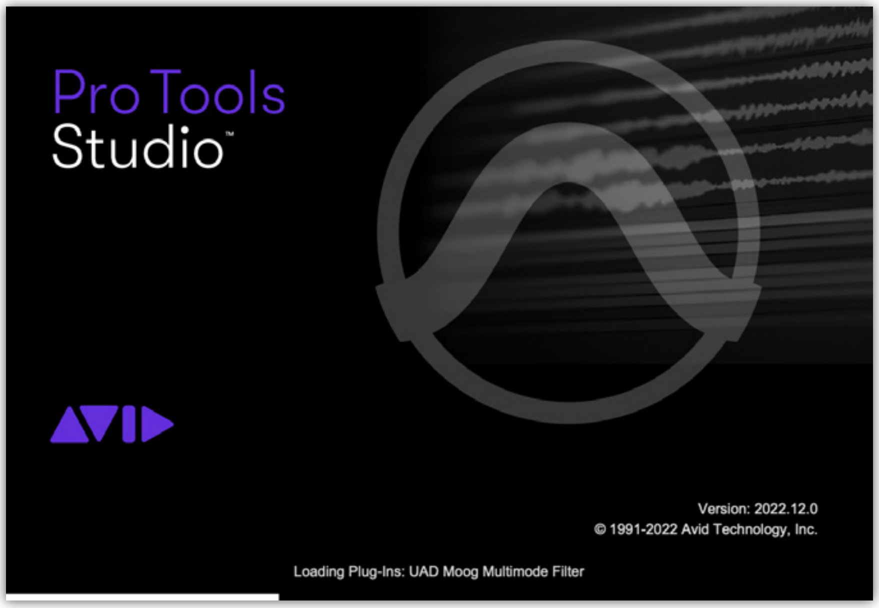

Pro Tools Studio bridges the gap between the somewhat limited Artist edition and the costly, HDX-oriented Pro Tools Ultimate.

With 2022 behind us, now seems a good time to reflect on what was an extremely positive year for Pro Tools, and also to consider what might be next.

Last year saw an overhaul of the line-up, which introduced the free Pro Tools Intro (a timely replacement for the discontinued Pro Tools First), and a new entry-level version: Pro Tools Artist.

This is a big shake-up. A familiar grumble among Pro Tools users was that the standard Pro Tools product lacked features important to serious users, making Pro Tools Ultimate the only viable option for many. The new product range brings options appropriate to far more users.

Pro Tools Intro is everything Pro Tools First should have been: a free version of Pro Tools that shares the same codebase as the rest of the family, works with the same Session file format and can be used with third-party AAX plug-ins. Pro Tools Artist, meanwhile, is accessibly priced, and with 32 audio tracks is more than capable of handling typical music productions. However it is the mid-tier Pro Tools Studio which is the stand-out product for many.

If you are a serious music user who doesn't use HDX, Pro Tools Studio is the product you've been waiting for. In years gone by, the distinction between LE and HD versions of Pro Tools was stable and reasonably well understood. If you wanted sufficient inputs to track a full band, enough tracks for big sessions, VCA tracks, advanced automation and surround mixing, you needed HD. Many of the most frustrating limitations were gradually lifted in the native Pro Tools Software which replaced LE, with an increase in the number of simultaneous inputs, and the availability of VCA tracks, but advanced automation features like Preview and Capture were a sticking point for serious mixers working in stereo. And if you wanted surround, you were still compelled to use Ultimate.

Replacing the old Pro Tools Software, Pro Tools Studio has surround and Dolby Atmos capabilities, and advanced automation. Unless you need DigiLink connectivity for HDX hardware, or are working in post-production, it will probably have everything you need, and is much more affordable than Pro Tools Ultimate.

While the top-tier Ultimate product hasn't changed, Avid did reverse their initial decision to provide Ultimate exclusively as part of the Flex package, which was aimed at large-scale facilities. This had the unintended consequence of overlooking users who needed Ultimate but fell outside this target group. By reversing this change and reducing the pricing of Ultimate at the end of 2022, Avid showed that they are listening to and responding to the needs of users.

New Features

Reorganising the product family is one thing, but there were also significant new features introduced in 2022, two of which addressed longstanding issues. The first was Aux I/O. The way the Pro Tools Playback Engine works has always presented issues for users looking for more flexibility, and even a task as apparently simple as connecting a USB microphone could be surprisingly difficult. Users of HDX systems have long had to make special arrangements to monitor system audio, and people working in Dolby Atmos have had to change to a dedicated playback engine, or use MADI or Dante to send their Atmos-format audio to a second system. Aux I/O brings routing flexibility between different audio interfaces on the same system.

Another long-overdue feature was ARA support. While at the moment only Celemony's Melodyne has received the ARA treatment with Pro Tools, this is expected to change.

The other big feature that's been missing from Pro Tools was native Apple Silicon support. Pro Tools has worked well on Apple Silicon Macs under Rosetta 2, but as part of the 2022.12 release a Silicon-native public beta was made available for all versions, including Intro users. Testing shows the anticipated leap in performance, and a fully qualified version is in the works.

What's Next?

So, 2022 brought many of the things we'd been asking for as Pro Tools users. The question is whether this trend will continue, and if it does, what might we see first? Entering the realm of speculation and future-gazing is always risky but here are some predictions: some safe bets, some a little more speculative.

One feature that's already been announced (and might even be available by the time you read this) is fully qualified Apple Silicon native support. With third-party plug-in support for Silicon growing rapidly, the transition can't come soon enough for owners of current Macs. Ventura support is already with us, and the speed with which Avid have achieved this (they haven't always been the fastest on this front) suggests that new operating systems might be supported more promptly in future.

Further ARA support is something we can also be reasonably confident about. Celemony were instrumental in the development of ARA in the first place (the protocol was a collaboration between Celemony and PreSonus), so it's understandable that Melodyne was the first ARA product to come to Pro Tools, but it's not unreasonable to predict that other applications will be integrated soon. The post-production community will be waiting for iZotope RX, but I'm personally waiting for Synchro Arts' ReVoice Pro and the related VocAlign and RePitch plug-ins to come over to Pro Tools ARA. You can get by without ARA, but once you've experienced it, it's hard to go back!

Something which was announced as part of the 2022.12 release was the new Scripting Software Development Kit, which Avid have now made available. This won't be relevant to people without proper coding chops, but developers will be able to use the API calls made available by Avid to develop tools that can automate tasks in Pro Tools, and as more and more API calls are released, the potential of this deep access will grow. The SDK is intended for use either by large facilities who wish to create scripts for use in-house, or for wider distribution. However it is used, it's looking like the beginning of something interesting.

Avid also doubled down on their return to making hardware interfaces in 2022 with the MBox Studio, and having had a unit for some months, I can confirm that it's a more than worthy successor to the MBoxes of old. At present, however, there's one MBox Studio feature that isn't yet fully realised: the set of four user-assignable buttons in the centre of the top panel. As it stands, these can be assigned to control up to eight actions within the MBox Control software. The plan for a future release of Pro Tools is that these user-definable buttons will be able to control Pro Tools itself. Exactly how deep this control will go is yet to be revealed, but transport control would be very welcome, particularly Back and Play, which are invaluable for automation passes. Toggling Groups, changing automation modes... the potential uses are endless, and if you want to try software control using hardware buttons but aren't ready for a EuCon surface, this is another reason to consider this excellent new interface.

The recent MBox Studio release suggests Avid intend to keep making hardware for some time to come.

There are a wealth of as-yet unrealised items on my wish list for Pro Tools itself. One small fix would be to have folders in the clips list (having used Avid's Media Composer, I've always wanted MC's system of media Bins to be available in Pro Tools, rather than the flat list we're currently stuck with). I also find the inability to change Playback Engine without restarting Pro Tools infuriating, and the restriction of only being able to have one Session open at a time is something else I'd love to see change. Proper support for 4K monitors is long overdue. And with the direct integration of Dolby Atmos into Logic Pro and Nuendo, Pro Tools' approach using the standalone Dolby Atmos Renderer looks increasingly out of date.

Returning to hardware, and definitely straying into conjecture, Avid recently announced that the grey-and-black series of HDX interfaces, which includes the HD IO, were being discontinued at the end of 2022. This, combined with the release of the MTRX Studio, MBox Studio, Carbon and Carbon Pre makes me wonder whether more hardware releases will further establish Avid as being very much still in the hardware game. Perhaps they'll even introduce a a replacement for the HDX platform itself? It could be an exciting year...

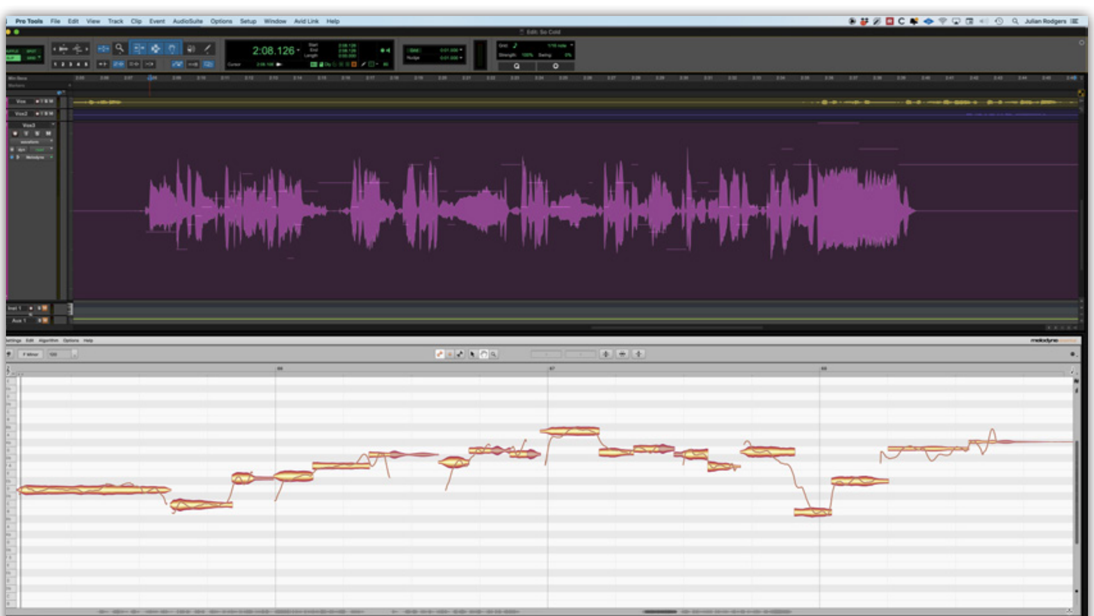

With Pro Tools now boasting ARA support for Melodyne, could other products soon be getting in on the ARA action?

Looking for musical inspiration? You might just find it in Cubase 12's Verve felt piano.

Verve has something of a niche sound set but offers easy sound editing (including hands-on control via the Focus Quick Controls) and some inspiring chord presets.

JOHN WALDEN

For many Cubase Artist and Pro users, a highlight amongst the new toys in Cubase 12 was Verve, a sample-based virtual instrument library for HALion Sonic SE (and, of course, the separately available full versions of HALion Sonic or HALion). Verve is built from a deeply sampled (nearly 9GB of samples) soft-sounding felt piano, alongside a range of textural sounds, and it allows the user to blend these two elements. The sound set provides some inspiring options for ballad or melancholic song styles, but it can work equally well in music-to-picture contexts. Interestingly, Steinberg have also made use of HALion Sonic's integrated chord pad feature, and each of Verve's 70 presets is populated with a set of chords. So if you're looking for some instant musical inspiration, Verve's sounds and the HS chord sets could be a great resource.

Sound Start

The piano sounds great on its own, with a beautifully soft tonality. You also get options of both close and distant mics, adjusting the overall tone, compression, elements of mechanical noise (for added realism) and adding reverb and delay. This is joined by a collection of nearly 70 texture types, spanning various sustained string, synth, metal and electronic sound types, the idea being to make it easy to blend a pad-like element in alongside the piano. And that choice can instantly transform the musical mood! For example (as demonstrated in the accompanying audio examples on the *SOS* website: https://sosm.ag/cubase-0323), take the same simple chord-based performance and swap out an analogue string texture for one of the mysterious-sounding electronic textures (eg. Transmission), and you can quickly go from singer-songwriter ballad territory into something more suitable for some gentle film-score suspense.

It's worth exploiting the Focus Quick Controls (see the *SOS* November 2022 column for a discussion on these as part of the new MIDI Remote feature set) for hands-on control of Verve's main sound-shaping options. This includes the Balance control, to adjust the blend between the piano and texture elements (on FQC 1). Usefully, Verve's FQCs are also preconfigured with the Texture's Variation (in general, higher values seem to result in greater complexity in the sound), Contour (a filter effect), Attack and Release assigned to FQC 3, 4, 5 and 6 respectively. A combination of these lets you coax plenty of additional movement from the texture component of your sound.

The FQC system aside, adjusting Verve's controls to create your own piano/texture blends is fast, effective and completely undaunting, with a minimalist control set spread across the three – Piano, Texture and Effects – pages of the UI. And, yes, while it's certainly a somewhat niche sound set, within that 'melancholic meets mysterious' niche, it does sound fabulous.

Keys To Your New Pad

Verve provides for speedy customisation of your sound, then. But it also offers an equally speedy (and easily overlooked) route to musical creation, courtesy of those presets including a set of chords for HALion Sonic's integrated chord pad system. As each chord, regardless of complexity, can be triggered from a single MIDI key, it becomes very easy to experiment with some harmonically tasty sequences. You can,

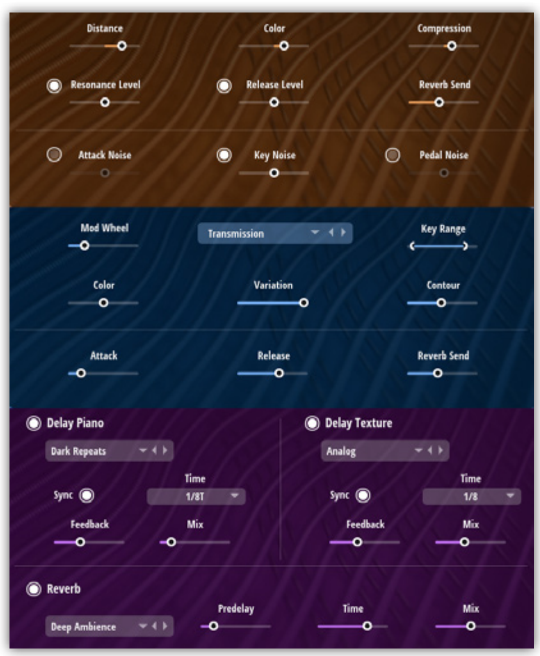

The control sets for each of Verve's three main pages (stacked into a single screenshot here for convenience) make it easy to customise your sound.

The included texture options span the organic to the other-worldly.

if you like, mouse click on the pads to audition the chords, but in order to trigger them properly you'll need to assign each pad to a MIDI key. To do this, simply right-click on each pad in turn and, on the pop-up menu choose the Assign Trigger Note option and select the desired MIDI note. Notes assigned to a chord pad in this fashion turn blue on the mini-keyboard display at the bottom of the GUI. Incidentally, once you've assigned all eight pads, it's worth popping that menu open again and using the Save Trigger Notes As Default option to save your settings for easy recall.

Once you've made assignments in HSSE you can (just as with the main Cubase Chord Pad system) then play and record a MIDI sequence of these trigger notes, building a full chord sequence from the preset chord voicings — even if your piano keyboard skills are somewhat limited!

Good Pad, Bad Pad

So, inspired by Verve, and driven by the HSSE chord pad presets, your next song or cue idea is ready to be created. There are plenty of musical opportunities to explore, but before we do I want to run through just a few additional details — some good, some not so good, but all worth noting if you want to build this Verve/HSSE combination into your musical workflow. On the positive side, all of the Verve HSSE Chord Pad presets are available independently of the main Verve preset system. Simply clicking on the small downward arrowhead icon to the left of the actual chord pads will open a chord pad preset browser. So it's easy to experiment with the different chord set combinations without having to change the overall Verve preset.

Rather oddly, while HSSE's chord pads and Cubase's own Chord Pad system share plenty of functional similarities, they don't seem able to share chord set presets! That's a missed opportunity, really, as it would be great to be able to exchange the underlying sets of chords between the two systems, and I hope Steinberg will address this at some point because it would make for a much easier workflow when you want to trigger the same chord sets on different virtual instruments. This wouldn't be an issue if HSSE offered MIDI out, of course, as you could pass the full chord parts (and not just the trigger notes) to a suitable destination (MIDI or Instrument) track outside of HSSE. But as far as I'm aware, that's currently not possible either (please feel free to enlighten me if I have missed something here!).

Of course, workarounds do exist. Recreating, within the main Cubase Chord Pad system, a set of chords contained in a Verve chord pad preset is easy enough (check the May 2015 column for a quick Chord Pad refresher). And if you want to match the exact voicing of each chord, the mini-keyboard display in HSSE can serve as a guide, allowing you to use the manual chord assignment process on the main Chord Pad system to replicate the exact combination of notes.

Alternatively, you could give Cubase 12's Audio To MIDI Chords function (covered in *SOS* July 2022) a spin; while this requires a few steps it can be achieved fairly swiftly. First,

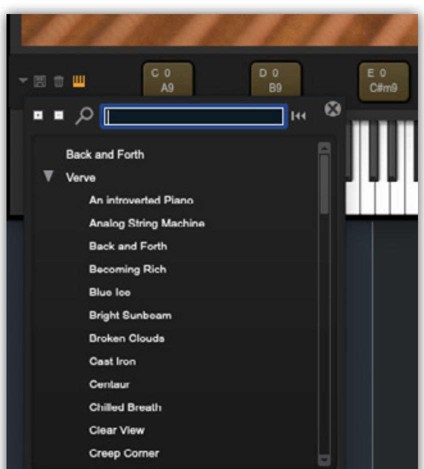

Usefully, the chord set presets provided in Verve can be loaded independently of the sound presets.

create a simple MIDI clip on your Verve track, containing trigger notes for each of the eight chords, with a short silence between each chord. Second, toggle off the Texture and Effects layers so that, on playback, you can only hear the Piano layer. Third, select the MIDI clip and then execute the Edit/Render In Place command. This will generate an audio clip from the same piano chord performance. Fourth, drag and drop this audio clip onto the Chord Track. This triggers the Audio To MIDI detection process and, hopefully, should populate the Chord Track with the same chords (or something very similar in terms of harmonic content) as in the original Verve chord pad preset.

Some tweaking of Audio To MIDI's 'best guess' chords might be required, but if my own experience (based upon a solo piano audio source) is anything to go on, Audio To MIDI will generally do a pretty good job. Finally, once your Chord Track entries are suitable matches for your Verve chords, the Functions menu of the Chord Pad's tab (in the Project window's Lower Zone) includes the option to Assign Pads From Chord Track; any chord events on the Chord Track will then be assigned to an empty chord pad (duplicate chords are ignored) and can then be used with any virtual instrument.

Added Verve

So, where can this take you? Well, hopefully, the audio examples that accompany this workshop hint at the sorts of musical applications Verve might suit. To my mind, it's so good that I hope Steinberg continue to develop the feature set! With my greedy hat on (and aside from addressing the chord pad technicalities mentioned above), I can think of a couple of additional features that I'd love to see. First, full amplifier ADSR options for both the Piano and Texture layers could be really useful. Second, with my music-for-video hat on and inspired by having used UAD's Ravel Grand Piano, it would be very cool to see a 'reverse the piano samples' feature, which could really increase the 'spooky and creepy' potential for media composers.

But even as it stands, Verve is a wonderful addition to the Cubase Pro/Artist virtual instrument collection. It combines great sounds, ease of use, and a very creative collection of chord sets to get the compositional muse moving. So why not see if you can add some Verve to your own music making?

Logic offers a number of ways to repeat your notes and phrases.

DAVID RICARD

Since repetition is a major component of music, it makes sense that it would make up a large portion of what goes into creating it. Let's look at some of the ways in which we handle repetition in Logic Pro.

One of the most essential functions we require in composing is the ability to repeat regions in our arrangements. By simply clicking on a region (or group of regions) and hitting Command+R, you can create a duplicate. Keep hitting this combination to continue making additional copies. If you know that you want an exact number of repeats, you can right-click on the selection, go to Edit and choose Repeat Regions / Events Multiple Times... From there, type in the number of copies you need.

As you would expect, this behaviour also works on notes in the Piano Roll window as well as in the Notation Editor.

Copy That

I'm sure you all know how to copy and paste regions, but I've been surprised to find that many Logic Pro users rely on only one method of doing so, when some of the alternatives might suit them better in certain situations.

Select any region in the arrangement window and hit Command+C. Now move the playhead to wherever you want the region copied to and hit Command+V. Basic stuff. Equally useful is clicking a region and Option-dragging it to create a copy.

If you want to copy an entire section of your song with all of the tracks included, the most efficient way is to place the locators inside the measures you want to copy, and go to Edit / Cut/Insert Time / Copy Selection Between Locators. You can now hit Command+V to paste wherever you place the playhead. This function eliminates the need to select all of the tracks and cut them before copying.

Record Repeat

If nailing the perfect take gives you fits of frustration, you should acquaint yourself with the Record Repeat function. Open up your Key Command window (Control+K) and search for it. Assign it a command that you can easily remember and, more importantly,

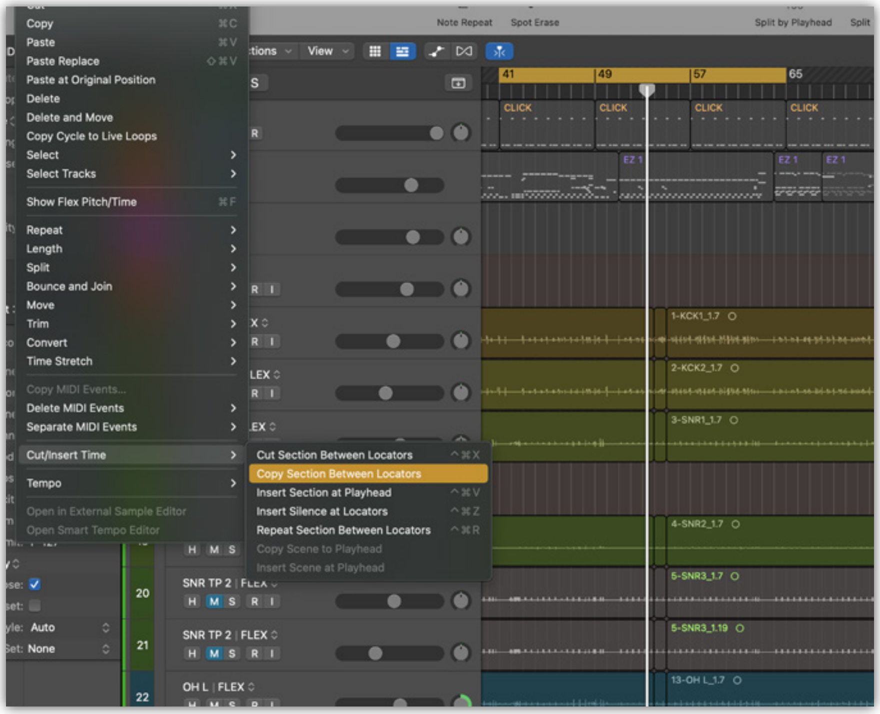

The Copy Section Between Locators function copies all the tracks without cutting the regions, and moves the selection to the clipboard.

type with one hand with minimal contorting. When you are recording yourself, you should consider adopting this as your main recording command.

Here's a scenario that I used to find myself repeating: I record a few bars of guitar into Logic Pro. The take is horrible. I hit the space bar to stop. I can either Command+Z to undo the recording, or I can click on the region and delete it (and have to click OK when the 'delete file?' warning appears) before hitting Record again.

Of course I could just press Record again after hitting the space bar and allow Logic to create a take folder. And while this could make sense if I planned on comping a part together from various takes, I'm usually merely trying to get a somewhat simple part down and would prefer to avoid the extra mess in the Arrange window. Yes, take folders are an invaluable tool; however, one drawback to them is that they are a bit cumbersome to deal with once Flex edits are needed for pitch or timing (which, in my case, they always are!). I'd rather just have one clean version or, at least, get something usable before adding additional takes.

This is where Record Repeat comes to the rescue. Start recording as you normally would (with or without Cycle enabled), but once you determine that you need another attempt, do not press the space bar or anything other than your new command for Record Repeat. Using this function returns the playhead to where you began, erases presumably your previous take, and starts recording again — all with one command! If you do hit the space bar, Record Repeat will still work but it will not delete your previous take.

If the idea of creating a new key command for basic recording is too much of an adjustment to make to your workflow, you have another option. Click and hold on the Record button in the Control Bar (not to be confused with the Toolbar) and select Record / Record Repeat. Now, once you're recording your take, you can just hit the

With Record/Record Repeat selected in the Record button drop-down menu, the button (or key command) now sends the playhead to where the recording starts (including count-in) while deleting the previous take.

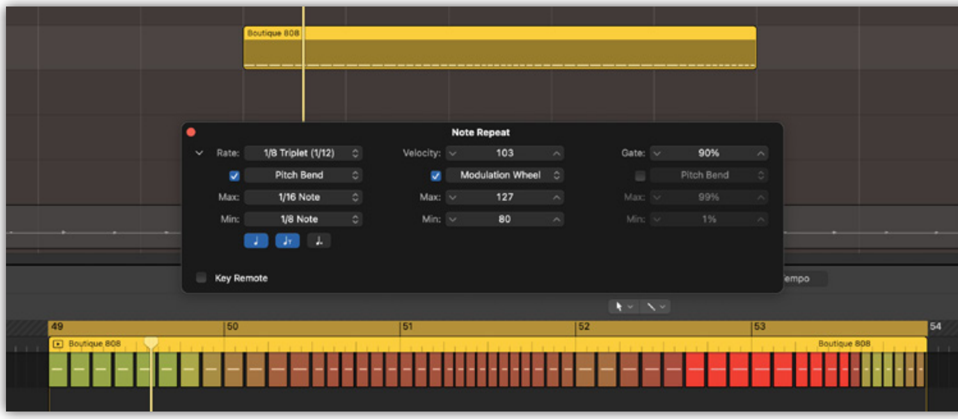

While it may appear that Note Repeat is a type of arpeggiator, it actually generates MIDI notes that get recorded.

on-screen Record button (without stopping) to achieve the same result.

Note Repeat

Whenever Logic has an upgrade and a list of new features is presented, I have a bad habit of completely ignoring anything that I don't anticipate using in the next 20 minutes. When I eventually discover a (new to me) feature, it's already old news. Fortunately, I've come to learn that many Logic users share this dicey habit. With that in mind, let's take a look at the new (since 2015!) feature called Note Repeat.

Note Repeat works as an intermediary between the notes you play and what's recorded into Logic, regardless of what track is selected. Probably the best example of what Note Repeat can do is when programming hi-hats.

With a drum instrument track selected, open up Note Repeat. This can be accomplished by hitting Option+Control+Return or, alternatively, you could access it from the Toolbar.

When you strike any note it will play back in sync at the note value indicated in the Rate field. With the Modulation Wheel box ticked, you can then set up a range of values to scroll between. Employing the note duration buttons gives you more steps along the range. By simply holding down your hi-hat key and moving the mod wheel, you'll quickly get a creative groove going. The mod wheel is an obvious choice for this controller, but I prefer to use pitch-bend as it makes it a bit more precise to dial in a performance, especially if your range is only a few values. If you're feeling especially ambitious, you can set a velocity range and use aftertouch (or any controller) to manage your dynamics.

What makes Note Repeat alluring is that it's not an arpeggiator; it's actually playing the notes for you. When you record your performance, the individual notes are captured and, thus, can be edited just like anything else in Logic. That means that you don't necessarily need to nail the mod wheel timing on the fly. If changing hi-hat rhythms from the mod wheel is a bit counter-intuitive, click the Key Remote box to enable a keyswitching method.

I doubt that Note Repeat is intended to help people record tight rhythms. Between quantising and basic editing, it's easy enough to get a programmed pattern that you're happy with. Note Repeat has more appeal as a live feature. This becomes increasingly evident if you have Logic Remote installed on an iPad. Simply play a hi-hat pad with two fingers and spread your digits to change the rate.

Don't Fear The Repeater

If you're looking for a cross between Note Repeat and an arpeggiator, look no further than the Note Repeater. And while we're at it, let's applaud Logic for coming up with some of the most literal and anodyne names for its features.

While Note Repeat is a function for inputting notes, Note Repeat*er* is a MIDI effect. You load it onto a software instrument track and it does its work on the incoming and recorded MIDI. It's a subtle distinction but worth noting. When you hold down one note with Note Repeat open, it will record a series of notes into Logic. Conversely, when you do the same with Note Repeater, just the one note is recorded.

When you open Note Repeater, it defaults to three repeats, which is actually a great starting point. For every one note you play, Note Repeater will play three (represented by the vertical bars) at the duration specified by the Delay slider. Technically, the first one isn't a repeat because it's the original note. If you wanted to omit the first note from your sequence, you could turn off Thru.

The Velocity Ramp percentage is where you make your echoes fade out (negative values) or ramp up (positive values). Transpose changes the pitch of the repeats. Because it uses the same interval for each repeat, the results (unless you set it 12) can be less useful... or more useful, depending on how weird you are!

Keep in mind that this is MIDI (note) pitch not actual pitch. So if you apply transpose to a hi-hat in a drum kit, Note Repeater will play back other drums in the kit. Surprisingly, this can yield some unexpectedly satisfying results.

At the risk of repeating myself, Logic offers numerous methods for repeating oneself.

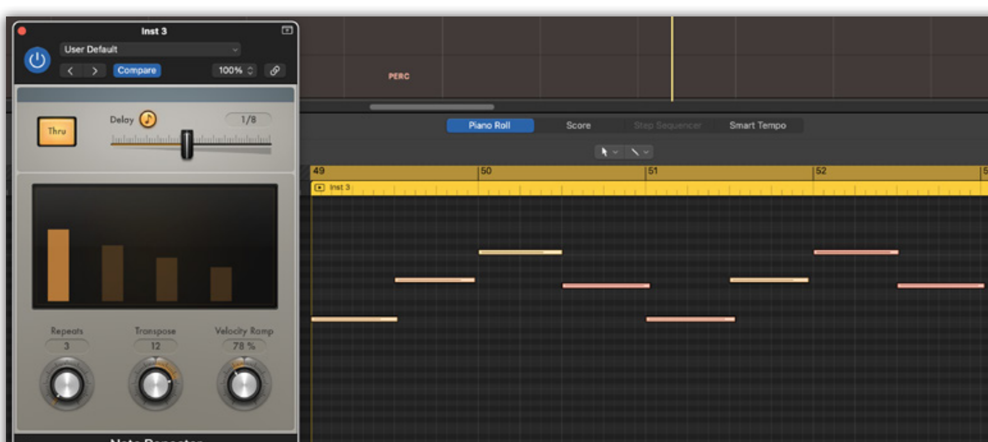

The Note Repeater MIDI effect has a simple interface and can add movement to lines, phrases, and pads.

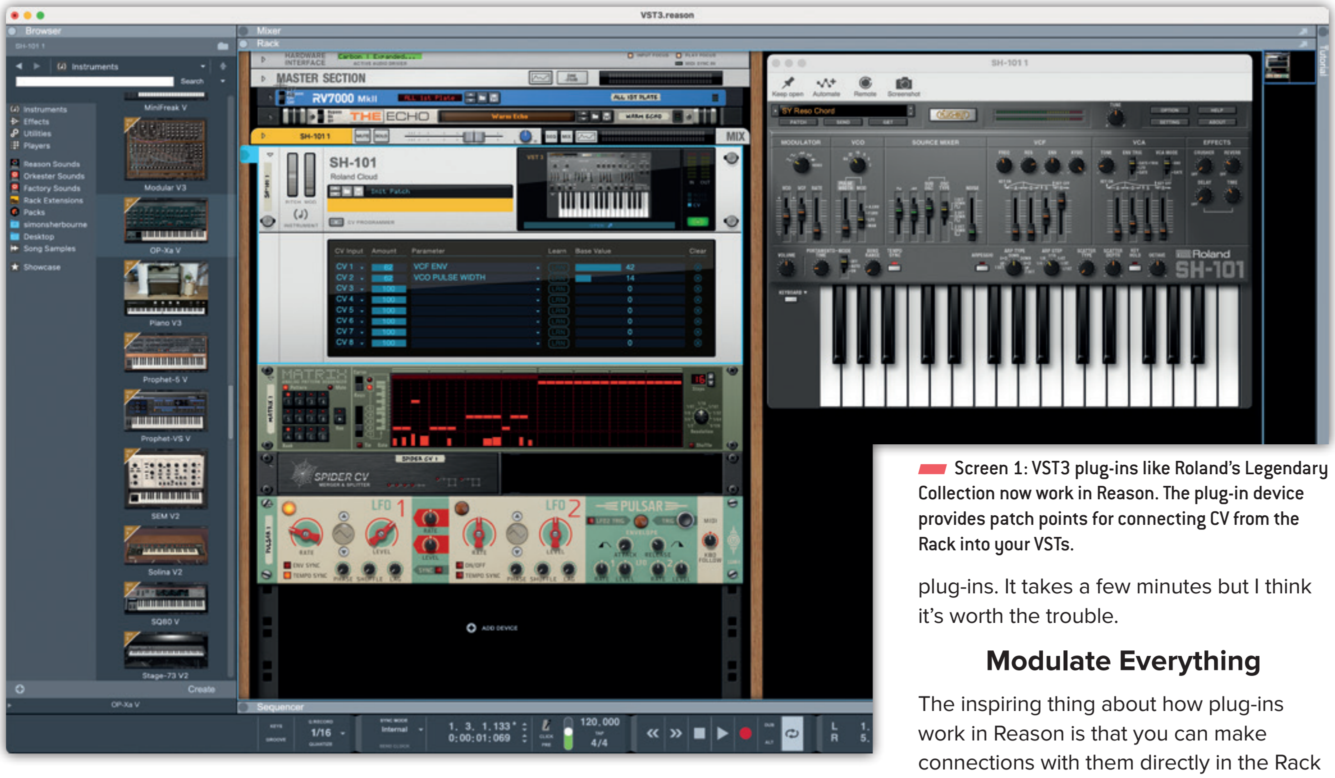

Screen 1: VST3 plug-ins like Roland's Legendary Collection now work in Reason. The plug-in device provides patch points for connecting CV from the Rack into your VSTs.

Reason 12.5 introduces full support for and integration with VST3 plug-ins.

SIMON SHERBOURNE

The recent Reason 12.5 update introduced support for VST3 plug-ins. Let's go over what this means and take a fresh look at plug-ins in Reason in general. Quick recap: VST plug-ins can be used in the standalone Reason app, but not when you're using the Rack as a plug-in in another DAW. Plug-ins are integrated into the Rack as modular devices, similar to the built-in instruments and effects. They are hosted inside a special device with connection points for CV and audio.

Housekeeping

Available VST plug-ins appear in Reason's browser, within the Instruments and Effects categories. They are listed after the built-in devices and grouped by manufacturer. If you already had plug-ins enabled, you'll most likely now see more choices populating the list. I was delighted to see the (VST3-only) Roland Cloud instrument collection appear in my list for the first time. I also noticed that two of each of the Arturia V-Collection instruments were now being shown, because I have both the VST2 and VST3 versions installed on my machine.

There are a couple of ways you can tidy things up. One is to go to Preferences / Folders and disable either the VST2 or 3 option. If all your plug-ins have VST3 versions it makes sense to disable the VST2 file paths altogether. This will save you quite a bit of manual sorting and also make Reason launch faster. However, if you have any plug-ins that are VST2-only, you need to go with plan B...

In the Windows menu select the Manage Plug-ins option to open a list of all the VST plug-ins that Reason has scanned (Screen 2). From here each plug-in can be individually selected and disabled or enabled. In my case I went through this list and disabled every VST2 plug-in that has a VST3 equivalent. Unfortunately, there's no way to select multiple items at once: you have to laboriously select, disable and confirm for every plug-in in turn. Get the kettle on first.

The other bit of housekeeping I'd recommend is to add pictures to each of your plug-ins so that you get graphical devices in the browser instead of grey text boxes. To do this, the first time you load any plug-in into the Rack open its window and click the Screenshot button. I got another cup of tea and did this for all my plug-ins. It takes a few minutes but I think it's worth the trouble.

Modulate Everything

The inspiring thing about how plug-ins work in Reason is that you can make connections with them directly in the Rack environment. In addition to Note and Gate inputs, each plug-in container has eight general-purpose CV inputs on the rear for patching in modulation from other devices. These can be mapped to just about any parameter on your device (everything that's automatable). This is a bit of an advantage over regular Rack Extensions which have pre-defined connections (although you can achieve the same thing by adding native devices to a Combinator).

As well as introducing VST3, the new version upgrades the plug-in audio output channel count to 64. This is great for drum machine plug-ins if you want to split each sound out to different processing chains or mixer channels. However, it's still not possible to send multiple MIDI channels into a plug-in in Reason for multitimbral operation. Hopefully the extra outputs hint that Reason Studios are working on this too. The other notable limitations are that you can't send MIDI out of a plug-in and there are no CV outputs: modulation and sequencing are a one-way street for plug-ins in Reason.

Let's take a deeper look at how to use CV modulation in plug-ins. In Screen 1 I've loaded up Roland's virtual SH-101 synth. I'm sequencing the plug-in from a Matrix unit which is auto-cabled to the Note and Gate CV connections round the back. Underneath the plug-in I've added a Pulsar modulation device, and on the rear panel

(Screen 3) connected the first LFO output to CV input 1 on the plug-in container.

Routing and scaling of the modulation is set up in the Programmer panel on the front. In the first column you'll see the list of eight CV inputs. Notice that these are pop-up selectors, enabling you to choose the same input more than once for patching a single source to more than one parameter. The second column is Amount, determining the depth of modulation. Next is Parameter: this is another pop-up list, showing every available mod destination. This list can be long, so it's usually quicker and easier to click the Learn button to the right, then wiggle the parameter you want to modulate on the plug-in's interface.

The final setting is Base Value. This is the nominal setting of the parameter before CV modulation is applied. Adjusting the Base Value is the same as moving the plug-in control on its panel — changes made from either place will be reflected on the other. There can be an offset if it's a stepped parameter like a switch: adjustments to the Base Value here will only have an effect on the plug-in at threshold points, but different settings will have an effect on how much modulation it takes to push a parameter to the next step.

Pushing The Envelope

In my example, I've mapped the Pulsar LFO to the Filter Envelope Amount on the SH-101, which is not a modulation that is available within the instrument itself. Note that both the Level setting on the Pulsar and the Amount setting on the plug-in CV Programmer will change the depth of modulation. A nice thing about modulation in plug-ins is that the changes can be seen animated on the plug-in control panel.

As a second example I'm using the Pulsar's envelope generator to modulate the SH-101's VCO Pulse Width. The envelope needs triggering, so I've used the Spider CV utility device to split the Gate output from my Matrix sequencer, taking this to both the Pulsar Envelope Gate input and the instrument plug-in's main Gate connection. The Pulsar's Envelope output is patched to CV input 2 on the SH-101. Again, I've set a Base Value and modulation amount, and can now see the PWM slider bouncing up and down as the sequence plays.

Remote Working

While we're going in-depth on plug-ins let's look at the other two key implementations: automation and remote control. There are buttons for both of these in the header bar at the top of any open plug-in window. The Remote button is for mapping MIDI controller sources to plug-in parameters. MIDI mapping is built into the native Rack environment, so for regular Rack Extension devices you can use the global Remote Override Edit Mode, or right-click any control on a device. Plug-ins are not included in this, so you need to do it by clicking the Remote button then clicking on a plug-in control. This brings up the familiar Remote Override window where you can assign a control source by manual selection or learning from a MIDI input.

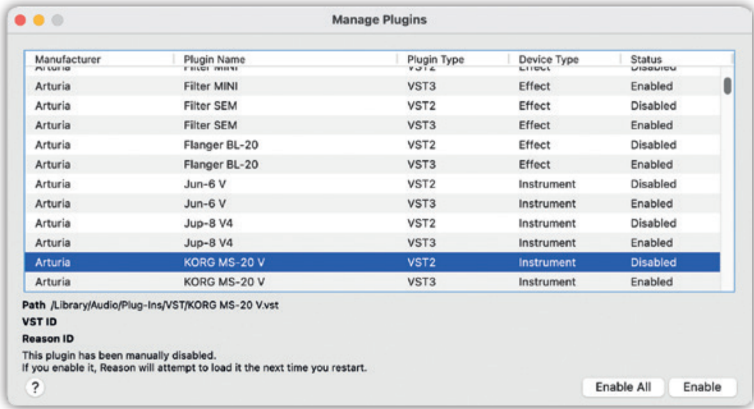

Screen 2: With both VST2 and 3 plug-ins now showing up in the list you may need to do some housekeeping in the Manage Plug-ins window.

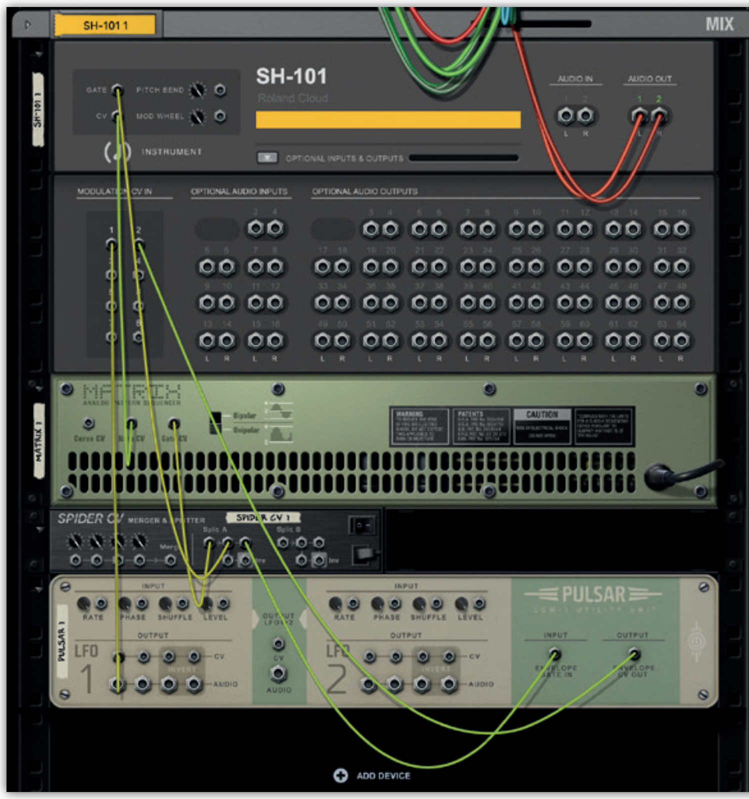

Screen 3: Here a Pulsar modulation generator is being used as both an LFO and Envelope source for the SH-101 plug-in synth.

The Automate button lets you pre-enable parameters for automation editing. Clicking this button followed by any automatable plug-in control will create a lane for that parameter in the plug-in's Sequencer track. However this is a redundant step if you just want to record automation on the fly: you can simply put Reason into Record and perform your parameter changes as you would with any other device. Automation is armed by default for whichever device or track is selected. You can also arm non-selected tracks by clicking the auto-arm button to the right of Mute and Solo in the track. After you've captured your moves a lane will be created in the main Sequencer automatically.

What's great is that track automation and CV Rack modulation of plug-ins work together in Reason rather than fighting each other. Modulation is effectively merged with automation, the latter simply moving the Base Value parameters over time.

INSIDE TRACK

SECRETS OF THE MIX ENGINEERS

Rob Bisel

It's no surprise that SZA's long-awaited second album has been a hit. More surprising was the key role of little-known engineer and producer Rob Bisel.

PAUL TINGEN

The big music-industry story of the end of 2022 and the beginning of 2023 has been the enormous success of SZA's single 'Kill Bill' and her second album, *SOS* — which, at the time of writing, is enjoying its fourth week at number one on the US *Billboard* album chart. In the five years since her 2017 debut album, *Ctrl*, SZA (pronounced 'Sizza') remained in the limelight through collaborations with Kendrick Lamar, the Weeknd, Travis Scott, Justin Timberlake, Doja Cat, Cardi B, Post Malone, DJ Khaled and many more, galvanising her newfound position as an A-list artist. She also let slip that she had worked on a second album with the likes of Mark Ronson, Tame Impala, Timbaland and Sia. Expectations were high, but there were many delays, causing considerable frustrations among her fans.

SOS was finally released on 9th December, but there was no trace of the aforementioned collaborators. Instead, big-name producers who had worked on the album included Babyface, Jeff Bhasker, Benny Blanco, Darkchild, Emile Haynie, the Rascals, Shellback and Michael Uzowuru. The album's credits also list ubiquitous top 40 hit mixers

Serban Ghenea, Manny Marroquin and Jaycen Joshua, as well as Shawn Everett, Dana Nielsen, Derek '206Derek' Anderson, and Jon Castelli.

Often described as alt-R&B and neo-soul, SZA's music is not typical top 40 fare, making the presence of so many top 40 hitmakers among the credits slightly surprising. However, a closer look at the credits throws a different light on the proceedings. It turns out that someone called Rob Bisel co-wrote and co-produced 17 of the album's 23 songs, has an engineering credit on all, and a mix credit on eight songs. These include three of the album's main singles: 'I Hate U', 'Shirt' and 'Kill Bill'.

Hot Sauce

"The success of the album is pretty cool," remarks Bisel. "I thought people would like it, but the strength of the reactions has been a surprise. It's really satisfying after all the hard work! For me, it's been a long journey. I'm 30, and I first worked in studios at the age of 17, when I interned at Studio 880 in Oakland, which was Green Day's unofficial studio. They were touring, so there was no action in the studio, but it was still a dream job for me to clean the spaces they would be in, including the bathrooms!"

Bisel's interest in music preceded his toilet-cleaning job at Studio 880, as he relates: "I played in bands as a teenager, mostly bass. I also did some pretty serious choral singing, that reared its head on the SZA album. But I was always interested in the recording side, and have been recording friends in my bedroom since I was 14. I never wanted to be a professional musician; for me, the studio side is the coolest thing ever!

"I'm from the Bay Area, but studied music at the University of Michigan, and during summers I interned for Mark Needham [*Fleetwood Mac, Imagine Dragons, the Killers — featured in* SOS *September 2017*]. I would show up to clean the studio and watch him mix, and definitely soaked up a lot of sauce from him. After I graduated from college I moved to LA, and when I was looking for a paid position, I stumbled on Dana Nielsen, who works a lot with Rick Rubin at Rubin's Shangri-La studio."

Learning From The Best

The connection with Rubin and Shangri-La turned out to become foundational to Bisel's career and skill set. "Dana connected me with the studio, and I started as an intern and a runner in 2014, and eventually became an assistant engineer and engineer.

"I ended up doing maybe 15 to 20 projects with Rick [*Rubin*], which really fuelled my hunger to dive deeper into production. It quickly became clear to me that his super power is taste, and watching him work was really inspiring for me. There are many intangible things I picked up from him, but possibly most of all Rick is a master at creating pressure-free moments where people can get out of their own way and do what comes natural to them.

"An added bonus of working at Shangri-La was that it's not just Rick and his projects, but others also book the studio. So I was a fly on the wall for Mark Ronson sessions, and Kendrick Lamar booked the studio for half a year, and I got to watch and assist his producers and engineers. If I were to try to describe my style and approach, it'd be a combination of so many people I have worked with, mostly Rick, but also Greg Fidelman, Jason Lader, Dana Nielsen, Ed Stasium, Ben Rice, Noah Goldstein, Caleb Laven and Ken Oriole, and Mark Needham for mixing."

Beats Working

At the beginning of 2020, after six years at the studio, Bisel started to feel that it was time for a change. As luck would have it, in her long-lasting journey to creating a second album, SZA had booked Rubin's Kauai home and studio, and needed an engineer to work with her. Rubin passed along Bisel's name, and two »

'Kill Bill'
Written by Solána Rowe, Carter Lang & Rob Bisel
Produced by Carter Lang & Rob Bisel

days later he was on a flight to Kauai. “We worked for a week together, with the two of us just vibing and coming up with song ideas. We returned to LA and worked for another two weeks, and then Covid hit and the world was shutting down. So she suggested that we stay and work at her house. So me working with her for a long period of time was sort of the product of the Covid situation.

“Among the first songs we did at her Malibu house were ‘Good Days’ and ‘Hit Different’. It was a pretty crazy situation, because she was used to going to big studios with tons of people coming in and out, and suddenly it was just the two of us, and I had to step into much more of a Rick Rubin producer role than just be the typical engineer. It gave me the space to present beats to her that I had made, alone or with others. If she liked the beat, she recorded over it. She also ended up really liking the way I record and mix her vocals, which led to me mixing songs as well.”

Much of *SOS* thus ended up being home-grown, because Bisel tends to start his beat ideas at his home studio, which is called Ponzu Studios, after his dog. Ponzu Studios is full of hardware instruments and outboard, and *SOS* features many ‘real’ instruments that were played rather than programmed. ‘Kill Bill’ is a prime example.

“The idea for ‘Kill Bill’ began with me messing on my [*Sequential*] Prophet 6. I bought it two years ago, and I played some basic chords using a flute-like sound. I recorded that into Ableton as audio, and added a bass line, using an electric guitar that I tuned down an octave. At a certain point, I didn’t really know where to go with it, and sent it to producer Carter Lang to see if it sparked something in him. He sent me back three or four different approaches, adding another bass, some guitars and layers of drum machines. The version we liked the most had a strong retro, almost boom-bap influence, sort of like Amy Winehouse.

“We played SZA five or six beats, including this one, and she immediately gravitated towards it. A week or so later she asked me to pull up the beat, so I put it on loop for her, and she did her thing where she goes very quiet sitting in a corner by herself. I know to stay out of the way and let her be. Five minutes later she said, ‘I have an idea. It might be a little too crazy, so let me know

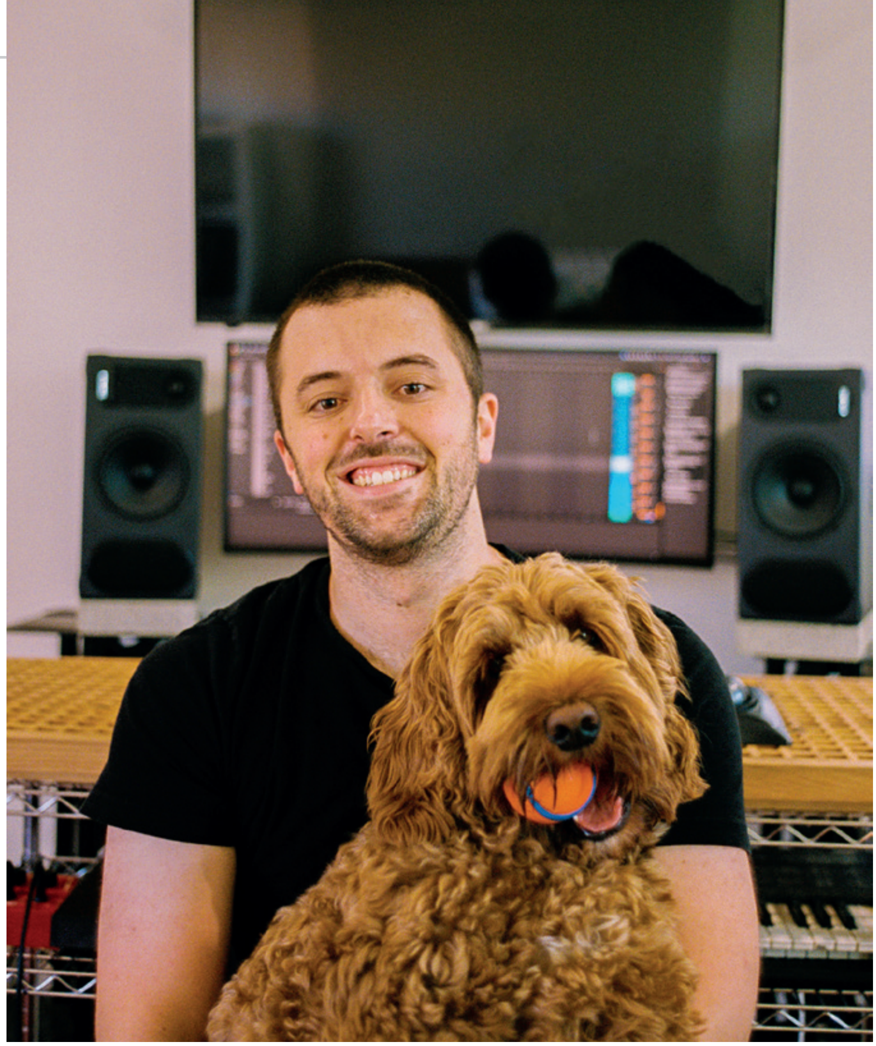

Rob Bisel’s studio is named after his dog, Ponzu.

Crucial to the genesis of ‘Kill Bill’ was a flute sound from Rob’s Sequential Prophet 6.

what you think.' And she sang the lyrics and the melody of the hook of the song, note for note. I couldn't believe it. If she had finished the song with passable verses, it would have still been great because the hook was so good, but of course she went on to write these incredible verses."

Boom-bap was a subgenre of '90s East Coast hip-hop, and the influences contribute to the general retro vibe of 'Kill Bill'. "The song came together really quickly and it felt fresh and like the type of music that hadn't been made before in today's world. As we were making the track, we were listening to the Everly Brothers and the Beach Boys and acts like that for vocal harmony ideas, and other things that you would not associate with a contemporary artist who critics often lump in with R&B. We were making something that felt really cool and exciting to us. It felt almost selfish to make something purely for our own enjoyment."

World Building

According to Bisel, the beat for 'Kill Bill' was just one of many "cool sonic worlds" he created with a variety of co-producers, in particular Lang. "He was one of the main producers on SZA's first album. We've known each other for several years through Shangri-La. For this album we teamed up whenever we felt we had a vision for a sonic place that SZA would like, and we'd then make batches of ideas within these sonic worlds. For example, there's an orchestral song called 'Blind', and we made probably about 20 different orchestral instrumentals, and we'd pick the three or four that we felt were the best, and played them to SZA. She then picked the one she liked best to work with."

Bisel worked for nearly three years with SZA, off and on, and on dozens of tracks, with many different versions. How did he manage to keep seeing the wood for the trees? "I'd say that's where my Rick Rubin education came into play," says Bisel. "I'd like to think I have a good ear for knowing when a change you make makes something better, or just different. I guess being able to keep an eye on the bigger picture is one of my strengths. Again, it's the result of the amount of time I spent with Rick, because he's a master at that.

"SZA also has a great ear for that, and a lot of it was the two of us bouncing off each other. We did millions of versions of almost every song, and it wasn't a waste of time. We were always trying something different — adding something, taking something away, trying a different effect, and then we had to judge whether it was harmful or helpful for the finished song. Someone else who was crucial in this respect was her manager and executive producer, Punch [*Top Dawg Entertainment President Terrence Henderson*]. We sent him stuff pretty much every day we did something, and he would be at the studio all the time sharing his thoughts. He has a great ear for the big picture."

Real Feel

Bisel starts his "sonic worlds" in Ableton, a DAW he "first tried about seven years ago. Before that I was making beats in Pro Tools, and it was pretty clunky. You can do it, but it's not the most organic workflow. It's not super-inspiring to me. I know amazing producers that are way better than me and that only work in Pro Tools, but for my workflow, I feel like I'm constantly hitting speed bumps. It's easier

Ponzu Studios has enough outboard to handle a decent-sized tracking session, connected to Rob's Mac using two UA Apollo x16 interfaces.

to audition sounds in Ableton, and being able to keep a loop going while adding tracks, and dragging and dropping sounds, and changing the duration of your loop in real time, are just some of the many perks of Ableton. Obviously Pro Tools has its strengths too, but for raw creation and jamming, Ableton tends to be smoother for me personally.

"Ableton is great for making stems, way better than Pro Tools, so that makes it easier to go between the two. When I'm recording SZA, I'm working with a two track of the beat in Pro Tools, in which I have a template to record her. I then may bring those vocals back into Ableton to spruce up the production, and then I'll go back into Pro Tools to finish the vocals and production. I also do rough and final mixes in Pro Tools. So I'm constantly bouncing between the two DAWs."

The retro feel of 'Kill Bill' is enhanced by live drums, bass, guitars and keyboards performed by Lang and Bisel. Although, in this case, most of these instruments were recorded at Lang's place, they could have all tracked at Ponzu Studio, which is equipped with all the necessary gear.

"I have guitars and basses and amps," explains Bisel, "a 1960s Ludwig kit, and a piano. I also have a bunch of synths, including a Roland HS-60, Korg DW-8000, Korg Polysix, Moog Grandmother, Prophet 6 and a handful of others things. Everything is miked up or plugged in, ready to go. It's another part of the Shangri-La philosophy that stuck with me. If you need 10 minutes to set up a mic pre, the spark of magic that was there may have evaporated. You don't want any hiccups over technical issues.

"I also use soft synths, by Arturia and Native Instruments, for example, and I love Ableton's Simpler. But I think instruments and hardware synths are just more creatively stimulating, especially when jamming up ideas with other people. It's easier to get more people creatively involved. They tend to cause more happy accidents too, which are priceless."

Fast Forward

Bisel's recording gear at Ponzu is also a mixture of 21st and 20th Centuries. "I have a BAE 1073 BAE mic pre, a Tube-Tech CL-1B compressor, and banks of mic pres, like the Focusrite ISA 828, and a Radial Workhorse with several API pres. It's a tasteful amount of craziness. My main microphone for vocals is the Shure SM7, which is my favourite mic. I have a handful of other mics on my drum kit, like the Shure Beta 91A on the kick, and a Neumann TLM102 on the piano. My good friend from college, Deni Mesanovic, makes some amazing mics through his company Mesanovic Microphones, and I have one on his Model 2 mics on top of my kit, and aim to get more, because they are incredible.

"My monitors are PMC twotwo 6s. I had the ATC SCM25As before, but I did an A/B and preferred the twotwo6s. I felt like they were a little more defined in the low mids. In the past I used a Pro Tools HDX system but I switched to two [*Universal Audio*] Apollo X16s. The HDX system wasn't working very well with Ableton, there would some latency and it would crash pretty frequently. When I switched to the Apollos, it resolved all those issues. It was a gamechanger for me.

"SZA has her own studio, and she owns a Neumann U47. So that's what we used for almost the entire project. Before doing this album, the 47 would not have been my favourite mic for vocals, but I've come to love them, and I intend to get one. She also has a BAE 1073 and a CL-1B, which is our default vocal chain. I never use EQ on the BAE 1073, because if we go to a studio with a 1073 reissue without an EQ, I don't want to suddenly have a huge part of our sound missing. I want things to be really super uniform and consistent.

"I'll rarely use the monitors when we work, because we cut vocals together in the same room, even when we're in big studios. We'll listen to what we have done on monitors at the end of the night or when we're taking a break, but so much of what we do is the two of us in the same room, with her three feet away from me, and me wearing Audio Technica ATH-M50x headphones. When we go to a big studio like Westlake or Conway with fancy live rooms it can feel a bit silly for both of us to be in the control room with headphones while we're at this big multimillion-dollar fancy facility.

Rob Bisel's Ludwig drum kit and upright piano are permanently set up and ready to record.

"Recording SZA's vocals is usually a matter of her doing 5 to 10 minutes freestyle over the beat, and I'll be making notes the whole time of what I think are magic moments and where they can fit in the song. She doesn't always agree and will have her own ideas of the structure of the song, and between the two of us we piece together the entire song. Getting the bare bones of the song in place can go really fast. We know we have a good song if it gets written in less than an hour. The songs we've laboured over and over-thought often aren't as potent as the songs that come naturally.

"With every recording we'll be talking about any issues with timing or tuning, and we tend to address those immediately. Sometimes something slips by after the first night, and we'll tweak it the next day. I also try to add effects during the writing process, like putting on crazy hard Auto-Tune or adding a cool delay or reverb or formant shift, because it can spark new ideas for her and inspire a lyric or a theme for the song, or it can »

heavy lifting, so I added probably 15-20 percent. I knew what the finish line sounded like with that song, and I knew we were not quite there with the rough.

"Vocally I did not change much, but I really wanted the instrumental side to hit harder, particularly the drums, and make the entire track sound more boom-bap-like. So unusually for me, I approached this mix as if I had never heard the song before, and went through part by part from top to bottom, starting with the drums. There were so many layers, I wanted to make sure each layer got the love and attention it needed. I also approached this mix as if I was doing it on a console at Shangri-La, with subgroups everything is sent to, which in turn are bussed together. So it was more of an old-school mix approach, that I might not necessarily do with a more modern-sounding song."

Bisel's mix session of 'Kill Bill' is roughly 120 tracks large, and structured in a conventional way, with drums at the top (purple), then bass (blue), guitars (green), keyboards (purple-blue), choir and backing vocals (purple), lead vocals (beige and green) and main group tracks (black). Many of these sections are sent to their own aux group tracks, and aux effect tracks are integrated alongside, rather than grouped at the bottom of the session.

» influence her vocal delivery. We'll often go in again later and do final takes and polish any weak spots."

Radio Ready

Bisel was doing rough mixes of the songs SZA and he were working on continuously. "Yeah, there's an expectation as we're making a song that whatever I send her will sound like something that could be heard on the radio. It's never, 'Just wait until it gets mixed.' I always treat a session so it could be released immediately. And this has happened, where she posted a snippet on Instagram or Twitter, and it becomes the sound that people expect. If we deviate too far from that with the final release, people pick up on it, and say, 'What happened to the mix?' It's been shocking to me how perceptive the ears of some of her fans are!"

It's a process that contributed to Bisel doing the final mixes on a number of songs. "I think in several cases, SZA liked the way I had mixed songs when we were recording them. With 'Kill Bill' I asked Punch if it was cool that I mixed that one, because I felt we were doing such a specific thing sonically, that it might get lost if we were to hand over the mix to someone else. We had made the beat, and recorded her vocals and helped find a sonic world for all that, and I really wanted to see it through all the way to the end. Luckily they were on board with this.

"I mixed 'Kill Bill' at my house, entirely in the box, because we were moving and things were changing so much, I needed to be flexible. It was probably the mix in which I went the furthest. Normally I don't stray far from my rough mixes. I might take them maybe 5-10 percent further. But with 'Kill Bill' I felt it needed a bit more

Drums

The live drum kit recording made at Lang's place contributes four tracks: kick, snare, hi-hat and overheads. "It's Carter playing, and it was a full pass of drums that we edited and comped, and during the mix I went in and chopped up individual kicks and snares so I could really get those isolated and have less bleed. I wanted more control sonically, and I probably did some minor timing adjustments as well, but nothing crazy, because I wanted to keep the human quality that you'd expect with a song like

Rob Bisel says the spring reverb module from NI's Guitar Rig is the best he's found in software.

this. Next to the Kick.02 track is the kick from a drum machine programmed by Carter, that adds more low end.

"The live kick has the Knock plug-in, which is made by Plugins That Knock. It's got a very simple layout, just a couple of knobs, Punch and Saturation, and I feel a bit lazy using it, but it's really cool and inspiring and can create a very powerful sound. I also have the UAD SPL Transient Designer for more attack, and I have three sends, one to Kick Side-Chain, one to Drum Crush and one to Drum Aphex. I wanted the bass and guitar in particular to duck a little when the kick hits, as if the kick was printed to tape too loud, emulating an analogue phenomenon. The Crush aux has the UAD Fatso for parallel compression, and the Waves Aphex Vintage adds more top end. In general, I wanted to add more snap and bring out drum textures that otherwise might get buried.

"The snare also has the Knock, and the Transient Master, because I wanted it to sound more pokey. The Waves Puig-tech EQP-1A adds more air. The snare is also sent to the Drum Crush aux, as is the overhead track. Below all this is an assortment of drum machine tracks, and many of them have the Native Instruments Guitar Rig plug-in on them, which I think is cool from a mix standpoint. Guitar Rig has the best spring reverb I've found. I'm using it to give more of an old-school '50s/'60s sound. There's also a Soundtoys Little Radiator on some of the drum machine tracks, to add some crispiness and drive, as if it's coming from a broken drum machine. It's all to give these tracks more life."

Bass, Guitars & Keys

"The 89 Bass track is the recording of my electric guitar, tuned down an octave, and there's a bass guitar track done by Carter and I. The sub-bass comes from the Waves RBass on the Bass Edit track. Both bass tracks are sent to the Bass Crush aux, which is doing a good amount of heavy lifting, with the Waves CLA-76, FabFilter Pro-Q 3 EQ and Pro-C 2 compressor. It is intentionally a very bass-forward song, to create a modern spin to an old-school song and feel, and the bass was the modern element. We tried to have it be almost like an 808.

"The guitars are also all played. I treat them on three aux tracks, GTR1, GTR2 and GTR3. Maybe it's lazy of me, but it's easier for me to wrap my head around treating things from a broader perspective, rather than getting into detail on each track. Having said that, Carter and I recorded each of those individual sounds the way we wanted them. Kind of like pre-mixing it in some ways.

"The Prophet tracks are the original idea that started the song. I wanted them to sound lo-fi and old school, so I added the XLN Audio RC20 Retro Color plug-in. I tend to think it's an easy option, but it sounded great on this. The Prophet tracks are sent to two effect aux tracks, Prophet Slap and Flute Verb, which are also part of me going for that vintage sound. The Slap has the UAD Galaxy for the slap, and the Decapitator for some dirt. The Flute Verb aux has the AudioThing Springs plug-in, that is great. AudioThing is possibly my favourite plug-in company right now. All their stuff is unique and super-helpful."

Soundtoys' EchoBoy was one of several delays and reverbs used on the lead vocals.

Vocals

"All the tracks marked RB are the Beach Boys-sounding choir thing. The individual tracks have [*Antares*] Auto-Tune Pro, [*FabFilter*] Pro-C 2 compression and Pro-Q 3 EQ, and the most important treatment is that they all have a send to the BGV Spring effects aux, which has the UAD AKG BX15, and the Valhalla VintageVerb. The Little Radiator is first in the chain, to give it more edge, so it feels more sampled. Those parts sound night and day without that reverb, it definitely brings it into the same sonic space as everything else.

"Track WL47.47 is the main lead vocal. There's a block of five lead vocal tracks, and WL47 3.45 is like a filtered sample vocal that comes in during the second pre-chorus. All lead vocals go to the Lead Vocal bus, which is where everything happens in terms of treatments, with a chain of the FabFilter Pro-DS, Pro-Q 3, Waves CLA-3A, Pro-C 2 and sends to five effect aux tracks, with reverb from Valhalla VintageVerb and Guitar Rig, delay from Soundtoys EchoBoy, and the Soundtoys Little MicroShift and [*Waves*] MetaFlanger. The reverbs are pretty prominent, but the rest of the auxes are all subtle brush strokes.

"These treatments are pretty standard when I work with SZA. I try to keep it in this zone, so when I do a bounce, it doesn't sound too different song by song. The backing vocals all go to a BGV aux, on which I have the Pro-Q 3, Waves RCompressor, Puig-Tech EQP-1A and MetaFlanger. The final group tracks mostly have just the Waves CLA-2A and Pro-C 2 for some control. Volume automation on these sections is the main goal of the group tracks.

"Finally, I tried some things on the master bus, like tape machine stuff, and decided it sounded better without it. On the Mix bus I have iZotope Ozone 9, FabFilter Pro-L 2 and the [*Oeksound*] Soothe 2. I need to turn my mixes in at a decent level, and I normally use the Pro-L 2 for volume. I also applied a little compression and EQ with Ozone, and Soothe tones down some resonances, nothing crazy, just some soft, soft brush strokes. I keep it as simple as I can."

AMD v Intel

CPUs On Test

What benefits have the latest crop of CPUs brought for audio users — and what might the coming year have in store for us?

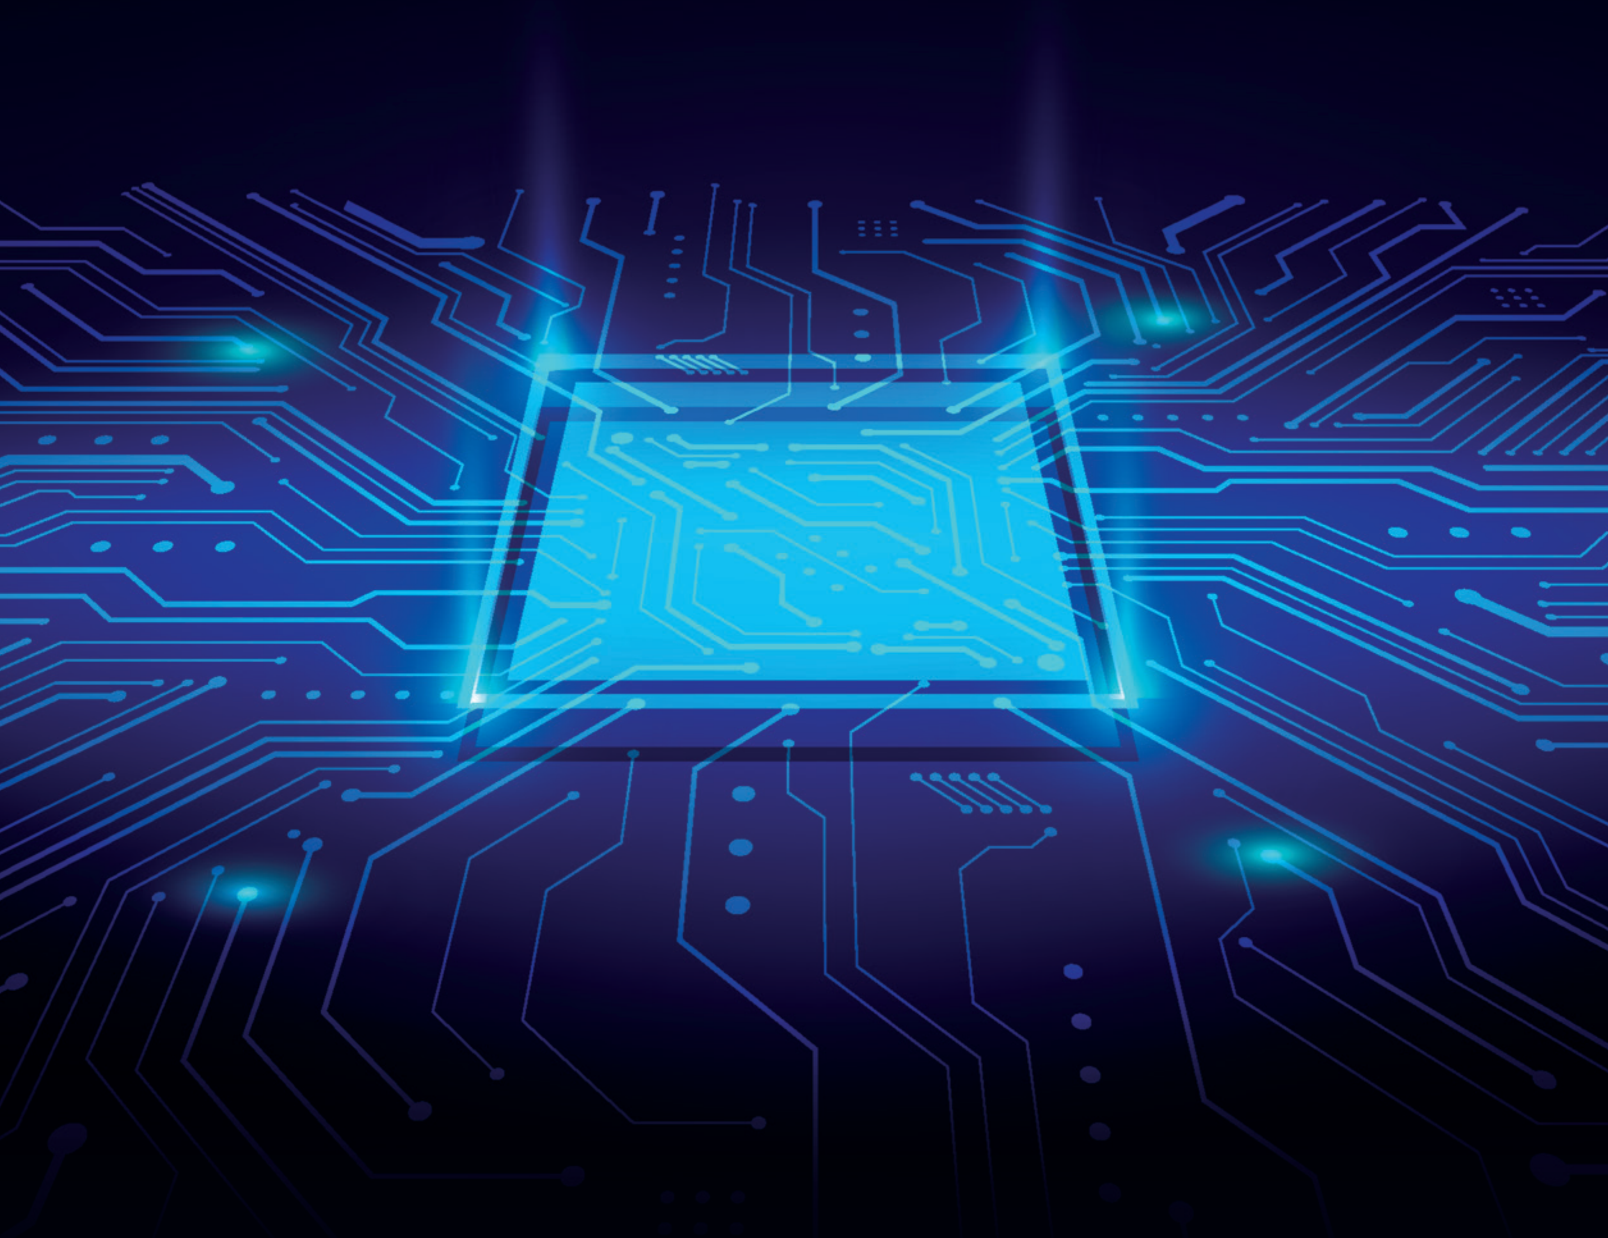

PETE GARDNER

For those who keep a keen eye on the computer hardware market, it's been an interesting few years when it comes to CPU development. Intel, the leaders for many years when it came to high-performance applications such as audio and video production, were dethroned a few CPU generations ago. Indeed, for a while they appeared almost to be heading out into the wilderness due to delays in their platform updates, but the last time we took a look at how the leading Intel and AMD chips compared (a shade under a year ago, in *SOS* April 2022: www.soundonsound.com/sound-advice/intel-v-amd), Intel were showing strong signs of a recovery, and we found that they and AMD performed comparably at that time.

Intel's step forward at that time was largely as a result of moving over to a so-called 'big.little' core arrangement, with the CPU as a whole using both high- and lower-performance cores in an arrangement similar to that in Apple's current hardware. A timely switch to Windows 11 helped to support this new way of working and, now various DAW and plug-in creators have had the chance to do a little tweaking and refinement, it seems largely to have proved successful. This, along with support for the newer DDR5 RAM, meant we saw large intergenerational gains in Intel's testing results. Still, as I noted at the time, all of this was just enough to bring them back to level pegging with the well-established high performance of the AMD platform.

Hot Gossip

Fast-forward to today, and the situation has reversed: it's Intel who are now looking to optimise an established platform, whereas AMD have brought out their new AM5 chipset which, along with any gain in raw performance, incorporates support for DDR5 RAM. While Intel's last generation Z690 chipset continues to support their newly released 13th-generation CPUs, the company have also brought out their new Z790 chipset, which includes PCIe 5.0 support. So too does the new AMD platform. All of which means that both companies' latest ranges boast a similar range of features.

Although dedicated PCIe 5.0 devices remain thin on the ground — the first next-generation storage options aren't due to arrive until later on this year — the upgrade to PCIe 5.0 doubles the available bandwidth, which means there are other potential benefits. A number of mainboards have already appeared on both platforms which can take advantage of these extra resources, and offer up to five current-generation PCIe 4.0 M.2 NVMe drives, along with the standard SATA drive connections. That will be a boon to anyone who's looking to fit a large amount of fast storage for their ever-growing collection of audio projects and sample libraries; those doing serious virtual orchestral work, for example.

Given the similar performance of their high-end chips a year or so ago, it's perhaps not surprising to discover that both companies have chosen to focus most strongly on headline improvements in the raw hardware performance. The performance gains are most notable with the companies' flagship Intel 13900K and AMD 7950X models, but a similar trend can be seen across both ranges. Of course, these attempts to snatch this generation's number one slot mean that power consumption has risen, and while that's not a huge concern for many users, there's a significant issue for audio users: while they might well be happy about the performance boost, higher power consumption invariably means more heat and thus potentially more system noise due to cooling.

Intel's Z790 boards play host to CPUs which are already pushing the envelope, with a 253W TDP (Thermal Design Power — the maximum amount of heat generated) rating on the Core i9 13900K for its Turbo mode at stock ratings. This has been amplified by a number of new mainboards which, by default, already apply a small degree of performance boosting — which results in the Intel-advised power limits being largely ignored! During our tests, we found that various boards apply chips' PL2 voltage settings (Performance Level 2, intended for turbo performance) to regular PL1 workloads, whilst also unlocking the PL2 mode's power draw completely, to allow the CPU to use as high a voltage as it wants when under heavier turbo-boosted loads. This resulted in peak loads occasionally spiking in excess of 325W total draw, a level which would prove challenging for even the very best cooling options available today.

When the power consumption was unlocked in this way, the 'real world' performance gains that were achieved » were surprisingly modest: typically it resulted in only an extra 100MHz being applied to each core. In terms of practical, everyday use, that seems a poor trade-off. Making a few BIOS adjustments to set things in line with Intel's specification vastly reduced the power spikes (to a much more manageable level) and brought the average temperatures down to a point that made quieter cooling far easier to achieve.

For the previous generation, AMD's Ryzen platform was already the more power-efficient solution, and although we've noted an increase in power usage this time around, the upper specified limit of 170W on the 7950X model ensures that it remains significantly less power-demanding than the Intel equivalent. AMD have, to some degree, changed how performance is managed: typically, with Intel and previous AMD generations, you set your CPU target speed in the BIOS and the board would attempt to use whatever power it took to achieve it, unless you specified otherwise, even to the point of overheating. AMD have now switched to an approach whereby you set a maximum target temperature under full load, and it works to maintain a steady degree of performance depending on how well the system is cooled. Our tests suggest that this works very well.

Again, we've found that the AMD board settings push slightly beyond the official limits, with the CPU peaking above 200W when the overhead was available, but this is where AMD's new Eco mode feature comes into its own. With a number of different low-power settings, all the way down to 65W, making a change to the setting within the BIOS allows you to limit the chip's overall power draw. With a workstation, you hardly want to be limiting your performance quite so dramatically, but the upper 170W Eco mode setting is of interest, as it quickly brings the CPU back into AMD's official spec. In stress testing of the 7950X, we observed that, whilst unlocked, it would swing between roughly 5200 and 5300 MHz across all cores when under a 100% load. With the 170W Eco mode applied, this stabilised in the middle, with a more constant 5250MHz, allowing the system to run slightly cooler overall. The benefit extended to running a few extra plug-ins in our DSP test (of which more later), because the removal of the larger dips in processor clock speed helped to achieve slightly greater performance overall.

With these observations in mind, we tested both companies' current flagship chips, in setups that adhered strictly to their official specs, and with the same cooling system used for each: a Be Quiet! Silent Loop 2 280mm AIO (all-in-one) water cooler. Interestingly, the performance of the AMD Ryzen 9 7950X on the 170W TDP setting remained similar when substituting the Silent Loop system for our regular Dark Rock Pro 4 fan-based cooler; switching to a triple-fan water AIO model may help push the CPU performance slightly further, but at the cost of noticeably increased fan noise. The Intel 13900K, though, with its higher TDP figure of around 253W under maximum load, would still very briefly peak at a level a little beyond the rating of even the best of air coolers, as the turbo boosted and then settled again. So for this CPU at least, a water cooler AIO is definitely recommended.

Moving away from the flagship chips, the difference was striking. Even though AMD's 7900X and Intel's 13700K models have similar 'on-paper' power usage ratings to those of their larger siblings, these less densely packed chips, with lower core counts and reduced maximum turbo ratings, typically proved much easier to cool — a good air cooler is an option for either chip (and those below them in either range).

On The RAM Page?

System memory is often of just as significant interest to music makers as the CPU itself — sample-library players such as NI Kontakt, which is used in our DAWBench VI test, can benefit from more and faster RAM — and it's another technology which has seen a significant jump in performance in recent years. If you read the previous article on this subject (the one linked to in the opening paragraph) you might recall that the 12th-generation Intel chips were compatible with a mixture of mainboards, some supporting DDR4 RAM and some DDR5, and that in testing it became evident that switching to DDR5 could prove beneficial for anyone working with a lot of audio and library data. Yet this was also a source of frustration, since the initial supplies of DDR5 sold very fast, and the supply chain problems at the time meant many early adopters were left waiting a few months before they could implement this new technology. Thankfully, both the supply and cost slowly improved over the course of 2022, and now that AMD's AM5-chipset-based boards support DDR5, the year ahead should see its more widespread availability and adoption.

When it comes to running the most effects and processing plug-ins, Intel appear to have reclaimed the top spot — for now!

The RAM for our latest tests was DDR5, and we ran comparable dual 32GB sticks for 64GB at a 5600MHz rated speed on each platform, with Intel XMP or AMD EXPO profiles as appropriate. Over the past few years, we've found that one way to give an easy bump to performance is to 'overclock' the RAM, simply by using kits that are faster than the chipset's base recommendation. It yields a small but worthwhile boost to performance without any major increase to the voltage or heat levels associated with CPU overclocking. This time around, there are some additional things to note from the manufacturer's recommendations which may affect anyone looking to maximise their available system memory.

For this latest generation, Intel advise RAM speeds of up to 5600MHz, while AMD lists support at 5200MHz. Note that, in both cases, these figures are for a pair of memory sticks running together in a dual-channel arrangement. Since DDR5's launch, ever-faster RAM kits have become available, with some of those starting to appear on mainboard QVL (Qualified Vendors List) sheets. At the time of writing, the very best is capable of running all the way up to 7000MHz, but that's only for dual-RAM-stick arrangements. For a full complement of four sticks, which is required if you wish to install the maximum allowed quantity of 128GB, the situation becomes a little murkier.

Once fully populated, the Intel chipset will tend to limit itself down to the platform's actual base RAM clock speed of 4800MHz. A search through mainboard QVL charts for pre-validated solutions suggests that this means there's currently not yet a 'real world' 128GB configuration available (at the time of writing — this situation could change quickly!). AMD has a similar lack of pre-validated options, specifying a supported limit of only 3600MHz when configuring a similar arrangement. That said, those who are willing to experiment

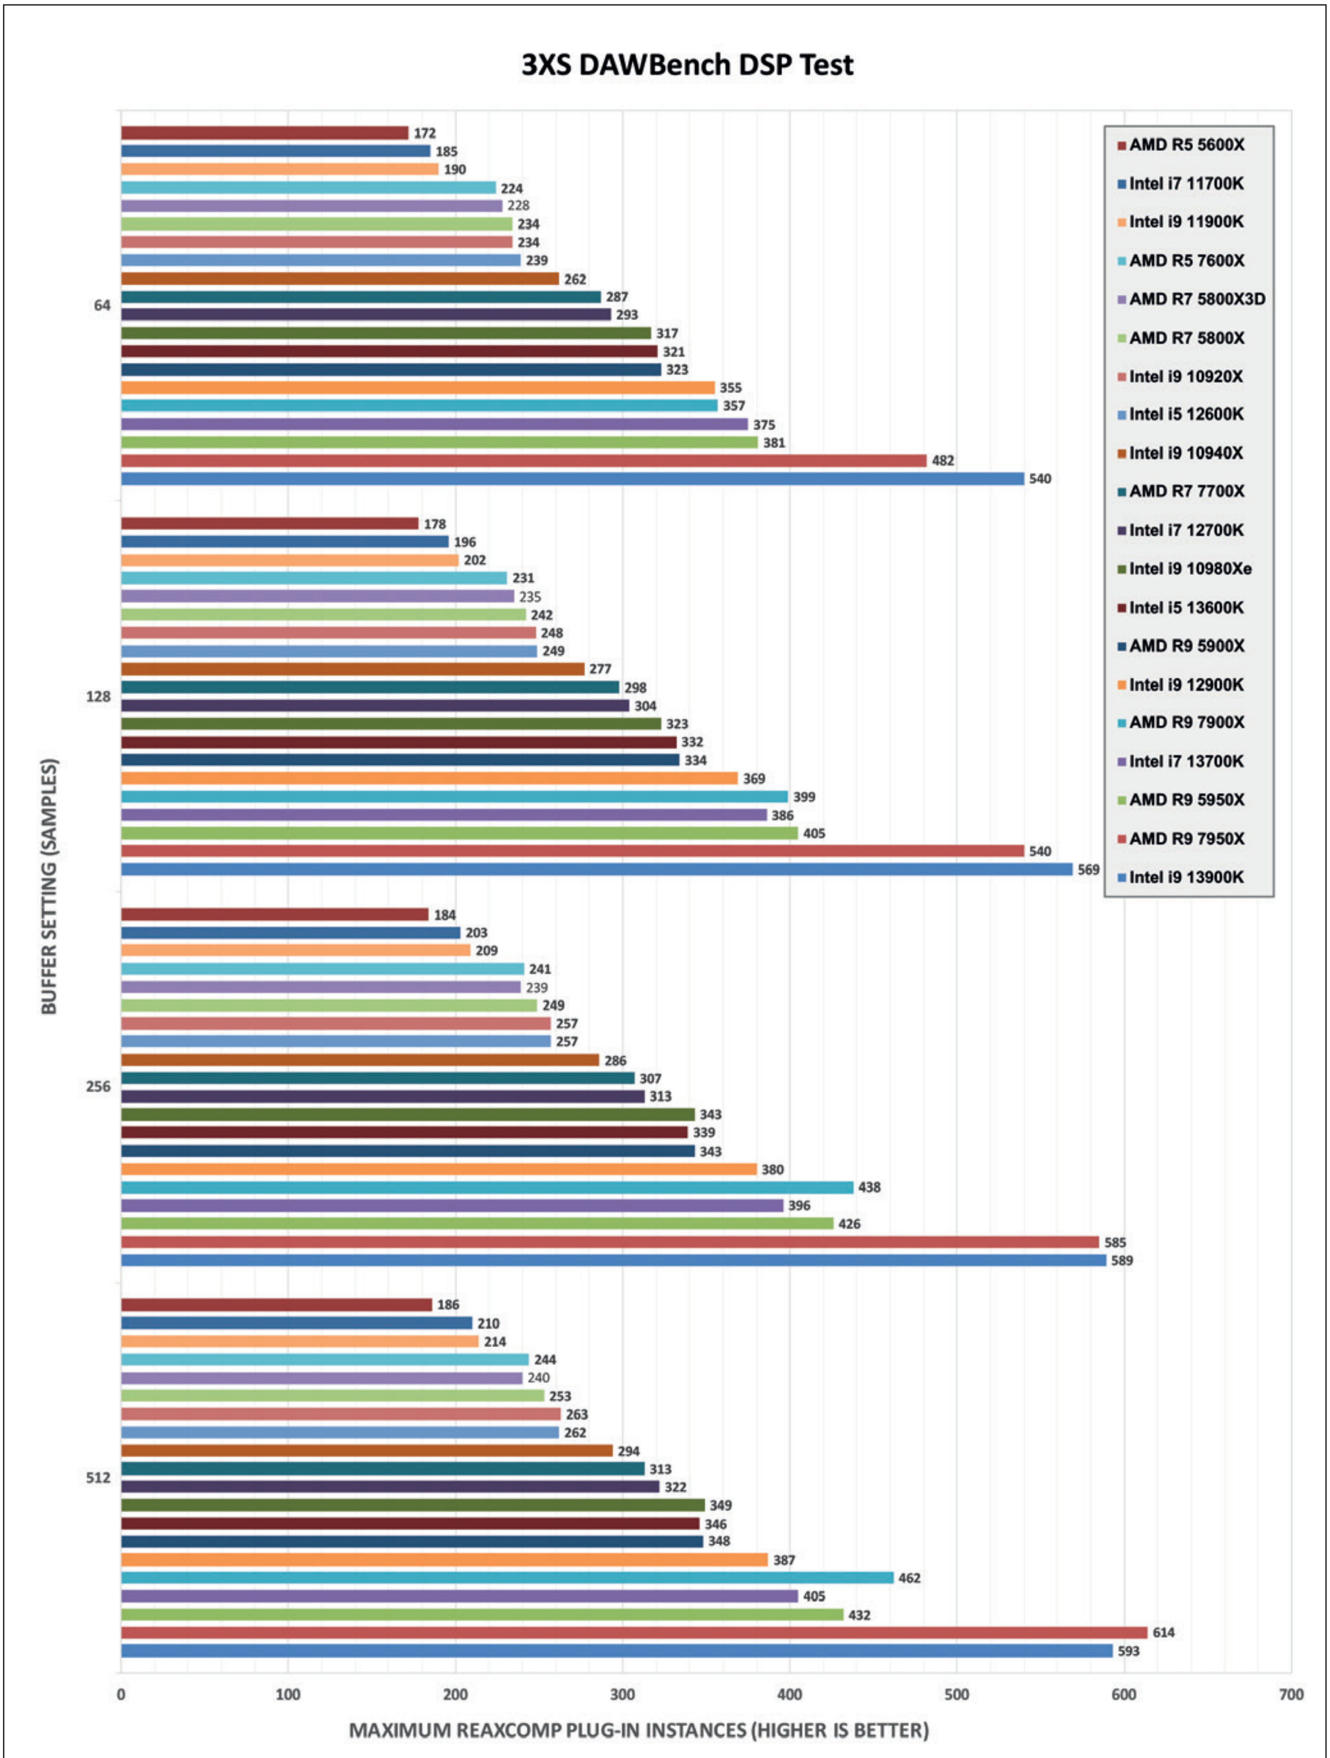
3XS DAWBench DSP Test
AMD R5 5600X
Intel i7 11700K
Intel i9 11900K
AMD R5 7600X
AMD R7 5800X3D
AMD R7 5800X
Intel i9 10920X
Intel i5 12600K
Intel i9 10940X
AMD R7 7700X
Intel i7 12700K
Intel i9 10980Xe
Intel i5 13600K
AMD R9 5900X
Intel i9 12900K
AMD R9 7900X
Intel i7 13700K
AMD R9 5950X
AMD R9 7950X
Intel i9 13900K
64
172 185 190 224 228 234 234 239 262 287 293 317 321 323 355 357 375 381 482 540
128
178 196 202 231 235 242 248 249 277 298 304 323 332 334 369 399 386 405 540 569
256
184 203 209 241 239 249 257 257 286 307 313 343 339 343 380 438 396 426 585 589
512
186 210 214 244 240 253 263 262 294 313 322 349 346 348 387 462 405 432 614 593
BUFFER SETTING (SAMPLES)
0 100 200 300 400 500 600 700
MAXIMUM REAXCOMP PLUG-IN INSTANCES (HIGHER IS BETTER)

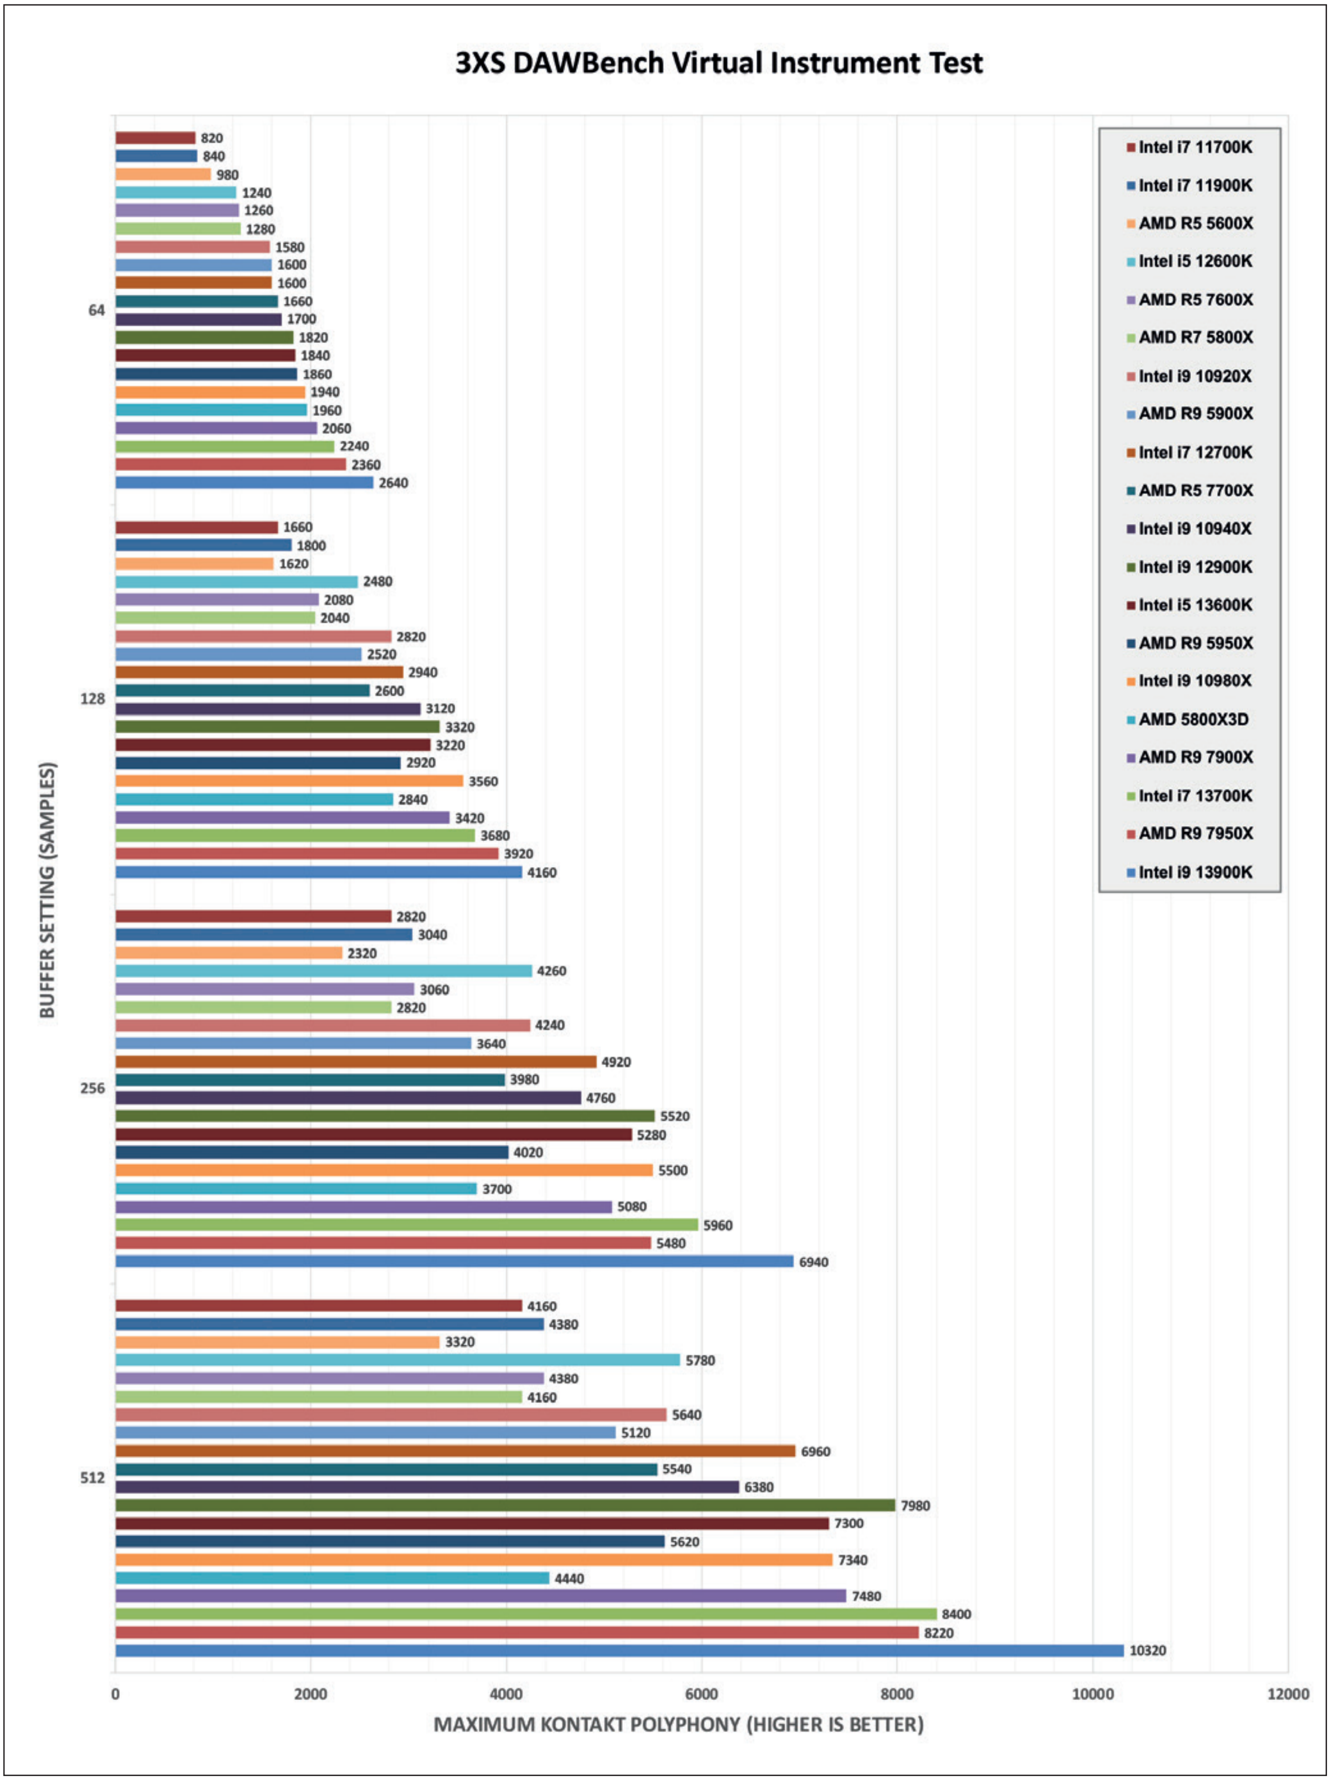
3XS DAWBench Virtual Instrument Test
Intel i7 11700K
Intel i7 11900K
AMD R5 5600X
Intel i5 12600K
AMD R5 7600X
AMD R7 5800X
Intel i9 10920X
AMD R9 5900X
Intel i7 12700K
AMD R5 7700X
Intel i9 10940X
Intel i9 12900K
Intel i5 13600K
AMD R9 5950X
Intel i9 10980X
AMD 5800X3D
AMD R9 7900X
Intel i7 13700K
AMD R9 7950X
Intel i9 13900K
BUFFER SETTING (SAMPLES)
64
820 840 980 1240 1260 1280 1580 1600 1600 1660 1700 1820 1840 1860 1940 1960 2060 2240 2360 2640
128
1660 1800 1620 2480 2080 2040 2820 2520 2940 2600 3120 3320 3220 2920 3560 2840 3420 3680 3920 4160
256
2820 3040 2320 4260 3060 2820 4240 3640 4920 3980 4760 5520 5280 4020 5500 3700 5080 5960 5480 6940
512
4160 4380 3320 5780 4380 4160 5640 5120 6960 5540 6380 7980 7300 5620 7340 4440 7480 8400 8220 10320
0 2000 4000 6000 8000 10000 12000
MAXIMUM KONTAKT POLYPHONY (HIGHER IS BETTER)

Intel also seem to be out in front again when it comes to sample libraries and virtual instruments — though signs point to some interesting developments by AMD in the coming months.

and stray beyond the official specification will potentially find a little more wiggle room with the right hardware choices.

A quirk of the AMD architecture resulted in previous generations performing more optimally when RAM overclocking as you reached 1:1 parity with the speed of the memory controller. For this generation, 6000MHz is that target figure. Larger 32GB-based kits that could run at 6000MHz were unavailable at the time of writing, but initial tests with 4x16GB-based kits have gone well — this suggests that when 32GB sticks (these are required to install the full 128GB) become available, they could be well worth exploring.

Top Spot?

Returning to the present, let's consider the latest set of test results. If you look over the charts for the ReaXcomp compressor-based test, which puts the spotlight on digital signal processing (DSP) performance, you'll notice straight away that competition for first place is fierce. As you might expect, the Core i9 13900K and Ryzen 9 7950X are firmly ahead of the pack, but there's a tighter spread for the range of Intel models across the various buffer sizes, while the AMDs tend to come into their own at higher buffer settings. Indeed, the Ryzen 9 7950X is ahead overall on the uppermost 512 buffer setting.

For the Kontakt-based virtual instrument test, which asks more of RAM performance, Intel offered the strongest performer last time around and that's repeated here. Not only does the Core i9 13900K put in a very strong showing, but the next-model-down Core i7 13700K holds its own against AMD's top-of-the-range R9 7950X. We saw in the comparison test last time that this type of workload would take advantage of DDR5 where it was available, but while this means the new AMDs certainly outperform their previous-generation CPUs it seems they can't yet quite keep pace with Intel here.

This pattern is repeated across both charts: Intel tend to appear stronger when placed up against AMD's competing model. And while AMD's CPUs used to be less expensive than the corresponding Intel models, that's no longer the case — so evaluating the best bang for buck for your particular requirements should be pretty straightforward. Of course, the power draw could also be a consideration. Intel really have gone all out to secure the top result, and that comes at the cost of greater power consumption, although, having said that, both platforms do scale the power draw according to the required usage, and those headline upper limits are likely to be fairly uncommon in daily use. Still, with less heat being generated overall, there may still be benefit to AMD's lower power consumption if configuring a particularly quiet system is a priority.

With the introduction of their new platform, AMD have lost perhaps the biggest competitive advantage they enjoyed last time around. Intel had suffered not just from stock availability but also the cost associated with the early adoption of new technologies like DDR5. Since both companies' ranges performed comparably last time, it made sense to a lot of people to pick up a well-established, lower-priced AMD AM4-based board along with the more freely available and cheaper DDR4 RAM. AMD's move to DDR5 has levelled the cost and availability playing fields — the first range of AM5 boards cost broadly the same as their Intel equivalents. AMD launched their most recent chips a month before Intel's went on sale, and the latter were able to price their range accordingly; priced similarly to the models they were replacing, Intel's new CPUs are more keenly priced than most people had expected. It's a move that has already caused AMD to revise their prices downward. (See? I told you an 'arms race' was on the cards!)

The Need For Speed

One result from this latest round of tests that sticks out like a sore thumb is the great low-latency performance of the AMD Ryzen 5800X3D: this CPU isn't even from AMD's latest generation of chips! it arrived towards the very end of the 5000 series' reign, and it could well be a sign of what's yet to come from AMD. Readers of past CPU round-ups may remember me remarking that the low-latency ASIO handling of AMD's earlier ranges tended to be weaker than Intel's (an issue for some music makers, though not all). AMD's new 3D cache design, which is used in the 5800X3D, helps to reduce the CPUs' internal latency, which is what lay at the heart of this issue. But while low-latency performance has improved greatly in recent generations, the AMD chips do still lag slightly behind the equivalent Intels when using the lowest (64-sample) buffer settings. The lower-clocked 5800X3D ranks, as expected, slightly below the regular 5800X, but its cache comes into play where AMD has often proven at its weakest — and, as a result, we see an absolutely storming virtual instrument test result that, with the lowest buffer settings, effortlessly exceeds the performance of what even that range's top-end 5950X model can deliver.

At the start of January 2023, in a bid to follow up on this 'experiment', AMD confirmed that a number of 7000-series 3D cache-enabled models would become available later in Q1 2023. This includes the 7950X3D, a revision of the top-end chip. This suggests that, even though it's only a few months since the launch of the current range, AMD are already preparing their latest challenger.

Intel's next development isn't too far off, either. They're looking to introduce their long-awaited successor to the X299 high-end chipset in the first half of 2023, in the form of the Fishhawk Falls chipset platform, as well as a workstation-grade Sapphire Falls range. Indications are that these will remove the 'efficiency' cores found in the established 13th-generation chips, and employ a more traditional array of full-performance cores across the whole chip. That, along with support for up to 256GB of DDR5 memory, has the potential to make this the ideal platform for anyone working on the largest of projects, such as scoring with large libraries to film or other mixed media.

Takeaways

It's clear that Intel have gone all out to deliver the highest performance. It's been a bold move, and one that looks, for now, to have paid off for them. Their willingness to adopt an aggressive pricing strategy demonstrates their keenness to regain market share, after the past few years' dip. AMD's reluctance to bring us their 3D-cache-enabled 7000 series at the launch of their new range may have left them on the back foot for now, but with a mid-cycle refresh already imminent, these could well shake up the charts once more. This is already shaping up to be another exciting year in tech!

How I Got THAT SOUND

Hear The Sound

W https://open.spotify.com/track/5huSgTVqBdIJTFhpShJfcg
W www.youtube.com/watch?v=MQcPB1WkISI

Ed Stasium • Living Colour 'Love Rears Its Ugly Head'

JOE MATERA

Since starting out in the early '70s, American producer, engineer and mixer Ed Stasium has worked with a *Who's Who* of the music world, from the Ramones, Talking Heads and Mick Jagger to Soul Asylum, Julian Cope, Motörhead and many, many others. From this stellar discography, Ed nominates the drum sound on Living Colour's 'Love Rears Its Ugly Head' as a particular favourite.

"I absolutely love the drum sound we achieved on this track because of its natural ambience, with no outboard reverb whatsoever. Being self-taught with no proper training everything I developed was through my own experimentation and experience. When I started out, my personal observation was that the drum sounds that were being recorded in studios all sounded dead. In November of '75 I was hired as a staff engineer at Le Studio Morin Heights in Quebec, Canada.

"One of the first projects that I worked on was Pilot's *Morin Heights* LP with producer Roy Thomas Baker. Before the band and Roy arrived, as it was popular in the day, my engineering partner the late great Nick Blagona and I set up the drums in the existing drum booth. When Roy and the band entered the studio for the first day of recording, Roy immediately said 'What are the drums doing in that tiny little booth? Bring them out here in the room and let's capture some of the lovely room ambience!' Roy then proceeded to play back a quarter-inch copy of Queen's 'Bohemian Rhapsody', and when I heard the huge drum sound on the track it was completely life-changing! This proved to be a revelation for me and would inspire me to use this method until this day."

Reading The Room

"In 1990 when we recorded Living Colour's second LP *Time's Up* that this song is included on, I suggested that we cut the backing tracks at A&M Studios in LA for the fact that Studio A was just incredible — absolutely the best room for drums I ever worked in. The main room was huge, being 39 by 38 feet, and 20 feet high. There were also two large isolation booths which we used for the bass and guitar amps. A&M had an incredible collection of vintage mics and outboard gear, not to mention the desk was the amazing AIR Montserrat custom Neve 4792 — the last desk designed by Rupert Neve — that was purchased by A&M after the disastrous volcano eruption on the island.

"When it came to tracking the band we placed the drums in the main room, using four room mics: two Neumann U87s on the largest mic stands in the building placed around 18 feet away in the far corners of the room near the ceiling, and two positioned up close to the kit. A Shure SM57 was used on the snare, Sennheiser MD421s on the tom toms, AKG C451s on the hi-hat and ride cymbal, AKG C414s on the overheads and an AKG D12 inside the kick. I also had a carpet rolled into a tube that was six feet long that we placed over and attached to the kick drum to keep it isolated. A Neuman FET U47 was placed approximately four feet back from the kit in the carpeted tube to capture the sound further away from the kick.

"So, what you are hearing on this song is the drums in their complete natural state, with no added outboard reverb whatsoever on them. It is only the room sound at A&M and it was a live take of the entire band with no edits. During mixing I would always split the snare and kick onto two separate channels and add very subtle compression, EQ and gating on the second channel, placing them behind the original signal for an added 'boost'. It was my intention to keep as much of the 'room sound' of A&M Studio A on the drums as possible. The *Time's Up* LP was mixed on an SSL 6000 G-series desk at Right Track Studios in New York City."

Classic TRACKS

Plastikman 'Consumed'

Richie Hawtin takes us through a seminal Plastikman album — and its unexpected sequel.

TOM DOYLE

As a world-renowned electronic music producer and DJ operating for more than three decades, Richie Hawtin has assumed various identities for his releases down the years, including F.U.S.E., Circuit Breaker and Concept 1. But he is perhaps best known for being a pioneer of minimal techno under the most famous of his trading names, Plastikman.

Hawtin's debut album using that pseudonym was the stripped-back acid electro of *Sheet One*, released in 1993. But it was with his third Plastikman album, 1998's *Consumed*, that as a producer he perfected a spacious, ambient and hypnotic dance music style that was entirely his own — not least with the polyrhythmic analogue pulsing of its near-12-minute-long title track, 'Consumed'. As Hawtin tells *SOS* today, this groundbreaking techno sound arrived as a result of his desire to move away from records made purely for the dancefloor.

"I was kind of bored of DJ'ing," he admits. "Around '97 to '99, there were so many derivative techno records coming out. I was feeling, like, 'It's all the same. When's something going to sound different?'"

For his sleek arrangements on *Consumed*, Hawtin was inspired by the Kraftwerk, Pink Floyd and Tangerine Dream albums that his electronics-obsessed father had played around the house as he was growing up, rather than the Detroit techno that fuelled his creativity as a teenager. "Those echoes are 100 percent there in the composition of my albums," Hawtin says. "The beauty of having an hour-long listening experience."

More recently, Hawtin has returned to 'Consumed', and its parent album, with 2022's *Consumed In Key*, a collaboration with Canadian musician/producer Chilly Gonzales that is described as more of a "reimagining" than a remix record, since it features the latter improvising piano parts over the original tracks. Gonzales has said that the shuffling techno grooves of the original *Consumed* sounded to him "like some science fiction version of jazz" and so he reacted to that idea in his playing, beginning work on the project before Hawtin was even made aware of it.

"Of course, *Consumed* is the inspiration for the project," says Hawtin, "because that's what Chilly decided to sit down and listen to and play the piano to. But I think somehow it's morphed into really its own unique space, which perhaps is a planet in the *Consumed* universe. But it's definitely not the planet where it first came to fruition."

Early Days

British-born Richie Hawtin grew up in Banbury, Oxfordshire and emigrated in

1979, aged nine, with his family to Canada when his father got a job working as a robotics engineer for General Motors. Living in Windsor, Ontario, just a few miles across the river from Detroit, he was in the right place at the right time to witness as a teenager the birth of techno in the US city in the mid-1980s.

Having been given an 8-bit Commodore VIC-20 computer as a kid, using it to initially program games before beginning to experiment with rudimentary sounds, he was comfortable with technology from an early age. As a teenager he began to listen to alternative and industrial music, his portal to dance music being the radio shows of Detroit DJ/musician Jeff Mills, who would mix records by European electronic acts such as Yello and Front 242 in with proto-house tracks.

Inspired, Hawtin began DJ'ing at clubs in Windsor, even booking Jeff Mills for one event, and meeting other older Detroit techno DJs and producers along the way, including Scott Gordon and Derrick May. His first experience of making dance music himself came when he hooked up with fellow DJ John Acquaviva and began messing around in the latter's home studio, which featured an AKAI S1000 sampler and — set to become a constant piece of kit in Hawtin's setups — the Roland TR-909 Rhythm Composer.

Funding a label, Plus 8 Records, from credit cards, the pair scored a European underground dance hit in 1990, 'Technarchy', as Cybersonik. At this point, Hawtin set up his own studio in his parents' home, naming it Under The Kitchen, and filling it with whatever gear he could afford.

"I was a victim of economics," he laughs. "I bought what I could find cheap. 303s. Of course, people don't think they were cheap, but they were 50 bucks. 808s, a Dave Smith Sequential Circuits [*Pro One*]. I didn't have Moogs, they were too expensive. I didn't have ARPs, they were too expensive. But I did have a [*Korg*] Wavestation keyboard, also a Dave Smith co-design, which gave me a kind of beautiful digital ambience."

Among the effects units that Hawtin managed to buy at the time were an ART Multiverb and a Yamaha SPX90. "The SPX had a dirty, great flange," he recalls, "and the ART had the gated reverb that nearly every clap that I ever used back then had on it."

From here, Richie Hawtin began releasing other records as F.U.S.E. and Circuit Breaker. These tracks, such as 'Substance Abuse' and 'Overkill', were made by him running sequences from various synths and drum machines chained together and driven by an Atari ST, while he live-manipulated and effected their sounds. He recorded the jams to DAT, before editing them digitally.

"My dad had an IBM computer," he remembers, "and there was this great software by Minnetonka called Fast Edit. It was really basic… you recorded it in, and you could just cut and throw away. It was complete destructive editing. But it was very, very fast.

"But I remember when I made the Circuit Breaker stuff, 'Overkill' was, like, two DATs long [*laughs*]. It took you longer to find the good part to edit."

Around this time, Hawtin was earning a reputation for his own warehouse-staged club nights in Detroit. But, even with his growing status as a DJ, he felt that his own records were becoming too conventional. "Like F.U.S.E. 'Substance Abuse'," he says, "it's a banger. But it's like, eight bars in, bring the clap and the acid line in. Okay, take a break, drop the kick, bring it back. Like, it's a bit more formulaic."

This led Hawtin towards the more angular and artful sound of Plastikman. "If you look at even the kind of thumping Plastikman tracks from the beginning, like 'Elektrostatik' or 'Krakpot', I think these are all pretty out-there tracks. They took a long time to develop."

New Studio, New Synth

In 1995, Richie Hawtin set up a new studio, The Building, and began to update his gear. He remembers having typical teething problems with the upgrade. "I moved in there and did what we all do… stupid shit," he laughs. "I sold my Allen & Heath mixer, a GS[*3000*] that I'd used to record all my stuff. I bought a Mackie mixer because it was kind of the cool hype thing, and everybody loved it. And I rewired everything to it, made one song, Plastikman 'Are Friends Electrik?', which I love, one of my favourite songs ever. But it was the only song that I recorded that sounded good. Everything else sounded shit. And it was just something to do with the bass undertones on that mixer.

"So, '95 was kind of me figuring out that new studio, selling my stuff, rebuying back another Allen & Heath [*laughs*], changing amps. I'd gone from a little Tannoy System 8 with dual concentrics to the System 15s with subwoofers. But then I realised that the subwoofers actually gave me too much bass. So after that was all done, it was '96 and that studio was just humming."

Track: 'Consumed'
Artist: Plastikman
Producer: Richie Hawtin
Label: Minus/Novamute
Year: 1998

Hawtin's main — and most expensive — purchase at this point was a Serge modular synth system. "That was a moment when I was looking at ARPs and everything," he recalls. "And I was like, 'Fuck this, if I'm going to spend a couple of thousand dollars, I should be buying an instrument that nobody else has, or nobody really knows.'"

Serge modular synths were originally built in the '70s, but Hawtin invested in the STS Serge rebuilds made by Rex Probe in the early 1990s. "I went and met Rex Probe down in Berkeley," Hawtin says. "I loved that it was built new, but it was still kind of old techie circuitry. It was something different. Now you say you have a Serge and people are like, 'Wow, that's so cool.' But back then, when you said you had a Serge, even the normal synth-head kind of looked blankly at you like you were speaking a different language."

Another important addition to Hawtin's setup was his acquisition of a Doepfer MAQ16/3 MIDI analogue sequencer (used by Kraftwerk and Jean-Michel Jarre). "I had one Doepfer MAQ16/3, an original black one. I loved it so much that I went out and I bought three of the newer editions. Those were usually used on separate lines. So that's nine different CV and MIDI 16-step sequences, which would be going, more often than not, via [*additional*] CV patches, to opening up Serge modular parameters." »

The Plastikman setup circa 1993.

Richie Hawtin's first project in his new studio was an ambitious one — a 12-inch release every month for a year under the name Concept 1. "Concept 1 was my departure from worrying about making dancefloor music," he says. "My challenge was to be able to make something every day. Actually, if you listen, they're very simple things, y'know, kicks and percussion going through Serge filters. There's even a couple of non-Roland drum sounds which were all coming from a Kawai XD-5 [*drum synth module*]. Sometimes it would be, like, 20-, 30-minute jams. The whole year of '96, Concept 1 gave me the confidence and experience for what *Consumed* would be."

Hypnotic Grooves

Even if the gear that Richie Hawtin used for Concept 1 and on Plastikman's *Consumed* was far more sophisticated than anything he'd ever previously utilised, he was still a stranger to multitrack recording. "*Consumed* comes from a time where many of us electronic musicians didn't have the knowledge of multitrack recording," he stresses. "Everything was recorded as a live jam to 2-track DAT."

Hawtin's daily working method during this time involved him firing up his equipment at the start of each session and trying to immediately find a hook.

"All my stuff is mostly created on a groove," he says. "Even if that's a melodic 303 line. I don't necessarily mean a groove that gets you dancing. It's something that is maybe more hypnotic. Whether it's a drum track or an acid line, you get something that you're kind of grooving to and listening to for 20 minutes, and you felt like it was two minutes. You just can't get enough of it.

"Then once you're in that moment, you grab it and record it, before you get tired of it. Before you think, 'Oh, could this be better?' and then you fuck it up and you've lost it. So, things happened very fast, as in I would write a track and record a track in the same day."

Central to his setup was the 909, by this point often used more as a sequencer than a sound unit. "Sometimes the kicks or other sounds were either coming from an [*Akai*] S3000 or from the Kawai [*XD-5*]. But those would have been sequenced from the 909. The 909 has a certain feeling in its rhythm. A big trick was using the 909 as a one-note sequencer. That came from the days when I was using the [*Roland SH-*]101 a lot and you would use a rimshot as a trigger out."

The shuffle function on the 909 was the key to some of the distinctive grooves on *Consumed*. "Well, y'know, I'd done a number of kind of house-y records, like Robotman 'Do Da Doo' [*1993*], and I always liked the 909 shuffle. It was usually shuffle at number three, so it's just slight. I was also using triplet delays to give a different kind of syncopation. Like, honestly at that moment, I was just looking to have the music sound and feel different than what I'd done before."

Hawtin was also particularly fond of the slight tempo fluctuations he felt he could hear happening between his hardware machines, believing they gave his grooves a more human feel. "This is the most difficult thing with going into computers as your main tool," he argues. "Because all the plug-ins basically end up — unless they've been made by really fucking cool people — all just following the same clock. And so everything is in time too much. Whereas when you're using hardware, sometimes these slight fluctuations — very, very, very slight, so something you may not even be able to measure — you can feel somehow. That gives your music a bit of its own heart and soul."

Hawtin also employed irregular step patterns, sent from his Doepfer MAQ16/3 to his Serge modular, which were to be a central feature of the rhythmic style of *Consumed*. "There would be a lot of cyclical five-note patterns, 12-note patterns that were all kind of moving and opening up filters and things at certain times," he explains. "*Consumed* is an album of feedback. Everything was cross-modulating everything else."

One Take

For *Consumed*, the tracks were recorded in the same sequence in which they appeared on the album, and all in a single day apiece. The one exception was its title track.

"I remember being so happy and excited with myself," Hawtin says of the initial creation of 'Consumed'. "The melody — bom-bom-bom, bom-bom-bom — was like a lot of my stuff where I don't really write melodies. I write rhythms which are kind of like melodies and that's just a tom/conga sound from the Kawai drum machine, which I slowly opened up and closed through a Serge filter with a lot of effects.

"It was all about all these reverbs and delays and you have this XD-5 line coming up with lots and lots of reverb. The Allen & Heath mixer had MIDI mute automation, and so I had a loop going in the Atari ST that was basically just eight or 16 bars. So, I'd let all the effects play, and then in one set instantly turn off the effects, and then eight bars later turn them back on.

"I was just listening to that and I had it cranked. I remember going upstairs, and downstairs [*in the studio*], every time the reverb turned off it was like, 'Wow, what a fucking powerful moment.' That was the last song I recorded for *Consumed* and it was the only song where I made it one day and I left it going all night. I left all the machines running, because I didn't want to lose anything, and then I recorded it the next morning to get the final cut. Everything was recorded live and in one take. There was a couple of edits here and there, but you had to nail it, or you were fucked. There was no going back."

By this stage, Hawtin was using Sonic Foundry's Sound Forge program to edit his DAT jams. "*Consumed* would have been edited on Sound Forge, but I didn't even like to use the playlisting function. It would just be like, 'Okay, this part is a good beginning, this middle part needs to be shorter, cut. And this ending over here is good... done.'"

Ten minutes into 'Consumed' there's a section where a Serge modular riff cracks up and distorts. Hawtin chose to leave the skronky sound in the final cut. "I just liked the way it was feeling," he says. "That was just too much resonance on a filter, and that's how it was recorded.

"It's funny, because when I was listening to it again [*recently*] and remixing, I did notice that that was really pushed. But that's the way it is. I will not say I'm the best mixer or engineer in the world. I've learned everything as I go, and I just have a particular ear and I try to make things that made me feel like, 'Fuck, that sounds great.' And there's not really a technical consideration. It's just like, 'Does that sound good to me?'"

Consumed In Key

For their 2022 *Consumed In Key* collaboration, Richie Hawtin and Chilly Gonzales worked remotely, their go-between being a mutual friend, Canadian producer/DJ Tiga. Gonzales had first begun by recording his piano parts over MP3s of the *Consumed* tracks and »

Chilly Gonzales and Richie Hawtin collaborated on 2022's *Consumed In Key* album.

Tiga then sent Hawtin three tracks for his consideration.

"The first three tracks were 'Contain', 'Consume' and 'Locomotion', which are kind of the core of *Consumed*," Hawtin says. "I thought they were respectful somehow and then extremely disrespectful sometimes because I was still listening to it like, 'Somebody is playing on top of my album [*laughs*].'

"But the big thing for me was that Chilly Gonzales I'd never met. I didn't know much music by him, but I knew the figure of Chilly. I knew he was well respected. He was a serious musician and he had done a lot of serious compositions. And I thought, 'If this guy sat down in his own time and recorded this, I need to respect and listen to it, and give it a lot of deep thought about where this is coming from, or where it may lead.'

"Then, more so out of the respect for Chilly as a musician, I said, 'Let's go with this.' I wouldn't say it was on the strength of those first demos. It was partly due to that, but it was as much weighted on who Chilly was and what I thought he would be able to deliver in the end result."

Hawtin's only stipulation was that he had control over the final mixes. Once he'd been sent pre-mixed WAV stems of Gonzales' piano contributions, he returned to his original equipment to try to blend the new parts in with the mastered stereo *Consumed* tracks.

"I hooked up the originals," he says, "and the only pieces of the puzzle that I could really remember, because I had some notes from the *Consumed* times. I had the effects patches for the Lexicon [*PCM90*] and the Roland SRV-330 [*Dimensional Space*] reverb and the [*Ensoniq*] DP/4 [*Parallel Effects Processor*]. So I hooked all those back up and started sending Chilly's piano through those.

"I'd never mixed piano or acoustic instruments and I told the guys, 'That can be my contribution... to give a bit of my sonic imprint into this new project,' because I mix and hear things in a certain style. So I sent everything through the effects, and I did a first version where it was, like, piano through reverbs. And it really sounded like the piano was in there from the very beginning. I thought it was amazing."

Unfortunately, Gonzales, who favours what Hawtin calls a "hyper-realistic piano style", didn't agree. "That became a learning point and even a struggle point between Chilly and I," Hawtin admits. "Because I was bringing his work into the world of digital effects and experimenting with some things which, in the end, I can see why they didn't work. They were necessary experiments for me in the mixing process, but they were things that Chilly strongly disagreed with at the beginning of that process [*laughs*].

"When the mixing happened, I really went to the other extreme. As I listened more and more in solo to Chilly's work, I started to learn about Chilly. And that, in the end, turned my effects idea upside down. There are still some residual effects that I left in there, although it's very subtle. But the process was really exciting, to learn by listening and find out how these two sensitive frequency arrangements came together in a harmonious way with basically nearly zero verbal communication."

Of the new versions, Hawtin singles out 'Locomotion (In Key)' as his favourite. "That always had a great momentum, but the way Chilly played the piano... it's like a steam engine. It's like there's these two trains next to each other... one's nudging and then the other's nudging. They're competing with each other, but they're on the same mission and on the same path. There's this duet going on. I think that's really a beautiful moment."

Studio Evolution

In terms of his studio setup these days, Hawtin is moving further away from hardware and more into software, Bitwig

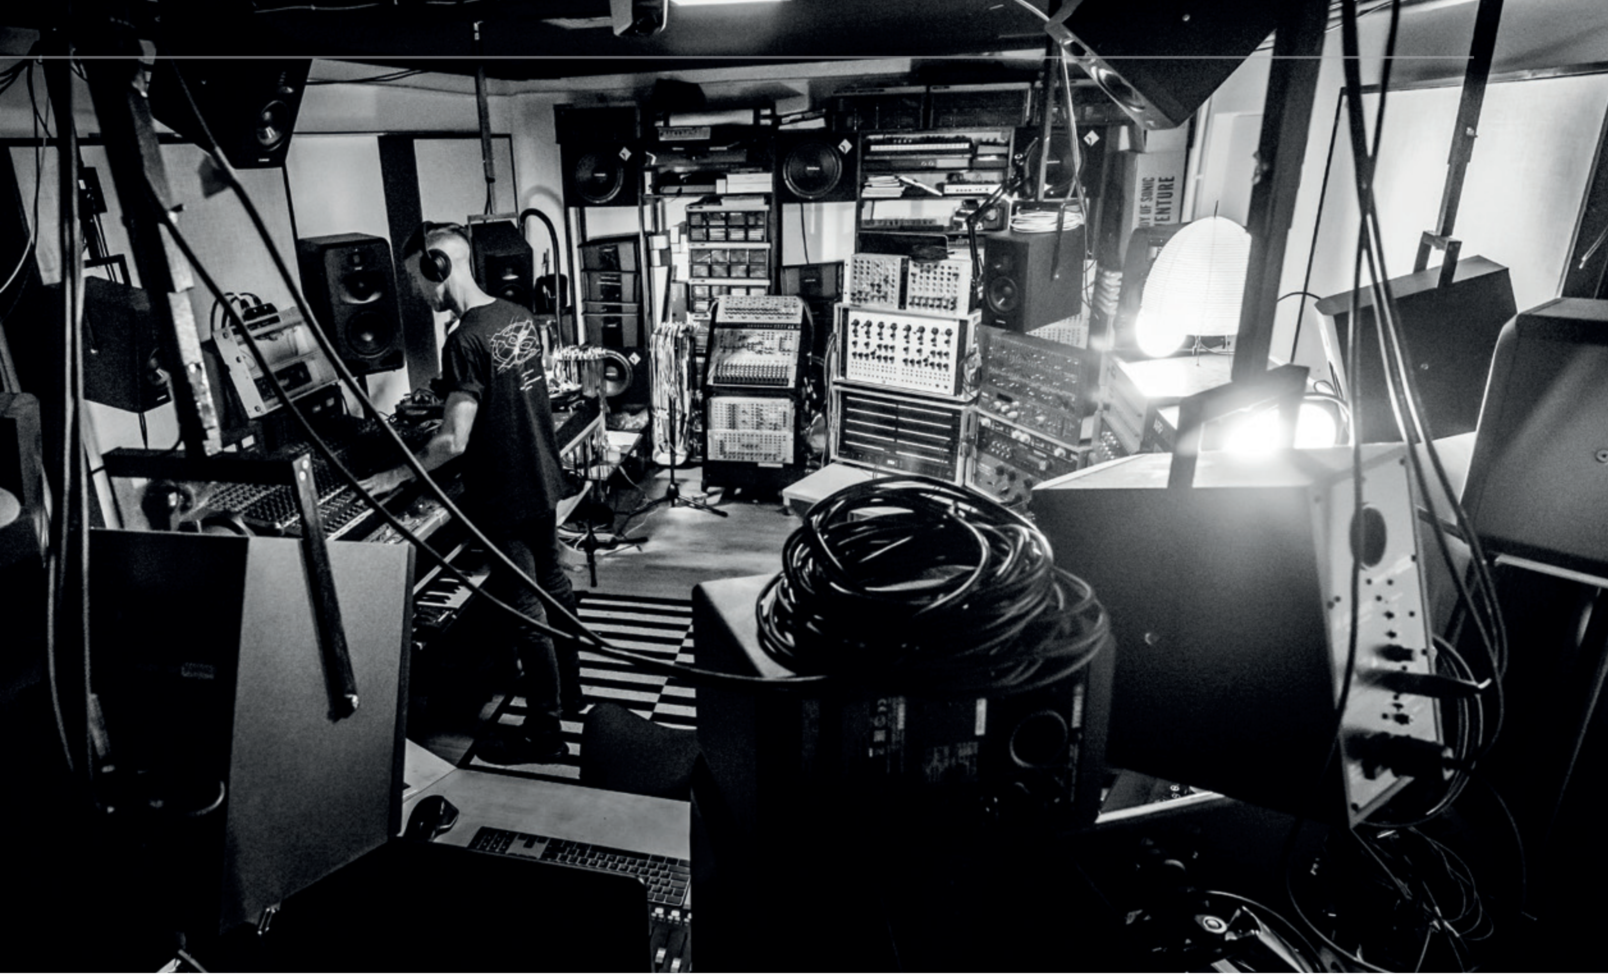

Despite being mainly computer-based today, Richie Hawtin's studio remains a home for his impressive collection of hardware.

Studio being his preferred DAW. In fact, he recently worked with Bitwig on custom scripts to enable him to better interface his Novation Launchpad Pro and Akai Fire with the software.

"Last year and the year before, as I moved from a mostly analogue and sometimes hybrid studio setup to a complete digital setup with Bitwig, I spent a couple of months researching and playing with my original setup with the 909s, and all of the things that we've spoken about... like using a 909 to sequence.

"I developed some custom scripts that gave me some of those flavours and some of those functions that I didn't find on any other controllers for Bitwig or Ableton. So, I'm hoping that these will also help other musicians find their own new or potentially better ways to connect with the computer and feel that there's a nice human interface."

Not that Hawtin has ditched all of his other hardware, though. "No, no," he insists, "it's all in my studio because I do find it quite alluring. And I do like all the lights and the warmth that it gives my studio in the winter months. Just sitting in front of a computer is still a little bit uninspiring for me. But I think you'll see more custom script design."

Similarly, Hawtin co-designed the Model 1 analogue DJ mixer, along with engineer Andy Rigby-Jones. Made with what he calls "a DJ/producer mindset", it enables the user to live-sculpt EQ moves using the high- and low-pass filters that have increasingly become essential to dance music.

"It was important that the Model 1 transcended just being a mixer," he says. "Because the original intention of a mixer was just to mix some signals together. But I think DJs took the mixer idea and helped develop it into what became an instrument. And with that in mind, I wanted to make, with Andy, something that was an instrument designed as a mixer rather than a mixer hoping to become an instrument."

Elsewhere, Hawtin has been busy exploring other synthesis platforms. "There's a really interesting project called Audioglyphs, which is a web interface into web audio components to make generative music out of a kind of a node-based graph system, which I'm learning and experimenting with now. And I'm also teaching myself the new synthesis engine in Unreal 5 called MetaSounds. I don't know where this will take me. But I think it's very interesting to experiment and learn new techniques.

"One of the most exciting times for me," he adds, "was when I was making *Consumed* or *Sheet One*, and I didn't know all the things I was doing. I was still learning as I was going, and there were surprises happening. People have asked me, 'Hey, why don't you just plug your 909 and 303 in?' I can't do that. I love those machines. But I need to feel like I'm sitting in front of a 909 and a 303 again for the first time. And so the only way to do that is to not sit in front of a 909 and 303 again."

The Future Of Plastikman

Richie Hawtin — who has now moved to Portugal after some years spent in Berlin — has periodically revived the Plastikman name down the decades, most recently with the 2014 album *EX*, recorded live at the Guggenheim Museum in New York. For him, Plastikman seems to represent a freewheeling creativity that he finds himself drawn to again and again.

"All the research and development I'm doing right now is based upon the assumption and the hope that I will be deep into a new, multi-faceted Plastikman album. I think Plastikman has always had more sort of freedom in there. It's just about letting something happen freely, spontaneously, and, in a way, introvertedly. It's like, 'Where's my head at?' and I try to record that.

"That's probably as close to a description of Plastikman as you're going to get."

TalkBack

Owain Fleetwood Jenkins

WILLIAM STOKES

Deep in rural Pembrokeshire, South West Wales, engineer and producer Owain Fleetwood Jenkins has built a paradisal residential studio in an abandoned chapel. StudiOwz (a portmanteau of 'studio' and Owain's childhood nickname 'Owz') is stuffed with lovingly maintained vintage instruments and equipment, pride of place going to the Cadac J-Type console.

"I first got into music through work experience in school, actually," says Jenkins. "I went to a local recording studio in Tenby, and then, kind of, never left that studio! Before I went there, I had no idea that that's how music came to be heard on the radio. Like, I hadn't even figured out that bands go to a studio, and that they do this thing called 'recording'!"

StudiOwz has hosted artists and producers from all over the world, from acclaimed singer-songwriter Emily Barker to Cate Le Bon and erstwhile Keane frontman Tom Chaplin. It has also welcomed many local musicians for sessions conducted entirely in Welsh, Jenkins' first language. "It's just all been word of mouth, really," he explains. "That's how the studio has grown."

At the moment I can't stop listening to

I always listen to vinyl in the house, to be honest. There's an artist called Erin Rae, actually. *Putting On Airs* is the album. I actually came across her while doing live sound in St Davids. It was one of those shows that just blew me away, so I ended up buying the record. It's an Americana, singer-songwriter album, but it's got really cool production as well. There are really interesting percussion parts in there. Like, the use of timpani and stuff, instead of toms. I love it when the orchestral world gets involved a little bit... but not much. Not enough to sound orchestral. It just sounds great! There are a few albums that have that kind of thing in it. And I love that. I'd love to buy some timps, actually, but they're huge and you'd only use them every now and then.

The project I'm most proud of

Probably Jodi Marie, who's also my partner. She trusts me with her music, a lot. Her latest album, *The Answer*, actually took quite a long time to record. It started off all live, and then we went in and did overdubs and generally replaced the vocals. And you can really hear that in the recordings, I think. We started in my old studio and finished it off in the new studio. It's really interesting to work on the same record in two different studios and see how the room changes things, even though you're using all the same gear.

Loads of spring reverbs were used — I'm a big lover of old springs! And just, like, misusing old analogue gear: driving spring units far too hard and then turning off the spring, so you get distortion from absolutely abusing a piece of analogue equipment. I love that. I had lots of fun with that on that record. I got into saturation and vintage springs and tape echo. Which is a fun world to start exploring deeper. It's a very '60s or '70s-sounding record, which is mainly the music I listen to, but I also tried to get that little hint of modern production in there so it stands up today as well. I've found that quite hard in the past, to get the balance right. It can either sound dated, or just a little bit shit.

The first thing I look for in a studio

The microphone cupboard! I just love the classics. Old valve mics, old ribbons... it's quite a boring answer, really! But U67s, the old 87s, the old 47s. I don't have all of those here at the studio. And that's probably why I'd go looking for them. They just get you so much closer. The closer you are to the start of the chain, the more important that piece of equipment in the chain is. Recently, I've started trusting in the microphone a little bit more, rather than immediately grabbing EQ. Because we love these microphones and how they sound. So, if it's in good condition, we should trust that the tone character of it is 'correct'. And we need to get more on board with that sometimes, rather than, you know, grabbing for 10kHz straight away and giving it that modern sheen, or whatever. Just trust the microphone. That it sounds great. Microphones are important to me.

The person I would consider my mentor

That would be Bruce Campbell, who I went on work experience with. To a certain degree he told me everything I know, really. And from an early age, especially in live sound, he was the person that taught me that It's not all about the bass drum. Because growing up, you go to gigs where all the engineers are trying to get that bass drum as loud and punchy as possible. They totally forget about the vocals and the whole point of the music. He taught me to hear what's important in the music, what's carrying the melody in the song. That was a big lesson early on, I think. I still hear it with young sound engineers now: you walk into a big venue and that bass drum is absolutely pounding. And, you know, sometimes I quite like that as well! But it's a shame when they miss the point with everything else.

With my record collection, I never really notice the bass drum on those records I love, because you're listening to how the acoustic guitars are double-tracked, or the vocals are right in your face. I don't feel like you fall in love with a song for the bass drum. So I don't know why some engineers spend so much time on it. Learning what's important in a song was a lesson I learned at, like, 15. Which was invaluable for me.

My go-to reference track or album

I definitely have an album for testing speakers. It's Blake Mills' album, *Heigh Ho*. It sounds like he's done lots of varispeed stuff with the drums, the players on it are

incredible and the recording is so good. And the low end is insane. Really interesting sounds, like you know, there's a great, really soft, close upright piano sound. Crazy reverse delays and things. A huge range, quite an unusual range in the arrangement of the songs, too. He'll like, hold out right till the end to like introduce a drum kit or something. Throughout a five-minute song, you'll have been waiting for the drums to kick in, and they'll kick in for the last eight bars — and then it's finished! But it feels incredible when it happens. And the stereo imaging on it is really insane too. It gives me chills when I listen to it on the big Tannoy Arden monitors in here. I love those speakers so much. They're from 1976. It's as if we've come so far in speaker design since then, yet we haven't got much further in their audio quality!

My top tip for a successful session

Good communication between everyone involved. Who's on the session, why they're there, what their role is — having clear roles for people. And then, the communication between artists and producers and engineers, between artists themselves. Being good at communicating when you're in the creative process is really important, I think, for how much you get done, you know. Keeping a good vibe. Good communication. When people communicate badly, they kind of tense up and lock up. Or they start having a grudge against each other. It's that kind of vibe that comes before arguing. It totally kills any creativity.

The studio session I wish I'd witnessed

I think it would be a home recording, actually: Paul and Linda McCartney's album, *Ram*. It's just incredible, isn't it? It's like a lesson in production. Just, all the sounds. The way he uses his voice as an instrument, as well. I would have loved just to see them doing takes. Like, he never uses his voice in the same way on different songs. It's like he's in character. it just sounds like such a fun album to make. I bet it was really fast paced and just, ultimately, fun. Often in the studio these days it can be so slow: another take, another take. It's important to not slip into that kind of monotonous 'studio-takeage'.

Originally, I had this theory that all the odd-numbered tracks on *Ram* were the best songs. But after listening to it over 100 times, I've gone against that theory. They're all just crackers.

The producer I'd most like to work with

I think I've worked with him actually. Which is good! Ethan Johns. When I was growing up, I listened to lots of records that he'd produced and played on. And they're still records I listen to now and that I think sound good. We worked together on the new Tom Chaplin record. They worked at four different studios. I think they did The Church first, then Real World. Then they came here, and then I think they went to Angel.

We did some drums, some acoustic, some electric, loads of Mellotron... Ethan's got, like, nine or something. It was Glyn Johns who recorded the original recordings that went on the Mellotron [*tapes*]. They were recorded at Abbey Road, apparently. And actually, it was a digital one that Ethan brought with him. It was really interesting, watching him play it. The way he performed with it. That was really cool. He was like, "You can't think like a keyboard player. You have to think like a French horn player, or violin player." The movements, and using the volume for swelling. It sounded so good.

The part of music creation I enjoy the most

The recording process! Massively! That's where all the fun is. You know, you've got a clean page to fill in. The foundation of the sounds. That's so much fun, I think. I like imagining and aiming for the end result, right when you place your first microphone. It's just a feeling I get. I can't describe the feeling.

The advice I'd give myself of 10 years ago

Just keep working hard. Because there's always going to be someone working harder than you, and better than you. So yeah: head down. I remember as a child, my dad used to say to me, "Every day is a Monday." There's no such thing as a weekend, or a day off. I grew up on a farm. And on a farm, you know, every day really is a Monday! Aim for things and just go for them. I wanted a studio by the time I was 30 when I was about 17 years old. And I'm sitting in it now. The only way I got that really was by aiming for it and keeping my head down.

SAMPLE LIBRARIES

Bunker Samples

Bunker Strings Vol 1 & 2

Kontakt Instruments

★★★★

Most people would prefer not to spend time in a bomb shelter, but for violist/composer Nikolaj Moeller Nielsen a Copenhagen WWII bunker came to be a home from home. Originally built as a civilian bomb shelter in 1945, it was later converted into a recording studio and rented to Mr Nielsen, who spent two years happily recording string samples there. The musician fondly recalls: "I loved that place so much... it was my creative heaven, so I decided to name my company Bunker Samples after it, since it's where it all began."

With Bunker Strings Nielsen's aim is not to pursue ultimate realism, but to explore the possibilities of real performances within the virtual domain. Its two volumes feature unusual articulations played by Nielsen himself on a collection of violins and violas, along with cellos and double basses played by friends. Each sample was multitracked nine times and mixed down to create a virtual ensemble of variable size controlled by a 'density' fader — there's also an option to use round-robin samples as an extra layer, thus doubling the band's virtual player numbers. Orchestral purists would throw up their hands in horror, but I liked the resulting full and flexible sound.

Bunker Strings Vol 1 consists mainly of interesting variations on the tremolo theme. Instead of the standard rapid back-and-forth bowing, the players perform 'pizzicato tremolo': this shower of repeated plucked notes produces an excited, bubbling murmuration, a cheerful texture which would make a great, optimistic curtain-raiser for an epic sci-fi movie. 'Drumstick tremolos' have a more intense, skittering flavour, while 'ricochet tremolos' up the ante with repeated agitated bow bounces, a superbly threatening effect when played as cluster chords. For rhythmic work, the dynamic 'drumstick shorts' range from light spiccato brushes to heavy col legno-style hits.

Volume 2 kicks off with two pad-friendly textures. The gently mobile 'sul tasto pulses' sound soft and serene, while 'tremolo bursts normale' is more stately and dramatic. The latter's sul ponticello variant introduces nervy horror-film edginess. The 'plectrum tremolo' style can evoke the sound of massed Mediterranean mandolins or a swarm of locusts, depending on how you play it; single-note plectrum plucks represent a return to normality, with unpitched hand-damped plectrum strums adding some welcome rock attitude. Also included are bonus synth patches which include some exquisite pads.

All articulations are separately performed by violins, violas, cellos and basses, with well-programmed ensemble patches providing playable full-range options throughout. I enjoyed these samples and found them to be an inspiring alternative to conventional string collections — it's good to know something positive can come out of a bunker! The two volumes (3.72GB and 4.91GB respectively) both require the full version of Kontakt 5.6.6 or later and will not work in the free Kontakt Player.
Dave Stewart

Vol 1 $59, Vol 2 $79, both $109.
www.bunkersamples.com

Sound Dust

Drift 001

Kontakt Instrument

★★★★★

There can be few instruments with a simpler interface than Sound Dust's Drift 001, but then that was the design ethos behind its creation. This Kontakt instrument requires the full version of Kontakt 5.8 or later and is based around 5GB of multi-samples arranged into 11 hybrid articulations. There are 100 snapshots based on these samples, but despite the obvious simplicity of the controls it's possible to coax a surprisingly wide array of sounds from them.

There are only two tabs on the screen, one for Drift and another, labelled RTFM, which takes you to the manual (which is just a copy of the panel with text descriptions attached to each control). The down arrow at the left of the control row selects a Snapshot, of which there are 11 defaults and several variations on a theme to make up the 100. There are no legends attached to the controls unless you hover the mouse above the knob, presumably to make the minimalist interface seem even more minimalist. The first knob, Grist, adds various degrees of dirt to the sound and as you turn it the distortion changes from sounding fairly analogue to somewhat digital and gritty. Next along is a single control that morphs through a range of ADSR permutations so you can get anything from a piano-type attack to a slow attack, slow release sound or even a pseudo reverse sound. The third knob, which is the largest, is modestly named EQ, but this is no simple tone control as it again morphs between a number of EQ characters, from deep and smooth via fizzy and bright to strongly resonant. Careful use of this control makes it possible to extract a range of quite different timbres from a single sample set.

Finally we get the Space knob, which ties in with the selector to its right offering a choice of Small, Medium or Large Spaces as well as Reverse and Spring. Each of these choices has its own sub-menu of variations such as Shimmer, Swell, Cloud and even Endless. These offer a lot of tonal and textural variation so when added in large amounts they can completely change the character of a sound.

Most of the core samples are generated using a combination of synthesizers and effects, and range from a woolly-sounding 'Bad Piano', orchestral strings and dense synth layers to abstract pads and electronic stabs. The same sample can produce warm sounds, aggressive sounds, smooth attacks, percussive attacks, pulsing tremolo and cavernous washes, so don't let the simplistic interface put you off — I managed to create several useful and very different-sounding string pads just starting from the default Sielchestra Snapshot. There's a lot to like here, especially if you like your gratification to be almost instant. *Paul White*

$45
www.sound-dust.com

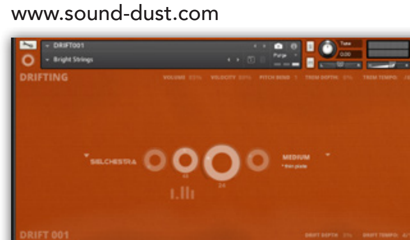

Osterhouse Sounds

Pathfinder Violin

Kontakt Instrument

★★★★½

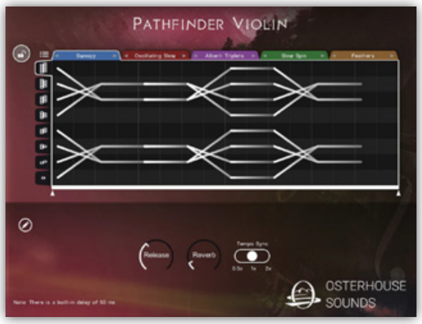

Arpeggiators have long proven themselves to be a useful addition to the synth World, but with the development of programmable arpeggiators, the lines between step sequencing and arpeggiation have become blurred. This concept hasn't always been one to make the leap to acoustic instrumentation, in no small part thanks to the intricacies of acoustic realism. Pathfinder Violin straddles this divide, in quite an inspired way. The basic concept is much like an arpeggiator; play a chord, and Pathfinder Violin will react to the notes that you play, corresponding to selected patterns or shapes.

Playing host to 2.6GB of compressed sample data, the 9000+ solo violin samples are organised in the background, according to the selected 'path'. These paths can be adapted and edited, in order to handle as many notes as your fingers can muster (well, eight in total) creating lines and patterns which sound natural and effective.

This instrument is very visual in its concept; the left of the instrument window indicates the notes being played. If you play a triad, the three notes are named, and ordered from bottom to top. If you add a fourth note mid-cycle, that note will be introduced on the next playing of the cycle loop. The path cycle defaults to a four-bar loop, but this can be altered to suit a shorter number of bars or note values, down to a single crotchet. There's also a Tempo Sync element, which will halve or double the cycle's speed, and hence length.

A large number of preset paths have been included with the instrument, ranging from plain sustained notes to arpeggio-based movements. The latter include the ability to move swiftly or legato-glide, with a highly effective sense of portamento. The net result is that the playing of a chord will yield beautiful polyphonic movement, which can be adapted by altering notes or switching up the chosen path via keyswitching. It's a clever concept, and one that yields results quickly and impressively. Editing may also be taken further, thanks to the editing mode where you can build and play your own paths. This can be incredibly exacting, as you prescribe exactly which note from your chord heads where, along with the ability to truncate notes and phrases.

The slower-moving part leading is very realistic indeed. There is often a sense of feathered bowing, which sounds exceptionally natural, although I did notice that with some of the faster phrase movements, there was a slight sense of note granulation. There is a built-in 50ms delay, which will require a degree of on-screen nudging; *de rigeur* for anyone working with orchestral samples!

Pathfinder Violin is a very desirable package, delivering quick and effective results for solo violin interest at an attractive price. *Dave Gale*

$81.63

www.lootaudio.com

Spitfire Audio

Fractured Strings

Plug-in Instrument

★★★★

Built in collaboration with Hans Zimmer's Bleeding Fingers Music collective and featured in the BBC's *Frozen Planet II*, Fractured Strings features a chamber ensemble of eight string players, along with patches featuring violin and cello soloists. Hosted in Spitfire's own dedicated plug-in, it has 12 mic and three macro mix options, and also includes the useful Evo Grid featured in other recent releases. Aesthetically, the instrument is quite focused: from glittering and refracting textures at home in a Terrence Malick or Kieslowski score, to the deeply expressive and stark intervallic work of an Arvo Pärt piece, the deeply moody nature of this library is worth exploring.

Opening Fractured Strings in the minimalist Spitfire host application, one is greeted with more than an aesthetically pleasing frozen background: Spitfire's new 'scale mode' is also there. This innovative and simple GUI allows users to conform performance patches to a scale or mode. So for instance, a patch such as 'Cello Rotations Up' — which features performances of icily expressive rising intervals — can be set to play all of its rising intervals within a customisable scale. While the mod wheel still controls dynamic levels, users can then vary the interval being played by using velocity, allowing the same chord pattern to sound from close to farther intervals without making a single change in fingering. Aside from being highly expressive and convincing, it is also a fun workshopping and aleatoric experimentation tool. The dynamic layers are very evocative, moving from warm pianissimos to icy borderline ponticello-like fortissimos. The musical intent and intensity of the performances is palpable, and with a bit of patience the results can be breathtaking.

The techniques themselves are quite varied. We begin with 42 types of violin section 'fractures', where different chord types and evocative flourishes can be controlled with a speed parameter. There are 54 violin and 51 cello techniques available: these include statements (shorter phrases which move away and return to the same starting note), natural harmonic patterns moving away from a central note, rotations (this writer's favourite, which are tempo-sync'ed small phrase performances moving up or down), and seven varied trill options called 'Dispersals'. There are also declamatory shorts (norm, harmonic and sul ponticello). Finally there are highly expressive 'Reaches' (a rise of a seventh with a passing tone), and several pizzicato strums. In a potential nod to the aforementioned Arvo Pärt, there are also nine available arpeggio performances for both solo violin and cello.

One wonders how Spitfire can continue to release the volume of projects that they are while maintaining their quality. And yet with Fractured Strings, they once again bring a compelling library full of distinct personality into the virtual instrument world. This crystalline effort comes highly recommended for film composers, and those just looking to explore new sonic vistas. *Mark Nowakowski*

$299

www.spitfireaudio.com

Audio examples of this month's libraries are available at www.soundonsound.com.

QUESTIONS ANSWERED, PROBLEMS SOLVED

Q Is gain staging important when calibrating my console and monitors using pink noise?

I've confused myself into a possibly needless tailspin! While mixing, I usually monitor with the control room level pot on my analogue console set between about four and five (out of 10). This usually equates to around 76dB SPL (C-weighted, slow response) from my main studio monitors. It's a comfortable yet loud enough level to work with, and my mixes have generally seemed consistent in terms of overall level.

I should probably have left it there! The trouble began when I decided to try to measure the output from the main monitors when playing pink noise from the DAW, and I think my confusion may be due to my convoluted signal flow and hybrid studio setup.

I'm sending one track of mono uncorrelated pink noise to one input channel of a Soundcraft Series 600 analogue console. The channel fader is set at unity, with no gain, and it's bussed to the main mix with the master fader at unity. The signal goes from the main mix to an RME ADI-2FS interface, which converts the stereo analogue signal to a S/PDIF signal and routes that back to the interface and a DAW stereo track running a Sonarworks plug-in. The signal then goes out through the D-A converter to a two-track input on the console, and that's what I am monitoring. It's a complex setup, but it allows me to hear the output of the console with Sonarworks 'correcting' for the deficiencies in my room.

Because of that signal flow, and the different maximum input and output levels of the DAW/interface, the analogue console and the RME converter (as well as the 7.3dB level loss from the Sonarworks correction software), I'm getting confused about what it's important to measure and where. For example, -20dBFS pink noise from the DAW doesn't come back at the end of the line and hit the stereo track running Sonarworks that I am monitoring at -20dBFS. It's lower — but I can increase the output of the pink noise from the DAW above -20dBFS so that it does come back in at -20dBFS, and then measure the SPL from the monitors at that point.

But is that what I should be measuring? If I do this, then, coincidentally, I'm back at around 4.5 on the console's control room pot again, and measuring around 76dB SPL.

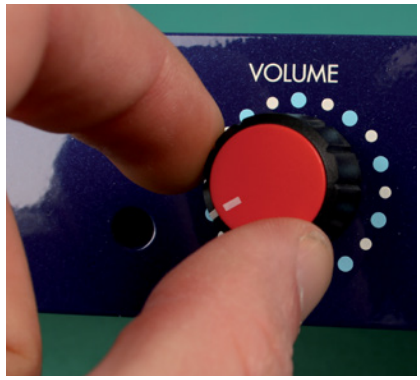

On top of all that, I've noticed that when I am getting -20dBFS out of the Sonarworks track and monitoring that on the console, the console's VU meters hover around -12 or -13 VU. Does that seem correct? In other words, it's as if 0VU on my console corresponds more to around -10dBFS than -20dBFS in the DAW. For what it's worth, I just sent a sine wave down the same path and when the VU meters read zero on the console I'm measuring that output at around 1.2V RMS, so that seems correct.

***SOS* Forum Post**

Hugh Robjohns, Technical Editor

I think 'tailspin' is the correct description — you have several different problems in the gain structure of your signal path, and that's clearly confusing you!

76dBC is a perfectly acceptable reference listening level in a small room (equating to -20dBFS in your DAW). So that's great. Unfortunately, having the monitoring volume control on the desk around 4/10 or 5/10 is not so great, as in that region you'll have stereo mistracking and poor level resolution issues. It would be far better if you could reduce the sensitivity of your monitors to allow that control to reside around 7/10 when delivering your comfortable (76dBC) reference level from the speakers. If you can't turn the speakers' input sensitivity down, you could buy or make inline attenuators to solve the problem, or perhaps increase the attenuation using Sonarworks.

As for the rest of the alignment process, all that matters is that -20dBFS on your final mix bus equates to your acoustic reference level. It doesn't matter what level you're pushing through the analogue console at this stage, nor what level the pink noise generator has to run at to achieve the desired goal. That said, you'd make life much easier if you could bypass the console part of the input signal path for this alignment process.

You can't measure the true RMS level of pink noise on a DAW's sample-peak meters, so it's better to use a file that has been properly designed. I use the Blue Sky calibration files, which I know to be correct: https://abluesky.com/support/blue-sky-calibration-test-files. All you need to do is load the file into your DAW and play it directly through your output monitoring chain (and Sonarworks) into the console's monitor section and your speaker, and adjust the speaker's sensitivity accordingly with the monitoring volume control at it's reference position (ideally around 7/10, as mentioned). You may well find that using a band-limited (500Hz-2.5kHz) pink noise file gives more reliable results with your SPL meter as it neatly avoids problems with LF room resonances and HF splashes from hard surfaces like consoles.

Once you've established a suitably calibrated relationship between the DAW's output and your speakers, you can do away with the pink noise, turn the monitors down, and align the rest of the signal path using sine-wave tones. (I prefer working with 400Hz to 1kHz, since it's less painful on the ears, though the BlueSky files only offer 1kHz.)

Then work through your interface I/O level settings and your console input gains to achieve a unity gain signal path all the way through the chain. This may or may not be possible, depending on the options in your interface, and what your console's VU meters are aligned for. In normal professional circles 0VU should equate to +4dBu and -20dBFS (meaning that 0dBFS = +24dBu), but there are various other 'standard' calibrations that are all perfectly workable. You mentioned an RME converter, but only the very latest ones can cope with +24dBu. Most of the older ones have an option for 0dBFS = +19dBu, and are designed to conform to the European broadcasting standard alignment where 0dBu = -18dBFS (or +4dBu (0VU) = -14dBFS). Anyway, check the alignment of your console meters, decide on a suitable alignment between the console meters and the DAW, and adjust gains and sensitivities accordingly so that the headroom margins through the whole chain are sensible and everything works as you expect it to.

SUBSCRIBE & SAVE

TABLET | WEB | REPLICA | PRINT

- Go **DIGITAL** - get 12 monthly Tablet issues + all locked Web articles + monthly Full Issue PDF download the day the issue goes live.
- Go **PRINT+DIGITAL** - get the best of both worlds and save over 50%.*

1 ICONIC MAGAZINE
6 GREAT WAYS TO READ IT

TABLET - PHONE - WEB - PDF - REPLICA - PRINT

All this for less than the price of a Coffee+Muffin each month!

iOS / ANDROID

Every Digital sub includes a 5-device licence to our App: mix and match across iOS/Android Tablet and Smartphone.

ONLY READING SOS IN PRINT?

Choose our full Print+Digital bundle for unlimited 24/7 access.

www.soundonsound.com/subscribe

** The 50% discount offer is based on the US retail price of the Print magazine, Tablet and Web subscription when sold separately. Subscriptions can be cancelled at any time and you will receive a pro-rata refund for all outstanding issues.*

SOUND ON SOUND

PODCASTS

Three dedicated channels each with their own series of episodes to keep you entertained throughout the month.

The **Recording & Mixing** channel takes the practical approach and will keep you inspired with expert hints and tips.

The **Electronic Music** channel is for everyone interested in synths, samplers and the world of electronic music.

The **People & Music** Industry channel features the great and good in engineering, production and manufacturing.

Follow our channels by subscribing to the shows on Apple Podcasts, Google Podcasts, Spotify, Amazon Music or wherever you get your podcasts.

Listen on Apple Podcasts | Listen on Google Podcasts | LISTEN ON Spotify | Listen on amazon music

Check out our website page for further details

www.soundonsound.com/podcasts

SOS SHOWCASE

David Carson (Regional Sales Associate): david.carson@soundonsound.com


The Aeolian®
Toll-Free Number:
888-765-2900
AcousticsFirst.com®
The first name in state-of-the-art sound diffusion.™


JACK OF ALL TRADES
FatKeys™
Quad VCA
SOUND
SEMICONDUCTOR
soundsemiconductor.com
SSI2164


Kludge Audio
506 Equalizer
THE ANALOGUE EQUALIZER FOR THE 96 KHZ AGE
http://www.kludgeaudio.com/500


DO IT YOUR WAY
"BUILD YOUR OWN" WEB PAGES
Full control for design, editing + more.
REDCO AUDIO
www.redco.com
1-203-502-7600


CRAS
BE IN DEMAND
Your career in audio is only 11 months away!
CRAS GRADS MAKE IT!
Over the past 3 years, CRAS grads have worked on hundreds of GRAMMY nominated recordings and multiple AMA, Billboard, Emmy, ACM, and Oscar winning recordings.
Struggling to create a recording career? Want to be in demand but don't know how to get from where you are to where you want to be?
At CRAS, our 1,200+ industry connections can get you in the door, and our one of a kind education will expose you to 13 different studios, a 6,000 sg/ft live sound venue, the latest professional audio gear, software and more.
Whether you're already into music and sound or looking to start from scratch, CRAS is the 11 month audio engineering school that provides the education and industry connections needed to succeed and be in demand.
CRAS.edu | 1.888.365.8312
Tempe Campus: 2300 E Broadway Road, Tempe, AZ 85282

SOS WHY I LOVE...

EARLY ADVENTURES IN RECORDING

JONN SAVANNAH

I suppose I must have been about nine years old when my father brought home a Dictaphone from work. He used to dictate letters to his secretary, recording his voice on to a brown, spinning, flat disc. He knew that I might be interested in having some fun with it, and he was right. I immediately got to grips with the simple way it worked and set about some field recordings, the obligatory burps and goonish voices. My friend Stephen came over shortly afterwards and we tried recording the TV, switching channels to achieve pre-razorblade, hard-cut editing. I still recall guffawing at some of the strange chance links of phrases matching across various channels.

Moving on to a cassette machine soon thereafter, I read the manual to discover that there were three separate heads: record, play, and erase. By some leap of imagination, I figured out that, if I could stop the erase head from doing its job, I might be able to layer more than one recording. By sliding a piece of paper in between the cassette tape and the erase head I actually achieved my goal. This was a primitive precursor to the method of bouncing between Revox tape-machine channels that I used years later. One run through recorded a piano track, a second run layered acoustic guitar and vocals. A final run allowed me to add some flute. The degrading of the quality of each recording was spectacular! The piano sounded underwater, vocals cottonwool-mouthed, acoustic guitar all but gone, and the flute warbled like an old Mellotron. But this was multitrack recording. On a cassette machine. In my bedroom. Pre-dating the Portastudio by over a decade. I was, to say the least, ecstatic. I proudly played the result to my sceptical Scottish mother:

"Who's playing the guitar and the flute?"

"Me, mum..."

"No, you're playing the piano and singing..."

No amount of explaining would make sense of, or clarify the situation — she thought I was messing with her.

Affordable multitrack recording systems were still a few years off. As soon as the Tascam Portastudio appeared I nabbed one. With rudimentary EQ, and a simple track-bouncing system, this double-speed cassette system allowed me to spread my wings — if only a little. Using a guitar pedal for some phasing effects, the reverb tank of a Peavey practice amp, and a Roland Dr Rhythm pedal for the drums, I was able to craft a track polished enough to reside on the B-side of a giant half-million-selling single hit I had in Germany under the name of 'The Catch'. And no-one batted an eyelid! No one in the record company questioned the quality of the recording. And it still sounds pretty good to this day.

I wonder how many other 'B' sides have exhibited such a complete lack of recording equipment pedigree. Hence the maxim: never say never!

AD INDEX


NOW SHIPPING
MOTU
M6
48V
MON
A/B
PLAYBACK
1-2 3-4 5-6 1-2 3-4
IN
OUT
3-4
DT 770 PRO

MIX WITH THE BEST

"I have been a reader of *Sound On Sound* for as long as I can remember. The amount of enlightening information, the intelligence of the writing, the beautiful print quality, and the personal stories make it as useful as the audio tools it covers. I'm a huge fan."

Jacquire King
Engineer, mixer, producer, Grammy Award winner (Kings of Leon, Norah Jones, Tom Waits, James Bay).

"*Sound On Sound* is informative from cover to cover, and I always read it. I always learn something from reading the interviews, and I rely on the new gear reviews. I look forward every month to the next edition."

Cenzo Townshend
Engineer, producer, mixer, two-time winner of Best Mix Engineer at Music Producers Guild Awards (U2, the Maccabees, Jungle).

"I love the magazine. There are certain magazines you read where you take everything with a grain of salt and you're not quite sure if you believe it. And *Sound On Sound*, when I read it, I believe it. It feels as though its only agenda is people reporting on stuff they're interested in."

Andrew Scheps
Engineer, producer, mixer, Grammy Award winner (Lana Del Rey, Hozier, Jay-Z).

SOUND ON SOUND

These articles were originally published in Sound On Sound magazine, **March 2023 edition**

Subscribe and Save Money!

Visit our subscriptions page at www.soundonsound.com/subscribe
for more information on the Sound On Sound App go to: www.soundonsound.com/app