

CEDAR CAMBRIDGE

AUDIO RESTORATION TOOLS



DAVE FOISTER runs through the new real-time, eight-channel mother of all noise reduction from CEDAR Audio.

I've been following developments at audio restoration specialist CEDAR for nearly ten years, and one theme shines through: the more sophisticated it gets on the inside, the simpler it gets on the outside. The palette of processes available to salvage the original material from damaged or compromised audio sources has, of course, grown, but it has also progressed from being very much a laboratory tool, running under DOS on a big, old PC platform, through dedicated hardware boxes in at least three generations, back to a computer host managing proprietary hardware in the form of CEDAR for Windows. This used the rapid advances in computing technology to integrate the processes on one platform, but in a configuration that allowed multiple simultaneous processes in a format vastly more elegant and manageable than the original system. Now the current generation of PC power makes it unnecessary to build separate hardware, so yet another step sees further refinements of the algorithms running as native processes on a dedicated rackmount computer, with the fastest, most intuitive front end yet to manage them, and all in eight simultaneous real-time channels. This is CEDAR Cambridge.

In a sense the hardware is the dull bit, but CEDAR's hardware has always been a good example of how things should be done; a glance under the lid of any of the hardware processors showed a standard of design and construction that made the insides a work of art. Just because they are now using a PC as the host doesn't mean those ideals have been compromised, as the central chassis is specially built with a server-grade motherboard, elaborate power supply arrangements and top-quality components. Compared to this, your average PC looks like a £5 clock radio inside.

However, just in case you're tempted to try a home build and somehow get hold of the software, forget it: the system relies on the separate Timecode Automation Controller for its operation and cannot function without it. The TAC is not just an option for timecode applications; it plays a crucial role in the operation of the whole system, making it, in a sense, the most sophisticated dongle you've ever seen.

Processing

But more of the TAC in a while; the crucial part of any CEDAR product is the processing at its heart and the ease with which it can be made to do its job. Cambridge offers seven separate processing modules, one of which combines the functions of two previous algorithms, as well as six core modules for configuring, controlling, and monitoring the processes. The central part of this control is the Process Manager, a deceptively simple and small window that allows the selection, ordering, and routing of the processing modules. Each chosen module appears as a little green screen, illuminated when enabled and dark when not, and can be switched in and out in this window for AB comparisons. There's also a status bar showing the selected I/O, sample rate, and bit depth settings, as well as the current CPU load. The hardware uses twin processors and great care has been taken to balance the load on the two, whatever the system is doing.


The restoration processes themselves are, by and large, familiar, although there are some new faces and all have been refined and improved. One surprise is the combination of the Declick and Decrackle algorithms, once said to be quite separate processes requiring different handling. In CEDAR terminology, clicks have always been large interruptions of the kind caused by a scratched or broken 78rpm record, while crackles have been the continuous barrage of smaller ticks of the kind caused by a mistreated LP surface. Crackles can also include some types of buzz and distortion, and CEDAR's Decrackle process has also been good at dealing with many of these. Now a totally new approach has resulted in the simplest process yet to deal with both types of disturbance, and the highly descriptive if slightly dodgy name given to it is Declickle. It can handle all the above problems, interpolating good signal to replace what has to be removed with the large clicks, and does it all with just three controls: a coarse/medium/fine detection switch, a Threshold setting, and a Modelling parameter. The first two allow very fast identification of the optimum setting to get rid of the unwanted effects, while Modelling affects Declickle's response to high harmonics in the original signal. The designers are particularly pleased with the improved behaviour in the presence of bright brass instruments, and the results now are spectacular, with no trace either of the original problem or of the processing that has removed it.

De-hiss-3 And NR-4

The other main bogeyman in recordings needing restoration is broadband noise, and Cambridge provides two separate modules for dealing with it. The simpler is Dehiss-3, the latest development in the line of powerful, fast algorithms. Setting it up consists simply of identifying the noise portion of the signal using the Threshold control, then deciding how much to attenuate it. Getting it wrong can introduce unwanted side-effects, but the right setting is usually pretty obvious and easy to find. The Brightness setting is not a tone control to compensate for loss of HF - Dehiss-3 does nothing to compromise the top end when set properly - but a means of telling the algorithm a little more about the nature of the signal. When something a little more powerful is required, NR-4 is the latest incarnation of perhaps the archetypal CEDAR process, the sophisticated noise removal system that



One of the options for Cambridge is a Dither output module.



Dehiss - rids your recordings of broadband.



Debuzz - little changed.




No more clicks and crackles with Declick.

uses a fingerprint of the noise to remove it. Extensive manipulation of its operation is possible; once a fingerprint has been sampled, EQ can be applied to it to affect how much noise is removed in the various areas of the spectrum, and the inverse of this is that EQ can be applied to the main noise-reduced signal without affecting the noise at all. Noise-free EQ is just that; having rescued a hissy recording you can then crank up the top end without putting the hiss back in. In fact, the EQ here is a major leap forward; whereas before it was fairly rudimentary corrective EQ, it's now a powerful, smooth, musical process. Both of these processes represent the pinnacle of what CEDAR can do - real noise removal without side-effects with results I've never heard matched or even approached by anything else.

Declip

An important newcomer to the system is Declip. Declip can deal with some types of distortion to a limited degree, but real salvaging of a truly clipped signal is beyond it. Declip is an altogether new process that can rescue a signal from hard digital clipping as well as other kinds of overload-generated distortion. The operation is typically simple, with a fascinating display to guide the user through the process. The display shows a histogram of the sample values in the signal, so that a healthy signal will show a kind of bell-shaped curve since the lower values, sitting in the middle of the curve, occur most frequently. If a signal is clipped, the maximum extremes occur unnaturally often, and these show as spikes at the sides of the curve with an immediate drop-off beyond as no values outside those points will occur at all. Red lines on the graph show where Declip comes into operation; when they are aligned just inside the clipping spikes, any sample values outside their range are replaced with interpolated good signal. This sounds like too much to ask even of CEDAR, but the example I heard, of a badly overmodulated piano recording, was dealt with incredibly well, all traces of distortion gone and a completely natural piano sound in its place. Other, softer types of clipping give a less clear display, but it's usually apparent where the problem is and easy to set Declip up to handle it.




No detail left out.



The centre of the processing monitor.

DNS

Also new to the package, although it's been available in hardware for some time, is DNS, which is now redesignated as Dialogue Noise Suppression. This is very specifically intended to improve dialogue and speech recordings by removing traffic noise, excessive reverberation, and noises from fans and motors, and is, in essence, a dynamic filter bank, applying sophisticated expansion to individual frequency bands right across the spectrum. The screen is very similar to the front panel of the hardware DNS box, and shows a fader for each of six bands together with an overall level control. The frequencies addressed by the actual analysis and noise reduction filters (many of which are controlled by each fader) are determined by the range selection, where the unit can be working across the full range, or just on lows or highs, or low plus mid, high plus mid - the combinations give it the flexibility to deal with most kinds of intrusive noise whatever the source. The band faders shape the response to the noise, and the level control decides how much the noise will be reduced. This is not intended as an alternative to Dehiss-3 or NR-4, but neither will they do the job of DNS. An example I heard of a big studio acoustic being converted to a small, dry living-room showed just what DNS can do, and it also has big advantages in the forensic audio work that forms a large part of



DNS - Digital Noise Suppressive.




The 0.2Hz resolution spectrum analyser.


what CEDAR is asked to do.

Debuzz

The final restorative process is Debuzz, which is a completely new algorithm, significantly upgraded from previous versions. It eliminates problems with a much wider range of fundamental frequencies, allows you to select the frequency range over which it operates, and offers an interesting display that shows how confidently the process has locked on to a buzz and how far its frequency is deviating from the one you've chosen. This doesn't use comb filtering in any way, but deals with a harmonic series based on a chosen fundamental to remove the buzz harmonics while leaving the wanted signal intact. It will track variations in the frequency over a $\pm 2\%$ range, allowing for fluctuating speed in the original recording medium.



A 12-band EQ anyone?



NR-4: Classic broadband noise reduction.

Many of these processes benefit from monitoring on a spectrum analyser, and the one built into Cambridge is a remarkable one. It is not an FFT analyser, but has far greater resolution controlled by the zoom level on the window - it's capable of 0.02Hz resolution even at low frequencies. Several measurement and hold tools are available, and a nice touch is a set of four Markers that, once set, appear on every other process window where appropriate, allowing chosen points to be transferred between algorithms.

One of these is the new EQ, which has no less than eight parametric bands, plus two low-shelving and two high-shelving bands. Each has a phenomenal range of control available, and the result is one of the most powerful EQ packages you're likely to encounter, capable of radical corrective filtering and smooth musical shaping. Like all other CEDAR Cambridge processes, it's snapshot-automatable against time, and these and manual changes are completely smooth and glitch-free. This is the main function of the Timecode Automation Controller, changing snapshots at predetermined times, but it also deals with timecode offsets to automatically correct for the latency inherent in Cambridge processes, and to realign for other delays elsewhere in the system.

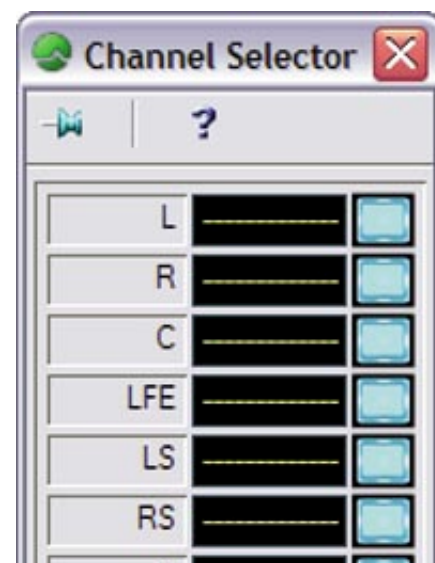
All the processes can operate across eight channels simultaneously, and the potentially complex routing and channel management is handled by a special module that can select individual channels and even group them together. All the processes allow different amounts of treatment to be applied to individual channels if required, giving massive flexibility, and a big bright metering window shows all the output channels simultaneously. Only at this point is metering and gain control really required as the whole package operates 64-bit floating point throughout; there's also a Dither module available for use on the output with a wide range of dither types and noise shapes.

Cambridge is a massive package whose capabilities need more space than we've got here to describe fully. Having said that, it's probably the most intuitive and friendly incarnation of the CEDAR technology yet produced, as well as the most powerful and refined. With its remarkable palette of treatments, its multi-channel, high-resolution capability, and its incredibly simple operation, there's nothing else like it on the market.

INFORMATION

£ CEDAR Cambridge from £7500 + VAT.

A CEDAR Audio Limited, 20 Home End, Fulbourn, Cambs. CB1



SBS, UK.

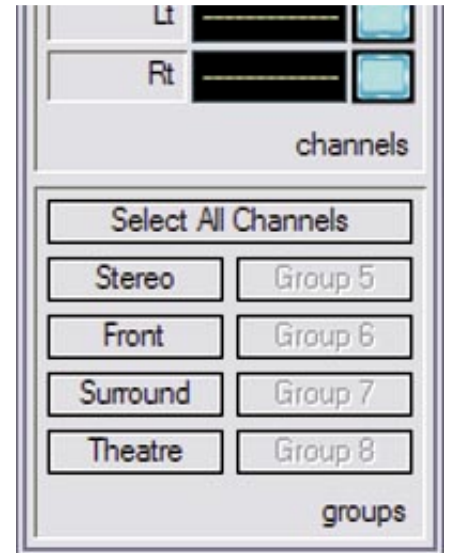
T +44 (0)1223 881 771.

E sales@cedar-audio.com

w www.cedar-audio.com

THE REVIEWER


Since doing the Surrey (UK) Tonmeister course, Dave Foister has produced, engineered, and played on innumerable recordings in the classical, commercial, folk, jazz, and TV music areas, as well as live sound mixing and theatre sound operation. He currently manages the Guildhall School of Music & Drama recording department, and spends his spare time(!) watching racing cars and wishing he was driving them.



Pick a channel, any channel.



Snapshot automation of treatments is a core part of the package.



Declip - rescues a file from hard digital clipping.

© 2003 IMAS Publishing UK Ltd. All rights reserved.