

CEDAR Cambridge V3

Restoration, cleaning up or the reconstitution of audio is a serious business be it in mastering, recording or forensics. User choice has certainly increased but if you're really serious about it then this is the product you need to look at.

ROB JAMES



BLACK MAGIC, WHITE MAGIC? Sometimes the dividing line is an indistinct and moving target. I've been rescuing dodgy audio for over 30 years from chopping clicks out with a razor blade and the Philips analogue declicker, all the way to today's digital wonders. It's an obsessive, esoteric arena. I wrote my first article about audio 'restoration' around 17 years ago and I still don't have any definitive answers. What I do know is that the leading purveyor of digital tools to exponents of the art remains CEDAR Audio. As a variety of restoration algorithms have become generally available all manner of pretenders have arrived at the palace gates, promising alchemical miracles. Some are good, lots are not. There is a lot more to it than simply picking an algorithm out of the grimoire.

Suitable cases for treatment can be divided into three broad categories; enhancing existing recordings for re-release, removing defects from original material for new productions, and forensic use. Priorities obviously differ across these areas. For remastering the most gentle touch must be applied with due regard for the original. Cleaning up material in the course of production can be more aggressive since other sounds, such as backgrounds and music, may well conceal artefacts. On the other hand, in forensic work, intelligibility is the principle criterion and the tools used to achieve this are often positively vicious. Here, provided that a transcriber can accurately interpret the speech, considerations of high fidelity are secondary.

CEDAR offers a wide range of standalone, real-time processors addressing specific problems. Its flagship PC-based system, Cambridge, can be specified to suit any or all of these areas with a wide variety of processes supporting many strategies.

The Version 3 demo system I had is a workstation class Cambridge 'Q' rackmount host PC equipped with two dual-core AMD Opteron processors. There is no proprietary DSP hardware. CEDAR optimises the Windows operating system at Registry level, therefore adding extra third-party software is definitely not encouraged.

An RME card provides the audio I-O in this instance. All Cambridge systems include the TAC

(Timecode Automation Controller) and communicating with the PC via USB this 1U provides full automation of the CEDAR processes internally and when chase synchronised against external code. TAC is also a giant dongle, storing the license keys.

In the Cambridge environment the Process Manager window is the command centre. From here you can select file or physical input and invoke the File Processor, the Set Up window and a variety of tools. This is also where processing modules are instantiated. This approach, with separate windows for logical components of the user interface, is well suited to the task. Windows can be arranged into a layout appropriate to the job in hand. The workspace thus created can be saved and any workspace can be set as the default for new sessions. Process Manager also relays information to the operator about CPU load, incoming sample rate, sample resolution and latency through the entire processing chain.

The Channel Selector determines which channels or channel groups are selected for adjustment and also reflects values for each channel when the mouse pointer hovers over a control in a Process Module. Global Settings along with Channel groups are defined in the Set Up window.

Automation is handled in the Event Manager window. An Event is a snapshot of all parameter values in every module loaded at the time the Event is created. Morphing between Events can be user determined from instantaneous up to the length of the entire Event.

A timeline display in the File Processor window shows currently loaded audio files as waveforms and enables regions to be defined. Cambridge can handle up to eight simultaneous channels. Thus surround formats up to 7.1, or multiple files with any programme content and bit depth but the same sample rate, can be processed, for example, four stereo files at once. File, transport and render controls are also to be found here.

Modules are added to the process chain by clicking on Insert Module and selecting a process. Clicking the blue block minimises/maximises the module's control window while the On/Off ideogram duplicates the

button function in the Module Control window. Signal flow is from left to right and clicking on a Module and dragging to a new position changes the process order.

One thing Cambridge does not do is record. All processing is performed either on existing files or live on the soundcard inputs. CEDAR says that the majority of Cambridge systems are on networks and few users wish to record in the Cambridge environment. Properly set up, network file handling is completely transparent. Temporary copies of files are always created and these can be local or on a remote disk as can rendered output files. Rendering can be undertaken in real time for monitoring or in non real-time mode at up to 30 times faster.


Cambridge is offered 'a la carte' and in four standard configurations aimed at Film, CD and DVD mastering and two levels of Forensic. Prices range from UK£6,250 to around £40,000 with a typical system coming in at around £20,000 (all plus VAT). In practice, the exact complement of modules supplied with even the 'standard' configurations varies according to the client's requirements. Every Cambridge system has sample rate conversion and dither/noise shaping and ships with a number of core modules under the heading of Utilities.




DC removal is normally the first step. Offering 12dB/octave roll-off in the user defined range 1Hz - 20Hz, a display shows the

extent of its action, if any, and which channels it is affecting. The exception to the rule is DeClipping. In this case removing DC offset first lessens the effectiveness of the process and DC removal, if required, should happen afterwards.

The Gain module provides gain control, solo in-place and muting. As with other controls this applies to all channels selected in the Channel Selector. Gain offsets can be achieved by simply deselecting channels you do not wish to affect. As with all other modules, multiple instances are allowed.



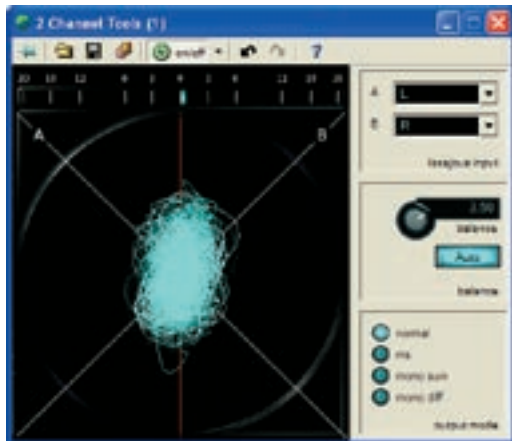
Metering offers four concurrent displays for each selected channel. Near instantaneous peak metering is shown by a horizontal bargraph with a fall time-constant of around 1.4s. RMS is indicated by a cyan bar and numeric. Peak Hold (1) is a grey bar with a hold time of 1s. Peak Hold (2) is numeric and is retained until the window is closed or the value reset by clicking on it.



Spectrum analyser is a high precision (0.02Hz resolution) graphic display of the selected channel(s). In contrast, the FFTs often found in DAWs have a typical resolution of 50Hz. All graphs in Cambridge share a common control set for averaging modes, zooming and Markers. Once you've zoomed in far

enough the display slows down to improve clarity.

Markers are a very powerful Cambridge tool. A frequency identified and marked in the analyser can be used to set a parameter in a process module with a mouse click.



The last Utilities module, 2-ChannelTools, provides a general toolkit for adjusting and monitoring 2-channel signals. A goniometer and +30dB balance bargraph give insight into the image. The balance control offers reversible gain shift between the channels and the Auto button keeps the image centred. Output controls enable the module to produce sum, difference or M&S outputs. Use of two modules enables M&S effects, such as width control, to be achieved.

There is no 'magic bullet'. If this is your Holy Grail, you will always be disappointed. All clean up and restoration tools are necessarily a compromise. Notwithstanding, the quest continues.

But, make no mistake, although the price can be steep there is no substitute for the real thing.

CEDAR processing continues to be the state of the art. 64-bit internal processing and meticulous software engineering endows freedom — freedom to use processes without fear of running out of headroom and tools to reduce the range for output. There is no real alternative.

Lest this seem like too much praise, workstation users may find a few things about the user interface rather odd. Mostly these could be viewed as charming eccentricities.

Within reason, whatever you do, the results can be heard, evaluated and modified without worrying about the limitations of the system. Given my film background, the favourite remains DNS, which still brings a smile at the miracles it can achieve. Linear EQ is highly desirable and, in the right circumstances, Debuzz-3, DeClickle-2 and NR-5 rewrite the definition of what is possible.

If you need an audio sorcerer, CEDAR has the spells.

THE PROCESSES

EQ — CEDAR offers two 64-bit 96kHz EQ processes for Cambridge, Precision EQ and Linear Phase EQ. Both can be automated.



Precision EQ is perhaps more applicable to forensic work, with seriously deep notches and brick-wall filtering, although in restoration it can help correct excesses perpetrated with analogue EQ.

Linear Phase EQ is something of a find. I discovered that both subtle and very heavy equalisation can be applied with none of the usual artefacts. The prodigious number of filter bands, two LF and two HF shelving

with 8 variable Q parametrics positively encourage experimentation. Multiple deep, narrow notches, which would result in horrible phase effects and even pre-echo with conventional EQs, are meat and drink to this CEDAR algorithm. Constant power mode means you don't have to worry about overcooking it.



constant power mode means you don't have to worry about overcooking it.



DYNAMICS — A suite of a Compressor, Upward Expander and Downward Expander. Positive and negative 'read-ahead' enables the dynamics to be

highly transparent or punchy when required. Each process has flexible linking options and side-chain EQ.



LIMITER — This Limiter is very different from a very high ratio compressor, employing a different signal path to minimise side effects.

Linking is available.



PHASE CORRECTOR — Sorts out static and varying phase discrepancies between channels with little effort.



DEBUZZ-3 — Buzz is defined as unwanted audio with many high-level, closely spaced harmonics. DeBbuzz-3 provides full-bandwidth processing, removing components across the entire audio spectrum, and is capable of removing all manner of buzzes and hums with fundamental frequencies as high as 500Hz. It can also track buzz varying by small percentages. Unlike alternative strategies it retains wanted signal even at the buzz harmonic frequencies by analysing the dynamics of the signal and buzz.

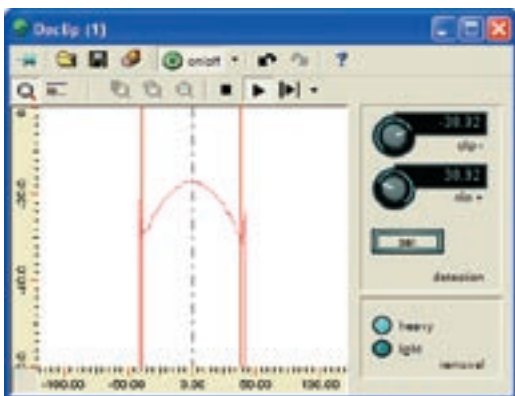
In the past, I've not been overly impressed by products purporting to remove buzz. DeBbuzz-3 has changed that perception and the results I heard were almost uncanny.



DECLICKLE-2 — Early digital declickers and decracklers could and did damage the sound of brass instruments, solo violins and sometimes singing. Declickle 2 has an improved noise detector and interpolator and offers improved performance by retaining useful information during clicks and crackles.



VINTAGE DECRACKLE — Uses the same algorithm as the award-winning CR-1 Decrackler. It will also remove some forms of buzz and amplitude distortion from material spanning cylinder recordings to live broadcasts with lighting buzz.



DECLIP — Clipping results in harsh distortion, increasing in intensity as clipping becomes heavier. In digital clipping the signal contains many false

samples. The analogue audio waveform shows 'flat tops' where the real signal has been destroyed and replaced by false samples at maximum amplitude. Declip enables the user to identify and remove most clipping in one pass. If the analysis window shows 'hard' vertical edges, the signal has been clipped at that sample value.

Analogue clipping is less well defined with a multitude of causes and exhibits a more rounded profile. Having spent many fruitless hours attempting to manually redraw audio waveforms damaged by clipping, the results obtainable here are quite remarkable.

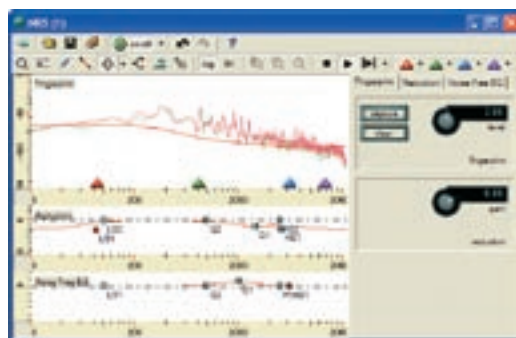
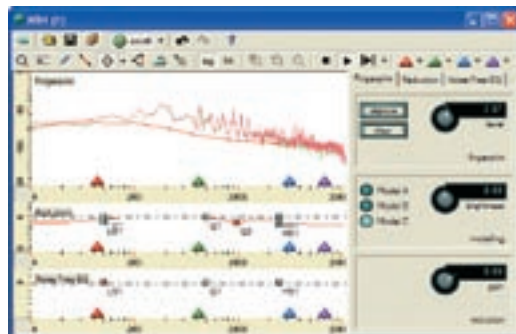
DEHISS 3 — Best suited for use when the unwanted broadband noise is consistent and uncoloured and when it is not possible or there is no time to use NR-4 or NR-5.

AUTO DEHISS — Embodies a more advanced algorithm. It has a unique 'Auto' mode that enables the software to determine the broadband noise content and remove it without introducing artefacts. Manual mode offers control over all parameters to fine-tune the process.

NR-4 AND NR-5 — NR-4 is based on developments of the original CEDAR de hiss algorithms while NR-5 uses the more recent Auto DeHiss algorithm. They share an almost identical interface. NR-5 lacks the modelling and brightness pane.

These algorithms are radically different. In most cases NR-5 will remove more noise with fewer artefacts. The exception is when the noise is buried deep within the wanted signal. Experimentation is the order of the day to determine which process does the better job.

As with all broadband denoise processes the controls are highly interactive and it takes time and experience to achieve the best results. An accurate fingerprint is useful. However, the interface enables the fingerprint to be adjusted or created manually using the drawing tools and the Noise Reduction EQ. This re-shapes the noise content contained in



the input signal and therefore allows the operator to concentrate the noise reduction in the range where the hiss is most intrusive, while leaving other areas relatively untouched. The Noise Free EQ acts only

upon that part of the signal identified as wanted. This can be used to introduce a mild high frequency boost to counteract loss of perceived brightness. NR-5 can produce the best results I've ever heard from this class of device in the shortest possible time.

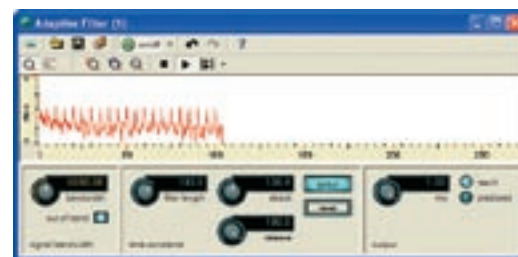


DNS — The modern world is a noisy place. Even at the poles there is no guarantee a jet plane won't happen by. CEDAR's Academy Award-winning DNS process splits the signal into a large number of well-defined bands.

Digital filters analyse each band, then independently suppress the noise in each. The really clever bit is presenting all this power to the user via relatively few controls, making it simple and quick to use.

Especially applicable to dialogue clean-up, DNS remains one of the most powerful tools ever offered for sound for picture clean up. The DNS defaults will seem strange to those brought up using Dolby Cat43a and Cat 430 processors. However, a moment taken to read the manual will reveal that the approach is necessarily different.

In forensic applications, DNS will remove motor noise, eliminate electrical interference and help clean up recordings made in unfavourable acoustics with poor microphone positioning.



ADAPTIVE FILTER PACK — The four Adaptive Filters are designed specifically for use in forensic audio. Adaptive filters are primarily intended to address constant unwanted sound. The unique CEDAR feature is the ability to retain out of band material free of processing, e.g. Sibilants and fricatives, improving intelligibility and listenability. Practically, this means more accurate transcription with less fatigue to the transcriber. The Attack time controls the rate at which the filter is allowed to react to changes. Speech can radically change inside a 100ms frame while the unwanted material may not. Filter length is the duration of the audio used to determine the filter coefficients. As filter length increases so does filter precision, however, long filters use a lot of CPU power. Where a reference recording of the unwanted signal is available, e.g. a broadcast, the results are truly astonishing. But don't even think about using adaptives for music or film. They're just not designed for it. ■

PROS Unbeatable performance; easy for novices but with the depth for experts.

CONS Price, which won't matter if you really need it; cannot do the impossible; not much else.

Contact

CEDAR AUDIO, UK:
Website: www.cedaraudio.com