CEDAR Series X

Declicker, Decrackler & Dehisser

- Signal Processors

By Hugh Robjohns

Published July 2000

CEDAR Audio's Series X units offer a trio of phenomenally powerful audio restoration tools in neat, stand-alone boxes with simple user interfaces. Hugh Robjohns cleans up..

CEDAR have been in the business of professional audio restoration tools for quite a few years now and are widely recognised as the leaders in this field. Their products span the immensely flexible and powerful CEDAR For Windows 2 PC-based system at the top end, all the way down to the relatively affordable 'entry-level' Series X products at the bottom. In absolute terms, even these CEDAR boxes may seem expensive, but you have to remember that they represent state-of-the-art noise-removal algorithms which are amazingly easy to set up and can perform a kind of sonic magic on the most sub-standard audio tracks.

The Series-X rackmount units were designed to package three of the most widely used processes — declicking, decrackling and dehissing — to appeal to broadcasters and mastering houses as an affordable alternative to their powerful PC-based flagship system or the more flexible 2U-rackmounting Series 2 processors. The design emphasis is really on achieving a good improvement in sound quality quickly and simply — the other CEDAR products may be able to produce better results under the most challenging circumstances, but require greater expertise to operate effectively and take far longer to set up properly.

Recently, a further pair of units were added to the Series X range, the BRX+ Debuzzer and the AZX+ Azimuth Corrector. Since these algorithms have greater complexity, they require a more powerful DSP engine and slightly more elaborate user controls than the original Series X units, and so have been christened the Series X+. Look out for a separate review of these units in a forthcoming *Sound On Sound*.

Clever Boxes

The three Series X units all share the same underlying hardware platform, running their processing on a 50 Mflops DSP engine. Each has an internal universal power supply which will accept mains voltages from 85 to 260V. They each occupy 1U of rack space, measuring 45 x 483 x 240mm, and weigh in at 2.5kg. The left-hand side of the control panel is identical for all three machines, but the right-hand side carries different controls depending on the specific role of the processor. There are no analogue I/O facilities at all, and only standard rate (44.1 or 48kHz) digital audio interfaces, via

In this article...

either AES-EBU or S/PDIF. However, the data path supports 24-bit resolution, and 40-bit floating-point signal processing is used internally.

The rear panel carries a pair of XLRs and very chunky isolated gold phono sockets for the digital audio interfaces at the left-hand end, with an IEC mains inlet at the right. There is no mains isolation switch — the unit is permanently powered as long as it is plugged in. The front panel is almost as simple, with a small rocker switch on the left to put the unit into standby mode (knocking 5W off the 15W operating power).

Three round buttons towards the centre all have associated LEDs to show their status. One selects the I/O connectors (AES-EBU or S/PDIF), the second attenuates the output by 3dB (to provide slightly greater operating headroom), and the third switches the restoration process on or off (effectively a bypass facility). When processing full amplitude signals the restoration algorithms can sometimes generate reconstructed peaks which are higher than the maximum allowed quantisation level for the output. Switching in the 3dB pad avoids this problem. The number and role of the rotary controls on the right-hand side of the front panel relates to the specific process and will be described in detail for each unit below.

Inherently, noise reduction is complex and requires a lot of processing, which means audio delays of 50, 41 and a whopping 180mS for declicking, decrackling and dehissing, respectively. These figures are for 44.1kHz sampling; naturally, the delay diminishes at 48kHz to 46, 38 and 135mS. In many applications, such as straight audio restoration, this may not be an issue, but if repairing the soundtrack of an old film, for example, a compensating delay may be required in the video path (a single PAL video frame has a duration of 40mS).

DCX Declicker

The declicker removes loud clicks and ticks such as those caused by scratches on vinyl records. The various CEDAR restoration tools work best in concert when the signal flow through the machines is arranged to optimise their respective noise-detection and modelling processes. The declicker should always be first in the chain, generally followed by the decrackler, with the dehisser at the end of the chain. This sequence is necessary since large clicks make it difficult for the decrackle process to identify and remove the tiny clicks and crackles of surface noise or dimmer buzz, and if clicks reach the dehiss process they can confuse it with very unmusical side-effects! Conversely, dehissing first will make it almost impossible to identify and remove clicks and scratches afterwards because important clues about the presence of clicks and crackles will have been removed by the dehiss process.

There is only one operational control on the DCX declicker and that sets the sensitivity of the algorithm. A low setting allows the machine to remove only the largest clicks, whereas a higher value encourages it to remove smaller clicks and ticks as well. As always, there is a tradeoff between click removal and audio quality — setting the sensitivity too high can induce distortion in the wanted audio as the algorithm becomes increasingly confused over what it needs to remove. The result of over-processing is a distinctive burbling sound, most easily heard after the starting transients of brass instruments such as trumpets and saxophones.

CEDAR's declick algorithm works in two stages: the first detects and removes the clicks, while the second fills in the gap! The latter is, of course, the tricky part and CEDAR use an elaborate technique to analyse the audio signal extending several milliseconds around a detected click. With this information it is possible to construct a model of the underlying resonant signal which can then be used to replace the section marred by the original click.

Signal modelling is not a perfect science, but CEDAR have refined and reinvented their algorithms since their first tools were developed 15 years ago. The current software is capable of removing up to 2,500 clicks per channel, per second and, with careful setting, most clicks can be removed extremely effectively with little or nothing to give the game away in the restored signal!

Crx Decrackler

The decrackler can be used to remove a surprisingly wide range of unwanted noises, not all of which would immediately be thought of as crackles! Aside from the obvious record surface noise, which is usually composed of high-intensity but low-level ticks, the decrackler can also remove buzzes and even some types of distortion — they all share similar, recognisable characteristics as far as the decrackle algorithm is concerned.

The disruptions that cause crackle tend to be small, and they add to the audio signal rather than corrupting it completely as clicks tend to. Therefore CEDAR's approach is to divide the crackly signal into two components: one containing the bulk of the undamaged, required signal, and the other containing all the degradation (inevitably with some of the clean signal). Again, separating the two elements is not a trivial process and requires some human instruction, of which more in a moment.

By careful analysis of the portion of signal containing the majority of crackles, it is far easier to detect the crackles and buzzes and to remove them, in this case with an interpolation system optimised for the high-density nature of crackles. The remaining signal is all wanted audio and is recombined with the undamaged part of the signal. This technique ensures that the maximum amount of audio information is retained, even during the brief periods when the interpolation process is active in removing crackles.

The CRX unit is provided with a pair of rotary controls (labelled Level and Sensitivity) together with a button which switches between a 'Detect' mode and the normal 'Decrackle' mode. To set up the process the Detect mode is selected and the level control adjusted to establish the lowest point at which all the offending crackles or buzzes have been removed — this is determining the undamaged portion of the original signal. Separating the signal in this way is most odd to listen too, the retained portion exhibiting a strange 'under-water' characteristic, but this is entirely irrelevant and will not manifest itself in the final output.

Once again, setting of the level control is critical for the final result — too low and not all the crackles will be removed (because some will have bypassed the decrackle algorithm by residing in the 'undamaged' portion of the signal); too high and the processed part of the signal will become distorted.

With the optimum setting found, the button is released to select the Decrackle mode (in which the two signal elements are recombined) and the sensitivity control can then be adjusted to determine how hard the decrackling algorithm has to work. A low setting only removes the most obvious crackles, whereas a higher setting will remove finer crackles, buzzes and even some kinds of peak distortion, if present. Turn the knob up too far, though, and you can end up adding distortion!

I found it necessary to keep switching the whole process in and out to check that I was only removing the unwanted crackles — it is very easy to get carried away with the astonishing effectiveness of these processes. I found it all too easy to throw the baby out with the bathwater by over-processing the signal. This results in no clicks, no crackles, but lots of very odd low-level artifacts! The thing to remember is that these processes are intended to *improve* the sound quality, not necessarily to render it perfect in every way. With care, significant and worthwhile improvements are easy to achieve in most cases and sometimes the results border on the miraculous!

DHX Dehisser

The dehissing technique used in CEDAR's PC-based systems is based on a spectral subtraction process where the audio band is divided into very narrow slices and each is processed individually to subtract an identified noise spectrum whilst leaving the wanted audio intact. Identifying the unwanted noise in the first place requires a 'noise fingerprint' to define its spectral characteristics — this can be obtained during an otherwise silent section of the wanted material.

Furthermore, since noise is random and continually changing, a static 'fingerprint' is not sufficient to produce worthwhile results and, consequently, the CEDAR algorithm effectively updates its noise fingerprint every 1,024 samples to keep track of the inevitable variations in the noise content — a technique which also helps to prevent the compression of incoming transients and distinguishes between true noise and, for example, reverberation.

However, deriving a noise fingerprint in a stand-alone product such as the Series-X dehisser would require a relatively complex user interface, inappropriate for this kind of application. Consequently, CEDAR have developed a noise-removal algorithm capable of completely autonomous determination of the noise content, requiring minimal user-input via just three rotary controls.

To set up the process, the Attenuation control is first set fully anticlockwise (which provides maximum noise removal), the Variance control is placed at its mid point and the Level control backed fully off. The Level knob is then increased until a point is reached where all the hiss suddenly disappears and a distinct 'twittering' effect becomes audible. Advancing the control a little further removes the twittering but introduces a different artefact known as 'glugging', recognisable for its obvious HF compression and gating effects. The optimum Level setting is where both the 'glugging' and 'twittering' are minimised, although on some material these points may overlap one another.

By adjusting the Level control in this way the dehiss algorithm is provided with all the information it requires to operate correctly, although at this stage the output will sound extremely over-processed! The Variance control interacts with the Level control to some degree in biasing the algorithm to work harder in removing a coloured or variable hiss (a high setting), or taking a more gentle approach (a low setting) which maintains optimal audio quality at the expense of less noise removal.

The Attenuation control is the final arbiter of how much noise the process removes. Too much

attenuation will, inevitably, reveal processing artifacts, and this process is not quite as powerful as the technique used in the flagship CEDAR systems. However, it can certainly improve a noisy recording very effectively and with minimal fiddling.

The Ears Have It

The kind of advanced noise-removal achieved by the CEDAR Series X boxes is, simply, amazing. They can not perform complete miracles (yet), but they can certainly salvage material which would otherwise be destined for the bin and produce extremely impressive, artifact-free results. All three units are very fast and simple to use, although the more care taken in fine-tuning their processing the better. I found it took quite a while to train my ears to recognise the sometimes subtle consequences of overprocessing on each machine, and it is far better to err slightly on the side of caution than to risk creating low-level artifacts.

If you are involved in regular restoration projects or have to routinely handle compromised source material, the Series X units represent very good value for money. However, if your requirement for a little CEDAR salvaging is occasional, the company also runs a 'bureau service' from its offices in Cambridge, where in-house experts will perform minor miracles for only a very modest fee!

Cedar Wood For The Trees

The CEDAR story started back in 1983 when the British Library National Sound Archive wanted to develop PC-based audio-restoration systems. By 1985 the Engineering Department at Cambridge University had become involved and two years later the first prototype CEDAR computer system was announced, with sophisticated algorithms designed to restore damaged recordings on decaying media such as wax, shellac, and celluloid.

The first commercial CEDAR systems were launched in 1990 using custom hardware running inside a standard PC. The first 16-bit CEDAR-2 System quickly lead on to the development of a range of 24-bit dedicated hardware processors and the PC-based 20-bit CEDAR-20 Production System. The first of these stand-alone hardware processors were the DC1 Declicker (released 1992) and the CR1 Decrackler (1993), but rapid developments in DSPs allowed an improved Series 2 range to be launched a year later in 1994. Besides updated DC1 and CR1 units, the new range introduced the AZ1 Azimuth Corrector and the DH2 Dehisser.

Back in the PC domain, the CEDAR For Windows restoration system was launched in 1996 with the ability to cascade processes within a single PC (or process up to 16 tracks simultaneously). These restoration tools were also made available to Mac users in the form of CEDAR For Pro Tools. Last year the CEDAR For Windows platform was upgraded to Mark II status incorporating the latest algorithms, refined user interfaces, and even more software processes.

The subjects of this review, the Series X processors were first revealed in late 1997 as a cost-effective range of rackmount processors providing the most frequently used algorithms in an easy-to-use format. CEDAR have also been developing software plug-ins for other manufacturers' platforms. The first was a Denoise algorithm for the SADiE audio editor in 1996; this was followed two years later by the Declick process, which was unique in that all of the processing decisions are made by the algorithm itself. Last year saw the Declick algorithm being ported to the Soundscape SSHDR1 and R.Ed systems. CEDAR's latest product is also one of its most innovative. The CEDAR Dethump algorithm runs off-line on the SADiE platform and is capable of removing low-frequency thumps lasting over a second.

Pros

- Astonishingly simple user controls.
- Very powerful and effective noise-reduction algorithms.
- Negligible side-effects when set up carefully.

Cons

- Rocket science ain't cheap.

Summary

Rocket science inside, simplicity outside. The CEDAR Series X processors are delightfully fast and easy to use and produce stunning results from the most dreadful source material, given a little care.