Cedar Series X+

Debuzzer & Corrector

- Signal Processors

By Hugh Robjohns

Published August 2000

We looked at CEDAR's three Series X audio-restoration processors last month. This issue sees Hugh Robjohns correcting his azimuth and removing his buzzes with the two latest additions to the company's catalogue...

In this article...

CEDAR's original Series X processors, reviewed last month, ported the advanced audio-restoration algorithms used in their computer-based products to stand-alone rackmount devices, providing straightforward solutions for declicking, dehissing and decrackling. Since their launch, DSP power has become even more affordable, allowing CEDAR to port two of their most processor-intensive algorithms over to stand-alone hardware as part of the X series. The newly available processes are the BRX+ debuzzer and the AZX+ azimuth corrector.

To differentiate the new, more powerful, hardware platforms that form the basis of these devices, which run 60Mflop processors (as opposed to the mere 50Mflop version in the original Series X units), CEDAR refer to these new processors as the 'Series X+', and they are recognisable by their inclusion of seven-segment LED displays and a large encoder knob. Much of the underlying hardware is the same as on the original units, but CEDAR have incorporated a field-prgrammable gate array (FPGA) to drive the front-panel LED display and, at the same time, remove some of the housekeeping workload from the DSP. In practice, this means the X+ units have close to 25 percent more processing power than their elder siblings.

Like the other processors in the series, the X+ units provide only standard-rate digital audio interfaces (44.1 or 48kHz), but both AES-EBU and S/PDIF are supported with full 24-bit resolution. Internal processing is to 40-bit, with floating-point number crunching. The mains power supply is a universal switched-mode type spanning input voltages from 85 to 260V with the obligatory IEC inlet socket, but there is no mains isolation switch — the front-panel rocker merely puts the machine into a low-power standby mode.

Brx+ Debuzzer

There can be few recording engineers who have not been plagued by hums and buzzes finding a way into the signal paths and spoiling the sound quality to some degree. CEDAR's debuzzer is a highly innovative restoration tool designed to remove a wide range of hums and buzzes from a stereo audio signal — anything from subtle underlying mains hum to harmonically rich interference buzzes, such as those sometimes caused by lighting dimmers, for example.

The debuzzer algorithm was originally released as part of the flagship CEDAR for Windows restoration system at the beginning of 1998, but the company have now been able to incorporate it in the popular stand-alone Series X+ hardware. The BRX+ unit carries the same arrangement of three front-panel buttons and associated LEDs as the other Series X units: an input format selector, a 3dB headroom switch, and a Process In/Out button. There are just three user controls specific to the BRX+: two knobs, a numeric frequency display, and an encoder wheel. The two control knobs adjust the sensitivity of the processing (Buzz Power) and the amount of hum/buzz to be removed (Attenuation). The frequency display and associated encoder wheel are used to instruct the system on the fundamental frequency of the hum or buzz. There are three possible operating modes: Selection of standard preset values, Coarse manual tuning, and Fine manual tuning.

Setting Up

As on the Series X modules, the Series X+ units offer only digital I/O, in either AES-EBU or S/PDIF format.

Connecting the BRX+ into a digital signal path is simple — the machine synchronises automatically to its input source — although there is a fixed processing delay of 41mS (equating to one television video frame). As with most of CEDAR's processes, the starting point is to set the controls to their maximum positions so that the processing becomes extremely obvious. The fundamental frequency of the buzz or hum must then be identified, either by selecting a preset frequency or by tuning the unit manually. Selecting presets is performed by pressing and rotating the encoder wheel until the desired frequency is displayed (most presets are centred around the common 50 and 60Hz mains frequencies and their overtones).

Once the nearest appropriate frequency has been found, releasing the wheel automatically engages the manual-tuning mode, which can then be toggled between coarse and fine operation by pressing the encoder wheel again. The current coarse or fine mode is indicated by small dots to the right of the numeric display. The manual tuning range extends down from 130Hz to 100Hz in increments of 0.1Hz, and then in increments of 0.01Hz below that, the degree of precision highlighting the importance of locating the correct fundamental frequency.

The debuzzing algorithm is designed to lock on to a consistent hum-like fundamental frequency anywhere within a 2 percent window centred on the frequency shown in the display. Once found, tracking indicators (Up and Down arrows to the left of the frequency readout) denote the relationship between the indicated and detected frequencies. Fine-tuning in the appropriate direction to centre the system on the detected fundamental results in a central green dot between the two arrows, and the algorithm can then track small frequency drifts automatically.

Locating the fundamental frequency of the hum or buzz is relatively easy whilst all the user controls are set to maximum as these unwanted noises all disappear! However, it's important to ensure the true fundamental has been located and not a harmonic or sub-harmonic. This could result in under- or over-processing respectively, which can create undesirable audible artifacts.

With the fundamental hum frequency correctly identified, the next step is to wind the Buzz Power control back until you reach the point where the unwanted hums or buzzes become audible once more — an action which instructs the algorithm about the division between the unwanted noises and the required musical information. The final step is simply to decide how much of the buzz or hum to remove, which is the role of the Attenuation control.

In practice, the system is very fast and simple to set up. In fact, once you are familiar with the controls and their effects it takes far less time to do than to read my description of the process! With the most evil test material I could find, the only time the system performed less than perfectly was when dealing with multiple simultaneous hums and buzzes with different fundamental frequencies. The BRX+ could be set up to sort out one set of hums or buzzes very well, but it could not clean up everything in one go. In this situation CEDAR recommend repeated passes through the unit (or processing via multiple BRX+ processors or multiple debuzzing algorithms on a CEDAR for Windows system) with different fundamental frequency settings to remove each set of hums as necessary. None of the other CEDAR restoration tools (such as the declicker, decrackler and dehisser) can be operated in this multi-pass manner, as too much of the information required to identify the unwanted material is removed during the first pass.

With the typical types of hums and buzzes from ground loops, dimmer noise, electromagnetic interference and the like, the effect of the BRX+ is, quite simply, amazing. Audio material marred to the extent that it is virtually unlistenable, and certainly not of commercial quality, can be completely saved with no detectable side effects whatsoever! Dealing with the more modest, but still annoying mains-related noises typical of electrified instruments is pure child's play. As with most of the CEDAR tools, it is easy to get carried away and over-process things at first, but I quickly began to recognise the subtle audible clues that inform just how far the system can be pushed in removing unwanted hums and buzzes. As always, careful use of the controls resulted in substantially cleaned audio with no side-effects at all.

Because I could, I even tried mixing low-frequency drones of various kinds from a synth into otherwise clean material to really challenge the efficacy of the machine in laboratory conditions! However, it became immediately apparent that, with careful setting, the processed version was completely indistinguishable from the original. For comparison, I also tried to remove the synth drone with high-Q notch filters in a digital console (a very commonly employed technique). Not unsurprisingly, this approach caused easily audible damage to the audio surrounding the notch frequencies and was infinitely less effective than the BRX+.

Azx+ Azimuth Corrector

The CEDAR azimuth corrector is intended to resolve the small inter-channel timing differences which can creep into stereo recordings, and the algorithm has already been seen on the CEDAR for Windows PC system as well as in a Series 2 hardware processor. Inter-channel timing variations are commonly caused by misaligned record cartridges or tape heads in the analogue world, and by incorrect conversion and signal-processing strategies in the digital domain. However, the audible result is the same no matter how the timing inaccuracy is caused: muddy bass and a general smearing of the stereo image or, if the two channels are combined to mono, a loss of high frequencies and a boxy or flanging quality imparted on the sound. It doesn't take *much* of a timing difference between channels, either. Back in the mid-'80s, those with 'golden ears' became quite agitated about some types of early A-D and D-A converters which introduced a half-sampling-period delay in the signal — all 11 microseconds of it!

The process employed in the AZX+ operates in two distinct stages. First, the algorithm detects the timing error between the two channels (or this can be entered manually) by looking for a common signal element in both channels, examining the programme material roughly 50 times a second. When the appropriate time-shift has been found, the audio data is realigned in each channel to remove the offset. This is a continuous process with minute time corrections corresponding to every calculation of the offset, which ensures a glitch-free output.

The nature of the algorithm used in the AZX+ azimuth corrector means that most of the standard selection facilities found on the other X-series units are not necessary. Consequently, the digital input selector button of the other units is retained but is joined by three other push buttons related to the operation of the process itself rather than the signal path through the machine. As on the BRX+ unit, a rotary encoder knob is used to determine the operating mode by holding it in whilst rotating it. The normal condition is Auto, in which the machine detects and corrects the timing error automatically. A Manual mode allows the user to determine a timing offset which the processor will then apply. The required time-shift is entered by rotating the encoder control and both coarse and fine adjustment modes are available.

The display is calibrated in samples rather than real time — fine mode allows a resolution of 1 percent of the sample period (0.2µS at 44.1kHz) over a 10-sample-period range (226µS), whereas the coarse mode provides 10 percent (2µS) increments over a 100-sample range (2mS). In a similar fashion to the BRX+, tracking indicators show the relationship between a manually entered time shift and the value calculated by the AZX+ algorithm. The third operating mode is, simply, a quick zeroing facility which cancels any applied time-shift and resets the algorithm's tracking mechanism.

With a manual or automatically assessed time-shift entered in to the processor, the signal is processed according to the conditions determined by three buttons on the left of the front-panel, alongside the digital input selector. The first of these is labelled Pre/Post and, effectively, bypasses the processing when in the Pre condition to enable comparison of the original and corrected signal.

Two further buttons select the output between Mono/Stereo and Sum/Difference. The normal Stereo mode simply passes each input channel to its respective output, through the time-correction process as appropriate. When selected to operate in the Mono mode, the Sum/Difference button becomes active and determines how the two input channels are combined after time-correction and before being output (identically) through both output channels (in the Mono mode the output level is reduced by 6dB to ensure that peak-level input signals do not overload the digital outputs).

The reason for this facility on the AZX+ is to aid with the restoration of mono programme material. It is often advantageous to replay such material with a stereo pickup cartridge or tape head and then recombine the two tracks. Since the noise collected on each channel is not coherent, summing it results in only about 3dB more level whereas the (time-aligned and therefore coherent) music signal will increase by 6dB — a 3dB overall improvement in the signal-to-noise ratio! Selecting the difference mode can also be very helpful in trying to establish the optimal time-shift for stereo material as signals which are identical on both channels (ie. those from centrally placed instruments) will cancel out when correctly aligned.

With most material, setting the azimuth corrector up is as simple as selecting the Auto mode and sitting back. However, with some stereo music (mainly that which lacked a continuous central element) the algorithm seemed to get confused and started to cycle the time-shift value up and down, desperately trying to find a consistent timing reference. In these cases manual intervention was required, which is where the Mono and Difference buttons became extremely helpful in locating the optimal time-shift setting (assuming it was possible to find some portion of the programme with a centrally placed instrument).

The AZX+ is one of those processors which can largely be plugged in and forgotten whilst it beavers away sorting out all those little timing errors which no one has noticed before. On most of the modern material I tried the benefits were often extremely subtle, if any at all, but the machine came into its own when processing some of my collection of old quarter-inch tapes dating back to the '50s and '60s, where I suspect that real tape-head azimuth errors detracted from the ultimate fidelity.

Mono recordings certainly benefited from slightly lower noise after the process (the worst also exhibited a slight but noticeable HF loss, proving that the replay azimuth alignment was not matched to that of the original recorder). Of more interest to me, though, was the effect of the processing on stereo recordings. On many recordings the stereo imaging seemed far sharper and, on some jazz tapes, the timing between the various musicians (and especially with the bassist) was perceptibly tighter. I had not previously felt that there was anything materially wrong with these tapes but, after routing them through the AZX+, the benefit of accurate time-alignment was clear to all who listened.

Conclusion

Inevitably, the highly specialised R&D and relatively small marketplace for such units as these make CEDAR's audio-restoration tools seem expensive; and for many, these restorative tools may be required so infrequently as to make them unjustifiable. For people like these, it might suffice to access them when needed via CEDAR's own bureau service, which is staffed by highly experienced operators. However, to put the price in perspective, the Series X and X+ units are very cost-effective compared with the (more powerful) Series 2 and CEDAR for Windows systems. If hums and buzzes plague your recordings — and cannot be resolved by good engineering practices — the BRX+ could present the ideal solution. Equally, correct time-alignment proved remarkably beneficial on some material I had previously thought was perfectly fine, and the AZX+ is an extraordinarily simple machine to use. The important question is whether you really need them. If you do, they do exactly what you expect them to do and without any undesirable side-effects. The choice is yours.

Hum Removal

Most of the hum-removing systems I have encountered rely either on simple notch or high-pass filtering, or some form of comb-filtering technique, but none have been particularly effective, particularly if the offending hum varied in any way. Notch filtering fails to deal with any harmonics and cannot cope with frequency drifts. The high-pass approach can accommodate some drift in the fundamental frequency of the hum but throws the baby out with the bath water since effective reduction of the fundamental also removes any wanted musical bass frequencies. Comb filtering can remove some of the hum's harmonics if carefully set up, but inherently sounds boxy and artificial and cannot cope with frequency drifts at all.

The CEDAR approach is fundamentally different, and uses a development of the company's sophisticated audio restoration algorithms to remove unwanted noise in the presence of wanted signals. A certain amount of user input is required to indicate the fundamental frequency of the hum or buzz, and to indicate the relative power of the unwanted signal. But with this information the algorithm is able to recognise consistent and continuous harmonically-related noises in the programme material as the unwanted hum or buzz and subsequently reduce or remove it without damaging the underlying audio in any way.

Pros

- Easy to use.
- Stunningly effective.
- Negligible side-effects when processed carefully.

Cons

- Rocket science remains as costly as ever!

Summary

The BRX+ and AZX+ are simple to use and stunningly effective audio-restoration tools capable of removing complex real-world hums and buzzes and realigning stereo material using state-of-the-art algorithms.