CEDAR

Series X

Digital Audio Restoration Modules

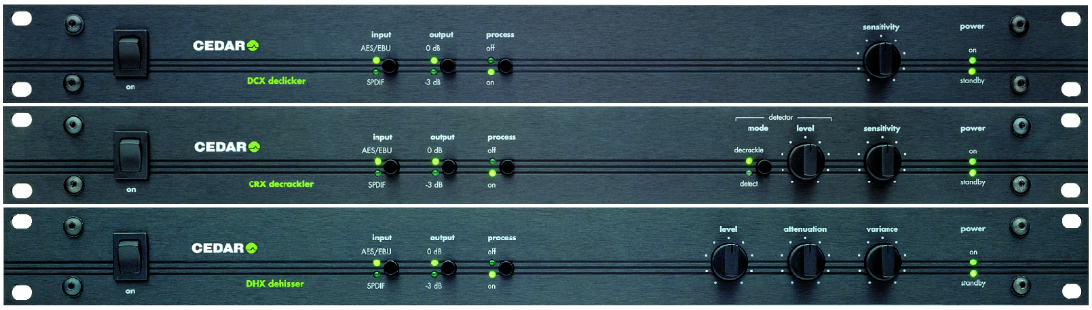

OWNER'S MANUAL

Manual Ver 1.00: 1997
Reformatted Nov 2010

TABLE OF CONTENTS

PART 1

PART 2

PART 1

Series X
Operations Manual

INTRODUCTION

Thank you for purchasing this CEDAR Series X audio restoration module. It is one of a family of the world's most advanced and most effective noise removal units, and it offers outstanding processing power and performance. Quality, speed and simplicity are paramount in the Series X design, and all the units in the family are intended for professional use. Nevertheless, they will work perfectly well in any digital audio environment. Series X features include the following:

- **Fidelity**
  Each Series X offers the world's latest and best noise identification and removal algorithms
- **Simplicity**
  High levels of artificial intelligence make each Series X extremely simple to use
- **Accuracy**
  All Series X interfaces offer 24-bit input and output resolution
- **Flexibility**
  Each model in the Series X range will handle a wide range of audio restoration requirements
- **Tolerance**
  Every Series X supports any digital sample rate from 30kHz to 50kHz
- **Immediacy**
  With a Series X there's no laborious installation into a hard disk editing platform, and no compatibility problems with your existing equipment
- **Transportability**
  Unlike a computer-based system, a Series X can quickly and easily be moved from site to site for use in a wide range of applications
- **Power**
  A 40-bit floating point DSP processor delivers 50MFlops so that the Series X will handle the most complex processing requirements
- **Conformity**
  Each Series X incorporates twin digital audio interfaces conforming to both the AES/EBU and SPDIF standards
- **Universality**
  Its universal power supply means that a Series X will work anywhere in the world

SAFETY INSTRUCTIONS

Read these instructions, follow them, and save them for future reference.

- **Water and Moisture**
  The unit must not be exposed to rain or moisture. Furthermore, if the unit is brought directly from a cold environment into a warm one, moisture may condense inside it. This, in itself, will not cause damage, but may cause electrical shorting. This could damage the unit, and even cause danger to life. ALWAYS allow a Series X to naturally reach ambient temperatures before connecting the mains power.

- **Mounting and Ventilation**
  The unit may be mounted in a 19" EIA rack, or placed on a flat, stable surface. Do not subject it to strong sunlight, excessive dust, mechanical vibration or periodic shocks. The Series X is not susceptible to heat build up, but should be installed away from heat sources such as radiators, and audio devices such as amplifiers that produce large amounts of heat. If used as a free-standing unit, the supplied rubber feet should be fixed to the base of the unit.

- **Power Sources**
  The Series X features a universal power supply which will work safely on any mains supply in the ranges 85V to 260V, 50Hz or 60Hz AC only. The unit should always be grounded (or 'earthed'), and power connectors should be routed so that they will not be walked on or pinched. The Series X is connected to the mains power as long as it is plugged in to a mains supply so, if it is not to be used for an extended period, unplug it from the wall. Pull the connector out by the plug, never by the cord itself.

- **Connections**
  Turn off the power to all equipment before making any connections.

- **Cleaning**
  Clean only with a dry cloth. Never use abrasive pads or liquid cleaners such as alcohol or benzene.

- **Damage Requiring Service**
  The Series X contains no user-serviceable parts and should on no account be opened or dismantled by unauthorised personnel. It should be returned to qualified service agents when it has been exposed to liquids, when it fails to function correctly, when it has been dropped, or when the case is damaged.

SET UP

Unpacking

Unpack the Series X carefully. Save the carton and all packing materials since you may need them to transport the unit in the future.

In addition to the Series X and the packaging, the carton should contain the following:

- mains connection lead
- this manual
- rack mounting hardware
- four self-adhesive feet

Installation Site

To maintain reliability and prolong operating life observe the following environmental considerations:

- nominal temperature should be between 5˚C - 35˚C (41˚F – 95˚F).

- relative humidity should be in the range 30% to 80% non-condensing.

- strong magnetic fields should not exist nearby.

AUDIO CONNECTIONS AND STANDARDS

A Series X may be connected to most professional digital audio equipment. Two types of digital audio input and output are provided, and these will satisfy most interconnection requirements.

Signal Lead Connections

The Series X offers two audio connection standards, and passes its signal to both outputs irrespective of the input used. The standards are:

- digital SPDIF format audio data
- digital AES/EBU format audio data

Digital SPDIF format audio data
SPDIF is used by domestic and semi-professional digital audio devices. You should connect the SPDIF output from your source to the SPDIF input of the Series X using a single cable terminated with an RCA (or 'phono') plug. The SPDIF output of the Series X should be connected to the SPDIF input of a recording device or external DAC.

Digital AES/EBU format audio data
The AES/EBU format is used by professional digital audio devices. You should connect the AES/EBU output from your source to the AES/EBU input of the Series X using a single cable terminated with an XLR plug. The AES/EBU output of the Series X should be connected to the AES/EBU input of a recording device or external DAC.

Channel Status

No Series X processes are affected by channel status data. All units will echo any such data directly to the outputs.

Attenuation

On occasion, a Series X will correct a signal that would have exceeded the digital I/O maximum amplitude had it not been damaged. In this case, restoration may lead to digital clipping.

Avoid clipping using the Output Attenuator.

FSD Signals

Commercial test CDs and signal generators may be used to test the operation of components within audio systems. The Series X may generate distortion if a digital FSD (full-scale deflection) sine wave is applied to either input. This does not imply that your unit is faulty, and the effect should be ignored.

Group Delays

A Series X delays any signal passed through it, as in the following table:

	ms @ 44.1kHz	ms @ 48kHz
DCX:	49.7	45.7
CRX:	41.4	38.0
DHX:	180.0	165.4

EMC Regulations

In order to fully comply with EMC regulations, the Series X should be connected using metal-shelled connectors and good quality shielded cable suitable for digital audio.

FRONT PANEL INDICATORS AND CONTROLS

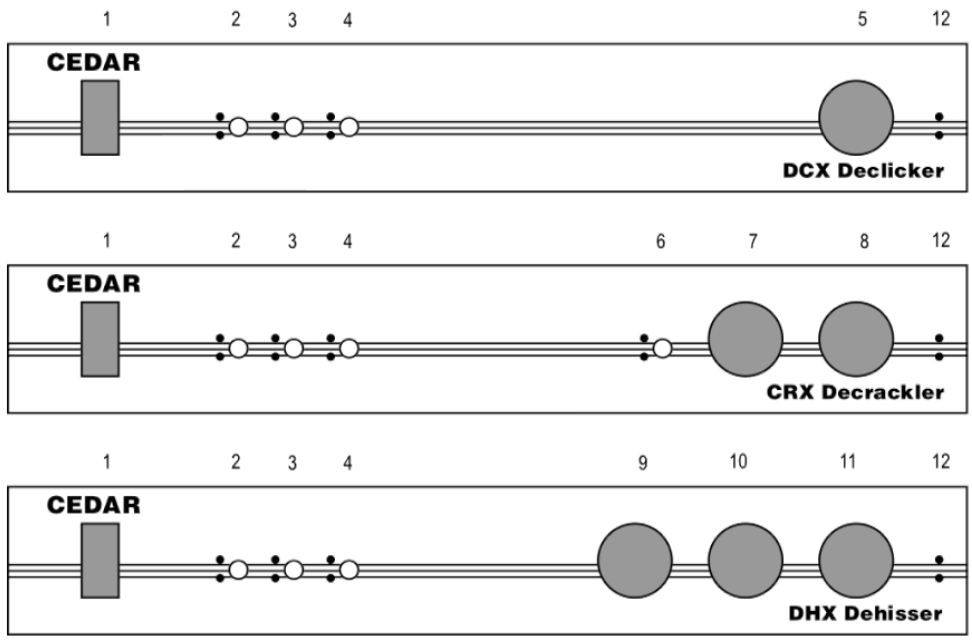

Refer to Figure 1 above for each Series X product:

1. Power Switch

- When OFF, the Series X is in standby mode.
- When ON, the Series X is powered-up and operational.

2. Input Selector and LEDs

- When OUT, the AES/EBU input is selected.
- When IN, the SPDIF input is selected.
- If a Series X fails to lock to an incoming signal, the LEDs will flash.

3. Output Attenuator and LEDs

- When OUT, no attenuation is applied.
- When IN, the signal is attenuated by 3dB.

4. Process ON/OFF and LEDs

- When OUT, the process is minimised, and the input signal is echoed at the outputs.
- When IN, the process is active and the processed signal is directed to the outputs.

5. Sensitivity (DCX)

Sets the sensitivity of the click detector.

6. Mode and LEDs (CRX)

- When OUT, the decrackled (final) signal is directed to the outputs.
- When IN, the unaffected part of the original signal (i.e. that which does not suffer from the degradation) is directed to the outputs.

7. Detect Level (CRX)

Determines the signal fed to the crackle detector.

8. Sensitivity (CRX)

Sets the sensitivity of the crackle detector.

9. Noise Level (DHX)

Determines the level of the noise contained within the original signal

10. Noise Attenuation (DHX)

Determines the amount of noise removed once detected

11. Variance (DHX)

Fine-tune the dehiss process for various types of signal and noise

12. Power Indicator LEDs

- When the Power Switch is OFF, the lower LED indicates that the Series X is in standby mode.
- When the Power Switch is ON, the upper LED indicates that the Series X is operational.

The DCX Declicker

OPERATING THE CEDAR DCX

The term ‘click’ is used to describe many different audio phenomena, including ticks, pops, clicks, crackle and thumps. Each of these displays different sonic characteristics, so a single process designed to remove all of them would be a compromise, incapable of the best possible repair of any single category.

We separate these degradations into three categories: thumps, clicks (including ticks and scratches), and crackle. The DCX has been designed to perform real-time click removal. It features just one control, and this determines the action of the declick algorithm. We represent this algorithm as follows:

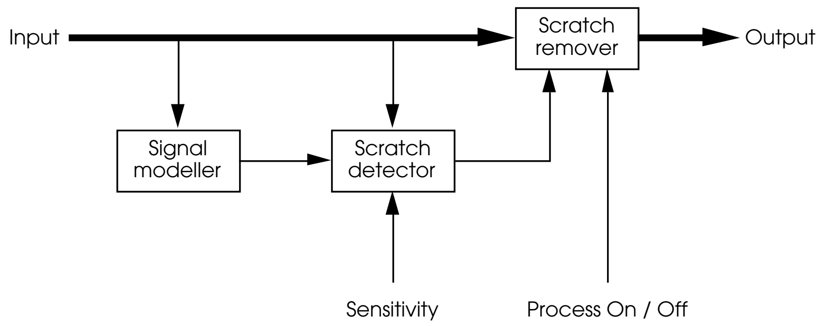

Figure 2: The DCX process overview

Sensitivity

This determines the type and number of clicks removed. A low SENSITIVITY removes only large clicks, while a higher SENSITIVITY also removes fine ticks and clicks.

Increase or decrease the SENSITIVITY by turning the control clockwise or anticlockwise (respectively).

Note: *Distortion of the genuine signal may result if the SENSITIVITY for a given piece of music is too high, and in these cases a lower SENSITIVITY is advised. If distortion is introduced, it will be most noticeable as a burbling sound, particularly noticeable following the transients of harsh sounds such as trumpets and other brass.*

The CRX Decrackler

OPERATING THE CEDAR CRX

The term 'crackle' is used to describe many different audio phenomena - low level high density ticks, surface noise, buzz, and some forms of amplitude distortion. Fortunately, these exhibit similar characteristics, and the unique CEDAR 'Split & Recombine' process is capable of removing all of them. Consequently, the CRX has been designed to remove crackle, buzz, and distortion (all of which we now refer to as 'crackle').

The CRX will only function correctly if there are no major ticks or clicks received at the input. The CRX is NOT a de-clicker, and will rarely give acceptable results when used as such. If clicks exist on your material they must be removed by the CEDAR DCX or another CEDAR De-clicker before the CRX will give best results.

We represent the decrackle algorithm as follows:

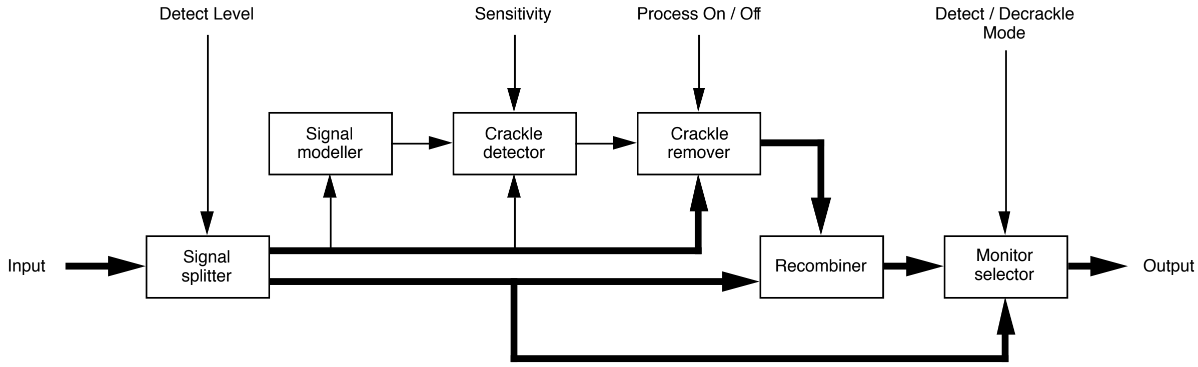

Figure 3: The CRX process overview

Mode

You may monitor the 'clean' signal that is output by the 'Signal splitter', bypassing the de-crackle algorithms. Alternatively, you may listen to the processed output. Press the MODE button to toggle between these settings.

Detect Level

This determines the degree to which the Splitter divides the input signal into 'clean genuine signal' and crackle.

Set the output monitor to 'DETECT'. Now increase or decrease the DETECT LEVEL by turning the knob clockwise or anticlockwise (respectively) until you find the lowest level at which the signal sounds smooth, and all the crackle has disappeared from the sound you are hearing. The audio will show many side effects at this stage - 'bubbliness', or an 'under-water' effect being the most common. Do not worry… this can be ignored because it will not mar the final output.

The Optimum DETECT LEVEL is the lowest level at which all the crackle is removed from the monitored signal in DETECT mode. If it is too low, crackle will remain in the processed output. If it is too high, distortion of the desired signal may result.

Sensitivity

Having set the DETECT LEVEL in DETECT mode, switch to DECRACKLE mode to monitor the final processed signal. Doing so will eliminate the bubbly side-effects generated by the signal splitter.

The SENSITIVITY now determines the amount of crackle removed by the CRX.

A low SENSITIVITY tells the system to remove only the most obvious crackle, while a higher SENSITIVITY also removes fine crackle, buzz, and distortion. To increase or decrease the SENSITIVITY turn knob clockwise or anticlockwise (respectively).

Note: *If the SENSITIVITY for a given piece of music is too high, distortion of the genuine signal may result.*

The DHX Dehisser

OPERATING THE CEDAR DHX

Broadband noise is a global degradation of a signal, and therefore differs significantly from clicks and crackle which, while often extending throughout the signal, are comprised of individual events localised in time. Consequently, dehissing requires a quite different process than those developed for click and crackle removal.

The DHX has three process controls that manipulate the noise reduction algorithm. This algorithm may be represented as follows:

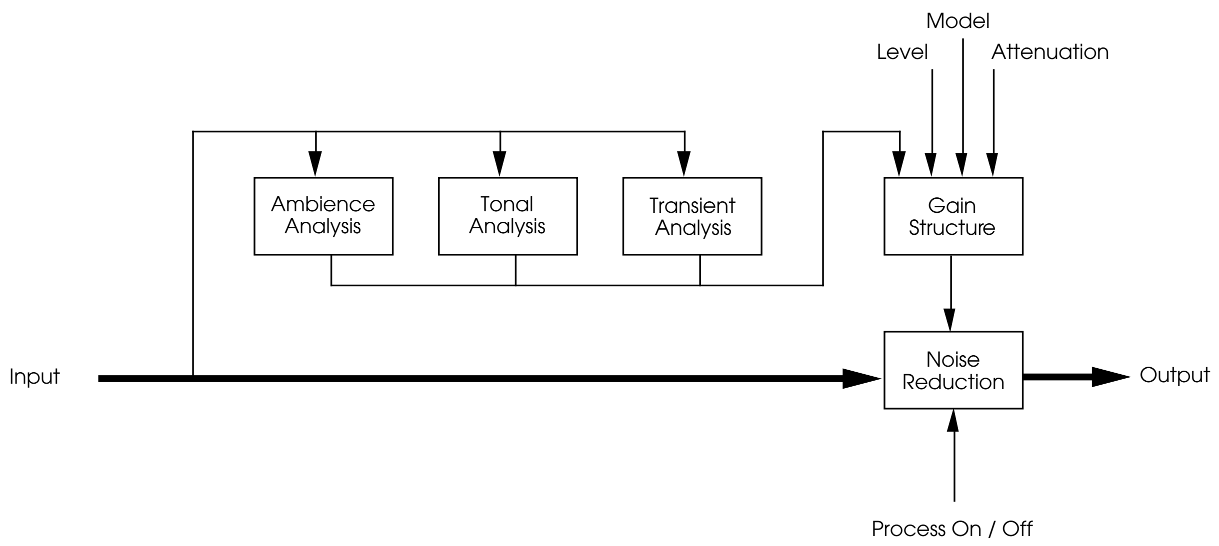

Figure 4: The DHX process overview

Noise Level

This tells the algorithm the amount of noise present in the signal. It is the most important control, and incorrect use will result in sub-standard results and/or unwanted side-effects.

Increase or decrease the NOISE LEVEL by turning it clockwise or anticlockwise (respectively).

- NOISE LEVEL = too low

  The DHX will not remove all the noise and may generate an artefact from the residual noise. This artefact is often described as 'twittering'.

- NOISE LEVEL = too high

  Some low-level signal will be treated as noise and will be attenuated. This results in compression and an artefact called 'glugging'.

Noise Reduction

This sets a limit on the amount of noise removed at any given frequency. Correct use will help ensure that none of the side-effects caused by over- or under- processing are heard in the output signal. You can often trade side-effects against the amount of hiss reduction.

Increase or decrease the NOISE REDUCTION by turning it clockwise or anticlockwise (respectively).

Note: *The side-effects introduced by inappropriate use of the NOISE LEVEL control are exaggerated when the NOISE REDUCTION is too high.*

Variance

(shown in the diagram as 'Model')

The optimum VARIANCE setting depends upon the nature of the signal and the hiss contained within it. Your choice of VARIANCE may also be influenced by any compromises you wish to make in order to achieve certain results.

You increase or decrease the VARIANCE by turning it clockwise or anticlockwise (respectively).

- VARIANCE = too high

  This will favour consistent compression instead of more unmusical or artificial-sounding artefacts. A high VARIANCE is appropriate if you want to remove a lot of noise.

- VARIANCE = too low

  This will minimise the compression of any low-level signals, but is susceptible to twitters and noise pumping. A low VARIANCE is suitable when you wish to remove a small amount of noise.

Note: *The NOISE LEVEL and the* VARIANCE *interact. Lower* VARIANCES *will usually require slightly higher NOISE LEVELS than higher* VARIANCES.

Tutorial

The DHX will only function correctly if there are no clicks or crackle received at the input. These must be removed by the CEDAR DCX and CRX (or other CEDAR processors) before the DHX will give best results.

1. Ensure that the DHX process is ON, and that the NOISE REDUCTION control is at its maximum.

2. Your first task is to find the appropriate setting for the NOISE LEVEL. This is the biggest influence on the quality of the processed signal.

   Starting with the NOISE LEVEL at its minimum, slowly increase the value. You will notice that little happens until, at some point, the amount of hiss decreases rapidly. (As this happens, you may also hear a side-effect known as 'twittering', and possible some gating of high frequencies.)

   The point at which this occurs is defined by the nature of the noise contained within the signal. With further increases, the twittering will disappear, but high frequency compression, gating, and a side-effect sometimes called 'glugging' will begin.

   The optimum value of value of the NOISE LEVEL is around the crossover point between the twitters and the glugging. Unfortunately, if the hiss is highly inconsistent there may be an overlap these effects, and no setting can be considered optimum.

3. If you want a particularly clean signal, or are processing material with highly coloured hiss, you should now set the VARIANCE to a high value. Similarly, if you want a to achieve an acceptable result quickly, use a high VARIANCE. If your defining criterion is signal quality, an intermediate VARIANCE (together with moderate attenuation) will usually offer best results. If you want a very clean signal and are prepared to accept the risk of low-level artefacts, you will obtain the best results using a low VARIANCE.

   Where possible, listen to the fidelity of any background instruments, and bear in mind the general feeling of 'air' or 'presence' in the signal, as well as the quality of any residual noise. In general, the DHX is less sensitive to the NOISE LEVEL setting when VARIANCE is high than it is when the VARIANCE is low.

4. Now adjust the NOISE REDUCTION to determine the amount of noise removed.

   Decrease the NOISE REDUCTION to its minimum, at which point the processed signal is identical to the unprocessed (because the NOISE REDUCTION control is limiting the noise removal to -0dB at every frequency - i.e. there is no effect).

   Now increase the NOISE REDUCTION to a level defined by the material and your taste. You will notice that, if the NOISE LEVEL is high, you can only increase the NOISE REDUCTION by a few dBs before the on-set of side-effects such as loss of transients and loss of high frequencies. If the NOISE LEVEL is low, the NOISE REDUCTION can be increased further, but with reduced effect.

As you increase the NOISE REDUCTION you may notice one of two detrimental effects:

- If there were ‘twitters’ present after step (2), and if you increase the NOISE REDUCTION beyond the optimal level for the material being processed, the twitters may be re-introduced as a form of high-frequency noise modulation.
- If there was a loss of high frequencies present after step (2), you will notice that this loss is gradually re-introduced as you increase the NOISE REDUCTION.

5. It is unlikely that the values of the three controls are already optimised, so you should return to step (2) to find a better NOISE LEVEL. Having done this you will, no doubt, wish to modify the VARIANCE and the NOISE REDUCTION settings.

   Continued fine-tuning of these controls will lead to excellent noise removal with few or no side-effects. However, the DHX is not a magic wand, and it may not be possible to restore some material beyond a certain point. Experience will enable you to judge whether you have removed as much noise as possible without unacceptable compromises.

Note: *In general, if the NOISE REDUCTION is small the NOISE LEVEL can be set lower and the VARIANCE can be set lower without introducing twitters or hiss modulation.*

PART 2:

An Introduction to Audio Restoration

THE CONCEPTS

There are an infinite number of processes that can affect the human perception of sound. For example, the live sound of an orchestra is dependent upon the venue, the audience and the local ambience. A recording of the same orchestra can be affected by a myriad more effects.

So what is the aim of a restoration engineer? The archival viewpoint suggests that such an engineer should present the listener with the most authentic reproduction of the original sound that can be obtained. But what about the creative influence of the recording engineer? With modern recordings, the ensemble sound often exists only on the recording medium, and many parts of it have probably never been through a microphone. Therefore the commercially minded engineer may, in contrast, attempt to generate a new recording more appropriate to its intended use. This use could be, for example, to please the public palate, or to represent accurately the sound of an era. Every restoration has its own criteria.

The algorithm designer is responsible for creating the facilities by which the restoration engineer generates new recordings from old. He or she does this by developing and implementing algorithms which remove unwanted sounds and/or effects present on the old recording.

CEDAR Audio provides a powerful set of restoration tools flexible enough to be used as the restoration engineer sees fit. However, it is the company's policy that the human ear should always be the final arbiter of sound quality; judgments based upon signal processing techniques are secondary considerations.

Algorithm Design

Any sound recording has been through a process history. Figure 5 shows an example of this.

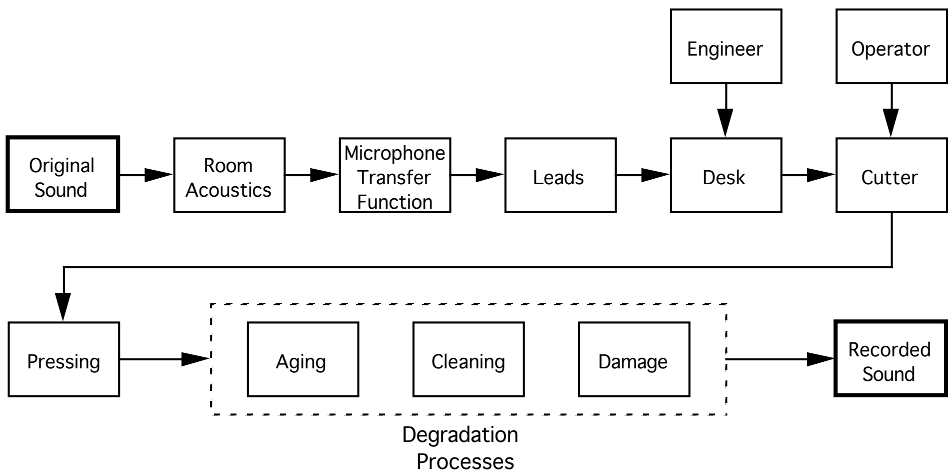

Figure 5: An example of the process history of a sound recording

A degradation process can be described by the elements shown in Figure 6:

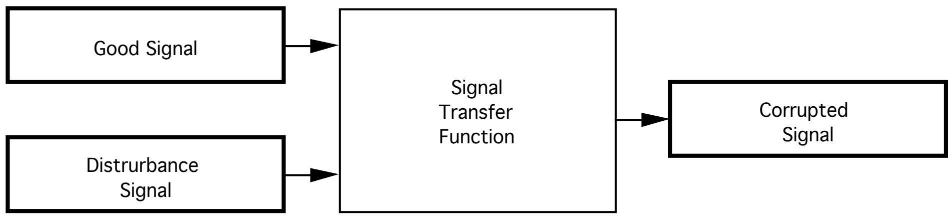

Figure 6: Degradation of a sound signal

A restoration algorithm uses assumptions about the behaviour of these elements to restore the good signal from the corrupted signal, and the quality of the restoration depends upon the quality of the assumptions. If you can glean more information about a recording, you can devise a better restoration algorithm.

For example, if the uncorrupted signal is a violin solo, you may wish to include in your algorithm the sonic characteristics of a violin. Such assumptions should then enable the algorithm to differentiate between the effect of the degradation process and the good signal. These differences can then be used to regenerate the good signal from the corrupt. But problems will arise when your assumptions fail: i.e. when the assumptions represent an incomplete picture of the true degradation process, or do not accurately represent the good signal being restored. The assumptions therefore limit the number of recordings to which your algorithm can be applied. You should not expect your solo violin algorithm to work for a full orchestra or a solo pan-pipe.

A musical signal is random in nature, as are most degrading processes. Information theory tells us that the mixing of two random signals represents a loss of information, and that a perfect restoration is then impossible. A restoration algorithm therefore has to generate the 'most likely' good signal given the information available. Curiously, such an algorithm represents an additional loss of information about the degrading action - it has removed most of it. This has important implications for further reprocessing should a better algorithm become available: it is almost always better to work from the original recording rather than from an earlier processed version. A good example of this is found when removing the crackle found on 78rpm records: while the recorded signal may only have a bandwidth of 12kHz, the crackle will exhibit the full bandwidth of the reproduction equipment, so low pass filtering (which may offer a subjective improvement in signal quality) removes a lot of information that would be of good use to a more advanced restoration algorithm.

The final test for any algorithm is the human ear. The questions to ask are:

- does the algorithm affect the perceived signal quality?
- upon what range of material will it work successfully?
- have the disturbances been removed/reduced?
- have any processing artefacts been introduced?
- is there an acceptable trade-off between the above points?

If the ear rejects the results as unsuitable for the intended application, then it will be necessary to redesign the algorithm.

Why Digital?

The public has now accepted digital sound, and the debates regarding the pros and cons of analogue vs. digital have waned. However, there are other considerations to be taken into account when designing a sonic processor.

While a digital sound signal can be made into a near-perfect reproduction of a band-limited analogue signal, there are sonic processes that are native to each domain, and each can only approximate the other. For simple ideas, an analogue implementation is often the most cost-effective solution. However, most audio concepts requiring higher mathematics are impractical in the analogue domain. Digital Signal Processing technology has been developed to implement these higher mathematical applications.

The advent of computers and hard disks enabled a restoration algorithm to take as much time as it needed to fulfill its purpose. Unfortunately, this was a mixed blessing because real-time applications have the advantage of allowing direct feed-back between the engineer's ear and the algorithm's user-controlled parameters. How quickly a track can be restored also has important commercial considerations.

CEDAR has, for some time, been implementing its most widely used processes as 'black boxes' with simple user-interfaces. The Series X, as the latest product of these developments, provides a range of powerful real-time digital restoration solutions of unparalleled simplicity.

Which Process?

The order in which restoration processes are carried out makes a great deal of difference to the quality of the result. The correct sequence is declick, decrackle, and then dehiss. This is because large clicks make it difficult for the de-crackle process to identify and remove the tiny clicks and crackles that constitute surface noise, buzz, and other such problems. Furthermore, all clicks and scratches are, in effect, tightly defined packets of noise. If clicks are presented to any dehiss process they confuse it and create unmusical side-effects. Furthermore, dehissing first will make it almost impossible to identify and remove clicks and scratches at a later time.

Decrackling should be the second process because small crackles will also cause problems for the dehisser.

Consequently, dehissing must always be the final process in the Series X restoration chain.

THE BACKGROUND TO SCRATCH REMOVAL

The most easily understood example of signal restoration is ‘scratch and click removal’. Many methods exist to remove these degradations: signal muting; channel swapping; Sample and Hold; linear interpolation; and complex interpolation.

Perhaps the most trivial de-clicker is the thorn needle sometimes used to replay old gramophone records. This, in essence, acts as a non-linear low-pass filter, but exhibits too many deficiencies to be of interest to a modern audio engineer.

The simplest analogue de-clicker is the attenuator. At the precise moment that a click is detected, this mutes both channels, thereby reducing its impact. The minimum duration of the mute (actually a high speed fade-out and fade-in) is typically 2.5ms, and in any case must be greater than the scratch length. Therefore, even a small number of mutes seriously affects the perceived sound quality. Also, since the method only seeks to make the clicks less obtrusive it does not in any way restore the underlying signal. In addition, it is only applicable when the energy contained within a click is very much greater than the energy within the signal. (Primitive though this concept is, it has also been utilised in the digital domain. When an error correction system is unable to cope with high density or long data errors many digital systems will mute.)

The most sophisticated analogue click-removal algorithm is used in a device known as a ‘switcher’. Using two sources of nearly identical signals (the opposite groove walls of a monaural record replayed using a stereophonic cartridge) this monitors for the cleaner signal (i.e. the lowest energy) and switches the output source between groove walls as appropriate. This method removes large clicks but, like all non-digital solutions, is unable to distinguish small ticks from genuine signal components. Also, the switcher assumes a perfect monaural source. If the groove walls differ significantly, or suffer degradation simultaneously, then the assumption (and therefore the restoration) fails.

Digital technology has made it possible to implement ideas that could not be realised using analogue electronics. The first of these is Sample & Hold (S&H) which, in many ways, is the same algorithm as used in a perfect muting system. However, instead of creating a signal plateau at zero amplitude, this method assumes that a plateau at the level of the most recent valid signal will be closer to the true signal. S&H removes the largest manifestations of clicks and scratches, but the resulting waveform contains unpleasant distortion and many audible ‘bumps’ and ‘pops’. While these low amplitude thumps may be preferable to the high amplitude clicks of the untreated data, the signal will show signs of severe break-up if the click density is high. Many listeners complain that these artefacts and side effects (as, for example, implemented in the error correction of domestic CD players) are more unpleasant than the clicks that they replace.

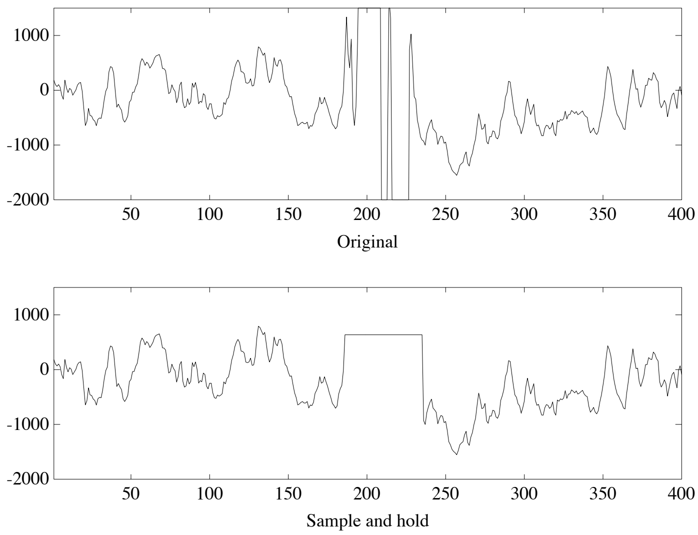

Figure 7: Sample and Hold

Linear interpolation is the next stage following sample and hold. In this algorithm the corrupted data is replaced by a straight line joining the last good sample and the next available good sample. It is impractical to implement this method in the analogue domain, but relatively straightforward in the digital domain.

Figure 8 shows such an interpolation. The audible result of this method is less offensive than S&H, but suffers from low frequency artefacts and a reduction in audio bandwidth over the interpolated region.

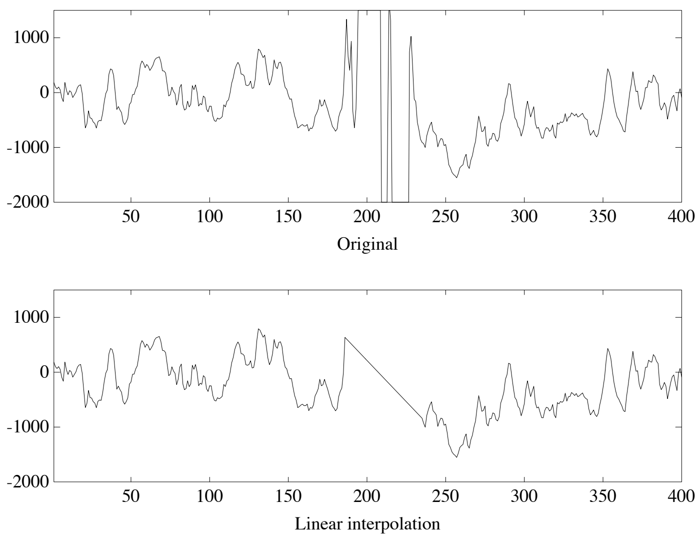

Figure 8: Straight line interpolation following sophisticated scratch detection algorithm

Note: Sample and hold takes account of one adjacent good sample, whereas linear interpolation takes account of 2 adjacent good samples and the elapsed time between them.

Musical and other real-world signals are, by their very nature, resonant. In contrast, clicks are not. We can use this distinction to design a click detector, and a signal interpolator (which will rebuild the signal once the clicks have been removed) using the technique of signal modelling.

CEDAR's signal modeller is an advanced algorithmic process that analyses the signal in the digital domain, then uses the information thus obtained to detect clicks (signal elements which do not fit the model), and replace them with generated signal which does fit the model. The more information that can be analysed (as defined by a parameter known as 'the order') the better the detection and interpolation can be.

A high order interpolation is nearly indistinguishable from the surrounding signal, and the results are good enough to fool the human ear in almost all cases. The CEDAR DCX uses a sophisticated scratch detection algorithm allied to a high order interpolator to detect and remove up to 2500 clicks per channel per second in real time. The performance of this system is so good that, in almost all cases, it is not possible to hear that the original signal was damaged prior to the restoration.

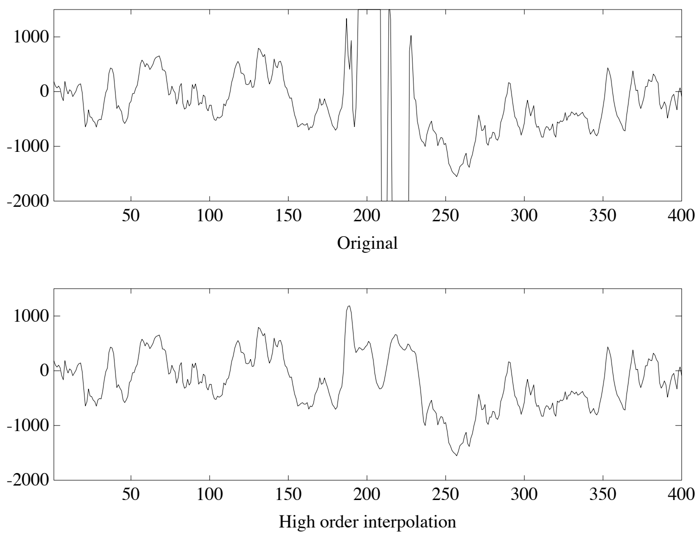

Figure 9: A high order interpolation following a sophisticated click detection algorithm

THE BACKGROUND TO CRACKLE REMOVAL

Whereas clicks and scratches are clearly visible to the eye when a signal containing them is shown on a waveform display, a crackly signal may often be severely degraded, but the waveform will appear to be clean. This is a consequence of the shapes, densities, and amplitudes of the events constituting the degradation. As a result, declicking techniques will not perform satisfactory decrackling. A new method must be found.

CEDAR's research has shown that, once digitised, a crackly signal may be divided into two components. The first of these contains the bulk of the clean (desired) signal. The second (the 'split' signal) contains all the degradation plus the residual of the clean signal.

Once the split signal has been obtained, the crackle becomes a much larger and more obvious element within the waveform to be processed. This means that the crackle detector has a much better chance of accurately identifying the undesired events, and the interpolator can more accurately estimate the genuine sound that existed before the damage occurred.

Once the split signal has been restored, it may be recombined with first component in order to recreate the original signal, but without the crackle.

Figure 10 shows a degraded (crackly) input signal, and the audibly clean restoration obtained using the Split and Recombine process. As you can see, and quite unlike the declicked examples shown before, the results of the restoration (while audibly very significant) are visually almost unnoticeable.

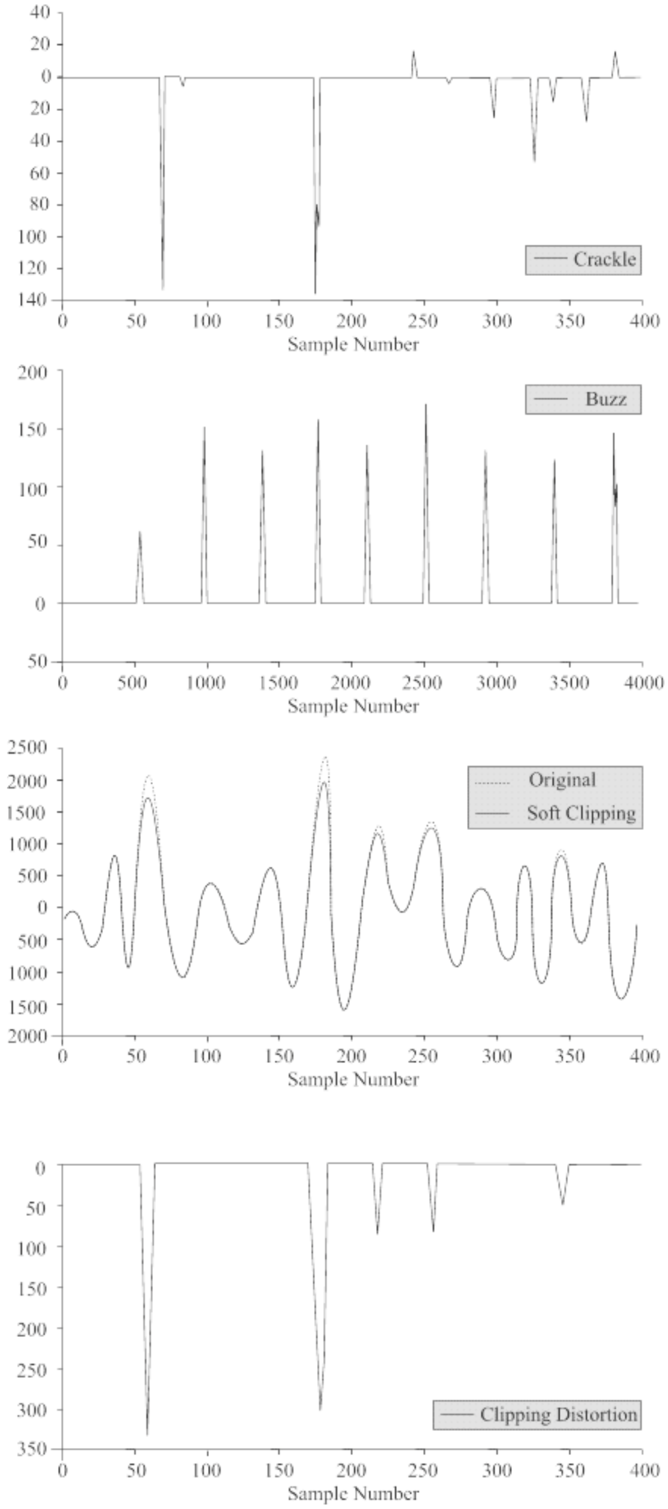

Figure 10: Removing crackle

THE BACKGROUND TO NOISE REMOVAL

Broadband noise adds (or subtracts) a random amplitude at all times to (or from) all frequencies within the audio spectrum. We cannot remove it from existing signals using dual-ended processes that encode during recording and decode upon playback. These limit the accumulation of extra noise added by the limitations of analogue recording tape, but do nothing to remove it from signals that already contain it.

Since broadband noise is most intrusive at high frequencies, the first stage in the evolution of noise reduction is the low-pass filter, which removes a proportion of the signal above its cut-off frequency. Unfortunately, if, at any given frequency, you reduce the amplitude of the noise by, say, 6dB, you also reduce the desired signal by the same amount. So a low-pass filter may clean antique 78s (which have little or no high frequency content) but even then, only at a cost.

Dynamic filters are devices in which the cut-off frequency moves dynamically according to the signal content, thus removing high frequencies when there is no signal present, but leaving them untouched when the noise is being masked by genuine high frequency signal. But such devices are limited because they only remove the noise that exists above the cut-off, which is itself an inaccurate representation of the highest frequencies contained in the genuine signal. Secondly, and in common with the simple filter, they have roll-offs of the order -12dB/octave or -6dB/octave, so they allow some high frequencies through. And thirdly, even though the filters are designed to track quickly, they still round off transients and dull the wanted signal.

Instead of altering the frequency response of the signal, how about changing the level in some way? Consider: if constant amplitude noise is present then, if the total amplitude drops to the noise level, we can assume that no genuine signal is present. While there are many flaws in this argument, it suggests a device that will eliminate some noise: a Noise Gate. This detects when the signal drops below a threshold set by the user, and then cuts off the signal entirely. There are many enhancements to the idea (added to limit damaging side-effects) but the principle holds: if the total signal drops below the threshold, the gate shuts and removes all the noise. Unfortunately, an ‘open’ gate removes no noise whatsoever.

An Expander is another device with a threshold control, but this applies a progressive gain reduction, the amount of which is determined by the user. For example, if a signal drops 3dB below the threshold, the Expander may reduce the signal volume by 6dB, 12dB, or any other figure, depending upon the expansion ratio. Unfortunately, the subjective difference between the gate and the expander is small. A multi-band expander separates the audio spectrum into a number of bands, treating each as an individual signal. But multi-band units are still unable to distinguish accurately between genuine signal and noise. They still act upon the inaccurate assumption that, if the total signal level approaches its noise floor, all that is present is noise. Consequently, even the most sophisticated expanders remove genuine signal. Furthermore, poor band separation filters severely limit performance. The consequences are loss of high frequencies, loss of ambience, and degradation of hard transients. Some units feature a combination of dynamic filtering, expansion, and even compression and excitation - effects which have been included in order to obscure some of the side-effects of the processes, but these are only partially successful.

All the processes so far described are ‘ratio’ operations - that is, if (at any given frequency) you remove half the noise, you remove half the signal; if you remove 3/4 of the noise, you remove 3/4 of the signal... and so on. Consider now a signal that has, at a given frequency, 100 units of ‘volume’ on an arbitrary scale. By measuring the noise content of that signal during an otherwise silent passage, you can determine that there are, say, 20 units of noise present at that frequency. It should be possible to remove this noise by removing 20% of the signal. But what if, a moment later, the total ‘volume’ of the signal drops to 40 units? An analogue filter, removing

20% of the signal, will remove 8 units. On the other hand, a subtractive filter (which is practical only in the digital domain) will still remove the full 20 units - a reduction of 50%. This is what we want, because the noise at this moment represents 50% of the total signal. (See figure 11.) This Spectral Subtraction becomes useful when a DSP is used to split the signal into hundreds of bands. You can then be very precise about how much noise you remove but, if this sounds too good to be true, it is. The noise spectrum (the sonic fingerprint) can only be measured if there is an otherwise silent passage within the music, and if it is not accurate you will hear unpleasant side-effects. But let's assume that you have obtained a perfect fingerprint. Even then, experience shows that spectral subtraction produces dry and dull results with unacceptable artefacts. This is, in part, because the fingerprint is a snapshot of the random noise, accurate only at the instant at which it is taken. Because the noise is constantly changing, the subtractive algorithm is deriving its result from inappropriate data.

Many researchers have investigated enhancements designed to overcome these pitfalls. CEDAR Audio's developments are embodied in a process that updates the noise fingerprint every few samples, thus tracking variations in the noise content. This prevents the compression of incoming transients, and distinguishes between noise and, for example, reverberation in the wanted signal. Other features help to avoid many of the problems of basic spectral subtraction, and they allow you to remove noise without undue damage to the genuine signal. But this algorithm still requires a complex user-interface, and cannot be implemented in a stand-alone box. Removing the requirement for a spectral fingerprint simplifies the interface greatly, but requires an algorithm capable of autonomous determination of the noise content. Which brings us, finally, to the CEDAR DHX, a stand-alone module that will analyses the noise content of a signal and applies noise reduction with minimal effort on the part of the user.

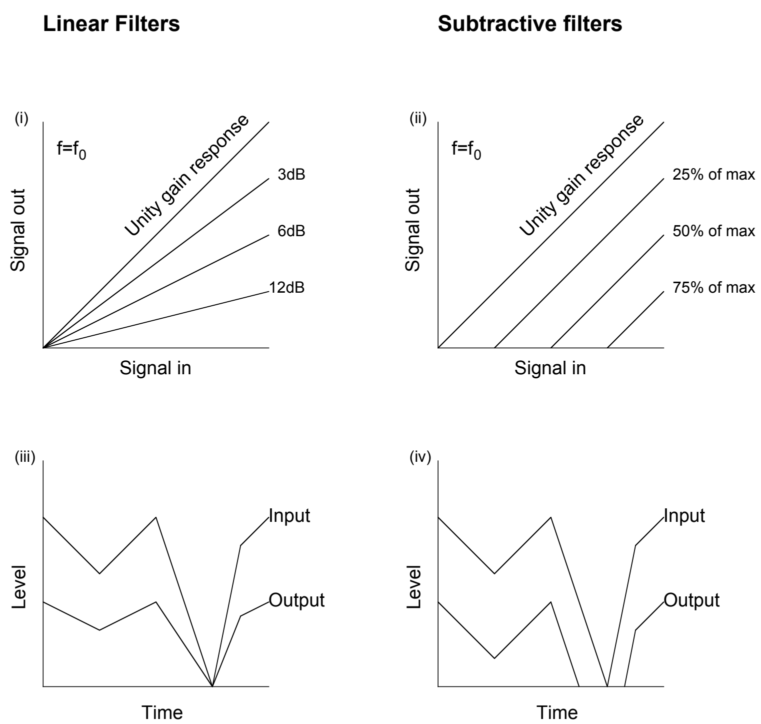

Figure 11: A comparison of 'ratio' filtering and Spectral Subtraction

CEDAR Series X

SERIAL NUMBER:	
INSPECTION DATE:	
QC ENGINEER:	

Designed and Manufactured by

CEDAR Audio Ltd
20 Home End
Fulbourn
Cambridge
CB21 5BS
United Kingdom

Tel: +44 (0) 1223 881771
Fax: +44 (0) 1223 881778
Email: info@cedaraudio.com