The Duo Series

declickle and auto dehiss

Version 1.2

19 September 2004

Copyright CEDAR Audio Ltd, © 2004

CEDAR is a registered trademark of CEDAR Audio Ltd

Duo, declickle and auto dehiss are trademarks of CEDAR Audio Ltd

E&OE, Subject to revision at the Company's sole discretion

interstage
Phistersvej 31, 2900 Hellerup, Danmark
Telefon 3946 0000, fax 3946 0040
www.interstage.dk
- pro audio with a smile

This page left blank

Table Of Contents

This page left blank

Introducing the CEDAR Duos

CEDAR Duo declickle module

CEDAR Duo auto dehiss module

Thank you for purchasing this CEDAR Duo 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 considerations in the Duo design, and both units are intended for professional use. Features include the following:

■ Flexibility

Both Duo models handle a wide range of audio restoration requirements.

■ Portability

Each Duo is a compact, lightweight unit designed to fit a standard 1U rack space.

■ Audio interfaces

Duos incorporate 96kHz, 24-bit digital audio interfaces conforming to the AES/EBU and SPDIF standards.

■ Power

Their universal power supplies means that Duos work anywhere in the world.

■ Processor

Two, powerful SHARC processors ensure that a Duo will handle the most complex processing requirements.

Specifications and CE Certificates

General:

Power supply: 85–260VAC; 50–60Hz

Power consumption: 15W (standby 1W)

Overall dimensions: 45 x 483 x 240mm

Weight: 3kg (net); 4kg (gross)

I/O type: Digital PCM

I/O resolution: 24 bits

Sample rate: 30–100kHz

Data format: SPDIF or AES/EBU

Processor: 400 Mflops (peak)

Group delay (milliseconds):

	declickle (processing)	**auto dehiss** (processing)
44.1kHz	190ms	463ms
48kHz	174ms	458ms
88.2kHz	95ms	461ms
96kHz	87ms	456ms

In Bypass, each unit has a latency of less than five samples, irrespective of sample rate.

EMC Regulations:

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

Declaration of conformity:

Date of issue: 1 April 2004

Equipment: CEDAR Duo

Manufacturer: CEDAR Audio Ltd

Address: 20 Home End, Fulbourn, Cambridge CB1 5BS, UK

This is to certify that the aforementioned equipment, when used in accordance with the instructions in this manual, fully conforms to the protection requirements of the following EC Council Directives: on the approximation of the laws of the member states relating to:

■ 89/336/EEC Electromagnetic Compatibility

Applicable standards: EN 50081-1:92

EN 50082-1:92

■ 73/23/EEC Low Voltage Equipment

Applicable standard: BSEN 60-065:1994

Contact Information

Headquarters:

Address:	CEDAR Audio Limited, 20 Home End, Fulbourn, Cambridge, CB1 5BS, United Kingdom
Telephone and Fax:	T: +44 1223 881771 F: +44 1223 881778
Email:	support@cedaraudio.com

USA Office:

Address:	CEDAR Audio USA, 43 Deerfield Road, Portland, ME 04101-1805, USA
Telephone and Fax:	T: +1 207 828 0024 F: +1 207 773 2422
Email:	cedarusa@cedaraudio.com

Asian Office:

Address:	CEDAR Asia (Mailing address) Attention: David William Robinson, Tanglin Post Office, PO Box 282, Singapore 912410
Telephone and Fax:	T: +66 1 822 9227 F: +66 2 619 1277
Email:	asia@cedaraudio.com

Web:

Web:	www.cedaraudio.com

Worldwide Dealer List:

For an up-to-date dealer list, please visit www.cedaraudio.com

Technical Support

We hope that you will never need to read this page, but should you experience difficulties with your CEDAR Duo, please contact your local dealer or CEDAR office. Alternatively, you may send an email to support@cedaraudio.com. In either case, please provide the following details:

▪ Your Duo hardware and software version numbers. (See General Set-up and Connections)

▪ A precise description of the problem

▪ All details shown in the status page.

Thank you.

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 Duo to reach ambient temperatures naturally 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. A Duo 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 it is used as a free-standing unit, the supplied rubber feet should be fixed to the base of the unit.

■ Power sources

Each Duo features a universal power supply that 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.

A Duo 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 the unit only with a dry cloth. Never use abrasive pads or liquid cleaners such as alcohol or benzene.

■ Damage requiring service

A Duo 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.

Getting Started

Unpacking

Unpack the Duo carefully. Save the carton and all packing materials since you may need them to transport the unit in the future. In addition to the Duo 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:

- the nominal temperature should be maintained between 5° and 35° Celsius
- relative humidity should be in the range 30% to 80% non-condensing
- strong magnetic fields should not exist nearby.

Front Panel Indicators and Controls

Power switch and power LEDs

■ When OFF, the Duo is in standby mode.

■ When ON, the Duo is powered up and operational.

The power status is reported by the power LEDs on the right of the front panel.

Warning: The unit is not powered off completely unless the power lead is removed.

Screen, soft buttons and spin wheel

The set up and operation of the Duo is determined using the illuminated 'soft' buttons and the spin wheel.

■ When a control or menu option is available, the button alongside it will be faintly illuminated.

■ When an option is selected for adjustment, the button alongside it will be fully illuminated. Adjustment is carried out either by (i) repeated pressing of the illuminated button, or (ii) rotating the spinwheel in either direction, as appropriate.

■ If no menu or control is associated with a button, it is not illuminated.

General Set-up and Connections

Basic Set-up

When you switch on the Duo, it will present one of the following screens to you. These are the Main screens for declickle and auto dehiss, respectively. To set up the unit and access a number of less often required parameters, press the MENU option.

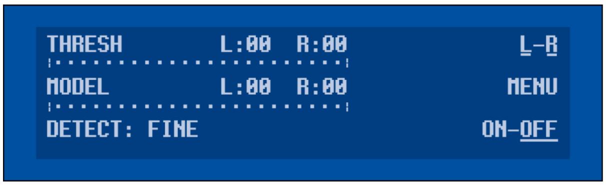

declickle main screen

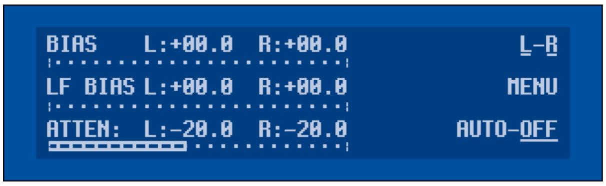

auto dehiss main screen (auto mode)

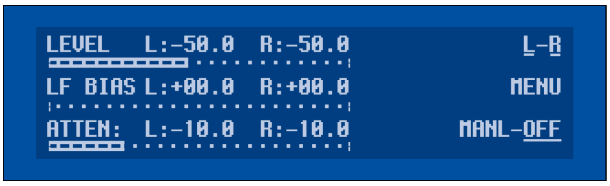

auto dehiss main screen (manual mode)

Menu

You will now be presented with the MENU screen:

declickle menu screen

auto dehiss menu screen

These offer five or six options, depending upon the product. Press the desired option to access the screen of your choice.

■ SETUP: Enter the set-up and reset screen

■ AUDIO I/O: Enter the audio set-up and monitoring screen

■ PROCESS MODE: Enter the AUTO DEHISS process mode screen.

■ MEMORY: Enter the memory screens to name, save, recall and delete user memories

■ CLOSE: Returns you to the main operation screen

■ STATUS: Provides information about the operation of the unit

Setup

■ Contrast

Press the CONTRAST option and use the spinwheel to alter the contrast of the display to suit your viewing angle.

■ MIDI Channel

Press the MIDI CH option and use the spinwheel to select the desired MIDI channel for remote control.

■ Reset

Press RESET to reset the Duo to the factory defaults. This will clear all user-defined parameters and all the memories. You will be asked to CANCEL or CONFIRM the RESET.

Pressing CANCEL will return you to the Setup screen.

Pressing CONFIRM will reset the unit to the factory defaults and return you to the Main screen.

■ Close

Press CLOSE to return to the Main screen.

■ Version number

Reports the version number of the software installed in this Duo.

■ Serial number

Reports the serial number of this unit.

Audio I/O

■ Output

Press the OUTPUT option repeatedly to cycle through 16, 18, 20 and 24-bit output wordlength options. You should choose the option most suitable for the next device in the audio chain.

The internal 32-bit wordlengths are dithered to the output wordlength using TPDF dithering.

■ Gain

Press the GAIN option and use the spinwheel to adjust the gain of the unit between -10dB and +10dB in steps of 0.1dB.

■ Level

This displays the signal level presented to the unit's input. A 'peak hold' is implemented, and the unit will display the word CLIP! if it detects digital ‘overs’.

Pressing the LEVEL button in this screen will reset the peak hold.

■ Bypass

Press the BYPASS option to bypass the Duo's processing. In this mode, the Duo is 24-bit transparent. The latency is reduced to less than 10 samples.

■ Close

Press CLOSE to return to the Main screen.

Physical Connections

Audio

Your Duo 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

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 Duo using a single cable terminated with an RCA (or 'phono') plug. The SPDIF output of the Duo 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 Duo using a single cable terminated with an XLR plug. The AES/EBU output of the Duo should be connected to the AES/EBU input of a recording device or external DAC.

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

Input Selection

The Duo will automatically switch to whichever input is receiving a valid signal. If both inputs are active when you switch on, the Duo may not select the input you require, so it is advisable only to present signal to one input at a time.

MIDI

The Duo offers MIDI In, Out and Thru. These operate in the conventional manner.

Use MIDI to store and recall the memories, and to control the Duo using external equipment. See MEMORIES.

USB

The Duo offers a USB socket. This is for factory use only.

Status

The **Menu/Status** page provides important information about the operation of your unit, and technical support information in the event of a failure.

Status

■ Input

There are three options:

■ LOCK XLR: A valid digital signal is being received at the XLR (AES/EBU) input and the Duo has locked to this

■ LOCK RCA: A valid digital signal is being received at the RCA/phono (SPDIF) input and the Duo has locked to this

■ OFF LOCK: No valid digital signal is being received at either the AES/EBU input or the RCA/phono (SPDIF) input and the Duo is not locked to incoming audio

■ FS (Sample rate)

This shows the measured sample rate of incoming audio, and is accurate to approximately 0.01%.

This field is blank if the Duo is OFF LOCK

■ DSPn: NN NN mm

This provides information about each of the SHARC processors within the Duo.

■ NN NN: These are status codes that should be reported to CEDAR Audio in the event of a failure

■ mm: This field reports two messages, "OK" or "ERROR"

■ Close

Press CLOSE to return to the Main screen.

Memories

Your Duo offers 99 memory locations in addition to the current workspace. The workspace is volatile, so it is useful to be able to store and recall set-ups. You do so using the MEMORY screens, as follows.

Press the MEMORY option in the MENU screen to access the memory screens. When you press this for the first time, the following screen will appear:

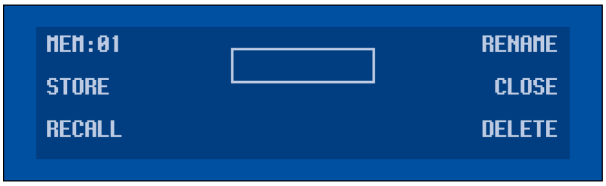

There are six options:

Mem:xx

The number XX shown here tells you which memory location is being addressed. To change memory location, press the soft button associated with this option and use the spinwheel to select other memories.

Store

Press this to save the current workspace in the location MEM:XX shown above.

Recall

Press this to load the settings in MEM:XX into the current workspace.

Rename

When you press RENAME, the following screen will appear:

■ CLEAR: Press this to clear the current name

■ NEXT: Moves the edit cursor to the right

■ DELETE: Deletes the current character (identified by an underline) and moves the edit cursor to the left

▪ CANCEL: Exits from this screen without making any changes, and returns you to the MEMORY screen

▪ DONE: Saves the name changes, exits from this screen, and returns you to the MEMORY screen

You may change the identified character by rotating the spinwheel clockwise or anticlockwise, as appropriate.

If you do not choose a name for a memory, it will be set to MEMORYXX by default:

If you clear a name and do not replace it, the Duo will insert UN-NAMED as the default.

Close

Closes this screen and returns you to the Main screen.

Delete

Deletes the contents of the current memory, MEM:XX. Use this with care; you will not be prompted to confirm the operation.

When the Duo is first switched on, it recalls the settings held in MEM:01. If you wish to create a default set-up, store it here.

Remote Control (MIDI) Implementation

Standard MIDI control messages may be used to control all the Duos' audio parameters. The only system parameters that can be altered from the front panel but which cannot be edited via MIDI are the display contrast and the MIDI channel.

- Audio parameters are adjusted by sending MIDI controller change messages.

- User memories may be stored and recalled using MIDI patch change messages.

MIDI Controller Numbers (All models)

Command	CC#	Comments
Bypass	16	■ 0-63 = Bypass OFF ■ 64-127 = Bypass ON
Makeup Gain	17	0-127 maps to -10dB to +10dB
Wordlength	18	■ 0-31 = 16-bit ■ 32-63 = 18-bit ■ 64-95 = 20-bit ■ 96-127 = 24-bit
On/Off	19	■ 0-63 = Processing OFF ■ 64-127 = Processing ON

MIDI Controller Numbers (Declickle only)

Command	CC#	Comments
Threshold (left)	20	0-127 is scaled to 0-99
Threshold (right)	21	0-127 is scaled to 0-99
Model (left)	22	0-127 is scaled to 0-99
Model (right)	23	0-127 is scaled to 0-99
Detect	24	■ 0-42 = fine ■ 43-84 = medium ■ 85-127 = coarse

MIDI Controller Numbers (Auto dehiss only)

Command	CC#	Comments
Process mode	15	■ 0-63 = Left/Right ■ 64-127 = Mid & Side
Level A (left or mid)	20	0-127 maps to -99.5dB to 0dB
Level B (right or side)	21	0-127 maps to -99.5dB to 0dB
Auto bias A	22	0-127 maps to +10.0 to -03.0
Auto bias B	23	0-127 maps to +10.0 to -03.0
LF Bias A	24	0-127 maps to +20.0 to -20.0
LF Bias B	25	0-127 maps to +20.0 to -20.0
Atten A	26	0-127 maps to 0.0 to -40.0
Atten B	27	0-127 maps to 0.0 to -40.0
Auto mode on/off	28	■ 0-63 = Auto mode OFF ■ 64-127 = Auto mode ON

Patch change messages

Ordinarily, a patch change message (Patch=NN) causes the user memory of the corresponding number (NN) to be recalled. There are 99 user memories, and patch numbers outside the range 1-99 are ignored, except for the following special cases:

■ Sending Patch=0 causes the audio parameters to return to their factory default values

■ Sending Patch=125 followed by Patch=NN causes memory NN to be deleted

■ Sending Patch=126 followed by Patch=NN saves the current settings to memory NN

■ Sending Patch=127 after patch 125 or 126 acts as a "CANCEL" function

Declickle: Removing impulsive noises

You will find all the parameters specific to declicking and decrackling on the Main screen:

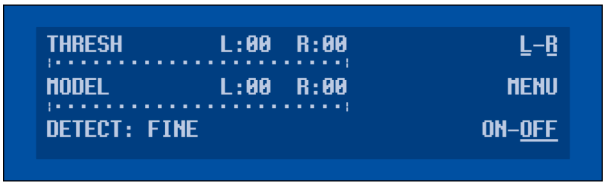

There are six options:

- THRESH
- MODEL
- DETECT
- L-R
- MENU
- ON-OFF

General operation:

- Select which channel or channels you are adjusting by repeatedly pressing the L-R button.

- Select whether processing is ON or OFF by pressing the soft button associated with the ON-OFF option.

When the Duo is in Bypass, this is indicated by **BYP** and the ON-OFF button is disabled. To switch Bypass off, use the Audio I/O page.

Threshold

The threshold controls the sensitivity of the process. With the threshold set high, Declickle will remove only the largest clicks and scratches. A lower threshold will also remove smaller ticks, clicks and crackle. If you reduce the value too far, the wanted audio may exhibit mild damage in the form of a slight instability in the background signal. You should, therefore, always attempt to find the highest value at which the unwanted noise is removed.

The threshold ranges from 00 to 99 on an arbitrary scale, and the current values are displayed both numerically and on the horizontal bar-graphs.

Model

The model value determines the way in which Declickle responds to the high frequency harmonics of incoming signals. With a high modelling value, the process models the high frequency content very accurately, making Declickle far better at removing clicks and crackle in the presence of instruments such as trumpets and violins.

The model value ranges from 00 to 99 on an arbitrary scale, and the current values are displayed both numerically and on the horizontal bar-graphs.

Detect

Allows you to select the most appropriate detection mode for the material being processed. The options are:

■ fine

Designed for audio signals that exhibit clearly defined clicks. It is not suitable for restoring material with 'grungier' crackle. This mode is the least likely to cause any damage when used inappropriately, and will maintain the brightness of the original signal. Use it for removing clicks from high bandwidth signals such as obtained from high quality vinyl and signals recorded directly to digital media.

■ medium

A compromise between fine and coarse settings, this is the most commonly used mode.

■ coarse

Designed to detect rounded, 'grungier' clicks and crackle, this mode will detect more problems, but with a greater risk of signal degradation when used inappropriately. This damage will be heard as a loss of high frequency content in very bright passages such as the rasp of a trumpet. This mode is particularly suitable for restoring low quality media that exhibit a great deal of damage, such as optical film soundtracks.

Declickle tutorial

Detect

Your first job is to choose the correct detection mode for the material to be restored. Listen to the audio, and decide whether it contains well-defined clicks, whether it suffers from grungy clicks and crackle, or whether the problems lie somewhere between. When you have done so, select the fine, coarse, or medium modes (respectively) as you see fit.

The best results for a given piece of material will be obtained when - if possible - you use a finer mode. So, for example, if you feel that the problems lie between 'coarse' and 'medium', you should in preference try to find settings that work with 'medium'. Likewise with 'medium' and 'fine'... try to use 'fine' whenever possible. This will ensure that the high frequency and transient content of the wanted audio is best preserved.

Threshold

Once you have chosen the detection mode, you must find a suitable value for the threshold. You can often leave this at its default value of 10, reducing it only when removing fine crackle and buzzes, or increasing it if you notice any dullness or distortion.

If you reduce the threshold too far, the wanted audio may exhibit mild damage in the form of a slight instability in background sounds and any background noise present. You should always attempt to find the highest value at which the unwanted clicks and crackle are removed.

The correct threshold will vary with the detector mode chosen. Fine mode will require a lower threshold than coarse mode if it is to detect the same number of clicks and crackle.

Model

Once you have selected the detection mode and threshold, you must find the most appropriate value for the modelling. For most purposes, you can leave this at its default value of 99, which allows Declickle to remove clicks and crackle accurately in the presence of difficult signals such as trumpets and violins.

However, if the audio contains a high-frequency buzz (regularly spaced clicks) the harmonic modelling within Declickle will recognise this inappropriately as wanted signal. and attempt to ensure that this is not removed. In this case, you should reduce the modelling parameter to a suitable value. Always use the highest value that removes the problem. This will minimise the risk of damage to the high frequency content of the genuine signal.

An Introduction to Click and Crackle Removal

Clicks

Clicks are impulsive noises of short duration.

The simplest audio declicker is an attenuator which, at the precise moment that a click is detected, mutes the audio, thereby reducing the impact of the click. The minimum duration of the mute (actually a high speed fade-out and fade-in) is typically 2.5ms, and even a small number of mutes seriously affect the perceived sound quality. Also, since the method only seeks to make the clicks less obtrusive, it does not 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.

A more 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 and switches between them as appropriate. This removes large clicks but, like all non-digital solutions, is unable to distinguish small ticks from genuine signal. 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 audio contains unpleasant distortion and many audible 'bumps' and 'pops'. While low amplitude thumps may be preferable to the high amplitude clicks of 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 are more unpleasant than the clicks that they replace.

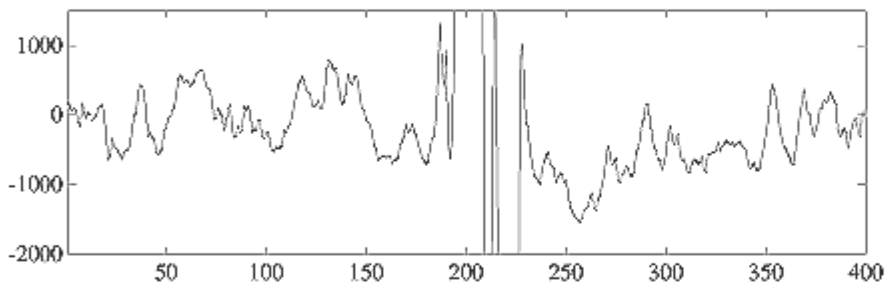

Original

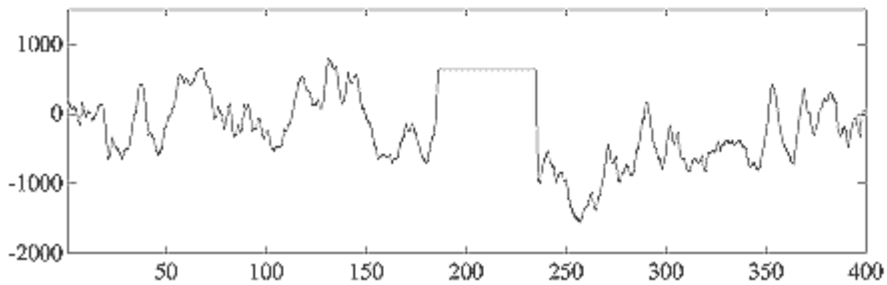

Sample and hold

The next stage is Linear Interpolation. Whereas Sample and Hold takes account of one good sample adjacent to the scratch, linear interpolation takes account of two adjacent good samples and the elapsed time between them. In this algorithm the corrupted data is replaced by a straight line joining the last good sample and the next available good sample. The result is less offensive than Sample and Hold, but suffers from low frequency artefacts and a reduction in audio bandwidth over the interpolated region.

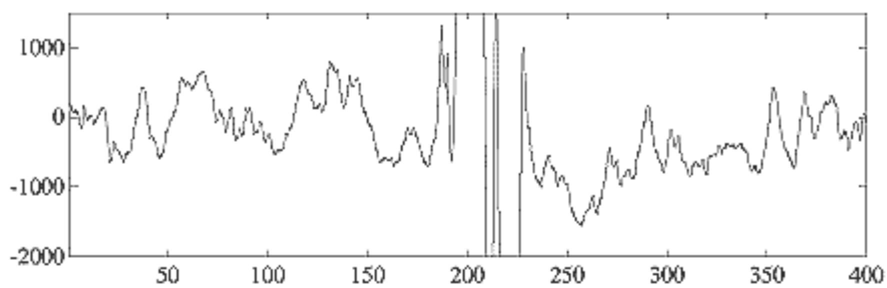

Original

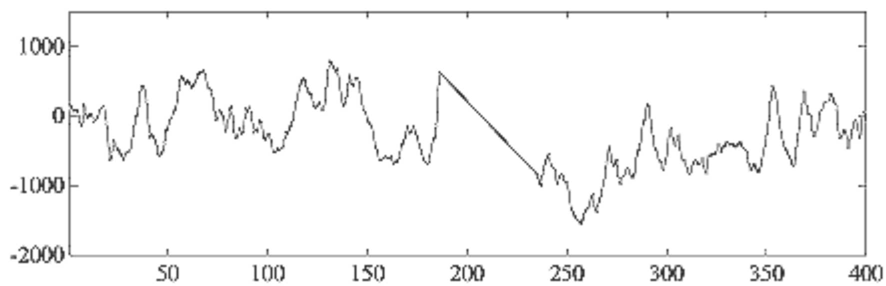

Linear interpolation

Modern interpolators are based on signal modelling. By analysing the signal over a region of several milliseconds the algorithm can build a model of the underlying signal and use the information to replace each click with an interpolation that fits this model.

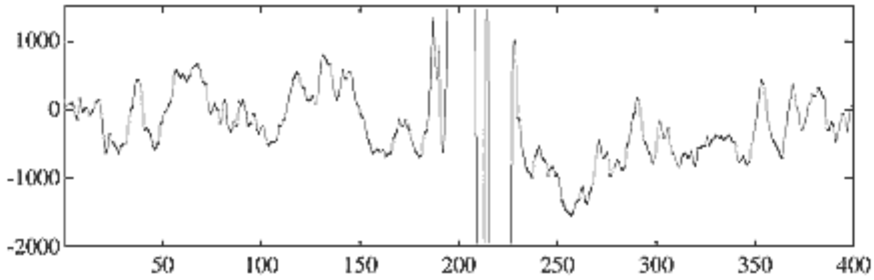

Original

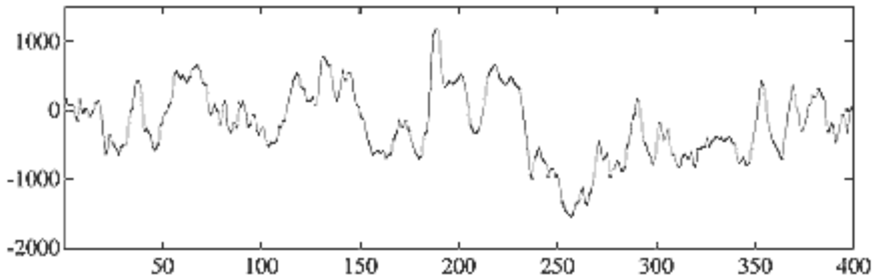

High order interpolation

Crackle

Unlike a standard declicker, Declickle is able to detect high density, low amplitude, short duration ticks within the signal.

This type of small impulsive disturbance is the cause of several common signal defects. If the disturbances are randomly distributed, the degradation will sound like crackle or a 'chip-fryer' type of noise. If they are regularly spaced, the degradation will sound like a buzz. Alternatively, if they are correlated with the signal, the degradation will sound like distortion. Declickle will correct all of these problems, so the term 'crackle removal' is slightly misleading.

There are other types of distortion (such as tape saturation) that Declickle does not address. This is because these distortions are not caused by short-duration impulsive disturbances, but by other defects that Declickle is not designed to detect.

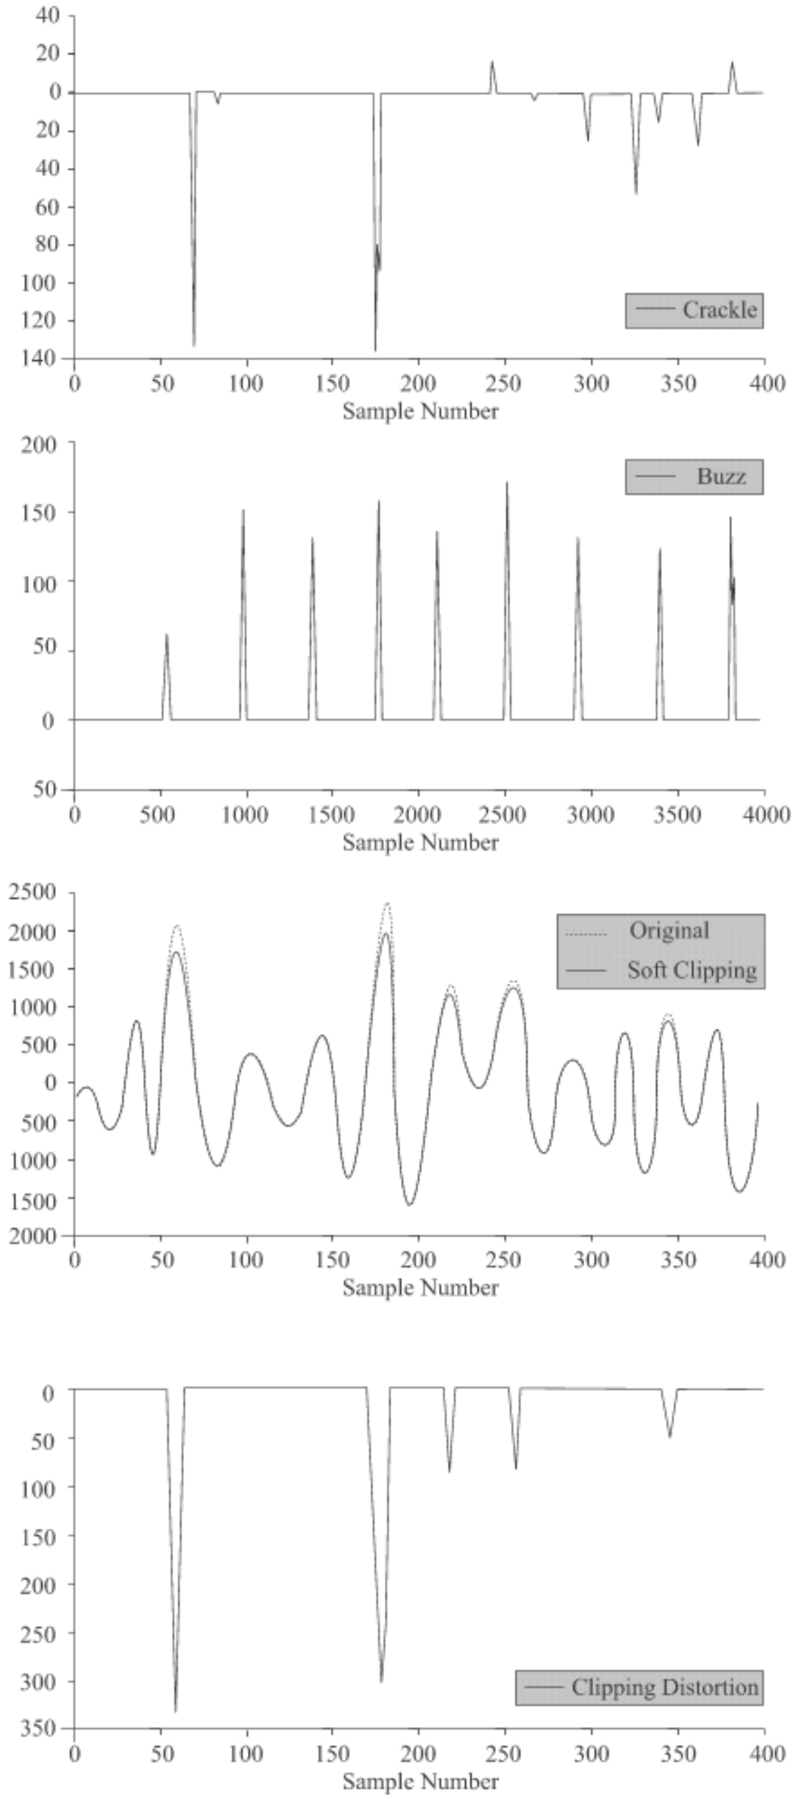

Various Forms of "Crackly" Signal Defects

Unfortunately, interpolating each impulse using a basic declicker is not satisfactory because of the fundamental nature of the crackle. A standard click-detection algorithm will miss smaller impulses, and it will fail to cope with the high density of those impulses. Furthermore, standard declick algorithms assume that a click or scratch has totally corrupted the audio and that there is no useful information about the signal during the click. Since the disruptions that cause crackle are very small they may be considered to be additive, and this assumption is no longer appropriate. Declicking a crackly signal discards a lot of useful information that an optimal algorithm can use to improve the quality of the processed audio.

Declickle

CEDAR's Declickle algorithm was born out of research into the dual problems of declicking and decrackling difficult signals such as those including overblown brass instruments, solo strings, and human singing. In attempting to solve these difficulties, our researchers discovered that a single, new algorithm yielded better results than applying separate declick and decrackle processes. The performance of the resulting module is so good that, in almost all cases, it is not possible to hear that the signal was damaged prior to restoration.

Auto dehiss: Reducing broadband noise

Process Mode

Auto dehiss offers four combinations of processing modes:

■ Noise detection mode may be either Auto or Manual

In auto mode, the Duo processor automatically determines the noise level contained within the signal, whereas in manual mode the noise level is set by the user.

■ Channel mode may be either Left/Right or Mid & Side

Left/Right mode processes the left and right channels of a stereo signal in the conventional manner. Mid & Side mode surrounds the process in a mid/side encoder and decoder. This allows you to process the signal in Mid & Side format whilst retaining conventional Left/Right stereo at the input and output.

Switching between Modes

You may switch between modes using the front panel buttons as follows:

■ Auto <-> Manual

Toggle between AUTO and MANUAL modes by pressing and holding the AUTO-OFF / MANL-OFF button for three seconds or more. This operation is also available in the MENU/PROCESS MODE screen.

■ Left/Right <–> Mid & Side

Toggle between Left/Right and Mid & Side modes by pressing and holding the L-R / M-S button for three seconds or more. This operation is also available in the MENU/PROCESS MODE screen.

Process Mode screen

You may also select the modes using the Process Mode screen selectable from the Main menu.

■ Auto-Manl

Toggle this to select between Auto and Manual modes.

■ L/R – M/S

Toggle this to select between Left/Right and Mid & Side modes.

Auto Mode

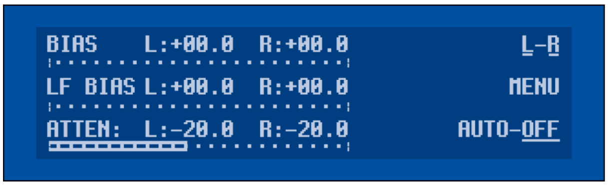

There are six options:

- BIAS
- LF BIAS
- ATTEN
- L-R / M-S (shown above in L-R mode)
- MENU
- AUTO-OFF

Bias (dB)

The Bias control allows you to tune the automatic algorithm for different applications and tastes; for example a broadcaster may want a very clean, dry result, but a for a CD remaster a more transparent sound (but with a little more residual noise) may be more appropriate.

A positive bias ensures that the noise is completely removed, but runs the risk of slight signal compression. A negative bias may retain more ambience in the signal, but at an increased risk of noise pumping.

This control operates across the whole signal spectrum and once set, the process will automatically apply a similar style of noise reduction to a wide variety of audio material.

The Bias control is a multiplying factor that affects the noise level calculated during the course of the signal. It is not the same as the Level control in Manual mode, which sets an absolute value for the noise content of the recording.

LF Bias (dB at 0Hz)

The LF Bias control allows fine control of the algorithm below 5kHz. It has no effect on higher frequencies.

Use positive values for material that exhibits prominent low- and mid-frequency noise. Examples of this may include 78rpm records and microgroove LPs. Use negative values when there is relatively little low- and mid-frequency noise, and the noise content is predominantly hissy.

Attenuation (dB)

This is the maximum amount of noise attenuation applied at any given frequency at any given moment. If the algorithm determines that the noise attenuation at any moment should be less than the ATTEN that you select, the lesser figure will be applied.

L-R / M-S

■ Select which channel or channels you are adjusting by repeatedly pressing the L-R / M-S button.

■ Toggle between Left/Right and Mid & Side modes by pressing and holding this button for three seconds or more.

AUTO-OFF

■ Switch processing ON or OFF by pressing the button associated with the AUTO-OFF option.

■ Toggle between AUTO and MANUAL modes by pressing and holding this button for three seconds or more.

When the Duo is in Bypass, this is indicated by **BYP** and the AUTO-OFF button is disabled. This also disables the press & hold mechanism for jumping between AUTO and MANL modes. You may press the BYPASS button in the Audio I/O page to cancel Bypass.

Manual Mode

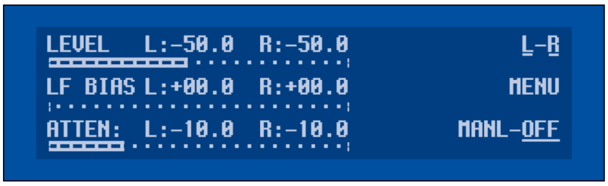

In Manual mode, the Bias control is replaced by a Level control, and the AUTO-OFF button becomes MANL-OFF. All other controls operate the same way as they do in Auto mode.

There are six options:

■ LEVEL

■ LF BIAS

■ ATTEN

■ L-R / M-S (shown above in L-R mode)

■ MENU

■ MANL-OFF

Level (dB)

The Level control allows you to determine the amount of noise contained within the signal. It operates across the entire signal spectrum. It is critical that the Level is set correctly, or undesirable side-effects will occur.

If the Level is slightly too low, auto dehiss will not remove all the noise and may generate short noise bursts that sound a little like noise pumping. If the Level is far too low, it will be difficult to obtain any noise reduction.

If the Level is too high, auto dehiss will treat some low-level signal as if it was noise and attenuate it. This may make the signal sound muffled.

LF Bias (dB at 0Hz)

The LF Bias control allows fine control of the algorithm below 5kHz. It has no effect on higher frequencies.

Use positive values for material that exhibits prominent low- and mid-frequency noise. Examples of this may include 78rpm records and microgroove LPs. Use negative values when there is relatively little low- and mid-frequency noise, and the noise content is predominantly hissy.

Attenuation (dB)

This is the maximum amount of noise attenuation applied at any given frequency at any given moment. If the algorithm determines that the noise attenuation at any moment should be less than the ATTEN that you select, the lesser figure will be applied.

L-R / M-S

▪ Select which channel or channels you are adjusting by repeatedly pressing the L-R / M-S button.

▪ Toggle between Left/Right and Mid & Side modes by pressing and holding this button for three seconds or more.

MANL-OFF

▪ Switch processing ON or OFF by pressing the button associated with the MANL-OFF option.

▪ Toggle between MANUAL and AUTO modes by pressing and holding this button for three seconds or more.

When the Duo is in Bypass, this is indicated by **BYP** and the MANL-OFF button is disabled. This also disables the press & hold mechanism for jumping between MANL and AUTO modes. You may press the BYPASS button in the Audio I/O page to cancel Bypass.

Auto dehiss tutorials

These Left/Right tutorials assume that the material exhibits virtually identical noise characteristics in each channel. If it does not, you should follow these instructions but adjust each channel independently to obtain satisfactory results.

Preparation

It is important that the audio presented to auto dehiss is free of clicks and crackle. This is because these degradations will interfere with the dehissing process and prevent you from reaching an optimal result. If necessary, you should pass the signal through the Duo declickle processor or other CEDAR declick and decrackle processes before applying the auto dehisser.

Manual Mode (L-R channel mode selected)

Set the auto dehiss processor in MANL (manual) mode, and ensure that the MANL-OFF control is in the MANL position. Now set the Level and LF Bias to zero, and the Attenuation to -40dB, and follow the next three tutorials in sequence.

Level Tutorial

Your first task will be to find the most appropriate setting for the Level. This will be the biggest influence on the quality of the processed signal.

Starting with the Level at zero and the Attenuation at -40dB, increase the value of the Level. At first, you will notice very little happening. Next, you enter a region in which there may be noise bursts and an artefact similar to noise pumping. Increase the Level further and these side-effects will begin to disappear and, at some point determined by the nature of the signal, the noise will decrease rapidly. Because the Attenuation is set to maximum, you will in all likelihood now find the signal to be somewhat muffled.

The optimal value of the Level is approximately the crossover point between the noise artefacts and the muffled sound.

There is a delay of a fraction of a second between adjusting the Level and hearing the effect at the output.

LF Bias Tutorial

At this point, you may wish to modify the processor's action to accommodate broadband noise with differing frequency characteristics. Your Duo allows you to do this by accentuating or suppressing the amount of noise identified at low+mid frequencies.

Switch the processing off, and listen to the source material. Does the noise appear to be present at all frequencies? If not, is it predominantly high-frequency (hiss) or low-frequency (rumble)? If the former, turn the processing on again, and reduce the LF Bias. This will reduce the Level below 5kHz, and reduce the amount of noise detected at these frequencies. As you do this, you should hear the low+mid frequency signal content 'open up'. This will allow you to increase the Level without risk of damaging the desired signal at low+mid frequencies.

If the noise appears to be predominantly in the lower frequencies (as is often the case with recordings obtained from discs) you should increase the LF Bias to tell the Duo that there is more noise below 5kHz. This will increase noise suppression at low+mid frequencies without increasing it at high frequencies, reducing the risk of high-frequency dullness when removing large amounts of low-frequency noise.

Adjust the LF Bias to obtain a good balance between high- and low- frequency noise attenuation.

There is a delay of a fraction of a second between adjusting the LF Bias and hearing the effect at the output.

Attenuation Tutorial

With the Level and LF Bias set to suitable values, you can now adjust the Attenuation to determine the amount of noise removed.

Decrease the amount of attenuation from -40.0 to 0.0, at which point you will hear that the processed signal is identical to the unprocessed signal. This is because the Attenuation control is limiting the amount of noise removal to 0dB - i.e. there is no effect.

Now increase the amount of Attenuation as defined by the material and your taste. You may find that, if the Level is too high, you can only increase the Attenuation by a few dBs before the signal starts to sound muffled. If the Level is too low, you can have more Attenuation, but with reduced effect.

Iterating the Procedure

It is unlikely that the values of the three controls are optimised, so you should now attempt to find a better value for the Level. Having done this you will probably wish to modify the LF Bias and Attenuation values. Fine-tuning of these controls will lead to excellent noise removal with few or no side-effects. However, the Duo is not a magic wand, and it may not be possible to restore some badly degraded material beyond a certain point. Experience will enable you to judge whether you have removed as much noise as possible without unacceptable consequences.

Auto Mode (L-R channel mode selected)

Set the auto dehiss processor in AUTO mode, and ensure that the AUTO-OFF control is in the AUTO position. Now set the Bias and LF Bias to zero, and the Attenuation to -40dB, and follow the next three tutorials in sequence.

Bias Tutorial

Since the auto dehisser is operating in Auto mode, it should in most cases be unnecessary to adjust the Bias. However, there are occasions where you may wish to bias the operation of the unit toward identifying greater or lesser amounts of noise in the signal (hence the name).

With the signal playing, leave the controls at the values selected above, and press the AUTO-OFF button a few times to compare the untreated and treated material. Now increase the Bias so that the auto dehisser errs on the side of identifying more noise in the signal. You will hear the amount of noise reduction increase, but possibly at the expense of making the wanted material sound a little muffled. Now decrease the Bias below zero, so that the dehisser identifies less noise, and compare. You will hear that the amount of noise reduction decreases, but with some audio you may find that the wanted signal benefits from increased ambience and 'air'. While the theoretical ideal Bias for any given signal is zero, you may find that you can achieve desirable results by adjusting the Bias in this way.

There is a delay of a fraction of a second between adjusting the Bias and hearing the effect at the output.

LF Bias Tutorial

At this point, you may wish to modify the processor's action to accommodate broadband noise with differing frequency characteristics. Your Duo allows you to do this by accentuating or suppressing the amount of noise identified at low+mid frequencies.

Switch the processing off, and listen to the source material. Does the noise appear to be present at all frequencies? If not, is it predominantly high-frequency (hiss) or low-frequency (rumble)? If the former, turn the processing on again, and reduce the LF Bias. This will reduce the Bias below 5kHz, and therefore reduce the amount of noise detected at these frequencies. As you do this, you should hear the low+mid frequency signal content 'open up'. This will allow you to increase the Bias without risk of damaging the desired signal at low+mid frequencies.

If the noise appears to be predominantly in the lower frequencies (as is often the case with recordings obtained from discs) you should increase the LF Bias to tell the Duo that there is more noise below 5kHz. This will increase noise suppression at low+mid frequencies without increasing it at high frequencies, reducing the risk of high-frequency dullness when removing large amount of low-frequency noise.

Adjust the LF Bias to obtain a good balance between high- and low- frequency noise attenuation.

There is a delay of a fraction of a second between adjusting the LF Bias and hearing the effect at the output.

Attenuation Tutorial

With the Bias and LF Bias set to suitable values, you can now adjust the Attenuation to determine the amount of noise removed.

Decrease the amount of attenuation from -40.0 to 0.0, at which point you will hear that the processed signal is identical to the unprocessed signal. This is because the Attenuation control is limiting the amount of noise removal to 0dB - i.e. there is no effect.

Now increase the amount of Attenuation as defined by the material and your taste. You may find that, if the Bias is too high, you can only increase the Attenuation by a few dBs before the signal starts to sound muffled. If the Bias is too low, you can have more Attenuation, but with reduced effect.

Iterating the Procedure

You may now wish to attempt to find a better value for the Bias. Having done this you will probably wish to modify the LF Bias and Attenuation values further. Fine-tuning of these controls will lead to excellent noise removal with few or no side-effects. However, the Duo is not a magic wand, and it may not be possible to restore some badly degraded material beyond a certain point. Experience will enable you to judge whether you have removed as much noise as possible without unacceptable consequences.

Noise Reduction in M-S Mode

Our research has shown that it is sometime beneficial to perform noise reduction in M-S mode. This means that, instead of adjusting the Level (or Bias), LF Bias and Attenuation values independently for the left and right channels, you may adjust them for the monophonic content within the signal (M) and for the content that is only present in the left or right channels (S).

Signal with a high Mid (M) noise component

If you perceive that the worst of the noise is coming from the centre of the stereo image, place the auto dehisser in M-S mode. Now toggle the M-S button so that you are adjusting the S component only, and reduce the Attenuation to zero. Toggle again so that you are adjusting the M component only. You are now performing noise reduction on the centre of the stereo image.

Follow the tutorials above, and listen to the effect that this has. You will find that you can apply a considerable amount of noise reduction on the M component without affecting the perceived width or ambience of the wanted signal.

Signal with a high Side (S) noise component

Sometimes, you will find that loud, wanted signals in the centre of the stereo image mask the noise, so that it is only noticeable at the sides of the image. Again, this can be handled effectively in M-S mode. Toggle the M-S button so that you are adjusting the M component only, and reduce the Attenuation to zero. Now toggle again so that you are adjusting the S component only. You are now performing noise reduction on the sides of the stereo image.

Follow the tutorials above, and listen to the effect that this has. You will find that you can apply a considerable amount of noise reduction without affecting the wanted signal at the centre of the soundstage.

A Short History of Noise Reduction

It would be best to describe first what we mean by the term 'broadband' noise, defining it to be an effect that adds (or subtracts) a random amplitude at all times to (or from) all frequencies within the audio spectrum. Thus, we do not include artefacts of limited duration such as clicks or crackles, both of which are removed by quite different methods to those described below.

Next, let's dispel any illusions regarding **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 noise from a signal that already contains it - they simply limit the amount you add when you commit that signal to tape and play it back again.

A **single-ended process** removes noise from your audio prior to committing it to tape, or at the very least, improves the signal to noise ratio without affecting the signal adversely. Which brings us neatly to the volume control... stunningly effective at removing noise, it does nothing to improve the S/N ratio, and has an all-too-noticeable side-effect. No noise, No signal.

Since broadband noise is most intrusive at high frequencies, the first stage in our evolutionary tale is the **Low-pass Filter**. Less damaging than the volume control, this 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 the low-pass filter will clean your 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 frequencies. 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 genuine signal.

Now, instead of altering the frequency response of the signal, how about changing the level in some way? Consider: if noise of relatively constant amplitude is always 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 which 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 remains the same: 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 unlike the gate, 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, the poor band separation filters (typically -6dB/oct or -12dB/oct) severely limit performance. The consequences of these problems 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 noise reduction 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.

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, subtracting a lot at (say) 8kHz, while leaving 8.1kHz virtually untouched.

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 the fingerprint is not accurate you will hear unpleasant side-effects. But let's assume that you have obtained a perfect fingerprint. You might then expect a good restoration, with few or no side-effects. Yet 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.

The background to CEDAR Duo auto-dehissing

Perfect restoration of hissy audio is theoretically impossible, but there are many techniques that CEDAR has developed to reduce the audible effect of hiss with as few side effects as possible.

Our earlier dehissers, such as the DH-1, DH-2, DHX, Dehiss-2 and Dehiss-3 (which were optimised for reasonably consistent broadband noises such as tape-hiss) were sometimes called ‘auto dehissers’ because they did not require a spectral fingerprint. These were able to achieve excellent results *if* the user dialled-in the correct values for the noise level, attenuation, and a third parameter called ‘brightness’ or ‘variance’. However, this method is not what we mean today when we use the term “auto dehissing”.

In our search for a truly automatic dehisser that you can set up and forget, we have devised an algorithm that - without human intervention - performs a remarkably good job of identifying the noise content of any signal presented to it. To use this, you need only set the maximum noise attenuation that you want, and you will be able to obtain excellent noise reduction without further adjustment. THIS is the essence of **auto dehiss**.

To make the Duo auto dehiss module as flexible as possible, we have retained a manual mode that allows you to fine-tune the noise reduction, and full remote control of all parameters so that you can automate its operation when desired.

Licence and Limited Warranty

1. DEFINITIONS

In this Licence and Limited Warranty the following words and phrases shall bear the following meanings:

'the Company' means CEDAR Audio Limited of 20 Home End, Fulbourn, Cambridge CB1 5BS, UK

'the System' means an instance of the sound-reprocessing system comprising hardware and software held on EPROM/FLASH memory ('firmware') developed by the Company and with respect to the software any updates or replacements thereof;

'this Document' means this Licence and Limited Warranty.

2. ISSUE AND USE OF THE SYSTEM

2.1 The terms and conditions of this Document are implicitly accepted by any person or body corporate who shall at any time use or have access to the System, and are effective from the date of supply of the System by CEDAR Audio Limited to its immediate customer.

2.2 The Company hereby grants to the Licensee and the Licensee agrees to accept a non-exclusive right to use the System.

3. PROPERTY AND CONFIDENTIALITY

3.1 The System contains confidential information of the Company and all copyright, trade marks, trade names, styles and logos and other intellectual property rights in the System including all documentation and manuals relating thereto are the exclusive property of the Company. The Licensee acknowledges that all such rights are the property of the Company and shall not question or dispute the ownership of any such rights nor use or adopt any trading name or style similar to that of the Company.

3.2 The Licensee shall not attempt to reverse engineer, modify, copy, merge or transcribe the whole or any part of the System or any information or documentation relating thereto.

3.3 The Licensee shall take all reasonable steps to protect the confidential information and intellectual property rights of the Company.

4. LIMITED WARRANTY AND POST-WARRANTY OBLIGATIONS

4.1 The Company warrants that the System will perform substantially in accordance with the appropriate section of its accompanying product manual for a period of one year from the date of supply to the Company's immediate customers.

4.2 The Company will make good at its own expenses by repair or replacement any defect or failure that develops in the System within one year of supply to the Company's immediate customer.

4.3 The Company shall have no liability to remedy any defect, failure, error or malfunction that arises as a result of any improper use, operation or neglect of the System, or any attempt to repair or modify the System by any person other than the Company or a person appointed with the Company's prior written consent.

4.4 In the case of any defect or failure in the System occurring more than twelve months after its supply to the Company's immediate customer the Company will at its option and for a reasonable fee make good such defect or failure by repair or replacement (at the option of the Company) subject to the faulty equipment having first been returned to the Company. The Company will use reasonable efforts to return repaired or replacement items promptly, all shipping, handling and insurance costs being for the account of the Licensee.

4.5 The above under takings 4.1 to 4.4 are accepted by the Licensee in lieu of any other legal remedy in respect of any defect or failure occurring during the said period and of any other obligations or warranties expressed or implied including but not limited to the implied warranties of saleability and fitness for a specific purpose.

4.6 The Licensee hereby acknowledges and accepts that nothing in this Document shall impose upon the Company any obligation to repair or replace any item after a time when it is no longer produced or offered for supply by the Company or which the Company certifies has been superseded by a later version or has become obsolete.

5. FORCE MAJEURE

The Company shall not be liable for any breach of its obligations here under resulting from causes beyond its reasonable control including, but not limited to, fires, strikes (of its own or other employees), insurrection or riots, embargoes, container shortages, wrecks or delays in transportation, inability to obtain supplies and raw materials, or requirements or regulations of any civil or military authority.

6. WAIVER

The waiver by either party of a breach of the provisions hereof by the other shall not be construed as a waiver of any succeeding breach of the same or other provisions, nor shall any delay or omission on the part of either party to exercise any right that it may have under this Licence operate as a waiver of any breach or default by the other party.

7. NOTICES

Any notices or instruction to be given hereunder shall be delivered or sent by first-class post or telecopier to the other party, and shall be deemed to have been served (if delivered) at the time of delivery or (if sent by post) upon the expiration of seven days after posting or (if sent by telecopier) upon the expiration of twelve hours after transmission.

8. ASSIGNMENT AND SUB-LICENSING

The Licensee may at his discretion assign the System and in doing so shall assign this Licence its rights and obligations to the purchaser who shall without reservation agree to be bound by this Licence. The original Licensee and any subsequent Licensees shall be bound by the obligations of this Licence in perpetuity.

9. LIMITATION OF LIABILITY

The Company's maximum liability under any claim including any claim in respect of infringement of the intellectual property rights of any third party shall be, at the option of the Company either:

(a) return of a sum calculated as the price received for the System by the Company from its immediate customer depreciated on a straight line basis over a one year write-off period; or

(b) repair or replacement of those components of the System that do not meet the warranties contained within this Document.

The foregoing states the entire liability of the Company to the Licensee.

10. CONSEQUENTIAL LOSS

Even if the Company has been advised of the possibility of such damages, and notwithstanding anything else contained herein the Company shall under no event be liable to the Licensee or to any other persons for loss of profits or contracts or damage (whether direct or consequential) arising in connection with the System or any modification, variation or enhancement thereof and including any documentation or data provided by the Company or for any other indirect or consequential loss.

11. ENTIRE AGREEMENT

The Company shall not be liable to the Licensee for any loss arising in connection with any representations, agreements, statements or undertakings made prior to the date of supply of the System to the Licensee.

12. TERMINATION

This Licence may be terminated forthwith by the Company if the Licensee commits any material breach of any terms of this Licence. Forthwith upon such termination the Company shall have immediate right of access to the System for the purpose of removing it.

13. SEVERABILITY

Notwithstanding that the whole or any part of any provision of this Document may prove to be illegal or unenforceable the other provisions of this Document and the remainder of the provision in question shall remain in full force and effect.

14. HEADINGS

The headings to the Clauses are for ease of reference only and shall not affect the interpretation or construction of this Document.

15. LAW

This Document shall be governed by and construed in accordance with English law and all disputes between the parties which cannot be resolved by negotiation shall be determined by arbitration in England in accordance with the Arbitration Act 1950 and 1979.

Serial Number:	
Inspected by:	
Date:	

Manufactured by:

CEDAR Audio Ltd

20 Home End

Fulbourn

Cambridge

CB1 5BS

UK

Web: www.cedaraudio.com

Email: info@cedaraudio.com

Tel: +44 1223 881771

interstage
Phistersvej 31, 2900 Hellerup, Danmark
Telefon 3946 0000, fax 3946 0040
www.interstage.dk
- pro audio with a smile