CEDAR for Pyramix 64

CEDAR Audio Limited

www.cedaraudio.com

CEDAR for Pyramix 64 manual v1.0 – 15 May 2015

This page left blank.

Table Of Contents

This page left blank.

Getting Started

CEDAR for Pyramix 64 incorporates award-winning processes from the flagship CEDAR Cambridge™ system. These can process your audio in many desirable ways and are capable of correcting all manner of problems without damaging the desired audio or introducing unwanted side-effects and artefacts.

Unpacking

This package should contain the following:

- A DVD or USB memory stick containing the installer for CEDAR for Pyramix 64
- A CEDAR Dongle
- An identification fob to attach to the dongle
- This manual

Assumed Knowledge

This manual assumes that you are fully conversant with your Pyramix system, and that you know how to operate the host software and operating system onto which you're loading CEDAR for Pyramix 64. It will refer to operations that are common to these products, but will not attempt to explain them.

Troubleshooting Non-CEDAR Components

If you encounter problems with your PC, Microsoft Windows® or Pyramix, please refer to the relevant manuals or contact the dealer that supplied these to you. Unless appointed independently as authorised dealers for the following products, CEDAR Audio's dealers will not attempt to provide technical support for:

- PCs of any description
- Microsoft Windows
- 3rd party host hardware
- 3rd party host software

Installing CEDAR for Pyramix 64

If you are reading this, you will have already inserted the DVD or USB memory stick as appropriate. First, remove any previous versions of CEDAR for Pyramix or Retouch for Pyramix from the system. Now double-click on the installer package: **CEDARforPyramix64.msi**

If you already have CEDAR for Pyramix 64 installed upon the PC when you run **CEDARforPyramix64.msi**, the installer will ask whether you wish to repair or remove it.

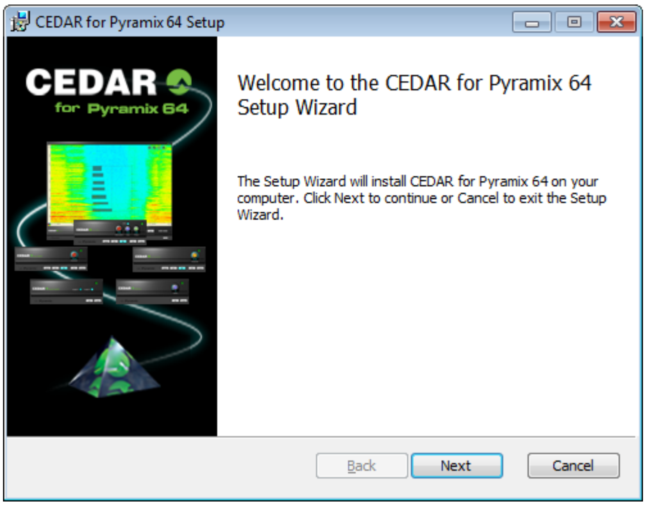

- Ensure that you have sufficient space on your drive and then press Install to install the complete package or press the Customise button to select which format and/or selection of modules are loaded.
- A message will appear to tell you that installation was completed successfully. Click on Close
- Insert the dongle into any available USB socket. If you do not do so, CEDAR for Pyramix 64 will either run an Evaluation Version that does not permit auditioning or saving of results, or provide you with information about the requested function, depending upon the module selected.

If the Wizard does not run automatically (the function may be switched off in your Windows settings) you may browse the DVD or memory stick and run the installer in the usual way.

Setting up Pyramix

To run the CEDAR processes optimally, you should check and if necessary make the following changes to your Pyramix setup:

- Edit>Editing Modes>Insert Mode>Overwrite should be ticked (selected)
- Edit>Auto-Ripple should be unticked (off)

Licence Manager

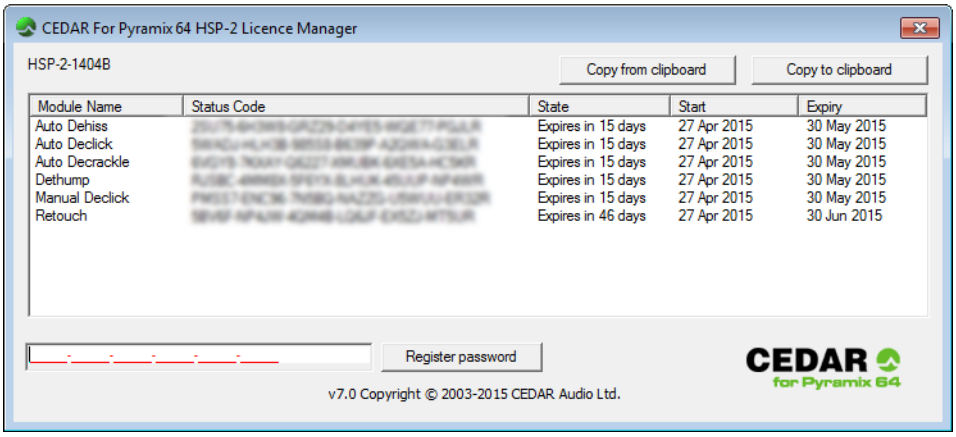

To install a licence that connects new software to the existing dongle, type the password supplied by CEDAR Audio into the **Register password** window. (You may also **Copy from clipboard** if appropriate.) When accepted, the password will change from red to black. Press the **Register password** button to complete the registration. The new software is now ready for use.

There may be occasions when CEDAR Audio requires your Status Codes. If so, use the **Copy to clipboard** button to copy and paste this information into an email or other suitable document.

Licence states

Click on the 'key' icon within the plug-ins' user interfaces to access the Licence Manager.

This icon has three states:

Image in GUI	Status
	The software is fully licensed.
	A time-out demo licence has been issued.
	There are six reasons why you might see the red icon. Hovering over the icon will display a tool tip that offers further information as follows: 1. No licence has been issued. 2. A timed licence has expired. 3. A timed licence has been disabled because of a problem. This should not affect full licences. *If you encounter this message, check that the system clock is set to the correct date and time. If it is not, correct it. If the clock is correct, it's possible that the dongle has failed and should be returned to CEDAR Audio if you wish to use timed (demo or rental) licences.* 4. DSP not loaded. You should relaunch the host system with the CEDAR dongle attached. If this does not cure the problem, you should reinstall CEDAR for Pyramix 64. 5. The HASP dongle has become physically detached after loading the host software and CEDAR for Pyramix 64. 6. Some software hosts do not update the status until the process is run in each session. Once run, the amber or green icon will appear, as appropriate.

General controls

CEDAR for Pyramix 64 processes use a selection from the common controls, as follows:

	**Play** Audition the audio processed using the current parameter values. This audio is not laid back into the track until you press OK.
	**Bypass** Bypass processing while auditioning the audio.
	**Stop** Stop auditioning.
	**OK** Lay the audio, processed using the current parameter values, back into the track. This will create a new clip that exactly matches the duration of the original audio.
	**Cancel** Close the process without applying any processing.

Auto processes

The standard method for processing using the 'auto' modules is as follows:

- Select the audio you wish to process. This can span up to eight tracks if wished.
- Select the desired CEDAR for Pyramix 64 process from the Render window.
- Press play to audition the audio.
- While listening to the looped audio, select suitable parameter values.
- If happy, press OK to lay the processed audio back into the track(s).
- If not happy, adjust the parameter values further or press Cancel to close the process without changing the audio.

Manual processes

The standard method for processing using Manual declick and Dethump is as follows:

- Select the audio you wish to process. This can span up to eight tracks if wished.
- Select the desired CEDAR for Pyramix 64 process from the Render window. Ensure that both Extra Handles values are set to at least 2 seconds.
- Choose appropriate parameter values.
- Press OK to lay the processed audio back into the track(s).
- Audition the audio. If you are not happy with the result, undo in the standard fashion and process again with different parameter values.

Retouch

Retouch provides a powerful spectral editor with seven types of processing within the Pyramix host.

- Select the audio you wish to process. This can span up to eight tracks if wished.
- Edit the audio within the Retouch environment.
- At any time, you may press play to audition the audio in its current state.
- Press stop to stop auditioning.
- When you have completed your editing within Retouch, press OK to lay the processed audio back into the track(s).
- Press cancel to discard all edits within Retouch and return to the Pyramix tracks window.

Removing clicks

Auto Declick

Auto declick removes clicks from a wide range of material without introducing unwanted distortion or artefacts.

For each piece of material that you process, you should find a suitable value for the threshold. You can often leave this at its default value.

Controls

	**Threshold** The threshold controls the sensitivity of the process. With the threshold set high, Auto declick will remove only the largest clicks and scratches. A lower threshold will also remove smaller ticks and clicks. The threshold ranges from 3 to 100 on an arbitrary scale. You will need to reduce the value when removing fine crackle and buzzes. 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 find the highest value at which the unwanted noise is removed.

Manual Declick

Manual declick should be used to remove extended clicks. It allows you to specify the audio that constitutes the click and restores the audio using an interpolator that is optimised for longer clicks and scratches. The maximum click length handled is approximately 0.1s irrespective of sample rate.

Load the audio that contains the click into the file processor, select the passage that includes it, and zoom in to identify the click itself. If you have correctly identified the problem, you will see an easily recognisable click superimposed upon the desired audio waveform.

Mark the click in the usual fashion. Be aware that the tail of the click may extend a little beyond the obvious region. Now invoke Manual Declick, select the appropriate model, and then press OK to remove the click. You may repeat the operation to remove as many clicks as desired.

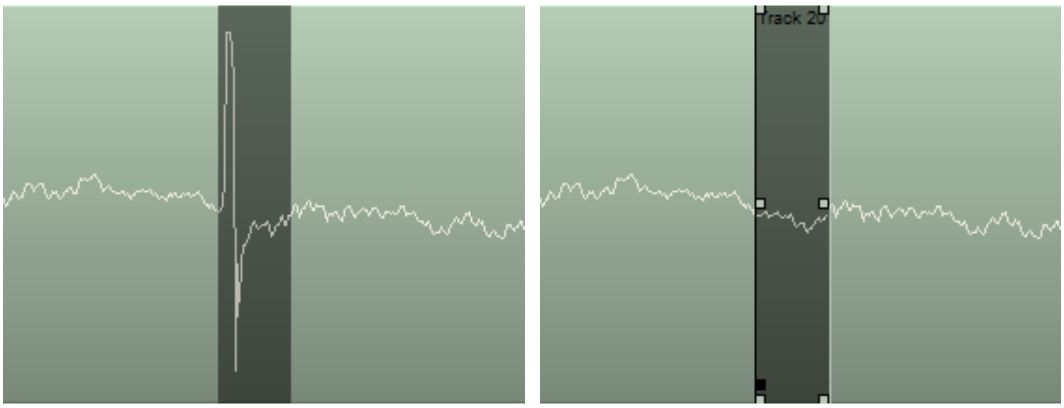

Note: There are maximum and minimum lengths to which you can apply manual declick. If the marked region exceeds the maximum length that can be addressed, an error message will appear. If you cannot mark the entire click in Manual Declick you should consider using Retouch and/or Dethump to restore the audio.

Controls

	**Model A** The general purpose manual declicker; the optimal choice for most situations.
	**Model B** Optimised for long clicks in the presence of a high noise floor.

This page left blank.

Removing crackle

Auto decrackle

Auto decrackle removes surface noise, crackles, and some types of buzz. You may also use it to reduce the audible effects of amplitude distortions caused by problems such as microphone diaphragm grounding and over-driven mixer inputs. It is an almost foolproof process. However, it will only function correctly if there are no major ticks or clicks received at the input. If obvious clicks exist within your material they must be removed before processing the audio using Auto decrackle.

Controls

	**Threshold** This determines the amount of crackle removed. A high threshold tells the system to remove only the most obvious crackles, while a lower threshold also removes fine crackle, buzz, and distortion. The threshold ranges from 0 to 40 on an arbitrary scale. Find the threshold that removes as much of the crackle/distortion as possible without introducing unwanted side-effects.

Removing thumps

Dethump

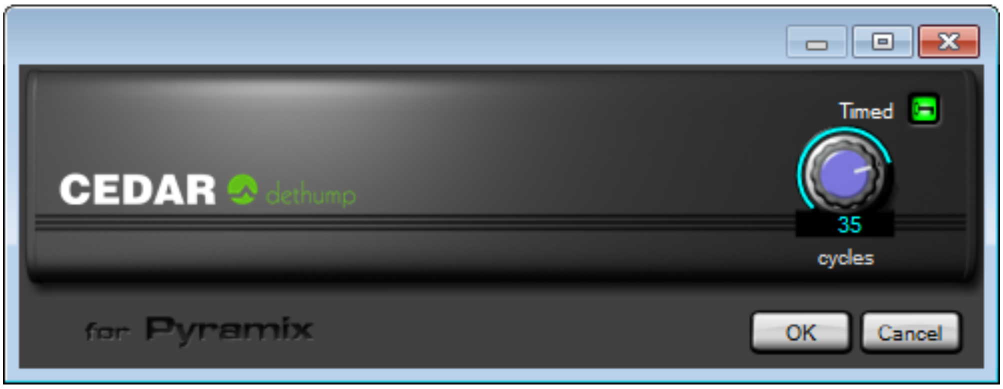

Dethump should be used to remove low frequency disturbances. The maximum thump length handled is approximately 1s irrespective of sample rate.

Select the passage of audio that includes the thump and zoom in to identify the thump clearly. If you have correctly identified the problem, you will see an easily recognisable thump superimposed upon the desired audio waveform.

Tell the process where the thump lies by dragging the cursor over the area of damaged audio. Fine-tune the selection by dragging either edge using the mouse. Alternatively (if you are working with no more than two channels) you may place Mark In as close as possible to the start of the thump, and Mark Out as closely as possible to the end of the thump.

Invoke Dethump. In this example, there are two, clearly visible cycles. Set the Cycles knob to this value and then press OK. Dethump will generate a new clip that replaces the selected audio.

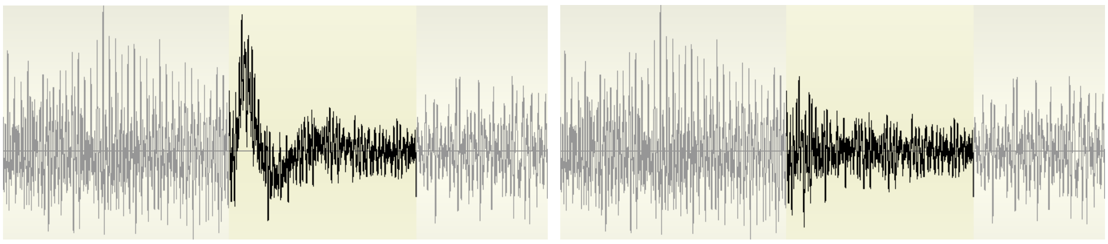

Note: There are maximum and minimum lengths to which you can apply Dethump. If the marked region is outside this range the system will present an error message.

Controls

	**Cycles** This tells the system how many cycles exist within the thump, and fine-tunes the process accordingly. For any given length of marked thump, a higher value suggests that the thump comprises a wider range of frequencies. Find the value that correctly removes the thump without introducing unwanted side-effects.

Reducing hiss

Auto dehiss

Auto Dehiss embodies an advanced algorithm that offers a unique Auto mode that enables the software to determine the noise content of the recording and remove this with minimal loss of high frequencies and without the introduction of unwanted artefacts. It offers two modes:

Auto

The process automatically determines the noise level contained within the signal.

Manual

The noise level is set by the user.

Auto Mode

There are three controls: Attenuation, Bias and LF bias.

Bias

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 for a CD remaster a more transparent sound (but with a little more residual noise) may be more appropriate. A positive Bias ensures that more noise is removed, but runs the risk of slight signal compression. A negative Bias may retain more ambience in the signal, but at a slightly increased risk of 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.

> Note: Bias is a multiplying factor that affects the noise level calculated during the course of the signal. It is not the same as the Threshold 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 attenuation that you select, the lesser figure will be applied.

Manual Mode

In Manual mode, the Bias control becomes inactive, and the Threshold control becomes active. The three controls are now:

Threshold

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

If the Threshold 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 Threshold is far too low, it will be difficult to obtain any noise reduction. If the Threshold 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 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 Attenuation that you select, the lesser figure will be applied.

Tutorials

It is important that the audio presented to Auto Dehiss is free from 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 Declick and Decrackle modules before applying the auto dehisser.

Auto Mode

Ensure that the process is active. Now set the Attenuation to -40dB and the Bias and LF Bias to zero, and follow the instructions in sequence.

Bias

Since Auto Dehiss 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 on/off button a few times to compare the untreated and treated material. Now increase the Bias so that Auto Dehiss 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.

LF Bias

At this point, you may wish to modify the processor's action to accommodate broadband noise with differing frequency characteristics. You 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 Auto Dehiss 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.

Attenuation

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, Auto Dehiss 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.

Manual mode

Threshold

Your first task will be to find the most appropriate setting for the Threshold. This will be the biggest influence on the quality of the processed signal. Starting with the Threshold at zero and the Attenuation at -40dB, increase the value of the Threshold.

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 Threshold 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 Threshold is approximately the crossover point between the noise artefacts and the muffled sound.

LF Bias

At this point, you may wish to modify the processor's action to accommodate broadband noise with differing frequency characteristics. Auto dehiss 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 Threshold 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 Threshold 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 process 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.

Attenuation

With the Threshold 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 Threshold is too high, you can only increase the Attenuation by a few dBs before the signal starts to sound muffled. If the Threshold 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 Threshold. 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, Auto Dehiss 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 an MS environment

Our research has shown that, with stereo material, it is sometime beneficial to perform noise reduction in MS mode. This means that, instead of adjusting the Threshold (or Bias), LF Bias and Attenuation values independently for the left and right channels, you 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).

Please refer to your host system documentation to determine whether you are able to set up an MS environment within it.

Emphasis and De-emphasis

On rare occasions, you can create an EQ curve and its inverse to set up emphasis and de-emphasis curves that can benefit the dehissing process.

Controls

<table>
<tr><td>
</td><td>

Mode

Choose whether to process in Manual or Auto mode.
</td></tr>
<tr><td>
</td><td>

Threshold (manual mode only)

This tells Auto Dehiss how much noise is present in the signal. It is important that the threshold is set correctly, or undesirable side-effects may occur.

Threshold too low:

The system will not remove all the noise and may generate an artefact from residual noise let through by the process.

Threshold too high:

Auto Dehiss may treat some low-level signal as if it was noise and attenuate it. This results in a muffled signal.
</td></tr>
<tr><td>
</td><td>

Bias (auto mode only)

Biases the operation of Auto Dehiss toward identifying greater or lesser amounts of noise in the signal. The theoretical ideal for any given signal is zero.
</td></tr>
<tr><td>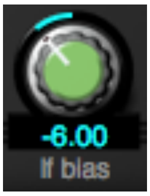
</td><td>

LF Bias

Allows you to accommodate broadband noise with differing frequency characteristics by accentuating or suppressing the amount of noise identified at Low+Mid frequencies.

Reduce the LF Bias value to reduce the amount of noise detected below 5kHz.

Increase the LF Bias value to increase the amount of noise detected below 5kHz.
</td></tr>
<tr><td>
</td><td>

Attenuation

This is the maximum amount of noise attenuation (in dB) 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 Attenuation that you select, the lesser figure will be applied.
</td></tr>
</table>

Spectral editing

Introduction to Retouch

In photography, retouching is the art of identifying and removing unwanted blemishes from a picture to produce a perfect image. Retouch brings this concept into the audio world, allowing you to remove unwanted acoustic events from a signal.

A tool for removing unwanted noises

Until CEDAR Audio invented the spectrographic editing system, audio restoration systems were limited as to the types of noises that they could remove: clicks and scratches, crackle and buzz, hiss, pops and thumps, and so on. But the release of Retouch offered a huge leap forward in processing technology, allowing users to identify and eliminate sounds as varied as coughs, record scuffs, squeaky chairs, page turns, the creak of a piano pedal, and even a car horn.

Before Retouch, some engineers attempted to use techniques such as severe EQ to remove these types of noise. Others used (and misused) harsh compression, tight editing, and signal interpolators that affect the whole frequency spectrum. All of these methods damage good signal that should be left untouched, and can introduce unwanted side-effects such as ringing and drop-outs.

In contrast, Retouch offers tools that localise unwanted sounds accurately in both time and frequency. Once identified, those sounds are replaced with audio derived from the surrounding, good signal. All other audio remains untouched.

A tool for editing sounds in both time and frequency space

Because of its ability to identify audio both in terms of time and frequency, Retouch is able to move precisely defined audio from one location to another, either overlaying the new audio over the old, or mixing the existing audio with the sound you are moving. There are many uses for this, including noise elimination, editing, correcting the pitch of (or eliminating) incorrectly played notes, and the creation of sound effects.

A tool for revealing wanted sounds and suppressing unwanted background

Retouch allows you to reveal individual sounds or utterances within a file, either by amplifying the wanted sounds or by suppressing the rest of the audio, or both simultaneously. You can also use this to retain only the sounds or words wanted within an audio file.

Getting started with Retouch

Retouch provides seven audio processing modes. These are:

The use of these is broadly similar but, for any given job, you should choose the appropriate mode to achieve optimal results.

This manual will guide you through the use of the first of these - the Interpolate mode - before proceeding to describe the use of each of the others.

Invoking Retouch

Select the audio that contains the problem that you wish to correct, plus a significant duration either side of the noise itself. As a guideline, you should mark at least three times as much signal as the extent of the unwanted sound itself, with the sound centred in the region. If you select no audio, the whole track will be loaded.

When you have done this, load Retouch in the usual fashion. It will handle up to eight tracks simultaneously.

The Transport Controls

Use the transport controls to play the audio. You may use the Spacebar in the common fashion to start/stop playback. (See Hot Keys.)

Undo and Redo

Step forward and backward through your actions in the usual fashion.

Reveal menus

As well as right-clicking in the axes and on the spectrogram itself, you can click on the three-dot icons in these locations to reveal the contextual menus.

Resizing

You can resize Retouch in the usual fashion.

Hot Keys

You may create a personalised set of Hot Keys to help you navigate and process more quickly and efficiently.

Invoke the Hot Key editor by pressing on the hot key icon.

To change a Hot Key definition, simply highlight the wanted action and then press the desired key or key combination. There is no need to save the definition file, which is contained within Retouch.

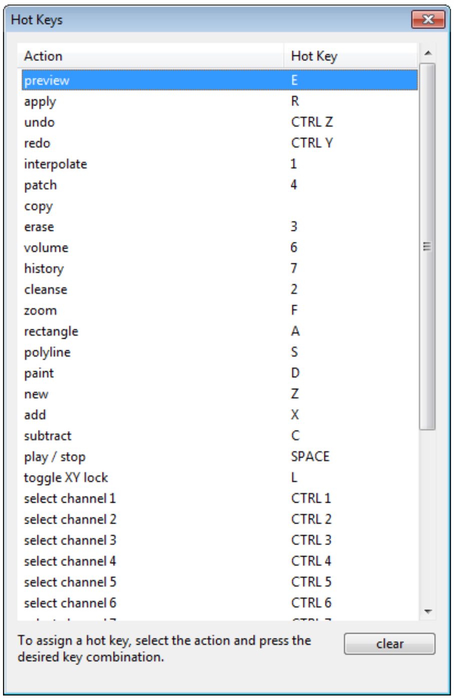

The Retouch Spectrogram

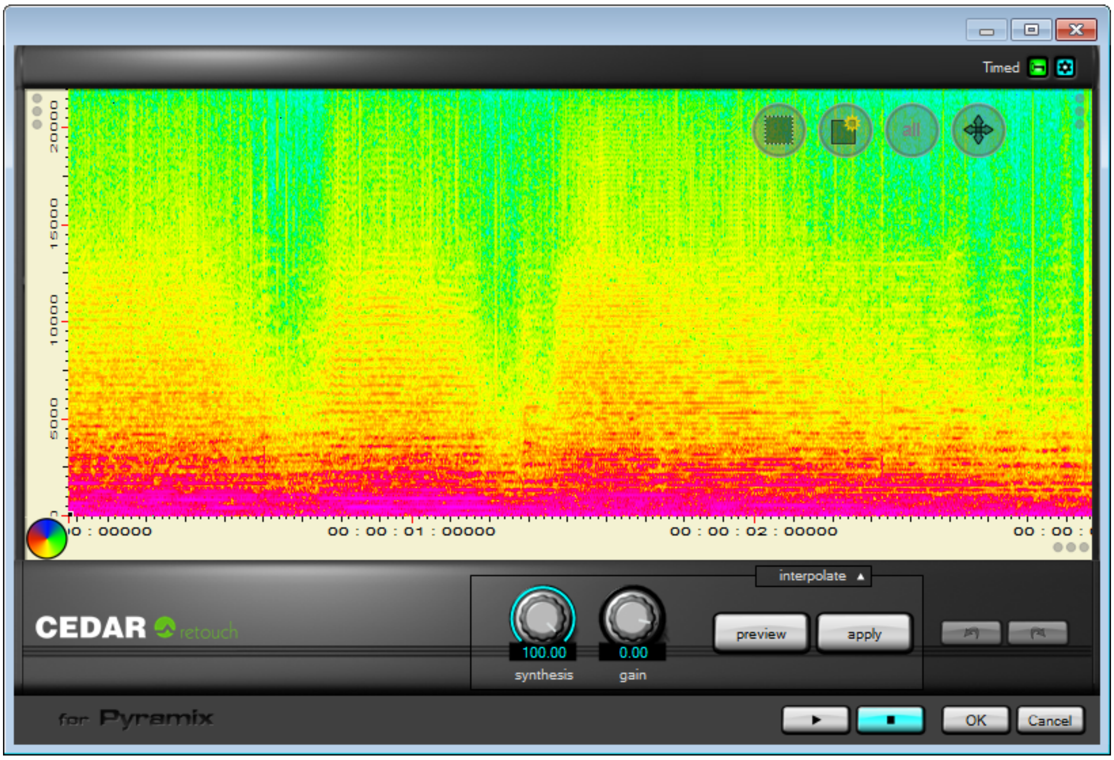

The spectrogram represents audio in three dimensions:

- The horizontal ('X') axis is time.
- The vertical axis ('Y') is frequency.
- The colours represent the Z-axis, and are amplitude in dB.

The spectrogram is, therefore, a representation of the amplitudes of all frequencies at all times within the region selected. (This is not a complete representation of the signal because no phase information is shown.)

Controlling the spectrogram

When you first invoke Retouch, the signal amplitude will be displayed using the standard colour-set. This choice is suitable for displaying signals with a wide dynamic range, but may not be ideal when signal amplitudes occupy a narrower range. To help you differentiate signals of similar amplitude, Retouch provides various means for changing the colour-set.

The Colour Wheel

Click and hold on the colour wheel. The colour map will be displayed on the Y axis and you can then drag your mouse to rotate the colour map. You may also click on the wheel and then drag the upper and lower extremes of the map upward or downward as desired, or slide the whole of the active area up and down the axis by clicking and dragging upon it.

As you change the colour map, you will notice that the detail on the spectrogram can change considerably. For some signals, events (i.e. areas of distinct colour changes) can appear and disappear depending on the position of the colour map. This is a valuable tool that allows you to identify the positions and frequencies of sounds that might otherwise be invisible.

Colour clipping

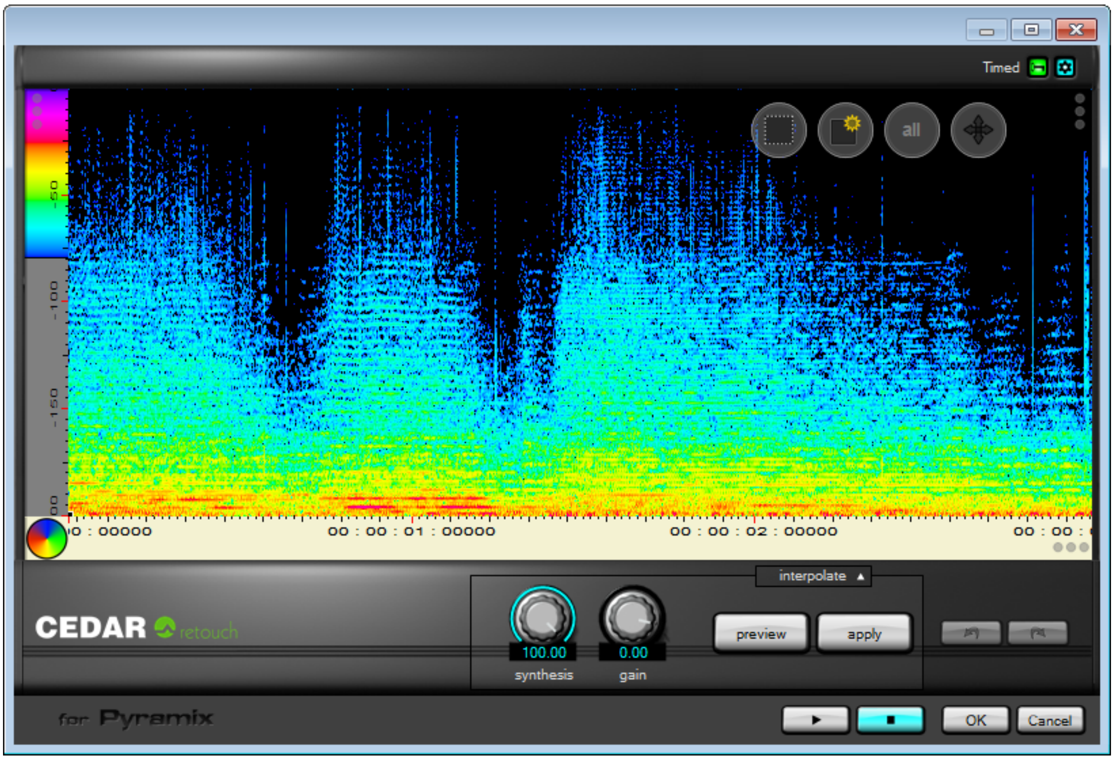

To maximise the resolution on the colour axis, it is useful to fit the entire range of the colour wheel within the highest and lowest amplitudes in the Retouch window. Furthermore, it is often useful to represent low amplitudes as black areas on the spectrogram, and high amplitudes in white. You may make these adjustments (which can be thought of as zooming on the Z axis) using the Amplitude Bounds controls and menu.

To access these, right-click on the colour wheel. The Y axis will now show the colour map. You may adjust the upper and lower bounds by left-clicking and dragging its upper and lower extremes.

Right-click within the map or click on the menu selector to reveal the colour map menu.

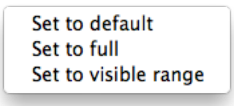

Set to default

A default setting – found by CEDAR to be useful in many instances – is loaded.

Set to full

The colour map is distributed between the maximum and minimum possible amplitudes.

Set to visible range

Sets the colour map such that the colours are distributed between the greatest and lowest amplitudes contained within the visible part of the signal.

> None of these operations affect the underlying audio data. They simply alter the visual representation to accentuate details within it.

Having selected the desired colour map, click again on the colour wheel to return the Y axis to a frequency display.

Zooming and scales

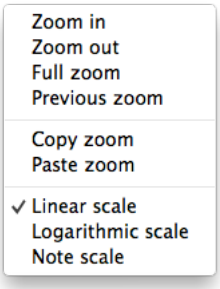

Zooming

Zooming is very important when identifying sounds. You can zoom using the mouse wheel, the trackpad (zooming is centred on the "+" cursor), the right-click commands, and by right-clicking and dragging within the axes to select the desired area of the audio. You may also copy a zoom level and paste this into another instance of Retouch.

Scales

You may display the frequency (Y) axis in linear or logarithmic modes, or as MIDI notes based upon A440. This parameter changes the visual representation of the audio, but does not affect it in any other way.

The play region

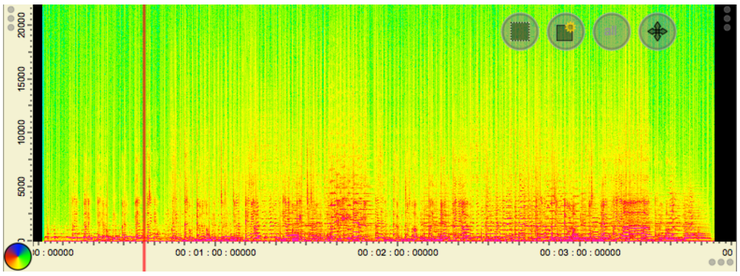

When you play audio within Retouch, it will loop around the region of audio selected in the time (X) axis at the moment that you press Play. You may change the zoom after this, and the same piece of audio will continue to play until stopped.

Harmonic markers

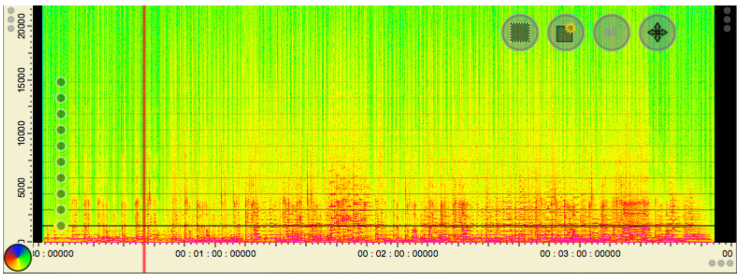

Invoked within the spectrogram menu, harmonic markers allow you to view the first ten harmonics of a given fundamental frequency. Drag any of the markers to adjust the fundamental and, therefore, the other nine harmonics. Markers can be very useful when, for example, identifying the harmonics of a tone that you wish to remove or shift in frequency to correct a wrongly played note.

Examples of sounds that can be Retouched

Recognising a Sound

You now need to identify the unwanted sound within the spectrogram. This can take the form of tonal events (for which you may need to remove individual harmonics or partials from the spectrum), short transient events such as clicks, reverberant events, and noisy events.

The following graphics show three common (and easily recognised) types of noise that can be addressed using Retouch:

Clicks:

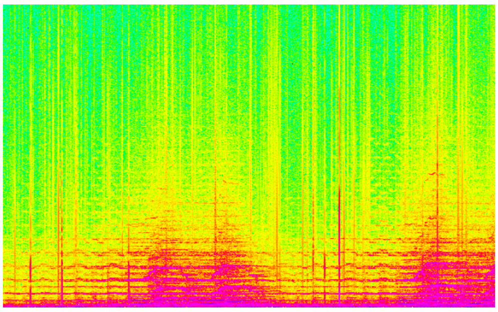

This dense pattern of thin, vertical bars is a series of loud clicks. Retouch may be useful for eliminating those that are resistant to standard methods of declicking.

Noise burst (or 'chuff'):

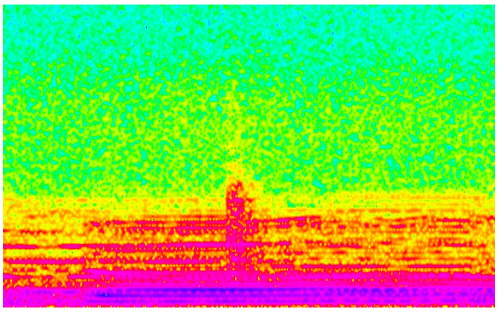

This vertical band is far too wide to be a well-defined click, but is a burst of broadband noise (sometimes called a 'chuff').

Tonal noise:

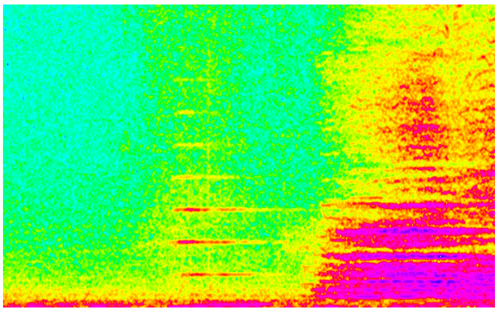

The harmonically related horizontal bars in the centre of this spectrogram demonstrate that there a tone – which may be wanted or unwanted – present. In this case, it is a car horn contaminating a quiet passage of choral music.

Defining sounds to be processed

When you have located a sound of interest, you will often find that it exhibits more detail than you may have thought, and you may then wish to zoom more closely to identify it more accurately. When you have identified it, you may define the region to be processed.

Defining a region

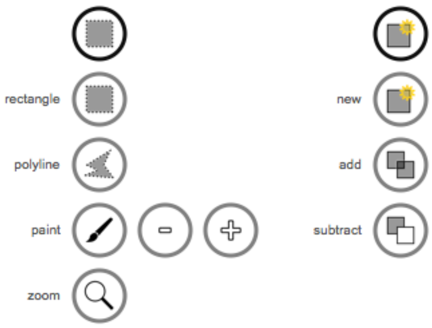

First, ensure that the Interpolate Tool is selected, then decide how you wish to define the region to be Retouched. There are two related menus that assist you; the selection tool menu and the selection mode menu.

Simple rectangular regions

You can mark many unwanted sounds using a simple rectangle. Click on the Rectangle and New Selection icons, then click and drag over the noise. Retouch will mark a rectangular region by greying out the area you describe. This is the area of audio that Retouch will process.

Complex regions

You may define complex shapes using three additional tools:

Use Polyline to click around a complex shape. You can left-double-click to close the loop, or right-click to cancel an incomplete shape.

Use the Paintbrush to paint a shape. Adjust the width of the brush using the + and – keys, by using <SHIFT> together with the mouse wheel, or by assigning Hot Keys to do so.

Click on this icon to create complex shapes by adding to an existing defined region. You may add rectangles using the Rectangle tool or add complex shapes defined using the Polyline and Paintbrush tools. Press and hold <SHIFT> while drawing on screen to add to an existing shape. This overrides the icon selection.

> It is advisable to work in Add mode when defining complex regions to avoid the risk of accidentally starting a new region and deleting the existing one.

Click on this icon to create complex shapes by subtracting from an existing defined region. You may subtract rectangles using the Rectangle tool or subtract complex shapes defined using the Polyline tool.

> Press and hold <CTRL-SHIFT> while drawing on screen to subtract from an existing shape. This overrides the icon selection.

Zooming (again)

Associated with the region tools, you'll find a further means to zoom within the spectrogram. When you click and drag across the audio, the upper and lower frequencies (or notes) and the start and end times of the selection will be displayed until you release the mouse button.

Resizing a region

Having defined a region, you may adjust its duration and the frequency range it covers by dragging its grab handles.

Multiple separate regions

It is possible to mark multiple, separate regions in the same way as adding to, or subtracting from, existing regions. This technique works well for multiple objects separated in frequency (along the Y axis) but is not appropriate for interpolating multiple objects - such as a succession of clicks - separated in time (along the X axis).

Defining the Wings

Depending upon the processing mode selected it's possible that, when you have defined the region to Retouch and released the mouse button, it will be greyed out, and 'wings' will appear on either side of it. These show the audio data that will be used to calculate the audio within the region. The example below shows a simple marquee around a click with the default wings on either side.

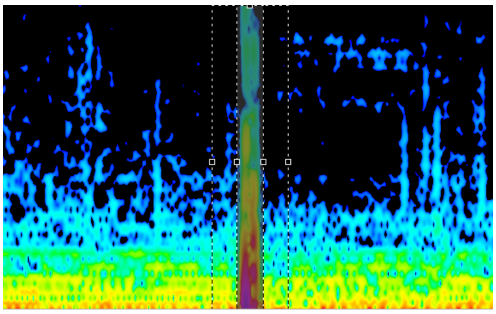

It is important that the wings contain audio that is representative of the audio that you wish to have fill the selected area. Therefore, they should not include other unwanted noises, or inappropriate – for the purposes of reconstruction – musical events such as the transients of following notes. In order to ensure that the wings contain suitable data, you may adjust their durations using the grab handles.

When first displayed, each wing occupies a theoretically optimal area. If you choose to adjust either or both of them, bear in mind that:

- if the wings are too small, there may be insufficient data to rebuild the signal correctly, and you may introduce an artefact.
- if the wings are too large, you may include inappropriate data that is not suitable for rebuilding the signal, and this too may introduce artefacts.

There is no reason why the wings need to be equal in duration. Defining them unequally can be useful when, for example, two unwanted events are close together.

Channel Activation

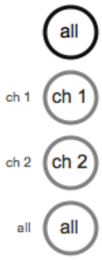

Click on a channel button to activate a channel for display and processing. Click on the ‘all’ button to activate all channels.

If a single channel is selected in the channel selector, the data contained within that channel alone is displayed in the spectrogram window. If multiple channels are selected, the spectrogram shows the average value of the data in the selected channels.

If you are processing a multichannel clip, the process region(s) and wings will be defined identically in all channels, and you should flip between channels to ensure that they are appropriate to all. You should perform multiple mono processes if you are unable to define an area that is appropriate for all channels.

Interpolating a Region

When you are happy with the selected region and the wings, you can choose appropriate parameters for Synthesis and Gain.

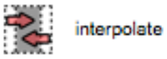

Firstly, ensure that the Interpolation mode is selected in the pop-up menu above the Preview and Apply buttons.

The controls

Synthesis

This determines how close the sound of the replaced signal is to the sound of the original, unwanted signal.

At Synthesis = 0 (minimum) the marked region is replaced with the original signal: i.e. there is no change.

At Synthesis = 100 (maximum) the marked region is replaced with the calculated samples determined from the data contained in the wings, and the original signal has no influence on the result.

At Synthesis settings between 0 and 100, the original data are included in the calculation to a greater or lesser extent. You may ask why you would want to use any of the information contained in the unwanted noise. The answer is that it can help to create synthesised audio that best matches the true audio either side of the unwanted noise. This can be particularly useful at low frequencies.

Gain

This applies gain to the replacement audio in the range -150dB to +20dB. However, the edges of the replacement audio remain correctly matched to the original signal either side of the process area, so that there are no clicks or pops introduced.

There are a number of reasons why you may wish to apply gain to the replacement audio. For example, the data in the wings may be louder than the required data in the process region, and a small attenuation may enable you to match the calculated audio to the desired result. You can also use reduction with the Synthesis set to zero to attenuate or accentuate small areas of signal such as individual harmonics or breath noise.

The default settings [synthesis = 100 and reduction = 0] remove the audio in the process area and replace it with synthesised data. These are appropriate for removing loud, unwanted noises.

Preview and Apply

When you are happy with your settings, press the Preview button to invoke processing, which will be performed on the active channel(s). The progress bar within the button will show how much of the calculation has been completed. The larger the process area or wings, the longer the process will take. You can abort the process by clicking on Preview a second time. Depending upon the size of the process area, it may take a few moments for the command to take effect.

Sometimes the progress bar will move backwards. This indicates that the process has rejected a given set of synthesised data, and is recalculating a better solution.

If you are not happy with the result, you can pick up and move the marked region or adjust it and its wings using the grab handles. You need not press Preview again - as soon as you have finished adjusting the region, the existing preview will be dumped and the audio will be reprocessed.

When you are happy with the result obtained, press the Apply button to lay the corrected audio back into the file processor.

> Note: You may ignore the Preview stage and immediately press Apply to process the audio and lay it back in a single action.

The figures below show the result of interpolating the marked click.

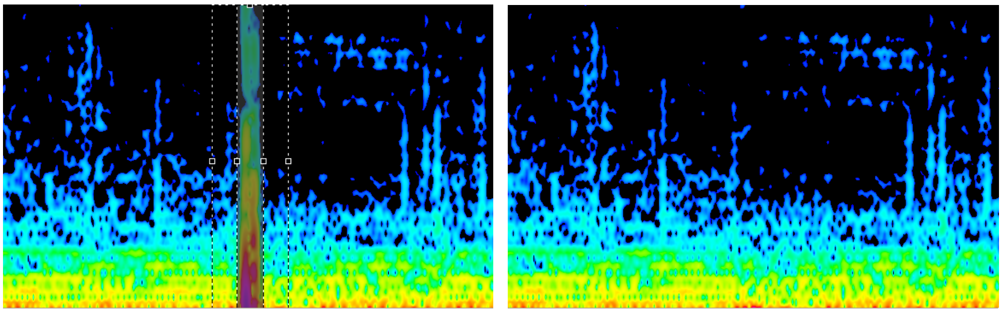

Interpolation example #1: Removing a Noise Burst

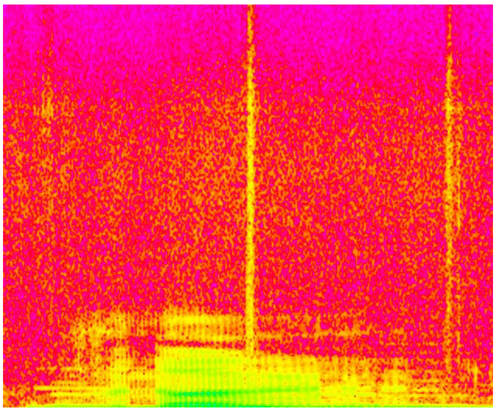

This spectrogram shows about two seconds of 44.1kHz audio. You can see two yellow vertical bars that represent two loud events, one in the centre, and one to the far right. We will concentrate on removing the one in the centre.

At first sight this appears to stretch across all frequencies. However, if you look at the low end, you can see the presence of two musical tones (which appear as the green areas). These tones do not appear to be affected by the disturbance, suggesting that, even though we can not see it, the lowest extent of the unwanted noise is between 1kHz and 3.5kHz.

Too diffuse and far too long to be a distinct click, experience tells us (even without hearing the audio) that this event is a "chuff" of mid- and high- frequency noise, probably from a transcribed disc. We will now use the information gained from this spectrogram to help remove the noise without damaging the desired signal.

Marking the Burst

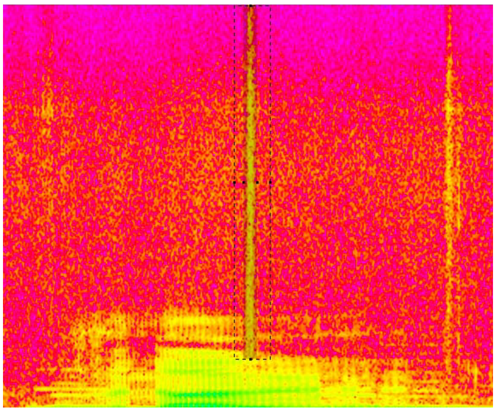

On this occasion, it is not necessary to zoom into the spectrogram to identify the chuff. However, if you wish to do so, you could zoom in on the time axis.

The chuff is short and has little internal structure, so we can expect a single application of Retouch to remove it successfully. This would not necessarily be the case if, for example, it showed a resonant tail.

Having located the chuff, mark it using the interpolate tool. In this example, the durations of the wings have been adjusted such that more audio is included after the chuff than before. This is because there is a signal before the chuff (in the region 4kHz to 5kHz) that could adversely affect the synthesis of the replacement audio.

Furthermore, the chuff has not been marked below 3kHz. This is because the low frequency content of the recording will in all likelihood mask any residual noise if it exists.

Removing the Burst

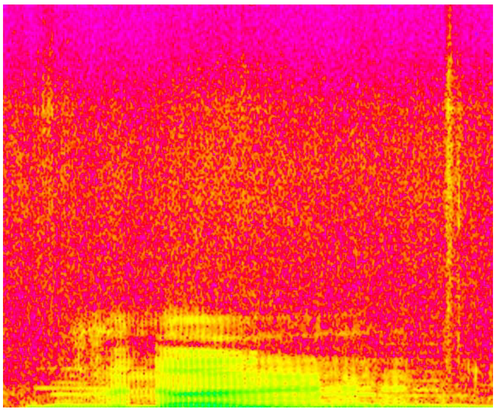

Select an appropriate value for Synthesis. In this example, the maximum value of 100 is correct because (i) there is no obvious meaningful signal within the marked region, and (ii) the audio adjacent to the marked region appears to be consistent and therefore a good model for the replacement.

Because the surrounding signal appears consistent in amplitude, and because we are not incorporating any of the original data in the calculation, it is unnecessary to match the replacement audio using the Gain control. Consequently, set the gain to 0dB.

Now click on Preview. Retouch will remove the chuff, substitute synthesised audio, and re-display the spectrogram with the new audio inserted. As you can see, the chuff has disappeared and the marked region contains audio that is indistinguishable from the surrounding area.

Click on the Play button to audition the result. If it is acceptable, click on Apply to lay the corrected audio back into the File Processor. You may now proceed to process other areas of the spectrogram, or close Retouch as desired.

Interpolation example #2: Removing a Horn

Not all noises are as obvious as the burst in the previous example, and not all can be removed using a single application of Retouch. The following example demonstrates how you may need to use multiple applications to eliminate sounds buried within the signal.

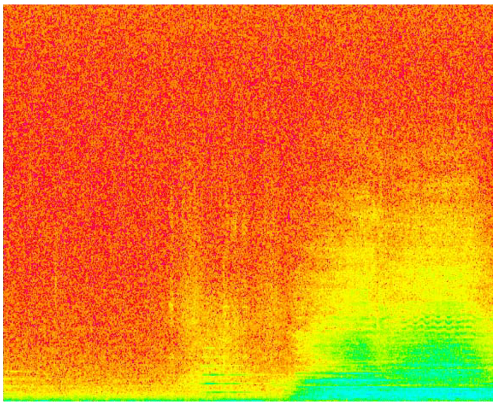

This spectrogram shows the sound of a car horn that occurred during a piece of otherwise wanted signal. You can see this as a small region of green horizontal lines in the centre/bottom of the display.

The sound underneath the car horn is that of a choir inhaling prior to singing the notes that appear as the large green and cyan area on the right of the display. It is important that the horn is removed without damaging the delicate breath sound represented by the diffuse yellow areas.

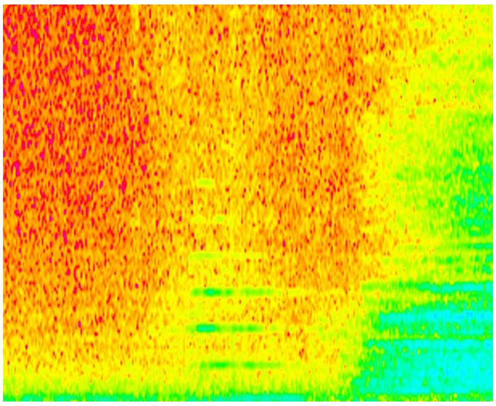

Zooming in on the region containing the horn yields the spectrogram above, and shows the harmonic structure of the noise more clearly. You can see a significant region of audio on either side of the horn; this allows you space to place and manipulate the wings.

The next spectrogram shows how you can remove each of the harmonics individually.

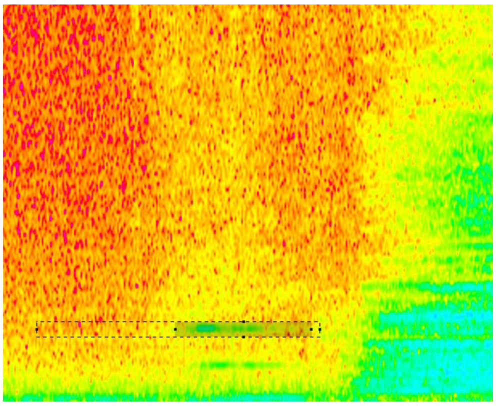

There is nothing to stop you from processing the unwanted noise in one application of Retouch, mapping out the whole area in a single sweep of the Interpolate tool, but you will achieve better results if you apply the process individually to each harmonic. This is because, if you mark the whole region as unwanted, you will be including much of the breath sound as 'bad', when it clearly is not. A more precise approach ensures that you remove the bad sound with minimum impact upon the good.

Removing the harmonics individually

Having decided to remove each harmonic individually, you can mark the first, choose the extent of the wings, choose suitable settings for Synthesis and Gain, and then apply Retouch. If the result of this appears satisfactory, you then select the second harmonic and repeat the process. Do this a third time, and a fourth, and so on until all the harmonics have been eliminated.

It is important that the wings are appropriate for each application of the process, and what is suitable for one harmonic may not be suitable for another (although it often is). If you study the spectrogram above you can see that the wing to the left has been extended, while the one to the right has been truncated. This ensures that none of the audio in the sung notes to the right is used in the synthesis. (The horn is superimposed only on breath sounds, not the notes.)

Removing the harmonics as a set of separated composite shapes

It is possible to remove all the harmonics in a single application of Retouch by marking them as a set of separated composite shapes.

Mark the first harmonic that you wish to remove. Then, while holding down the <SHIFT> key, carefully mark the second, third... and so on until all of them are marked. Retouch will now treat each mark as a separate area of audio to restore, using the information contained *at the appropriate frequencies* in a single pair of wings. You should adjust the extent of the wings as before to ensure that no inappropriate audio is included in the interpolation.

This technique works well for multiple objects separated in frequency (along the Y axis) but is not appropriate for restoring multiple objects - such as a succession of clicks - separated in time (along the X axis). To use Retouch as a declicker, you should mark and process each event individually.

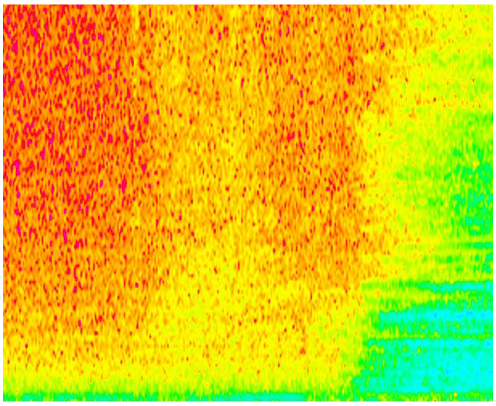

This spectrogram shows the audio with all the harmonics removed.

Patch mode

Patch mode allows you to replace an area of audio with another of the same duration and the same range of frequencies (although not necessarily of the same frequencies). Think of this as copying the data from elsewhere in the spectrogram to the region that you initially defined.

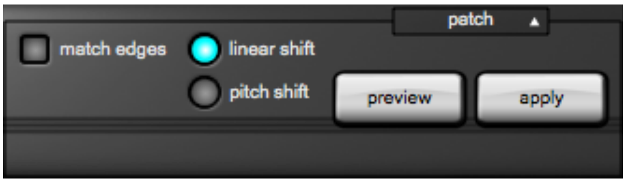

To replace an unwanted burst of noise with some audio from nearby, select Patch mode and define the region that you are going to patch. No wings will appear because they are not appropriate for this operation.

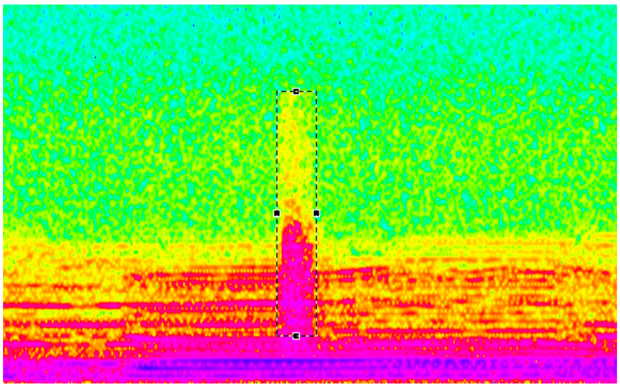

Having marked the audio you will find that the cursor turns into a hand icon when you move it over the marked area. You can now move the selection area to a new position to determine what audio will be substituted into the marked area.

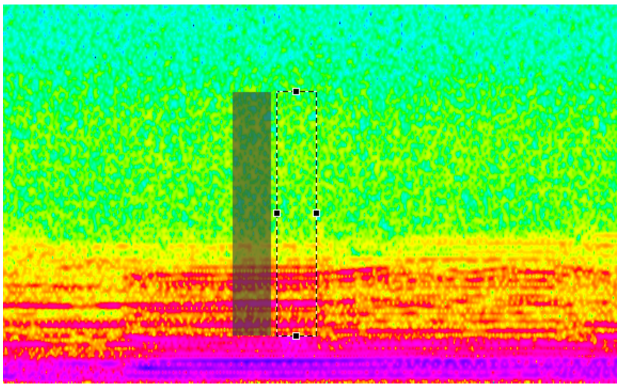

Click on Preview to perform the substitution and, if you are happy with the result, click on Apply to lay the audio back.

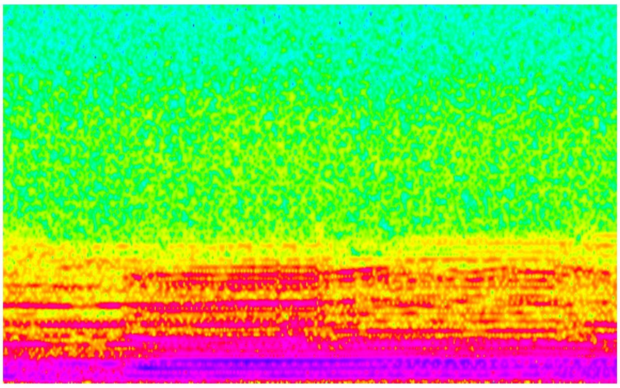

You may click anywhere outside the source or destination areas to clear the marquee. Otherwise, click within the greyed out area to move the source and/or the marqueed area to find a more appropriate substitution.

Lock direction

Often, you will want to patch a block of audio from one point in time to another, but with the same range of frequencies. Click on the lock icon to constrain the movement of the marked area to either up/down or left/right. The direction will be determined when you first move the area.

When lock is On, the arrows change colour, as below.

Linear shift

The source and destination marquees cover the same number of frequencies. For example, a marked area covering the range 1,000Hz – 1,100Hz will cover the range 10,000Hz – 10,100Hz if patched to those frequencies.

Selecting this option is appropriate when the material has no discernible tonal content.

Pitch shift

The marked areas take account of the relative pitch of the components within them. For example, a marked area covering the range 1,000Hz – 1,100Hz will cover the range 10,000Hz – 11,000Hz if patched to those frequencies.

This option invokes a pitch shift algorithm appropriate for harmonically rich material.

Match edges

The amplitude of the patched audio is adjusted to best fit the amplitude of the surrounding audio. Retouch achieves this by 'tilting' the new audio to accommodate changes on the time axis (for example, a fade out) and changes on the frequency axis (for example, to match audio that has been subjected to low-pass filtering).

Copy mode

 copy

Copy mode allows you to move an area of audio to another position. You can think of this as copying the data from a defined region to anywhere else in the spectrogram. In this example, we will use a region of typical background to overwrite an unwanted audio event (which is a noise burst).

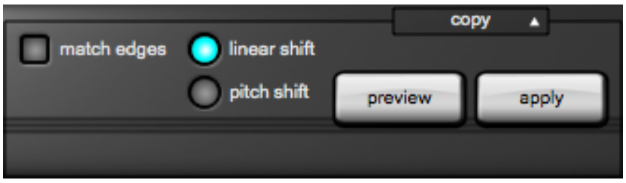

To begin, select Copy mode, and ensure that Match Edges is Off. Now define the region that you are going to copy. No wings will appear because they are not appropriate for this operation.

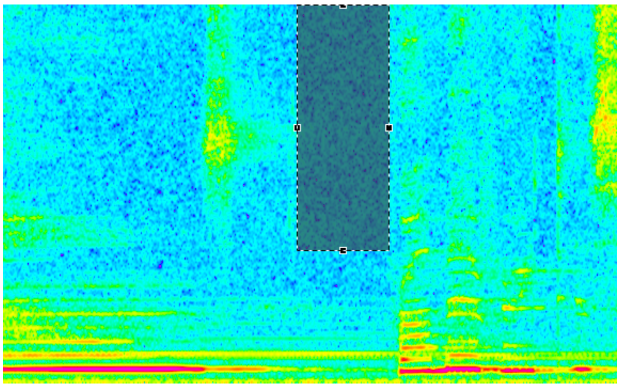

Having marked the audio you will find that the cursor turns into a hand icon when you move it over the marked area. You can now move the audio, depositing it in a new position as you choose.

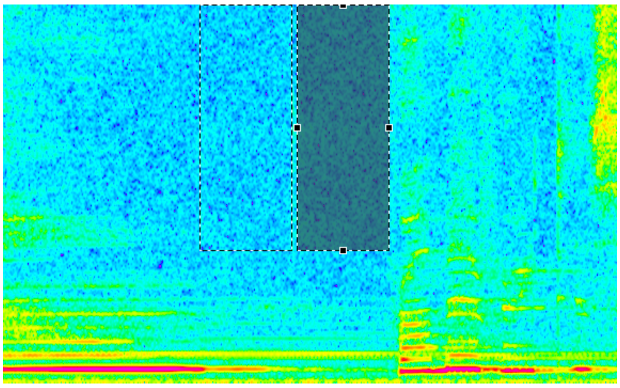

Click on Preview to perform the substitution and, if you are happy with the result, click on Apply to confirm the change.

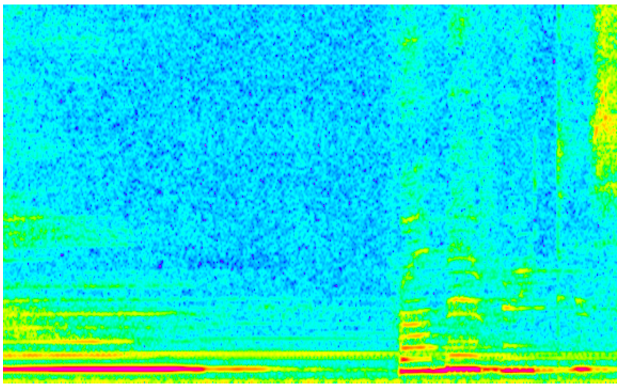

You may click anywhere outside the source or destination areas to clear the marquees. Otherwise, click within the greyed out area to move the source and/or the marqueed area to find a more appropriate substitution.

Lock direction

Often, you will want to copy a block of audio from one point in time to another, but with the same range of frequencies. Click on the lock icon to constrain the movement of the marked area to either up/down or left/right. The direction will be determined when you first move the area.

When lock is On, the arrows change colour, as below.

Linear shift

The source and destination marquees cover the same number of frequencies. For example, a marked area covering the range 1,000Hz – 1,100Hz will cover the range 10,000Hz – 10,100Hz if copied to those frequencies.

Selecting this option is appropriate when the material has no discernible tonal content.

Pitch shift

The marked areas take account of the relative pitch of the components within them. For example, a marked area covering the range 1,000Hz – 1,100Hz will cover the range 10,000Hz – 11,000Hz if patched to those frequencies.

This option invokes a pitch shift algorithm appropriate for harmonically rich material.

Match edges

The amplitude of the copied audio is adjusted to best fit the amplitude of the surrounding audio. Retouch achieves this by 'tilting' the new audio to accommodate changes on the time axis (for example, a fade out) and changes on the frequency axis (for example, to match audio that has been subjected to low-pass filtering).

Erase mode

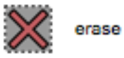

This mode provides a quick and simple way to erase unwanted events and replace them with background (atmos) calculated from the surrounding audio.

In this example, you will replace the harmonics of a car horn with background generated by Retouch.

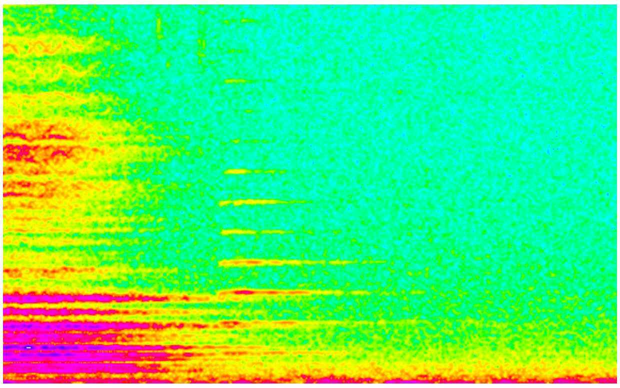

Having identified the harmonics by eye, mark them as shown using the Add to Existing Selection tool or by pressing <SHIFT> while marking them individually.

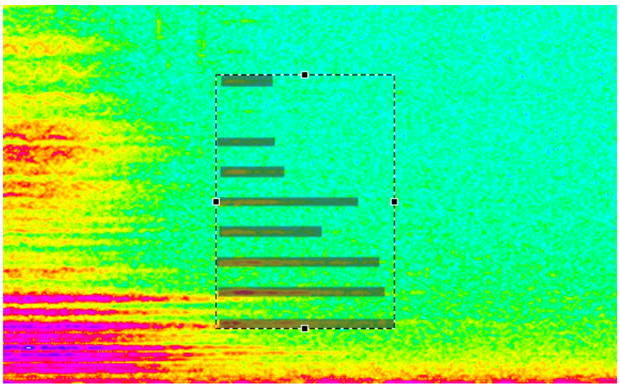

Having done so, click on Preview to erase the audio in the marked regions and substitute background calculated by the process.

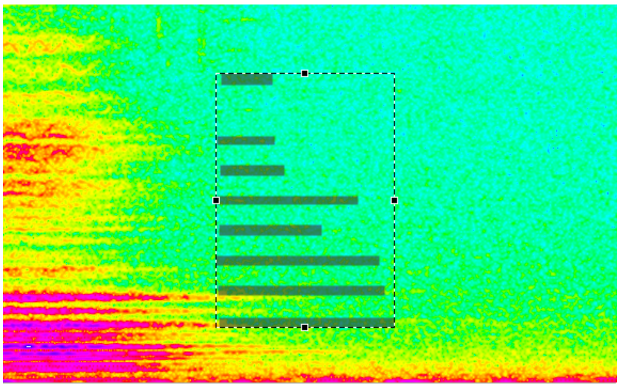

If you are happy with the result of the erase, click on Apply to confirm the change.

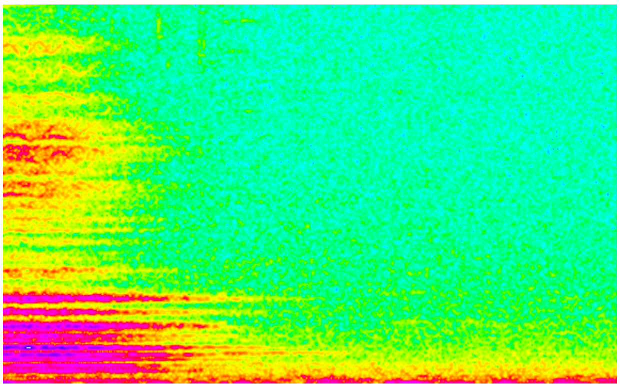

Flat

The amplitude of the audio generated to replace the event in the marked region will be flat on both the time and frequency axes.

Tilted

The amplitude of the audio generated to replace the event in the marked region may have a gradient on the X-axis and/or on the Y-axis. This is calculated such that the edges of the generated audio best match the amplitude of the surrounding audio.

Volume mode

volume

You may affect the amplitude of the signal within a region, and independently affect the amplitude of the signal lying outside that region. This allows you, for example, to reveal individual sounds or utterances within a file, either by amplifying the wanted sounds or by suppressing the rest of the audio, or both simultaneously. By marking multiple regions simultaneously and reducing the amplitude of the 'outside' to its minimum, you can also use this to retain only the sounds or words wanted in an audio file.

The example below shows a number of wanted sounds within a noisy background.

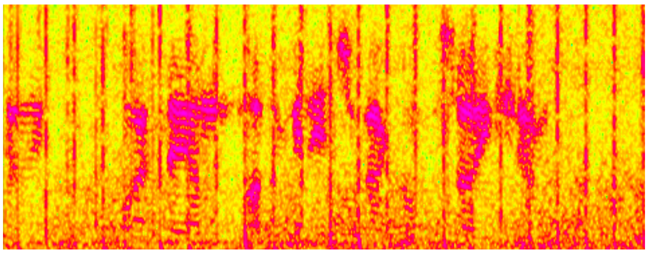

Use the Polyline or paintbrush tools to select these in the usual fashion.

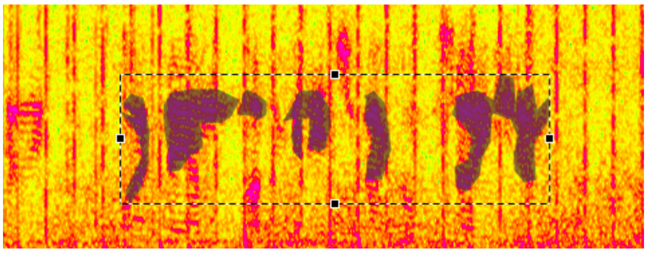

Now choose appropriate value for the Inner (marked) and Outer (unmarked) volumes. In this example, we have left the amplitude of the wanted sounds unaffected but attenuated the background by 50dB.

Press Preview and audition the result. Note that, to eliminate the introduction of artefacts, a short fade is applied to the edges of each of the regions. If you are not happy with this, leave Preview On, adjust the volume controls, wait for Preview to recalculate the audio, and audition again. When you are happy, click on Apply to confirm the change.

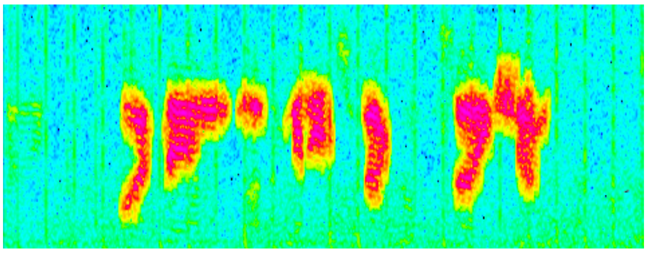

Now click on the spectrogram window to release the marquee.

Cleanse mode

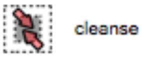

Despite the power of Interpolation, Patching and Copying, there are times when those modes are unable to eliminate unwanted sounds quickly and efficiently. Example of this include a chuff of noise in a dense cluster of notes, or wind roar during speech. Cleanse mode was developed specifically to help in these situations, and may achieve results that cannot be obtained using the other methods.

In this example, we drastically reduce low frequency wind roar in the present of speech. The image below shows the original audio.

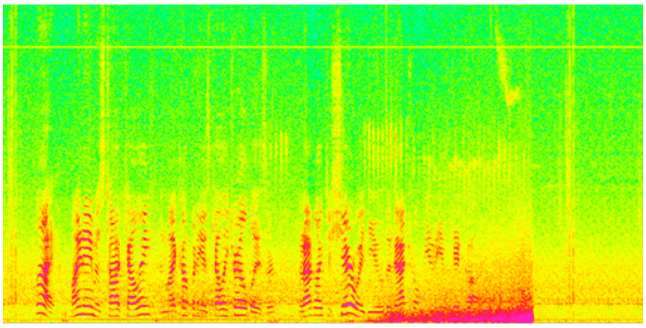

In the next image, the roar has been identified and marked.

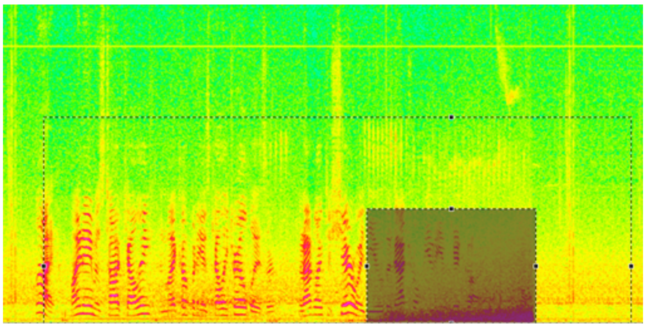

A large area has been marked around the region containing the roar. Retouch uses the information in this to build a model of 'good' audio which will be used to retain those elements in the marked region that are also deemed to be good, while attenuating those that are not.

Press Preview to audition the result. Click on Apply to confirm the change.

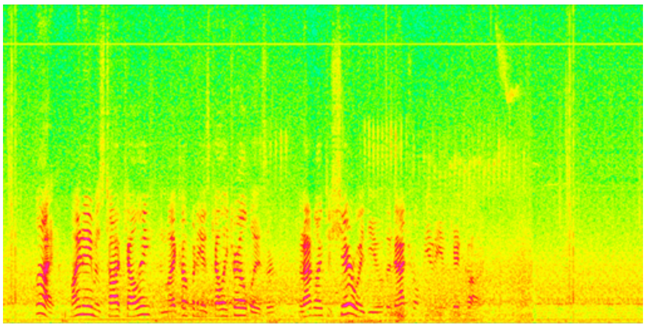

Revert mode

Undoing and redoing allows you step through the history of your actions. However, there may be many occasions when you want to undo an action (or actions) at one point in the timeline while leaving later actions elsewhere unscathed. Revert allows you to do so by replacing the current audio within a region with the audio from the original file.

In this example, we will reintroduce a transient sound that has been mistaken for a click. The image below shows audio that appears to be affected by four clicks.

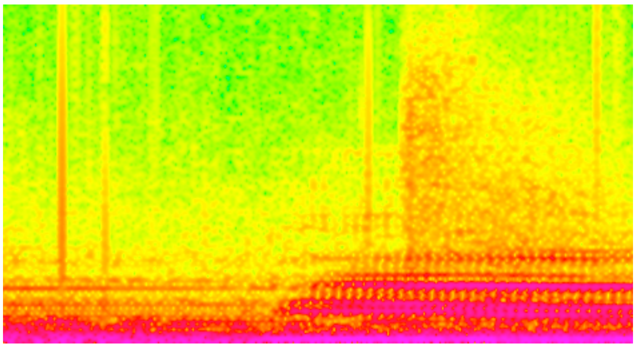

In the next image, these have been Retouched out from left to right using Interpolate mode.

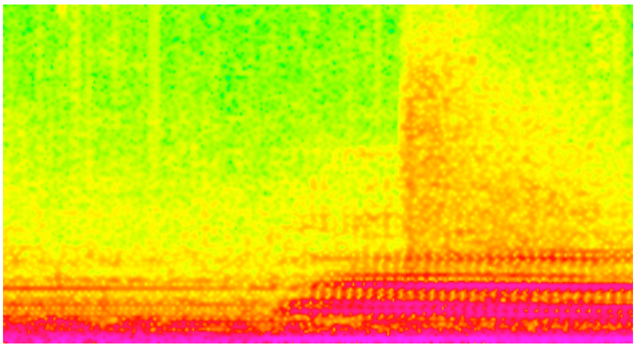

Let’s now suppose that, on playback, you have discovered that the third Retouch removed a genuine sound. You could Undo two stages to reintroduce this, and then reprocess the fourth sound as before. However, it’s much quicker and more efficient to cause that area of the spectrogram to Revert to the original audio.

To do so, mark the region that you would like to Revert.

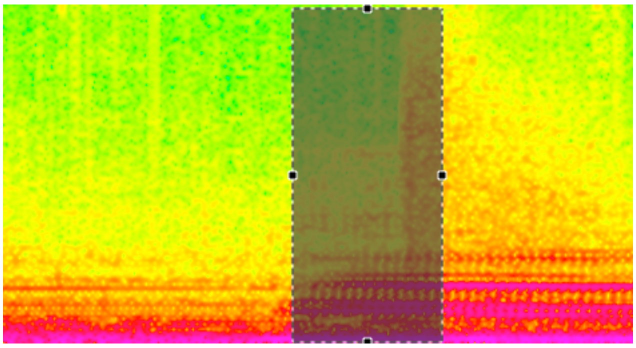

Press Preview to audition the result. Click on Apply to confirm the change.

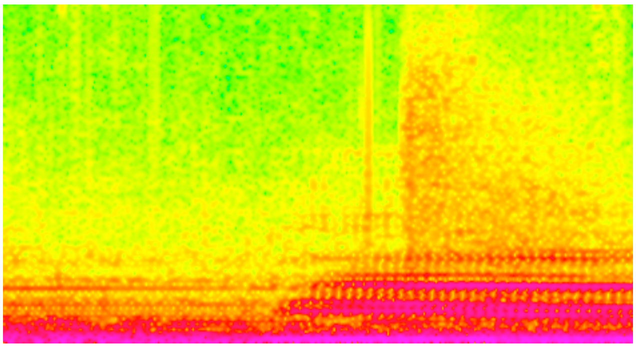

Note: The area reverted does not need to contain all of a previous Retouch (or selection of Retouches). You can cause part of a Retouched region to revert to its original audio content.

Additional information

Licence and Limited Warranty

1. DEFINITIONS

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

'the Company' is CEDAR Audio Limited of 20 Home End, Fulbourn, Cambridge, CB21 5BS, UK;

'the System' means any instance of any CEDAR for Pyramix products developed by the Company, including but not limited to Auto Dehiss, Auto declick, Manual declick, Auto decrackle, Dethump and Retouch;

'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 undertakings 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 hereunder 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 shall be determined in England in accordance with the Arbitration Act 1950 and 1979.

CEDAR for Pyramix 64

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

t: +44 1223 881771
e: info@cedaraudio.com
w: www.cedaraudio.com